
SCAFFLSA: Quantifying and Eliminating Heterogeneity Bias in Federated Linear Stochastic Approximation and Temporal Difference Learning

Paul Mangold^{*1} Sergey Samsonov^{*2} Safwan Labbi¹ Ilya Levin² Reda Alami³ Alexey Naumov²
Eric Moulines^{1,4}

Abstract

In this paper, we perform a non-asymptotic analysis of the federated linear stochastic approximation (FedLSA) algorithm. We explicitly quantify the bias introduced by local training with heterogeneous agents, and investigate the sample complexity of the algorithm. We show that the communication complexity of FedLSA scales polynomially with the desired precision ϵ , which limits the benefits of federation. To overcome this, we propose SCAFFLSA, a novel variant of FedLSA, that uses control variates to correct the bias of local training, and prove its convergence without assumptions on statistical heterogeneity. We apply the proposed methodology to federated temporal difference learning with linear function approximation, and analyze the corresponding complexity improvements.

1. Introduction

Federated Learning (FL) (McMahan et al., 2017; Konečný et al., 2016) represents a paradigm shift in machine learning, leveraging decentralized datasets across different devices or agents, while ensuring that the data remains at its source. This federated approach to model training involves the iterative refinement of a global model through a sequence of interactions with local oracles. These oracles, working with individual agent datasets, provide updates to a centralized server. The server then integrates these updates to incrementally improve and evolve the global model.

The predominant research in federated learning (FL) focuses on stochastic gradient methods tailored to distributed

loss minimization tasks. These tasks often benefit from a particular structure (e.g. finite-sum, overparameterization), that allow the use of *ad-hoc* techniques to handle gradients' stochasticity. In this study, the focus is shifted to the stochastic approximation (SA) paradigm (Robbins & Monro, 1951). Unlike stochastic gradient scenarios, the stochastic oracles in SA do not originate from the gradient of a loss function. Consequently, analyzing SA techniques requires a novel set of analytical tools. These tools, despite their importance in federated settings, have yet to be explored.

In this manuscript, we dive into the field of federated linear stochastic approximation (federated LSA) with a particular focus on the *heterogeneity* of participating agents. Heterogeneity poses a major challenge: if not carefully managed, it can severely affect the performance of federated algorithms, limiting the practical utility of such methods in real-world applications. The main goal of federated LSA is to solve a linear system of equations, where (i) the system matrix and the corresponding target are only accessible via stochastic oracles, and (ii) these oracles are distributed over an ensemble of heterogeneous agents. This problems can be solved using the FedLSA procedure, which performs LSA locally, with occasional consensus steps. An important application of federated LSA is federated *policy evaluation* in reinforcement learning. This is usually carried out using temporal difference learning (TD); see Sutton (1988). In the TD paradigm, the stochastic oracles involved in the process are characterized by high variance, a property that highlights the potential benefits of promoting cooperation between agents to speed up evaluation. Consequently, it is the stochastic nature of these oracles that often proves to be the most critical constraint in such environments.

In particular, there is a notable strand of research, initiated by Karimireddy et al. (2020), that deals with heterogeneity using control variates. However, the effectiveness of these approaches remains unclear when the local oracle has high variance; see Mishchenko et al. (2022). Since such noisy environments are particularly relevant for TD learning, this raises a critical question: *Should the usage of control variates provably improve the communication complexity of the Federated SA procedure?*

^{*}Equal contribution ¹École Polytechnique, Palaiseau, France ²HSE University, Russia ³Technology Innovation Institute, 9639 Masdar City, Abu Dhabi, United Arab Emirates ⁴Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, United Arab Emirates. Correspondence to: Paul Mangold <paul.mangold@polytechnique.edu>, Sergey Samsonov <svsamsonov@hse.ru>.

In this paper, we attempt to provide a positive answer to this question in the specific setting of federated LSA with heterogeneous agents. Our primary contribution is threefold:

- Inspired by the study of Wang et al. (2022), we investigate the non-asymptotic behavior of the FedLSA algorithm. Our analysis delineates the exact relationship between the mean squared error (MSE) of the FedLSA procedure and three key factors: the number of local updates, the magnitude of the local step size, and the number of agents. Furthermore, we derive an analytical formulation that captures the intrinsic heterogeneity bias.
- To correct the heterogeneity bias, we propose and analyze SCAFFLSA, a variant of FedLSA that leverages control variates, with two different communication strategies: first, the randomized communication scheme, which is similar to the framework used in proximal federated methods (see e.g. Mishchenko et al. (2022)), and second, the deterministic-temporal communication pattern, where communication is systematically activated at the end of blocks of H local step. This ensures compatibility with different operational paradigms typical of federated learning environments.
- We show how the proposed methodology can be applied to federated policy evaluation problems in reinforcement learning in heterogeneous settings. We focus on the situation of federated TD learning with linear functional approximation and establish tight finite-sample upper bounds of the mean square error, improving over existing analyses.

This paper is organized as follows. We first discuss related work in Section 2. We then introduce federated LSA and TD in Section 3, and analyze it in Section 4. Section 5 then introduces a novel strategy for mitigating the bias. Finally, we illustrate our results numerically in Section 6.

Notations and definitions. In the rest of the paper, we will use the following notations. For matrix A we denote by $\|A\|$ its operator norm. Setting N for the number of communicating clients, we use the notation $\mathbb{E}_c[a_c] = \frac{1}{N} \sum_{c=1}^N a_c$ for the average over different clients. For the matrix $A = A^\top \succeq 0$, $A \in \mathbb{R}^{d \times d}$ and vector $x \in \mathbb{R}^d$ we define the corresponding norm $\|x\|_A = \sqrt{x^\top A x}$. For the sequences a_n and b_n we write $a_n \lesssim b_n$ if there exist a constant $c > 0$ such that $a_n \leq cb_n$ for $c > 0$.

2. Related Work

Federated Learning. Apart from the seminal work from Doan (2020) on federated stochastic approximation, most of the FL literature is devoted to federated gradient methods.

Overall, a strong focus has been put on the celebrated Federated Averaging (*FedAvg*) algorithm (McMahan et al., 2017). *FedAvg* aims to reduce communication through local training, which causes *local drift* when agents are heterogeneous (Zhao et al., 2018). Sample and communication complexity of *FedAvg* have been studied under different assumptions, with homogeneous agents (Li et al., 2020; Haddadpour & Mahdavi, 2019) and heterogeneous agents (Khaled et al., 2020; Koloskova et al., 2020). It was also shown to yield linear speed-up in the number of agents when gradient are stochastic (Qu et al., 2021). In practice, *FedAvg* exhibits much better performance than predicted by most theoretical analyses when agents are heterogeneous (Collins et al., 2021; Reddi et al., 2020). This phenomenon has been studied theoretically by Wang et al. (2022); Patel et al. (2023), who proposed new measures of heterogeneity, that are more adapted to *FedAvg*. Our analysis of federated linear stochastic approximation extends these ideas to the linear stochastic approximation setting.

Another line of work aims to correct *FedAvg*'s bias by introducing control variates, that compensate for heterogeneity. This was proposed by Karimireddy et al. (2020), with the *Scaffold* method, and was then studied by Gorbunov et al. (2021); Mitra et al. (2021), recovering the convergence rate of gradient descent (with a single local step), regardless of heterogeneity. It was later shown with *ProxSkip* (Mishchenko et al., 2022) that such methods indeed *accelerate* the training. However, and contrarily to *Scaffold*, they lose the linear speed-up in the number of agents when gradients are stochastic. Multiple other methods with accelerated rates have been proposed (Malinovsky et al., 2022; Condat et al., 2022; Condat & Richtárik, 2022; Grudziński et al., 2023; Hu & Huang, 2023), although all lose the linear speed-up. Inspired from these works, we propose a new method for federated linear stochastic approximation, with improved rates at the cost of the linear speed-up.

Single-agent and Federated TD-learning. Temporal difference (TD) learning have a long history in policy evaluation (Sutton, 1988; Dann et al., 2014), with classical asymptotic analyses in the linear function approximation setting (Tsitsiklis & Van Roy, 1997; Sutton et al., 2009). Multiple non-asymptotic analyses of the mean squared error (MSE) of TD-type algorithms have been proposed (Bhandari et al., 2018; Dalal et al., 2018; Patil et al., 2023; Li et al., 2023b; Samsonov et al., 2023). Recently, significant attention has been drawn to federated reinforcement learning (Lim et al., 2020; Qi et al., 2021; Xie & Song, 2023) and to Federated TD learning with linear function approximation. Khodadadian et al. (2022); Dal Fabbro et al. (2023); Liu & Olshevsky (2023) analysed it with very little communication, under the strong assumption that agents are identical. Federated TD was studied with heterogeneous agents, first without local training (Doan, 2020), then with local training, but no linear

speed-up (Doan et al., 2019; Jin et al., 2022). More recently, Wang et al. (2023) analyzed federated TD with heterogeneous agents, local training, and linear speed-up in number of agents in a low-heterogeneity setting. Unfortunately, their analysis suffers from unavoidable bias. Additionally, its dependence on noise levels can be pessimistic, and requires the server to project aggregated iterates to a ball of unknown radius. In contrast, we give sharp bounds on the convergence of federated TD, with linear speed-up, and explicit characterization of the bias. We also propose a strategy to mitigate this bias, allowing for extended local training.

3. Federated LSA and TD learning

In this paper, we study the federated linear stochastic approximation problem, where N agents aim to collaboratively solve a system with the following finite sum structure

$$\left(\frac{1}{N} \sum_{c=1}^N \bar{\mathbf{A}}^c\right) \theta_\star = \frac{1}{N} \sum_{c=1}^N \bar{\mathbf{b}}^c, \quad (1)$$

where for $c \in [N]$, $\bar{\mathbf{A}}^c \in \mathbb{R}^{d \times d}$, $\bar{\mathbf{b}}^c \in \mathbb{R}^d$, and we assume the solution θ_\star to be unique. We set the system matrix and corresponding right-hand side, respectively, as

$$\bar{\mathbf{A}} = \frac{1}{N} \sum_{c=1}^N \bar{\mathbf{A}}^c, \quad \bar{\mathbf{b}} = \frac{1}{N} \sum_{c=1}^N \bar{\mathbf{b}}^c.$$

In linear stochastic approximation, neither matrices $\bar{\mathbf{A}}^c$ nor vectors $\bar{\mathbf{b}}^c$ are observed directly. Instead, each agent $c \in [N]$ has access to its own observation sequence $\{(\mathbf{A}^c(Z_k^c), \mathbf{b}^c(Z_k^c))\}_{k \in \mathbb{N}}$. For each $c \in [N]$, $\mathbf{A}^c : \mathcal{Z} \rightarrow \mathbb{R}^{d \times d}$ and $\mathbf{b}^c : \mathcal{Z} \rightarrow \mathbb{R}^d$ are measurable functions, and $(Z_k^c)_{n \in \mathbb{N}}$ is an i.i.d. sequence, with values in a state space $(\mathcal{Z}, \mathcal{Z})$ and distribution π_c satisfying $\mathbb{E}[\mathbf{A}^c(Z_1^c)] = \bar{\mathbf{A}}^c$ and $\mathbb{E}[\mathbf{b}^c(Z_1^c)] = \bar{\mathbf{b}}^c$. In the following, we assume that agents' observation sequences are independent from each other, and that each local system $\bar{\mathbf{A}}^c \theta_\star^c = \bar{\mathbf{b}}^c$ has a unique solution θ_\star^c .

In the federated setting, agents can only communicate through a central server, and such communications are generally costly. This makes problem (1) challenging, as it must be solved collaboratively by all agents. To this end, we propose FedLSA, where agents perform local updates, that are aggregated periodically to limit the communication burden. During round $t > 0$, agents start from a shared value θ_t and perform $H > 0$ local updates, for $h = 1$ to H ,

$$\theta_{t,h}^c = \theta_{t,h-1}^c - \eta(\mathbf{A}^c(Z_{t,h}^c) - \mathbf{b}^c(Z_{t,h}^c)),$$

where we set $\theta_{t,0}^c = \theta_t$, and use the alias $Z_{t,h}^c \equiv Z_{Ht+h-1}^c$ to simplify notations. Agents then send $\theta_{t,H}^c$ to the central server, that aggregates them as $\theta_t = \frac{1}{N} \sum_{c=1}^N \theta_{t-1,H}^c$ and sends it back to all agents. We describe the complete FedLSA procedure in Algorithm 1.

Stochastic expansions for local updates. We use the error expansion framework (see e.g. Aguech et al., 2000;

Algorithm 1 FedLSA

Input: $\eta > 0, \theta_0 \in \mathbb{R}^d, T, N, H > 0$
for $t = 0$ to $T - 1$ **do**
 Initialize $\theta_{t,0} = \theta_t$
 for $c = 1$ to N **do**
 for $h = 1$ to H **do**
 Receive $Z_{t,h}^c$ and perform local update:
 $\theta_{t,h}^c = \theta_{t,h-1}^c - \eta(\mathbf{A}^c(Z_{t,h}^c) - \mathbf{b}^c(Z_{t,h}^c))$ (2)
 end for
 end for
 Average: $\theta_{t+1} = \frac{1}{N} \sum_{c=1}^N \theta_{t,H}^c$ (3)
end for

Durmus et al., 2022) for LSA to analyze the mean squared error of the estimates of Algorithm 1. For this purpose, we rewrite local updates (2) as

$$\theta_{t,h}^c - \theta_\star^c = (\mathbf{I} - \eta \mathbf{A}^c(Z_{t,h}^c))(\theta_{t,h-1}^c - \theta_\star^c) - \eta \varepsilon^c(Z_{t,h}^c), \quad (4)$$

where we have set $\tilde{\mathbf{b}}^c(z) = \mathbf{b}^c(z) - \bar{\mathbf{b}}^c$, $\tilde{\mathbf{A}}^c(z) = \mathbf{A}^c(z) - \bar{\mathbf{A}}^c$, and defined the noise variable at each θ_\star^c ,

$$\varepsilon^c(z) = \tilde{\mathbf{A}}^c(z) \theta_\star^c - \tilde{\mathbf{b}}^c(z).$$

Running the recursion (4) until the start of local training, we obtain that

$$\theta_{t,H}^c - \theta_\star^c = \Gamma_{t,1:H}^{(c,\eta)} \{\theta_{t,0}^c - \theta_\star^c\} - \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \varepsilon^c(Z_{t,h}^c),$$

where we introduced the notation

$$\Gamma_{t,m:n}^{(c,\eta)} = \prod_{h=m}^n (\mathbf{I} - \eta \mathbf{A}^c(Z_{t,h}^c)), \quad 1 \leq m \leq n \leq H,$$

with the convention $\Gamma_{t,m:n}^{(c,\eta)} = \mathbf{I}$ for $m > n$. Now we proceed with the classical assumptions, under which we analyze the LSA error dynamics:

A1. For each agent c the observations $(Z_k^c)_{k \in \mathbb{N}}$ are i.i.d. random variables taking values in $(\mathcal{Z}, \mathcal{Z})$ with a distribution π_c satisfying $\mathbb{E}_{\pi_c}[\mathbf{A}^c(Z_1^c)] = \bar{\mathbf{A}}^c$ and $\mathbb{E}_{\pi_c}[\mathbf{b}^c(Z_1^c)] = \bar{\mathbf{b}}^c$. Moreover, for each $c \in [N]$ the following matrices exist

$$\Sigma_\varepsilon^c = \int_{\mathcal{Z}} \varepsilon^c(z) \varepsilon^c(z)^\top d\pi_c(z), \quad (5)$$

$$\Sigma_{\tilde{\mathbf{A}}}^c = \int_{\mathcal{Z}} (\tilde{\mathbf{A}}^c(z))^\top \tilde{\mathbf{A}}^c(z) d\pi_c(z), \quad (6)$$

and $C_{\mathbf{A}} = \sup_{c \in [N]} \|\bar{\mathbf{A}}^c\| < \infty$.

A2. There exist $a > 0, \eta_\infty > 0$, such that $\eta_\infty a \leq 1/2$, and for $\eta \in (0; \eta_\infty)$, $c \in [N]$, $u \in \mathbb{R}^d$, $h \in \mathbb{N}$,

$$\mathbb{E}^{1/2} [\|\Gamma_{t,1:h}^{(c,\eta)} u\|^2] \leq (1 - \eta a)^h \|u\|. \quad (7)$$

Assumption A1 is classical in finite-time studies of LSA iterates; see e.g. Srikant & Ying (2019); Durmus et al. (2022).

In the considered application to the Federated TD learning, the random matrices $\tilde{\mathbf{A}}^c(z)$ and vectors $\varepsilon^c(z)$ are almost sure bounded, so they automatically admit finite moments. Note that the matrix Σ_ε^c from (5) measures the noise level at the *local* optimum θ_\star^c . Verifying the exponential stability assumption A2 is crucial for the convergence of LSA algorithm since it automatically implies that the transient component of the error $\tilde{\theta}_t^{(\text{tr})}$ decreases exponentially fast; see Guo & Ljung (1995); Priouret & Veretenikov (1998). In practical applications, this assumption requires careful verification, with particular attention to the way the parameter a in (7) scales with instance-dependent quantities. Importantly, A2 implies the stability of the deterministic matrix product, that is, for any $u \in \mathbb{R}^d$ we have

$$\|(I - \eta \tilde{\mathbf{A}}^c)^h u\| \leq (1 - \eta a)^h \|u\| .$$

In Section 5, we will sometimes require a finer assumption, which we state here for ease of reference.

A3. There exist constants $a, L > 0$, such that for any $\eta \in (0, 1/L)$, $c \in [N]$, it holds for $Z \sim \pi_c$, that

$$aI \preceq \mathbb{E}[\frac{1}{2}(\mathbf{A}^c(Z) + \mathbf{A}^c(Z)^\top)] \preceq \frac{1}{L}\mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)] .$$

This assumption is slightly more restrictive than A2, since whenever A3 holds, A2 also holds with the same constant a (see Patil et al., 2023; Samsonov et al., 2023).

Federated temporal-difference learning. We now consider the particular example of federated TD-learning with linear function approximation. In this setting we observe N Markov Decision Processes $(\mathcal{S}, \mathcal{A}, \mathbb{P}_{\text{MDP}}^c, r^c, \gamma)$ with common state and action spaces. Here, γ is a discounting factor, and each of the agents has its own environment dynamics $\mathbb{P}_{\text{MDP}}^c$. The state space \mathcal{S} is assumed to be a complete metric space. The dynamics of each MDP is encapsulated by the Markov kernel $\mathbb{P}_{\text{MDP}}^c(B|s, a)$, which specifies the probability of transitioning from state s to a set $B \in \mathcal{B}(\mathcal{S})$ upon taking action a . We denote by $r^s : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ the reward function (assumed to be deterministic for simplicity) and by $\pi(\cdot|s)$ an agent’s policy (assumed to be the same for all agents). On the contrary, the dynamics $\mathbb{P}_{\text{MDP}}^c$ and reward function r^s are *specific to each agent*. The objective is to compute the agent’s value function under the policy π

$$V^{c,\pi}(s) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r^c(S_t^c, A_t^c)] ,$$

where $S_0^c = s$, and for $t \in \mathbb{N}$, $A_t^c \sim \pi(\cdot|S_t^c)$ and $S_{t+1}^c \sim \mathbb{P}_{\text{MDP}}^c(\cdot|S_t^c, A_t^c)$. We consider linear functional approximation for $V^{c,\pi}(s)$ with a feature mapping $\varphi : \mathcal{S} \rightarrow \mathbb{R}^d$, i.e. we aim to approximate the value function by $V_{\theta}^{c,\pi}(s) = \varphi^\top(s)\theta$, $\theta \in \mathbb{R}^d$. For $c \in [N]$, we set μ^c the invariant distribution over \mathcal{S} induced by $(\pi, \mathbb{P}_{\text{MDP}}^c)$, and define the design matrix Σ_φ^c as

$$\Sigma_\varphi^c = \mathbb{E}_{\mu^c}[\varphi(S_0^c)\varphi(S_0^c)^\top] \in \mathbb{R}^{d \times d} .$$

Properties of Σ_φ^c are extremely important for our analysis, since we define the optimal approximation parameter θ_\star^c as

$$\theta_\star^c \in \arg \min_{\theta \in \mathbb{R}^d} \mathbb{E}_{s \sim \mu^c} [\|V_{\theta}^{c,\pi}(s) - V^{c,\pi}(s)\|^2] . \quad (8)$$

With the linear functional approximation (LFA) we can write the federated TD(0) as a particular setting of FedLSA algorithm. Here we follow other papers of TD(0) with LFA (Patil et al., 2023; Wang et al., 2023), and leave detailed description to Appendix D. We put the assumptions on Σ_φ^c and generative mechanism, which are classical in TD (Patil et al., 2023; Li et al., 2023a; Samsonov et al., 2023).

TD 1. Tuples $(S_t^c, A_t^c, S_{t+1}^c)$ are generated i.i.d. with $S_t^c \sim \mu^c$, $A_t^c \sim \pi(\cdot|s)$, $S_{t+1}^c \sim \mathbb{P}_{\text{MDP}}^c(\cdot|S_t^c, A_t^c)$.

TD 2. Matrices Σ_φ^c are non-degenerate with the minimal eigenvalue $\nu = \min_{c \in [N]} \lambda_{\min}(\Sigma_\varphi^c) > 0$. Moreover, the feature mapping $\varphi(\cdot)$ satisfies $\sup_{s \in \mathcal{S}} \|\varphi(s)\| \leq 1$.

The generative model assumption **TD 1** is used in many previous works; see, e.g. Dalal et al. (2018); Li et al. (2023a); Patil et al. (2023). We note that it is possible to generalize this assumption to the more realistic setting of on-policy evaluation over a single trajectory leveraging the Markovian noise dynamics, following Wang et al. (2023). However, we leave it as a direction for the future work. Assumption **TD 2** allows to ensure the uniqueness of the optimal parameter θ_\star^c in (8). Under **TD 1** and **TD 2** we check the general LSA assumptions A1-A3, and the following statement holds.

Lemma 3.1. *Assume **TD 1-TD 2**. Then the sequence of TD(0) updates satisfies assumptions A1-A3 with*

$$C_{\mathbf{A}} = 1 + \gamma , \quad \|\Sigma_{\tilde{\mathbf{A}}}^c\| \leq 2(1 + \gamma)^2 , \quad (9)$$

$$\text{Tr}(\Sigma_\varepsilon^c) \leq 2(1 + \gamma)^2 (\|\theta_\star^c\|^2 + 1) , \quad (10)$$

$$a = \frac{(1-\gamma)\nu}{2} , \quad \eta_\infty = \frac{(1-\gamma)}{4} , \quad L = \frac{1+\gamma}{(1-\gamma)^2\nu} . \quad (11)$$

Proof of the bounds (9)–(11) can be found in Samsonov et al. (2023) or Patil et al. (2023). For reader’s convenience we provide a detailed argument in Appendix D. In the considered setting the local optimum parameter θ_\star^c corresponds to the (unique) solution of the system $\tilde{\mathbf{A}}^c \theta_\star^c = \tilde{\mathbf{b}}^c$, where we set, respectively,

$$\tilde{\mathbf{A}}^c = \mathbb{E}_{s \sim \mu^c, s' \sim P^{\pi, c}(\cdot|s)} [\phi(s) \{\phi(s) - \gamma \phi(s')\}^\top] \quad (12)$$

$$\tilde{\mathbf{b}}^c = \mathbb{E}_{s \sim \mu^c, a \sim \pi(\cdot|s)} [\phi(s) r(s, a)] .$$

In case of a single-agent TD(0) the quantity of interest, which measures the quality of parameter θ compared to θ_\star^c is arguably (see Patil et al., 2023; Wang et al., 2023) not the norm $\|\theta - \theta_\star^c\|^2$, but rather

$$\|\theta - \theta_\star^c\|_{\Sigma_\varphi^c}^2 = \mathbb{E}_{s \sim \mu^c} [\|V_{\theta}^{c,\pi}(s) - V_{\theta_\star^c}^{c,\pi}(s)\|^2] .$$

In our setting, the optimal parameter θ_\star corresponds to the averaged system $\frac{1}{N} \sum_{c=1}^N \tilde{\mathbf{A}}^c \theta_\star = \frac{1}{N} \sum_{c=1}^N \tilde{\mathbf{b}}^c$. This system does not naturally relate to the iterates of TD(0) algorithm with dynamics from some transition kernel $P^{\pi, c}(\cdot|s)$,

$c \in [N]$. For this reason, we will present our results using the euclidean norm $\|\cdot\|$.

4. Analysis of the FedLSA algorithm

In this section, we analyze the FedLSA algorithm. First, we express θ_T as a function of θ_0 . Then, we study the convergence rate of the algorithm. Finally, we discuss sample and communication complexity of the algorithm.

Stochastic expansion for FedLSA. To derive an expression of θ_T as a function of θ_0 , we start from the expansions from Section 3. Using the fact $\theta_{t,0}^c = \theta_{t-1}$, and employing the global averaging procedure (3), we obtain that

$$\theta_t - \theta_\star = \bar{\Gamma}_{t,H}^{(\eta)} \{\theta_{t-1} - \theta_\star\} + \bar{\rho}_H + \bar{\tau}_{t,H} + \eta \bar{\varphi}_{t,H}, \quad (13)$$

where we have defined

$$\begin{aligned} \bar{\Gamma}_{t,H}^{(\eta)} &= \frac{1}{N} \sum_{c=1}^N \Gamma_{t,1:H}^{(c,\eta)}, \quad (14) \\ \bar{\rho}_H &= \frac{1}{N} \sum_{c=1}^N (\mathbf{I} - (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^H) \{\theta_\star^c - \theta_\star\}, \\ \bar{\tau}_{t,H} &= \frac{1}{N} \sum_{c=1}^N \{(\mathbf{I} - \eta \bar{\mathbf{A}}^c)^H - \Gamma_{t,1:H}^{(c,\eta)}\} \{\theta_\star^c - \theta_\star\}, \\ \bar{\varphi}_{t,H} &= -\frac{1}{N} \sum_{c=1}^N \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \Gamma_{t,h}^c \varepsilon^c(Z_{t,h}^c). \end{aligned}$$

The *transient* term $\bar{\Gamma}_{t,H}^{(\eta)} (\theta_{t-1} - \theta_\star)$, responsible for the rate of forgetting the previous iteration error $\theta_{t-1} - \theta_\star$, and the *fluctuation* term $\eta \bar{\varphi}_{t,H}$, reflecting the oscillations of the iterates around θ_\star , are similar to the ones from the standard LSA error decomposition. The two additional terms in (13) reflect the *heterogeneity bias*. This bias is composed of two parts: the true bias $\bar{\rho}_H$, which is non-random, and its fluctuations $\bar{\tau}_{t,H}$. Note that, if $H = 1$, then the bias term vanishes, as

$$\bar{\rho}_1 = \frac{1}{N} \sum_{c=1}^N \bar{\mathbf{A}}^c \{\theta_\star^c - \theta_\star\} = 0, \quad (15)$$

which shows that, when doing a single local step, the bias due to local training disappears. When $H \geq 2$, we show in Lemma C.3 that the bias term can be bounded by

$$\|\bar{\rho}_H\| \leq \frac{\eta^2 H^2}{N} \sum_{c=1}^N \exp(\eta H \|\bar{\mathbf{A}}^c\|) \|\theta_\star^c - \theta_\star\|. \quad (16)$$

This bound cannot be improved without further assumptions. To analyze the complexity and communication complexity of FedLSA, we run the recurrence (13) to obtain

$$\theta_t - \theta_\star = \tilde{\theta}_t^{(\text{tr})} + \tilde{\theta}_t^{(\text{bi,bi})} + \tilde{\theta}_t^{(\text{fl,bi})} + \tilde{\theta}_t^{(\text{fl})}, \quad (17)$$

where we have defined

$$\begin{aligned} \tilde{\theta}_t^{(\text{tr})} &= \prod_{s=1}^t \bar{\Gamma}_{s,H}^{(\eta)} \{\theta_0 - \theta_\star\}, \\ \tilde{\theta}_t^{(\text{bi,bi})} &= \sum_{s=1}^t (\bar{\Gamma}_H^{(\eta)})^{t-s} \bar{\rho}_H, \\ \tilde{\theta}_t^{(\text{fl,bi})} &= \sum_{s=1}^t \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{\tau}_{s,H} + \Delta_{H,s,t}^{(\eta)} \bar{\rho}_H, \\ \tilde{\theta}_t^{(\text{fl})} &= \eta \sum_{s=1}^t \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{\varphi}_{s,H}, \end{aligned}$$

with the notations $\bar{\Gamma}_H^{(\eta)} = \mathbb{E}[\bar{\Gamma}_{s,H}^{(\eta)}] = \frac{1}{N} \sum_{c=1}^N (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^H$ and $\Delta_{H,s,t}^{(\eta)} = \{\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)}\} - (\bar{\Gamma}_H^{(\eta)})^{t-s}$. The first term, $\tilde{\theta}_t^{(\text{tr})}$ gives the rate at which the initial error is forgotten. The terms $\tilde{\theta}_t^{(\text{bi,bi})}$ and $\tilde{\theta}_t^{(\text{fl,bi})}$ represent the bias and fluctuation due to statistical heterogeneity across agents. Note that in the special case where agents are homogeneous (i.e. $\bar{\mathbf{A}}^c = \bar{\mathbf{A}}$ for all $c \in [N]$), these two terms vanish. Finally, the term $\tilde{\theta}_t^{(\text{fl})}$ depicts the fluctuations of θ_t around the solution θ_\star .

Convergence rate of FedLSA. First, we analyze the rate at which FedLSA converges to $\tilde{\theta}_t^{(\text{bi,bi})} + \theta_\star$. The two following quantities, that are due to heterogeneity and stochasticity of local estimators, will play a central role in this rate,

$$\tilde{v}_{\text{heter}} = \mathbb{E}_c[\|\sum_{\mathbf{A}}^c \|\theta_\star^c - \theta_\star\|^2], \quad \bar{\sigma}_\varepsilon = \mathbb{E}_c[\text{Tr}(\Sigma_\varepsilon^c)]. \quad (18)$$

The quantities \tilde{v}_{heter} and $\bar{\sigma}_\varepsilon$ corresponds to the different sources of noise in the error decomposition (17). Indeed, $\bar{\sigma}_\varepsilon$ is related to the variance of *local* LSA iterates on each of the agents, while \tilde{v}_{heter} controls bias fluctuation term $\tilde{\theta}_t^{(\text{fl,bi})}$. Note also that in the centralized setting (i.e. $N = 1$) the term \tilde{v}_{heter} vanishes, but not the term $\bar{\sigma}_\varepsilon$. We now proceed with the analysis of the MSE of FedLSA's iterates.

Theorem 4.1. *Assume A1 and A2. Then for any step size $\eta \in (0, \eta_\infty)$ it holds that*

$$\begin{aligned} \mathbb{E}^{\frac{1}{2}} [\|\theta_t - \tilde{\theta}_t^{(\text{bi,bi})} - \theta_\star\|^2] &\lesssim \sqrt{\frac{\eta \tilde{v}_{\text{heter}}}{aN}} + \sqrt{\frac{\eta \bar{\sigma}_\varepsilon}{aN}} \\ &+ \sqrt{\frac{\mathbb{E}_c[\|\sum_{\mathbf{A}}^c \|\bar{\rho}_H\|]}{HN}} \frac{\|\bar{\rho}_H\|}{a} + (1 - \eta a)^{tH} \|\theta_0 - \theta_\star\|, \end{aligned}$$

where the bias $\tilde{\theta}_t^{(\text{bi,bi})}$ converges to $(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1} \bar{\rho}_H$ at a rate $\|\tilde{\theta}_t^{(\text{bi,bi})} - (\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1} \bar{\rho}_H\| \leq (1 - \eta a)^{tH} \|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\| \|\bar{\rho}_H\|$.

The proof of this result relies on bounding carefully each of the terms from (17). We give the detailed proof with explicit constants in Appendix A. Importantly, the fluctuation and heterogeneity terms scale in linearly with N , which allows to use larger step-size than in the single-agent setting.

Remark 4.2. When $N = 1$, FedLSA reverts to the centralized algorithm: the bias term $\bar{\rho}_H$ and its fluctuation \tilde{v}_{heter} vanish in Theorem 4.1, yielding the last-iterate bound

$$\mathbb{E}^{1/2} [\|\theta_t - \theta_\star\|^2] \lesssim \sqrt{\eta \bar{\sigma}_\varepsilon / a} + (1 - \eta a)^{tH} \|\theta_0 - \theta_\star\|,$$

which is known to be sharp in its dependence on η for single-agent LSA (see Theorem 5 in Durmus et al. (2021)).

These bounds can be instantiated for federated TD(0): using the bounds from Lemma 3.1, we obtain the following result.

Corollary 4.3. *Assume TD 1 and TD 2. Then for any step size $\eta \in (0, \frac{1-\gamma}{4})$, the iterates of federated TD(0) satisfy*

$$\begin{aligned} \mathbb{E}^{1/2} [\|\theta_t - \tilde{\theta}_t^{(\text{bi,bi})} - \theta_\star\|^2] &\lesssim \sqrt{\frac{\eta (\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|^2] \vee (1 + \mathbb{E}_c[\|\theta_\star^c\|^2]))}{(1-\gamma)\nu N}} \\ &+ \sqrt{\frac{1}{HN}} \frac{\|\bar{\rho}_H\|}{(1-\gamma)\nu} + (1 - \frac{\eta(1-\gamma)\nu}{2})^{tH} \|\theta_0 - \theta_\star\|. \end{aligned}$$

The right hand side of Corollary 4.3 scales linearly with N , allowing for linear speed-up. This is in line with recent results on federated TD(0), that show linear speed-up either without local training (Dal Fabbro et al., 2023) or up to a possibly large bias term (Wang et al., 2023) (see analysis of their Theorem 2). Next, we will see that our tighter analysis of the bias allows to prove convergence of federated TD(0) even under heterogeneity.

Sample and communication complexity of FedLSA.

We begin with the analysis without local training (that is, $H = 1$). There, the bias term disappears, and above results directly give a simplified sample complexity bound.

Corollary 4.4. *Assume A1 and A2. Let $H = 1$, $0 < \epsilon < 1$. Set $\eta = \frac{aN\epsilon^2}{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}} \wedge \eta_\infty$. Then, to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] \leq \epsilon^2$, the required number of communications is*

$$T = \mathcal{O}\left(\left(\frac{1}{a\eta_\infty} \vee \frac{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}}{Na^2\epsilon^2}\right) \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right).$$

Here, we obtain a linear speed-up in terms of sample complexity. This is expected, since when $H = 1$, FedLSA amounts to using a stochastic oracle with reduced variance, allowing for larger step-size. Now, we proceed with the more delicate setting where $H > 1$, and heterogeneity bias does not vanish. We obtain the following bounds.

Corollary 4.5. *Assume A1 and A2. Let $H > 1$, and $0 < \epsilon < \frac{(\sqrt{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]})^{2/5}}{a} \vee \frac{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{aC_A}$. Set the step size $\eta = \mathcal{O}\left(\frac{aN\epsilon^2}{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}} \wedge \eta_\infty\right)$. Then, to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] < \epsilon^2$ the required number of communications is*

$$T = \mathcal{O}\left(\left(\frac{1}{a\eta_\infty} \vee \frac{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{a^2\epsilon}\right) \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right), \quad (19)$$

where the number of local steps is set to

$$H = \mathcal{O}\left(\frac{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}}{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]} \frac{1}{N\epsilon}\right).$$

In Corollary 4.5, the total number of oracle calls scales as

$$TH = \mathcal{O}\left(\frac{\tilde{v}_{heter}\sqrt{\bar{\sigma}_\epsilon}}{Na^2\epsilon^2} \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right),$$

which corresponds to the number of iterations of the synchronous version of the LSA method predicted by Corollary 4.4. Thus, despite the bias due to heterogeneity, it is still possible to achieve linear speed-up. This is notably the case when communicating regularly enough, so that the "physical time" of the local iterations, ηH , scales as $\eta H = \mathcal{O}(\epsilon)$. Therefore, to achieve precision ϵ^2 on the MSE, the number of communications T in (19) must scale polynomially with ϵ^{-1} . In the next section, we will show that this drawback can be overcome by using appropriate control variates. This will allow to de-bias the algorithm, which will in turn allow to choose ηH of constant order with respect to ϵ .

Now we state the communication bound of federated TD(0).

Corollary 4.6. *Assume TD 1 and TD 2. Then for any $0 < \epsilon < \frac{g_1(\theta_\star^c, \theta_\star)}{(1-\gamma)^\nu}$ with $g_1 = \mathcal{O}((1 + \|\theta_\star\|)\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|])$.*

Then, to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] < \epsilon^2$ the required number of communications for federated TD(0) is

$$T = \mathcal{O}\left(\left(\frac{1}{(1-\gamma)^{2\nu}} \vee \frac{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{(1-\gamma)^{2\nu}\epsilon}\right) \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right).$$

Corollary 4.6 shows that, even when agents are heterogeneous, it is possible for federated TD(0) to converge to θ_\star with arbitrary precision. This follows from our precise characterization of the bias of federated TD(0) with local training, and its dependence on the product ηH . Nonetheless, the number of rounds T still has to scale polynomially in ϵ^{-1} . We propose a way of fixing this in the next section.

5. Bias-Corrected Federated LSA

We now introduce the *Stochastic Controlled Averaging for Federated LSA* algorithm (SCAFFLSA), an improved version of FedLSA that mitigates client drift using control variates. This method is inspired by *ProxSkip* (see Mishchenko et al., 2022), and more specifically *Scaffnew*, its instance tailored to minimize sums of strongly convex functions. In SCAFFLSA, each agent $c \in [N]$ keeps a local variable ξ_k^c , that remains constant between successive communication rounds. At timestep k , agents perform a local update on the current estimates of the parameters $(\theta_{k-1}^c)_{c=1}^N$ as follows

$$\hat{\theta}_k^c = \hat{\theta}_{k-1}^c - \eta(\mathbf{A}^c(Z_k^c)\hat{\theta}_{k-1}^c - \mathbf{b}^c(Z_k^c) - \xi_{k-1}^c).$$

Then, they communicate according to the value of a variable $B_k \in \{0, 1\}$, that is set by one of the following rules.

H1. $B_k \sim \text{Bernoulli}(p)$ are i.i.d., with $p \in [0, 1]$.

H2. $B_k = 1$ if k is a multiple of $H > 0$, $B_k = 0$ otherwise.

When $B_k = 1$, agents (i) communicate to the central server (CS), then (ii) the CS averages local iterates, and (iii) agents update their local control variates. We describe the procedure in Algorithm 2. In the following, we study SCAFFLSA under **H 1** and **H 2**. We establish finite-time bounds for the MSE of the parameters of interest $\mathbb{E}[\frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2]$. To this end, we define the *ideal* control variates at the global solution, given by $\xi_\star^c = \bar{\mathbf{A}}^c \theta_\star - \bar{\mathbf{b}}^c = \bar{\mathbf{A}}^c(\theta_\star - \theta_\star^c)$. Using ξ_\star^c , we can rewrite the local update as

$$\hat{\theta}_k^c - \theta_\star = (\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))(\theta_{k-1}^c - \theta_\star) + \eta(\xi_{k-1}^c - \xi_\star^c) - \eta \omega^c(Z_k^c), \quad (20)$$

where we defined $\omega^c(z) = \tilde{\mathbf{A}}^c(z)\theta_\star - \tilde{\mathbf{b}}^c(z)$. Under **A 1**, it has finite covariance $\Sigma_\omega^c = \int_{\mathcal{Z}} \omega^c(z)\omega^c(z)^\top d\pi_c(z)$. Similarly to (18), we use the notation $\bar{\sigma}_\omega = \mathbb{E}_c[\text{Tr}(\Sigma_\omega^c)]$.

Analysis under H 1. In the stochastic communication paradigm described in **H 1**, the agents average their respective local parameters with a predefined probability p during each iteration. It is noteworthy that this framework,

Algorithm 2 SCAFFLSA: Stochastic Controlled FedLSA

Input: $\eta > 0, \theta_0, \xi_0^c \in \mathbb{R}^d, T, N, H, p > 0$
 Set: $K = T/p$ (if **H 1**), $K = TH$ (if **H 2**)
for $k = 1$ to K **do**
 for $c = 1$ to N **do**
 Receive Z_k^c and perform local update:
 $\hat{\theta}_k^c = \hat{\theta}_{k-1}^c - \eta(\mathbf{A}^c(Z_k^c)\hat{\theta}_{k-1}^c - \mathbf{b}^c(Z_k^c) - \xi_{k-1}^c)$
 end for
 Choose B_k according to **H 1** or **H 2**
if $B_k = 1$ **then**
 Average local iterates: $\theta_k^c = \frac{1}{N} \sum_{c=1}^N \hat{\theta}_k^c$
 Update: $\begin{cases} \xi_k^c = \xi_{k-1}^c + \frac{p}{\eta}(\theta_k^c - \hat{\theta}_k^c) & \text{if } \mathbf{H 1} \\ \xi_k^c = \xi_{k-1}^c + \frac{1}{\eta H}(\theta_k^c - \hat{\theta}_k^c) & \text{if } \mathbf{H 2} \end{cases}$
else
 Set: $\theta_k^c = \hat{\theta}_k^c, \xi_k^c = \xi_{k-1}^c$
end if
end for

as specified in Algorithm 2, does not include an inner loop. This structural simplification, originally introduced in Mishchenko et al. (2022), significantly streamlines the analysis: it can simply be performed by formulating a descent lemma that revolves around a particular Lyapunov function,

$$\psi_k = \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2.$$

Examining the expected one-step improvement of this function effectively captures the dynamics of both the local parameter updates and the control variates. This provides a clear way to study the convergence properties and general behavior of the algorithm. However, this one-step analysis requires the stronger assumption A3 on $\bar{\mathbf{A}}^c$. We now establish finite time-bounds, from which we deduce sample and communication complexity of Algorithm 2 under **H 1**.

Theorem 5.1 (MSE bound). *Assume A1, A3, and H1. Then, for any step size $\eta \leq 1/2L$ and number of global iteration $K = T/p > 0$, it holds, with $\zeta = \min(\eta a, p^2)$,*

$$\mathbb{E}[\frac{1}{N} \sum_{c=1}^N \|\theta_K^c - \theta_\star\|^2] \leq (1 - \zeta)^K \psi_0 + \frac{2\eta^2}{\zeta} \bar{\sigma}_\omega,$$

where $\psi_0 = \|\theta_0 - \theta_\star\|^2 + \frac{\eta^2}{p^2} \mathbb{E}_c[\|\bar{\mathbf{A}}^c(\theta_\star^c - \theta_\star)\|^2]$.

Corollary 5.2 (Iteration complexity). *Assume A1, A3, and H1. Let $\epsilon > 0$. Set $\eta = \mathcal{O}(\min(\frac{1}{L}, \frac{\epsilon^2 a}{\bar{\sigma}_\omega}))$ and $p = \sqrt{\eta a}$ (so that $\zeta = \eta a$). Then, $\mathbb{E}[\frac{1}{N} \sum_{c=1}^N \|\theta_K^c - \theta_\star\|^2] \leq \epsilon^2$ for an expected number of communication rounds*

$$T = \mathcal{O}\left(\max\left(\sqrt{\frac{L}{a}}, \frac{\sqrt{\bar{\sigma}_\omega}}{\epsilon a}\right) \log\left(\frac{\psi_0}{\epsilon^2}\right)\right), \quad (21)$$

where ψ_0 is defined in Theorem 5.1, with an expected number of local updates

$$1/p = \mathcal{O}\left(\max\left(\sqrt{\frac{L}{a}}, \frac{\sqrt{\bar{\sigma}_\omega}}{\epsilon a}\right)\right), \text{ for } t \in [T].$$

We provide the detailed proof of these two statements in Appendix B.1. In Corollary 5.2 it is shown that the inclusion of control variates in local updates reduces the required communication rounds by a significant amount, quantifiable as $\mathcal{O}(\min(\sqrt{L/a}, 1/a\epsilon))$. Compared to the scenarios described in Section 4, this improvement means a gain by a factor of $1/a$. At the same time, however, this also means a loss of the linear acceleration factor in the form of $1/N$. This phenomenon is primarily due to the inherent properties of the analytical method used, as we show in Appendix B.1. Note that, under **TD 2** and **TD 1** the expected number of communication rounds for SCAFFLSA applied to federated TD(0) algorithm scales, following (21), as

$$T = \mathcal{O}\left(\max\left(\frac{1}{(1-\gamma)^{3/2\nu}}, \frac{1+\sqrt{\mathbb{E}_c[\|\theta_\star^c\|^2]}}{\epsilon(1-\gamma)\nu}\right) \log\left(\frac{\psi_0}{\epsilon^2}\right)\right),$$

which improves uniformly over the number of communications required by FedLSA that we presented in the previous section. Nonetheless, the number of communication rounds for SCAFFLSA still scales with the inverse of ϵ . This seems to be due to the fact that our step-by-step analysis is loose in high-precision regimes, where the step size is very small. For this reason, we study in the rest of this section another variant of SCAFFLSA, which is more adapted to a block-by-block analysis, alike FedLSA.

Analysis under H2. In the deterministic-temporal communication paradigm **H 2**, agents average their local parameters after a *fixed* number of local updates H . In this scenario, the simplified analysis in Section 5 is not applicable. Instead, we decompose the algorithm block-by-block. Similarly to the analysis of FedLSA, we use (20) to describe the sequence of aggregated iterates as

$$\begin{aligned} \theta_{t+1} - \theta_\star &= \frac{1}{N} \sum_{c=1}^N [\Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star) \\ &+ \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)}(\xi_t^c - \xi_\star^c) - \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)], \end{aligned}$$

where t is the integer part of k/H , and $h = k \bmod H$ (we restate the algorithm with these notations in Appendix B.2). Using this expression, we can analyze SCAFFLSA under **H 2**, by studying the following Lyapunov function

$$\psi_t = \|\theta_t - \theta_\star\|^2 + \frac{\eta^2 H^2}{N} \sum_{c=1}^N \|\xi_t^c - \xi_\star^c\|^2,$$

which is defined as the error in θ_\star estimation on *communication rounds*, as well as the average error on control variates. We now state the convergence rate and iteration complexity of Algorithm 2 under the more general Assumption A2.

Theorem 5.3. *Assume A1, A2, H2. Let $\eta, H > 0$ such that $\eta \leq \min(1/C_A, \eta_\infty)$, and $H \leq \min(a/32\eta C_A^2, 1/\eta, 1/\eta a)$. Set $\xi_0^c = 0$ for all $c \in [N]$. Then we have*

$$\mathbb{E}[\|\theta_T - \theta_\star\|^2] \leq \left(1 - \frac{\eta a H}{4}\right)^T \mathbb{E}[\psi_0] + \frac{16\eta}{a} \bar{\sigma}_\omega,$$

where $\psi_0 = \|\theta_0 - \theta_\star\|^2 + \eta^2 H^2 \mathbb{E}_c[\|\bar{\mathbf{A}}^c(\theta_\star^c - \theta_\star)\|^2]$.

Corollary 5.4. Let $\epsilon > 0$. Assume $\epsilon^2 \leq \frac{32\bar{\sigma}_\omega}{aC_A^2}$, $a^2 \leq \min(1, 32C_A^2)$. Set $\eta = \mathcal{O}(\min(\eta_\infty, a\epsilon^2/\bar{\sigma}_\omega))$. We have $\mathbb{E}[\|\theta_T - \theta_\star\|^2] \leq \epsilon^2$ when the number of communication is

$$T = \mathcal{O}\left(\frac{C_A^2}{a^2} \log\left(\frac{\psi_0}{\epsilon^2}\right)\right),$$

with $\psi_0 = \mathcal{O}(\|\theta_0 - \theta_\star\|^2 + \frac{a^2}{C_A^2} \mathbb{E}_c[\|\bar{\mathbf{A}}^c(\theta_\star^c - \theta_\star)\|^2])$, and the number of local updates

$$H = \mathcal{O}\left(\max\left(\frac{a}{\eta_\infty C_A^2}, \frac{\bar{\sigma}_\omega}{\epsilon^2 C_A^2}\right)\right).$$

We prove these two statements in Appendix B.2. In Corollary 5.4, we show that with small enough step-size, the total number of communication depends only logarithmically on the precision ϵ . This is in stark contrast with Algorithm 1, where the necessity of controlling the bias’ magnitude prevents from scaling H with $1/\epsilon^2$. Thus, in high precision regime (i.e. small ϵ and η), using control variates reduces communication complexity compared to FedLSA. However, we note that, as in Corollary 5.2, we lose the linear speed-up in $1/N$. This result translates into the following communication complexity bound for federated TD(0).

Corollary 5.5. Assume TD 1 and TD 2 and let $0 < \epsilon \leq \sqrt{8\mathbb{E}_c[1 + \|\theta_\star^c\|^2]/((1-\gamma)\nu)}$. Then, in order to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] \leq \epsilon^2$, the required number of communication for federated TD(0) algorithm is

$$T = \mathcal{O}\left(\frac{1}{(1-\gamma)^2\nu^2} \log\left(\frac{\psi_0}{\epsilon^2}\right)\right),$$

where $\psi_0 = \mathcal{O}(\|\theta_0 - \theta_\star\|^2 + (1-\gamma)^2\nu^2\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|^2])$.

Corollary 5.5 confirms that, when applied to TD(0), SCAFFLSA’s communication complexity depends only logarithmically on heterogeneity and on the desired precision. In contrast with existing methods for federated TD(0) (Doan et al., 2019; Jin et al., 2022; Wang et al., 2023), it converges to the solution of system (1) even with many local steps.

6. Numerical Experiments

In this section, we illustrate the practical performance of our two algorithms for different levels of heterogeneity. To this end, we consider the classical Garnet problem (Archibald et al., 1995), in the simplified version proposed by Geist et al. (2014). These problems are characterized by the number of states n , number of actions a , and branching factor b (i.e. the number of neighbors of each state in the MDP). We set these values to $n = 30$, $a = 2$ and $b = 2$, and aim, for illustration, to evaluate the value function of the uniform policy, that chooses a random action at each time step.

Low heterogeneity regime. We generate an instance of Garnet with mentioned parameters, and perturb it by adding small uniform noise $\varepsilon_{s,a,s'} \sim \mathcal{U}([0, \epsilon])$ to each non-zero element of the transition matrix. Then we normalize each

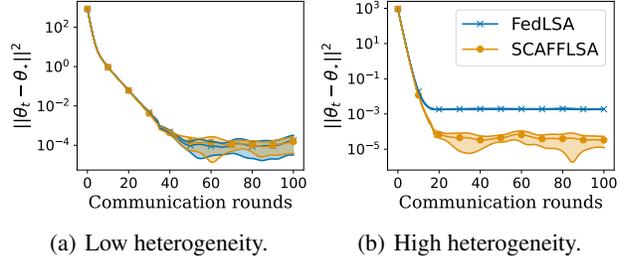


Figure 1. Mean squared error for estimation of the global optimal parameter of the value function, for FedLSA and SCAFFLSA instantiated for TD(0) on Garnet problems. For each algorithm, we plot the average MSE and standard deviation over 5 runs.

row so that they sum to 1. We generate $N = 100$ such environments, with $\epsilon = 0.0002$, and run Algorithms 1 and 2 with parameters $\eta = 0.01$, $H = 10000$ and $N = 100$. We give the results in Figure 1(a), where we see that both algorithms behave similarly. This is due to the fact that in the low heterogeneity regime local iterations cause little client drift, making bias small: there, the noise from stochastic oracles dominates (i.e. \tilde{v}_{heter} and $\bar{\rho}_H$ are smaller than $\bar{\sigma}_\epsilon$).

High heterogeneity regime. To simulate highly heterogeneous environments, we generate $N = 100$ independent Garnet environments, with the same parameters as above. We show the results in Figure 1(b), where the FedLSA stops making progress due to heterogeneity bias, while SCAFFLSA continues towards the global solution of the problem, until noise starts to dominate. Note that, according to the theory, we could always decrease the step size of FedLSA to obtain more precise results. However, this would require more communication rounds, which is not always possible.

7. Conclusion

In this paper, we studied federated linear stochastic approximation, with local training and heterogeneous agents. We analyzed the complexity of FedLSA algorithm. We carefully analyzed its bias, induced by local training, and showed how to parameterize the algorithm to control it. We then proposed SCAFFLSA, an algorithm that corrects FedLSA’s bias with control variates, and studied its complexity. Both methods are applied to federated TD learning with linear function approximation, for which we give explicit communication complexity bounds. With our analysis of SCAFFLSA, we improve iteration and communication complexity compared to FedLSA in terms of key problem parameters and desired precision, but we lose the linear speed-up in number of agents. The same problem appears in many existing analyses of gradient-based federated learning algorithms, making it a challenging future research direction.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgement

The work of P.M. and S.L has been supported by Technology Innovation Institute (TII), project Fed2Learn. The work of Eric Moulines has been partly funded by the European Union (ERC-2022-SYG-OCEAN-101071601). Ilya Levin, Alexey Naumov and Sergey Samsonov were supported by HSE University Basic Research Program. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Aguech, R., Moulines, E., and Priouret, P. On a perturbation approach for the analysis of stochastic tracking algorithms. *SIAM Journal on Control and Optimization*, 39(3):872–899, 2000.
- Archibald, T., McKinnon, K., and Thomas, L. On the generation of markov decision processes. *Journal of the Operational Research Society*, 46(3):354–361, 1995.
- Bhandari, J., Russo, D., and Singal, R. A finite time analysis of temporal difference learning with linear function approximation. In *Conference On Learning Theory*, pp. 1691–1692, 2018.
- Collins, L., Hassani, H., Mokhtari, A., and Shakkottai, S. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, pp. 2089–2099. PMLR, 2021.
- Condat, L. and Richtárik, P. Randprox: Primal-dual optimization algorithms with randomized proximal updates. *arXiv preprint arXiv:2207.12891*, 2022.
- Condat, L., Agarsky, I., and Richtárik, P. Provably doubly accelerated federated learning: The first theoretically successful combination of local training and compressed communication. *arXiv preprint arXiv:2210.13277*, 2022.
- Dal Fabbro, N., Mitra, A., and Pappas, G. J. Federated td learning over finite-rate erasure channels: Linear speedup under markovian sampling. *IEEE Control Systems Letters*, 2023.
- Dalal, G., Szörényi, B., Thoppe, G., and Mannor, S. Finite sample analyses for TD(0) with function approximation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Dann, C., Neumann, G., Peters, J., et al. Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Research*, 15:809–883, 2014.
- Doan, T., Maguluri, S., and Romberg, J. Finite-Time Analysis of Distributed TD(0) with Linear Function Approximation on Multi-Agent Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 1626–1635. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/doan19a.html>. ISSN: 2640-3498.
- Doan, T. T. Local stochastic approximation: A unified view of federated learning and distributed multi-task reinforcement learning algorithms. *arXiv preprint arXiv:2006.13460*, 2020.
- Durmus, A., Moulines, E., Naumov, A., Samsonov, S., Scaman, K., and Wai, H.-T. Tight high probability bounds for linear stochastic approximation with fixed stepsize. In Ranzato, M., Beygelzimer, A., Nguyen, K., Liang, P. S., Vaughan, J. W., and Dauphin, Y. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 30063–30074. Curran Associates, Inc., 2021.
- Durmus, A., Moulines, E., Naumov, A., and Samsonov, S. Finite-time high-probability bounds for polyak-ruppert averaged iterates of linear stochastic approximation. *arXiv preprint arXiv:2207.04475*, 2022.
- Geist, M., Scherrer, B., et al. Off-policy learning with eligibility traces: a survey. *J. Mach. Learn. Res.*, 15(1): 289–333, 2014.
- Gorbunov, E., Hanzely, F., and Richtárik, P. Local sgd: Unified theory and new efficient methods. In *International Conference on Artificial Intelligence and Statistics*, pp. 3556–3564. PMLR, 2021.
- Grudzień, M., Malinovsky, G., and Richtárik, P. Can 5th generation local training methods support client sampling? yes! In *International Conference on Artificial Intelligence and Statistics*, pp. 1055–1092. PMLR, 2023.
- Guo, L. and Ljung, L. Exponential stability of general tracking algorithms. *IEEE Transactions on Automatic Control*, 40(8):1376–1387, 1995.
- Haddadpour, F. and Mahdavi, M. On the Convergence of Local Descent Methods in Federated Learning, December 2019. URL <http://arxiv.org/abs/1910.14425>. arXiv:1910.14425 [cs, stat].

- Hu, Z. and Huang, H. Tighter analysis for proxskip. In *International Conference on Machine Learning*, pp. 13469–13496. PMLR, 2023.
- Jin, H., Peng, Y., Yang, W., Wang, S., and Zhang, Z. Federated Reinforcement Learning with Environment Heterogeneity. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pp. 18–37. PMLR, May 2022. URL <https://proceedings.mlr.press/v151/jin22a.html>. ISSN: 2640-3498.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pp. 5132–5143. PMLR, 2020.
- Khaled, A., Mishchenko, K., and Richtárik, P. Tighter theory for local sgd on identical and heterogeneous data. In *International Conference on Artificial Intelligence and Statistics*, pp. 4519–4529. PMLR, 2020.
- Khodadadian, S., Sharma, P., Joshi, G., and Maguluri, S. T. Federated reinforcement learning: Linear speedup under markovian sampling. In *International Conference on Machine Learning*, pp. 10997–11057. PMLR, 2022.
- Koloskova, A., Loizou, N., Boreiri, S., Jaggi, M., and Stich, S. A unified theory of decentralized sgd with changing topology and local updates. In *International Conference on Machine Learning*, pp. 5381–5393. PMLR, 2020.
- Konečný, J., McMahan, H. B., Ramage, D., and Richtárik, P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
- Li, G., Wu, W., Chi, Y., Ma, C., Rinaldo, A., and Wei, Y. Sharp high-probability sample complexities for policy evaluation with linear function approximation, 2023a.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3):50–60, 2020.
- Li, T., Lan, G., and Pananjady, A. Accelerated and instance-optimal policy evaluation with linear function approximation. *SIAM Journal on Mathematics of Data Science*, 5(1):174–200, 2023b. doi: 10.1137/21M1468668. URL <https://doi.org/10.1137/21M1468668>.
- Lim, H.-K., Kim, J.-B., Heo, J.-S., and Han, Y.-H. Federated reinforcement learning for training control policies on multiple iot devices. *Sensors*, 20(5):1359, 2020.
- Liu, R. and Olshevsky, A. Distributed td (0) with almost no communication. *IEEE Control Systems Letters*, 2023.
- Malinovsky, G., Yi, K., and Richtárik, P. Variance reduced proxskip: Algorithm, theory and application to federated learning. *Advances in Neural Information Processing Systems*, 35:15176–15189, 2022.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.
- Mishchenko, K., Malinovsky, G., Stich, S., and Richtárik, P. Proxskip: Yes! local gradient steps provably lead to communication acceleration! finally! In *International Conference on Machine Learning*, pp. 15750–15769. PMLR, 2022.
- Mitra, A., Jaafar, R., Pappas, G. J., and Hassani, H. Linear convergence in federated learning: Tackling client heterogeneity and sparse gradients. *Advances in Neural Information Processing Systems*, 34:14606–14619, 2021.
- Patel, K. K., Glasgow, M., Wang, L., Joshi, N., and Srebro, N. On the still unreasonable effectiveness of federated averaging for heterogeneous distributed learning. In *Federated Learning and Analytics in Practice: Algorithms, Systems, Applications, and Opportunities*, 2023. URL <https://openreview.net/forum?id=vhS68bKv7x>.
- Patil, G., Prashanth, L., Nagaraj, D., and Precup, D. Finite time analysis of temporal difference learning with linear function approximation: Tail averaging and regularisation. In *International Conference on Artificial Intelligence and Statistics*, pp. 5438–5448. PMLR, 2023.
- Priouret, P. and Veretenikov, A. A remark on the stability of the LMS tracking algorithm. *Stochastic analysis and applications*, 16(1):119–129, 1998.
- Qi, J., Zhou, Q., Lei, L., and Zheng, K. Federated reinforcement learning: Techniques, applications, and open challenges. *arXiv preprint arXiv:2108.11887*, 2021.
- Qu, Z., Lin, K., Li, Z., and Zhou, J. Federated learning’s blessing: Fedavg has linear speedup. In *ICLR 2021-Workshop on Distributed and Private Machine Learning (DPML)*, 2021.
- Reddi, S., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., and McMahan, H. B. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*, 2020.
- Robbins, H. and Monro, S. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400 – 407, 1951. doi: 10.1214/aoms/1177729586. URL <https://doi.org/10.1214/aoms/1177729586>.

-
- Samsonov, S., Tiapkin, D., Naumov, A., and Moulines, E. Finite-sample analysis of the Temporal Difference Learning. *arXiv preprint arXiv:2310.14286*, 2023.
- Srikant, R. and Ying, L. Finite-time error bounds for linear stochastic approximation and TD learning. In *Conference on Learning Theory*, pp. 2803–2830. PMLR, 2019.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3:9–44, 1988.
- Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvári, C., and Wiewiora, E. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th annual international conference on machine learning*, pp. 993–1000, 2009.
- Tsitsiklis, J. N. and Van Roy, B. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690, May 1997. ISSN 2334-3303. doi: 10.1109/9.580874.
- Wang, H., Mitra, A., Hassani, H., Pappas, G. J., and Anderson, J. Federated temporal difference learning with linear function approximation under environmental heterogeneity. *arXiv preprint arXiv:2302.02212*, 2023.
- Wang, J., Das, R., Joshi, G., Kale, S., Xu, Z., and Zhang, T. On the unreasonable effectiveness of federated averaging with heterogeneous data. *arXiv preprint arXiv:2206.04723*, 2022.
- Xie, Z. and Song, S. Fedkl: Tackling data heterogeneity in federated reinforcement learning by penalizing kl divergence. *IEEE Journal on Selected Areas in Communications*, 41(4):1227–1242, 2023.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.

A. Analysis of Federated Linear Stochastic Approximation

For the analysis we need to define two filtration: $\mathcal{F}_{s,h}^+ := \sigma(Z_{t,k}^c, t \geq s, k \geq h, 1 \leq c \leq N)$ (future events) and $\mathcal{F}_{s,h}^- := \sigma(Z_{t,k}^c, t \leq s, k \leq h, 1 \leq c \leq N)$ (preceding events). Recall that the local LSA updates are written as

$$\theta_{t,h}^c - \theta_*^c = (\mathbf{I} - \eta \mathbf{A}(Z_{t,h}^c))(\theta_{t,h-1}^c - \theta_*^c) - \eta \varepsilon^c(Z_{t,h}^c).$$

Performing H local steps and taking average, we end up with the decomposition (17), which we duplicate here for user's convenience:

$$\theta_t - \theta_* = \tilde{\theta}_t^{(\text{tr})} + \tilde{\theta}_t^{(\text{bi,bi})} + \tilde{\theta}_t^{(\text{fl,bi})} + \tilde{\theta}_t^{(\text{fl})}, \quad (22)$$

where we have defined

$$\begin{aligned} \tilde{\theta}_t^{(\text{tr})} &= \prod_{s=1}^t \bar{\Gamma}_{s,H}^{(\eta)} \{\theta_0 - \theta_*\}, \\ \tilde{\theta}_t^{(\text{bi,bi})} &= \sum_{s=1}^t (\bar{\Gamma}_H^{(\eta)})^{t-s} \bar{\rho}_H, \\ \tilde{\theta}_t^{(\text{fl,bi})} &= \sum_{s=1}^t \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{r}_{s,H} + \Delta_{H,s,t}^{(\eta)} \bar{\rho}_H, \\ \tilde{\theta}_t^{(\text{fl})} &= \eta \sum_{s=1}^t \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{\varphi}_{s,H}. \end{aligned} \quad (23)$$

Now we need to upper bound each of the terms in decomposition (22). This is done in a sequence of lemmas below: $\tilde{\theta}_t^{(\text{fl})}$ is bounded in Lemma A.1, $\tilde{\theta}_t^{(\text{fl,bi})}$ in Lemma A.2, $\tilde{\theta}_t^{(\text{tr})}$ in Lemma A.3, and $\tilde{\theta}_t^{(\text{bi,bi})}$ in Lemma A.4. Then we combine the bounds in order to state a version of Theorem 4.1 with explicit constants in Theorem A.5.

Lemma A.1. *Assume A1 and A2. Then, for any step size $\eta \in (0, \eta_\infty)$ it holds*

$$\mathbb{E}[\|\tilde{\theta}_t^{(\text{fl})}\|^2] \leq \frac{\eta \bar{\sigma}_\varepsilon}{aN(1 - e^{-2})}.$$

Proof. We start from the decomposition (23). With the definition of $\tilde{\theta}_t^{(\text{fl})}$ and $\mathbb{E}^{\mathcal{F}_{s+1,1}^+}[\{\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)}\} \bar{\varphi}_{s,H}] = 0$, we obtain that

$$\mathbb{E}[\|\tilde{\theta}_t^{(\text{fl})}\|^2] = \eta^2 \sum_{s=1}^t \mathbb{E}[\|\{\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)}\} \bar{\varphi}_{s,H}\|^2].$$

Now, using the assumption A2 and Minkowski's inequality, we obtain that

$$\mathbb{E}^{1/2}[\|\{\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)}\} \bar{\varphi}_{s,H}\|^2] \leq \frac{1}{N} \sum_{c=1}^N \mathbb{E}^{1/2}[\|\bar{\Gamma}_{t,H}^{(c,\eta)}\{\prod_{i=s+1}^{t-1} \bar{\Gamma}_{i,H}^{(\eta)}\} \bar{\varphi}_{s,H}\|^2] \stackrel{(a)}{\leq} (1 - \eta a)^H \mathbb{E}^{1/2}[\|\{\prod_{i=s+1}^{t-1} \bar{\Gamma}_{i,H}^{(\eta)}\} \bar{\varphi}_{s,H}\|^2]. \quad (24)$$

In (a) applied A2 conditionally on $\mathcal{F}_{t-1,H}^-$. Hence, by induction we get from the previous formulas that

$$\mathbb{E}[\|\tilde{\theta}_t^{(\text{fl})}\|^2] \leq \eta^2 \sum_{s=1}^t (1 - \eta a)^H \mathbb{E}[\|\bar{\varphi}_{s,H}\|^2]. \quad (25)$$

Now we proceed with bounding $\mathbb{E}[\|\bar{\varphi}_{s,H}\|^2]$. Indeed, since the clients are independent, we get using (14) that

$$\begin{aligned} \mathbb{E}[\|\bar{\varphi}_{s,H}\|^2] &= \frac{1}{N^2} \sum_{c=1}^N \mathbb{E}[\|\sum_{h=1}^H \Gamma_{s,h+1:H}^{(c,\eta)} \varepsilon^c(Z_{s,h}^c)\|^2] \\ &= \frac{1}{N^2} \sum_{c=1}^N \left[\sum_{h=1}^H \mathbb{E}[\|\Gamma_{s,h+1:H}^{(c,\eta)} \varepsilon^c(Z_{s,h}^c)\|^2] \right] \\ &\leq \frac{1}{N^2} \sum_{c=1}^N \sum_{h=1}^H (1 - \eta a)^{2(H-h)} \mathbb{E}[\|\varepsilon^c(Z_{s,h}^c)\|^2]. \end{aligned}$$

Therefore, using (5) and the following inequality,

$$\sum_{h=0}^{H-1} (1 - \eta a)^{2h} \leq H \wedge \frac{1}{\eta a}, \quad \text{for all } \eta \geq 0, \text{ such that } \eta a \leq 1,$$

we get

$$\mathbb{E}[\|\tilde{\varphi}_{s,H}\|^2] \leq \frac{1}{N} \left(H \wedge \frac{1}{\eta a} \right) \bar{\sigma}_\varepsilon.$$

Plugging this inequality in (25), we get

$$\begin{aligned} \mathbb{E}[\|\tilde{\theta}_t^{(\text{fl})}\|^2] &\leq \frac{\mathbb{E}_c[\text{Tr}(\Sigma_\varepsilon^c)]}{N} \left(\eta^2 H \wedge \frac{\eta}{a} \right) \sum_{s=1}^t [(1-\eta a)^{2H(t-s)}] \\ &\leq \frac{\bar{\sigma}_\varepsilon}{N} \left(\eta^2 H \wedge \frac{\eta}{a} \right) \frac{1}{1 - (1-\eta a)^{2H}} \\ &\leq \frac{\eta \bar{\sigma}_\varepsilon}{aN} (\eta a H \wedge 1) \frac{1}{1 - e^{-2\eta a H}}, \end{aligned}$$

where we used additionally

$$e^{-2x} \leq 1 - x \leq e^{-x}, \quad (26)$$

which is valid for $x \in [0; 1/2]$. Now it remains to notice that

$$\frac{x \wedge 1}{1 - e^{-2x}} \leq \frac{1}{1 - e^{-2}}$$

for any $x > 0$. □

We proceed with analyzing the fluctuation of the true bias component of the error θ_t defined in (23). The first step towards this is to obtain the respective bound for $\bar{\tau}_{s,H}$, $s \in \{1, \dots, T\}$, where $\bar{\tau}_{s,H}$ is defined in (14). Now we provide an upper bound for $\tilde{\theta}_t^{(\text{fl,bi})}$:

Lemma A.2. *Assume A1 and A2. Then, for any step size $\eta \in (0, \eta_\infty)$ it holds*

$$\mathbb{E}^{1/2}[\|\tilde{\theta}_t^{(\text{fl,bi})}\|^2] \leq \sqrt{\frac{2\eta \tilde{v}_{\text{heter}}}{Na}} + \frac{2\sqrt{\mathbb{E}_c\|\Sigma_{\tilde{\mathbf{A}}}^c\|}\|\bar{\rho}_H\|}{aH^{1/2}N^{1/2}}.$$

Proof. Recall that $\tilde{\theta}_t^{(\text{fl,bi})}$ is given (see (23)) by

$$\tilde{\theta}_t^{(\text{fl,bi})} = \underbrace{\sum_{s=1}^t \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{\tau}_{s,H}}_{T_1} + \underbrace{\left(\sum_{s=1}^t \left\{ \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \right\} - (\bar{\Gamma}_H^{(\eta)})^{t-s} \right) \bar{\rho}_H}_{T_2}, \quad (27)$$

where $\bar{\tau}_{s,H}$ and $\bar{\rho}_H$ are defined in (14). We begin with bounding T_1 . In order to do it we first need to bound $\bar{\tau}_{s,H}$. Since the different agents are independent, we have

$$\mathbb{E}[\|\bar{\tau}_{s,H}\|^2] = \frac{1}{N^2} \sum_{c=1}^N \mathbb{E}[\|((I - \eta \bar{\mathbf{A}}^c)^H - \Gamma_{s,1:H}^{(c,\eta)})\{\theta_\star^c - \theta_\star\}\|^2]. \quad (28)$$

Applying Lemma C.1 and the fact that $\{(I - \eta \bar{\mathbf{A}}^c)^{h-1} \tilde{\mathbf{A}}^c(Z_{s,h}^c) \Gamma_{s,(h+1):H}^{(c,\eta)} (\theta_\star^c - \theta_\star)\}_{h=1}^H$ is a martingale-difference w.r.t. $\mathcal{F}_{s,h}^-$, we get that

$$\begin{aligned} &\mathbb{E}[\|((I - \eta \bar{\mathbf{A}}^c)^H - \Gamma_{s,1:H}^{(c,\eta)})\{\theta_\star^c - \theta_\star\}\|^2] \\ &= \eta^2 \mathbb{E}[\| \sum_{h=1}^H (I - \eta \bar{\mathbf{A}}^c)^{h-1} \tilde{\mathbf{A}}^c(Z_{s,h}^c) \Gamma_{s,(h+1):H}^{(c,\eta)} \{\theta_\star^c - \theta_\star\} \|^2] \\ &= \eta^2 \sum_{h=1}^H \mathbb{E}[\| (I - \eta \bar{\mathbf{A}}^c)^{h-1} \tilde{\mathbf{A}}^c(Z_{s,h}^c) \Gamma_{s,(h+1):H}^{(c,\eta)} \{\theta_\star^c - \theta_\star\} \|^2] \\ &\leq \eta^2 \sum_{h=1}^H (1 - \eta a)^{2(h-1)} \{\theta_\star^c - \theta_\star\}^\top \mathbb{E}[(\Gamma_{s,(h+1):H}^{(c,\eta)})^\top (\tilde{\mathbf{A}}^c(Z_{s,h}^c))^\top \tilde{\mathbf{A}}^c(Z_{s,h}^c) \Gamma_{s,(h+1):H}^{(c,\eta)}] \{\theta_\star^c - \theta_\star\}. \end{aligned}$$

Using the tower property conditionally on $\mathcal{F}_{s,h+1}^+$, we get

$$\mathbb{E}[(\Gamma_{s,(h+1):H}^{(c,\eta)})^\top (\tilde{\mathbf{A}}^c(Z_{s,h}^c))^\top \tilde{\mathbf{A}}^c(Z_{s,h}^c) \Gamma_{s,(h+1):H}^{(c,\eta)}] = \mathbb{E}[(\Gamma_{s,(h+1):H}^{(c,\eta)})^\top \Sigma_{\tilde{\mathbf{A}}}^c \Gamma_{s,(h+1):H}^{(c,\eta)}],$$

where $\Sigma_{\tilde{\mathbf{A}}}^c$ is the noise covariance matrix defined in (6). Since for any vector $u \in \mathbb{R}^d$ we have $\|u\|_{\Sigma_{\tilde{\mathbf{A}}}^c} \leq \|\Sigma_{\tilde{\mathbf{A}}}^c\|^{1/2}\|u\|$, we get

$$\begin{aligned} \mathbb{E}[\|((\mathbf{I} - \eta \tilde{\mathbf{A}}^c)^H - \Gamma_{s,1:H}^{(c,\eta)})\{\theta_\star^c - \theta_\star\}\|^2] &\leq \eta^2 \sum_{h=1}^H (1 - \eta a)^{2(h-1)} \mathbb{E}[\|\Gamma_{s,(h+1):H}^{(c,\eta)}\{\theta_\star^c - \theta_\star\}\|_{\Sigma_{\tilde{\mathbf{A}}}^c}^2] \\ &\leq \eta^2 \|\Sigma_{\tilde{\mathbf{A}}}^c\| \sum_{h=1}^H (1 - \eta a)^{2(h-1)} \mathbb{E}[\|\Gamma_{s,(h+1):H}^{(c,\eta)}\{\theta_\star^c - \theta_\star\}\|^2] \\ &\leq H\eta^2 (1 - \eta a)^{2(H-1)} \|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2. \end{aligned} \quad (29)$$

Combining the above bounds in (28) yields that

$$\mathbb{E}[\|\bar{\tau}_{s,H}\|^2] \leq \frac{H\eta^2 (1 - \eta a)^{2(H-1)} \sum_{c=1}^N \|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2}{N^2}. \quad (30)$$

Thus, proceeding as in (24) together with (30), we get

$$\begin{aligned} \mathbb{E}[\|T_1\|^2] &= \sum_{s=1}^t \mathbb{E}[\|\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \bar{\tau}_{s,H}\|^2] \\ &\leq \sum_{s=1}^t \frac{H\eta^2 (1 - \eta a)^{2(H-1)} \sum_{c=1}^N \|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2}{N^2} (1 - \eta a)^{2H(t-s)} \\ &\leq \frac{H\eta^2 (1 - \eta a)^{2(H-1)}}{(1 - (1 - \eta a)^{2H})N} \mathbb{E}_c[\|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2] \\ &\leq \frac{\eta}{aN(1 - \eta a)^2} \frac{Ha\eta e^{-2Ha\eta}}{1 - e^{-2Ha\eta}} \mathbb{E}_c[\|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2] \\ &\leq \frac{2\eta}{Na} \mathbb{E}_c[\|\Sigma_{\tilde{\mathbf{A}}}^c\| \|\theta_\star^c - \theta_\star\|^2]. \end{aligned}$$

In the bound above we used (26) together with the bound

$$\frac{xe^{-2x}}{1 - e^{-2x}} \leq \frac{1}{2}, x \geq 0.$$

Now we bound the second part of $\tilde{\theta}_t^{(f,bi)}$ in (27), that is, T_2 . To begin with, we start with applying Lemma C.1 and we get for any $s \in \{1, \dots, t\}$ and $i \in \{s+1, \dots, t\}$, that

$$\left\{ \prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)} \right\} - (\bar{\Gamma}_H^{(\eta)})^{t-s} \bar{\rho}_H = \sum_{i=s+1}^t \left\{ \prod_{r=i+1}^t \bar{\Gamma}_{r,H}^{(\eta)} \right\} (\bar{\Gamma}_{i,H}^{(\eta)} - \bar{\Gamma}_H^{(\eta)}) (\bar{\Gamma}_H^{(\eta)})^{i-s-1} \bar{\rho}_H.$$

Note that,

$$\mathbb{E}^{\mathcal{F}_{i+1,1}^+} \left[\left\{ \prod_{r=i+1}^t \bar{\Gamma}_{r,H}^{(\eta)} \right\} (\bar{\Gamma}_{i,H}^{(\eta)} - \bar{\Gamma}_H^{(\eta)}) (\bar{\Gamma}_H^{(\eta)})^{i-s-1} \bar{\rho}_H \right] = 0 \quad (31)$$

Proceeding as in (29), we get using independence between agents for any $u \in \mathbb{R}^d$,

$$\begin{aligned} \mathbb{E}[\|(\bar{\Gamma}_{i,H}^{(\eta)} - \bar{\Gamma}_H^{(\eta)})u\|^2] &= \frac{1}{N^2} \mathbb{E}[\|\sum_{c=1}^N (\Gamma_{s,1:H}^{(c,\eta)} - (\mathbf{I} - \eta \tilde{\mathbf{A}}^c)^H)u\|^2] \\ &= \frac{1}{N^2} \sum_{c=1}^N \mathbb{E}[\|(\Gamma_{s,1:H}^{(c,\eta)} - (\mathbf{I} - \eta \tilde{\mathbf{A}}^c)^H)u\|^2] \\ &\leq \frac{H\eta^2 (1 - \eta a)^{2(H-1)}}{N} \left(\frac{1}{N} \sum_{c=1}^N \|\Sigma_{\tilde{\mathbf{A}}}^c\| \right) \|u\|^2. \end{aligned}$$

Hence, using (31), we get

$$\mathbb{E}[\|(\{\prod_{i=s+1}^t \bar{\Gamma}_{i,H}^{(\eta)}\} - (\bar{\Gamma}_H^{(\eta)})^{t-s})\bar{\rho}_H\|^2] = \frac{H\eta^2(1-\eta a)^{2H(t-s)-2}\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}{N}\|\bar{\rho}_H\|^2.$$

Combining the above estimates in (27), and using Minkowski's inequality, we get

$$\begin{aligned}\mathbb{E}^{1/2}[\|T_2\|^2] &\leq \frac{H^{1/2}\eta}{(1-\eta a)N^{1/2}}\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\|\sum_{s=1}^{t-1}(1-\eta a)^{H(t-s)} \\ &\leq \frac{2}{aH^{1/2}N^{1/2}}\frac{Ha\eta e^{-Ha\eta}}{1-e^{-Ha\eta}}\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\| \\ &\leq \frac{2}{aH^{1/2}N^{1/2}}\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\|,\end{aligned}$$

where we used that $\eta a \leq 1/2$ and

$$\frac{xe^{-x}}{1-e^{-x}} \leq 1, \quad x \geq 0.$$

and the statement follows. \square

Lemma A.3. Assume A1 and A2. Then for any step size $\eta \in (0, \eta_\infty)$ we have

$$\mathbb{E}^{1/2}[\|\tilde{\theta}_t^{(\text{tr})}\|^2] \leq (1-\eta a)^{tH}\|\theta_0 - \theta_\star\|$$

Proof. Proceeding as in (24) for any $u \in \mathbb{R}^d$ we have

$$\mathbb{E}^{1/2}[\|\prod_{s=1}^t \bar{\Gamma}_{s,H}^{(\eta)} u\|^2] \leq (1-\eta a)^{tH}\|u\|$$

Using this result for $u = \theta_0 - \theta_\star$ we get the statement. \square

Lemma A.4. Assume A1 and A2. Then for any $\eta \in (0, \eta_\infty)$ we have

$$\|\tilde{\theta}_t^{(\text{bi,bi})} - (\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\bar{\rho}_H\| \leq (1-\eta a)^{tH}\|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\|\|\bar{\rho}_H\|$$

Proof. Using A2 and Minkowski's inequality, we get

$$\begin{aligned}\|\tilde{\theta}_t^{(\text{bi,bi})} - (\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\bar{\rho}_H\| &= \|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}(\bar{\Gamma}_H^{(\eta)})^t\bar{\rho}_H\| \\ &\leq \|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\|\|(\bar{\Gamma}_H^{(\eta)})^t\bar{\rho}_H\| \\ &\leq \|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\|\frac{1}{N}\sum_{c=1}^N\|(\mathbf{I} - \eta\mathbf{A}^c)^H(\bar{\Gamma}_H^{(\eta)})^{t-1}\bar{\rho}_H\| \\ &\leq (1-\eta a)^H\|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\|\|(\bar{\Gamma}_H^{(\eta)})^{t-1}\bar{\rho}_H\|\end{aligned}$$

and the statement follows. \square

Theorem A.5. Assume A1 and A2. Then for any step size $\eta \in (0, \eta_\infty)$ it holds that

$$\mathbb{E}^{1/2}[\|\theta_t - \tilde{\theta}_t^{(\text{bi,bi})} - \theta_\star\|^2] \leq \sqrt{\frac{\eta\bar{\sigma}_\varepsilon}{aN(1-e^{-2})}} + \sqrt{\frac{2\eta\tilde{v}_{\text{heter}}}{Na}} + \frac{2\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\|}{aH^{1/2}N^{1/2}} + (1-\eta a)^{tH}\|\theta_0 - \theta_\star\|, \quad (32)$$

where the bias $\tilde{\theta}_t^{(\text{bi,bi})}$ converges to $(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\bar{\rho}_H$ at a rate

$$\|\tilde{\theta}_t^{(\text{bi,bi})} - (\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\bar{\rho}_H\| \leq (1-\eta a)^{tH}\|(\mathbf{I} - \bar{\Gamma}_H^{(\eta)})^{-1}\|\|\bar{\rho}_H\|.$$

Proof. Proof follows by combining the results Lemma A.1-Lemma A.4 above. \square

Corollary A.6. Assume A1 and A2. Let $H = 1$, then for any $0 < \epsilon < 1$, in order to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] \leq \epsilon^2$ the required number of communications is

$$T = \mathcal{O}\left(\frac{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}{Na^2\epsilon^2} \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right)$$

number of communications, setting the step size

$$\eta_0 = \frac{aN\epsilon^2}{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}. \quad (33)$$

Proof. Bounding the first two terms in decomposition (32) we get that the step size should satisfy

$$\eta \leq \frac{aN\epsilon^2}{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}.$$

From the last term we have

$$t \geq \frac{1}{a\eta} \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon} \geq \frac{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}{Na^2\epsilon^2} \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}$$

\square

Corollary A.7. Assume A1 and A2. For any $0 \leq \epsilon \leq \frac{C_{\mathbf{A}}^{-1} \mathbb{E}_c \|\theta_\star^c - \theta_\star\|}{a} \vee \left(\frac{\sqrt{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon} \mathbb{E}_c \|\theta_\star^c - \theta_\star\|}{a}\right)^{2/5}$ in order to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] < \epsilon^2$ the required number of communications is

$$T = \mathcal{O}\left(\frac{\mathbb{E}_c \|\theta_\star^c - \theta_\star\|}{a^2\epsilon} \log \frac{\|\theta_0 - \theta_\star\|}{\epsilon}\right), \quad (34)$$

setting the step size

$$\eta = \mathcal{O}\left(\frac{aN\epsilon^2}{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}\right)$$

and number of local iterations

$$H = \mathcal{O}\left(\frac{\tilde{v}_{heter} \vee \bar{\sigma}_\epsilon}{N\epsilon \mathbb{E}_c \|\theta_\star^c - \theta_\star\|}\right)$$

Proof. We aim to bound separately all the terms in the r.h.s. of Theorem 4.1. Note that it requires to set $\eta \in (0; \eta_0)$ with η_0 given in (33) in order to fulfill the bounds

$$\sqrt{\frac{\eta \tilde{v}_{heter}}{aN}} \lesssim \epsilon, \quad \sqrt{\frac{\eta \bar{\sigma}_\epsilon}{aN}} \lesssim \epsilon.$$

Now, we should bound the bias term

$$\mathbb{E}^{1/2}[\|\tilde{\theta}_t^{(\text{bi, bi})}\|^2] \leq (1 + (1 - \eta a)^{tH}) \|(I - \bar{\Gamma}_H^{(\eta)})^{-1} \bar{\rho}_H\| \leq 2 \|(I - \bar{\Gamma}_H^{(\eta)})^{-1} \bar{\rho}_H\|.$$

Thus, using the Neuman series, we can bound the norm of the term above as

$$\|(I - \bar{\Gamma}_H^{(\eta)})^{-1} \bar{\rho}_H\| = \left\| \sum_{k=0}^{\infty} (\bar{\Gamma}_H^{(\eta)})^k \bar{\rho}_H \right\| \leq \sum_{k=0}^{\infty} (1 - \eta a)^{Hk} \|\bar{\rho}_H\| \leq \frac{\|\bar{\rho}_H\|}{1 - (1 - \eta a)^H}.$$

Hence, using the bound of (16), we get

$$\begin{aligned} \mathbb{E}^{1/2}[\|\tilde{\theta}_t^{(\text{bi, bi})}\|^2] &\leq \frac{2\|\bar{\rho}_H\|}{1 - (1 - \eta a)^H} \leq \frac{\eta a H}{1 - (1 - \eta a)^H} \frac{\eta H \mathbb{E}_c[\exp(\eta H \|\bar{\mathbf{A}}^c\|) \|\theta_\star^c - \theta_\star\|]}{a} \\ &\leq \frac{2\eta H \mathbb{E}_c[\exp(\eta H \|\bar{\mathbf{A}}^c\|) \|\theta_\star^c - \theta_\star\|]}{a} \lesssim \frac{\eta H \mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{a}, \end{aligned}$$

Algorithm 3 SCAFFLSA-H 1: Stochastic Controlled FedLSA with probabilistic communication

Input: $\eta > 0, \theta_0, \xi_0^c \in \mathbb{R}^d, T, N, H, p > 0$

Set: $K = T/p$

for $k = 1$ to K **do**

for $c = 1$ to N **do**

 Receive Z_k^c and perform local update:

$$\hat{\theta}_k^c = \hat{\theta}_{k-1}^c - \eta(\mathbf{A}^c(Z_k^c)\hat{\theta}_{k-1}^c - \mathbf{b}^c(Z_k^c) - \xi_{k-1}^c)$$

end for

 Draw $B_k \sim \text{Bernoulli}(p)$

if $B_k = 1$ **then**

 Average local iterates: $\theta_k^c = \frac{1}{N} \sum_{c=1}^N \hat{\theta}_k^c$

 Update: $\xi_k^c = \xi_{k-1}^c + \frac{p}{\eta}(\theta_k^c - \hat{\theta}_k^c)$

else

 Set: $\theta_k^c = \hat{\theta}_k^c, \xi_k^c = \xi_{k-1}^c$

end if

end for

where we used the fact that the step size η is chosen in order to satisfy $\eta H C_{\mathbf{A}} \leq 1$. Thus in order to fulfill $\mathbb{E}^{1/2}[\|\tilde{\theta}_t^{(\text{bi}, \text{bi})}\|^2] \lesssim \varepsilon$ we need to choose η and H such that

$$\eta H \mathbb{E}_c[\|\theta_\star^c - \theta_\star\|] \leq \varepsilon a.$$

It remains to bound the term $\frac{\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\|}{aH^{1/2}N^{1/2}}$. Using the bound of (16), we get

$$\frac{\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}\|\bar{\rho}_H\|}{aH^{1/2}N^{1/2}} \leq \sqrt{\frac{\eta}{N}} \times \frac{\sqrt{\mathbb{E}_c\|\Sigma_{\mathbf{A}}^c\|}(\eta H)^{3/2}}{a} \lesssim \varepsilon^{5/2} \sqrt{\frac{1}{\tilde{v}_{\text{heter}} \vee \bar{\sigma}_\varepsilon}} \frac{a}{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}.$$

Hence, it remains to combine the bounds above in order to get the sample complexity result (34). □

Corollary A.8. Assume TD 1 and TD 2. Then for any

$$0 \leq \varepsilon \leq \frac{2\left(\sqrt{2}(1+\gamma)\sqrt{\mathbb{E}_c\|\theta_\star^c - \theta_\star\|^2 \vee (1+\mathbb{E}_c[\|\theta_\star\|^2])\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}\right)^{2/5}}{(1-\gamma)\nu} \sqrt{\frac{2\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{(1-\gamma)\nu(1+\gamma)}},$$

in order to achieve $\mathbb{E}[\|\theta_T - \theta_\star\|^2] < \varepsilon^2$ the required number of communications for federated TD(0) algorithm is

$$T = \mathcal{O}\left(\left(\frac{1}{(1-\gamma)^2\nu} \vee \frac{\mathbb{E}_c[\|\theta_\star^c - \theta_\star\|]}{(1-\gamma)^2\nu^2\varepsilon}\right) \log \frac{\|\theta_0 - \theta_\star\|}{\varepsilon}\right).$$

B. Federated Linear Stochastic Approximation with Control Variates

B.1. Probabilistic communication (Assumption H 1)

To mitigate the bias caused by local training, we may use control variates. We assume in this section that at each iteration we choose, with probability p , whether agents should communicate or not. Consider the following algorithm, where for $k = 1, \dots, T/p$, we compute

$$\hat{\theta}_k^c = \theta_{k-1}^c - \eta(\mathbf{A}^c(Z_k^c)\theta_{k-1}^c - \mathbf{b}^c(Z_k^c) - \xi_{k-1}^c)$$

i.e. we update the local parameters with LSA adjusted with a control variate ξ_{k-1}^c . This control variate is initialized to zero, and updated after each communication round. We draw a Bernoulli random variable B_k with success probability p and then update the parameter as follows:

$$\theta_k^c = \begin{cases} \bar{\theta}_k = \frac{1}{N} \sum_{c=1}^N \hat{\theta}_k^c & B_k = 1 \\ \hat{\theta}_k^c & B_k = 0 \end{cases}$$

We then update the control variate

$$\xi_k^c = \xi_{k-1}^c + \frac{p}{\eta}(\theta_k^c - \hat{\theta}_k^c). \quad (35)$$

where we have set $\xi_0^c = 0$. For clarity, we restate Algorithm 2 under this assumption as Algorithm 3.

By construction, θ_k^c is the current value of parameter the parameter at time k . The control variate stays constant between two successive consensus steps. Note that, for all $k \in \mathbb{N}$, $\sum_{c=1}^N \xi_k^c = 0$. Indeed, if $B_k = 0$, for any $c \in \{1, \dots, N\}$, $\xi_k^c = \xi_{k-1}^c$. Additionally, if $B_k = 1$, we have from (35) that

$$\frac{1}{N} \sum_{c=1}^N \xi_k^c = \frac{1}{N} \sum_{c=1}^N \xi_{k-1}^c + \frac{p}{\eta}(\bar{\theta}_k - \bar{\theta}_k) = 0.$$

We now proceed to the proof, which amounts to constructing a common Lyapunov function for the sequences $\{\theta_k^c\}_{k \in \mathbb{N}}$ and $\{\xi_k^c\}_{k \in \mathbb{N}}$. Define the Lyapunov function,

$$\psi_k = \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2,$$

where θ_\star is the solution of $\bar{\mathbf{A}}\theta_\star = \bar{\mathbf{b}}$, and $\xi_\star^c = \bar{\mathbf{A}}^c(\theta_\star - \theta_\star^c)$. A natural measure of heterogeneity is then given by

$$\Delta_{\text{heter}} = \frac{1}{N} \sum_{c=1}^N \|\xi_\star^c\|^2 = \frac{1}{N} \sum_{c=1}^N \|\bar{\mathbf{A}}^c(\theta_\star - \theta_\star^c)\|^2.$$

Lemma B.1 (One step progress). *Assume A1 and A2(2). Assume that $\eta \leq \frac{1}{2L}$. The iterates of the algorithm described above satisfy*

$$\mathbb{E}[\psi_k] \leq \left(1 - \min(\eta a, p^2)\right) \mathbb{E}[\psi_{k-1}] + \frac{2\eta^2}{N} \sum_{c=1}^N \text{Tr}(\Sigma_\varepsilon^c).$$

Proof. Decomposition of the update. Remark that the update can be reformulated as

$$\hat{\theta}_k^c - \theta_\star = (\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))(\theta_{k-1}^c - \theta_\star) + \eta(\xi_{k-1}^c - \xi_\star^c) - \eta \omega^c(Z_k^c), \quad (36)$$

where $\omega^c(z) = \tilde{\mathbf{A}}^c(z)\theta_\star - \tilde{\mathbf{b}}^c(z)$. This comes from the fact that, for all z ,

$$\begin{aligned} \mathbf{b}^c(z) + \xi_{k-1}^c &= \bar{\mathbf{b}}^c + \tilde{\mathbf{b}}^c(z) + \xi_{k-1}^c \\ &= \bar{\mathbf{A}}^c \theta_\star + \tilde{\mathbf{b}}^c(z) + \xi_{k-1}^c \\ &= \bar{\mathbf{A}}^c \theta_\star + \tilde{\mathbf{b}}^c(z) + \xi_{k-1}^c - \xi_\star^c \\ &= \mathbf{A}^c(z)\theta_\star - \tilde{\mathbf{A}}^c(z)\theta_\star + \tilde{\mathbf{b}}^c(z) + \xi_{k-1}^c - \xi_\star^c \\ &= \mathbf{A}^c(z)\theta_\star - \omega^c(z) + \xi_{k-1}^c - \xi_\star^c. \end{aligned}$$

Expression of communication steps. Using that $\sum_{c=1}^N \xi_{k-1}^c = 0$ and $\sum_{c=1}^N \xi_\star^c = 0$, we get

$$\begin{aligned} \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 &= \mathbf{1}_{\{1\}}(B_k) \|\bar{\theta}_k - \theta_\star\|^2 + \mathbf{1}_{\{0\}}(B_k) \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \theta_\star\|^2 \\ &= \mathbf{1}_{\{1\}}(B_k) \left\| \frac{1}{N} \sum_{c=1}^N \left(\hat{\theta}_k^c - \frac{\eta}{p} \xi_{k-1}^c \right) - \theta_\star \right\|^2 + \mathbf{1}_{\{0\}}(B_k) \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \theta_\star\|^2. \end{aligned}$$

The first term can be upper bounded by using Lemma C.4, which gives

$$\begin{aligned} \mathbf{1}_{\{1\}}(B_k) \|\bar{\theta}_k - \theta_\star\|^2 &= \mathbf{1}_{\{1\}}(B_k) \left\{ \frac{1}{N} \sum_{c=1}^N \left\| \hat{\theta}_k^c - \frac{\eta}{p} (\xi_{k-1}^c - \xi_\star^c) - \theta_\star \right\|^2 - \frac{1}{N} \sum_{c=1}^N \left\| \bar{\theta}_k - \left(\hat{\theta}_k^c - \frac{\eta}{p} \xi_{k-1}^c \right) + \frac{\eta}{p} \xi_\star^c \right\|^2 \right\} \\ &= \mathbf{1}_{\{1\}}(B_k) \left\{ \frac{1}{N} \sum_{c=1}^N \left\| \hat{\theta}_k^c - \frac{\eta}{p} (\xi_{k-1}^c - \xi_\star^c) - \theta_\star \right\|^2 - \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2 \right\}. \end{aligned}$$

We now expand the first term in the right-hand side of the previous equation. This gives

$$\frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \frac{\eta}{p}(\xi_{k-1}^c - \xi_\star^c) - \theta_\star\|^2 = \frac{1}{N} \sum_{c=1}^N \left\{ \|\hat{\theta}_k^c - \theta_\star\|^2 - \frac{2\eta}{p} \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle + \frac{\eta^2}{p^2} \|\xi_{k-1}^c - \xi_\star^c\|^2 \right\},$$

which yields

$$\begin{aligned} \mathbf{1}_{\{1\}}(B_k) \{\psi_k\} &= \mathbf{1}_{\{1\}}(B_k) \left\{ \|\bar{\theta}_k - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2 \right\} \\ &= \mathbf{1}_{\{1\}}(B_k) \left\{ \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \theta_\star\|^2 - \frac{2\eta}{p} \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2 \right\}. \end{aligned} \quad (37)$$

On the other hand, note that

$$\begin{aligned} \mathbf{1}_{\{0\}}(B_k) \{\psi_k\} &= \mathbf{1}_{\{0\}}(B_k) \left\{ \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2 \right\} \\ &= \mathbf{1}_{\{0\}}(B_k) \left\{ \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2 \right\}. \end{aligned} \quad (38)$$

By combining (38) and (37), we get

$$\begin{aligned} \psi_k &= \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2 \\ &= \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_k^c - \theta_\star\|^2 - 2\frac{\eta}{p} \mathbf{1}_{\{1\}}(B_k) \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2. \end{aligned} \quad (39)$$

Progress in local updates. We now bound the first term of the sum in (39). For $c \in [N]$, (36) gives

$$\begin{aligned} \|\hat{\theta}_k^c - \theta_\star\|^2 &= \|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))(\theta_{k-1}^c - \theta_\star) + \eta(\xi_{k-1}^c - \xi_\star^c) - \eta \omega^c(Z_k^c)\|^2 \\ &= \|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\} - \eta \omega^c(Z_k^c)\|^2 + \eta^2 \|\xi_{k-1}^c - \xi_\star^c\|^2 \\ &\quad + 2\eta \langle \xi_{k-1}^c - \xi_\star^c, (\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\} - \eta \omega^c(Z_k^c) \rangle \\ &= \underbrace{\|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\} - \eta \omega^c(Z_k^c)\|^2}_{T_1} + 2\eta \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle - \eta^2 \|\xi_{k-1}^c - \xi_\star^c\|^2. \end{aligned} \quad (40)$$

Define the σ -algebra $\mathcal{G}_{k-1} = \sigma(B_s, s \leq k-1, Z_s^c, s \leq k-1, c \in [N])$. We now bound the conditional expectation of T_1

$$\begin{aligned} \mathbb{E}^{\mathcal{G}_{k-1}} [T_1] &= \mathbb{E}^{\mathcal{G}_{k-1}} \left[\|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\} - \eta \omega^c(Z_k^c)\|^2 - 2\eta \langle (\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\}, \omega^c(Z_k^c) \rangle + \eta^2 \|\omega^c(Z_k^c)\|^2 \right] \\ &= \mathbb{E}^{\mathcal{G}_{k-1}} \left[\|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\}\|^2 + 2\eta^2 \langle \mathbf{A}^c(Z_k^c)\{\theta_k^c - \theta_\star\}, \omega^c(Z_k^c) \rangle + \eta^2 \|\omega^c(Z_k^c)\|^2 \right], \end{aligned}$$

where we used the fact that $\langle \mathbf{I}, \omega^c(Z_k^c) \rangle = 0$. Using Young's inequality for products, and Lemma C.5 with $\eta \leq \frac{1}{2L}$ and $u = \theta_k^c - \theta_\star$, we then obtain

$$\begin{aligned} \mathbb{E}^{\mathcal{G}_{k-1}} [T_1] &\leq \mathbb{E}^{\mathcal{G}_{k-1}} \left[\|(\mathbf{I} - \eta \mathbf{A}^c(Z_k^c))\{\theta_k^c - \theta_\star\}\|^2 + \eta^2 \|\mathbf{A}^c(Z_k^c)\{\theta_k^c - \theta_\star\}\|^2 + \eta^2 \|\omega^c(Z_k^c)\|^2 + \eta^2 \|\omega^c(Z_k^c)\|^2 \right] \\ &\leq (1 - \eta a) \|\theta_k^c - \theta_\star\|^2 - \eta \left(\frac{1}{L} - 2\eta \right) \mathbb{E}^{\mathcal{G}_{k-1}} \left[\|\mathbf{A}^c(Z_k^c)\{\theta_k^c - \theta_\star\}\|^2 \right] + 2\eta^2 \mathbb{E}^{\mathcal{G}_{k-1}} \left[\|\omega^c(Z_k^c)\|^2 \right]. \end{aligned} \quad (41)$$

Plugging (41) in (40) and using the assumption $\eta \leq \frac{1}{2L}$, we obtain

$$\mathbb{E}^{\mathcal{G}_{k-1}} \left[\|\hat{\theta}_k^c - \theta_\star\|^2 - 2\eta \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle \right] \leq (1 - \eta a) \|\theta_k^c - \theta_\star\|^2 - \eta^2 \|\xi_{k-1}^c - \xi_\star^c\|^2 + 2\eta^2 \text{Tr}(\Sigma_\varepsilon^c). \quad (42)$$

Bounding the Lyapunov function. Taking the conditional expectation of (39) and using (42) for $c = 1$ to N , we obtain the following bound on the Lyapunov function

$$\begin{aligned}\mathbb{E}^{\mathcal{G}^{k-1}}[\psi_k] &= \frac{1}{N} \sum_{c=1}^N \mathbb{E}^{\mathcal{G}^{k-1}} \left[\|\hat{\theta}_k^c - \theta_\star\|^2 - 2\eta \langle \xi_{k-1}^c - \xi_\star^c, \hat{\theta}_k^c - \theta_\star \rangle \right] + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2 \\ &\leq \frac{1}{N} \sum_{c=1}^N \left[(1 - \eta a) \|\theta_k^c - \theta_\star\|^2 - \eta^2 \|\xi_{k-1}^c - \xi_\star^c\|^2 + 2\eta^2 \text{Tr}(\Sigma_\varepsilon^c) \right] + \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2 \\ &= (1 - \eta a) \frac{1}{N} \sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + (1 - p^2) \frac{\eta^2}{p^2} \frac{1}{N} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2 + \frac{2\eta^2}{N} \sum_{c=1}^N \text{Tr}(\Sigma_\varepsilon^c),\end{aligned}$$

and the result of the Lemma follows from the Tower property. \square

Theorem B.2 (Convergence rate). *Assume A1 and A2(2). Then, for any $\eta \leq \frac{1}{2L}$ and $T > 0$, it holds*

$$\mathbb{E}[\psi_K] \leq (1 - \zeta)^K \left(\|\theta_0 - \theta_\star\|^2 + \frac{\eta^2}{p^2} \Delta_{\text{heter}} \right) + \frac{2\eta^2}{\zeta} \frac{1}{N} \sum_{c=1}^N \text{Tr}(\Sigma_\varepsilon^c),$$

where $\zeta = \min(\eta a, p^2)$.

Corollary B.3 (Iteration complexity). *Let $\epsilon > 0$. Set $\eta = \min\left(\frac{1}{2L}, \frac{\epsilon^2 a}{8\bar{\sigma}_\varepsilon}\right)$ and $p = \sqrt{\eta a}$ (so that $\zeta = \eta a$). Then, $\mathbb{E}[\psi_K] \leq \epsilon^2$ as long as the number of iterations is*

$$K \geq \max\left(\frac{2L}{a}, \frac{4\bar{\sigma}_\varepsilon}{\epsilon^2 a^2}\right) \log\left(\frac{\|\theta_0 - \theta_\star\|^2 + \min\left(\frac{1}{2aL}, \frac{\epsilon^2}{8\bar{\sigma}_\varepsilon}\right) \Delta_{\text{heter}}}{2\epsilon^2}\right),$$

which corresponds to an expected number of communication rounds

$$T \geq \max\left(\sqrt{\frac{2L}{a}}, \sqrt{\frac{4\bar{\sigma}_\varepsilon}{\epsilon^2 a^2}}\right) \log\left(\frac{\|\theta_0 - \theta_\star\|^2 + \min\left(\frac{1}{2aL}, \frac{\epsilon^2}{8\bar{\sigma}_\varepsilon}\right) \Delta_{\text{heter}}}{2\epsilon^2}\right).$$

Theorem B.4 (No linear speedup in the probabilistic communication setting with control variates). *The bounds obtained in Theorem B.2 are minimax optimal up to constants that are independent from the problem. Precisely, for every (p, η) there exists a FLSA problem such that*

$$\mathbb{E}[\psi_K] = (1 - \zeta)^K \left(\|\theta_0 - \theta_\star\|^2 + \frac{\eta^2}{p^2} \Delta_{\text{heter}} \right) + \frac{2\eta^2}{\zeta} \bar{\sigma}_\varepsilon,$$

where we have defined $\zeta = \min(2\eta a, p^2)$.

Proof. Define for all $c \in [N]$,

$$\bar{\mathbf{A}}^c = a\mathbf{I}, \quad \bar{\mathbf{b}}^c = b_c u,$$

where u is a vector whom all coordinates are equal to 1. We also consider the sequence of i.i.d random variables (Z_k^c) such that that for all $c \in [N]$ and $0 \leq t \leq T$, Z_k^c follows a Rademacher distribution. Moreover, we define

$$\mathbf{A}^c(Z_k^c) = \bar{\mathbf{A}}^c, \quad \mathbf{b}^c(Z_k^c) = \bar{\mathbf{b}}^c + Z_k^c u.$$

In particular this implies

$$\omega^c(z) = Z_k^c u.$$

Algorithm 4 SCAFFLSA-H 2: Stochastic Controlled FedLSA with deterministic communication

Input: $\eta > 0, \theta_0, \xi_0^c \in \mathbb{R}^d, T, N, H$

Set: $K = TH$

for $t = 1$ to T **do**

for $c = 1$ to N **do**

for $h = 1$ to H **do**

 Receive $Z_{t,h}^c$ and perform local update:

$$\hat{\theta}_{t,h}^c = \hat{\theta}_{t,h-1}^c - \eta(\mathbf{A}^c(Z_{t,h}^c)\hat{\theta}_{t,h-1}^c - \mathbf{b}^c(Z_{t,h}^c) - \xi_t^c)$$

end for

 Average local iterates: $\theta_{t+1}^c = \frac{1}{N} \sum_{c=1}^N \hat{\theta}_{t,H}^c$

 Update: $\xi_{t+1}^c = \xi_t^c + \frac{1}{\eta H}(\theta_{t+1}^c - \hat{\theta}_{t,H}^c)$

end for

end for

We follow the same proof of Lemma B.1 until the chain of equalities breaks. Thereby, we start from

$$\begin{aligned} \mathbb{E}[\psi_k] &= \mathbb{E}\left[\sum_{c=1}^N \|\theta_k^c - \theta_\star\|^2 + \frac{\eta^2}{p^2} \sum_{c=1}^N \|\xi_k^c - \xi_\star^c\|^2\right] \\ &= \mathbb{E}\left[\sum_{c=1}^N \|(I - \eta \mathbf{A}^c(Z_k^c))\{\theta_{k-1}^c - \theta_\star\} - \eta \omega^c(Z_k^c)\|^2 + (1-p^2) \frac{\eta^2}{p^2} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2\right] \\ &= \mathbb{E}\left[\sum_{c=1}^N \|(I - \eta \bar{\mathbf{A}}^c)\{\theta_{k-1}^c - \theta_\star\} - \eta \omega^c(Z_k^c)\|^2 + (1-p^2) \frac{\eta^2}{p^2} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2\right] \\ &= \mathbb{E}\left[\sum_{c=1}^N (1 - \eta a)^2 \|\theta_{k-1}^c - \theta_\star\|^2 + \eta^2 \|\omega^c(Z_k^c)\|^2 + (1-p^2) \frac{\eta^2}{p^2} \sum_{c=1}^N \|\xi_{k-1}^c - \xi_\star^c\|^2\right] \end{aligned}$$

where we used that $\mathbf{A}^c(Z_k^c) = \bar{\mathbf{A}}^c$. Unrolling the recursion gives the desired result. \square

B.2. Deterministic communication (Assumption H 2)

Consider the same algorithm as in the previous part, except that the variables B_k are chosen such that $B_k = 1$ if k is a multiple of H and $B_k = 0$ otherwise. This amounts to making blocks of local updates of constant size H . For clarity, we thus define t as the integer part of k/H and $h = k \bmod H$. Thus, for a given $t \geq 0$, each agent $c \in [N]$ performs H local updates as follows

$$\hat{\theta}_{t,h}^c = \hat{\theta}_{t,h-1}^c - \eta(\mathbf{A}^c(Z_{t,h})\hat{\theta}_{t,h-1}^c - \mathbf{b}^c(Z_{t,h}) - \xi_t^c),$$

where $\hat{\theta}_{t,0}^c = \theta_t$ and control variates ξ_t^c remain constant. At the end of each block, the global iterate and control variates are updated as

$$\theta_{t+1} = \frac{1}{N} \sum_{c=1}^N \theta_{t,H}^c, \quad \xi_{t+1}^c = \xi_t^c + \frac{1}{\eta H}(\theta_{t+1} - \hat{\theta}_{t,H}^c).$$

For clarity, we restate Algorithm 2 under this assumption in Algorithm 4. We Consider the Lyapunov function,

$$\psi_t = \|\theta_t - \theta_\star\|^2 + \frac{\eta^2 H^2}{N} \sum_{c=1}^N \|\xi_t^c - \xi_\star^c\|^2,$$

which is naturally defined as the error in θ_\star estimation on communication rounds, and the average error on the control variates.

Theorem B.5. Assume A1 and A2. Let η, H such that $\eta a H \leq 1, \eta H \leq 1, \eta \|\bar{\mathbf{A}}^c\| \leq 1$, and $\eta \leq \frac{a}{16H \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}}$, and set $\xi_0^c = 0$ for all $c \in [N]$. Then, the sequence $(\psi_t)_{t \in \mathbb{N}}$ satisfies, for all $t \geq 0$,

$$\mathbb{E}[\psi_t] \leq \left(1 - \frac{\eta a H}{4}\right)^t \mathbb{E}[\psi_0] + \frac{16\eta}{a} \bar{\sigma}_\varepsilon, \quad (43)$$

where $\psi_0 = \|\theta_0 - \theta_\star\|^2 + \frac{\eta^2 H^2}{N} \sum_{c=1}^N \|\bar{\mathbf{A}}^c(\theta_\star^c - \theta_\star)\|^2$.

Proof. Expression of local updates. Similarly to (36), we rewrite the local updates as, for $t \geq 0, h \geq 0, c \in [N]$,

$$\hat{\theta}_{t,h}^c - \theta_\star = (\mathbf{I} - \eta \mathbf{A}^c(Z_{t,h}^c))(\hat{\theta}_{t,h-1}^c - \theta_\star) + \eta(\xi_t^c - \xi_\star^c) - \eta \omega^c(Z_{t,h}^c),$$

where $\omega^c(z) = \tilde{\mathbf{A}}^c(z)\theta_\star - \tilde{\mathbf{b}}^c(z)$. Unrolling this recursion, we get, for each $c \in [N]$,

$$\theta_{t,H}^c - \theta_\star = \Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star) + \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)}(\xi_t^c - \xi_\star^c) - \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c). \quad (44)$$

Expression of the Lyapunov function. Since the sum control variates is $\sum_{t=1}^N \xi_t^c = \sum_{t=1}^N \xi_\star^c = 0$, we have $\theta_{t+1} = \frac{1}{N} \sum_{c=1}^N \theta_{t,H}^c = \frac{1}{N} \sum_{c=1}^N \theta_{t,H}^c - \eta H(\xi_t^c - \xi_\star^c)$. Applying Lemma C.4, we obtain

$$\begin{aligned} \|\theta_{t+1} - \theta_\star\|^2 &= \left\| \frac{1}{N} \sum_{c=1}^N \hat{\theta}_{t,H}^c - \theta_\star - \eta H(\xi_t^c - \xi_\star^c) \right\|^2 \\ &= \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_{t,H}^c - \theta_\star - \eta H(\xi_t^c - \xi_\star^c)\|^2 - \frac{1}{N} \sum_{c=1}^N \|\theta_{t+1} - \theta_{t,H}^c + \eta H(\xi_t^c - \xi_\star^c)\|^2 \\ &= \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_{t,H}^c - \theta_\star - \eta H(\xi_t^c - \xi_\star^c)\|^2 - \frac{\eta^2 H^2}{N} \sum_{c=1}^N \|\xi_{t+1}^c - \xi_\star^c\|^2, \end{aligned}$$

since $\xi_{t+1}^c = \xi_t^c + \frac{1}{\eta H}(\theta_{t+1} - \theta_{t,H}^c)$. Adding $\frac{\eta^2 H^2}{N} \sum_{c=1}^N \|\xi_{t+1}^c - \xi_\star^c\|^2$ on both sides and using (44), we obtain

$$\begin{aligned} \psi_{t+1} &= \frac{1}{N} \sum_{c=1}^N \|\hat{\theta}_{t,H}^c - \theta_\star - \eta H(\xi_t^c - \xi_\star^c)\|^2 \\ &= \frac{1}{N} \sum_{c=1}^N \|\Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star) + \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)}(\xi_t^c - \xi_\star^c) - \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) - \eta H(\xi_t^c - \xi_\star^c)\|^2 \\ &= \frac{1}{N} \sum_{c=1}^N \|\Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star) - \eta H(\mathbf{I} - \frac{1}{H} C_{\eta,H}^c)(\xi_t^c - \xi_\star^c) - \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2, \end{aligned}$$

where we defined $C_{\eta,H}^c = \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)}$. Expanding the norm gives

$$\begin{aligned} \psi_{t+1} &= \frac{1}{N} \sum_{c=1}^N \underbrace{\|\Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star) - \eta H(\mathbf{I} - \frac{1}{H} C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2}_{T_0^c} + \underbrace{\|\eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2}_{T_1^c} \\ &\quad - \underbrace{2\eta \langle \Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star), \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) \rangle}_{T_2^c} + \underbrace{2\eta H \langle (\mathbf{I} - \frac{1}{H} C_{\eta,H}^c)(\xi_t^c - \xi_\star^c), \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) \rangle}_{T_3^c}. \end{aligned} \quad (45)$$

In the following, we will use the filtration of all events up to step t , $\mathcal{F}_t := \sigma(Z_{s,h}^c, 0 \leq s \leq t, 0 \leq h \leq H, 1 \leq c \leq N)$.

Bounding T_0^c . Using Young's inequality with $\alpha_0 > 0$, and Assumption A2, we can bound

$$\begin{aligned} \mathbb{E}[T_0^c] &\leq \mathbb{E}[(1 + \alpha_0)\|\Gamma_{t,1:H}^{(c,\eta)}(\theta_t - \theta_\star)\|^2] + (1 + \alpha_0^{-1})\eta^2 H^2 \mathbb{E}[\|(I - \frac{1}{H}C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2] \\ &\leq (1 + \alpha_0)(1 - \eta a)^{2H} \mathbb{E}[\|\theta_t - \theta_\star\|^2] + (1 + \alpha_0^{-1})\eta^2 H^2 \mathbb{E}[\|(I - \frac{1}{H}C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2]. \end{aligned}$$

Using Lemma C.6 and Young's inequality again, and since $\eta\|\bar{\mathbf{A}}^c\| \leq 1$ we can bound

$$\begin{aligned} \mathbb{E}[\|(I - \frac{1}{H}C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2] &= \mathbb{E}[\mathbb{E}^{\mathcal{F}_{t-1}}[\|(I - \frac{1}{H}C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2]] \\ &\leq \mathbb{E}[\mathbb{E}^{\mathcal{F}_{t-1}}[\|I - \frac{1}{H}C_{\eta,H}^c\|^2]] \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2] \\ &\leq \left(\frac{1}{2}[\exp((H-1)\eta\|\bar{\mathbf{A}}^c\|) - 1]^2 + 2\eta^2(H-1)^2\|\Sigma_{\bar{\mathbf{A}}}^c\|\right) \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2] \\ &\leq \left(2\eta^2(H-1)^2\|\bar{\mathbf{A}}^c\|^2 + 2\eta^2(H-1)^2\|\Sigma_{\bar{\mathbf{A}}}^c\|\right) \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2]. \end{aligned}$$

We thus obtain:

$$\mathbb{E}[T_0^c] \leq (1 + \alpha_0)(1 - \eta a)^{2H} \mathbb{E}[\|\theta_t - \theta_\star\|^2] + 2(1 + \alpha_0^{-1})(\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|)\eta^4 H^4 \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2]. \quad (46)$$

Bounding T_1^c . Since the Z_{t,h^c} are all independent, we have:

$$\mathbb{E}[T_1^c] = \mathbb{E}[\|\eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2] = \eta^2 \sum_{h=1}^H \mathbb{E}[\|\Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2].$$

Using Assumption A2, we bound

$$\mathbb{E}[T_1^c] \leq \eta^2 \sum_{h=1}^H \mathbb{E}[\|\omega^c(Z_{t,h}^c)\|^2] = \eta^2 H \text{Tr}(\Sigma_{\bar{\mathbf{A}}}^c). \quad (47)$$

Bounding T_2^c . Since $(I - \eta\bar{\mathbf{A}}^c)^H$ is independent from all $Z_{t,h}$, we have for any $\alpha_2 > 0$,

$$\begin{aligned} |\mathbb{E}^{\mathcal{F}_{t-1}}[T_2^c]| &= |2\eta \mathbb{E}^{\mathcal{F}_{t-1}} \left[\langle (\Gamma_{t,1:H}^{(c,\eta)} - (I - \eta\bar{\mathbf{A}}^c)^H)(\theta_t - \theta_\star), \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) \rangle \right]| \\ &\leq \alpha_2 \mathbb{E}^{\mathcal{F}_{t-1}} \left[\|\Gamma_{t,1:H}^{(c,\eta)} - (I - \eta\bar{\mathbf{A}}^c)^H\|^2 \right] + \frac{\eta^2}{\alpha_2} \mathbb{E}^{\mathcal{F}_{t-1}} \left[\left\| \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) \right\|^2 \right]. \quad (48) \end{aligned}$$

To bound the expectation of the first term in (48), we bound the following conditional expectation for $u \in \mathbb{R}^d$, using Minkowski's inequality and Lemma C.1,

$$\begin{aligned} \mathbb{E}^{1/2}[\|(\Gamma_{t,1:H}^{(c,\eta)} - (I - \eta\bar{\mathbf{A}}^c)^H)u\|^2] &\leq \mathbb{E}^{1/2}[\|\eta \sum_{h=1}^H \Gamma_{t,1:h-1}^{(c,\eta)} (\mathbf{A}^c(Z_{t,h}) - \bar{\mathbf{A}}^c)(I - \eta\bar{\mathbf{A}}^c)^{H-h-1} u\|^2] \\ &\leq \eta \sum_{h=1}^H \mathbb{E}^{1/2}[\|\Gamma_{t,1:h-1}^{(c,\eta)} (\mathbf{A}^c(Z_{t,h}) - \bar{\mathbf{A}}^c)(I - \eta\bar{\mathbf{A}}^c)^{H-h-1} u\|^2]. \end{aligned}$$

From assumption A2, we then obtain

$$\begin{aligned} \mathbb{E}^{1/2}[\|(\Gamma_{t,1:H}^{(c,\eta)} - (I - \eta\bar{\mathbf{A}}^c)^H)u\|^2] &\leq \eta \sum_{h=1}^H (1 - \eta a)^{h-1} \mathbb{E}^{1/2}[\|(\mathbf{A}^c(Z_{t,h}) - \bar{\mathbf{A}}^c)(I - \eta\bar{\mathbf{A}}^c)^{H-h-1} u\|^2] \\ &\leq \eta \sum_{h=1}^H (1 - \eta a)^{h-1} \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \|(I - \eta\bar{\mathbf{A}}^c)^{H-h-1} u\|^2 \\ &\leq \eta \sum_{h=1}^H (1 - \eta a)^H \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \|u\|^2 \leq \eta H \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \|u\|. \quad (49) \end{aligned}$$

We then bound the second term of (48), using the independence of the $Z_{t,h}^c$,

$$\mathbb{E}[\|\sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2] = \sum_{h=1}^H \mathbb{E}[\|\Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\|^2] \leq \sum_{h=1}^H \mathbb{E}[\|\omega^c(Z_{t,h}^c)\|^2] = H \text{Tr}(\Sigma_\varepsilon^c), \quad (50)$$

which leads to the following inequality

$$\mathbb{E}[T_2^c] \leq \alpha_2 \eta^2 H^2 \|\Sigma_{\bar{\mathbf{A}}}^c\| \mathbb{E}[\|\theta_t - \theta_\star\|^2] + \frac{\eta^2 H}{\alpha_2} \text{Tr}(\Sigma_\varepsilon^c). \quad (51)$$

Bounding T_3^c . Let $\bar{C}_{\eta,H}^c = \sum_{h=1}^H (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^{H-h}$. Since \mathbf{I} and $\bar{C}_{\eta,H}^c$ are independent from all $Z_{t,h}$, we have, for any $\alpha_3 > 0$,

$$\begin{aligned} |\mathbb{E}^{\mathcal{F}_{t-1}}[T_3^c]| &= 2|\mathbb{E}^{\mathcal{F}_{t-1}}\left[\left\langle \eta(\bar{C}_{\eta,H}^c - C_{\eta,H}^c)(\xi_t^c - \xi_\star^c), \eta \sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c) \right\rangle\right]| \\ &\leq \alpha_3 \eta^2 \mathbb{E}^{\mathcal{F}_{t-1}}[\|(\bar{C}_{\eta,H}^c - C_{\eta,H}^c)(\xi_t^c - \xi_\star^c)\|^2] + \frac{\eta^2}{\alpha_3} \mathbb{E}^{\mathcal{F}_{t-1}}\left[\left\|\sum_{h=1}^H \Gamma_{t,h+1:H}^{(c,\eta)} \omega^c(Z_{t,h}^c)\right\|^2\right]. \end{aligned} \quad (52)$$

From (50), we can bound the expectation of the second term by $\frac{\eta^2 H}{\alpha_3} \text{Tr}(\Sigma_\varepsilon^c)$. We now proceed as in the derivation of (49) to bound the first term from (52). For any vector $u \in \mathbb{R}^d$, Minkowski's inequality and Lemma C.2 give

$$\begin{aligned} \mathbb{E}^{1/2}[\|(\bar{C}_{\eta,H}^c - C_{\eta,H}^c)u\|^2] &= \mathbb{E}^{1/2}\left[\left\|\sum_{h=1}^H (\Gamma_{t,h+1:H}^{(c,\eta)} - (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^{H-h})u\right\|^2\right] \\ &= \eta \mathbb{E}^{1/2}\left[\left\|\sum_{h=1}^H \sum_{h'=h+1}^H \Gamma_{t,h+1:h'-1}^{(c,\eta)} (\mathbf{A}^c(Z_{t,h'}) - \bar{\mathbf{A}}^c)(\mathbf{I} - \eta \bar{\mathbf{A}}^c)^{H-h'}u\right\|^2\right] \\ &= \eta \sum_{h=1}^H \sum_{h'=h+1}^H \mathbb{E}^{1/2}\left[\left\|\Gamma_{t,h+1:h'-1}^{(c,\eta)} (\mathbf{A}^c(Z_{t,h'}) - \bar{\mathbf{A}}^c)(\mathbf{I} - \eta \bar{\mathbf{A}}^c)^{H-h'}u\right\|^2\right]. \end{aligned}$$

Then, using assumption A2, we obtain

$$\mathbb{E}^{1/2}[\|(\bar{C}_{\eta,H}^c - C_{\eta,H}^c)u\|^2] \leq \eta \sum_{h=1}^H \sum_{h'=h+1}^H (1 - \eta a)^{H-h-1} \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \|u\| \leq \frac{\eta}{2} H(H-1) \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \|u\|.$$

We thus have the following bound

$$\mathbb{E}[T_3^c] \leq \frac{\eta^4}{4} \alpha_3 H^2 (H-1)^2 \|\Sigma_{\bar{\mathbf{A}}}^c\| \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2] + \frac{\eta^2 H}{\alpha_3} \text{Tr}(\Sigma_\varepsilon^c). \quad (53)$$

Bounding ψ_{t+1} . We can now bound ψ_{t+1} by plugging (46), (47), (51) and (53) in the expectation of (45)

$$\begin{aligned} \mathbb{E}[\psi_{t+1}] &= \frac{1}{N} \sum_{c=1}^N \mathbb{E}[T_0^c + T_1^c + T_2^c + T_3^c] \\ &\leq \left((1 + \alpha_0)(1 - \eta a)^{2H} + \alpha_2 \eta^2 H^2 \mathbb{E}_c[\|\Sigma_{\bar{\mathbf{A}}}^c\|] \right) \mathbb{E}[\|\theta_t - \theta_\star\|^2] \\ &\quad + \frac{1}{N} \sum_{c=1}^N \left(2(1 + \alpha_0^{-1})(\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|) + \frac{\alpha_3}{4} \|\Sigma_{\bar{\mathbf{A}}}^c\| \right) \eta^4 H^4 \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2] + \left(1 + \frac{1}{\alpha_2} + \frac{1}{\alpha_3}\right) \eta^2 H \bar{\sigma}_\varepsilon \\ &\leq \left((1 + \alpha_0)(1 - \eta a)^{2H} + \alpha_2 \eta^2 H^2 \mathbb{E}_c[\|\Sigma_{\bar{\mathbf{A}}}^c\|] \right) \mathbb{E}[\|\theta_t - \theta_\star\|^2] \\ &\quad + \left(2(1 + \alpha_0^{-1}) \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\} + \frac{\alpha_3}{4} \max_{c \in [N]} \{\|\Sigma_{\bar{\mathbf{A}}}^c\|\} \right) \eta^4 H^4 \frac{1}{N} \sum_{c=1}^N \mathbb{E}[\|\xi_t^c - \xi_\star^c\|^2] \\ &\quad + \left(1 + \frac{1}{\alpha_2} + \frac{1}{\alpha_3}\right) \eta^2 H \bar{\sigma}_\varepsilon. \end{aligned} \quad (54)$$

Now, we set $\alpha_0 = \frac{\eta a H}{2}$, $\alpha_2 = \frac{a}{4\eta H \mathbb{E}_c\{\|\Sigma_{\mathbf{A}}^c\|\}}$, $\alpha_3 = 1$. Since $\eta a H \leq 1$, $\eta H \leq 1$, and $\eta \leq \frac{a}{16H \max_{c \in [N]}\{\|\mathbf{A}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}}$ we have

$$\begin{aligned} (1 + \alpha_0)(1 - \eta a)^{2H} + \alpha_2 \eta^2 H^2 \mathbb{E}_c\{\|\Sigma_{\mathbf{A}}^c\|\} &\leq (1 + \frac{\eta a H}{2})(1 - \eta a H) + \frac{\eta a H}{4} \leq 1 - \frac{\eta a H}{4}, \\ \eta^2 H^2 (2(1 + \alpha_0^{-1}) \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\} + \frac{\alpha_3}{4} \max_{c \in [N]}\{\|\Sigma_{\mathbf{A}}^c\|\}) &\leq \eta^2 H^2 (3 + \frac{4}{\eta a H}) \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\} \leq \frac{1}{2}, \\ 1 + \frac{1}{\alpha_2} + \frac{1}{\alpha_3} &\leq 2 + \frac{2\eta H \mathbb{E}_c\{\|\Sigma_{\mathbf{A}}^c\|\}}{a} \leq 4. \end{aligned}$$

Using these inequalities, we can simplify (54) as

$$\begin{aligned} \mathbb{E}[\psi_{t+1}] &\leq \left(1 - \frac{\eta a H}{4}\right) \mathbb{E}[\|\theta_t - \theta_\star\|^2] + \frac{1}{2} \eta^2 H^2 \mathbb{E}[\|\xi - \xi_\star\|^2] + 4\eta^2 H \bar{\sigma}_\epsilon \\ &\leq \left(1 - \min\left(\frac{1}{2}, \frac{\eta a H}{4}\right)\right) \mathbb{E}[\psi_t] + 4\eta^2 H \bar{\sigma}_\epsilon \\ &= \left(1 - \frac{\eta a H}{4}\right) \mathbb{E}[\psi_t] + 4\eta^2 H \bar{\sigma}_\epsilon, \end{aligned}$$

where the last inequality comes from $\eta a H \leq 1$. Applying this inequality iteratively gives

$$\mathbb{E}[\psi_T] \leq \left(1 - \frac{\eta a H}{4}\right)^T \mathbb{E}[\psi_0] + \frac{16\eta}{a} \bar{\sigma}_\epsilon,$$

which is the result of the theorem. \square

Corollary B.6. Let $\epsilon > 0$. Assume $\epsilon^2 \leq \frac{32\bar{\sigma}_\epsilon}{a \max_{c \in [N]}\{\|\mathbf{A}^c\|\}}$, $a \leq 1$ and $a^2 \leq 16 \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}$. In order to achieve $\|\theta_T - \theta_\star\|^2 \leq \epsilon^2$, the required number of communication is

$$T = \frac{64 \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}}{a^2} \log \left(\frac{2\|\theta_0 - \theta_\star\|^2 + \frac{a^2}{256 \max_{c \in [N]}\{\|\mathbf{A}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}^2} \Delta_{\text{heter}}}{\epsilon^2} \right),$$

using the step size

$$\eta = \min \left(\eta_\infty, \frac{a\epsilon^2}{32\bar{\sigma}_\epsilon} \right),$$

and the number of local iterations

$$H = \frac{a \max \left(\frac{1}{16\eta_\infty}, \frac{2\bar{\sigma}_\epsilon}{a\epsilon^2} \right)}{\max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}}.$$

Proof. First, we require that the second term of the right-hand side of (43) is bounded by $\frac{\epsilon^2}{2}$, that is $\frac{16\eta}{a} \bar{\sigma}_\epsilon \leq \frac{\epsilon^2}{2}$. This requires that

$$\eta \leq \min \left(\eta_\infty, \frac{a\epsilon^2}{32\bar{\sigma}_\epsilon} \right),$$

which also ensures that $\eta \leq \frac{1}{\|\bar{\mathbf{A}}^c\|}$ for all $c \in [N]$ since $\epsilon^2 \leq \frac{32\bar{\sigma}_\epsilon}{a \max_{c \in [N]}\{\|\mathbf{A}^c\|\}}$. We thus set $\eta = \min \left(\eta_\infty, \frac{a\epsilon^2}{32\bar{\sigma}_\epsilon} \right)$, and look at the largest H possible. The conditions $\eta a H \leq 1$, $\eta H \leq 1$, and $\eta \leq \frac{a}{16H \max_{c \in [N]}\{\|\mathbf{A}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}}$ give that

$$\begin{aligned} H &\leq \min \left(\frac{1}{\eta}, \frac{1}{a\eta}, \frac{a}{16\eta \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}} \right) \\ &= \frac{1}{\eta} \min \left(1, \frac{1}{a}, \frac{a}{16 \max_{c \in [N]}\{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\mathbf{A}}^c\|\}} \right). \end{aligned}$$

Since $a \leq 1$ and $a^2 \leq 16 \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}$, the smallest term of the minimum is the third one and we can take

$$H = \frac{a}{16\eta \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}} = \frac{a \max\left(\frac{1}{16\eta_\infty}, \frac{2\bar{\sigma}_\epsilon}{a\epsilon^2}\right)}{\max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}}.$$

To obtain $\mathbb{E}[\psi_T] \leq \epsilon^2$, we need to take T so that $(1 - \frac{\eta a H}{4})^T \psi_0 \leq \frac{\epsilon^2}{2}$, which gives

$$T \geq \frac{4}{\eta a H} \log\left(\frac{2\psi_0}{\epsilon^2}\right) = \frac{64 \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}}{a^2} \log\left(\frac{2\|\theta_0 - \theta_\star\|^2 + \frac{a^2}{256 \max_{c \in [N]} \{\|\bar{\mathbf{A}}^c\|^2 + \|\Sigma_{\bar{\mathbf{A}}}^c\|\}^2} \Delta_{\text{heter}}}{\epsilon^2}\right),$$

and the result follows. \square

C. Technical proofs

Lemma C.1. *For any matrix-valued sequences $(U_n)_{n \in \mathbb{N}}$, $(V_n)_{n \in \mathbb{N}}$ and for any $M \in \mathbb{N}$, it holds that:*

$$\prod_{k=1}^M U_k - \prod_{k=1}^M V_k = \sum_{k=1}^M \left\{ \prod_{j=1}^{k-1} U_j \right\} (U_k - V_k) \left\{ \prod_{j=k+1}^M V_j \right\}.$$

Lemma C.2 (Stability of the deterministic product). *Assume A2. Then, for any $u \in \mathbb{R}^d$ and $h \in \mathbb{N}$,*

$$\|(\mathbf{I} - \eta \bar{\mathbf{A}}^c)^h u\| \leq (1 - \eta a)^h \|u\|.$$

Proof. Since $(Z_{t,h}^c)_{1 \leq h \leq H}$ are i.i.d, we get

$$\mathbb{E}[\Gamma_{t,1:h}^{(c,\eta)} u] = \mathbb{E}[\prod_{l=1}^h (\mathbf{I} - \eta \mathbf{A}(Z_{t,l}^c)) u] = \prod_{l=1}^h \mathbb{E}[\mathbf{I} - \eta \mathbf{A}(Z_{t,l}^c)] u = (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^h u.$$

The proof then follows from the elementary inequality: for any square-integrable random vector U , $\|\mathbb{E}[U]\| \leq (\mathbb{E}[\|U\|^2])^{1/2}$. \square

Lemma C.3. *Recall that $\bar{\rho}_H = \frac{1}{N} \sum_{c=1}^N (\mathbf{I} - (\mathbf{I} - \eta \bar{\mathbf{A}}^c)^H) \{\theta_\star^c - \theta_\star\}$, it satisfies*

$$\|\bar{\rho}_H\| \leq \frac{\eta^2 H^2}{N} \sum_{c=1}^N \exp(\eta H \|\bar{\mathbf{A}}^c\|) \|\theta_\star^c - \theta_\star\| \quad (55)$$

Proof. Using the identity,

$$1 - (1 - u)^H = H u - u^2 \sum_{k=0}^{H-2} (-1)^k \binom{H}{k+2} u^k$$

and the inequality $\binom{H}{k+2} \leq \binom{H-2}{k} H^2$, we get that

$$\left| \sum_{k=0}^{H-2} (-1)^k \binom{H}{k+2} u^k \right| \leq \frac{H^2}{2} \sum_{k=0}^{H-2} \binom{H-2}{k} |u|^k \leq \frac{H^2}{2} \exp((H-2)|u|)$$

Using (15), we finally get (55). \square

Lemma C.4. *Let $(x_i)_{i=1}^N$, and $(y_i)_{i=1}^N$ be N vectors of \mathbb{R}^d . Denote $\bar{x}_N = (1/N) \sum_{i=1}^N x_i$ and $\bar{y}_N = (1/N) \sum_{i=1}^N y_i$. Then,*

$$N \|\bar{x}_N - \bar{y}_N\|^2 = \sum_{i=1}^N \|x_i - y_i\|^2 - \sum_{i=1}^N \|x_i - \bar{x}_N - (y_i - \bar{y}_N)\|^2$$

Proof. Define $\mathbf{x} = [x_1^\top, \dots, x_N^\top]^\top$ and $\mathbf{y} = [y_1^\top, \dots, y_N^\top]^\top \in \mathbb{R}^{Nd}$. Define by \mathbf{P} the orthogonal projector on

$$\mathcal{E} = \{\mathbf{x} \in \mathbb{R}^{Nd} : \mathbf{x} = [x^\top, \dots, x^\top]^\top, x \in \mathbb{R}^d\}.$$

We show that $\mathbf{P}\mathbf{x} = [\bar{x}_N^\top, \dots, \bar{x}_N^\top]^\top$. Note indeed that for any $\mathbf{z} = [z^\top, \dots, z^\top]^\top \in \mathcal{E}$, we get (with a slight abuse of notations, $\langle \cdot, \cdot \rangle$ denotes the scalar product in \mathbb{R}^{Nd} and \mathbb{R}^d)

$$\langle \mathbf{x} - \mathbf{P}\mathbf{x}, \mathbf{z} \rangle = \sum_{i=1}^N \{\langle x_i, z \rangle - \langle \bar{x}_N, z \rangle\} = 0.$$

The proof follows from Pythagoras identity which shows that

$$\|\mathbf{P}\mathbf{x} - \mathbf{P}\mathbf{y}\|^2 = \|\mathbf{x} - \mathbf{y}\|^2 - \|(\mathbf{x} - \mathbf{P}\mathbf{x}) - (\mathbf{y} - \mathbf{P}\mathbf{y})\|^2$$

□

Lemma C.5. *Assume A3. Let Z be a random variable taking values in a state space (Z, \mathcal{Z}) with distribution π_c . Set $\eta \geq 0$, then for any vector $u \in \mathbb{R}^d$, we have*

$$\mathbb{E}[\|(I - \eta \mathbf{A}^c(Z))u\|^2] \leq (1 - \eta a)\|u\|^2 - \eta\left(\frac{1}{L} - \eta\right)\mathbb{E}[\|\mathbf{A}^c(Z)u\|^2].$$

Proof. First, remark that

$$\begin{aligned} \|(I - \eta \mathbf{A}^c(Z))u\|^2 &= u^\top (I - \eta \mathbf{A}^c(Z))^\top (I - \eta \mathbf{A}^c(Z))u \\ &= u^\top \left(I - 2\eta \left(\frac{1}{2}(\mathbf{A}^c(Z) + \mathbf{A}^c(Z)^\top) \right) + \eta^2 \mathbf{A}^c(Z)^\top \mathbf{A}^c(Z) \right) u. \end{aligned}$$

Since we have $\mathbb{E}[\frac{1}{2}(\mathbf{A}^c(Z) + \mathbf{A}^c(Z)^\top)] \succcurlyeq aI$ and $\mathbb{E}[\frac{1}{2}(\mathbf{A}^c(Z) + \mathbf{A}^c(Z)^\top)] \succcurlyeq \frac{1}{L}\mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)]$, we obtain

$$\begin{aligned} \mathbb{E}[\|(I - \eta \mathbf{A}^c(Z))u\|^2] &= u^\top u - 2\eta u^\top \mathbb{E}[\frac{1}{2}(\mathbf{A}^c(Z) + \mathbf{A}^c(Z)^\top)]u + \eta^2 u^\top \mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)]u \\ &\leq \|u\|^2 - \eta a \|u\|^2 - \frac{\eta}{L} u^\top \mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)]u + \eta^2 u^\top \mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)]u \\ &= (1 - \eta a)\|u\|^2 - \eta\left(\frac{1}{L} - \eta\right)u^\top \mathbb{E}[\mathbf{A}^c(Z)^\top \mathbf{A}^c(Z)]u, \end{aligned}$$

which gives the result. □

Lemma C.6. *Let $H > 0$, $\eta > 0$ such that $0 \leq \eta \|\bar{\mathbf{A}}^c\| \leq 1$ for all $c \in [N]$. Then it holds that*

$$\mathbb{E}^{1/2}[\|I - C_{\eta, H}^{(t, c)}\|^2] \leq \frac{1}{2}[\exp((H-1)\eta \|\bar{\mathbf{A}}^c\|) - 1] + \eta(H-1)\|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2}.$$

Proof. Using Minkowski's inequality, we get

$$\begin{aligned} \mathbb{E}[\|I - C_{\eta, H}^{(t, c)}\|^2] &= \mathbb{E}[\|I - \frac{1}{H} \sum_{h=1}^H \Gamma_{t, h+1: H}^{(c, \eta)}\|^2] \\ &\leq \|I - \frac{1}{H} \sum_{h=1}^H (I - \eta \bar{\mathbf{A}}^c)^{h-1}\| + \frac{1}{H} \mathbb{E}^{1/2}[\|\sum_{h=1}^H (I - \eta \bar{\mathbf{A}}^c)^{H+1-h} - \Gamma_{t, h+1: H}^{(c, \eta)}\|^2]. \end{aligned} \tag{56}$$

Let's now establish an upper bound for the first term in the previous inequality

$$\begin{aligned} \|I - \frac{1}{H} \sum_{h=1}^H (I - \eta \bar{\mathbf{A}}^c)^{h-1}\| &= \|I - \frac{(\bar{\mathbf{A}}^c)^{-1}}{\eta H} (I - (I - \eta \bar{\mathbf{A}}^c)^H)\| \\ &= \|I - \frac{(\bar{\mathbf{A}}^c)^{-1}}{\eta H} (I - \sum_{h=0}^H \binom{H}{h} (-\eta \bar{\mathbf{A}}^c)^h)\| \\ &= \|I - \frac{(\bar{\mathbf{A}}^c)^{-1}}{\eta H} \sum_{h=1}^H \binom{H}{h} (-\eta \bar{\mathbf{A}}^c)^h\| \\ &= \|I - \frac{1}{H} \sum_{h=0}^{H-1} \binom{H}{h+1} (-\eta \bar{\mathbf{A}}^c)^h\| \\ &= \|\frac{1}{H} \sum_{h=1}^{H-1} \binom{H}{h+1} (-\eta \bar{\mathbf{A}}^c)^h\|. \end{aligned}$$

Using triangle inequality, we obtain

$$\|I - \frac{1}{H} C_{\eta, H}^c\| \leq \frac{1}{H} \sum_{h=1}^{H-1} \binom{H}{h+1} \|\eta \bar{\mathbf{A}}^c\|^h.$$

Furthermore

$$\frac{1}{H} \sum_{h=1}^{H-1} \binom{H}{h+1} \|\eta \bar{\mathbf{A}}^c\|^h \leq \frac{1}{2} \sum_{h=1}^{H-1} \binom{H-1}{h} \|\eta \bar{\mathbf{A}}^c\|^h \leq \frac{1}{2} [\exp((H-1)\eta \|\bar{\mathbf{A}}^c\|) - 1].$$

Applying Minkowski's inequality and Lemma C.1 to the second term on the right side of the inequality (56) yields

$$\begin{aligned} & \frac{1}{H} \mathbb{E}^{1/2} [\|\sum_{h=1}^H (I - \eta \bar{\mathbf{A}}^c)^{H+1-h} - \Gamma_{t, h+1: H}^{(c, \eta)}\|^2] \\ & \leq \frac{1}{H} \sum_{h=1}^H \mathbb{E}^{1/2} [\|(I - \eta \bar{\mathbf{A}}^c)^{H+1-h} - \Gamma_{t, h+1: H}^{(c, \eta)}\|^2] \\ & = \frac{\eta}{H} \sum_{h=1}^H \mathbb{E}^{1/2} [\|\sum_{k=h+1}^H (I - \eta \bar{\mathbf{A}}^c)^{k-1-(h+1)} (\mathbf{A}^c(Z_{t, k}^c) - \bar{\mathbf{A}}^c) \Gamma_{t, k+1: H}^{(c, \eta)}\|^2] \\ & \leq \frac{\eta}{H} \sum_{h=1}^H \sum_{k=h+1}^H \mathbb{E}^{1/2} [\|(I - \eta \bar{\mathbf{A}}^c)^{k-1-(h+1)} (\mathbf{A}^c(Z_{t, k}^c) - \bar{\mathbf{A}}^c) \Gamma_{t, k+1: H}^{(c, \eta)}\|^2] \\ & \leq \frac{\eta}{H} \sum_{h=1}^H (H-h)(1-\eta a)^{H-h-1} \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2} \\ & \leq \eta(H-1) \|\Sigma_{\bar{\mathbf{A}}}^c\|^{1/2}. \end{aligned}$$

□

D. TD learning as a federated LSA problem

In this section we specify TD(0) as a particular instance of the LSA algorithm. In the setting of linear functional approximation the problem of estimating $V^\pi(s)$ reduces to the problem of estimating $\theta_* \in \mathbb{R}^d$, which can be done via the LSA procedure. For the agent $c \in [N]$ the k -th step randomness is given by the tuple $Z_k^c = (S_k^c, A_k^c, S_{k+1}^c)$. With slight abuse of notation, we write $\mathbf{A}_{t, h}^c$ instead of $\mathbf{A}(Z_{t, h}^c)$, and $\mathbf{b}_{t, h}^c$ instead of $\mathbf{b}(Z_{t, h}^c)$. Then the corresponding LSA update equation with constant step size η can be written as

$$\theta_{t, h}^c = \theta_{t, h-1}^c - \eta (\mathbf{A}_{t, h}^c \theta_{t, h-1}^c - \mathbf{b}_{t, h}^c),$$

where $\mathbf{A}_{t, h}^c$ and $\mathbf{b}_{t, h}^c$ are given by

$$\begin{aligned} \mathbf{A}_{t, h}^c &= \phi(S_{t, h}^c) \{\phi(S_{t, h}^c) - \gamma \phi(S_{t, h+1}^c)\}^\top, \\ \mathbf{b}_{t, h}^c &= \phi(S_{t, h}^c) r^c(S_{t, h}^c, A_{t, h}^c). \end{aligned} \tag{57}$$

Respective specialisation of FedLSA algorithm is stated in Algorithm 5.

The corresponding local agent's system writes as $\bar{\mathbf{A}}^c \theta_*^c = \bar{\mathbf{b}}^c$, where we have, respectively,

$$\begin{aligned} \bar{\mathbf{A}}^c &= \mathbb{E}_{s \sim \mu^c, s' \sim P^\pi(\cdot|s)} [\phi(s) \{\phi(s) - \gamma \phi(s')\}^\top] \\ \bar{\mathbf{b}}^c &= \mathbb{E}_{s \sim \mu^c, a \sim \pi(\cdot|s)} [\phi(s) r^c(s, a)]. \end{aligned}$$

The authors of (Wang et al., 2023) study the corresponding virtual MDP dynamics with $\tilde{\mathbb{P}} = N^{-1} \sum_{c=1}^N \mathbb{P}_{\text{MDP}}^c$, $\tilde{r} = N^{-1} \sum_{c=1}^N r^c$. Next, introducing the invariant distribution of the kernel $\tilde{\mu}$ of the averaged state kernel

$$\tilde{\mathbb{P}}_\pi(B|s) = N^{-1} \sum_{c=1}^N \int_{\mathcal{A}} \mathbb{P}_{\text{MDP}}^c(B|s, a) \pi(da|s),$$

we have $\tilde{\theta}$ as an optimal parameter corresponding to the system $\tilde{\mathbf{A}} \tilde{\theta} = \tilde{\mathbf{b}}$. Here

$$\begin{aligned} \tilde{\mathbf{A}} &= \mathbb{E}_{s \sim \tilde{\mu}, s' \sim \tilde{\mathbb{P}}_\pi(\cdot|s)} [\phi(s) \{\phi(s) - \gamma \phi(s')\}^\top] \\ \tilde{\mathbf{b}} &= \mathbb{E}_{s \sim \tilde{\mu}, a \sim \pi(\cdot|s)} [\phi(s) \tilde{r}(s, a)]. \end{aligned}$$

Algorithm 5 FedLSA applied to TD(0) setting with linear functional approximation

Input: $\eta > 0, \theta_0 \in \mathbb{R}^d, T, N, H > 0$

for $t = 0$ to $T - 1$ **do**

 Initialize $\theta_{t,0} = \theta_t$

for $c = 1$ to N **do**

for $h = 1$ to H **do**

 Receive tuple $(S_{t,h}^c, A_{t,h}^c, S_{t,h+1}^c)$ following **TD 1** and perform local update:

$$\theta_{t,h}^c = \theta_{t,h-1}^c - \eta(\mathbf{A}_{t,h}^c \theta_{t,h-1}^c - \mathbf{b}_{t,h}^c),$$

 where $\mathbf{A}_{t,h}^c$ and $\mathbf{b}_{t,h}^c$ are given in (57)

end for

end for

Average: $\theta_{t+1} = \frac{1}{N} \sum_{c=1}^N \theta_{t,H}^c$

(58)

end for

D.1. Proof of Lemma 3.1.

The proof below closely follows (Patil et al., 2023) (Lemma 7) and (Samsonov et al., 2023) (Lemma 1). Indeed, with **TD 2** and (12) we get

$$\|\mathbf{A}_1^c\| \leq (1 + \gamma)$$

almost surely, which implies $\|\bar{\mathbf{A}}^c\| \leq 1 + \gamma$ for any $c \in [N]$. This implies, using the definition of $\Sigma_{\bar{\mathbf{A}}}^c$, that

$$\|\Sigma_{\bar{\mathbf{A}}}^c\| = \|\mathbb{E}[\{\mathbf{A}_1^c\}^\top \mathbf{A}_1^c] - \{\bar{\mathbf{A}}^c\}^\top \bar{\mathbf{A}}^c\| \leq 2(1 + \gamma)^2,$$

and the bound (9) follows. Next we observe that

$$\begin{aligned} \text{Tr}(\Sigma_\varepsilon^c) &= \mathbb{E}[\|(\mathbf{A}_1^c - \bar{\mathbf{A}}^c)\theta_\star^c - (\mathbf{b}_1^c - \bar{\mathbf{b}}^c)\|^2] \leq 2\{\theta_\star^c\}^\top \mathbb{E}[\{\mathbf{A}_1^c\}^\top \mathbf{A}_1^c] \theta_\star^c + 2\mathbb{E}[(r^s(S_0^s, A_0^c))^2 \text{Tr}(\varphi(S_0^c)\varphi^\top(S_0^c))] \\ &\leq 2(1 + \gamma)^2 \{\theta_\star^c\}^\top \Sigma_\varphi[c] \theta_\star^c + 2 \leq 2(1 + \gamma)^2 (\|\theta_\star^c\|^2 + 1), \end{aligned}$$

where the latter inequality follows from **TD 2**, and thus (10) holds. In order to check the last equation (11), we note first that the bound for a and η_∞ readily follows from the ones presented in (Patil et al., 2023)[Lemma 5] and (Patil et al., 2023)[Lemma 7]. To check assumption A3, note first that, with $s \sim \mu^c, s' \sim P^\pi(\cdot|s)$, we have

$$\begin{aligned} \mathbf{A}^c + \{\mathbf{A}^c\}^\top &= \varphi(s)\{\varphi(s) - \gamma\varphi(s')\}^\top + \{\varphi(s) - \gamma\varphi(s')\}\varphi(s)^\top = 2\varphi(s)\varphi(s)^\top - \gamma\{\varphi(s)\varphi(s')^\top + \varphi(s')\varphi(s)^\top\} \\ &\preceq (2 + \gamma)\varphi(s)\varphi(s)^\top + \gamma\varphi(s')\varphi(s')^\top, \end{aligned}$$

where we additionally used that

$$-(uu^\top + vv^\top) \preceq uv^\top + vu^\top \preceq (uu^\top + vv^\top)$$

for any $u, v \in \mathbb{R}^d$. Thus, we get that

$$\mathbb{E}[\mathbf{A}^c + \{\mathbf{A}^c\}^\top] \preceq 2(1 + \gamma)\Sigma_\varphi^c.$$

The rest of the proof follows from the fact that

$$\mathbb{E}[\{\mathbf{A}_1^c\}^\top \mathbf{A}_1^c] \succeq \{\bar{\mathbf{A}}^c\}^\top \bar{\mathbf{A}}^c \succeq (1 - \gamma)^2 \lambda_{\min} \Sigma_\varphi^c,$$

which is proven e.g. in (Li et al., 2023a) (Lemma 5) or (Samsonov et al., 2023) (Lemma 7).