
Stochastic Gradient Push for Distributed Deep Learning

Mahmoud Assran^{1,2} Nicolas Loizou^{1,3} Nicolas Ballas¹ Mike Rabbat¹

Abstract

Distributed data-parallel algorithms aim to accelerate the training of deep neural networks by parallelizing the computation of large mini-batch gradient updates across multiple nodes. Approaches that synchronize nodes using exact distributed averaging (e.g., via ALLREDUCE) are sensitive to stragglers and communication delays. The PUSH-SUM gossip algorithm is robust to these issues, but only performs approximate distributed averaging. This paper studies Stochastic Gradient Push (SGP), which combines PUSH-SUM with stochastic gradient updates. We prove that SGP converges to a stationary point of smooth, non-convex objectives at the same sub-linear rate as SGD, and that all nodes achieve consensus. We empirically validate the performance of SGP on image classification (ResNet-50, ImageNet) and machine translation (Transformer, WMT’16 En-De) workloads. Our code will be made publicly available.

1. Introduction

Deep Neural Networks (DNNs) are the state-of-the-art machine learning approach in many application areas, including computer vision (He et al., 2016) and natural language processing (Vaswani et al., 2017). Stochastic Gradient Descent (SGD) is the current workhorse for training neural networks. The algorithm optimizes the network parameters, \boldsymbol{x} , to minimize a loss function, $f(\cdot)$, through gradient descent, where the loss function’s gradients are approximated using a subset of training examples (a mini-batch). DNNs often require large amounts of training data and trainable parameters, necessitating non-trivial computational requirements (Wu et al., 2016; Mahajan et al., 2018).

¹Facebook AI Research, Montréal, QC, Canada ²Department of Electrical and Computer Engineering, McGill University, Montréal, QC, Canada ³School of Mathematics, University of Edinburgh, Edinburgh, Scotland. Correspondence to: Mahmoud Assran <mahmoud.assran@mail.mcgill.ca>.

Large mini-batch parallel SGD is usually adopted for distributed training of deep networks (Goyal et al., 2017; Li et al., 2014). Worker nodes compute local mini-batch gradients of the loss function on different subsets of the data, and then calculate an exact inter-node average gradient using either the ALLREDUCE communication primitive, in synchronous implementations (Goyal et al., 2017; Akiba et al., 2017), or using a central parameter server, in asynchronous implementations (Dean et al., 2012). Using a parameter server to aggregate gradients introduces a potential bottleneck and a central point of failure (Lian et al., 2017). The ALLREDUCE primitive computes the exact average gradient at all workers in a decentralized manner, avoiding issues associated with centralized communication and computation. However, exact averaging algorithms like ALLREDUCE are not robust in communication-constrained settings, *i.e.*, where the network bandwidth is a significant bottleneck.

This issue motivates the investigation of a decentralized and inexact version of SGD to reduce the communication overhead associated with distributed training. There have been numerous decentralized optimization algorithms proposed and studied in the control-systems literature that leverage gossip-based approaches for the computation of aggregate information; see the survey of Nedić et al. (2018) and references therein. State-of-the-art gossip-based optimization methods build on the PUSH-SUM algorithm for distributed averaging (Kempe et al., 2003; Nedić et al., 2018). Rather than computing exact averages (as with ALLREDUCE), this line of work uses less-coupled message passing and computes approximate averages. The tradeoff is that approximate distributed averaging also injects additional noise in the average gradient estimate.

In this work we study Stochastic Gradient Push (SGP), an algorithm blending parallel SGD and PUSH-SUM. SGP enables the use of generic communication topologies that may be directed (asymmetric), sparse, and time-varying. In contrast, existing gossip-based approaches explored in the context of training DNNs (Lian et al., 2017; Jiang et al., 2017) are constrained to use symmetric communication (*i.e.*, if node i sends to j , then i must also receive from j before proceeding) and thus inherently require deadlock-avoidance, and more synchronization, making them slower and more sensitive to stragglers. Moreover, SGP can be seen as a generalization of parallel SGD and these previous approaches.

SGP was first proposed in the control systems literature for minimizing the sum of *strongly-convex* functions (Nedić & Olshevsky, 2016). We make three main contributions. 1) We propose and analyze a novel variant of SGP, called Overlap SGP, which overlaps communication and computation to hide communication overhead. 2) We provide novel theoretical guarantees, proving that SGP (and Overlap SGP) converges to a stationary point of smooth *non-convex* functions at an $\mathcal{O}(1/\sqrt{nK})$ rate, for an appropriately chosen step-size, where n is the number of nodes and K is the number of iterations. 3) We conduct experiments on image classification (ResNet50, ImageNet) and neural machine translation tasks (Transformer, WMT16 En-De), demonstrating that SGP and Overlap SGP can substantially accelerate training of deep neural networks, by reducing communication overhead and mitigating the effects of stragglers. Given a fixed runtime budget, we find that SGP and Overlap SGP can train models that achieve better final train/validation accuracies than ALLREDUCE SGD in communication-constrained settings.

For example, we train a ResNet-50 on ImageNet using 256 GPUs spread across 32 compute nodes (8 GPUs / node), where communication between nodes is over 10 Gbps Ethernet. In this setting ALLREDUCE SGD achieves 76.2% top-1 validation accuracy, while Overlap SGP achieves 76.2% accuracy in only 1/3 of the time, and 77.1% accuracy in 1/2 the time. Similarly, when training a Transformer network on the WMT’16 En-De translation task, SGP runs approximately $1.5\times$ faster than ALLREDUCE when using 8 GPUs and achieves a BLEU score that is 0.6 points higher. While our theory focuses on analyzing SGP, which combines PUSH-SUM with SGD, our experiments illustrate that PUSH-SUM can similarly be efficiently combined with other optimizers like Nesterov’s accelerated gradient method (Nesterov, 1983) and Adam (Kingma & Ba, 2015).

2. Preliminaries

Problem formulation. We consider the setting where a network of n nodes cooperates to solve the stochastic consensus optimization problem

$$\begin{aligned} \min_{\mathbf{x}_i \in \mathbb{R}^d, i=1, \dots, n} \quad & \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\xi_i \sim D_i} F_i(\mathbf{x}_i; \xi_i) \\ \text{subject to} \quad & \mathbf{x}_i = \mathbf{x}_j, \forall i, j = 1, \dots, n. \end{aligned} \quad (1)$$

Each node has local data following a distribution D_i , and the nodes wish to cooperate to find the parameters \mathbf{x} of a DNN that minimizes the average loss with respect to their data, where F_i is the loss function at node i . Moreover, the goal codified in the constraints is for the nodes to reach agreement (*i.e.*, consensus) on the solution they report. We assume that nodes can locally evaluate stochastic gradients $\nabla F_i(\mathbf{x}_i; \xi_i)$, $\xi_i \sim D_i$, but they must communicate to access information about the objective functions at other nodes.

Distributed averaging. The problem described above en-

compasses distributed training based on data parallelism, where the canonical approach is parallel stochastic gradient descent: for an overall mini-batch of size nb , each node computes a local stochastic mini-batch gradient using b samples, and then the nodes use the ALLREDUCE communication primitive to compute the average gradient at every node. Let $f_i(\mathbf{x}_i) = \mathbb{E}_{\xi_i \sim D_i} F_i(\mathbf{x}_i; \xi_i)$ denote the objective at node i , and let $f(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x})$ denote the overall objective. Since $\nabla f(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \nabla f_i(\mathbf{x})$, averaging gradients via ALLREDUCE provides an exact stochastic gradient of f .

Approximate distributed averaging. In this work we explore the alternative approach of using a gossip algorithm for approximate distributed averaging—specifically, the PUSH-SUM algorithm (Kempe et al., 2003). Let $\mathbf{y}_i^{(0)} \in \mathbb{R}^d$ be a vector at node i , and consider the goal of computing the average vector $\frac{1}{n} \sum_{i=1}^n \mathbf{y}_i^{(0)}$ at all nodes. Concatenate the initial vectors into a matrix $\mathbf{Y}^{(0)} \in \mathbb{R}^{n \times d}$ with one row per node. Typical gossip iterations have the form $\mathbf{Y}^{(k+1)} = \mathbf{P}^{(k)} \mathbf{Y}^{(k)}$ where $\mathbf{P}^{(k)} \in \mathbb{R}^{n \times n}$ is referred to as the mixing matrix, and conforms to the underlying communication topology. This corresponds to the update $\mathbf{y}_i^{(k+1)} = \sum_{j=1}^n p_{i,j}^{(k)} \mathbf{y}_j^{(k)}$ at node i . To implement this update, node i only needs to receive messages from other nodes j for which $p_{i,j}^{(k)} \neq 0$, so sparser matrices $\mathbf{P}^{(k)}$ correspond to less communications.

Drawing inspiration from the theory of Markov chains (Seneta, 1981), the mixing matrices $\mathbf{P}^{(k)}$ are designed to be column stochastic (each column must sum to 1). Then, under mild conditions (ensuring that information from every node eventually reaches all other nodes) one can show that $\lim_{K \rightarrow \infty} \prod_{k=0}^K \mathbf{P}^{(k)} = \boldsymbol{\pi} \mathbf{1}^\top$, where $\boldsymbol{\pi}$ is the ergodic limit of the chain and $\mathbf{1}$ is a vector with all entries equal to 1. Consequently, the gossip iterations converge to a limit $\mathbf{Y}^{(\infty)} = \boldsymbol{\pi} (\mathbf{1}^\top \mathbf{Y}^{(0)})$; *i.e.*, the value at node i converges to $\mathbf{y}_i^{(\infty)} = \pi_i \sum_{j=1}^n \mathbf{y}_j^{(0)}$.

When the matrices $\mathbf{P}^{(k)}$ are symmetric, it is straightforward to design the algorithm so that $\pi_i = 1/n$ for all i by making $\mathbf{P}^{(k)}$ doubly-stochastic (each row and each column must sum to 1). However, symmetric $\mathbf{P}^{(k)}$ has strong practical ramifications, such as requiring care in the implementation to avoid deadlocks. The PUSH-SUM algorithm only requires that $\mathbf{P}^{(k)}$ be column-stochastic, and not necessarily symmetric (so node i may send to node j , but not necessarily vice versa). However, when the matrices $\mathbf{P}^{(k)}$ are asymmetric, it is very difficult, and often not possible, to design the algorithm so that $\pi_i = 1/n$. Instead, one additional scalar parameter $w_i^{(k)}$ is maintained at each node. The parameter is initialized to $w_i^{(0)} = 1$ for all i , and updated using the same linear iteration, $\mathbf{w}^{(k+1)} = \mathbf{P}^{(k)} \mathbf{w}^{(k)}$. Consequently, the parameter converges to $\mathbf{w}^{(\infty)} = \boldsymbol{\pi} (\mathbf{1}^\top \mathbf{w}^{(0)})$, or $w_i^{(\infty)} = \pi_i n$ at node i . Thus each node can recover the average of the initial vectors by computing the *de-biased* ratio $\mathbf{y}_i^{(\infty)} / w_i^{(\infty)}$.

Algorithm 1 Stochastic Gradient Push (SGP)

Require: Initialize $\gamma > 0$, $\mathbf{x}_i^{(0)} = \mathbf{z}_i^{(0)} \in \mathbb{R}^d$ and $w_i^{(0)} = 1$ for all nodes $i \in \{1, 2, \dots, n\}$

- 1: **for** $k = 0, 1, 2, \dots, K$, at node i , **do**
- 2: Sample new mini-batch $\xi_i^{(k)} \sim \mathcal{D}_i$ from local distribution
- 3: Compute mini-batch gradient at $\mathbf{z}_i^{(k)}$: $\nabla F_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 4: $\mathbf{x}_i^{(k+\frac{1}{2})} = \mathbf{x}_i^{(k)} - \gamma \nabla F_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 5: Send $(p_{j,i}^{(k)} \mathbf{x}_i^{(k+\frac{1}{2})}, p_{j,i}^{(k)} w_i^{(k)})$ to out-neighbors;
 receive $(p_{i,j}^{(k)} \mathbf{x}_j^{(k+\frac{1}{2})}, p_{i,j}^{(k)} w_j^{(k)})$ from in-neighbors
- 6: $\mathbf{x}_i^{(k+1)} = \sum_j p_{i,j}^{(k)} \mathbf{x}_j^{(k+\frac{1}{2})}$
- 7: $w_i^{(k+1)} = \sum_j p_{i,j}^{(k)} w_j^{(k)}$
- 8: $\mathbf{z}_i^{(k+1)} = \mathbf{x}_i^{(k+1)} / w_i^{(k+1)}$
- 9: **end for**

In practice, we stop after a finite number of gossip iterations K and compute $\mathbf{y}_i^{(K)} / w_i^{(K)}$.

PUSHSUM provides a mechanism to approximately synchronize parameters across a network of nodes. In the next section we describe the Stochastic Gradient Push algorithm, where PUSHSUM is used to approximately synchronize (average) parameters across nodes running stochastic gradient descent locally. The same approach can easily be modified to obtain decentralized versions of other popular optimizers such as SGD with momentum or Adam, as illustrated in the experimental section.

3. Stochastic Gradient Push

The *Stochastic Gradient Push* (SGP) method (Nedić & Olshevsky, 2016) for solving equation 1 is obtained by interleaving one local stochastic gradient descent update at each node with one iteration of PUSHSUM. Pseudocode is shown in Alg. 1. Each node maintains three variables: the model parameters $\mathbf{x}_i^{(k)}$ at node i , the scalar PUSHSUM weight $w_i^{(k)}$, and the de-biased parameters $\mathbf{z}_i^{(k)} = \mathbf{x}_i^{(k)} / w_i^{(k)}$. The vector $\mathbf{x}_i^{(0)}$ can be initialized arbitrarily. At each iteration, every node performs a local SGD step (lines 2–4) followed by one step of PUSHSUM for approximate distributed averaging (lines 5–8). Note that the gradients are evaluated at the de-biased parameters $\mathbf{z}_i^{(k)}$ in line 3, and they are then used to update $\mathbf{x}_i^{(k)}$, in line 4. All communication takes place in line 5, and each message contains two parts, the PUSHSUM numerator $\mathbf{x}_i^{(k)}$ and PUSHSUM weight $w_i^{(k)}$.

The non-zero entries in the mixing matrix $\mathbf{P}^{(k)}$ define the communication topology at each iteration k . SGP can leverage various communication topologies including sparse, asymmetric or time-varying networks. We are mainly interested in the case where the mixing matrices $\mathbf{P}^{(k)}$ are sparse in order to have low communication overhead. However, we point out that when the nodes’ initial values are identi-

cal, $\mathbf{x}_i^{(0)} = \mathbf{x}_j^{(0)}$ for all $i, j \in [n]$, and every entry of $\mathbf{P}^{(k)}$ is equal to $1/n$, then SGP is mathematically equivalent to parallel SGD using ALLREDUCE.

Individual nodes do not need to know the entire mixing matrix at each time step. Each node i must only know/choose its outgoing mixing weights, which correspond to the i^{th} column of $\mathbf{P}^{(k)}$. Each node can consequently choose its mixing weights independently of the other nodes in the network. In Appendix A we describe one way to design a sequence of mixing matrices that satisfies the requirements of our theory (described in the next section) and for which each node sends and receives exactly one message at every iteration; all appendices are in the Supplementary Material.

Overlapping communication and computation. Although SGP does not use network-wide collective communication primitives like ALLREDUCE, the implementation of Alg. 1 requires using blocking sends and receives; *i.e.*, nodes do not proceed to line 6 until they have received messages from all in-neighbors at that iteration. To hide the communication overhead, we can overlap gradient computation with communication. For a given $\tau \geq 0$, nodes send messages to their out-neighbours every τ iterations (non-blocking), and can receive incoming messages at any time in-between communication intervals. If a node hasn’t received messages from its in-neighbors after τ iterations, then it blocks and waits to receive the messages.

Specifically, the communication in line 5 in Alg. 1 is made non-blocking, and each node may perform τ gradient update steps while it occurs. This may result in the gossip updates in lines 6 and 7 incorporating outdated messages, $(p_{j,i}^{(k')} \mathbf{x}_i^{(k'+\frac{1}{2})}, p_{j,i}^{(k')} w_i^{(k')})$, where $k - k' \leq \tau$. However, as long as the delay $k - k'$ remains bounded, SGP is still guaranteed to converge (see Theorem 2 below). We refer to this method, with delay bound τ as τ -overlap SGP (τ -OSGP). SGP is equivalent to τ -OSGP with $\tau = 0$. In practice we find that taking τ to be 1 or 2 is sufficient to hide effectively all of the communication overhead. Pseudocode for τ -OSGP is provided in Appendix B.

4. Theoretical guarantees.

SGP was first proposed and analyzed in (Nedić & Olshevsky, 2016) assuming the local objectives $f_i(\mathbf{x})$ are strongly convex. Here we provide convergence results in the more general setting of smooth, non-convex objectives with arbitrary, but bounded, message delays. We make the following four assumptions:

1. (L -smooth) There exists a constant $L > 0$ such that $\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$.
2. (Bounded variance) There exist finite positive constants

σ^2 and ζ^2 such that

$$\mathbb{E}_{\xi \sim D_i} \|\nabla F_i(\mathbf{x}; \xi) - \nabla f_i(\mathbf{x})\|^2 \leq \sigma^2 \quad \forall i, \forall \mathbf{x}, \text{ and}$$

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

Thus σ^2 bounds the variance of stochastic gradients at each node, and ζ^2 quantifies the similarity of data distributions at different nodes.

3. (Bounded delay) There exists a finite constant $\tau \in \mathbb{Z}_+$, such that the delay, if overlapping communication and computation, satisfies $k' - k \leq \tau$.
4. (Mixing connectivity) To each mixing matrix $\mathbf{P}^{(k)}$ we can associate a graph with vertex set $\{1, \dots, n\}$ and edge set $E^{(k)} = \{(i, j) : p_{i,j}^{(k)} > 0\}$; i.e., with edges (i, j) from j to i if i receives a message from j at iteration k . By convention, we take each node to be an in-neighbor of itself (each node in the graph has a self-loop), and we assume that there exists finite, positive integers, B and Δ , such that the graph with edge set $\bigcup_{k=lB}^{(l+1)B-1} E^{(k)}$ is strongly connected and has diameter at most Δ for every $l \geq 0$.¹

Let $\bar{\mathbf{x}}^{(k)} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^{(k)}$. Lian et al. (2017) define that a decentralized algorithm for solving equation 1 converges if, for any $\epsilon > 0$, it eventually satisfies

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(k)})\|^2 \leq \epsilon. \quad (2)$$

We show that SGP converges in this sense.

Theorem 1. *Suppose that Assumptions 1–4 hold, and run SGP for K iterations with step-size $\gamma = \sqrt{n/K}$. Let $f^* = \inf_{\mathbf{x}} f(\mathbf{x})$ and assume that $f^* > -\infty$. There exist constants C and $q \in [0, 1)$, which depend on the diameter of the network Δ , the upper bound on the delays τ , and the sequence of mixing matrices $\mathbf{P}^{(k)}$, such that when the total number of iterations satisfies*

$$K \geq \max \left\{ \frac{nL^4 C^4 60^2}{(1-q)^4}, \frac{L^4 C^4 P_1^2 n}{(1-q)^4 (f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L\sigma^2}{2})^2}, \frac{L^2 C^2 n P_2}{(1-q)^2 (f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L\sigma^2}{2})}, n \right\} \quad (3)$$

¹For the purpose of analysis, we model delays by augmenting the mixing-matrices $\mathbf{P}^{(k)}$, and the corresponding graph topologies, with virtual nodes and edges that store the state of information that was transmitted, but not yet received. We omit this description from the main text to simplify the discussion, and relegate this discussion to the supplementary material *Modeling message delays* in Appendix E.

where $P_1 = 4(\sigma^2 + 3\zeta^2)n + \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n}$ and $P_2 = \sigma^2 + 3\zeta^2 L^2 C^2 + 2 \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n}$, then

$$\frac{\sum_{k=0}^{K-1} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(k)})\|^2}{K} \leq \frac{12(f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L\sigma^2}{2})}{\sqrt{nK}}.$$

The proof is given in Appendix E, where we also provide precise expressions for the constants C and q .

Theorem 1 shows that the average of the nodes' parameters, $\bar{\mathbf{x}}^{(k)}$, converges, but it does not directly say anything about the parameters at each node. In fact, we can show that:

Theorem 2. *Under the same assumptions as in Theorem 1,*

$$\frac{1}{nK} \sum_{k=0}^{K-1} \sum_{i=1}^n \mathbb{E} \|\bar{\mathbf{x}}^{(k)} - \mathbf{z}_i^{(k)}\|^2 \leq \mathcal{O} \left(\frac{1}{K} + \frac{1}{K^{3/2}} \right),$$

and

$$\frac{1}{nK} \sum_{k=0}^{K-1} \sum_{i=1}^n \mathbb{E} \|\nabla f(\mathbf{z}_i^{(k)})\|^2 \leq \mathcal{O} \left(\frac{1}{\sqrt{nK}} + \frac{1}{K} + \frac{1}{K^{3/2}} \right).$$

The proof is also given in Appendix E. This result shows that as K grows, the de-biased variables $\mathbf{z}_i^{(k)}$ converge to the node-wise average $\bar{\mathbf{x}}^{(k)}$, and hence the de-biased variables at each node also converge to a stationary point. Note that for fixed n and large K , the $1/\sqrt{nK}$ term will dominate the other factors.

5. Related Work

A variety of approaches have been proposed to accelerate distributed training of DNNs in the communication bound setting, including quantizing gradients (Alistarh et al., 2007; Wen et al., 2007; Jia et al., 2018) and performing multiple local SGD steps at each node before averaging (McMahan et al., 2017). These approaches are complementary to the approach considered in this paper, which advocates for using approximate rather than exact distributed averaging. Quantizing gradients, performing multiple local SGD steps between averaging, and using approximate distributed averaging can all be seen as injecting additional noise (due to approximations) into SGD, leading to a tradeoff between reducing communication (towards training faster) and potentially obtaining worse predictive accuracy (due to approximations). Combining these approaches (quantized, infrequent, and inexact averaging) is an interesting direction for future work.

Blot et al. (2016) report initial experimental results on small-scale experiments with an SGP-like algorithm. Jin et al. (2016) make a theoretical connection between PUSHSUM-based methods and Elastic Averaging SGD (Zhang et al.,

2015). Relative to those previous works, we provide the first convergence analysis for a PUSHSUM-based method in the smooth non-convex case. Moreover, our analysis also holds in the presence of bounded message delays.

Lian et al. (2017) and Jiang et al. (2017) study synchronous gossip-based versions of SGD. Those methods involve symmetric message passing which inherently involves blocking (if i sends to j at iteration k , then j also sends to i before both nodes update). Consequently, they are slower in communication-constrained settings, in comparison to PUSHSUM-based SGP which may use directed message passing (i can send to j without needing a response). The work of Jakovetić et al. (2014) studies gossip-based versions of the Nesterov gradient method for smooth strongly-convex functions with deterministic gradients. The method requires performing two rounds of symmetric message passing per gradient update, and consequently is also slower in communication-constrained settings.

Decentralized parallel SGD (D-PSGD) (Lian et al., 2017) produces iterates whose node-wise average, $\bar{x}^{(k)}$, converges in the sense of equation 2. Our results in Sec. 3 show that SGP converges in the same sense and go beyond to show that the individual values at each node also converge to a stationary point, since the values at each node converge to the network-wide average. In fact, SGP is a generalization of D-PSGD: when the communication topology is static, undirected, and connected at every iteration, and when nodes use symmetric mixing weights ($p_{j,i}^{(k)} = p_{i,j}^{(k)}$ for all (i, j)), then the push-sum weights $w_i^{(k)}$ are always equal to 1 and SGP is mathematically equivalent to D-PSGD. We compare SGP with D-PSGD experimentally in Section 6 and find that SGP is consistently faster and the two methods find solutions of comparable accuracy.

Jin et al. (2016) and Lian et al. (2018) study asynchronous gossip-based methods for training DNNs. Lian et al. (2018) analyzes an asynchronous version of D-PSGD and proves that its node-wise averages also converge to a stationary point. In general, these contributions focusing on asynchrony can be seen as orthogonal to the use of a PUSHSUM based protocol for approximate distributed averaging. Moreover, we find that synchronous Overlap SGP runs faster than asynchronous state-of-the-art AD-PSGD, and produces models with better training/validation performance.

6. Experiments

Next we experimentally compare SGP with ALLREDUCE SGD (AR-SGD), D-PSGD, and asynchronous D-PSGD (AD-PSGD). We aim to study the relationship between runtime and predictive accuracy as a function of the number of nodes used and the network bandwidth. In low-bandwidth experiments the servers communicate over 10 Gbps Ethernet

links (typical in data centers) and in high-bandwidth experiments they communicate over 100 Gbps InfiniBand (typical in high-performance computing clusters). To illustrate the versatility of SGP, we consider two typical workloads: image classification and machine translation. All algorithms are implemented in PyTorch (Paszke et al., 2017).

While our SGP analysis focuses solely on the combination of PUSHSUM with SGD, we leverage Nesterov momentum or Adam in practice. Each node sends and receives one message per iteration in our SGP baseline implementation, and the destination and source of these messages changes from one iteration to the next. We refer the reader to Appendix A for implementation details, including how we design the sequence of mixing matrices $P^{(k)}$.

6.1. Image classification

We train a ResNet-50 (He et al., 2016) on the ImageNet classification task (Russakovsky et al., 2015).² Our experiments use 32 NVIDIA DGX-1 servers. Each server has 8 V100 GPUs. To investigate scaling we run experiments using 4, 8, 16, and 32 servers (*i.e.*, 32, 64, 128, and 256 GPUs). We follow the experimental protocol of Goyal et al. (2017). Every node uses a mini-batch size of 256, so using more nodes corresponds to larger effective mini-batch size. Unless indicated otherwise, all experiments are run for 90 epochs, the learning rate warms up to $n \times 0.1$ during the first five epochs following Goyal et al. (2017) and is decayed by a factor of 10 at epochs 30, 60, and 80. All methods use Nesterov momentum.

Scaling and convergence. The first set of experiments studies the scaling and convergence properties of our baseline SGP implementation, where every node sends and receives one message at every iteration (1-peer) and we do not overlap communication and computation (*i.e.*, $\tau = 0$).

Figure 1 (a) shows the iteration-wise training convergence of SGP, ALLREDUCE-SGD, and D-PSGD. Note that when we increase the number of nodes by a factor of 2, we also decrease the total number of iterations by the same factor. Figure 1 (b) shows the time-wise training convergence over 8 and 16 node Ethernet networks. In all cases, SGP completes 90 epochs in less time than ALLREDUCE-SGD and D-PSGD. Figures 1 (c) and (d) show the scaling efficiency of the methods on both 10 Gbps Ethernet and 100 Gbps InfiniBand networks. In the case of the InfiniBand network, all methods exhibit a near linear scaling (constant time per iteration), which is expected since communication is not a bottleneck in this setting. On the other hand, over the 10 Gbps Ethernet network, as we increase the number of

²ImageNet was only used for the non-commercial research purposes of this paper and not for training networks deployed in production or for other commercial purposes.

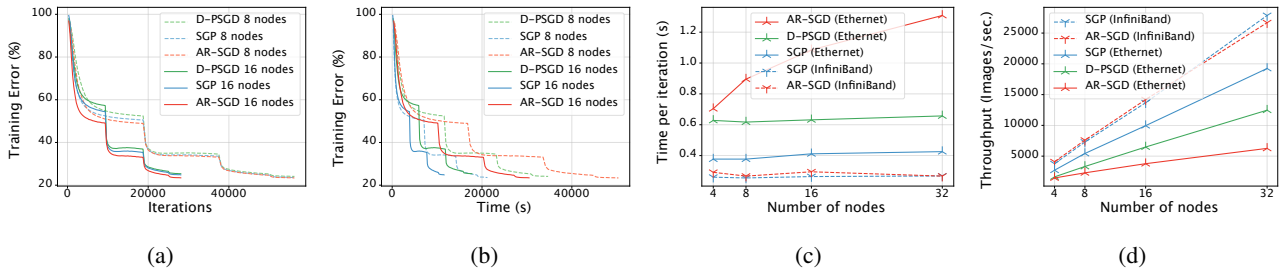


Figure 1: Scaling and convergence results on 4–32 nodes interconnected via 10 Gbps Ethernet and 100Gbps-InfiniBand for ALLREDUCE-SGD (AR-SGD), SGP and D-PSGD. (a)/(b): Iteration-wise convergences and time-wise convergence over 10 Gbps Ethernet. (c)/(d): Time-wise scaling efficiency over Ethernet and InfiniBand networks.

	4 nodes (32 GPUs)		8 nodes (64 GPUs)		16 nodes (128 GPUs)		32 nodes (256 GPUs)	
AR-SGD	76.2%	22.0 hrs.	76.4%	14.0 hrs.	76.3%	8.5 hrs.	76.2%	5.1 hrs.
D-PSGD	76.4%	19.7 hrs.	76.1%	9.7 hrs.	75.9%	5.0 hrs.	74.4%	2.6 hrs.
SGP	76.3%	11.8 hrs.	76.4%	5.9 hrs.	75.9%	3.2 hrs.	75.0%	1.7 hrs.

Table 1: Top-1 validation accuracy (%) and training time (hours), when communicating over 10 Gbps Ethernet for ALLREDUCE-SGD (AR-SGD), SGP and D-PSGD. SGP and D-PSGD are using 1-peer communication topologies.

nodes, the average iteration time stays almost constant for SGP and D-PSGD, while the per-iteration time of ALLREDUCE-SGD significantly increases, resulting in an overall slower training time. Moreover, although D-PSGD and SGP both exhibit strong scaling, SGP is roughly $1.5\times$ faster over 10 Gbps Ethernet, supporting the claim that it involves less communication overhead.

Table 1 shows the total training time and top-1 validation accuracy of the different runs over the 10 Gbps Ethernet network using the baseline 1-peer topologies. For any number of nodes used in our experiments, we observe that SGP consistently outperforms D-PSGD and ALLREDUCE-SGD in terms of total training time. In particular, for 32 node networks (256 GPUs), SGP training takes approximately 1.7 hours, while D-PSGD and ALLREDUCE-SGD require roughly 2.6 and 5.1 hours respectively.

To get a sense of the difference in runtimes between Ethernet and InfiniBand, and also to illustrate the variability, Table 2 shows the mean training time and top-1 validation accuracy along with the maximum absolute deviation from the mean, for 4- and 16-node experiments run on 100 Gbps InfiniBand networks. The max. absolute deviations are calculated based on five runs of each algorithm, using five different seeds. Even with the high-bandwidth InfiniBand network, SGP exhibits less variation in training time across different runs, supporting the claim that SGP helps reduce the effects of stragglers or other sources of latency.

All methods achieve roughly the same validation accuracy for smaller 1-peer topologies (4 and 8 nodes), and the accu-

	4 nodes (32 GPUs)	16 nodes (128 GPUs)
AR-SGD	$76.3 \pm 0.2\%$ 8.8 ± 0.4 hrs.	$76.2 \pm 0.2\%$ 2.5 ± 0.3 hrs.
SGP	$76.3 \pm 0.2\%$ 8.2 ± 0.1 hrs.	$75.8 \pm 0.2\%$ 2.2 ± 0.1 hrs.

Table 2: Top-1 validation accuracy and training time statistics (mean \pm max. abs. deviation), across 5 different runs of each algorithm using 5 different seeds. All methods complete 90 epochs, communicating over 100 Gbps InfiniBand.

racy of D-PSGD and SGP degrades for larger topologies (16 and 32 nodes). We hypothesize that this is due to the error introduced by approximate distributed averaging—models at different nodes are only approximately identical, and for larger networks the divergence between models at different nodes is larger. We investigate this further next.

Parameter deviations. We track the deviations between the model parameters at different nodes while training with SGP. Our convergence result (cf. Lemma 3 in Appendix E) states that each node converges exponentially fast to a neighborhood of the average, where the size of the neighborhood is proportional to the step-size and the connectivity of the graph topology.

Figure 2 shows the average Euclidean distance between the parameters at individual nodes and the node-wise average, at the end of each epoch, for two different network configurations. The sparse graph corresponds to the time-varying 1-




Figure 2: Parameter Deviations for ResNet-50 trained on ImageNet (using SGP) over 16 nodes. Figures show the average Euclidean distance between the parameters at individual nodes and the node-wise average; recorded at the end of each training epoch, after the last gradient step, but before the last gossip step. Shaded regions show the max. and min. parameter deviations across all nodes. Nodes’ parameters exhibit much greater deviation from the average when communicating over sparse topologies, than when communicating over dense topologies.

peer communication strategy used in the experiments above. The dense graph is fully-connected (all-to-all), and the deviations are computed just after line 4 in Alg. 1, before communicating.

Indeed, we see that distance from the mean model is proportional to the magnitude of the step-size, and each node’s parameters are approximately equidistant from the average. At epochs 30, 60, and 80, the parameter deviations drop by an order of magnitude, in concert with the learning rate. We also see the parameter deviations gradually increase in the first five epochs, following the learning rate warm-up. The relation between the parameter deviations and the communication topology is also evident. Indeed, the dense topology exhibits significantly less parameter deviations than the time-varying sparsely connected topology (1-peer communication topology). Therefore, one can directly control the parameter deviations by adjusting the learning rate and/or graph topology, as predicted by Lemma 3.

Communication and the speed-accuracy tradeoff.

Next we explore the effect of the communication topology on the speed-accuracy tradeoff when training with 16 and 32 nodes (128 and 256 GPUs) over 10 Gbps Ethernet. Table 3 shows the validation accuracy and wall-clock time for SGP using a 1-peer topology (1P-SGP), and SGP using a 2-peer topology (2P-SGP), *i.e.*, each node sends and receives to two peers at each iteration. Using just this one additional neighbor improves the validation accuracy of SGP to 76.2% in the 16 nodes case and to 75.7% for 32 nodes, while retaining much of the speed advantages. We also experiment with hybrid communication schemes that use more communication at the start of training, to mitigate the parameter deviations

	16 nodes (128 GPUs)		32 nodes (256 GPUs)	
AR-SGD	76.3%	8.5 hrs.	76.2%	5.2 hrs.
2P-SGP	76.2%	5.1 hrs.	75.7%	2.5 hrs.
1P-SGP	75.9%	3.2 hrs.	75.0%	1.7 hrs.
AR/1P-SGP	76.2%	4.8 hrs.	75.4%	2.8 hrs.
2P/1P-SGP	76.0%	3.5 hrs.	75.1%	1.8 hrs.

Table 3: Top-1 validation accuracies (%) and training time (hours) for 1P-SGP (1-peer topology); 2P-SGP (2-peer topology), AR-SGD (ALLREDUCE SGD), AR/1P-SGP (ALLREDUCE first 30 epochs, 1-peer topology last 60 epochs), and 2P/1P-SGP (2-peer topology first 30 epochs, 1-peer topology last 60 epochs), all communicating over 10 Gbps Ethernet.

	Train Acc.	Val. Acc.	Train Time
AR-SGD	76.9%	76.3%	8.5 hrs.
D-PSGD	75.6%	75.9%	4.9 hrs.
AD-PSGD	74.7%	75.5%	2.9 hrs.
SGP	75.6%	75.9%	3.2 hrs.
biased 1-OSGP	75.4%	75.3%	1.8 hrs.
1-OSGP	77.1%	75.7%	1.8 hrs.

Table 4: Comparing state-of-the-art synchronous and asynchronous gossip-based approaches to 1-OSGP, an implementation of synchronous SGP where communication is overlapped with 1 gradient step (all messages are always received with 1-iteration of staleness). 1-OSGP is also compared with a biased implementation of 1-OSGP that directly incorporates delayed messages without accounting for the bias in the push-sum weight. Experiments are run for 90 epochs over 16 nodes (128 GPUs) interconnected via 10 Gbps Ethernet.

when they are largest (cf. Figure 2). Table 3 compares 1P-SGP, 2P-SGP, and ALLREDUCE SGD (mathematically equivalent to running SGP with a fully-connected topology) with two hybrid methods: AR/1P-SGP, which uses ALLREDUCE for the first 30 epochs and 1-peer SGP for the last 60 epochs, and 2P/1P-SGP, which uses 2-peer SGP for the first 30-epochs and 1-peer SGP for the remainder. We find that these hybrid communication schemes provide a balance between speed and accuracy, and that communicating more during the first few epochs of training can mitigate accuracy tradeoffs.

Overlap SGP. Table 4 compares 1-OSGP using a 1-peer communication topology, with AD-PSGD, D-PSGD, and a biased implementation of 1-OSGP that directly incorporates delayed messages without accounting for the bias in the push-sum weight. Overlapping communication and computation greatly speeds up training and results in no accuracy degradation relative to non-overlap SGP. In contrast, the bi-

	Train Acc.	Val. Acc.	Train Time
AR-SGD	76.9%	76.2%	5.1 hrs. (90 epochs)
AD-PSGD	80.3%	76.9%	4.7 hrs. (270 epochs)
SGP	80.0%	77.1%	4.6 hrs. (270 epochs)
1-OSGP	81.8%	77.1%	2.7 hrs. (270 epochs)


Table 5: Comparing ALLREDUCE SGD (AR-SGD) and SGP under a fix runtime budget. Given a similar runtime, SGP outperforms SGD for both training and validation accuracy. Running 1-OSGP for the same number of epochs than SGP outperforms SGD while improving the overall training efficiency. Experiments are run over 1-peer graph topologies, using 32 nodes (256 GPUs) interconnected via 10 Gbps Ethernet.

ased implementation of 1-OSGP (not using the PUSHSUM weight) has significantly worse accuracy, supporting our theoretical development that the bias tracked in the push-sum weight facilitates convergence. Moreover, synchronous 1-OSGP runs faster than the state-of-the-art asynchronous AD-PSGD, and achieves better training and validation accuracy.

Fixed runtime budget. The results so far highlight that SGP completes 90 epochs of training faster than ALLREDUCE SGD, especially in communication-constrained settings, and this comes at the cost of reduced accuracy. However, if we compare the algorithms based on a runtime budget rather than an epoch budget, SGP achieves superior results. Specifically, since the baseline SGP is roughly $3\times$ faster than ALLREDUCE SGD when run on 32 nodes over 10 Gbps Ethernet, we run SGP for 270 epochs (instead of 90 epochs) using a scaled learning rate schedule (similar warm-up, but decaying the learning rate by a factor of 10 at epochs 90, 180, and 240). With this setup (32 nodes/256 GPUs and 10 Gbps Ethernet), SGP surpasses the best ALLREDUCE SGD accuracy (76.2% after 90 epochs/5.1 hrs.), and achieves a top-1 validation accuracy of 77.1% at then end of the 270 epochs (4.6 hrs.). Similarly, overlap SGP achieves a top-1 validation accuracy of 77.1% at then end of the 270 epochs (2.7 hrs.). When run longer, AD-PSGD achieves better accuracy than ALLREDUCE SGD, but not as fast nor as high of an accuracy as achieved by 1-OSGP.

6.2. Neural Machine Translation

We train a transformer network (Vaswani et al., 2017) on WMT16-En-De using our baseline implementation of SGP, and utilizing the same hyperparameters as (Vaswani et al., 2017). We train models using both 25K token batches (Vaswani et al., 2017), and 400K token batches (Ott et al., 2018). All methods use Adam. Training is performed on 8 NVIDIA V100 GPUs, located on 8 different machines,



(a) Iteration-wise convergence (b) Time-wise convergence

Figure 3: Neural Machine Translations experiments run over $8\times V100$ GPUs located on 8 different machines, interconnected via 10 Gbps Ethernet. Adam-SGP and ALLREDUCE Adam-SGD iteration- and time-wise convergence in both small- and large-batch settings. SGP makes slightly more progress per iteration in both small- and large-batch settings, and runs approximately $1.5\times$ faster in the large-batch setting and $2\times$ faster in the small batch setting.

interconnected via a 10 Gbps Ethernet link. This setup is similar to the AWS P3.2XLARGE machines commonly used for distributed training (Bernstein et al., 2019).

Figures 3 (a) and (b) show the iteration-wise and time-wise validation curves (respectively) of (Adam) SGP and ALLREDUCE (Adam) SGD using small- and large-batch training. We find that SGP makes slightly more progress per iteration in both small- and large-batch settings, and runs approximately $1.5\times$ faster than ALLREDUCE SGD in the large-batch setting, and $2\times$ faster in the small-batch setting. We also evaluate the test-set BLEU scores for models trained using 400K tokens batch. We investigate the performance obtained after the same number of iterations ($\sim 14K$), and after training for the same amount of time (3 days). We evaluate BLEU scores using a beam search of 4, and a length penalty of 0.6, following (Vaswani et al., 2017). SGP achieves a superior BLEU score to ALLREDUCE SGD, both after a fix number of iterations (26.4 for SGP vs 25.8 for AR-SGD), and for fix runtime budget (27.5 for SGP vs 26.9 for AR-SGD).

7. Conclusion

We propose SGP and OSGP for accelerating distributed training of DNNs. We provide theoretical convergence guarantees in the smooth non-convex setting, matching known convergence rates for parallel SGD. We also empirically study the methods over several computing infrastructures, and provide assessments on image classification (ImageNet, ResNet-50) and neural machine translation (Transformer, WMT16 EN-DE) tasks. We find that SGP and OSGP can run significantly faster than parallel SGD in communication-bound settings, and can train better models in less time.

ACKNOWLEDGMENTS

We thank Shubho Sengupta and Teng Li for useful discussions and for maintaining the computing infrastructure used to conduct these experiments.

References

- Akiba, T., Suzuki, S., and Fukuda, K. Extremely large mini-batch sgd: training resnet-50 on imagenet in 15 minutes. *arXiv preprint arXiv:1711.04325*, 2017.
- Alistarh, D., Grubic, D., Li, J. Z., Tomioka, R., and Vojnovic, M. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, pp. 1709–1720, 2007.
- Assran, M. and Rabbat, M. Asynchronous subgradient-push. *arXiv preprint arXiv:1803.08950*, 2018.
- Bernstein, J., Zhao, J., Azizzadenesheli, K., and Anandkumar, A. signSGD with majority vote is communication efficient and fault tolerant. In *International Conference on Learning Representations*, 2019.
- Blot, M., Picard, D., Cord, M., and Thome, N. Gossip training for deep learning. In *NIPS Workshop on Optimization for Machine Learning*, 2016.
- Charalambous, T., Yuan, Y., Yang, T., Pan, W., Hadjicostis, C. N., and Johansson, M. Distributed finite-time average consensus in digraphs in the presence of time delays. *IEEE Transactions on Control of Network Systems*, 2(4): 370–381, 2015.
- 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*, pp. 1223–1231, 2012.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1243–1252, 2017.
- 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 preprint arXiv:1706.02677*, 2017.
- Hadjicostis, C. N. and Charalambous, T. Average consensus in the presence of delays in directed graph topologies. *IEEE Transactions on Automatic Control*, 59(3):763–768, 2014.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jakovetić, D., Xavier, J., and Moura, J. M. Fast distributed gradient methods. *IEEE Transactions on Automatic Control*, 59(5):1131–1146, 2014.
- Jia, X., Song, S., He, W., Wang, Y., Rong, H., Zhou, F., Xie, L., Guo, Z., Yang, Y., Yu, L., et al. Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes. *arXiv preprint arXiv:1807.11205*, 2018.
- Jiang, Z., Balu, A., Hegde, C., and Sarkar, S. Collaborative deep learning in fixed topology networks. In *Advances in Neural Information Processing Systems*, pp. 5904–5914, 2017.
- Jin, P. H., Yuan, Q., Iandola, F., and Keutzer, K. How to scale distributed deep learning? In *NIPS ML Systems Workshop*, 2016.
- Kempe, D., Dobra, A., and Gehrke, J. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science.*, pp. 482–491. IEEE, 2003.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- Li, M., Andersen, D. G., Park, J. W., Smola, A. J., Ahmed, A., Josifovski, V., Long, J., Shekita, E. J., and Su, B.-Y. Scaling distributed machine learning with the parameter server. In *OSDI*, volume 14, pp. 583–598, 2014.
- 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*, pp. 5330–5340, 2017.
- Lian, X., Zhang, W., Zhang, C., and Liu, J. Asynchronous decentralized parallel stochastic gradient descent. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 3049–3058, 2018.
- Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., and van der Maaten, L. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 181–196, 2018.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., and Agüera y Arcas, B. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.


- Nedić, A. and Olshevsky, A. Stochastic gradient-push for strongly convex functions on time-varying directed graphs. *IEEE Trans. Automatic Control*, 61(12):3936–3947, 2016.
- Nedić, A., Olshevsky, A., and Rabbat, M. G. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.
- Nesterov, Y. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
- Ott, M., Edunov, Sergey, G. D., and Auli, M. Scaling neural machine translation. *arXiv preprint arXiv:1806.00187*, 2018.
- Paszke, A., Chintala, S., Collobert, R., Kavukcuoglu, K., Farabet, C., Bengio, S., Melvin, I., Weston, J., and Mariethoz, J. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration, 2017.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3): 211–252, 2015.
- Seneta, E. *Non-negative Matrices and Markov Chains*. Springer, 1981.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 5998–6008, 2017.
- 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*, pp. 1509–1519, 2007.
- Wolfowitz, J. Products of indecomposable, aperiodic, stochastic matrices. *Proceedings of the American Mathematical Society*, 14(5):733–737, 1963.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- Zhang, S., Choromanska, A., and LeCun, Y. Deep learning with elastic averaged SGD. In *Advances in Neural Information Processing Systems*, pp. 685–693, 2015.

Supplementary Material

Stochastic Gradient Push for Distributed Deep Learning

A. Communication Topology

Directed exponential graph. For the SGP experiments we use a time-varying directed graph to represent the inter-node connectivity. Thinking of the nodes as being ordered sequentially, according to their rank, $0, \dots, n - 1$,³ each node periodically communicates with peers that are $2^0, 2^1, \dots, 2^{\lfloor \log_2(n-1) \rfloor}$ hops away. Fig. A.1 shows an example of a directed 8-node exponential graph. Node 0's 2^0 -hop neighbour is node 1; node 0's 2^1 -hop neighbour is node 2; and node 0's 2^2 -hop neighbour is node 4.



(a) Directed Exponential Graph highlighting node 0's out-neighbours

Figure A.1: Example of an 8-node exponential graph used in experiments

In the one-peer-per-node experiments, each node cycles through these peers, transmitting, only, to a single peer from this list at each iteration. E.g., at iteration k , all nodes transmit messages to their 2^0 -hop neighbours, at iteration $k + 1$ all nodes transmit messages to their 2^1 -hop neighbours, and so on, eventually returning to the beginning of the list before cycling through the peers again. This procedure ensures that each node only sends and receives a single message at each iteration. By using full-duplex communication, sending and receiving can happen in parallel.

In the two-peer-per-node experiments, each node cycles through the same set of peers, transmitting to two peers from the list at each iteration. E.g., at iteration k , all nodes transmit messages to their 2^0 -hop and 2^1 -hop neighbours, at iteration $k + 1$ all nodes transmit messages to their 2^1 -hop and 2^2 neighbours, and so on, eventually returning to the beginning of the list before cycling through the peers again. Similarly, at each iteration, each node also receives, in a full-duplex manner, two messages from some peers that are unknown to the receiving node ahead of time. Thereby performing the send and receive operations in parallel.

Definition of $\mathbf{P}^{(k)}$. We choose the mixing matrices such that they are column stochastic (all columns sum to 1), and conform to the graph topology described above. Recall that each node i can choose its mixing weights (i^{th} column of $\mathbf{P}^{(k)}$), independently of the other nodes in the network. To minimize the number of floating point operations in each iteration, we choose to use uniform mixing weights, meaning that nodes assign uniform message weights to all neighbours. In the one-peer-per-node experiments, each node sends a message to one neighbor, and “sends a message” to itself at every iteration, and so each column of $\mathbf{P}^{(k)}$ has exactly two non-zero entries, both of which are equal to $1/2$. The first set of non-zero entries corresponds to the diagonals. At all time steps k , the diagonal entries satisfy $p_{i,i}^{(k)} = 1/2$ for all i . The

³We use indices $0, \dots, n - 1$ rather than $1, \dots, n$ only in this section, to simplify the discussion.

second set of non-zero entries correspond to the neighbor indices. At time step k , each node sends to a neighbor that is $h_k := 2^{k \bmod \lfloor \log_2(n-1) \rfloor}$ hops away. That is, at each time step k , each node i sends a message to node $(i + h_k) \bmod n$. Thus, we get that

$$p_{j,i}^{(k)} = \begin{cases} 1/2, & \text{if } j = (i + h_k) \bmod n \\ 0, & \text{otherwise.} \end{cases}$$

Note that, with this design, in fact each node sends to one peer and receives from one peer at every iteration, so the communication load is balanced across the network.

In the two-peer-per-node experiments, the definition is similar, but now there will be three non-zero entries in each column of $\mathbf{P}^{(k)}$, all of which will be equal to $1/3$; these are the diagonal, and the entries corresponding to the two neighbors to which the node sends at that iteration. In addition, each node will send two messages and receive two messages at every iteration, so the communication load is again balanced across the network.

Undirected exponential graph. For the D-PSGD experiments we use a time-varying undirected bipartite exponential graph to represent the inter-node connectivity. Odd-numbered nodes send messages to peers that are $2^1 - 1, 2^2 - 1, \dots, 2^{\lfloor \log_2(n-1) \rfloor} - 1$ (even-numbered nodes), and wait to receive a message back in return. Each odd-numbered node cycles through the peers in the list in a similar fashion to the one-peer-per-node SGP experiments. Even-numbered nodes wait to receive a message from some peer (unknown to the receiving node ahead of time), and send a message back in return.

We adopt these graphs to be consistent with the experimental setup used in Lian et al. (2017) and Lian et al. (2018).

Note also that these graphs are all regular, in that all nodes have the same number of in-coming and out-going connections.

Decentralized averaging errors. To further motivate our choice of using the directed exponential graph with SGP, let us forget about optimization for a moment and focus on the problem of distributed averaging, described in Section 2, using the PUSHSUM algorithm. Recall that each node i starts with a vector $\mathbf{y}_i^{(0)}$, and the goal of the agents is to compute the average $\bar{\mathbf{y}} = \frac{1}{n} \sum_i \mathbf{y}_i^{(0)}$. Then, since $\mathbf{y}_i^{(k+1)} = \sum_{j=1}^n p_{i,j}^{(k)} \mathbf{y}_j^{(k)}$, after k steps we have

$$\mathbf{Y}^{(k)} = \mathbf{P}^{(k-1)} \mathbf{P}^{(k-2)} \dots \mathbf{P}^{(1)} \mathbf{P}^{(0)} \mathbf{Y}^{(0)},$$

where $\mathbf{Y}^{(k)}$ is a $n \times d$ matrix with $\mathbf{y}_i^{(k)}$ as its i th row.

Let $\mathbf{P}^{(k-1:0)} = \mathbf{P}^{(k-1)} \mathbf{P}^{(k-2)} \dots \mathbf{P}^{(1)} \mathbf{P}^{(0)}$. The worst-case rate of convergence can be related to the second-largest singular value of $\mathbf{P}^{(k-1:0)}$ (Nedić et al., 2018). In particular, after k iterations we have

$$\sum_i \|\mathbf{y}_i^{(k)} - \bar{\mathbf{y}}\|_2^2 \leq \lambda_2(\mathbf{P}^{(k-1:0)}) \sum_i \|\mathbf{y}_i^{(0)} - \bar{\mathbf{y}}\|_2^2,$$

where $\lambda_2(\mathbf{P}^{(k-1:0)})$ denotes the second largest singular value of $\mathbf{P}^{(k-1:0)}$.

For the scheme proposed above, cycling deterministically through neighbors in the directed exponential graph, one can verify that after $k = \lfloor \log_2(n-1) \rfloor$ iterations, we have $\lambda_2(\mathbf{P}^{(k-1:0)}) = 0$, so all nodes exactly have the average. Intuitively, this happens because the directed exponential graph has excellent mixing properties: from any starting node in the network, one can get to any other node in at most $\log_2(n)$ hops. For $n = 32$ nodes, after 5 iterations averaging has converged using this strategy. In comparison, if one were to cycle through edges of the complete graph (where every node is connected to every other node), then for $n = 32$, after 5 consecutive iterations one would have still have $\lambda_2(\mathbf{P}^{(k-1:0)}) \approx 0.6$; i.e., nodes could be much further from the average (and hence, much less well-synchronized).

Similarly, one could consider designing the matrices $\mathbf{P}^{(k)}$ in a stochastic manner, where each node randomly samples one neighbor to send to at every iteration. If each node samples a destination uniformly from its set of neighbors in the directed exponential graph, then $\mathbb{E} \lambda_2(\mathbf{P}^{(k-1:0)}) \approx 0.4$, and if each node randomly selected a destination uniformly among all other nodes in the network (i.e., randomly from neighbors in the complete graph), then $\mathbb{E} \lambda_2(\mathbf{P}^{(k-1:0)}) \approx 0.2$. Thus, random schemes are still not as effective at quickly averaging as deterministically cycling through neighbors in the directed exponential graph. Moreover, with randomized schemes, we are no longer guaranteed that each node receives the same number of messages at every iteration, so the communication load will not be balanced as in the deterministic scheme.

The above discussion focused only on approximate distributed averaging, which is a key step within decentralized optimization. When averaging occurs less quickly, this also impacts optimization. Specifically, since nodes are less well-synchronized

Algorithm 2 Overlap Stochastic Gradient Push (SGP)

Require: Initialize $\tau \geq 0$, `count_since_last` = 0, $\gamma > 0$, $\mathbf{x}_i^{(0)} = \mathbf{z}_i^{(0)} \in \mathbb{R}^d$ and $w_i^{(0)} = 1$ for all nodes $i \in \{1, 2, \dots, n\}$

- 1: **for** $k = 0, 1, 2, \dots, K$, at node i , **do**
- 2: Sample new mini-batch $\xi_i^{(k)} \sim \mathcal{D}_i$ from local distribution
- 3: Compute mini-batch gradient at $\mathbf{z}_i^{(k)}$: $\nabla \mathbf{F}_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 4: $\mathbf{x}_i^{(k+\frac{1}{2})} = \mathbf{x}_i^{(k)} - \gamma \nabla \mathbf{F}_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 5: **if** $k \bmod \tau = 0$ **then**
- 6: Non-blocking send $(p_{j,i}^{(k)} \mathbf{x}_i^{(k+\frac{1}{2})}, p_{j,i}^{(k)} w_i^{(k)})$ to out-neighbors
- 7: $\mathbf{x}_i^{(k+1)} = p_{i,i} \mathbf{x}_i^{(k+\frac{1}{2})}$
- 8: $w_i^{(k+1)} = p_{i,i} w_i^{(k)}$
- 9: **else**
- 10: $\mathbf{x}_i^{(k+1)} = \mathbf{x}_i^{(k+\frac{1}{2})}$
- 11: $w_i^{(k+1)} = w_i^{(k)}$
- 12: **end if**
- 13: **if** `count_since_last` = τ **then**
- 14: Block until $(p_{i,j}^{(k-\tau)} \mathbf{x}_j^{(k-\tau+\frac{1}{2})}, p_{i,j}^{(k-\tau)} w_j^{(k-\tau)})$ is received for all in-neighbors j
- 15: `count_since_last` $\leftarrow 0$
- 16: **else**
- 17: `count_since_last` \leftarrow `count_since_last` + 1
- 18: **end if**
- 19: **if** Receive buffer non-empty **then**
- 20: **for** $(p_{i,j}^{(k')} \mathbf{x}_j^{(k'+\frac{1}{2})}, p_{i,j}^{(k')} w_j^{(k')})$ in the receive buffer **do**
- 21: $\mathbf{x}_i^{(k+1)} \leftarrow \mathbf{x}_i^{(k+1)} + p_{i,j}^{(k')} \mathbf{x}_j^{(k'+\frac{1}{2})}$
- 22: $w_i^{(k+1)} \leftarrow p_{i,j}^{(k')} w_j^{(k')}$
- 23: **end for**
- 24: **end if**
- 25: $\mathbf{z}_i^{(k+1)} = \mathbf{x}_i^{(k+1)} / w_i^{(k+1)}$
- 26: **end for**

(i.e., further from a consensus), each node will be evaluating its local mini-batch gradient at a different point in parameter space. Averaging these points (rather than updates based on mini-batch gradients evaluated at the same point) can be seen as injecting additional noise into the optimization process, and in our experience this can lead to worse performance in terms of train error.

B. Overlap SGP

Although SGP does not use network-wide collective communication primitives like ALLREDUCE, the implementation of Alg. 1 requires using blocking sends and receives; *i.e.*, nodes do not proceed to until they have received messages from all neighbors at that iteration. In this section we present the pseudocode of Overlap-SGP (OSGP) in Alg. 2 that overlaps gradient computation with communication to hide the communication cost. In line 25 in Algorithm 2, nodes compute the de-biased estimate of their model parameters. In lines 19 to 24, nodes aggregate all messages received in that iteration. Lines 13 to 18 ensure that the message delays are bounded, and that the nodes remain synchronized. In particular, Algorithm 2 is *synchronous* because of lines 13 to 18. If a node hasn't received a message from its in-neighbours in τ iterations, it will block and wait to received said messages. Note that if $\tau = 0$, vanilla SGP, then nodes block and wait to receive all incoming messages in each iteration. In lines 5 to 6, nodes send messages to their neighbours every τ iterations. Once again, note that if $\tau = 0$, vanilla SGP, then nodes send messages to their neighbours every iteration. In lines 2 to 4 the nodes take a stochastic gradient step. If $\tau = 1$ (1-overlap SGP), nodes transmit messages to their neighbours in every iteration, but don't wait to receive messages until the subsequent iteration.

We provide a lot of detail in Algorithm 2 to make it easier to implement the method; however, in essence, τ -overlap SGP is simply vanilla SGP with delayed communication. *i.e.*, where nodes only send a message to their neighbours every τ iterations, and can receive messages at any time in-between communication intervals.

³We define $(k \bmod 0) := 0$.

C. Implementation Details

In all of our experiments, we minimize the number of floating-point operations performed in each iteration, k , by using the mixing weights

$$p_{j,i}^{(k)} = 1/|\mathcal{N}_i^{\text{out}(k)}|$$

for all $i, j = 1, 2, \dots, n$. In words, each node assigns mixing weights uniformly to all of its out-neighbors in each iteration. Recalling our convention that each node is an in- and out-neighbor of itself, it is easy to see that this choice of mixing-weight satisfies the column-stochasticity property. It may very well be that there is a different choice of mixing-weights that lead to better spectral properties of the gossip algorithm; however we leave this exploration for future work. We denote node i 's uniform mixing weights at iteration k by $p_i^{(k)}$ — dropping the other subscript, which identifies the receiving node.

To leverage the highly efficient NVLink interconnect within each server, we treat each machine as one node in all of our experiments. In our implementation of SGP, each node computes a local mini-batch in parallel using all 8 GPUs via a local ALLREDUCE, which is efficiently implemented via the NVIDIA Collective Communications Library. Then inter-node averaging is accomplished using PUSHSUM either over Ethernet or InfiniBand. In the InfiniBand experiments, we leverage GPUDirect to directly send/receive messages between GPUs on different nodes and avoid transferring the model back to host memory. In the Ethernet experiments this is not possible, so the model is transferred to host memory after the local ALLREDUCE, and then PUSHSUM messages are sent over Ethernet.

To maximize the utility of the resources available on each server, each node (occupying a single server exclusively) runs two threads, a gossip thread and a computation thread. The computation thread executes the main logic used to train the local model on the GPUs available to the node, while the communication thread is used for inter-node network I/O. In particular, the communication thread is used to gossip messages between nodes. When using Ethernet-based communication, the nodes communicate their parameter tensors over CPUs. When using InfiniBand-based communication, the nodes communicate their parameter tensors using GPUDirect RDMA, thereby avoiding superfluous device to pinned-memory transfers of the model parameters.

Each node initializes its model on one of its GPUs, and initializes its scalar push-sum weight to 1. At the start of training, each node also allocates a *send*- and a *receive*- communication-buffer in pinned memory on the CPU (or equivalently on a GPU in the case of GPUDirect RDMA communication).

In each iteration, the communication thread waits for the send-buffer to be filled by the computation thread; transmits the message in the send-buffer to its out-neighbours; and then aggregates any newly-received messages into the receive-buffer.

In each iteration, the computation thread blocks to retrieve the aggregated messages (in the non-overlap case) in the receive-buffer; directly adds the received parameters to its own model parameters; and directly adds the received push-sum weights to its own push-sum weight. The computation thread then converts the model parameters to the *de-biased* estimate by dividing by the push-sum weight; executes a forward-backward pass of the *de-biased model* in order to compute a stochastic mini-batch gradient; converts the model parameters back to the *biased estimate* by multiplying by the push-sum weight; and applies the newly-computed stochastic gradients to the biased model. The updated model parameters are then multiplied by the mixing weight, $p_i^{(k)}$, and asynchronously copied back into the send-buffer for use by the communication thread. The push-sum weight is also multiplied by the same mixing weight and concatenated into the send-buffer.

In short, gossip is performed on the biased model parameters (push-sum numerators); stochastic gradients are computed using the de-biased model parameters; stochastic gradients are applied back to the biased model parameters; and then the biased-model and the push-sum weight are multiplied by the same uniform mixing-weight and copied back into the send-buffer.

C.1. Hyperparameters

For the ImageNet experiments, we follow the experimental protocol described in (Goyal et al., 2017). When we “apply the stochastic gradients” to the biased model parameters, we actually carry out an SGD step with nesterov momentum (see Alg. 3). For the 32, 64, and 128 GPU experiments we use the same exact learning-rate, schedule, momentum, and weight decay as those suggested in (Goyal et al., 2017) for SGD. In particular, we use a reference learning-rate of 0.1 with respect to a 256 sample batch, and scale this linearly with the batch-size; we decay the learning-rate by a factor of 10 at epochs 30, 60, 80; we use a Nesterov momentum parameter of 0.9, and we use weight decay 0.0001.

Algorithm 3 Stochastic Gradient Push with Momentum

Require: Initialize $\gamma > 0$, $m \in (0, 1)$, $\mathbf{x}_i^{(0)} = \mathbf{z}_i^{(0)} \in \mathbb{R}^d$ and $w_i^{(0)} = 1$ for all nodes $i \in [n]$

- 1: **for** $k = 0, 1, 2, \dots, K$, at node i , **do**
- 2: Sample new mini-batch $\xi_i^{(k)} \sim \mathcal{D}_i$ from local distribution
- 3: Compute mini-batch gradient at $\mathbf{z}_i^{(k)}$: $\nabla \mathbf{F}_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 4: $\mathbf{u}_i^{(k+1)} = m\mathbf{u}_i^{(k)} + \nabla \mathbf{F}_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})$
- 5: $\mathbf{x}_i^{(k+\frac{1}{2})} = \mathbf{x}_i^{(k)} - \gamma(m\mathbf{u}_i^{(k+1)} + \nabla \mathbf{F}_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)}))$
- 6: Send $(p_{j,i}^{(k)} \mathbf{x}_i^{(k+\frac{1}{2})}, p_{j,i}^{(k)} w_i^{(k)})$ to out-neighbors;
receive $(p_{i,j}^{(k)} \mathbf{x}_j^{(k+\frac{1}{2})}, p_{i,j}^{(k)} w_j^{(k)})$ from in-neighbors
- 7: $\mathbf{x}_i^{(k+1)} = \sum_{j \in \mathcal{N}_i^{\text{in}}(k)} p_{i,j}^{(k)} \mathbf{x}_j^{(k+\frac{1}{2})}$
- 8: $w_i^{(k+1)} = \sum_{j \in \mathcal{N}_i^{\text{in}}(k)} p_{i,j}^{(k)} w_j^{(k)}$
- 9: $\mathbf{z}_i^{(k+1)} = \mathbf{x}_i^{(k+1)} / w_i^{(k+1)}$
- 10: **end for**

For the machine translation experiment, we follow (Vaswani et al., 2017) and combine Stochastic Gradient Push with the Adam preconditioner. In particular, we make use of the FAIRSEQ code (Gehring et al., 2017), and train the transformer networks via SGP by replacing the built-in PyTorch parallel SGD model wrapper with our SGP model wrapper.

D. Additional Experiments

D.1. Additional Training Curves




Figure D.1: Training on Ethernet 10Gbit/s




Figure D.2: Training/Validation accuracy per iteration for SGP (Ethernet 10Gbit/s). Each time we double the number of node in the network, we half the total number of iterations.

The curves in Figure D.1 show the time-wise train- and validation-accuracies for the different runs performed on Ethernet 10Gbit/s. Figure D.2 reports the iteration-wise training and validation accuracy of SGP when using 10Gbits/s Ethernet.

D.2. Discrepancy across different nodes




Figure D.3: Resnet50, trained with SGP, training and validation errors for 4 and 32 nodes experiments. The solid and dashed lines in each figure show the mean training and validation error, respectively, over all nodes. The shaded region shows the maximum and minimum error attained at different nodes in the same experiment. Although there is non-trivial variability across nodes early in training, all nodes eventually converge to similar validation errors, achieving consensus in the sense that they represent the same function.

Here, we investigate the performance variability across nodes during training for SGP. In figure D.3, we report the minimum, maximum and mean error across the different nodes for training and validation. In an initial training phase, we observe that nodes have different validation errors; their local copies of the Resnet-50 model diverge. As we decrease the learning, the variability between the different nodes diminish and the nodes eventually converging to similar errors. This suggests that all models ultimately represent the same function, achieving consensus.

D.3. SGP Scaling Analysis




Figure D.4: SGP throughput on Ethernet (a) and InfiniBand (b). SGP exhibits 88.6% scaling efficiency on Ethernet 10Gbit/s and 92.4% on InfiniBand. Comparison of SGD vs SGP throughput in Figure (c) shows that SGP exhibit better scaling and is more robust to high-latency interconnect.

Figure D.4 highlights SGP input images throughput as we scale up the number of cluster node on both Ethernet 10Gbit/s and Infiniband 100Gbit/s. SGP exhibits 88.6% scaling efficiency on Ethernet 10Gbit/s and 92.4% on InfiniBand and stay close to the ideal scaling in both cases. In addition Figure (c) shows that SGP exhibit better scaling as we increase the network size and is more robust to high-latency interconnect.




Figure E.1: (a) Example of a 4-node network, with mixing-weights drawn on edges. (b) Example of a 4-node network, augmented with virtual nodes and edges, with mixing-weights draw on edges. The virtual nodes/edges are used to model the fact that messages from node 3 to node 0 can experience a delay of at most 2 iterations. In this particular example, we model the fact that node 3 sends a message to node 0 with a delay of 2 iterations. All virtual nodes always forward all of their messages to their out-neighbor.

E. Proofs of Theoretical Guarantees

Our convergence rate analysis is divided into three main parts. In the first one (subsection E.1) we present upper bounds for three important expressions that appear in our computations. In subsection E.2 we focus on proving the important for our analysis Lemma 8 based on which we later build the proofs of our main Theorems. Finally in the third part (subsection E.3) we provide the proofs for Theorems 1 and 2.

Preliminary results. In our analysis, two preliminary results are extensively used. We state them here for future reference.

- Let $a, b \in \mathbb{R}$. Since $(a - b)^2 \geq 0$, it holds that

$$2ab \leq a^2 + b^2. \quad (4)$$

Thus, $\|\mathbf{x}\| \|\mathbf{y}\| \leq (\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2)/2$.

- Let $r \in (0, 1)$ then from the summation of geometric sequence and for any $K \leq \infty$ it holds that

$$\sum_{k=0}^K r^k \leq \sum_{k=0}^{\infty} r^k = \frac{1}{1-r}. \quad (5)$$

Modeling message delays. To model message delays we follow the procedure used in Assran & Rabbat (2018) (which we will reiterate here). In essence, we augment the communication topology (and the mixing matrices) with virtual nodes that store messages that were transmitted, but not yet received. Similar graph augmentations have been used in Charalambous et al. (2015) and Hadjicostis & Charalambous (2014).

We commence by presenting a brief example of the delay-model before formalizing the discussion. Figure E.1 (a) shows an example of a 4-node network at some arbitrary iteration k . Suppose each node communicates to each of its out-neighbors with uniform mixing weights. These mixing weights are labeled on the corresponding edges in Figure E.1 (a). Then, the

mixing matrix $\mathbf{P}^{(k)} \in \mathbb{R}^{4 \times 4}$ is given by

$$\mathbf{P}^{(k)} = \begin{pmatrix} 1/2 & 0 & 0 & 1/3 \\ 1/2 & 1/2 & 0 & 1/3 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 1/2 & 1/3 \end{pmatrix}.$$

Column indices correspond to sending nodes, and row indices correspond to receiving nodes. Recall that sending nodes choose the mixing weights (columns of $\mathbf{P}^{(k)}$) used to pre-weight outgoing messages. Note that the matrix $\mathbf{P}^{(k)}$ is column stochastic (all columns sum to 1) — the crucial requirement of our analysis. Thus at time $k + 1$, we have the following parameter updates

$$\begin{aligned} \mathbf{x}_0^{(k+1)} &= \frac{1}{2}\mathbf{x}_0^{(k)} + \frac{1}{3}\mathbf{x}_3^{(k)} \\ \mathbf{x}_1^{(k+1)} &= \frac{1}{2}\mathbf{x}_0^{(k)} + \frac{1}{2}\mathbf{x}_1^{(k)} + \frac{1}{3}\mathbf{x}_3^{(k)} \\ \mathbf{x}_2^{(k+1)} &= \frac{1}{2}\mathbf{x}_1^{(k)} + \frac{1}{2}\mathbf{x}_2^{(k)} \\ \mathbf{x}_3^{(k+1)} &= \frac{1}{2}\mathbf{x}_2^{(k)} + \frac{1}{3}\mathbf{x}_3^{(k)}. \end{aligned}$$

In particular, each node updates its variables with the most recent information from its in-neighbours. Similar equations can be written for the push-sum weights w .

Now suppose that node 3 sends messages to its neighbors, nodes 0 and 1, at iteration k , but the message to node 0 doesn't arrive until iteration $k + 2$. To model this delay, we augment the graph topology with virtual nodes $0_1, 0_2$ (cf. Figure E.1 (b)). The virtual nodes are initialized with parameters $\mathbf{x}^{(0)} = \mathbf{0}$ and push-sum weight $w^{(0)} = 0$. Given this model, node 3 can send its pre-weighted message to virtual node 0_2 (instead of node 0) at iteration k , while the rest of communication proceeds business as usual. At the subsequent iteration, $k + 1$, node 0_2 forwards this message to node 0_1 . Subsequently, at iteration, $k + 2$, node 0_1 forwards this message to node 0, thereby modeling a 2-iteration message delay. The corresponding mixing matrix at iteration k is given by

$$\mathbf{P}^{(k)} = \begin{array}{c} \begin{array}{c} 0_1 \quad 0_2 \\ \left(\begin{array}{cccc|cc} 1/2 & 0 & 0 & 1/2 & 1 & 0 \\ 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 1/3 & 0 & 0 & 0 \\ 0 & 0 & 1/3 & 1/2 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1/3 & 0 & 0 & 0 \end{array} \right) \end{array} \end{array}.$$

Note that we have added two extra rows and columns corresponding to the virtual nodes 0_1 and 0_2 . As intended, node 2 sends a message to node 0_2 (instead of node 0) at iteration k . Node 0_2 always forwards any and all information it receives to node 0_1 , and node 0_1 always forwards any and all information it receives to node 0. Since all virtual nodes are initialized with parameters $\mathbf{x}^{(k)} = \mathbf{0}$ and push-sum weight $w^{(0)} = 0$, they do not have any impact on the final consensus value. The sole purpose of the virtual nodes is to store messages that are in-transit (transmitted but not yet received).

If the message delays at every node are upper-bounded by τ , then we can generalize this procedure, and add τ virtual nodes for every (non-virtual) node in the network. Thus, the augmented graph has $n(\tau + 1)$ nodes in total. The corresponding

augmented mixing matrix, $\mathbf{P}^{(k)} \in \mathbb{R}^{n(\tau+1) \times n(\tau+1)}$, in block matrix form is written as

$$\mathbf{P}^{(k)} = \begin{pmatrix} & (0_1, 1_1 \dots) & (0_2, 1_2, \dots) & & (0_\tau, 1_\tau, \dots) \\ \tilde{\mathbf{P}}_0^{(k)} & \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \tilde{\mathbf{P}}_1^{(k)} & \mathbf{0} & \mathbf{I} & & \vdots \\ \vdots & \vdots & & \ddots & \mathbf{0} \\ \tilde{\mathbf{P}}_{\tau-1}^{(k)} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{I} \\ \tilde{\mathbf{P}}_\tau^{(k)} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

where each block is of size $n \times n$. In particular, if node i sends a message to node j with weight $p_{j,i}^{(k)}$ at iteration k , and that message is received with delay r (i.e., received at iteration $k+r$), then

$$[\mathbf{P}_r^{(k)}]_{j,i} = p_{j,i}^{(k)},$$

otherwise

$$[\mathbf{P}_r^{(k)}]_{j,i} = 0.$$

The off-diagonal of block identity matrices \mathbf{I} denote the fact that the virtual nodes always forward all of their messages to the next node in the delay daisy-chain. It is straightforward to verify that these augmented mixing matrices are still column stochastic at all iterations k . We refer the curious reader to [Assran & Rabbat \(2018\)](#); [Charalambous et al. \(2015\)](#); [Hadjicostis & Charalambous \(2014\)](#) for a deeper discussion of the augmented delay model.

Matrix Representation. In Algorithm 1, SGP was presented from node i 's perspective (for all $i \in [n]$). However, we can actually write the SGP update at each iteration from a global viewpoint. To see this, first define the following matrices, for all $r = 1, 2, \dots, \tau$,

$$\mathbf{X}_r^{(k)} = [\mathbf{x}_{1_r}^{(k)}, \mathbf{x}_{2_r}^{(k)}, \dots, \mathbf{x}_{n_r}^{(k)}] \in \mathbb{R}^{d \times n}.$$

The matrix $\mathbf{X}_r^{(k)}$ denotes a concatenation of all the delay- r nodes' parameters at iteration k . For the purpose of notational consistency, we let the matrix $\mathbf{X}_0^{(k)}$ denote the concatenation of all the non-virtual nodes' parameters. We generalize this notation to other variables as well. In block-matrix form, we can define the augmented parameter matrix

$$\mathbf{X}^{(k)} = [\mathbf{X}_0^{(k)}, \mathbf{X}_1^{(k)}, \dots, \mathbf{X}_\tau^{(k)}] \in \mathbb{R}^{d \times n(\tau+1)},$$

which denotes a concatenation of *all* (virtual and non-virtual) nodes' parameters at iteration k . Recall that the we initialize all virtual nodes with parameters $\mathbf{x}^{(k)} = \mathbf{0}$ and push-sum weight $w^{(0)} = 0$. Additionally, since the virtual nodes are only used to model delays, and do not compute any gradient updates, we use the convention that $\mathbf{z}^{(k)} = \mathbf{0}$, $\boldsymbol{\xi}^{(k)} = \mathbf{0}$, and $\nabla F(\mathbf{z}^{(k)}; \boldsymbol{\xi}^{(k)}) = \mathbf{0}$ for all virtual nodes at all times k . Therefore, we define the augmented de-biased parameter matrix and stochastic-seed matrix as follows

$$\mathbf{Z}^{(k)} = [\mathbf{Z}_0^{(k)}, \mathbf{0}, \dots, \mathbf{0}] \in \mathbb{R}^{d \times n(\tau+1)}; \quad \boldsymbol{\xi}^{(k)} = [\boldsymbol{\xi}_0^{(k)}, \mathbf{0}, \dots, \mathbf{0}] \in \mathbb{R}^{n(\tau+1)}.$$

Similarly, we define the augmented stochastic-gradient matrix as

$$\nabla \mathbf{F}(\mathbf{Z}^{(k)}; \boldsymbol{\xi}^{(k)}) = [\nabla F_0(\mathbf{Z}_0^{(k)}; \boldsymbol{\xi}_0^{(k)}), \mathbf{0}, \dots, \mathbf{0}] \in \mathbb{R}^{d \times n(\tau+1)},$$

where the block matrix $\nabla F_0(\mathbf{Z}_0^{(k)}; \boldsymbol{\xi}_0^{(k)})$ denotes the concatenation of all non-virtual nodes' stochastic gradients at iteration k . Precisely

$$\nabla F_0(\mathbf{Z}_0^{(k)}; \boldsymbol{\xi}_0^{(k)}) = [\nabla F_1(\mathbf{z}_1^{(k)}; \boldsymbol{\xi}_1^{(k)}), \nabla F_2(\mathbf{z}_2^{(k)}; \boldsymbol{\xi}_2^{(k)}), \dots, \nabla F_n(\mathbf{z}_n^{(k)}; \boldsymbol{\xi}_n^{(k)})] \in \mathbb{R}^{d \times n}.$$

We also define the augmented expected gradient matrix (with respect to local node data distributions) as

$$\nabla \mathbf{F}(\mathbf{Z}^{(k)}) = [\nabla F_0(\mathbf{Z}_0^{(k)}), \mathbf{0}, \dots, \mathbf{0}] \in \mathbb{R}^{d \times n(\tau+1)},$$

where the block matrix $\nabla \mathbf{F}_0(\mathbf{Z}_0^{(k)})$ denotes the concatenation of all non-virtual nodes' expected stochastic gradients at iteration k . Precisely

$$\nabla \mathbf{F}_0(\mathbf{Z}_0^{(k)}) = \left[\mathbb{E}_{\xi_1^{(k)} \sim \mathcal{D}_1} [\nabla F_1(\mathbf{z}_1^{(k)}; \xi_1^{(k)})], \mathbb{E}_{\xi_2^{(k)} \sim \mathcal{D}_2} [\nabla F_2(\mathbf{z}_2^{(k)}; \xi_2^{(k)})], \dots, \mathbb{E}_{\xi_n^{(k)} \sim \mathcal{D}_n} [\nabla F_n(\mathbf{z}_n^{(k)}; \xi_n^{(k)})] \right] \in \mathbb{R}^{d \times n}.$$

For notational convenience, we simply write $\nabla f_i(\mathbf{z}_i^{(k)}) := \mathbb{E}_{\xi_i^{(k)} \sim \mathcal{D}_i} [\nabla F_i(\mathbf{z}_i^{(k)}; \xi_i^{(k)})]$. Using the above matrices, the 6^{th} step of SGP in Algorithm 1 (lines 19 to 24 in OSGP Algorithm 2) can be expressed from a global perspective as follows

$$\mathbf{X}^{(k+1)} = \left(\mathbf{X}^{(k)} - \gamma \nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) \right) [\mathbf{P}^{(k)}]^T, \quad (6)$$

where $[\mathbf{P}^{(k)}]^T \in \mathbb{R}^{n(\tau+1) \times n(\tau+1)}$ is the transpose of the augmented mixing matrix.

Lastly, let $\bar{n} := n(\tau + 1)$, and let $\bar{\mathbf{x}}^{(k)} = (1/n)\mathbf{X}^{(k)}\mathbf{1}_{\bar{n}}$ denote the average of all nodes' parameters at iteration k . Note that this definition incorporates parameters that are in-transit.

Bound for the mixing matrices. Next we state a known result from the control literature studying gossip-based optimization which allows us to bound the distance between the de-biased parameters at each node and the node-wise average.

Recall that we have assumed that the sequence of communication topologies is B -strongly connected. A directed graph is called *strongly connected* if every pair of vertices is connected with a directed path (*i.e.*, following the direction of edges). A sequence of directed graphs is called B -strongly connected if the graph with edge set $\bigcup_{k=lB}^{(l+1)B-1} E^{(k)}$ is strongly connected, for every $l \geq 0$. Recall that we have also assumed that the upper bound on the message delays is τ iterations. In particular, we assume all messages reach their destination within τ -iterations from transmission. *i.e.*, a message in-transit does not get dropped when the communication topology changes.

If the maximum message delay is τ , and all non-zero mixing weights are at least ϵ large, and the diameter of the graph with edge set $\bigcup_{k=lB}^{(l+1)B-1} E^{(k)}$ has diameter at most Δ , then the product

$$\mathbf{A}^{(k)} := \mathbf{P}^{(k+(\tau+1)\Delta B-1)} \dots \mathbf{P}^{(k+1)} \mathbf{P}^{(k)}$$

has no non-zero entries in the first n -rows (corresponding to non-virtual agents). Moreover, every entry in the first n -rows of $\mathbf{A}^{(k)}$ is at least $\epsilon^{(\tau+1)\Delta B}$.

If we further assume that all nodes have at most D out-neighbors in any iteration, and that all nodes always assign mixing weights uniformly, then $\epsilon = D^{-1}$, and every entry in the first n -rows of $\mathbf{A}^{(k)}$ is at least $D^{-(\tau+1)\Delta B}$.

Lemma 3. *Suppose that Assumption 3 (mixing connectivity) holds. Let $\lambda = 1 - nD^{-(\tau+1)\Delta B}$ and let $q = \lambda^{1/((\tau+1)\Delta B+1)}$. Then there exists a constant*

$$C < \frac{2\sqrt{d}D^{(\tau+1)\Delta B}}{\lambda^{\frac{(\tau+1)\Delta B+2}{(\tau+1)\Delta B+1}}},$$

where d is the dimension of $\bar{\mathbf{x}}^{(k)}$, $z_i^{(k)}$, and $x_i^{(0)}$, such that, for all $i = 1, 2, \dots, n$ (non-virtual nodes) and $k \geq 0$,

$$\left\| \bar{\mathbf{x}}^{(k)} - z_i^{(k)} \right\|_2 \leq Cq^k \left\| x_i^{(0)} \right\|_2 + \gamma C \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(z_i^{(s)}; \xi_i^{(s)}) \right\|_2.$$

This particular lemma follows after a small adaptation to Theorem 1 in Assran & Rabbat (2018) and its proof is based on Wolfowitz (1963). Similar bounds appear in a variety of other papers, including Nedić & Olshevsky (2016).

E.1. Important Upper Bounds

Lemma 4 (Bound of stochastic gradient). *We have the following inequality under Assumptions 1 and 2:*

$$\mathbb{E} \left\| \nabla f_i(\mathbf{z}_i^{(k)}) \right\|^2 \leq 3L^2 \mathbb{E} \left\| \mathbf{z}_i^{(k)} - \bar{\mathbf{x}}^{(k)} \right\|^2 + 3\zeta^2 + 3\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2$$

Proof.

$$\begin{aligned}
 \mathbb{E} \left\| \nabla f_i(\mathbf{z}_i^{(k)}) \right\|^2 &\leq 3\mathbb{E} \left\| \nabla f_i(\mathbf{z}_i^{(k)}) - \nabla f_i(\bar{\mathbf{x}}^{(k)}) \right\|^2 + 3\mathbb{E} \left\| \nabla f_i(\bar{\mathbf{x}}^{(k)}) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + 3\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\stackrel{\text{L-smooth}}{\leq} 3L^2 \mathbb{E} \left\| \mathbf{z}_i^{(k)} - \bar{\mathbf{x}}^{(k)} \right\|^2 + 3\mathbb{E} \left\| \nabla f_i(\bar{\mathbf{x}}^{(k)}) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + 3\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\stackrel{\text{Bounded Variance}}{\leq} 3L^2 \mathbb{E} \left\| \mathbf{z}_i^{(k)} - \bar{\mathbf{x}}^{(k)} \right\|^2 + 3\zeta^2 + 3\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2
 \end{aligned}$$

□

Lemma 5. *Let Assumptions 1-3 hold. Then,*

$$\begin{aligned}
 Q_i^{(k)} = \mathbb{E} \left\| \bar{\mathbf{x}}^{(k)} - \mathbf{z}_i^{(k)} \right\|^2 &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2} + \gamma \frac{q^k C^2}{1-q} \right) \sigma^2 + \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \gamma \frac{q^k 3C^2}{1-q} \right) \zeta^2 \\
 &+ \left(\gamma^2 \frac{12L^2 C^2}{1-q} + \gamma q^k 3L^2 C^2 \right) \sum_{j=0}^k q^{k-j} Q_i^{(j)} \\
 &+ \left(\gamma^2 \frac{12C^2}{1-q} + \gamma q^k 3C^2 \right) \sum_{j=0}^k q^{k-j} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(j)}) \right\|^2 \\
 &+ \left(q^{2k} C^2 + \gamma q^k \frac{2C^2}{1-q} \right) \left\| \mathbf{x}_i^{(0)} \right\|^2.
 \end{aligned} \tag{7}$$

Proof.

$$\begin{aligned}
 Q_i^{(k)} &= \mathbb{E} \left\| \bar{\mathbf{x}}^{(k)} - \mathbf{z}_i^{(k)} \right\|^2 \\
 &\stackrel{\text{Lemma 3}}{\leq} \mathbb{E} \left(Cq^k \left\| \mathbf{x}_i^{(0)} \right\| + \gamma C \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) \right\| \right)^2 \\
 &= \mathbb{E} \left(Cq^k \left\| \mathbf{x}_i^{(0)} \right\| + \gamma C \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) + \nabla f_i(\mathbf{z}_i^{(s)}) \right\| \right)^2 \\
 &\leq \mathbb{E} \left(\underbrace{Cq^k \left\| \mathbf{x}_i^{(0)} \right\|}_a + \underbrace{\gamma C \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|}_b + \underbrace{\gamma C \sum_{s=0}^k q^{k-s} \left\| \nabla f_i(\mathbf{z}_i^{(s)}) \right\|}_c \right)^2
 \end{aligned} \tag{8}$$

Thus, using the above expressions of a , b and c we have that $Q_i^{(k)} \leq \mathbb{E}(a^2 + b^2 + c^2 + 2ab + 2bc + 2ac)$. Let us now obtain bounds for all of these quantities:

$$\begin{aligned}
 a^2 &= C^2 \left\| \mathbf{x}_i^{(0)} \right\|^2 q^{2k} \\
 b^2 &= \gamma^2 C^2 \sum_{j=0}^k q^{2(k-j)} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \\
 &+ 2\gamma^2 C^2 \underbrace{\sum_{j=0}^k \sum_{s=j+1}^k q^{2k-j-s} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\| \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|}_{b_1} \\
 c^2 &= \gamma^2 C^2 \sum_{j=0}^k q^{2(k-j)} \left\| \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 + 2\gamma^2 C^2 \underbrace{\sum_{j=0}^k \sum_{s=j+1}^k q^{2k-j-s} \left\| \nabla f_i(\mathbf{z}_i^{(j)}) \right\| \left\| \nabla f_i(\mathbf{z}_i^{(s)}) \right\|}_{c_1}
 \end{aligned}$$

$$\begin{aligned}
 2ab &= 2\gamma C^2 q^k \left\| \mathbf{x}_i^{(0)} \right\| \left\| \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\| \right\| \\
 2ac &= 2\gamma C^2 q^k \left\| \mathbf{x}_i^{(0)} \right\| \left\| \sum_{s=0}^k q^{k-s} \left\| \nabla f_i(\mathbf{z}_i^{(s)}) \right\| \right\| \\
 2bc &= 2\gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-j-s} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\| \left\| \nabla f_i(\mathbf{z}_i^{(s)}) \right\|.
 \end{aligned}$$

The expression b_1 is bounded as follows:

$$\begin{aligned}
 b_1 &= \gamma^2 C^2 \sum_{j=0}^k \sum_{s=j+1}^k q^{2k-j-s} 2 \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\| \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\| \\
 &\stackrel{(4)}{\leq} \gamma^2 C^2 \sum_{j=0}^k \sum_{s=j+1}^k q^{2k-s-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \\
 &\quad + \gamma^2 C^2 \sum_{j=0}^k \sum_{s=j+1}^k q^{2k-s-j} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|^2 \\
 &\leq \gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-s-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \\
 &\quad + \gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-s-j} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|^2 \\
 &= \gamma^2 C^2 \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \sum_{s=0}^k q^{k-s} \\
 &\quad + \gamma^2 C^2 \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|^2 \sum_{j=0}^k q^{k-j} \\
 &\stackrel{(5)}{\leq} \frac{1}{1-q} \gamma^2 C^2 \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \\
 &\quad + \frac{1}{1-q} \gamma^2 C^2 \sum_{s=0}^k q^{k-s} \left\| \nabla F_i(\mathbf{z}_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(\mathbf{z}_i^{(s)}) \right\|^2 \\
 &= \frac{2}{1-q} \gamma^2 C^2 \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2. \tag{9}
 \end{aligned}$$

Thus,

$$\begin{aligned}
 b^2 &= \gamma^2 C^2 \sum_{j=0}^k q^{2(k-j)} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 + b_1 \\
 &\leq \frac{\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 + b_1 \\
 &\stackrel{(9)}{\leq} \frac{3\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(\mathbf{z}_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(\mathbf{z}_i^{(j)}) \right\|^2 \tag{10}
 \end{aligned}$$

where in the first inequality above we use the fact that for $q \in (0, 1)$, we have $q^k < \frac{1}{1-q}, \forall k > 0$.

By identical construction we have

$$c^2 \leq \frac{3\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla f_i(z_i^{(j)}) \right\|^2.$$

Now let us bound the products $2ab$, $2ac$ and $2bc$.

$$\begin{aligned} 2ab &= \gamma C^2 q^k \sum_{s=0}^k q^{k-s} 2 \left\| x_i^{(0)} \right\| \left\| \nabla F_i(z_i^{(s)}; \xi_i^{(s)}) - \nabla f_i(z_i^{(s)}) \right\| \\ &\stackrel{(4)}{\leq} \gamma C^2 q^k \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 + \gamma C^2 q^k \sum_{j=0}^k q^{k-j} \left\| x_i^{(0)} \right\|^2 \\ &\stackrel{(5)}{\leq} \gamma C^2 q^k \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 + \frac{\gamma C^2 \left\| x_i^{(0)} \right\|^2}{1-q} q^k \end{aligned} \quad (11)$$

By similar procedure,

$$2ac \leq \gamma C^2 q^k \sum_{s=0}^k q^{k-s} \left\| \nabla f_i(z_i^{(s)}) \right\|^2 + \frac{\gamma C^2 \left\| x_i^{(0)} \right\|^2}{1-q} q^k \quad (12)$$

Finally,

$$\begin{aligned} 2bc &= \gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-j-s} 2 \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\| \left\| \nabla f_i(z_i^{(s)}) \right\| \\ &\stackrel{(4)}{\leq} \gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-j-s} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 + \gamma^2 C^2 \sum_{j=0}^k \sum_{s=0}^k q^{2k-j-s} \left\| \nabla f_i(z_i^{(s)}) \right\|^2, \\ &= \gamma^2 C^2 \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 \sum_{s=0}^k q^{k-s} + \gamma^2 C^2 \sum_{s=0}^k q^{k-s} \left\| \nabla f_i(z_i^{(s)}) \right\|^2 \sum_{j=0}^k q^{k-j}, \\ &\stackrel{(5)}{\leq} \frac{\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 + \frac{\gamma^2 C^2}{1-q} \sum_{s=0}^k q^{k-s} \left\| \nabla f_i(z_i^{(s)}) \right\|^2 \end{aligned} \quad (13)$$

By combining all of the above bounds together we obtain:

$$\begin{aligned} Q_i^{(k)} &\leq \mathbb{E}(a^2 + b^2 + c^2 + 2ab + 2bc + 2ac) \\ &\leq \mathbb{E} \frac{4\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2 \\ &\quad + \mathbb{E} \frac{4\gamma^2 C^2}{1-q} \sum_{j=0}^k q^{k-j} \left\| \nabla f_i(z_i^{(j)}) \right\|^2 \\ &\quad + C^2 \left\| x_i^{(0)} \right\|^2 q^{2k} \\ &\quad + \frac{2\gamma C^2 \left\| x_i^{(0)} \right\|^2}{1-q} q^k \\ &\quad + \mathbb{E} \gamma C^2 q^k \sum_{j=0}^k q^{k-j} \left\| \nabla f_i(z_i^{(j)}) \right\|^2 \end{aligned}$$

$$+ \mathbb{E} \gamma C^2 q^k \sum_{j=0}^k q^{k-j} \left\| \nabla F_i(z_i^{(j)}; \xi_i^{(j)}) - \nabla f_i(z_i^{(j)}) \right\|^2. \quad (14)$$

After grouping terms together and using the upper bound of Lemma 4, we obtain

$$\begin{aligned} Q_i^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2} + \gamma \frac{q^k C^2}{1-q} \right) \sigma^2 + \left(q^{2k} C^2 + \gamma q^k \frac{2C^2}{1-q} \right) \left\| x_i^{(0)} \right\|^2 \\ &+ \left(\gamma^2 \frac{4C^2}{1-q} + \gamma q^k C^2 \right) \sum_{j=0}^k q^{k-j} \mathbb{E} \left\| \nabla f_i(z_i^{(j)}) \right\|^2 \\ &\stackrel{\text{Lemma 4}}{\leq} \left(\gamma^2 \frac{4C^2}{(1-q)^2} + \gamma \frac{q^k C^2}{1-q} \right) \sigma^2 + \left(q^{2k} C^2 + \gamma q^k \frac{2C^2}{1-q} \right) \left\| x_i^{(0)} \right\|^2 \\ &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \frac{\gamma q^k 3C^2}{1-q} \right) \zeta^2 \\ &+ \left(\gamma^2 \frac{12L^2 C^2}{1-q} + \gamma q^k 3L^2 C^2 \right) \sum_{j=0}^k q^{k-j} Q_i^{(j)} \\ &+ \left(\gamma^2 \frac{12C^2}{1-q} + \gamma q^k 3C^2 \right) \sum_{j=0}^k q^{k-j} \mathbb{E} \left\| \nabla f(\bar{x}^{(j)}) \right\|^2 \end{aligned} \quad (15)$$

This completes the proof. \square

Having found a bound for the quantity $Q_i^{(k)}$, let us now present a lemma for bounding the quantity $\sum_{k=0}^{K-1} M^{(k)}$ where $K > 1$ is a constant and $M^{(k)}$ is the average $Q_i^{(k)}$ across all (non-virtual) nodes $i \in [n]$. That is, $M^{(k)} = \frac{1}{n} \sum_{i=1}^n Q_i^{(k)}$.

Lemma 6. *Let Assumptions 1-3 hold and let us define $D_2 = 1 - \frac{\gamma^2 12L^2 C^2}{(1-q)^2} - \frac{\gamma 3L^2 C^2}{(1-q)^2}$. Then,*

$$\begin{aligned} \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2 D_2} \right) \sigma^2 K + \left(\gamma \frac{C^2}{(1-q)^2 D_2} \right) \sigma^2 \\ &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2 D_2} \right) \zeta^2 K + \left(\frac{\gamma 3C^2}{(1-q)^2 D_2} \right) \zeta^2 \\ &+ \left(\frac{C^2}{(1-q)^2 D_2} + \gamma \frac{2C^2}{(1-q)^2 D_2} \right) \frac{\sum_{i=1}^n \left\| x_i^{(0)} \right\|^2}{n} \\ &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2 D_2} + \gamma \frac{3C^2}{(1-q)^2 D_2} \right) \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{x}^{(k)}) \right\|^2 \end{aligned} \quad (16)$$

Proof. Using the bound for $Q_i^{(k)}$ let us first bound its average across all nodes $M^{(k)}$

$$\begin{aligned}
 M^{(k)} &= \frac{1}{n} \sum_{i=1}^n Q_i^{(k)} \\
 &\stackrel{\text{Lemma 5}}{\leq} \left(\gamma^2 \frac{4C^2}{(1-q)^2} + \gamma \frac{q^k C^2}{1-q} \right) \sigma^2 + \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \frac{\gamma q^k 3C^2}{1-q} \right) \zeta^2 \\
 &\quad + \left(\gamma^2 \frac{12C^2}{1-q} + \gamma q^k 3C^2 \right) \sum_{j=0}^k q^{k-j} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(j)}) \right\|^2 \\
 &\quad + \left(\gamma^2 \frac{12L^2 C^2}{1-q} + \gamma q^k 3L^2 C^2 \right) \sum_{j=0}^k q^{k-j} M^{(j)} \\
 &\quad + \left(q^{2k} C^2 + \gamma q^k \frac{2C^2}{1-q} \right) \frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n}. \tag{17}
 \end{aligned}$$

At this point note that for any $\lambda \in (0, 1)$, non-negative integer $K \in \mathbb{N}$, and non-negative sequence $\{\beta^{(j)}\}_{j=0}^k$, it holds that

$$\begin{aligned}
 \sum_{k=0}^K \sum_{j=0}^k \lambda^{k-j} \beta^{(j)} &= \beta^{(0)} (\lambda^K + \lambda^{K-1} + \dots + \lambda^0) + \beta^{(1)} (\lambda^{K-1} + \lambda^{K-2} + \dots + \lambda^0) + \dots + \beta^{(K)} (\lambda^0) \\
 &\leq \frac{1}{1-\lambda} \sum_{j=0}^K \beta^{(j)}. \tag{18}
 \end{aligned}$$

Similarly,

$$\sum_{k=0}^K \lambda^k \sum_{j=0}^k \lambda^{k-j} \beta^{(j)} = \sum_{k=0}^K \sum_{j=0}^k \lambda^{2k-j} \beta^{(j)} \leq \sum_{k=0}^K \sum_{j=0}^k \lambda^{2(k-j)} \beta^{(j)} \stackrel{(18)}{\leq} \frac{1}{1-\lambda^2} \sum_{j=0}^K \beta^{(j)} \tag{19}$$

Now by summing from $k = 0$ to $K - 1$ and using the bounds of (18) and (19) we obtain:

$$\begin{aligned}
 \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2} \right) \sigma^2 K + \left(\gamma \frac{C^2}{(1-q)^2} \right) \sigma^2 \\
 &\quad + \left(\gamma^2 \frac{12C^2}{(1-q)^2} \right) \zeta^2 K + \left(\frac{\gamma 3C^2}{1-q} \right) \zeta^2 \\
 &\quad + \left(\frac{C^2}{1-q^2} + \gamma \frac{2C^2}{(1-q)^2} \right) \frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n} \\
 &\quad + \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \gamma \frac{3C^2}{1-q^2} \right) \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\quad + \left(\gamma^2 \frac{12L^2 C^2}{(1-q)^2} + \gamma \frac{3L^2 C^2}{1-q^2} \right) \sum_{k=0}^{K-1} M^{(k)}.
 \end{aligned}$$

By rearranging:

$$\begin{aligned}
 \left(1 - \gamma^2 \frac{12L^2C^2}{(1-q)^2} - \gamma \frac{3L^2C^2}{1-q^2}\right) \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2}\right) \sigma^2 K + \left(\gamma \frac{C^2}{(1-q)^2}\right) \sigma^2 \\
 &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2}\right) \zeta^2 K + \left(\frac{\gamma 3C^2}{(1-q)^2}\right) \zeta^2 \\
 &+ \left(\frac{C^2}{1-q^2} + \gamma \frac{2C^2}{(1-q)^2}\right) \frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n} \\
 &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \gamma \frac{3C^2}{1-q^2}\right) \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2
 \end{aligned}$$

Note that since $q \in (0, 1)$ it holds that $\frac{1}{1-q^2} \leq \frac{1}{(1-q)^2}$.⁴ Thus,

$$\begin{aligned}
 \left(1 - \gamma^2 \frac{12L^2C^2}{(1-q)^2} - \gamma \frac{3L^2C^2}{(1-q)^2}\right) \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2}\right) \sigma^2 K + \left(\gamma \frac{C^2}{(1-q)^2}\right) \sigma^2 \\
 &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2}\right) \zeta^2 K + \left(\frac{\gamma 3C^2}{(1-q)^2}\right) \zeta^2 \\
 &+ \left(\frac{C^2}{(1-q)^2} + \gamma \frac{2C^2}{(1-q)^2}\right) \frac{\sum_{i=1}^n \|x_i^{(0)}\|^2}{n} \\
 &+ \left(\gamma^2 \frac{12C^2}{(1-q)^2} + \gamma \frac{3C^2}{(1-q)^2}\right) \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2
 \end{aligned}$$

Dividing both sides with $D_2 = 1 - \frac{\gamma^2 12L^2C^2}{(1-q)^2} - \frac{\gamma 3L^2C^2}{(1-q)^2}$ completes the proof. \square

E.2. Towards the proof of the main Theorems

The goal of this section is the presentation of Lemma 8. It is the main lemma of our convergence analysis and based on which we build the proofs of Theorems 1 and 2.

Let us first state a preliminary lemma that simplifies some of the expressions that involve expectations with respect to the random variable $\xi_i^{(t)}$.

Lemma 7. *Under the definition of our problem and the Assumptions 1-3 we have that:*

(i)

$$\mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)})}{n} \right\|^2 = \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})}{n} \right\|^2 + \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2$$

(ii)

$$\mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \left[\nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)}) \right]}{n} \right\|^2 \leq \frac{\sigma^2}{n}$$

⁴This step is used to simplified the expressions involve the parameter q . One can still obtain similar results by keeping the expression $\frac{1}{1-q^2}$ in the definition of D_2 .

Proof.

$$\begin{aligned}
 \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)})}{n} \right\|^2 &= \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})}{n} + \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \\
 &= \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})}{n} \right\|^2 \\
 &\quad + \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \\
 &\quad + 2 \left\langle \frac{\sum_{i=1}^n \mathbb{E}_{\xi_i^{(k)}} \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})}{n}, \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\rangle \\
 &= \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})}{n} \right\|^2 \\
 &\quad + \mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2. \tag{20}
 \end{aligned}$$

where in the last equality the inner product becomes zero from the fact that $\mathbb{E}_{\xi_i^{(k)}} \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) = \nabla f_i(z_i^{(k)})$.

$$\begin{aligned}
 &\mathbb{E}_{\xi_i^{(k)}} \left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \\
 &= \frac{1}{n^2} \mathbb{E}_{\xi_i^{(k)}} \left\| \sum_{i=1}^n [\nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)})] \right\|^2 \\
 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}_{\xi_i^{(k)}} \left\| \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)}) \right\|^2 \\
 &\quad + \frac{2}{n^2} \sum_{i \neq j} \left\langle \mathbb{E}_{\xi_i^{(k)}} \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)}), \mathbb{E}_{\xi_j^{(k)}} \nabla F_j(z_j^{(k)}; \xi_j^{(k)}) - \nabla f_j(z_j^{(k)}) \right\rangle \\
 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}_{\xi_i^{(k)}} \left\| \nabla F_i(z_i^{(k)}; \xi_i^{(k)}) - \nabla f_i(z_i^{(k)}) \right\|^2 \\
 &\stackrel{\text{Bounded Variance}}{\leq} \frac{1}{n^2} \sum_{i=1}^n \sigma^2 = \frac{\sigma^2}{n}, \tag{21}
 \end{aligned}$$

□

Before presenting the proof of next lemma let us define the conditional expectation

$$\mathbb{E}[\cdot | \mathcal{F}_k] := \mathbb{E}_{\xi_i^{(k)} \sim \mathcal{D}_i \forall i \in [n]}[\cdot] = \mathbb{E}_{\xi_i^{(k)} \forall i \in [n]}[\cdot].$$

The expectation in this expression is *only* with respect to the random choices $\xi_i^{(k)}$ for all nodes $i \in [n]$ at the k^{th} iteration. In addition, we should highlight that the choices of random variables $\xi_i^k \sim \mathcal{D}_i$, $\xi_j^k \sim \mathcal{D}_j$ at the step t of the algorithm, are independent for any two nodes $i \neq j \in [n]$. This is also true in the case that the two nodes follow the same distribution $\mathcal{D} = \mathcal{D}_i = \mathcal{D}_j$.

Lemma 8. *Let Assumptions 1-3 hold and let*

$$D_1 = \frac{1}{2} - \frac{L^2}{2} \left(\frac{12\gamma^2 C^2 + 3\gamma C^2}{(1-q)^2 D_2} \right) \quad \text{and} \quad D_2 = 1 - \frac{\gamma^2 12L^2 C^2}{(1-q)^2} - \frac{\gamma 3L^2 C^2}{(1-q)^2}.$$

Here $C > 0$ and $q \in (0, 1)$ are the two non-negative constants defined in Lemma 3. Let $\{\mathbf{X}_k\}_{k=0}^\infty$ be the random sequence produced by (6) (Matrix representation of Algorithm 1). Then,

$$\begin{aligned} & \frac{1}{K} \left(D_1 \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + \frac{1-L\gamma}{2} \sum_{k=0}^{K-1} \mathbb{E} \left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right) \\ & \leq \frac{f(\bar{\mathbf{x}}^{(0)}) - f^*}{\gamma K} + \frac{L\gamma\sigma^2}{2n} + \frac{4L^2\gamma^2 C^2 \sigma^2 + 12L^2\gamma^2 C^2 \zeta^2}{2(1-q)^2 D_2} + \frac{\gamma L^2 C^2 \sigma^2 + 3L^2\gamma C^2 \zeta^2}{2K(1-q)^2 D_2} \\ & \quad + \left(\frac{L^2 C + 2L^2\gamma C^2}{2(1-q)^2 D_2 K} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n}. \end{aligned}$$

Proof.

$$\begin{aligned} f(\bar{\mathbf{x}}^{(k+1)}) &= f\left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n}\right) \stackrel{(6)}{=} f\left(\frac{\mathbf{X}^{(k)} [\mathbf{P}^{(k)}]^\top \mathbf{1}_{\bar{n}} - \gamma \nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) [\mathbf{P}^{(k)}]^\top \mathbf{1}_{\bar{n}}}{n}\right) \\ &= f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} - \frac{\gamma \nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) \mathbf{1}_{\bar{n}}}{n}\right) \\ &\stackrel{L\text{-smooth}}{\leq} f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \gamma \left\langle \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right), \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\rangle \\ &\quad + \frac{L\gamma^2}{2} \left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \end{aligned} \tag{22}$$

Taking expectations of both sides conditioned on \mathcal{F}_k :

$$\begin{aligned} \mathbb{E} \left[f\left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n}\right) \middle| \mathcal{F}_k \right] &\leq f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \gamma \left\langle \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right), \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\rangle \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}, \boldsymbol{\xi}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \middle| \mathcal{F}_k \right] \\ &\stackrel{\text{Lemma 7[i]}}{=} f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \gamma \left\langle \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right), \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\rangle \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\sum_{i=1}^n \nabla F_i(z_i^{(k)}; \boldsymbol{\xi}_i^{(k)}) - \sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \middle| \mathcal{F}_k \right] \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \middle| \mathcal{F}_k \right] \\ &\stackrel{\text{Lemma 7[ii]}}{\leq} f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \gamma \left\langle \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right), \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\rangle \\ &\quad + \frac{L\gamma^2\sigma}{2n} + \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \middle| \mathcal{F}_k \right] \\ &= f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \frac{\gamma}{2} \left\| \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) \right\|^2 - \frac{\gamma}{2} \left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \\ &\quad + \frac{\gamma}{2} \left\| \nabla f\left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n}\right) - \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 + \frac{L\gamma^2\sigma^2}{2n} \\ &\quad + \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\sum_{i=1}^n \nabla f_i(z_i^{(k)})}{n} \right\|^2 \middle| \mathcal{F}_k \right] \end{aligned} \tag{23}$$

where in the last step above we simply expand the inner product.

Taking expectations with respect to \mathcal{F}_k and using the tower property, we get

$$\begin{aligned}
 \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n} \right) \right] &\leq \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] - \frac{\gamma}{2} \mathbb{E} \left[\left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right], \\
 &+ \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) - \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] + \frac{L\gamma^2\sigma^2}{2n} \\
 &+ \frac{L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\sum_{i=1}^n \nabla f_i(\mathbf{z}_i^{(k)})}{n} \right\|^2 \right] \\
 &= \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] - \frac{\gamma - L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right], \\
 &+ \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) - \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] + \frac{L\gamma^2\sigma^2}{2n} \tag{24}
 \end{aligned}$$

Let us now focus on find an upper bound for the quantity $\mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) - \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right]$.

$$\begin{aligned}
 \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) - \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] &= \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}) - \frac{\sum_{i=1}^n \nabla f_i(\mathbf{z}_i^{(k)})}{n} \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_i \nabla f_i(\bar{\mathbf{x}}) - \frac{\sum_{i=1}^n \nabla f_i(\mathbf{z}_i^{(k)})}{n} \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \frac{\sum_i \nabla f_i(\bar{\mathbf{x}}) - \sum_{i=1}^n \nabla f_i(\mathbf{z}_i^{(k)})}{n} \right\|^2 \right] \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_i [\nabla f_i(\bar{\mathbf{x}}) - \nabla f_i(\mathbf{z}_i^{(k)})] \right\|^2 \right] \\
 &\stackrel{\text{Jensen}}{\leq} \frac{1}{n} \sum_i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}) - \nabla f_i(\mathbf{z}_i^{(k)}) \right\|^2 \right] \\
 &\stackrel{L\text{-smooth}}{\leq} \frac{L^2}{n} \sum_{i=1}^n \mathbb{E} \left[\left\| \bar{\mathbf{x}} - \mathbf{z}_i^{(k)} \right\|^2 \right] \\
 &= \frac{L^2}{n} \sum_{i=1}^n Q_i^{(k)} \tag{25}
 \end{aligned}$$

Thus we have that:

$$\begin{aligned}
 \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n} \right) \right] &\leq \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] - \frac{\gamma - L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\nabla \mathbf{F}(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right], \\
 &+ \frac{\gamma L^2}{2n} \sum_{i=1}^n Q_i^{(k)} + \frac{L\gamma^2\sigma^2}{2n} \tag{26}
 \end{aligned}$$

By rearranging:

$$\begin{aligned} \frac{\gamma}{2} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] + \frac{\gamma - L\gamma^2}{2} \mathbb{E} \left[\left\| \frac{\nabla F(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] &\leq \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n} \right) \right] \\ &+ \frac{L\gamma^2 \sigma^2}{2n} + \frac{\gamma L^2}{2n} \sum_{i=1}^n Q_i^{(k)} \end{aligned} \quad (27)$$

Let us now sum from $k = 0$ to $k = K - 1$:

$$\begin{aligned} \frac{\gamma}{2} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] + \frac{\gamma - L\gamma^2}{2} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \frac{\nabla F(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] &\leq \sum_{k=0}^{K-1} \left[\mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(k+1)} \mathbf{1}_{\bar{n}}}{n} \right) \right] \right] \\ &+ \sum_{k=0}^{K-1} \frac{L\gamma^2 \sigma^2}{2n} + \frac{\gamma L^2}{2n} \sum_{k=0}^{K-1} \sum_{i=1}^n Q_i^{(k)} \\ &\leq \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(0)} \mathbf{1}_{\bar{n}}}{n} \right) \right] - \mathbb{E} \left[f \left(\frac{\mathbf{X}^{(K)} \mathbf{1}_{\bar{n}}}{n} \right) \right] \\ &+ \frac{LK\gamma^2 \sigma^2}{2n} + \frac{\gamma L^2}{2} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n Q_i^{(k)} \\ &\leq f(\bar{\mathbf{x}}^{(0)}) - f^* \\ &+ \frac{LK\gamma^2 \sigma^2}{2n} + \underbrace{\frac{\gamma L^2}{2} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n Q_i^{(k)}}_{M_k} \end{aligned} \quad (28)$$

For the last inequality above, recall that we let f^* denote the global infimum of our problem.

Using the bound for the expression $\sum_{k=0}^{K-1} M_k$ from Lemma 6 we obtain:

$$\begin{aligned} &\frac{\gamma}{2} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \nabla f \left(\frac{\mathbf{X}^{(k)} \mathbf{1}_{\bar{n}}}{n} \right) \right\|^2 \right] + \frac{\gamma - L\gamma^2}{2} \sum_{k=0}^{K-1} \mathbb{E} \left[\left\| \frac{\nabla F(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right] \\ &\leq f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{LK\gamma^2 \sigma^2}{2n} \\ &+ \frac{\gamma L^2}{2} \frac{4\gamma^2 C^2 \sigma^2 K + \gamma C^2 \sigma^2}{(1-q)^2 D_2} + \frac{\gamma L^2}{2} \frac{12\gamma^2 C^2 \zeta^2 K + 3\gamma C^2 \zeta^2}{(1-q)^2 D_2} \\ &+ \frac{\gamma L^2}{2} \left(\frac{12\gamma^2 C^2 + 3\gamma C^2}{(1-q)^2 D_2} \right) \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\ &+ \frac{\gamma L^2}{2} \left(\frac{C^2 + 2\gamma C^2}{(1-q)^2 D_2} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n}. \end{aligned}$$

By rearranging and dividing all terms by γK we obtain:

$$\begin{aligned} &\frac{1}{K} \left(\left[\frac{1}{2} - \frac{L^2}{2} \left(\frac{12\gamma^2 C^2 + 3\gamma C^2}{(1-q)^2 D_2} \right) \right] \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + \frac{1 - L\gamma}{2} \sum_{k=0}^{K-1} \mathbb{E} \left\| \frac{\nabla F(\mathbf{Z}^{(k)}) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right) \\ &\leq \frac{f(\bar{\mathbf{x}}^{(0)}) - f^*}{\gamma K} + \frac{L\gamma \sigma^2}{2n} + \frac{4L^2 \gamma^2 C^2 \sigma^2 + 12L^2 \gamma^2 C^2 \zeta^2}{2(1-q)^2 D_2} + \frac{\gamma L^2 C^2 \sigma^2 + 3L^2 \gamma C^2 \zeta^2}{2K(1-q)^2 D_2} \\ &+ \left(\frac{L^2 C^2 + 2L^2 \gamma C^2}{2(1-q)^2 D_2 K} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n}. \end{aligned}$$

By defining $D_1 = \left[\frac{1}{2} - \frac{L^2}{2} \left(\frac{12\gamma^2 C^2 + 3\gamma C^2}{(1-q)^2 D_2} \right) \right]$ the proof is complete. \square

E.3. Proofs of Main Theorems

Having present all of the above Lemmas we are now ready to provide the proofs of main Theorems 1 and 2.

E.3.1. PROOF OF THEOREM 1

Let $\gamma \leq \min \left\{ \frac{(1-q)^2}{60L^2C^2}, 1 \right\}$. Then:

$$D_2 = 1 - \frac{\gamma^2 12L^2C^2}{(1-q)^2} - \frac{\gamma 3L^2C^2}{(1-q)^2} \stackrel{(\gamma^2 < \gamma)}{\geq} 1 - \frac{\gamma 15L^2C^2}{(1-q)^2} \geq 1 - \frac{1}{4} \geq \frac{1}{2}$$

and

$$D_1 = \frac{1}{2} - \frac{L^2}{2} \left(\frac{12\gamma^2 C^2 + 3\gamma C^2}{(1-q)^2 D_2} \right) \stackrel{(\gamma^2 < \gamma)}{\geq} \frac{1}{2} - \frac{15\gamma C^2 L^2}{2(1-q)^2 D_2} \geq \frac{1}{2} - \frac{1}{8D_2} \geq \frac{1}{4}$$

By substituting the above bounds into the result of Lemma 8 and by removing the second term of left hand side we obtain:

$$\begin{aligned} \frac{1}{4} \frac{\sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2}{K} &= \frac{1}{K} \left(\frac{1}{4} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + \frac{1-L\gamma}{2} \sum_{k=0}^{K-1} \mathbb{E} \left\| \frac{\nabla F(\mathbf{Z}_k) \mathbf{1}_{\bar{n}}}{n} \right\|^2 \right) \\ &\leq \frac{f(\bar{\mathbf{x}}^{(0)}) - f^*}{\gamma K} + \frac{L\gamma\sigma^2}{2n} + \frac{4L^2\gamma^2 C^2 \sigma^2 + 12L^2\gamma^2 C^2 \zeta^2}{(1-q)^2} + \frac{\gamma L^2 C^2 \sigma^2 + 3L^2\gamma C^2 \zeta^2}{K(1-q)^2} \\ &\quad + \left(\frac{L^2 C + 2L^2\gamma C^2}{(1-q)^2 K} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \end{aligned} \quad (29)$$

Let us now substitute in the above expression $\gamma = \sqrt{\frac{n}{K}}$. This can be done due to the lower bound (see equation 3) on the total number of iterations K where guarantees that $\sqrt{\frac{n}{K}} \leq \min \left\{ \frac{(1-q)^2}{60L^2C^2}, 1 \right\}$.

$$\begin{aligned} \frac{1}{4} \frac{\sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2}{K} &\leq \frac{f(\bar{\mathbf{x}}^{(0)}) - f^*}{\gamma K} + \frac{L\gamma\sigma^2}{2n} + \gamma^2 \frac{4L^2 C^2 \sigma^2 + 12L^2 C^2 \zeta^2}{(1-q)^2} + \gamma \frac{L^2 C^2 \sigma^2 + 3L^2 C^2 \zeta^2}{K(1-q)^2} \\ &\quad + \frac{L^2 C}{(1-q)^2 K} \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} + \gamma \frac{2L^2 C^2}{(1-q)^2 K} \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \\ &\stackrel{\gamma = \sqrt{\frac{n}{K}}}{=} \frac{f(\bar{\mathbf{x}}^{(0)}) - f^*}{\sqrt{nK}} + \frac{L\sigma^2}{2\sqrt{nK}} + \frac{n}{K} \frac{4L^2 C^2 \sigma^2 + 12L^2 C^2 \zeta^2}{(1-q)^2} + \sqrt{\frac{n}{K}} \frac{L^2 C^2 \sigma^2 + 3L^2 C^2 \zeta^2}{K(1-q)^2} \\ &\quad + \frac{L^2 C^2}{(1-q)^2 K} \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} + \sqrt{\frac{n}{K}} \frac{2L^2 C^2}{(1-q)^2 K} \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \\ &= \frac{f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L}{2}\sigma^2}{\sqrt{nK}} + \frac{L^2 C^2}{K(1-q)^2} \left[(4\sigma^2 + 12\zeta^2)n + \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \right] \\ &\quad + \frac{\sqrt{n}L^2 C^2}{\sqrt{K}(1-q)^2 K} \left[\sigma^2 + 3L^2 C^2 \zeta^2 + 2 \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \right] \end{aligned} \quad (30)$$

Using again the assumption on the lower bound (3) of the total number of iterations K , the last two terms of the above expression are bounded by the first term. Thus,

$$\frac{1}{4} \frac{\sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2}{K} \leq 3 \frac{f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L}{2}\sigma^2}{\sqrt{nK}} \quad (31)$$

E.3.2. PROOF OF THEOREM 2

Proof. From Lemma 6 we have that:

$$\begin{aligned}
 \frac{1}{K} \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\gamma^2 \frac{4C^2}{(1-q)^2 D_2} \right) \sigma^2 + \left(\gamma \frac{C^2}{(1-q)^2 D_2} \right) \frac{\sigma^2}{K} \\
 &\quad + \left(\gamma^2 \frac{12C^2}{(1-q)^2 D_2} \right) \zeta^2 + \left(\frac{\gamma 3C^2}{(1-q)^2 D_2} \right) \frac{\zeta^2}{K} \\
 &\quad + \left(\frac{C^2}{(1-q)^2 D_2 K} + \gamma \frac{2C^2}{(1-q)^2 D_2 K} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \\
 &\quad + \left(\gamma^2 \frac{12C^2}{(1-q)^2 D_2} + \gamma \frac{3C^2}{(1-q)^2 D_2} \right) \frac{\sum_{k=0}^{K-1} \mathbb{E} \|\nabla f(\bar{\mathbf{x}}^{(k)})\|^2}{K}
 \end{aligned} \tag{32}$$

Using the assumptions of Theorem 1 and stepsize $\gamma = \sqrt{\frac{n}{K}}$:

$$\begin{aligned}
 \frac{1}{K} \sum_{k=0}^{K-1} M^{(k)} &\leq \left(\frac{n}{K} \frac{4C^2}{(1-q)^2 D_2} \right) \sigma^2 + \left(\sqrt{\frac{n}{K}} \frac{C^2}{(1-q)^2 D_2} \right) \frac{\sigma^2}{K} \\
 &\quad + \left(\frac{n}{K} \frac{12C^2}{(1-q)^2 D_2} \right) \zeta^2 + \left(\frac{\sqrt{\frac{n}{K}} 3C^2}{(1-q)^2 D_2} \right) \frac{\zeta^2}{K} \\
 &\quad + \left(\frac{C^2}{(1-q)^2 D_2 K} + \sqrt{\frac{n}{K}} \frac{2C^2}{(1-q)^2 D_2 K} \right) \frac{\sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n} \\
 &\quad + \left(\frac{n}{K} \frac{12C^2}{(1-q)^2 D_2} + \sqrt{\frac{n}{K}} \frac{3C^2}{(1-q)^2 D_2} \right) \frac{12 \left[f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L}{2} \sigma^2 \right]}{\sqrt{nK}} \\
 &= \frac{1}{K} \left[\frac{4nC^2\sigma^2}{(1-q)^2 D_2} + \frac{12nC^2\zeta^2}{(1-q)^2 D_2} + \frac{C^2 \sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{n(1-q)^2 D_2} + \frac{3\sqrt{n}C^2 12 \left[f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L}{2} \sigma^2 \right]}{\sqrt{n}(1-q)^2 D_2} \right] \\
 &\quad + \frac{1}{K\sqrt{K}} \left[\frac{n\sigma^2 C^2}{(1-q)^2 D_2} + \frac{\frac{n}{3} C^2 \zeta^2}{(1-q)^2 D_2} + \frac{2C^2 \sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2}{(1-q)^2 D_2 \sqrt{n}} + \frac{144\sqrt{n}C^2 \left[f(\bar{\mathbf{x}}^{(0)}) - f^* + \frac{L}{2} \sigma^2 \right]}{(1-q)^2 D_2} \right] \\
 &= O\left(\frac{1}{K} + \frac{1}{K\sqrt{K}} \right)
 \end{aligned} \tag{33}$$

where the Big O notation swallows all constants of our setting $(n, L, \sigma, \zeta, C, q, \sum_{i=1}^n \|\mathbf{x}_i^{(0)}\|^2 \text{ and } f(\bar{\mathbf{x}}^{(0)}) - f^*)$.

Now using the above upper bound equation 33 and result of Theorem 1 we obtain:

$$\begin{aligned}
 \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \|\nabla f(\mathbf{z}_i^k)\|^2 &= \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\| \nabla f(\mathbf{z}_i^k) + \nabla f(\bar{\mathbf{x}}^{(k)}) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\leq \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n 2\mathbb{E} \left\| \nabla f(\mathbf{z}_i^k) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + 2\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &= \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n 2\mathbb{E} \left\| \nabla f(\mathbf{z}_i^k) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n 2\mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &= 2\frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\| \nabla f(\mathbf{z}_i^k) - \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 + 2\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\stackrel{L\text{-smooth}}{=} 2L^2 \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\| \mathbf{z}_i^k - \bar{\mathbf{x}}^{(k)} \right\|^2 + 2\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{x}}^{(k)}) \right\|^2 \\
 &\stackrel{(33)+(31)}{\leq} O\left(\frac{1}{\sqrt{nK}} + \frac{1}{K} + \frac{1}{K^{3/2}}\right) \tag{34}
 \end{aligned}$$

where again the Big O notation swallows all constants of our setting $(n, L, \sigma, \zeta, C, q, \sum_{i=1}^n \|x_i^{(0)}\|^2 \text{ and } f(\bar{\mathbf{x}}^{(0)}) - f^*)$. \square