
Distributed Learning over Unreliable Networks

Chen Yu¹ Hanlin Tang¹ Cedric Renggli² Simon Kassing² Ankit Singla² Dan Alistarh³ Ce Zhang² Ji Liu^{4,1}

Abstract

Most of today’s distributed machine learning systems assume *reliable networks*: whenever two machines exchange information (e.g., gradients or models), the network should guarantee the delivery of the message. At the same time, recent work exhibits the impressive tolerance of machine learning algorithms to errors or noise arising from relaxed communication or synchronization. In this paper, we connect these two trends, and consider the following question: *Can we design machine learning systems that are tolerant to network unreliability during training?* With this motivation, we focus on a theoretical problem of independent interest—given a standard distributed parameter server architecture, if every communication between the worker and the server has a non-zero probability p of being dropped, does there exist an algorithm that still converges, and at what speed? The technical contribution of this paper is a novel theoretical analysis proving that distributed learning over unreliable network can achieve comparable convergence rate to centralized or distributed learning over reliable networks. Further, we prove that the influence of the packet drop rate diminishes with the growth of the number of parameter servers. We map this theoretical result onto a real-world scenario, training deep neural networks over an unreliable network layer, and conduct network simulation to validate the system improvement by allowing the networks to be unreliable.

1. Introduction

Distributed learning has attracted significant interest from both academia and industry. Over the last decade, re-

¹Department of Computer Science, University of Rochester, USA ²Department of Computer Science, ETH Zurich ³Institute of Science and Technology Austria ⁴Seattle AI Lab, FeDA Lab, Kwai Inc. Correspondence to: Chen Yu <cyu28@ur.rochester.edu>.




Figure 1: An illustration of the communication pattern of distributed learning with three parameter servers and four workers — each server serves a partition of the model, and each worker holds a replica of the whole model. In this paper, we focus on the case in which every communication between the worker and the server has a non-zero probability p of being dropped.

searchers have come with up a range of different designs of more efficient learning systems. An important subset of this work focuses on understanding the impact of different system relaxations to the convergence and performance of distributed stochastic gradient descent, such as the compression of communication, e.g (Seide & Agarwal, 2016), decentralized communication (Lian et al., 2017a; Sirb & Ye, 2016; Lan et al., 2017; Tang et al., 2018a; Stich et al., 2018), and asynchronous communication (Lian et al., 2017b; Zhang et al., 2013; Lian et al., 2015). Most of these works are motivated by real-world system bottlenecks, abstracted as general problems for the purposes of analysis.

In this paper, we focus on the centralized SGD algorithm in the distributed machine learning scenario implemented using standard AllReduce operator or parameter server architecture and we are motivated by a new type of system relaxation—the reliability of the communication channel. We abstract this problem as a theoretical one, conduct a novel convergence analysis for this scenario, and then validate our results in a practical setting.

The *Centralized SGD* algorithm works as Figure 1 shows. Given n machines/workers, each maintaining its own local model, each machine alternates local SGD steps with global communication steps, in which machines exchange their *local models*. In this paper, we covers two standard

distributed settings: the Parameter Server model¹ (Li et al., 2014; Abadi et al., 2016), as well as standard implementations of the AllReduce averaging operation in a decentralized setting (Seide & Agarwal, 2016; Renggli et al., 2018). There are two components in the communication step:

1. **Step 0 - Model Partitioning (Only Conducted Once).** In most state-of-the-art implementations of AllReduce and parameter servers (Li et al., 2014; Abadi et al., 2016; Thakur et al., 2005), models are partitioned into n blocks, and each machine is the owner of one block (Thakur et al., 2005). The rationale is to increase parallelism over the utilization of the underlying communication network. The partitioning strategy does not change during training.
2. **Step 1.1 - Reduce-Scatter.** In the Reduce-Scatter step, for each block (model partition) i , the machines average their model on the block by sending it to the corresponding machine.
3. **Step 1.2 - All-Gather.** In the subsequent All-Gather step, each machine broadcasts its block to all others, so that all machines have a consistent model copy.

In this paper, we focus on the scenario in which the communication is unreliable — The communication channel between any two machines has a probability p of not delivering a message, as the Figure 1 shows, where the grey arrows represent the dropping message and the black arrows represent the success-delivering message. We change the two aggregation steps as follows. In the Reduce-Scatter step, a *uniform random* subset of machines will average their model on each model block i . In the All-Gather step, it is again a *uniform random* subset of machines which receive the resulting average. Specifically, machines not chosen for the Reduce-Scatter step do not contribute to the average, and all machines that are not chosen for the All-Gather will not receive updates on their model block i . This is a realistic model of running an AllReduce operator implemented with Reduce-Scatter/All-Gather on unreliable network. We call this revised algorithm the RPS algorithm.

Our main technical contribution is characterizing the convergence properties of the RPS algorithm. To the best of our knowledge, this is a novel theoretical analysis of this faulty communication model. We will survey related work in more details in Section 2.

We then apply our theoretical result to a real-world use case, illustrating the potential benefit of allowing an unreliable network. We focus on a realistic scenario where the network is shared among multiple applications or tenants, for

instance in a data center. Both applications communicate using the same network. In this case, if the machine learning traffic is tolerant to some packet loss, the other application can potentially be made faster by receiving priority for its network traffic. Via network simulations, we find that tolerating a 10% drop rate for the learning traffic can make a simple (emulated) Web service up to 1.2x faster (Even small speedups of 10% are significant for such services; for example, Google actively pursues minimizing its Web services’ response latency). At the same time, this degree of loss does not impact the convergence rate for a range of machine learning applications, such as image classification and natural language processing.

Organization The rest of this paper is organized as follows. We first review some related work in Section 2. Then we formulate the problem in Section 3 and describe the RPS algorithm in Section 4, with its theoretical guarantee stated in Section 5. We evaluate the scalability and accuracy of the RPS algorithm in Section 6 and study an interesting case of speeding up colocated applications in Section 7. At last, we conclude the paper in Section 8. The proofs of our theoretical results can be found in the supplementary material.

2. Related Work

Distributed Learning There has been a huge number of works on distributing deep learning, e.g. Seide & Agarwal (2016); Abadi et al. (2016); Goyal et al. (2017); Colin et al. (2016). Also, many optimization algorithms are proved to achieve much better performance with more workers. For example, Hajinezhad et al. (2016) utilize a primal-dual based method for optimizing a finite-sum objective function and proved that it’s possible to achieve a $O(n)$ speedup corresponding to the number of the workers. In Xu et al. (2017), an adaptive consensus ADMM is proposed and Goldstein et al. (2016) studied the performance of transpose ADMM on an entire distributed dataset. In Scaman et al. (2017), the optimal convergence rate for both centralized and decentralized distributed learning is given with the time cost for communication included. In Lin et al. (2018); Stich (2018), they investigate the trade off between getting more mini-batches or having more communication. To save the communication cost, some sparse based distributed learning algorithms is proposed (Shen et al., 2018b; Wu, 2018; Wen et al., 2017; McMahan et al., 2016; Wang et al., 2016). Recent works indicate that many distributed learning is delay-tolerant under an asynchronous setting (Zhou et al., 2018; Lian et al., 2015; Sra et al., 2015; Leblond et al., 2016). Also, in Blanchard et al. (2017); Yin et al. (2018); Alistarh et al. (2018) They study the Byzantine-robust distributed learning when the Byzantine worker is included in the network. In Drumond et al. (2018), authors proposed a compressed DNN training strategy in order to save the computational cost of floating

¹Our modeling above fits the case of n workers and n parameter servers, although our analysis will extend to any setting of these parameters.

point.

Centralized parallel training Centralized parallel (Agarwal & Duchi, 2011; Recht et al., 2011) training works on the network that is designed to ensure that all workers could get information of all others. One communication primitive in centralized training is to average/aggregate all models, which is called *a collective communication operator* in HPC literature (Thakur et al., 2005). Modern machine learning systems rely on different implementations, e.g., parameter server model (Li et al., 2014; Abadi et al., 2016) or the standard implementations of the AllReduce averaging operation in a decentralized setting (Seide & Agarwal, 2016; Renggli et al., 2018). In this work, we focus on the behavior of centralized ML systems under unreliable network, when this primitive is implemented as a distributed parameter servers (Jiang et al., 2017), which is similar to a Reduce-Scatter/All-Gather communication paradigm. For many implementations of collective communication operators, partitioning the model is one key design point to reach the peak communication performance (Thakur et al., 2005).

Decentralized parallel training Another direction of related work considers decentralized learning. Decentralized learning algorithms can be divided into *fixed* topology algorithms and *random* topology algorithms. There are many work related to the *fixed* topology decentralized learning. Specifically, Jin et al. (2016) proposes to scale the gradient aggregation process via a gossip-like mechanism. Lian et al. (2017a) provided strong convergence bounds for a similar algorithm to the one we are considering, in a setting where the communication graph is fixed and regular. In Tang et al. (2018b), a new approach that admits a better performance than decentralized SGD when the data among workers is very different is studied. Shen et al. (2018a) generalize the decentralized optimization problem to a monotone operator. In He et al. (2018), authors study a decentralized gradient descent based algorithm (**CoLA**) for learning of linear classification and regression model. For the *random* topology decentralized learning, the weighted matrix for randomized algorithms can be time-varying, which means workers are allowed to change the communication network based on the availability of the network. There are many works (Boyd et al., 2006; Li & Zhang, 2010; Lobel & Ozdaglar, 2011; Nedic et al., 2017; Nedić & Olshevsky, 2015) studying the random topology decentralized SGD algorithms under different assumptions. Blot et al. (2016) considers a more radical approach, called GoSGD, where each worker exchanges gradients with a random subset of other workers in each round. They show that GoSGD can be faster than Elastic Averaging SGD (Zhang et al., 2015) on CIFAR-10, but provide no large-scale experiments or theoretical justification. Recently, Daily et al. (2018) proposed GossipGrad, a more complex gossip-based scheme with upper bounds on the time for workers to communicate indirectly, periodic rotation of partners and shuffling of the input data, which

provides strong empirical results on large-scale deployments. The authors also provide an informal justification for why GossipGrad should converge.

In this paper, we consider a general model communication, which covers both Parameter Server (Li et al., 2014) and AllReduce (Seide & Agarwal, 2016) distribution strategies. We specifically include the uncertainty of the network into our theoretical analysis. In addition, our analysis highlights the fact that the system can handle additional packet drops as we increase the number of worker nodes.

3. Problem Setup

We consider the following distributed optimization problem:

$$\min_{\mathbf{x}} f(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \underbrace{\mathbb{E}_{\xi} D_i F_i(\mathbf{x}; \cdot)}_{=: f_i(\mathbf{x})}, \quad (1)$$

where n is the number of workers, D_i is the local data distribution for worker i (in other words, we do not assume that all nodes can access the same data set), and $F_i(\mathbf{x}; \cdot)$ is the local loss function of model \mathbf{x} given data \cdot for worker i .

Unreliable Network Connection Nodes can communicate with all other workers, but with packet drop rate p (here we do not use the common-used phrase “packet loss rate” because we use “loss” to refer to the loss function). That means, whenever any node forwards models or data to any other model, the destination worker *fails* to receive it, with probability p . For simplicity, we assume that all packet drop events are independent, and that they occur with the same probability p .

Definitions and notations Throughout, we use the following notation and definitions:

$\nabla f(\cdot)$ denotes the gradient of the function f .

$\lambda_i(\cdot)$ denotes the i th largest eigenvalue of a matrix.

$\mathbf{1}_n = [1, 1, \dots, 1]^T \in \mathbb{R}^n$ denotes the full-one vector.

$A_n := \frac{\mathbf{1}_n \mathbf{1}_n^T}{n}$ denotes the all $\frac{1}{n}$'s n by n matrix.

$\|\cdot\|_k$ denotes the ℓ_2 norm for vectors.

$\|\cdot\|_F$ denotes the Frobenius norm of matrices.

4. Algorithm

In this section, we describe our algorithm, namely RPS — *Reliable Parameter Server* — as it is robust to package loss in the network layer. We first describe our algorithm in detail, followed by its interpretation from a global view.

4.1. Our Algorithm: RPS

In the RPS algorithm, each worker maintains an individual local model. We use $\mathbf{x}_t^{(i)}$ to denote the local model on worker i at time step t . At each iteration t , each worker first performs a regular SGD step

$$\mathbf{v}_t^{(i)} = \mathbf{x}_t^{(i)} - \gamma \nabla F_i(\mathbf{x}_t^{(i)}; \mathcal{D}_t^{(i)});$$

where γ is the learning rate and $\mathcal{D}_t^{(i)}$ are the data samples of worker i at iteration t .

We would like to reliably average the vector $\mathbf{v}_t^{(i)}$ among all workers, via the RPS procedure. In brief, the RS step performs communication-efficient model averaging, and the AG step performs communication-efficient model sharing.

The Reduce-Scatter (RS) step: In this step, each worker i divides $\mathbf{v}_t^{(i)}$ into n equally-sized blocks.

$$\mathbf{v}_t^{(i)} = \left(\left(\mathbf{v}_t^{(i,1)} \right)^\triangleright, \left(\mathbf{v}_t^{(i,2)} \right)^\triangleright, \dots, \left(\mathbf{v}_t^{(i,n)} \right)^\triangleright \right)^\triangleright. \quad (2)$$

The reason for this division is to reduce the communication cost and parallelize model averaging since we only assign each worker for averaging one of those blocks. For example, worker 1 can be assigned for averaging the first block while worker 2 might be assigned to deal with the third block. For simplicity, we would proceed our discussion in the case that worker i is assigned for averaging the i th block.

After the division, each worker sends its i th block to worker i . Once receiving those blocks, each worker would average all the blocks it receives. As noted, some packets might be dropped. In this case, worker i averages all those blocks using

$$\mathbf{v}_t^{(i)} = \frac{1}{j \mathcal{N}_t^{(i)}} \sum_{j \in \mathcal{N}_t^{(i)}} \mathbf{v}_t^{(i,j)},$$

where $\mathcal{N}_t^{(i)}$ is the set of the packages worker i receives (may including the worker i 's own package).

The AllGather (AG) step: After computing $\mathbf{v}_t^{(i)}$, each worker i attempts to broadcast $\mathbf{v}_t^{(i)}$ to all other workers, using the averaged blocks to recover the averaged original vector $\mathbf{x}_t^{(i)}$ by concatenation:

$$\mathbf{x}_{t+1}^{(i)} = \left(\left(\mathbf{v}_t^{(i,1)} \right)^\triangleright, \left(\mathbf{v}_t^{(i,2)} \right)^\triangleright, \dots, \left(\mathbf{v}_t^{(i,n)} \right)^\triangleright \right)^\triangleright.$$

Note that it is entirely possible that some workers in the network may not be able to receive some of the averaged blocks. In this case, they just use the original block. Formally,

$$\mathbf{x}_{t+1}^{(i)} = \left(\left(\mathbf{x}_{t+1}^{(i,1)} \right)^\triangleright, \left(\mathbf{x}_{t+1}^{(i,2)} \right)^\triangleright, \dots, \left(\mathbf{x}_{t+1}^{(i,n)} \right)^\triangleright \right)^\triangleright, \quad (3)$$

Algorithm 1 RPS

- 1: **Input:** Initialize all $\mathbf{x}_1^{(i)}$, $\mathcal{S}_i \subseteq [n]$ with the same value, learning rate γ , and number of total iterations T .
- 2: **for** $t = 1, 2, \dots, T$ **do**
- 3: Randomly sample $\mathcal{D}_t^{(i)}$ from local data of the i th worker, $\mathcal{S}_i \subseteq [n]$.
- 4: Compute a local stochastic gradient based on $\mathcal{D}_t^{(i)}$ and current optimization variable $\mathbf{x}_t^{(i)}$: $-\nabla F_i(\mathbf{x}_t^{(i)}; \mathcal{D}_t^{(i)})$, $\mathcal{S}_i \subseteq [n]$
- 5: Compute the intermediate model $\mathbf{v}_t^{(i)}$ according to

$$\mathbf{v}_t^{(i)} = \mathbf{x}_t^{(i)} - \gamma \nabla F_i(\mathbf{x}_t^{(i)}; \mathcal{D}_t^{(i)}),$$

and divide $\mathbf{v}_t^{(i)}$ into n blocks $\left(\left(\mathbf{v}_t^{(i,1)} \right)^\triangleright, \left(\mathbf{v}_t^{(i,2)} \right)^\triangleright, \dots, \left(\mathbf{v}_t^{(i,n)} \right)^\triangleright \right)^\triangleright$.

- 6: For any $i \in [n]$, randomly choose one worker $b_t^{(i)}$ without replacement. Then, every worker attempts to send their i th block of their intermediate model to worker $b_t^{(i)}$. Then each worker averages all received blocks using

$$\mathbf{v}_t^{(i)} = \frac{1}{j \mathcal{N}_t^{(i)}} \sum_{j \in \mathcal{N}_t^{(i)}} \mathbf{v}_t^{(i,j)}.$$

- 7: Worker $b_t^{(i)}$ broadcast $\mathbf{v}_t^{(i)}$ to all workers (maybe dropped due to packet drop), $\mathcal{S}_i \subseteq [n]$.
- 8: $\mathbf{x}_{t+1}^{(i)} = \left(\left(\mathbf{x}_{t+1}^{(i,1)} \right)^\triangleright, \left(\mathbf{x}_{t+1}^{(i,2)} \right)^\triangleright, \dots, \left(\mathbf{x}_{t+1}^{(i,n)} \right)^\triangleright \right)^\triangleright$, where

$$\mathbf{x}_{t+1}^{(i,j)} = \begin{cases} \mathbf{v}_t^{(i,j)} & j \in \tilde{\mathcal{N}}_t^{(i)} \\ \mathbf{x}_t^{(i,j)} & j \notin \tilde{\mathcal{N}}_t^{(i)} \end{cases},$$

for all $i \in [n]$.

- 9: **end for**
- 10: **Output:** $\mathbf{x}_T^{(i)}$

^aHere $b_t^{(i)} \in \{1, 2, \dots, n\}$ indicates which worker is assigned for averaging the i th block.

where

$$\mathbf{x}_{t+1}^{(i,j)} = \begin{cases} \mathbf{v}_t^{(i,j)} & j \in \tilde{\mathcal{N}}_t^{(i)} \\ \mathbf{x}_t^{(i,j)} & j \notin \tilde{\mathcal{N}}_t^{(i)} \end{cases}$$

We can see that each worker just replace the corresponding blocks of $\mathbf{v}_t^{(i)}$ using received averaged blocks. The complete algorithm is summarized in Algorithm 1.

4.2. RPS From a Global Viewpoint

It can be seen that at time step t , the j th block of worker i 's local model, that is, $\mathbf{x}_t^{(i,j)}$, is a linear combination of j th

block of all workers' intermediate model $\mathbf{v}_t^{(k,j)}$ ($k \geq [n]$),

$$X_{t+1}^{(j)} = V_t^{(j)} W_t^{(j)}, \quad (4)$$

where

$$\begin{aligned} X_{t+1}^{(j)} &:= (\mathbf{x}_{t+1}^{(1,j)}, \mathbf{x}_{t+1}^{(2,j)}, \dots, \mathbf{x}_{t+1}^{(n,j)}) \\ V_t^{(j)} &:= (\mathbf{v}_t^{(1,j)}, \mathbf{v}_t^{(2,j)}, \dots, \mathbf{v}_t^{(n,j)}) \end{aligned}$$

and $W_t^{(j)}$ is the coefficient matrix indicating the communication outcome at time step t . The (m, k) th element of $W_t^{(j)}$ is denoted by $[W_t^{(j)}]_{m,k}$. $[W_t^{(j)}]_{m,k} \notin 0$ means that worker k receives worker m 's individual j th block (that is, $\mathbf{v}_t^{(m,j)}$), whereas $[W_t^{(j)}]_{m,k} = 0$ means that the package might be dropped either in RS step (worker m fails to send) or AG step (worker k fails to receive). So $W_t^{(j)}$ is time-varying because of the randomness of the package drop. Also $W_t^{(j)}$ is not doubly-stochastic (in general) because the package drop is independent between RS step and AG step.




Figure 2: α_2 under different number of workers n and package drop rate p .




Figure 3: $(\alpha_1 - \alpha_2)$ under different number of workers n and package drop rate p .

The property of $W_t^{(j)}$ In fact, it can be shown that all

$W_t^{(j)}$'s ($\delta j, \delta t$) satisfy the following properties

$$\left(\mathbb{E}(W_t^{(j)}) \right) A_n = A_n$$

$$\mathbb{E} \left[W_t^{(j)} \left(W_t^{(j)} \right)^{\triangleright} \right] = \alpha_1 I_n + (1 - \alpha_1) A_n \quad (5)$$

$$\mathbb{E} \left[W_t^{(j)} A_n \left(W_t^{(j)} \right)^{\triangleright} \right] = \alpha_2 I_n + (1 - \alpha_2) A_n \quad (6)$$

for some constants α_1 and α_2 satisfying $0 < \alpha_2 < \alpha_1 < 1$ (see Lemmas 6, 7, and 8 in Supplementary Material). Since the exact expression is too complex, we plot the α_1 and α_2 related to different n in Figure 2 and Figure 3 (detailed discussion is included in **Section D** in Supplementary Material.). Here, we do not plot α_2 , but plot $\alpha_1 - \alpha_2$ instead. This is because $\alpha_1 - \alpha_2$ is an important factor in our Theorem (See Section 5 where we define $\alpha_1 - \alpha_2$ as β).

5. Theoretical Guarantees and Discussion

Below we show that, for certain parameter values, RPS with unreliable communication rates admits the same convergence rate as the standard algorithms. In other words, the impact of network unreliability may be seen as negligible.

First let us make some necessary assumptions, that are commonly used in analyzing stochastic optimization algorithms.

Assumption 1. We make the following commonly used assumptions:

1. **Lipschitzian gradient:** All function $f_i(\cdot)$'s are with L -Lipschitzian gradients, which means

$$\|k r f_i(\mathbf{x}) - r f_i(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$$

2. **Bounded variance:** Assume the variance of stochastic gradient

$$\mathbb{E}_{\xi} \left[D_i k r F_i(\mathbf{x}; \xi) - r f_i(\mathbf{x}) \right]^2 \leq \sigma^2, \quad \delta i, \delta \mathbf{x},$$

$$\frac{1}{n} \sum_{i=1}^n k r f_i(\mathbf{x}) - r f(\mathbf{x}) \|^2 \leq \zeta^2, \quad \delta i, \delta \mathbf{x},$$

is bounded for any x in each worker i .

3. **Start from 0:** We assume $X_1 = 0$ for simplicity w.l.o.g.

Next we are ready to show our main result.

Theorem 1 (Convergence of Algorithm 1). *Under Assumption 1, choosing γ in Algorithm 1 to be small enough that satisfies $1 - \frac{6L\beta\gamma^2}{(1-\beta)^2} > 0$, we have the following convergence rate for Algorithm 1*

$$\frac{1}{T} \sum_{t=1}^T \left(\mathbb{E} k r f(\bar{\mathbf{x}}_t) k^2 + (1 - L\gamma) \mathbb{E} \|\bar{r} f(X_t)\|^2 \right)$$

$$\begin{aligned} & \frac{2f(\mathbf{0}) - 2f(\mathbf{x})}{\gamma T} + \frac{\gamma L \sigma^2}{n} + 4\alpha_2 L \gamma (\sigma^2 + 3\zeta^2) \\ & + \frac{(2\alpha_2 L \gamma + L^2 \gamma^2 + 12\alpha_2 L^3 \gamma^3) \sigma^2 C_1}{(1 - \frac{p}{\beta})^2} \\ & + \frac{3(2\alpha_2 L \gamma + L^2 \gamma^2 + 12\alpha_2 L^3 \gamma^3) \zeta^2 C_1}{(1 - \frac{p}{\beta})^2}, \end{aligned} \quad (7)$$

where

$$\begin{aligned} r f(\bar{\mathbf{x}}_t) &= f\left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_t^{(i)}\right), \\ \bar{r} f(\mathbf{X}_t) &= \sum_{i=1}^n r f_i(\mathbf{x}_t^{(i)}), \\ \beta &= \alpha_1 - \alpha_2, \\ C_1 &= \left(1 - \frac{6L^2 \gamma^2}{(1 - \frac{p}{\beta})^2}\right)^{-1}, \end{aligned}$$

and α_1, α_2 follows the definition in (5) and (6).

To make the result more clear, we appropriately choose the learning rate as follows:

Corollary 2. Choose $\gamma = \frac{1}{6L + 3(\sigma + \zeta)} \frac{1 - \frac{p}{\beta}}{\alpha_2 T + \frac{p}{n}}$ in Algorithm 1, under Assumption 1, we have the follow convergence rate for Algorithm 1

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} k r f(\bar{\mathbf{x}}_t) k^2 & \leq \frac{(\sigma + \zeta) \left(1 + \frac{p}{n\alpha_2}\right)}{(1 - \frac{p}{\beta})^2 \frac{p}{nT}} + \frac{1}{T} \\ & + \frac{n(\sigma^2 + \zeta^2)}{(1 + n\alpha_2)\sigma^2 T + n\alpha_2 T \zeta^2}, \end{aligned}$$

where $\beta, \alpha_1, \alpha_2, r f(\bar{\mathbf{x}})$ follow to the definitions in Theorem 1, and we treat $f(0), f'$, and L to be constants.

We discuss our theoretical results below

(Comparison with centralized SGD and decentralized SGD) The dominant term in the convergence rate is $O(1/\sqrt{nT})$ (here we use $\alpha_2 = O(p(1-p)/n)$ and $\beta = O(p)$ which is shown by Lemma 8 in Supplement), which is consistent with the rate for centralized SGD and decentralized SGD (Lian et al., 2017a).

(Linear Speedup) Since the the leading term of convergence rate for $\frac{1}{T} \sum_{t=1}^T \mathbb{E} k r f(\bar{\mathbf{x}}_t) k^2$ is $O(1/\sqrt{nT})$. It suggests that our algorithm admits the linear speedup property with respect to the number of workers n .

(Better performance for larger networks) Fixing the package drop rate p (implicitly included in **Section D**), the convergence rate under a larger network (increasing

n) would be superior, because the leading terms' dependence of the $\alpha_2 = O(p(1-p)/n)$. This indicates that the affection of the package drop ratio diminishes, as we increase the number of workers and parameter servers.

(Why only converges to a ball of a critical point)

This is because we use a constant learning rate, the algorithm could only converges to a ball centered at a critical point. This is a common choice to make the statement simpler, just like many other analysis for SGD. Our proved convergence rate is totally consistent with SGD's rate, and could converge (in the same rate) to a critical point by choosing a decayed learning rate such as $O(1/\sqrt{T})$ like SGD.

6. Experiments: Convergence of RPS

We now validate empirically the scalability and accuracy of the RPS algorithm, given reasonable message arrival rates.

6.1. Experimental Setup

Datasets and models We evaluate our algorithm on two state of the art machine learning tasks: (1) image classification and (2) natural language understanding (NLU). We train ResNet (He et al., 2016) with different number of layers on CIFAR-10 (Krizhevsky & Hinton, 2009) for classifying images. We perform the NLU task on the Air travel information system (ATIS) corpus on a one layer LSTM network.

Implementation We simulate packet losses by adapting the latest version 2.5 of the Microsoft Cognitive Toolkit (Seide & Agarwal, 2016). We implement the RPS algorithm using MPI. During training, we use a local batch size of 32 samples per worker for image classification. We adjust the learning rate by applying a linear scaling rule (Goyal et al., 2017) and decay of 10 percent after 80 and 120 epochs, respectively. To achieve the best possible convergence, we apply a gradual warmup strategy (Goyal et al., 2017) during the first 5 epochs. We deliberately do not use any regularization or momentum during the experiments in order to be consistent with the described algorithm and proof. The NLU experiments are conducted with the default parameters given by the CNTK examples, with scaling the learning rate accordingly, and omit momentum and regularization terms on purpose. The training of the models is executed on 16 NVIDIA TITAN Xp GPUs. The workers are connected by Gigabit Ethernet. We use each GPU as a worker. We describe the results in terms of training loss convergence, although the validation trends are similar.

Convergence of Image Classification We perform convergence tests using the analyzed algorithm, model averaging SGD, on both ResNet110 and ResNet20 with CIFAR-10. Figure 4(a,b) shows the result. We vary probabilities for each packet being correctly delivered at each worker be-




Figure 4: Convergence of RPS on different datasets.

tween 80%, 90%, 95% and 99%. The baseline is 100% message delivery rate. The baseline achieves a training loss of 0.02 using ResNet110 and 0.09 for ResNet20. Dropping 1% doesn’t increase the training loss achieved after 160 epochs. For 5% the training loss is identical on ResNet110 and increased by 0.02 on ResNet20. Having a probability of 90% of arrival leads to a training loss of 0.03 for ResNet110 and 0.11 for ResNet20 respectively.

Convergence of NLU We perform full convergence tests for the NLU task on the ATIS corpus and a single layer LSTM. Figure 4(c) shows the result. The baseline achieves a training loss of 0.01. Dropping 1, 5 or 10 percent of the communicated partial vectors result in an increase of 0.01 in training loss.

Comparison with Gradient Averaging We conduct experiments with identical setup and a probability of 99 percent of arrival using a *gradient* averaging methods, instead of model averaging. When running data distributed SGD, gradient averaging is the most widely used technique in practice, also implemented by default in most deep learning frameworks (Abadi et al., 2016; Seide & Agarwal, 2016). As expected, the baseline (all the transmissions are successful) converges to the same training loss as its model averaging counterpart, when omitting momentum and regularization terms. As seen in figures 5(a,b), having a loss in communication of even 1 percentage results in worse convergence in terms of accuracy for both ResNet architectures on CIFAR-10. The reason is that the error of package drop will accumulate over iterations but never decay, because the model is the sum of all early gradients, so the model never converges to the optimal one. Nevertheless, this insight suggests that one should favor a model averaging algorithm over gradient averaging, if the underlying network connection is unreliable.

Extended Analysis We report additional experiments with a brief analysis of the generalization properties regarding the test accuracy in Section E in the supplementary material. We test a wider range of probabilities for each packet being correctly delivered, and finally analyze the impact of varying number of nodes at a fixed arrival rate. We use the

default hyperparameter values given by the deep learning frameworks for all the experiments conducted. As a further work, one might envisage tuning the hyper-parameters such as learning rate for a given number of nodes and arrival probability.

7. Case study: Speeding up Colocated Applications

Our results on the resilience of distributed learning to losses of model updates open up an interesting use case. That model updates can be lost (within some tolerance) without the deterioration of model convergence implies that model updates transmitted over the physical network can be de-prioritized compared to other more “inflexible,” delay-sensitive traffic, such as for Web services. Thus, we can colocate other applications with the training workloads, and reduce infrastructure costs for running them. Equivalently, workloads that are colocated with learning workers can benefit from prioritized network traffic (at the expense of some model update losses), and thus achieve lower latency.

To demonstrate this in practice, we perform a packet simulation over 16 servers, each connected with a 1 Gbps link to a network switch. Over this network of 16 servers, we run two workloads: (a) replaying traces from the machine learning process of ResNet110 on CIFAR-10 (which translates to a load of 2.4 Gbps) which is sent *unreliably*, and (b) a simple emulated Web service running on all 16 servers. Web services often produce significant background traffic between servers within the data center, consisting typically of small messages fetching distributed pieces of content to compose a response (e.g., a Google query response potentially consists of advertisements, search results, and images). We emulate this intra data center traffic for the Web service as all-to-all traffic between these servers, with small messages of 100 KB (a reasonable size for such services) sent reliably between these servers. The inter-arrival time for these messages follows a Poisson process, parametrized by the expected message rate, λ (aggregated across the 16 servers).

Different degrees of prioritization of the Web service traffic over learning traffic result in different degrees of loss in

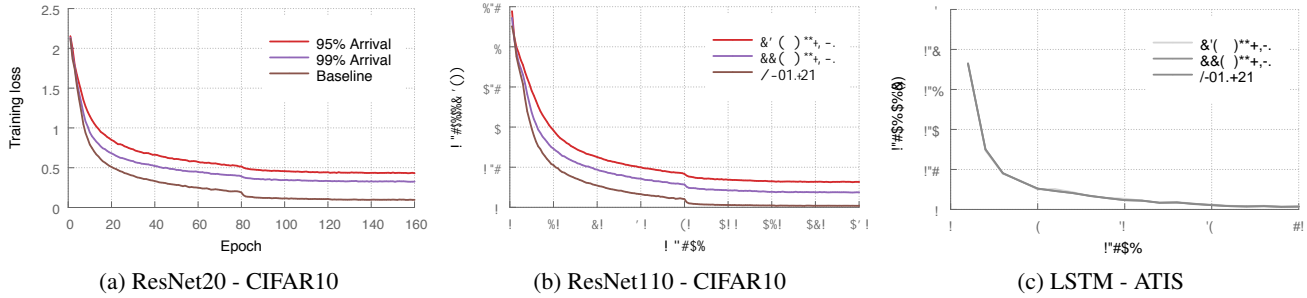


Figure 5: Why RPS? The Behavior of Standard SGD in the Presence of Message Drop.




Figure 6: Allowing an increasing rate of losses for model updates speeds up the Web service.

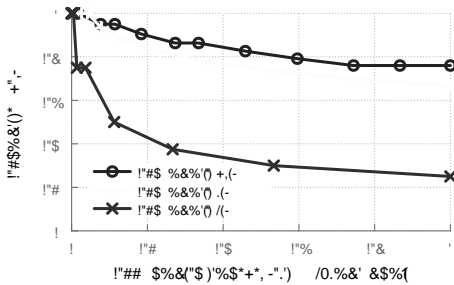


Figure 7: Allowing more losses for model updates reduces the cost for the Web service.

learning updates transmitted over the network. As the Web service is prioritized to a greater extent, its performance improves – its message exchanges take less time; we refer to this reduction in (average) completion time for these messages as a speed-up. Note that even small speedups of 10% are significant for such services; for example, Google actively pursues minimizing its Web services’ response latency. An alternative method of quantifying the benefit for the colocated Web service is to measure how many additional messages the Web service can send, while maintaining a fixed average completion time. This translates to running more Web service queries and achieving more throughput over the same infrastructure, thus reducing cost per request / message.

Fig. 6 and Fig. 7 show results for the above described Web service speedup and cost reduction respectively. In Fig. 6, the arrival rate of Web service messages is fixed ($\lambda = \{2000, 5000, 10000\}g$ per second). As the network prioritizes the Web service more and more over learning update traffic, more learning traffic suffers losses (on the x -axis), but performance for the Web service improves. With just 10% losses for learning updates, the Web service can be sped up by more than 20% (*i.e.*, 1.2).

In Fig. 7, we set a target average transmission time (2, 5, or 10 ms) for the Web service’s messages, and increase the message arrival rate, λ , thus causing more and more losses for learning updates on the x -axis. But accommodating higher λ over the same infrastructure translates to a lower cost of running the Web service (with this reduction shown on the y -axis).

Thus, tolerating small amounts of loss in model update traffic can result in significant benefits for colocated services, while not deteriorating convergence.

8. Conclusion

In this paper, we present a novel analysis for a general model of distributed machine learning, under a realistic unreliable communication model. We present a novel theoretical analysis for such a scenario, and evaluated it while training neural networks on both image and natural language datasets. We also provided a case study of application collocation, to illustrate the potential benefit that can be provided by allowing learning algorithms to take advantage of unreliable communication channels.

Acknowledgements

Chen Yu, Hanlin Tang, and Ji Liu are in part supported by NSF CCF1718513, IBM faculty award, and NEC fellowship. Ce Zhang and the DS3Lab gratefully acknowledge the support from the Swiss National Science Foundation (Project Number 200021_184628), Swiss Data Science Center, Google Focused Research Award, Huawei, MeteoSwiss, Oracle Labs, Swisscom, Zurich Insurance, and the Chinese Scholarship Council.

References

- Error compensated quantized sgd and its applications to large-scale distributed optimization. *arXiv preprint arXiv:1806.08054*, 2018.
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pp. 265–283, 2016.
- Agarwal, A. and Duchi, J. C. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*, pp. 873–881, 2011.
- Alistarh, D., Allen-Zhu, Z., and Li, J. Byzantine stochastic gradient descent. *arXiv preprint arXiv:1803.08917*, 2018.
- Blanchard, P., Guerraoui, R., Stainer, J., et al. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Advances in Neural Information Processing Systems*, pp. 119–129, 2017.
- Blot, M., Picard, D., Cord, M., and Thome, N. Gossip training for deep learning. *arXiv preprint arXiv:1611.09726*, 2016.
- Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. Randomized gossip algorithms. *IEEE transactions on information theory*, 52(6):2508–2530, 2006.
- Colin, I., Bellet, A., Salmon, J., and Cléménçon, S. Gossip dual averaging for decentralized optimization of pairwise functions. *arXiv preprint arXiv:1606.02421*, 2016.
- Daily, J., Vishnu, A., Siegel, C., Warfel, T., and Amatya, V. Gossipgrad: Scalable deep learning using gossip communication based asynchronous gradient descent. *arXiv preprint arXiv:1803.05880*, 2018.
- Drumond, M., LIN, T., Jaggi, M., and Falsafi, B. Training dnns with hybrid block floating point. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 451–461. Curran Associates, Inc., 2018.
- Goldstein, T., Taylor, G., Barabın, K., and Sayre, K. Unwrapping admm: efficient distributed computing via transpose reduction. In *Artificial Intelligence and Statistics*, pp. 1151–1158, 2016.
- 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.
- Hajinezhad, D., Hong, M., Zhao, T., and Wang, Z. Nestt: A nonconvex primal-dual splitting method for distributed and stochastic optimization. In *Advances in Neural Information Processing Systems*, pp. 3215–3223, 2016.
- 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.
- He, L., Bian, A., and Jaggi, M. Cola: Decentralized linear learning. In *Advances in Neural Information Processing Systems*, pp. 4541–4551, 2018.
- Jiang, J., Cui, B., Zhang, C., and Yu, L. Heterogeneity-aware distributed parameter servers. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD ’17, pp. 463–478, New York, NY, USA, 2017. ACM. ISBN 978-1-4503-4197-4. doi: 10.1145/3035918.3035933. URL <http://doi.acm.org/10.1145/3035918.3035933>.
- Jin, P. H., Yuan, Q., Iandola, F., and Keutzer, K. How to scale distributed deep learning? *arXiv preprint arXiv:1611.04581*, 2016.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. 2009.
- Lan, G., Lee, S., and Zhou, Y. Communication-efficient algorithms for decentralized and stochastic optimization. *arXiv preprint arXiv:1701.03961*, 2017.
- Leblond, R., Pedregosa, F., and Lacoste-Julien, S. Asaga: asynchronous parallel saga. *arXiv preprint arXiv:1606.04809*, 2016.
- 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.
- Li, T. and Zhang, J.-F. Consensus conditions of multi-agent systems with time-varying topologies and stochastic communication noises. *IEEE Transactions on Automatic Control*, 55(9):2043–2057, 2010.
- Lian, X., Huang, Y., Li, Y., and Liu, J. Asynchronous parallel stochastic gradient for nonconvex optimization. In *Advances in Neural Information Processing Systems*, pp. 2737–2745, 2015.
- Lian, X., Zhang, C., Zhang, H., Hsieh, C.-J., Zhang, W., and Liu, J. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pp. 5336–5346, 2017a.

- Lian, X., Zhang, W., Zhang, C., and Liu, J. Asynchronous decentralized parallel stochastic gradient descent. *arXiv preprint arXiv:1710.06952*, 2017b.
- Lin, T., Stich, S. U., and Jaggi, M. Don't use large mini-batches, use local sgd. *arXiv preprint arXiv:1808.07217*, 2018.
- Lobel, I. and Ozdaglar, A. Distributed subgradient methods for convex optimization over random networks. *IEEE Transactions on Automatic Control*, 56(6):1291, 2011.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., et al. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.
- Nedić, A. and Olshevsky, A. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615, 2015.
- Nedic, A., Olshevsky, A., and Shi, W. Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4):2597–2633, 2017.
- Recht, B., Re, C., Wright, S., and Niu, F. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in neural information processing systems*, pp. 693–701, 2011.
- Renggli, C., Alistarh, D., and Hoeffler, T. Sparcml: High-performance sparse communication for machine learning. *arXiv preprint arXiv:1802.08021*, 2018.
- Scaman, K., Bach, F., Bubeck, S., Lee, Y. T., and Massoulié, L. Optimal algorithms for smooth and strongly convex distributed optimization in networks. *arXiv preprint arXiv:1702.08704*, 2017.
- Seide, F. and Agarwal, A. Cntk: Microsoft's open-source deep-learning toolkit. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2135–2135. ACM, 2016.
- Shen, Z., Mokhtari, A., Zhou, T., Zhao, P., and Qian, H. Towards more efficient stochastic decentralized learning: Faster convergence and sparse communication. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4624–4633, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018a. PMLR. URL <http://proceedings.mlr.press/v80/shen18a.html>.
- Shen, Z., Mokhtari, A., Zhou, T., Zhao, P., and Qian, H. Towards more efficient stochastic decentralized learning: Faster convergence and sparse communication. *arXiv preprint arXiv:1805.09969*, 2018b.
- Sirb, B. and Ye, X. Consensus optimization with delayed and stochastic gradients on decentralized networks. In *Big Data (Big Data), 2016 IEEE International Conference on*, pp. 76–85. IEEE, 2016.
- Sra, S., Yu, A. W., Li, M., and Smola, A. J. Adadelay: Delay adaptive distributed stochastic convex optimization. *arXiv preprint arXiv:1508.05003*, 2015.
- Stich, S. U. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.
- Stich, S. U., Cordonnier, J.-B., and Jaggi, M. Sparsified sgd with memory. In *Advances in Neural Information Processing Systems*, pp. 4452–4463, 2018.
- Tang, H., Gan, S., Zhang, C., Zhang, T., and Liu, J. Communication compression for decentralized training. In *Advances in Neural Information Processing Systems*, pp. 7663–7673, 2018a.
- Tang, H., Lian, X., Yan, M., Zhang, C., and Liu, J. D2: Decentralized training over decentralized data. *arXiv preprint arXiv:1803.07068*, 2018b.
- Thakur, R., Rabenseifner, R., and Gropp, W. Optimization of collective communication operations in mpich. *Int. J. High Perform. Comput. Appl.*, 19(1):49–66, February 2005. ISSN 1094-3420. doi: 10.1177/1094342005051521. URL <http://dx.doi.org/10.1177/1094342005051521>.
- Wang, J., Kolar, M., Srebro, N., and Zhang, T. Efficient distributed learning with sparsity. *arXiv preprint arXiv:1605.07991*, 2016.
- 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*, 2017.
- Xu, Z., Taylor, G., Li, H., Figueiredo, M., Yuan, X., and Goldstein, T. Adaptive consensus admm for distributed optimization. *arXiv preprint arXiv:1706.02869*, 2017.
- Yin, D., Chen, Y., Ramchandran, K., and Bartlett, P. Byzantine-robust distributed learning: Towards optimal statistical rates. *arXiv preprint arXiv:1803.01498*, 2018.
- Zhang, S., Zhang, C., You, Z., Zheng, R., and Xu, B. Asynchronous stochastic gradient descent for dnn training. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pp. 6660–6663. IEEE, 2013.
- Zhang, S., Choromanska, A. E., and LeCun, Y. Deep learning with elastic averaging sgd. In *Advances in Neural Information Processing Systems*, pp. 685–693, 2015.

Zhou, Z., Mertikopoulos, P., Bambos, N., Glynn, P., Ye, Y., Li, L.-J., and Fei-Fei, L. Distributed asynchronous optimization with unbounded delays: How slow can you go? In *International Conference on Machine Learning*, pp. 5965–5974, 2018.

Supplemental Materials

A. Notations

In order to unify notations, we define the following notations about gradient:

$$\mathbf{g}^{(i)}(\mathbf{x}_t^{(i)}; \mathbf{t}^{(i)}) = \nabla F_i(\mathbf{x}_t^{(i)}; \mathbf{t}^{(i)})$$

We define I_n as the $n \times n$ identity matrix, $\mathbf{1}_n$ as $(1, 1, \dots, 1)^\top$ and A_n as $\frac{1}{n}\mathbf{1}\mathbf{1}^\top$. Also, we suppose the packet drop rate is p .

The following equation is used frequently:

$$\text{Tr}(XA_nX^\top) = \text{Tr}\left(X\frac{\mathbf{1}\mathbf{1}^\top}{n}X^\top\right) = n \text{Tr}\left(\left(X\frac{\mathbf{1}}{n}\right)^\top X\frac{\mathbf{1}}{n}\right) = n\left(X\frac{\mathbf{1}}{n}\right)^\top X\frac{\mathbf{1}}{n} = n\left\|X\frac{\mathbf{1}}{n}\right\|^2 \quad (8)$$

A.1. Matrix Notations

We aggregate vectors into matrix, and using matrix to simplify the proof.

$$\begin{aligned} X_t &:= \left(\mathbf{x}_t^{(1)}, \mathbf{x}_t^{(2)}, \dots, \mathbf{x}_t^{(n)}\right) \\ V_t &:= \left(\mathbf{v}_t^{(1)}, \mathbf{v}_t^{(2)}, \dots, \mathbf{v}_t^{(n)}\right) \\ \mathbf{t} &:= \left(\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \dots, \mathbf{t}^{(n)}\right) \\ G(X_t; \mathbf{t}) &:= \left(\mathbf{g}^{(1)}(\mathbf{x}_t^{(1)}; \mathbf{t}^{(1)}), \dots, \mathbf{g}^{(n)}(\mathbf{x}_t^{(n)}; \mathbf{t}^{(n)})\right) \end{aligned}$$

A.2. Averaged Notations

We define averaged vectors as follows:

$$\bar{\mathbf{x}}_t := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_t^{(i)} \quad (9)$$

$$\bar{\mathbf{v}}_t := \frac{1}{n} \sum_{i=1}^n \mathbf{v}_t^{(i)} \quad (10)$$

$$\bar{\mathbf{g}}(X_t; \mathbf{t}) := \sum_{i=1}^n \mathbf{g}^{(i)}(\mathbf{x}_t^{(i)}; \mathbf{t}^{(i)}) \quad (11)$$

$$\bar{f}(X_t) := \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}_t^{(i)})$$

$$\bar{\mathbf{x}}_t := \bar{\mathbf{x}}_{t+1} - \bar{\mathbf{x}}_t$$

A.3. Block Notations

Remember in (2) and (3), we have divided models in blocks:

$$\begin{aligned} \mathbf{v}_t^{(i)} &= \left(\left(\mathbf{v}_t^{(i,1)}\right)^\top, \left(\mathbf{v}_t^{(i,2)}\right)^\top, \dots, \left(\mathbf{v}_t^{(i,n)}\right)^\top\right)^\top \\ \mathbf{x}_t^{(i)} &= \left(\left(\mathbf{x}_t^{(i,1)}\right)^\top, \left(\mathbf{x}_t^{(i,2)}\right)^\top, \dots, \left(\mathbf{x}_t^{(i,n)}\right)^\top\right)^\top, \forall i \in [n]. \end{aligned}$$

We do the some division on some other quantities, see following (the dimension of each block is the same as the corresponding block in $\mathbf{v}_t^{(i)}$):

$$\bar{\mathbf{x}}_t = \left(\left(\bar{\mathbf{x}}_t^{(1)}\right)^\top, \left(\bar{\mathbf{x}}_t^{(2)}\right)^\top, \dots, \left(\bar{\mathbf{x}}_t^{(n)}\right)^\top\right)^\top$$

$$\begin{aligned}
 \bar{\mathbf{v}}_t &= \left(\left(\bar{\mathbf{v}}_t^{(1)} \right)^\succ, \left(\bar{\mathbf{v}}_t^{(2)} \right)^\succ, \dots, \left(\bar{\mathbf{v}}_t^{(n)} \right)^\succ \right)^\succ \\
 \bar{\mathbf{x}}_t &= \left(\left(\bar{\mathbf{x}}_t^{(1)} \right)^\succ, \left(\bar{\mathbf{x}}_t^{(2)} \right)^\succ, \dots, \left(\bar{\mathbf{x}}_t^{(n)} \right)^\succ \right)^\succ \\
 \bar{\mathbf{g}}^{(i)}(\cdot) &= \left(\left(\bar{\mathbf{g}}^{(i,1)}(\cdot) \right)^\succ, \left(\bar{\mathbf{g}}^{(i,2)}(\cdot) \right)^\succ, \dots, \left(\bar{\mathbf{g}}^{(i,n)}(\cdot) \right)^\succ \right)^\succ \\
 \bar{\mathbf{g}}(X_{t; \cdot}) &= \left(\left(\bar{\mathbf{g}}^{(1)}(X_{t; \cdot}) \right)^\succ, \dots, \left(\bar{\mathbf{g}}^{(n)}(X_{t; \cdot}) \right)^\succ \right)^\succ \\
 \bar{r} f_i(\mathbf{x}_t^{(i)}) &= \left(\left(\bar{r}^{(1)} f_i(\mathbf{x}_t^{(i)}) \right)^\succ, \left(\bar{r}^{(2)} f_i(\mathbf{x}_t^{(i)}) \right)^\succ, \dots, \left(\bar{r}^{(n)} f_i(\mathbf{x}_t^{(i)}) \right)^\succ \right)^\succ \\
 \bar{r} f(X_t) &= \left(\left(\bar{r}^{(1)} f(X_t) \right)^\succ, \left(\bar{r}^{(2)} f(X_t) \right)^\succ, \dots, \left(\bar{r}^{(n)} f(X_t) \right)^\succ \right)^\succ \\
 \bar{r} f(\bar{\mathbf{x}}_t) &= \left(\left(\bar{r}^{(1)} f(\bar{\mathbf{x}}_t) \right)^\succ, \left(\bar{r}^{(2)} f(\bar{\mathbf{x}}_t) \right)^\succ, \dots, \left(\bar{r}^{(n)} f(\bar{\mathbf{x}}_t) \right)^\succ \right)^\succ.
 \end{aligned}$$

A.4. Aggregated Block Notations

Now, we can define some additional notations throughout the following proof

$$\begin{aligned}
 X_t^{(j)} &:= (\mathbf{x}_t^{(1,j)}, \mathbf{x}_t^{(2,j)}, \dots, \mathbf{x}_t^{(n,j)}) \\
 V_t^{(j)} &:= (\mathbf{v}_t^{(1,j)}, \mathbf{v}_t^{(2,j)}, \dots, \mathbf{v}_t^{(n,j)}) \\
 G^{(j)}(X_{t; \cdot}) &:= \left(\mathbf{g}^{(1,j)}(\mathbf{x}_t^{(1, \cdot)}; \cdot), \dots, \mathbf{g}^{(n,j)}(\mathbf{x}_t^{(n, \cdot)}; \cdot) \right)
 \end{aligned}$$

A.5. Relations between Notations

We have the following relations between these notations:

$$\bar{\mathbf{x}}_t^{(j)} = X_t^{(j)} \frac{\mathbf{1}}{n} \quad (12)$$

$$\bar{\mathbf{v}}_t^{(j)} = V_t^{(j)} \frac{\mathbf{1}}{n} \quad (13)$$

$$\bar{\mathbf{g}}^{(j)}(X_{t; \cdot}) = G^{(j)}(X_{t; \cdot}) \frac{\mathbf{1}}{n} \quad (14)$$

$$A_n A_n = A_n \quad (15)$$

$$V_t = X_t \quad \gamma G(X_{t; \cdot}) \quad (16)$$

$$V_t^{(j)} = X_t^{(j)} \quad \gamma G^{(j)}(X_{t; \cdot}) \quad (17)$$

A.6. Expectation Notations

There are different conditions when taking expectations in the proof, so we list these conditions below: $\mathbb{E}_{t,G}[\cdot]$ Denote taking the expectation over the **computing stochastic Gradient** procedure at t th iteration on condition of the history information before the t th iteration. $\mathbb{E}_{t,P}[\cdot]$ Denote taking the expectation over the **Package drop in sending and receiving blocks** procedure at t th iteration on condition of the history information before the t th iteration and the SGD procedure at the t th iteration. $\mathbb{E}_t[\cdot]$ Denote taking the expectation over all procedure during the t th iteration on condition of the history information before the t th iteration. $\mathbb{E}[\cdot]$ Denote taking the expectation over all history information.

B. Proof to Theorem 1

The critical part for a decentralized algorithm to be successful, is that local model on each node will converge to their average model. We summarize this critical property by the next lemma.

Lemma 3. From the updating rule (4) and Assumption 1, we have

$$\sum_{s=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_{s+1}^{(i)} - \bar{\mathbf{x}}_{s+1} \right\|^2 \leq \frac{2\gamma^2 n \sigma^2 T C_1}{(1 - \frac{\rho}{\beta})^2} + \frac{6n\zeta^2 T C_1}{(1 - \frac{\rho}{\beta})^2}, \quad (18)$$

where $C_1 := \left(1 - \frac{6L_P^2 \gamma^2}{(1 - \frac{\rho}{\beta})^2}\right)^{-1}$ and $\beta = \alpha_1 - \alpha_2$.

We will prove this critical property first. Then, after proving some lemmas, we will prove the final theorem. During the proof, we will use properties of weighted matrix $W_t^{(j)}$ which is showed in **Section D**.

B.1. Proof of Lemma 3

Proof to Lemma 3. According to updating rule (4) and Assumption 1, we have

$$\begin{aligned} X_{t+1}^{(j)} &= V_t^{(j)} W_t^{(j)} \\ &= \left(X_t^{(j)} - \gamma G^{(j)}(X_t; t) \right) W_t^{(j)} \\ &= X_1^{(j)} \prod_{r=1}^t W_r^{(j)} - \gamma \sum_{s=1}^t G^{(j)}(X_s; s) \prod_{r=s}^t W_r^{(j)} \\ &= -\gamma \sum_{s=1}^t G^{(j)}(X_s; s) \prod_{r=s}^t W_r^{(j)}. \quad (\text{due to } X_1 = 0) \end{aligned} \quad (19)$$

We also have

$$\sum_{i=1}^n \left\| \mathbf{x}_{t+1}^{(i,j)} - \bar{\mathbf{x}}_{t+1}^{(j)} \right\|^2 = \left\| X_{t+1}^{(j)} - X_{t+1}^{(j)} \frac{\mathbf{1}}{n} \mathbf{1}_n \right\|_F^2 = \left\| X_{t+1}^{(j)} - X_{t+1}^{(j)} A_n \right\|_F^2 \quad (20)$$

Combing (19) and (20) together, and define

$$H_{t,s}^{(j)} := G^{(j)}(X_s; s) \prod_{r=s}^t W_r^{(j)}$$

we get

$$\begin{aligned} \sum_{i=1}^n \left\| \left(\mathbf{x}_{t+1}^{(i,j)} - \bar{\mathbf{x}}_{t+1}^{(j)} \right) \right\|^2 &= \left\| X_{t+1}^{(j)} (I_n - A_n) \right\|_F^2 \\ &= \gamma^2 \left\| \sum_{s=1}^t H_{t,s}^{(j)} (I_n - A_n) \right\|_F^2 \\ &= \gamma^2 \text{Tr} \left((I_n - A_n) \sum_{s=1}^t \left(H_{t,s}^{(j)} \right)^\succ \sum_{s^0=1}^t H_{t,s^0}^{(j)} (I - A_n) \right) \\ &= \gamma^2 \sum_{s,s^0=1}^t \text{Tr} \left((I_n - A_n) \left(H_{t,s}^{(j)} \right)^\succ H_{t,s^0}^{(j)} (I - A_n) \right) \\ &= \frac{\gamma^2}{2} \sum_{s,s^0=1}^t \left(k_{s,s^0} \left\| H_{t,s}^{(j)} (I_n - A_n) \right\|_F^2 + \frac{1}{k_{s,s^0}} \left\| H_{t,s^0}^{(j)} (I_n - A_n) \right\|_F^2 \right), \end{aligned} \quad (21)$$

where k_{s,s^0} is a scale factor that is to be computed later. The last inequality is because $2 \text{Tr}(A^\succ B) \leq k k_A k_B^2 + \frac{1}{k} k_B k_A^2$ for any matrix A and B .

For $\left\|H_{t,s}^{(j)}(I_n - A_n)\right\|_F^2$, we have

$$\begin{aligned} & \mathbb{E} \left\|H_{t,s}^{(j)}(I_n - A_n)\right\|_F^2 \\ &= \mathbb{E} \operatorname{Tr} \left(G^{(j)}(X_{s^0}; s) W_s^{(j)} - W_t^{(j)}(I_n - A_n) \left(W_t^{(j)}\right)^{\succ} \left(W_s^{(j)}\right)^{\succ} \left(G^{(j)}(X_{s^0}; s)\right)^{\succ} \right). \end{aligned} \quad (22)$$

Now we can take expectation from time $t - 1$ back to time $s - 1$. When taking expectation of time t , we only need to compute $\mathbb{E} \left[W_t^{(j)}(I_n - A_n) \left(W_t^{(j)}\right)^{\succ} \right]$. From Lemma 7 and Lemma 8, this is just $(\alpha_1 - \alpha_2)(I_n - A_n)$. Applying this to (22), we can get the similar form except replacing t by $t - 1$ and multiplying by factor $\alpha_1 - \alpha_2$. Therefore, we have the following:

$$\begin{aligned} \mathbb{E} \left\|H_{t,s}^{(j)}(I_n - A_n)\right\|_F^2 &= (\alpha_1 - \alpha_2)^{t-s+1} \mathbb{E} \left\|G^{(j)}(X_{s^0}; s)(I_n - A_n)\right\|_F^2 \\ &\quad (\alpha_1 - \alpha_2)^{t-s} \mathbb{E} \left\|G^{(j)}(X_{s^0}; s)(I_n - A_n)\right\|_F^2. \end{aligned}$$

The last inequality comes from $\alpha_2 < \alpha_1$ and $\beta = \alpha_1 - \alpha_2$ is defined in Theorem 1.

Then (21) becomes

$$\begin{aligned} & \sum_{i=1}^n \left\| \left(\mathbf{x}_{t+1}^{(i,j)} - \bar{\mathbf{x}}_{t+1}^{(j)} \right) \right\|^2 \\ & \frac{\gamma^2}{2} \sum_{s,s^0=1}^t \left(k_{s,s^0} \beta^{t-s} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 + \frac{1}{k_{s,s^0}} \beta^{t-s^0} \left\| G^{(j)}(X_{s^0}; \xi_{s^0})(I_n - A_n) \right\|_F^2 \right). \end{aligned}$$

So if we choose $k_{s,s^0} = \beta^{\frac{s-s^0}{2}}$, the above inequality becomes

$$\begin{aligned} & \sum_{i=1}^n \left\| \left(\mathbf{x}_{t+1}^{(i,j)} - \bar{\mathbf{x}}_{t+1}^{(j)} \right) \right\|^2 \\ & \frac{\gamma^2}{2} \sum_{s,s^0=1}^t \left(\beta^{\frac{2t-s-s^0}{2}} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 + \beta^{\frac{2t-s^0-s}{2}} \left\| G^{(j)}(X_{s^0}; \xi_{s^0})(I_n - A_n) \right\|_F^2 \right) \\ &= \frac{\gamma^2 \beta^t}{2} \sum_{s,s^0=1}^t \beta^{-\frac{s-s^0}{2}} \left(\left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 + \left\| G^{(j)}(X_{s^0}; \xi_{s^0})(I_n - A_n) \right\|_F^2 \right) \\ &= \gamma^2 \beta^t \sum_{s,s^0=1}^t \beta^{-\frac{s-s^0}{2}} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 \\ &= \gamma^2 \sum_{s=1}^t \beta^{\frac{t-s}{2}} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 \sum_{s^0=1}^t \beta^{\frac{t-s^0}{2}} \\ & \frac{\gamma^2}{1-\beta} \sum_{s=1}^t \beta^{\frac{t-s}{2}} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 \end{aligned} \quad (23)$$

We also have:

$$\begin{aligned} & \sum_{j=1}^n \mathbb{E} \left\| G^{(j)}(X_{s^0}; s)(I_n - A_n) \right\|_F^2 \\ &= \sum_{j=1}^n \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{s^0}; t) \right\|^2 \end{aligned}$$

$$\begin{aligned}
 & \sum_{j=1}^n \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(j)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\
 & + 3 \sum_{j=1}^n \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 + 6L \sum_{j=1}^n \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 \quad (\text{using (31)}) \\
 & = \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - r f^{(i)}(\mathbf{x}_t^{(i)}) \right\|^2 \\
 & + 3 \sum_{i=1}^n \left\| r f^{(i)}(\bar{\mathbf{x}}_t) - r f(\bar{\mathbf{x}}_t) \right\|^2 + 6L \sum_{i=1}^n \left\| \mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right\|^2 \\
 & n\sigma^2 + 6L \sum_{i=1}^n \left\| \bar{\mathbf{x}}_s - \mathbf{x}_s^{(i)} \right\|^2 + 3n\zeta^2
 \end{aligned}$$

From the inequality above and (23) we have

$$\begin{aligned}
 & \sum_{j=1}^n \sum_{s=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_{s+1}^{(i,j)} - \bar{\mathbf{x}}_{s+1}^{(j)} \right\|^2 \\
 & \left(\frac{\gamma^2 n \sigma^2}{1 - \rho/\beta} + \sum_{j=1}^n \frac{3n\zeta^2}{1 - \rho/\beta} \right) \sum_{s=1}^T \sum_{r=1}^s \beta^{\frac{s-r}{2}} + \frac{6L^2 \gamma^2}{1 - \rho/\beta} \sum_{s=1}^T \sum_{r=1}^s \sum_{i=1}^n \beta^{\frac{s-r}{2}} \left\| \bar{\mathbf{x}}_r - \mathbf{x}_r^{(i)} \right\|^2 \\
 & \frac{\gamma^2 n \sigma^2 T}{(1 - \rho/\beta)^2} + \frac{3n\zeta^2 T}{(1 - \rho/\beta)^2} + \frac{6L^2 \gamma^2}{1 - \rho/\beta} \sum_{i=1}^n \sum_{s=1}^T \sum_{r=1}^s \beta^{\frac{s-r}{2}} \left\| \bar{\mathbf{x}}_r - \mathbf{x}_r^{(i)} \right\|^2 \\
 & = \frac{\gamma^2 n \sigma^2 T}{(1 - \rho/\beta)^2} + \frac{3n\zeta^2 T}{(1 - \rho/\beta)^2} + \frac{6L^2 \gamma^2}{1 - \rho/\beta} \sum_{i=1}^n \sum_{r=1}^T \sum_{s=r}^T \beta^{\frac{s-r}{2}} \left\| \bar{\mathbf{x}}_r - \mathbf{x}_r^{(i)} \right\|^2 \\
 & = \frac{\gamma^2 n \sigma^2 T}{(1 - \rho/\beta)^2} + \frac{3n\zeta^2 T}{(1 - \rho/\beta)^2} + \frac{6L^2 \gamma^2}{1 - \rho/\beta} \sum_{i=1}^n \sum_{r=1}^T \sum_{s=0}^r \beta^{\frac{s}{2}} \left\| \bar{\mathbf{x}}_r - \mathbf{x}_r^{(i)} \right\|^2 \\
 & \frac{\gamma^2 n \sigma^2 T}{(1 - \rho/\beta)^2} + \frac{3n\zeta^2 T}{(1 - \rho/\beta)^2} + \frac{6L^2 \gamma^2}{(1 - \rho/\beta)^2} \sum_{i=1}^n \sum_{r=1}^T \left\| \bar{\mathbf{x}}_r - \mathbf{x}_r^{(i)} \right\|^2.
 \end{aligned}$$

If γ is small enough that satisfies $\left(1 - \frac{6L^2 \gamma^2}{(1 - \rho/\beta)^2}\right) > 0$, then we have

$$\left(1 - \frac{6L^2 \gamma^2}{(1 - \rho/\beta)^2}\right) \sum_{s=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_s^{(i)} - \bar{\mathbf{x}}_s \right\|^2 \leq \frac{\gamma^2 n \sigma^2 T}{(1 - \rho/\beta)^2} + \frac{3n\zeta^2 T}{(1 - \rho/\beta)^2}.$$

Denote $C_1 := \left(1 - \frac{6L^2 \gamma^2}{(1 - \rho/\beta)^2}\right)^{-1}$, then we have

$$\sum_{s=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_s^{(i)} - \bar{\mathbf{x}}_s \right\|^2 \leq \frac{2\gamma^2 n \sigma^2 T C_1}{(1 - \rho/\beta)^2} + \frac{6n\zeta^2 T C_1}{(1 - \rho/\beta)^2}.$$

□

B.2. Proof to Theorem 1

Lemma 4. From the updating rule (4) and Assumption 1, we have

$$\begin{aligned}
 \mathbb{E}_{t,P} [k \|\bar{\mathbf{x}}_t\|^2] & = \frac{\alpha_2}{n} \sum_{j=1}^n \text{Tr} \left(\left(V_t^{(j)} \right)^{\succ} (I_n - A_n) V_t^{(j)} \right) + \gamma^2 k \bar{\mathbf{g}}(X_t; \mathbf{x}_t)^2, \\
 \mathbb{E}_{t,P} [\|\bar{\mathbf{x}}_t\|] & = \gamma \bar{\mathbf{g}}(X_t; \mathbf{x}_t).
 \end{aligned}$$

Proof. We begin with $\mathbb{E}_{t,P}[k \|\bar{\mathbf{x}}_t\|^2]$:

$$\begin{aligned}
 \mathbb{E}_{t,P}[k \|\bar{\mathbf{x}}_t\|^2] &= \sum_{j=1}^n \mathbb{E}_{t,P} \left[\left\| \begin{matrix} (j) \\ \bar{\mathbf{x}}_t \end{matrix} \right\|^2 \right] \\
 &\stackrel{(12)}{=} \sum_{j=1}^n \mathbb{E}_{t,P} \left[\left\| \begin{pmatrix} X_{t+1}^{(j)} & X_t^{(j)} \\ \mathbf{1}_n \end{pmatrix} \right\|^2 \right] \\
 &= \sum_{j=1}^n \mathbb{E}_{t,P} \left[\left\| \begin{pmatrix} V_t^{(j)} W_t^{(j)} & X_t^{(j)} \\ \mathbf{1}_n \end{pmatrix} \right\|^2 \right] \\
 &\stackrel{(8)}{=} \frac{1}{n} \sum_{j=1}^n \mathbb{E}_{t,P} \left[\text{Tr} \left(\begin{pmatrix} V_t^{(j)} W_t^{(j)} & X_t^{(j)} \\ \mathbf{1}_n \end{pmatrix} A_n \left(\begin{matrix} (W_t^{(j)})^\succ & (V_t^{(j)})^\succ & (X_t^{(j)})^\succ \end{matrix} \right) \right) \right] \\
 &= \frac{1}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} \mathbb{E}_{t,P} \left[W_t^{(j)} A_n \left(W_t^{(j)} \right)^\succ \right] \left(V_t^{(j)} \right)^\succ \right) \\
 &\quad + \frac{2}{n} \sum_{j=1}^n \text{Tr} \left(X_t^{(j)} A_n \mathbb{E}_{t,P} \left[\left(W_t^{(j)} \right)^\succ \right] \left(V_t^{(j)} \right)^\succ \right) + \frac{1}{n} \sum_{j=1}^n \text{Tr} \left(X_t^{(j)} A_n \left(X_t^{(j)} \right)^\succ \right) \\
 &= \frac{\alpha_2}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} \left(I_n \quad A_n \right) \left(V_t^{(j)} \right)^\succ \right) + \frac{1}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} A_n \left(V_t^{(j)} \right)^\succ \right) \\
 &\quad + \frac{2}{n} \sum_{j=1}^n \text{Tr} \left(X_t^{(j)} A_n \left(V_t^{(j)} \right)^\succ \right) + \frac{1}{n} \sum_{j=1}^n \text{Tr} \left(X_t^{(j)} A_n \left(X_t^{(j)} \right)^\succ \right), \tag{24}
 \end{aligned}$$

where for the last two equations, we use Lemma (8), Lemma(6), and (15). From (16), we can obtain the following equation:

$$\begin{aligned}
 V_t^{(j)} A_n \left(V_t^{(j)} \right)^\succ &= X_t^{(j)} A_n \left(X_t^{(j)} \right)^\succ - \gamma G^{(j)}(X_{t; t}) A_n \left(X_t^{(j)} \right)^\succ - \gamma X_t^{(j)} A_n \left(G^{(j)}(X_{t; t}) \right)^\succ \\
 &\quad + \gamma^2 G^{(j)}(X_{t; t}) A_n \left(G^{(j)}(X_{t; t}) \right)^\succ \\
 X_t^{(j)} A_n \left(V_t^{(j)} \right)^\succ &= X_t^{(j)} A_n \left(X_t^{(j)} \right)^\succ - \gamma G^{(j)}(X_{t; t}) A_n \left(X_t^{(j)} \right)^\succ, \tag{25}
 \end{aligned}$$

From the property of trace, we have:

$$\text{Tr} \left(G^{(j)}(X_{t; t}) A_n \left(X_t^{(j)} \right)^\succ \right) = \text{Tr} \left(X_t^{(j)} A_n^\succ \left(G^{(j)}(X_{t; t}) \right)^\succ \right) = \text{Tr} \left(X_t^{(j)} A_n \left(G^{(j)}(X_{t; t}) \right)^\succ \right). \tag{26}$$

Combing (24), (25) and (26), we have

$$\begin{aligned}
 \mathbb{E}_{t,P}[k \|\bar{\mathbf{x}}_t\|^2] &= \frac{\alpha_2}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} \left(I_n \quad A_n \right) \left(V_t^{(j)} \right)^\succ \right) + \frac{\gamma^2}{n} \sum_{j=1}^n \text{Tr} \left(G^{(j)}(X_{t; t}) A_n \left(G^{(j)}(X_{t; t}) \right)^\succ \right) \\
 &\stackrel{(8)}{=} \frac{\alpha_2}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} \left(I_n \quad A_n \right) \left(V_t^{(j)} \right)^\succ \right) + \gamma^2 \sum_{j=1}^n \left\| G^{(j)}(X_{t; t}) \frac{\mathbf{1}_n}{n} \right\|^2 \\
 &\stackrel{(14)}{=} \frac{\alpha_2}{n} \sum_{j=1}^n \text{Tr} \left(V_t^{(j)} \left(I_n \quad A_n \right) \left(V_t^{(j)} \right)^\succ \right) + \gamma^2 k \bar{\mathbf{g}}(X_{t; t}) k^2.
 \end{aligned}$$

For $\mathbb{E}_{t,P}[\bar{\mathbf{x}}_t]$, we first compute $\mathbb{E}_{t,P}[\begin{matrix} (j) \\ \bar{\mathbf{x}}_t \end{matrix}]$, ($j \in [n]$).

$$\begin{aligned}
 \mathbb{E}_{t,P}[\begin{matrix} (j) \\ \bar{\mathbf{x}}_t \end{matrix}] &= \mathbb{E}_{t,P}[\bar{\mathbf{x}}_{t+1}^{(j)}] - \mathbb{E}_{t,P}[\bar{\mathbf{x}}_t^{(j)}] \\
 &= \mathbb{E}_{t,P} \left[X_t^{(j)} \right] \frac{\mathbf{1}}{n} - \bar{\mathbf{x}}_t^{(j)}
 \end{aligned}$$

$$\begin{aligned}
 &= V_t^{(j)} \mathbb{E}_{t,P} \left[W_t^{(j)} \right] \frac{\mathbf{1}}{n} \bar{\mathbf{x}}_t^{(j)} \\
 &\stackrel{\text{Lemma (6)}}{=} V_t^{(j)} (\alpha_1 I_n + (1 - \alpha_1) A_n) \frac{\mathbf{1}}{n} \bar{\mathbf{x}}_t^{(j)} \\
 &= V_t^{(j)} \frac{\mathbf{1}}{n} \bar{\mathbf{x}}_t^{(j)} \\
 &= \mathbf{v}_t^{(j)} \bar{\mathbf{x}}_t^{(j)} \\
 &= \gamma \bar{\mathbf{g}}^{(j)}(X_{t; t}),
 \end{aligned}$$

which immediately leads to $\mathbb{E}_{t,P}[\bar{\mathbf{x}}_t] = \gamma \bar{\mathbf{g}}(X_{t; t})$. \square

Lemma 5. *From the updating rule (4) and Assumption 1, we have*

$$\sum_{j=1}^n \mathbb{E}_{t,G} \left[\text{Tr} \left(V_t^{(j)} (I_n - A_n) (V_t^{(j)})^\top \right) \right] (2 + 12\gamma^2 L) \sum_{i=1}^n \left\| \left(\mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right) \right\|^2 + 6n\gamma^2 \zeta^2 + 2n\gamma^2 \sigma^2.$$

Proof.

$$\begin{aligned}
 \text{Tr} \left(V_t^{(j)} (I_n - A_n) (V_t^{(j)})^\top \right) &= \text{Tr} \left(V_t^{(j)} (V_t^{(j)})^\top \right) - \text{Tr} \left(V_t^{(j)} A_n (V_t^{(j)})^\top \right) \\
 &\stackrel{(8)}{=} \left\| V_t^{(j)} \right\|_F^2 - n \left\| V_t^{(j)} \frac{\mathbf{1}_n}{n} \right\|^2 \\
 &= \sum_{i=1}^n \left(\left\| \mathbf{v}_t^{(i,j)} \right\|^2 - \left\| \bar{\mathbf{v}}_t^{(j)} \right\|^2 \right) \\
 &= \sum_{i=1}^n \left\| \mathbf{v}_t^{(i,j)} - \bar{\mathbf{v}}_t^{(j)} \right\|^2,
 \end{aligned} \tag{27}$$

the last equation above is because

$$\sum_{i=1}^n k \mathbf{a}_i k^2 - \left\| \sum_{i=1}^n \frac{\mathbf{a}_i}{n} \right\|^2 = \sum_{i=1}^n \left\| \mathbf{a}_i - \sum_{k=1}^n \frac{\mathbf{a}_k}{n} \right\|^2. \tag{28}$$

Since

$$\begin{aligned}
 \mathbf{v}_t^{(j)} - \bar{\mathbf{v}}_t^{(j)} &= \bar{\mathbf{x}}_t^{(j)} - \gamma \bar{\mathbf{g}}^{(j)}(X_{t; t}) \\
 \mathbf{v}_t^{(i,j)} - \bar{\mathbf{v}}_t^{(j)} &= \left(\mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right) - \gamma \left(\mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{t; t}) \right),
 \end{aligned}$$

we have the following:

$$\begin{aligned}
 \sum_{i=1}^n \left\| \mathbf{v}_t^{(i,j)} - \bar{\mathbf{v}}_t^{(j)} \right\|^2 &= \sum_{i=1}^n \left\| \left(\mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right) - \gamma \left(\mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{t; t}) \right) \right\|^2 \\
 &= 2 \sum_{i=1}^n \left\| \left(\mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right) \right\|^2 + 2\gamma^2 \sum_{i=1}^n \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{t; t}) \right\|^2,
 \end{aligned} \tag{29}$$

where the last inequality comes from $\|k\mathbf{a} + b\|^2 \leq 2k\mathbf{a}k^2 + 2kbk^2$.

For $\left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{t; t}) \right\|^2$, we have

$$\sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(X_{t; t}) \right\|^2$$

$$\begin{aligned}
 & \stackrel{(28)}{=} \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) \right\|^2 - n \mathbb{E}_{t,G} \left\| \bar{\mathbf{g}}^{(j)}(X_t; \bar{\mathbf{x}}_t) \right\|^2 \\
 & = \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \left(\mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right) + r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\
 & \quad - n \mathbb{E}_{t,G} \left\| \left(\bar{\mathbf{g}}^{(j)}(X_t; \bar{\mathbf{x}}_t) - \bar{r}^{(j)} f(X_t) \right) + \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & = \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 + \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 - n \mathbb{E}_{t,G} \left\| \bar{\mathbf{g}}^{(j)}(X_t; \bar{\mathbf{x}}_t) - \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & \quad - n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 + 2 \sum_{i=1}^n \mathbb{E}_{t,G} \left[\left\langle \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}), r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\rangle \right] \\
 & \quad - 2n \mathbb{E}_{t,G} \left[\left\langle \bar{\mathbf{g}}^{(j)}(X_t; \bar{\mathbf{x}}_t) - \bar{r}^{(j)} f(X_t), \bar{r}^{(j)} f(X_t) \right\rangle \right] \\
 & = \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 + \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 - n \mathbb{E}_{t,G} \left\| \bar{\mathbf{g}}^{(j)}(X_t; \bar{\mathbf{x}}_t) - \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & \quad - n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & = \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \bar{\mathbf{x}}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 + \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 - n \left\| \bar{r}^{(j)} f(X_t) \right\|^2. \tag{30}
 \end{aligned}$$

For $\sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 - n \left\| \bar{r}^{(j)} f(X_t) \right\|^2$, we have

$$\begin{aligned}
 & \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 - n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & \stackrel{(28)}{=} \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) - \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & = \sum_{i=1}^n \left\| \left(r^{(j)} f_i(\mathbf{x}_t^{(i)}) - r^{(j)} f_i(\bar{\mathbf{x}}_t) \right) - \left(\bar{r}^{(j)} f(X_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right) + \left(r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right) \right\|^2 \\
 & = 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\mathbf{x}_t^{(i)}) - r^{(j)} f_i(\bar{\mathbf{x}}_t) \right\|^2 + 3 \sum_{i=1}^n \left\| \bar{r}^{(j)} f(X_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & \quad + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \quad \left(\text{due to } \|\mathbf{a} + \mathbf{b} + \mathbf{c}\|^2 = \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 + \|\mathbf{c}\|^2 + 2\langle \mathbf{a}, \mathbf{b} \rangle + 2\langle \mathbf{a}, \mathbf{c} \rangle + 2\langle \mathbf{b}, \mathbf{c} \rangle \right) \\
 & = 3L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 3n \left\| \bar{r}^{(j)} f(X_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & = 3L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + \frac{3}{n} \left\| \sum_{k=1}^n \left(r^{(j)} f_k(\mathbf{x}_t^{(k)}) - r^{(j)} f_k(\bar{\mathbf{x}}_t) \right) \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & = 3L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 3 \sum_{k=1}^n \left\| r^{(j)} f_k(\mathbf{x}_t^{(k)}) - r^{(j)} f_k(\bar{\mathbf{x}}_t) \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & = 3L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 3L^2 \sum_{k=1}^n \left\| \mathbf{x}_t^{(k,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & = 6L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2.
 \end{aligned}$$

Taking the above inequality into (30), we get

$$\begin{aligned}
 & \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - \bar{\mathbf{g}}^{(j)}(\mathbf{X}_t; \mathbf{x}_t) \right\|^2 \\
 & \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 + 3 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & + 6L^2 \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2.
 \end{aligned} \tag{31}$$

Combining (29) and (31) together we have

$$\begin{aligned}
 \mathbb{E}_{t,G} \sum_{i=1}^n \left\| \mathbf{v}_t^{(i,j)} - \bar{\mathbf{v}}_t^{(j)} \right\|^2 & (2 + 12L^2\gamma^2) \sum_{i=1}^n \left\| \mathbf{x}_t^{(i,j)} - \bar{\mathbf{x}}_t^{(j)} \right\|^2 + 6\gamma^2 \sum_{i=1}^n \left\| r^{(j)} f_i(\bar{\mathbf{x}}_t) - r^{(j)} f(\bar{\mathbf{x}}_t) \right\|^2 \\
 & + 2\gamma^2 \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i,j)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2
 \end{aligned}$$

Summing j from 1 to n , we obtain the following:

$$\begin{aligned}
 \mathbb{E}_{t,G} \sum_{i=1}^n \left\| \mathbf{v}_t^{(i)} - \bar{\mathbf{v}}_t \right\|^2 & (2 + 12L^2\gamma^2) \sum_{i=1}^n \left\| \mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right\|^2 + 6\gamma^2 \sum_{i=1}^n k r f_i(\bar{\mathbf{x}}_t) - r f(\bar{\mathbf{x}}_t) k^2 \\
 & + 2\gamma^2 \sum_{i=1}^n \mathbb{E}_{t,G} \left\| \mathbf{g}^{(i)}(\mathbf{x}_t^{(i)}; \mathbf{x}_t^{(i)}) - r f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\
 & (2 + 12L^2\gamma^2) \sum_{i=1}^n \left\| \left(\mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right) \right\|^2 + 6n\gamma^2\zeta^2 + 2n\gamma^2\sigma^2.
 \end{aligned} \tag{32}$$

From (27) and (32), we have

$$\sum_{j=1}^n \mathbb{E}_{t,G} \text{Tr} \left(\mathbf{V}_t^{(j)} (I_n - A_n) \left(\mathbf{V}_t^{(j)} \right)^{\succ} \right) (2 + 12\gamma^2 L) \sum_{i=1}^n \left\| \left(\mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right) \right\|^2 + 6n\gamma^2\zeta^2 + 2n\gamma^2\sigma^2.$$

□

Proof to Theorem 1. From Lemma 6 and Lemma 8, we have

$$\begin{aligned}
 \mathbb{E}_{t,P} (W_t^{(j)}) & = \alpha_1 I_n + (1 - \alpha_1) A_n \\
 \mathbb{E}_{t,P} \left(W_t^{(j)} A_n \left(W_t^{(j)} \right)^{\succ} \right) & = \alpha_2 I_n + (1 - \alpha_2) A_n
 \end{aligned}$$

From the updating rule (4) and L-Lipschitz of f , we have

$$\begin{aligned}
 \mathbb{E}_{t,P} f(\bar{X}_{t+1}) - f(\bar{X}_t) & + \mathbb{E}_{t,P} h r f(\bar{X}_t), \quad \bar{\mathbf{x}}_t + \mathbb{E}_{t,P} \frac{L}{2} k \bar{\mathbf{x}}_t k^2 \\
 & \stackrel{\text{Lemma 4}}{=} f(\bar{X}_t) - \gamma h r f(\bar{X}_t), \gamma \bar{\mathbf{g}}(\mathbf{X}_t; \mathbf{x}_t) + \mathbb{E}_{t,P} \frac{L}{2} k \bar{\mathbf{x}}_t k^2 \\
 & \stackrel{\text{Lemma 4}}{=} f(\bar{X}_t) - \gamma h r f(\bar{X}_t), \gamma \bar{\mathbf{g}}(\mathbf{X}_t; \mathbf{x}_t) + \frac{\alpha_2 L}{2n} \sum_{j=1}^n \text{Tr} \left(\mathbf{V}_t^{(j)} (I_n - A_n) \left(\mathbf{V}_t^{(j)} \right)^{\succ} \right) \\
 & + \sum_{j=1}^n \frac{\gamma^2 L}{2} \mathbb{E}_{t,G} \left\| \bar{\mathbf{g}}^{(j)}(\mathbf{X}_t; \mathbf{x}_t) \right\|^2.
 \end{aligned}$$

So

$$\begin{aligned} \mathbb{E}_t f(\bar{X}_{t+1}) - f(\bar{X}_t) &= \gamma h r f(\bar{X}_t), \bar{r} f(X_t) i + \frac{\alpha_2 L}{2n} \sum_{j=1}^n \mathbb{E}_{t,G} \text{Tr} \left(V_t^{(j)} (I_n - A_n) (V_t^{(j)})^\top \right) \\ &+ \frac{\gamma^2 L}{2} \sum_{j=1}^n \mathbb{E}_{t,G} \left\| \bar{g}^{(j)}(X_t; t) \right\|^2. \end{aligned} \quad (33)$$

Since

$$\begin{aligned} \mathbb{E}_{t,G} \left\| \bar{g}^{(j)}(X_t; t) \right\|^2 &= \mathbb{E}_{t,G} \left\| \left(\bar{g}^{(j)}(X_t; t) - \bar{r}^{(j)} f(X_t) \right) + \bar{r}^{(j)} f(X_t) \right\|^2 \\ &= \mathbb{E}_{t,G} \left\| \bar{g}^{(j)}(X_t; t) - \bar{r}^{(j)} f(X_t) \right\|^2 + \mathbb{E} \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \\ &\quad + 2 \mathbb{E}_{t,G} \left\langle \bar{g}^{(j)}(X_t; t) - \bar{r}^{(j)} f(X_t), \bar{r}^{(j)} f(X_t) \right\rangle \\ &= \mathbb{E}_{t,G} \left\| \bar{g}^{(j)}(X_t; t) - \bar{r}^{(j)} f(X_t) \right\|^2 + \left\| \bar{r}^{(j)} f(X_t) \right\|^2, \end{aligned}$$

and

$$\begin{aligned} &\sum_{j=1}^n \mathbb{E}_{t,G} \left\| \bar{g}^{(j)}(X_t; t) - \bar{r}^{(j)} f(X_t) \right\|^2 \\ &= \frac{1}{n^2} \sum_{j=1}^n \mathbb{E}_{t,G} \left\| \sum_{i=1}^n \left(g^{(i,j)}(\mathbf{x}_t^{(i)}; t) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right) \right\|^2 \\ &= \frac{1}{n^2} \sum_{j=1}^n \sum_{i=1}^n \mathbb{E}_{t,G} \left\| g^{(i,j)}(\mathbf{x}_t^{(i)}; t) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\ &\quad + \frac{1}{n^2} \sum_{j=1}^n \mathbb{E}_{t,G} \sum_{i \neq i^0} \left\langle g^{(i,j)}(\mathbf{x}_t^{(i)}; t) - r^{(j)} f_i(\mathbf{x}_t^{(i)}), g^{(i^0,j)}(\mathbf{x}_t^{(i^0)}; t) - r^{(j)} f_{i^0}(\mathbf{x}_t^{(i^0)}) \right\rangle \\ &= \frac{1}{n^2} \sum_{j=1}^n \sum_{i=1}^n \mathbb{E}_{t,G} \left\| g^{(i,j)}(\mathbf{x}_t^{(i)}; t) - r^{(j)} f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E}_{t,G} \left\| g^{(i)}(\mathbf{x}_t^{(i)}; t) - r f_i(\mathbf{x}_t^{(i)}) \right\|^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \sigma^2 \\ &= \frac{\sigma^2}{n} \end{aligned}$$

then we have

$$\mathbb{E}_{t,G} \sum_{j=1}^n \left\| \bar{g}^{(j)}(X_t; t) \right\|^2 = \frac{\sigma^2}{n} + \sum_{j=1}^n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \quad (34)$$

Combining (33) and (34), we have

$$\begin{aligned} \mathbb{E}_t f(\bar{X}_{t+1}) - f(\bar{X}_t) &= \gamma h r f(\bar{X}_t), \bar{r} f(X_t) i + \frac{\alpha_2 L}{2n} \sum_{j=1}^n \mathbb{E}_{t,G} \text{Tr} \left(V_t^{(j)} (I_n - A_n) (V_t^{(j)})^\top \right) \\ &+ \frac{\gamma^2 L \sigma^2}{2n} + \frac{\gamma^2 L}{2} \sum_{j=1}^n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \end{aligned}$$

$$\begin{aligned}
 & f(\bar{X}_t) - \gamma h r f(\bar{X}_t), \bar{r} f(X_t) \Big| + \frac{\gamma^2 L \sigma^2}{2n} + \frac{\gamma^2 L}{2} \sum_{j=1}^n \left\| \bar{r}^{(j)} f(X_t) \right\|^2 \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2)}{2n} \sum_{i=1}^n \sum_{j=1}^n \left\| \begin{pmatrix} \mathbf{x}_t^{(i,j)} \\ \bar{\mathbf{x}}_t^{(j)} \end{pmatrix} \right\|^2 + 2\alpha_2 L^2 \gamma^2 \sigma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \quad (\text{due to Lemma 5}) \\
 & = f(\bar{X}_t) - \gamma h r f(\bar{X}_t), \bar{r} f(X_t) \Big| + \frac{\gamma^2 L \sigma^2}{2n} + \frac{\gamma^2 L}{2} \left\| \bar{r} f(X_t) \right\|^2 \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2)}{2n} \sum_{i=1}^n \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2 + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \\
 & = f(\bar{X}_t) - \frac{\gamma}{2} \left\| r f(\bar{X}_t) \right\|^2 - \frac{\gamma}{2} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\gamma}{2} \left\| r f(\bar{X}_t) - \bar{r} f(X_t) \right\|^2 + \frac{\gamma^2 L \sigma^2}{2n} \\
 & + \frac{\gamma^2 L}{2} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\alpha_2 L (2 + 12L^2 \gamma^2)}{2n} \sum_{i=1}^n \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2 + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2. \quad (35)
 \end{aligned}$$

Since

$$\begin{aligned}
 \left\| r f(\bar{X}_t) - \bar{r} f(X_t) \right\|^2 & = \frac{1}{n^2} \left\| \sum_{i=1}^n \begin{pmatrix} r f_i(\bar{X}_t) \\ r f_i(\mathbf{x}_t^{(i)}) \end{pmatrix} \right\|^2 \\
 & = \frac{1}{n} \sum_{i=1}^n \left\| \begin{pmatrix} r f_i(\bar{X}_t) \\ r f_i(\mathbf{x}_t^{(i)}) \end{pmatrix} \right\|^2 \\
 & = \frac{L^2}{n} \sum_{i=1}^n \left\| \begin{pmatrix} \bar{X}_t \\ \mathbf{x}_t^{(i)} \end{pmatrix} \right\|^2.
 \end{aligned}$$

So (35) becomes

$$\begin{aligned}
 \mathbb{E}_t f(\bar{X}_{t+1}) - f(\bar{X}_t) & = \frac{\gamma}{2} \left\| r f(\bar{X}_t) \right\|^2 - \frac{\gamma}{2} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\gamma^2 L \sigma^2}{2n} \\
 & + \frac{\gamma^2 L}{2} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\alpha_2 L (2 + 12L^2 \gamma^2) + L^2 \gamma}{2n} \sum_{i=1}^n \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2 + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \\
 & = f(\bar{X}_t) - \frac{\gamma}{2} \left\| r f(\bar{X}_t) \right\|^2 - \frac{\gamma(1 - L\gamma)}{2} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\gamma^2 L \sigma^2}{2n} + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2) + L^2 \gamma}{2n} \sum_{i=1}^n \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2.
 \end{aligned}$$

Taking the expectation over the whole history, the inequality above becomes

$$\begin{aligned}
 \mathbb{E} f(\bar{X}_{t+1}) - \mathbb{E} f(\bar{X}_t) & = \frac{\gamma}{2} \mathbb{E} \left\| r f(\bar{X}_t) \right\|^2 - \frac{\gamma(1 - L\gamma)}{2} \mathbb{E} \left\| \bar{r} f(X_t) \right\|^2 + \frac{\gamma^2 L \sigma^2}{2n} + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2) + L^2 \gamma}{2n} \sum_{i=1}^n \mathbb{E} \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2,
 \end{aligned}$$

which implies

$$\begin{aligned}
 \frac{\gamma}{2} \mathbb{E} \left\| r f(\bar{X}_t) \right\|^2 + \frac{\gamma(1 - L\gamma)}{2} \mathbb{E} \left\| \bar{r} f(X_t) \right\|^2 - \mathbb{E} f(\bar{X}_t) - \mathbb{E} f(\bar{X}_{t+1}) & + \frac{\gamma^2 L \sigma^2}{2n} + 2\alpha_2 L \sigma^2 \gamma^2 + 6\alpha_2 L \zeta^2 \gamma^2 \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2) + L^2 \gamma}{2n} \sum_{i=1}^n \mathbb{E} \left\| \begin{pmatrix} \mathbf{x}_t^{(i)} \\ \bar{\mathbf{x}}_t \end{pmatrix} \right\|^2. \quad (36)
 \end{aligned}$$

Summing up both sides of (36), it becomes

$$\begin{aligned}
 & \sum_{t=1}^T \left(\mathbb{E} \|r f(\bar{X}_t)\|^2 + (1 - L\gamma) \mathbb{E} \|\bar{r} f(X_t)\|^2 \right) \\
 & \frac{2(f(\mathbf{x}_0) - \mathbb{E}f(\mathbf{x}_{T+1}))}{\gamma} + \frac{\gamma L \sigma^2 T}{n} + 4\alpha_2 L \sigma^2 \gamma T + 12\alpha_2 L \zeta^2 \gamma T \\
 & + \frac{\alpha_2 L (2 + 12L^2 \gamma^2) + L^2 \gamma}{2n\gamma} \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right\|^2. \tag{37}
 \end{aligned}$$

According to Lemma 3, we have

$$\sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left\| \mathbf{x}_t^{(i)} - \bar{\mathbf{x}}_t \right\|^2 \leq \frac{2\gamma^2 n \sigma^2 T C_1}{(1 - \rho/\beta)^2} + \frac{6n\zeta^2 T C_1}{(1 - \rho/\beta)^2},$$

where $C_1 = \left(1 - \frac{6L^2\gamma^2}{(1 - \rho/\beta)^2}\right)^{-1}$. Combing the inequality above with (37) we get

$$\begin{aligned}
 & \frac{1}{T} \sum_{t=1}^T \left(\mathbb{E} \|r f(\bar{X}_t)\|^2 + (1 - L\gamma) \mathbb{E} \|\bar{r} f(X_t)\|^2 \right) \\
 & \frac{2f(\mathbf{0}) - 2f(\mathbf{x})}{\gamma T} + \frac{\gamma L \sigma^2}{n} + 4\alpha_2 L \gamma (\sigma^2 + 3\zeta^2) \\
 & + \frac{(\alpha_2 L \gamma (2 + 12L^2 \gamma^2) + L^2 \gamma^2) \sigma^2 C_1}{(1 - \rho/\beta)^2} + \frac{3(\alpha_2 L \gamma (2 + 12L^2 \gamma^2) + L^2 \gamma^2) \zeta^2 C_1}{(1 - \rho/\beta)^2}.
 \end{aligned}$$

□

C. Proof to Corollary 2

Proof to Corollary 2. Setting

$$\gamma = \frac{1 - \rho/\beta}{6L + 3(\sigma + \zeta) \rho \frac{\rho/\beta}{\alpha_2 T} + \frac{\rho/\beta}{\alpha_2 T} + \frac{\rho/\beta}{\alpha_2 T}},$$

then we have

$$\begin{aligned}
 & 1 - L\gamma & 0 \\
 & C_1 & 2 \\
 & 2 + 12L^2\gamma^2 & 4
 \end{aligned}$$

So (7) becomes

$$\begin{aligned}
 & \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|r f(\bar{X}_t)\|^2 \leq \frac{(2f(\mathbf{0}) - 2f(\mathbf{x}) + L)\sigma}{\rho \frac{\rho/\beta}{nT} (1 - \rho/\beta)} + \frac{(2f(\mathbf{0}) - 2f(\mathbf{x}) + L)(\sigma + \zeta)}{1 - \rho/\beta} \sqrt{\frac{\alpha_2}{T}} \\
 & + \frac{(2f(\mathbf{0}) - 2f(\mathbf{x}))L}{T} + \frac{L^2(\sigma^2 + \zeta^2)}{(\frac{T}{n} + \alpha_2 T)\sigma^2 + \alpha_2 T \zeta^2}, \\
 & \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|r f(\bar{X}_t)\|^2 \leq \frac{\sigma + \zeta}{(1 - \rho/\beta) \rho \frac{\rho/\beta}{nT}} + \frac{\sigma + \zeta}{(1 - \rho/\beta)} \sqrt{\frac{\alpha_2}{T}} + \frac{1}{T} + \frac{n(\sigma^2 + \zeta^2)}{(1 + n\alpha_2)\sigma^2 T + n\alpha_2 T \zeta^2}.
 \end{aligned}$$

□

D. Properties of Weighted Matrix $W_t^{(j)}$

In this section, we will give three properties of $W_t^{(j)}$, described by Lemma 6, Lemma 7 and Lemma 8.

Throughout this section, we will frequently use the following two fact: **Fact 1:** $\frac{1}{m+1} \binom{n}{m} = \frac{1}{n+1} \binom{n+1}{m+1}$. **Fact 2:** $\frac{1}{(m+1)(m+2)} \binom{n}{m} = \frac{1}{(n+1)(n+2)} \binom{n+2}{m+2}$.

Lemma 6. *Under the updating rule (4), there exists $\alpha_1 \in [0, 1]$, s.t., $\forall j \in [n]$, \forall time t ,*

$$\mathbb{E}_{t,P}[W_t^{(j)}] = \alpha_1 I_n + (1 - \alpha_1) A_n.$$

Proof. Because of symmetry, we will fix j , say, $j = 1$. So for simplicity, we omit superscript (j) for all quantities in this proof, the subscript t for W , and the subscript t, P for \mathbb{E} , because they do not affect the proof.

First we proof: $\forall \alpha_1, s.t.$

$$\mathbb{E}[W] = \alpha_1 I_n + (1 - \alpha_1) A_n. \quad (38)$$

Let us understand the meaning of the element of W . For the (k, l) th element W_{kl} . From $X_{t+1} = V_t W$, we know that, W_{kl} represents the portion that $v_t^{(l)}$ will be in $x_{t+1}^{(k)}$ (the block number j has been omitted, as stated before). For $v_t^{(l)}$ going into $x_{t+1}^{(k)}$, it should first sent from k , received by node b_t (also omit j), averaged with other j th blocks by node b_t , and at last sent from b_t to l . For all pairs (k, l) satisfied $k \neq l$, the expectations of W_{kl} are equivalent because of the symmetry (the same packet drop rate, and independency). For the same reason, the expectations of W_{kl} are also equivalent for all pairs (k, l) satisfied $k = l$. But for two situations that $k = l$ and $k \neq l$, the expectation need not to be equivalent. This is because when the sending end l is also the receiving end k , node l (or k) will always keep its own portion $v_t^{(l)}$ if l is also the node dealing with block j , which makes a slight different. □

Lemma 7. *Under the updating rule (4), there exists $\alpha_1 \in [0, 1]$, s.t., $\forall j \in [n]$, \forall time t ,*

$$\mathbb{E}_{t,P}[W_t^{(j)} W_t^{(j)\top}] = \alpha_1 I_n + (1 - \alpha_1) A_n.$$

Moreover, α_1 satisfies:

$$\alpha_1 = \frac{np + (1 - p)^n + nT_1 + nT_2 - 1}{n - 1},$$

where

$$T_1 = \frac{2(1 - p)^{n+1} (n+1)(1 - p)p^n (n+1)n(1 - p)^2 p^{n-1}/2 (1 - p)^{n+1}}{n(n+1)(1 - p)^2},$$

$$T_2 = \frac{1 - p^n - n(1 - p)p^{n-1} (1 - p)^n}{(n-1)(1 - p)}.$$

Proof. Similar to Lemma (6), we fix $j = 1$, and omit superscript (j) for all quantities in this proof, the subscript t for W and the subscript t, P for \mathbb{E} .

Also similar to Lemma (6), there exists α , s.t.

$$\mathbb{E}[W W^\top] = \alpha I_n + (1 - \alpha) A_n \quad (39)$$

The only thing left is to bound α . From (39), we know that $\alpha = \frac{n\mathbb{E}[W W^\top]_{(1,1)} - 1}{n - 1}$. After simple compute, we have

$\mathbb{E}[W W^\top]_{(1,1)} = \mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2\right]$. So we have the following:

$$\alpha = \frac{n\mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2\right] - 1}{n - 1}.$$

Therefore, bounding α_1 equals bounding $\mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2\right]$. Similar to Lemma (6), we denote the event "node 1 deal with the first block" by A .

Case 1: node 1 deal with the first block In this case, let's understand W again. node 1 average the 1st blocks it has received, then broadcast to all nodes. Therefore, for every node i who received this averaged block, $\mathbf{x}_{t+1}^{(i)}$ has the same value, in other words, the column i of W equals, or, $W_{1,i}$ equals to $W_{1,1}$. On the other hand, for every node i who did not receive this averaged block, they keep their origin model $v_i^{(i)}$. But $i \neq 1$ (because node 1 deal with this block, itself must receive its own block), which means $W_{1,i} = 0$.

Therefore, for $i \neq 1, i \in [n]$, if node i receive the averaged model, $W_{1,i} = W_{1,1}$. Otherwise, $W_{1,i} = 0$. Based on this fact, we can define the random variable B_i for $i \neq 1, i \in [n]$. $B_i = 1$ if node i receive the averaged block., $B_i = 0$ if node i does not receive the averaged block. Immediately, we can obtain the following equation:

$$\mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2 j A\right] = \mathbb{E}\left[W_{1,1}^2 \left(\sum_{i=2}^n B_i + 1\right) j A\right] = \mathbb{E}[W_{1,1}^2 j A] (1 + (n-1)\mathbb{E}[B_n j A]). \quad (40)$$

The last equation results from that A, B_2, \dots, B_n are independent and B_2, \dots, B_n are from identical distribution.

First let's compute $\mathbb{E}[W_{1,1}^2 j A]$. If node i received the 1st block from m ($0 \leq m \leq n-1$) nodes (except itself), then $W_{1,1} = 1/(m+1)$. The probability of this event is $\binom{n-1}{m}(1-p)^m p^{n-1-m}$. So we can obtain:

$$\begin{aligned} \mathbb{E}[W_{1,1}^2 j A] &= \sum_{m=0}^{n-1} \frac{1}{(m+1)^2} \binom{n-1}{m} (1-p)^m p^{n-1-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + \sum_{m=1}^{n-2} \frac{1}{(m+1)^2} \binom{n-1}{m} (1-p)^m p^{n-1-m} \\ &\quad \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + \sum_{m=1}^{n-2} \frac{2}{(m+1)(m+2)} \binom{n-1}{m} (1-p)^m p^{n-1-m} \\ &\stackrel{\text{Fact 2}}{=} \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + \sum_{m=1}^{n-2} \frac{2}{n(n+1)} \binom{n+1}{m+2} (1-p)^m p^{n-1-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + \frac{2}{n(n+1)} \sum_{m=3}^n \binom{n+1}{m} (1-p)^m p^{n+1-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + \frac{2}{n(n+1)(1-p)^2} \sum_{m=3}^n \binom{n+1}{m} (1-p)^m p^{n+1-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} \\ &\quad + \frac{2(1-p^{n+1} - (n+1)(1-p)p^n - (n+1)n(1-p)^2 p^{n-1}/2 - (1-p)^{n+1})}{n(n+1)(1-p)^2} \\ &= \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + T_1, \end{aligned}$$

where we denote $T_1 := \frac{2(1-p^{n+1} - (n+1)(1-p)p^n - (n+1)n(1-p)^2 p^{n-1}/2 - (1-p)^{n+1})}{n(n+1)(1-p)^2}$.

Next let's compute $\mathbb{E}[B_n j A]$. B_n is just a $0-1$ distribution, with success probability $1-p$. Therefore, $\mathbb{E}[B_n j A] = 1-p$.

Applying all these equations into (40), we can get:

$$\begin{aligned} \mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2 j A\right] &= \left(\frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + T_1\right) (1 + (n-1)(1-p)) \\ &\quad \frac{(1-p)^{n-1}}{n} + p^{n-1} + nT_1 \end{aligned} \quad (41)$$

Case 2: node 1 does not deal with the first block and node 1 does not receive averaged block We define a new event C , representing that node 1 does not receive the averaged block. So, **Case 2** equals the event $A \setminus C$. In this case, node 1 keeps its origin block $v_t^{(1)}$, which means $W_{1,1} = 1$.

Again due to symmetry, we can suppose that node n deal with the first block. Then we can use the method in **Case 1**. But in this case, we only use B_2, \dots, B_{n-1} , because node n must receive its own block and node 1 does not receive averaged block, and we use $W_{1,n}$ instead of $W_{1,1}$. Then, we obtain:

$$\mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2 j A, C\right] = 1 + \mathbb{E}\left[W_{1,n}^2 \left(\sum_{i=2}^{n-1} B_i + 1\right) j A, C\right] \quad (42)$$

$$= 1 + \mathbb{E}[W_{1,n}^2 j A, C] (1 + (n-2)\mathbb{E}[B_{n-1} j A, C]). \quad (43)$$

Here, we similarly have $\mathbb{E}[B_{n-1} j A, C] = 1-p$, but we need to compute $\mathbb{E}[W_{1,n}^2 j A, C]$. When the 1st block from node 1 is not received by node n , $W_{1,n} = 0$. If node 1's block is received, together with other m , ($0 \leq m \leq n-2$) nodes' blocks, then $W_{1,n} = 1/(m+2)$ (node n 's block is always received by itself). The probability of this event is $\binom{n-2}{m}(1-p)^{m+1}p^{n-2-m}$. Therefore,

$$\begin{aligned} \mathbb{E}[W_{1,n}^2 j A, C] &= \sum_{m=0}^{n-2} \frac{1}{(m+2)^2} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{(m+2)^2} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{(m+2)(m+1)} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &\stackrel{\text{Fact 2}}{=} \frac{1}{n^2} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{n(n-1)} \binom{n-2}{m+2} (1-p)^{m+1} p^{n-2-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n(n-1)} \sum_{m=2}^{n-1} \binom{n}{m} (1-p)^m p^{n-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n(n-1)(1-p)} \sum_{m=2}^{n-1} \binom{n}{m} (1-p)^m p^{n-m} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \frac{1-p^n - n(1-p)p^{n-1} - (1-p)^n}{n(n-1)(1-p)} \\ &= \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2, \end{aligned}$$

where $T_2 := \frac{1-p^n - n(1-p)p^{n-1} - (1-p)^n}{(n-1)(1-p)}$.

Applying these equations into (42), we get:

$$\begin{aligned} \mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2 j A, C\right] &= 1 + \left(\frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2\right) (1 + (n-2)(1-p)) \\ &= 1 + \frac{(1-p)^{n-1}}{n} + T_2 \end{aligned} \quad (44)$$

Case 3: node 1 does not deal with the first block and node 1 receives averaged block This is the event $C \setminus A$. Similar to the analysis above, we have:

$$\mathbb{E}\left[\sum_{i=1}^n W_{1,i}^2 j A, C\right] = \mathbb{E}\left[W_{1,1}^2 \left(\sum_{i=2}^n B_i + 2\right) j A, C\right]$$

$$= \mathbb{E} \left[W_{1,1}^2 j A, C \right] (2 + (n-2)\mathbb{E}[B_2 j A, C])$$

Similarly, we have $\mathbb{E}[B_2 j A, C] = 1 - p$. For $\mathbb{E} \left[W_{1,1}^2 j A, C \right]$, the argument is the same as $\mathbb{E} \left[W_{1,n}^2 j A, C \right]$ in **Case 2**. So, we have:

$$\mathbb{E} \left[W_{1,1}^2 j A, C \right] = \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2.$$

Applying these together, we can obtain:

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 j A, C \right] &= \left(\frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2 \right) (2 + (n-2)(1-p)) \\ &= \frac{(1-p)^{n-1}}{n} + T_2 \end{aligned} \quad (45)$$

Combined with three cases, $P(A) = 1/n$, $P(A, C) = p(n-1)/n$, and $P(A, C) = (1-p)(n-1)/n$, we have

$$\mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 \right] = \frac{1}{n} \mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 j A \right] + \frac{p(n-1)}{n} \mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 j A, C \right] + \frac{(1-p)(n-1)}{n} \mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 j A, C \right].$$

Combing the inequality above and (41) (44) (45) together, we get

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n W_{1,i}^2 \right] &\leq \frac{(1-p)^n}{n^2} + \frac{p^n}{n} + T_1 + \frac{p(n-1)}{n} + \frac{p(1-p)^{n-1}}{n} + T_2 p + \frac{(1-p)^n (n-1)}{n^2} + T_2 (1-p) \\ &= p + \frac{(1-p)^n}{n} + T_1 + T_2 \\ \alpha_1 &= \frac{np + (1-p)^n + nT_1 + nT_2 - 1}{n-1} \end{aligned}$$

□

Lemma 8. Under the updating rule (4), there exists $\alpha_2 \in [0, 1]$, s.t., $\forall j \in [n]$, \forall time t ,

$$\mathbb{E}_{t,P} \left[W_t^{(j)} A_n W_t^{(j)} \right] = \alpha_2 I_n + (1 - \alpha_2) A_n.$$

Moreover, α_2 satisfies:

$$\alpha_2 = \frac{p(1 + 2T_3) + (1-p)^{n-1}}{n} + \frac{2p(1-p)^n}{n} + \frac{p^n(1-p)}{n^2} + T_1 + T_2,$$

where

$$\begin{aligned} T_1 &= \frac{2(1-p)^{n+1} - (n+1)(1-p)p^n - (n+1)n(1-p)^2 p^{n-1}/2 - (1-p)^{n+1}}{n(n+1)(1-p)^2}, \\ T_2 &= \frac{1-p^n - n(1-p)p^{n-1} - (1-p)^n}{(n-1)(1-p)}, \\ T_3 &= \frac{n}{n-1} (1-p^{n-1} - (1-p)^{n-1}) + (1-p)^{n-1}. \end{aligned}$$

Proof. Similar to Lemma (6) and Lemma (7), we fix $j = 1$, and omit superscript (j) for all quantities in this proof, the subscript t for W and the subscript t, P for \mathbb{E} . And we still use A to denote the event "node 1 deal with the first block", use the binary random variable B_i to denote whether node i receive the averaged block. The definitions is the same to them in Lemma 7.

Again similar to Lemma (6), there exists $\alpha_2, s.t.$

$$\mathbb{E}[W A_n W^>] = \alpha_2 I_n + (1 - \alpha_2) A_n. \quad (46)$$

The only thing left is to bound α_2 . From (46), we know that $\alpha_2 = \frac{n \mathbb{E}[W A_n W^>]_{(1,1)}}{n-1}$. After simple compute, we have

$$\mathbb{E}[W A_n W^>]_{(1,1)} = \mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \right] / n. \text{ So we have the following:}$$

$$\alpha_2 = \frac{\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \right] - 1}{n-1}.$$

Therefore, bounding α_2 equals bounding $\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \right]$.

Case 1: node 1 deal with the first block In this case, $\sum_{i=1}^n W_{1,i} = W_{1,1} \left(1 + \sum_{i=2}^n B_i \right)$, which means, $\left(\sum_{i=1}^n W_{1,i} \right)^2 = W_{1,1}^2 \left(1 + \sum_{i=2}^n B_i \right)^2$. Similar to Lemma 7, A and $fB_i g_{i=2}^n$ are independent, so we have:

$$\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A \right] = \mathbb{E} [W_{1,1}^2 j A] \mathbb{E} \left[\left(1 + \sum_{i=2}^n B_i \right)^2 j A \right].$$

From Lemma 7, we have

$$\mathbb{E} [W_{1,1}^2 j A] = \frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + T_1.$$

For $\mathbb{E} \left[\left(1 + \sum_{i=2}^n B_i \right)^2 j A \right]$, since $fB_i g_{i=2}^n$ are independent, we have the following:

$$\begin{aligned} \mathbb{E} \left[\left(1 + \sum_{i=2}^n B_i \right)^2 j A \right] &= \left(\mathbb{E} \left[1 + \sum_{i=2}^n B_i j A \right] \right)^2 + \text{Var} \left[1 + \sum_{i=2}^n B_i j A \right] \\ &= (1 + (n-1)(1-p))^2 + (n-1)p(1-p) \end{aligned}$$

Combined these together, we obtain:

$$\begin{aligned} &\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A \right] \\ &= \left(\frac{1}{n^2} (1-p)^{n-1} + p^{n-1} + T_1 \right) \left((1 + (n-1)(1-p))^2 + (n-1)p(1-p) \right) \\ &= \frac{(1-p)^{n-1} (1 + (n-1)(1-p))^2}{n^2} + p^{n-1} (1 + (n-1)(1-p))^2 + \frac{p^n (1-p) + (1-p)^n p}{n} + n^2 T_1 \end{aligned}$$

Case 2: node 1 does not deal with the first block and node 1 does not receive averaged block In this case, $\sum_{i=1}^n W_{1,i} =$

$1 + W_{1,n} \left(\sum_{i=2}^n B_i + 1 \right)$ (suppose node n deal with the first block). So we have:

$$\left(\sum_{i=1}^n W_{1,i} \right)^2 = 1 + 2W_{1,n} \left(\sum_{i=2}^n B_i + 1 \right) + W_{1,n}^2 \left(\sum_{i=2}^n B_i + 1 \right)^2,$$

which means (notice $W_{1,n}$ and $fB_i \mathcal{G}_{i=2}^{n-1}$ are independent),

$$\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \middle| j A, C \right] = 1 + 2\mathbb{E} [W_{1,n} \middle| j A, C] \mathbb{E} \left[\sum_{i=2}^{n-1} B_i + 1 \middle| j A, C \right] + \mathbb{E} [W_{1,n}^2 \middle| j A, C] \mathbb{E} \left[\left(\sum_{i=2}^{n-1} B_i + 1 \right)^2 \middle| j A, C \right].$$

First let's consider $\mathbb{E}[W_{1,n} \middle| j A, C]$. Similar to the analysis of **Case 2** in Lemma 7 except instead first moment of second moment, we have:

$$\begin{aligned} \mathbb{E}[W_{1,n} \middle| j A, C] &= \sum_{m=0}^{n-2} \frac{1}{m+2} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &= \frac{1}{n} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{m+2} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &\quad - \frac{1}{n} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{m+1} \binom{n-2}{m} (1-p)^{m+1} p^{n-2-m} \\ &\stackrel{\text{Fact 1}}{=} \frac{1}{n} (1-p)^{n-1} + \sum_{m=0}^{n-3} \frac{1}{n-1} \binom{n-1}{m+1} (1-p)^{m+1} p^{n-2-m} \\ &= \frac{1}{n} (1-p)^{n-1} + \frac{1}{n-1} \sum_{m=1}^{n-2} \binom{n-1}{m} (1-p)^m p^{n-1-m} \\ &= \frac{1}{n} (1-p)^{n-1} + \frac{1}{n-1} (1-p^{n-1} - (1-p)^{n-1}) \\ &= \frac{T_3}{n}, \end{aligned}$$

where we denote $T_3 := \frac{n-1}{n} (1-p^{n-1} - (1-p)^{n-1}) + (1-p)^{n-1}$.

Next, from Lemma 7, we have

$$\mathbb{E} [W_{1,n}^2 \middle| j A, C] = \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2.$$

Next we deal with item with B_i . We have the following:

$$\begin{aligned} \mathbb{E} \left[\sum_{i=2}^{n-1} B_i + 1 \middle| j A, C \right] &= (n-2)(1-p) + 1 \\ \mathbb{E} \left[\left(\sum_{i=2}^{n-1} B_i + 1 \right)^2 \middle| j A, C \right] &= 1 + 2(n-2)\mathbb{E}[B_2 \middle| j A, C] + \mathbb{E} \left[\left(\sum_{i=2}^{n-1} B_i \right)^2 \middle| j A, C \right] \\ &= 1 + 2(n-2)(1-p) + \left(\mathbb{E} \left[\sum_{i=2}^{n-1} B_i \middle| j A, C \right] \right)^2 + \text{Var} \left[\sum_{i=2}^{n-1} B_i \middle| j A, C \right] \\ &= 1 + 2(n-2)(1-p) + (n-2)^2(1-p)^2 + (n-2)p(1-p) \end{aligned}$$

Combining those terms together we get

$$\begin{aligned} &\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \middle| j A, C \right] \\ &= 1 + \frac{(1-p)^{n-1}}{n^2} + \frac{(n-2)(1-p)^n}{n} + \frac{p(1-p)^n}{n} + nT_2 + 2T_3 \end{aligned}$$

Case 3: node 1 does not deal with the first block and node 1 receives averaged block In this case, $\sum_{i=1}^n W_{1,i} = W_{1,1} \left(\sum_{i=2}^{n-1} B_i + 2 \right)$. So we have:

$$\left(\sum_{i=1}^n W_{1,i} \right)^2 = W_{1,1}^2 \left(\sum_{i=2}^{n-1} B_i + 2 \right)^2,$$

which means (notice that $W_{1,1}$ and $fB_i G_{i=2}^{n-1}$ are independent)

$$\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A, C \right] = \mathbb{E} [W_{1,1}^2 j A, C] \mathbb{E} \left[\left(\sum_{i=2}^{n-1} B_i + 2 \right)^2 j A, C \right].$$

Similar to Lemma 7, $\mathbb{E} [W_{1,1}^2 j A, C]$ is the same as $\mathbb{E} [W_{1,n}^2 j A, C]$.

$$\mathbb{E} [W_{1,1}^2 j A, C] = \frac{1}{n^2} (1-p)^{n-1} + \frac{1}{n} T_2$$

Also, we have the following:

$$\begin{aligned} \mathbb{E} \left[\left(\sum_{i=2}^{n-1} B_i + 2 \right)^2 j A, C \right] &= \left(\mathbb{E} \left[\sum_{i=2}^{n-1} B_i + 2 j A, C \right] \right)^2 + \text{Var} \left[\sum_{i=2}^{n-1} B_i + 2 j A, C \right] \\ &= [(n-2)(1-p) + 2]^2 + (n-2)p(1-p) \end{aligned}$$

So we have

$$\begin{aligned} &\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A, C \right] \\ &= \frac{(1-p)^{n-1} ((n-2)(1-p) + 2)^2}{n^2} + \frac{p(1-p)^n}{n} + nT_2 \end{aligned}$$

Combining these inequalities together, we have the following:

$$\begin{aligned} &\mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 \right] \\ &= \mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A \right] P(A) + \mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A, C \right] P(A, C) \\ &\quad + \mathbb{E} \left[\left(\sum_{i=1}^n W_{1,i} \right)^2 j A, C \right] P(A, C) \\ &= \frac{(1-p)^{n-1} (1 + (n-1)(1-p))^2}{n^3} + \frac{p^{n-1} (1 + (n-1)(1-p))^2}{n} + \frac{2p^n (1-p) + (1-p)^n p}{n^2} \\ &\quad + nT_1 + \frac{(n-1)p}{n} \left(1 + \frac{(1-p)^{n-1}}{n^2} + \frac{(n-2)(1-p)^n}{n} + \frac{p(1-p)^n}{n} + nT_2 + 2T_3 \right) \\ &\quad + \frac{(n-1)(1-p)}{n} \left(\frac{(1-p)^{n-1} ((n-2)(1-p) + 2)^2}{n^2} + \frac{p(1-p)^n}{n} + nT_2 \right) \\ &= \frac{(1-p)^{n-1}}{n} + \frac{p^{n-1} (1 + (n-1)(1-p))^2}{n} + \frac{p(1-p)^n + p^n (1-p)}{n^2} + nT_1 \end{aligned}$$




Figure 8: ResNet20 on CIFAR10 - Varying arrival probability on 16 nodes. Impact on test accuracy.

$$\begin{aligned}
 &+ p + \frac{p(1-p)^n}{n} + nT_2 + 2T_3p + \frac{(n-1)(1-p)^n}{n} \\
 &p(1+2T_3) + (1-p)^{n-1} + \frac{2p(1-p)^n}{n} + \frac{p^n(1-p)}{n^2} + n(T_1 + T_2),
 \end{aligned}$$

and

$$\alpha_2 \frac{p(1+2T_3) + (1-p)^{n-1}}{n} + \frac{2p(1-p)^n}{n} + \frac{p^n(1-p)}{n^2} + T_1 + T_2.$$

□

E. Extended Experimental Analysis

In this section we conduct an extended analysis of the experiments on the CIFAR-10 dataset using the ResNet20 architecture. We focus on this setting with default hyper-parameters in order to get some insight about the generalization properties when lowering the packet arrival probabilities or changing the number of nodes for a fixed probability of packet arrival. Not surprisingly, the effect of increasing the number of nodes for a fixed arrival rate (Figure 9) on the train loss is similar to lowering the arrival probability for a fixed number of nodes (Figure 8). When looking at the test accuracy though, the gap when lowering the arrival rate is less than 1 point in percentage, whereas the difference in test accuracy between 2 and 16 nodes is around 2 points in percentage.

To show that having too low arrival probabilities might effect convergence on the training loss or generalization, we further lower the parameter to an extreme value of 0% arrival rate (Figure 10). Although the training loss still converges for this scenario, the trained model does generalize very badly as visible in the test accuracy. Interestingly, this extreme scenario could in certain cases still offer proper convergence and generalization. Imagining having a very large training dataset. Training a model on n nodes with $p = 0.0\%$, and reporting the final model at a specific node is identical to training a model on $\frac{1}{n}$ -th of the data on a single node.

Lastly, we want to re-emphasize that we make use of default values given by the underlying deep learning framework for all the hyper-parameters (learning rate, decay strategy and minibatch size per nodes) regardless of the number of nodes or arrival probability. We believe that carefully tuning the hyper-parameters based on the given number of nodes and arrival probability might help further lowering the training loss while maintaining generalization with regard to the test accuracy.




Figure 9: ResNet20 on CIFAR10 - Varying number of nodes for 80% arrival probability.

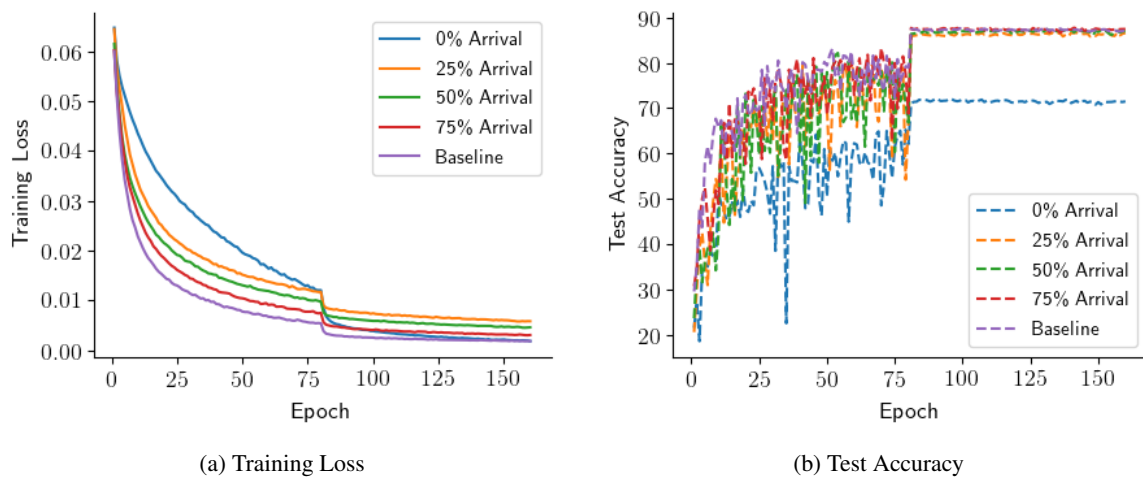


Figure 10: ResNet20 on CIFAR10 - Extreme arrival probabilities on 16 nodes.