
Federated Reinforcement Learning: Linear Speedup Under Markovian Sampling

Sajad Khodadadian¹ Pranay Sharma² Gauri Joshi² Siva Theja Maguluri¹

Abstract

Since reinforcement learning algorithms are notoriously data-intensive, the task of sampling observations from the environment is usually split across multiple agents. However, transferring these observations from the agents to a central location can be prohibitively expensive in terms of the communication cost, and it can also compromise the privacy of each agent’s local behavior policy. In this paper, we consider a federated reinforcement learning framework where multiple agents collaboratively learn a global model, without sharing their individual data and policies. Each agent maintains a local copy of the model and updates it using locally sampled data. Although having N agents enables the sampling of N times more data, it is not clear if it leads to proportional convergence speedup. We propose federated versions of on-policy TD, off-policy TD and Q -learning, and analyze their convergence. For all these algorithms, to the best of our knowledge, we are the first to consider Markovian noise and multiple local updates, and prove a *linear convergence speedup with respect to the number of agents*. To obtain these results, we show that federated TD and Q -learning are special cases of a general framework for federated stochastic approximation with Markovian noise, and we leverage this framework to provide a unified convergence analysis that applies to all the algorithms.

1. Introduction

Reinforcement Learning (RL) is an online sequential decision-making paradigm that is typically modeled as a Markov Decision Process (MDP) (Sutton & Barto, 2018). In an RL task, the agent aims to learn the optimal policy of the MDP that maximizes long-term reward, without knowledge of its parameters. The agent performs this task by repeatedly interacting with the environment according to a behavior policy, which in turn provides data samples that can be used to improve the policy. This MDP-based RL framework has a vast array of applications including self-driving cars (Yurtsever et al., 2020), robotic systems (Kober et al., 2013), games (Silver et al., 2016), UAV-based surveillance (Yun et al., 2022), and Internet of Things (IoT) (Lim et al., 2020).

Due to the high-dimensional state and action spaces that are typical in these applications, RL algorithms are extremely data hungry (Duan et al., 2016; Kalashnikov et al., 2018; Akkaya et al., 2019), and training RL models with limited data can result in low accuracy and high output variance (Islam et al., 2017; Xu et al., 2021). However, generating massive amounts of training data sequentially can be extremely time consuming (Nair et al., 2015). Hence, many practical implementations of RL algorithms from Atari domain to Cyber-Physical Systems rely on parallel sampling of the data from the environment using multiple agents (Mnih et al., 2016; Espeholt et al., 2018; Chen et al., 2021a; Xu et al., 2021). It was empirically shown in (Mnih et al., 2016) that the federated version of these algorithms yields faster training time and improved accuracy. A naive approach would be to transfer all the agents’ locally collected data to a central server that uses it for training. However, in applications such as IoT (Chen & Giannakis, 2018), autonomous driving (Shalev-Shwartz et al., 2016) and robotics (Kalashnikov et al., 2018), communicating high-dimensional data over low bandwidth network link can be prohibitively slow. Moreover, sharing individual data of the agents with the server might also be undesirable due to privacy concerns (Yang et al., 2019; Mothukuri et al., 2021).

Federated Learning (FL) (Kairouz et al., 2019) is an emerging distributed learning framework, where multiple agents seek to collaboratively train a shared model, while complying with the privacy and data confidentiality requirements

¹H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 30332, USA
²Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213, USA,. Correspondence to: Sajad Khodadadian <skhodadadian3@gatech.edu>.

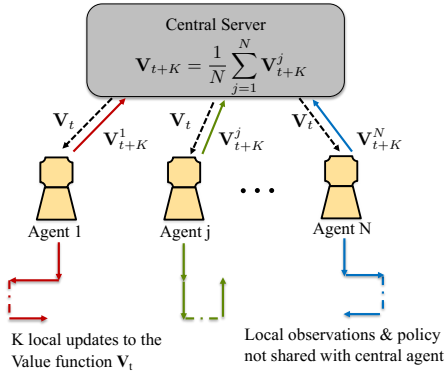


Figure 1. Schematic representation of FRL where N agents follow a Markovian trajectory and synchronize their parameter every K time steps.

(Qi et al., 2021; Yang et al., 2019). The key idea is that the agents collect data, use on-device computation capabilities to locally train the model, and only share the model updates with the central server. Not sharing data reduces communication cost and also alleviates privacy concerns.

Recently, there is a growing interest in employing FL for RL algorithms (also known as FRL) (Nadiger et al., 2019; Liu et al., 2019; Ren et al., 2019; Xu et al., 2021; Zhang et al., 2022). Unlike standard supervised learning where data is collected before training begins, in FRL, each agent collects data by following its own Markovian trajectory, while simultaneously updating the model parameters.

To ensure convergence, after every K time steps, the agents communicate with the central server to synchronize their parameters (see Figure 1). Intuitively, using more agents and a higher synchronization frequency should improve the convergence of training algorithm. However, the following questions remain to be concretely answered:

1. With N agents, do we get an N -fold (linear) speedup in the convergence of FRL algorithms?
2. How does the convergence speed and the final error scale with synchronization frequency K ?

While these questions are well-studied (Wang & Joshi, 2021; Stich, 2018; Qu et al., 2020; Li et al., 2019) in federated supervised learning, only a few works (Wai, 2020; Shen et al., 2020) have attempted to answer them in the context of FRL. However, none of them have established the convergence analysis of FRL algorithms by considering Markovian local trajectories and multiple local updates (see Table 1).

In this paper, we tackle this challenging open problem and answer both the questions listed above. We propose communication-efficient federated versions of on-policy TD,

off-policy TD, and Q -learning algorithms. In addition, we are the first to establish the convergence bounds for these algorithms in the realistic Markovian setting, showing a linear speedup in the number of agents. Previous works (Liu & Olshesky, 2021; Shen et al., 2020) on distributed RL have only shown such a speedup by assuming i.i.d. noise. Moreover, based on experiments, (Shen et al., 2020) conjectures that linear speedup may be possible under the realistic Markov noise setting, which we establish analytically. The main contributions and organization of the paper are summarized below.

- In the on-policy setting, in Section 4 we propose and analyze federated TD-learning with linear function approximation, where the agents’ goal is to evaluate a common policy using on-policy samples collected in parallel from their environments. The agents only share the updated value function (not data) with the central server, thus saving communication cost. We prove a linear convergence speedup with the number of agents and also characterize the impact of communication frequency on the convergence.
- In the off-policy setting, in Section 5 we propose and analyze the federated off-policy TD-learning and federated Q -learning algorithms. Again, we establish a linear speedup in their convergence with respect to the number of agents and characterize the impact of synchronization frequency on the convergence. Since every agent samples data using a private policy and only communicates the updated value or Q -function, off-policy FRL helps keep both the data as well as the policy private.
- In Section 6, we propose a general Federated Stochastic Approximation framework with Markovian noise (FedSAM) which subsumes both federated TD-learning and federated Q -learning algorithms proposed above. Considering Markovian sampling noise poses a significant challenge in the analysis of this algorithm. The convergence result for FedSAM serves as a workhorse that enables us to analyze both federated TD-learning and federated Q -learning. We characterize the convergence of FedSAM with a refined analysis of general stochastic approximation algorithms, fundamentally improving upon prior work.

2. Related Work

Single node TD-learning and Q -learning. Most existing RL literature is focused on designing and analyzing algorithms that run at a single computing node. In the on-policy setting, the asymptotic convergence of TD-learning was established in (Tsitsiklis & Van Roy, 1997; Tadić, 2001; Borkar, 2009), and the finite-sample bounds were studied

Table 1. Comparison of sample complexity results for federated supervised learning (local SGD) and reinforcement learning algorithms. The possible distributed architectures are: 1) Worker-server, with a central server that coordinates with N agents; 2) Decentralized, where each agent directly communicates with its neighboring agents, without a central server; and 3) Shared memory, where each agent modifies a subset of the parameters of a global model held in a shared memory, that is accessible to all agents. (Recht et al., 2011).

| Algorithm | Architecture | References | Local Updates | Markov Noise | Linear Speedup |
|--------------------------|---------------|----------------------------|---------------|--------------|----------------|
| Local SGD | Worker-server | (Khaled et al., 2020) | ✓ | ✗ | ✓ |
| Local SGD | Worker-server | (Spiridonoff et al., 2021) | ✓ | ✗ | ✓ |
| TD(0) | Worker-server | (Liu & Olshevsky, 2021) | ✓ | ✗ | ✓ |
| Stochastic Approximation | Decentralized | (Wai, 2020) | ✗ | ✓ | ✗ |
| A3C-TD(0) | Shared memory | (Shen et al., 2020) | ✗ | ✗ | ✓ |
| A3C-TD(0) | Shared memory | (Shen et al., 2020) | ✗ | ✓ | ✗ |
| TD & Q -learning | Worker-server | This paper | ✓ | ✓ | ✓ |

in (Dalal et al., 2018; Lakshminarayanan & Szepesvari, 2018; Bhandari et al., 2018; Srikant & Ying, 2019; Hu & Syed, 2019; Chen et al., 2021c). In the off-policy setting, (Maei, 2018; Zhang et al., 2020) study the asymptotic and (Chen et al., 2020a; 2021c) characterize the finite time bound of TD-learning. The Q -learning algorithm was first proposed in (Watkins & Dayan, 1992). There has been a long line of work to establish the convergence properties of Q -learning. In particular, (Tsitsiklis, 1994; Jaakkola et al., 1994; Bertsekas & Tsitsiklis, 1996b; Borkar & Meyn, 2000; Borkar, 2009) characterize the asymptotic convergence of Q -learning, (Beck & Srikant, 2012b; 2013; Wainwright, 2019; Chen et al., 2020a; 2021c) study the finite-sample convergence bound in the mean-square sense, and (Even-Dar & Mansour, 2004; Li et al., 2020; Qu & Wierman, 2020) study the high-probability convergence bounds of Q -learning.

Federated Learning with i.i.d. Noise. When multiple agents are used to expedite sample collection, transferring the samples to a central server for the purpose of training can be costly in applications with high-dimensional data (Shao et al., 2019) and it may also compromise the agents’ privacy. Federated Learning (FL) is an emerging distributed optimization paradigm (Konečný et al., 2016; Kairouz et al., 2019) that utilizes local computation at the agents to train models, such that only model updates, not data, is shared with the central server. In local Stochastic Gradient Descent (Local SGD or FedAvg) (McMahan et al., 2017; Stich, 2018), the core algorithm in FL, locally trained models are periodically averaged by the central server in order to achieve consensus among the agents at a reduced communication cost. While the convergence of local SGD has been extensively studied in prior work (Khaled et al., 2020; Spiridonoff et al., 2021; Qu et al., 2020; Koloskova et al., 2020), these works assume i.i.d. noise in the gradients, which is acceptable for SGD but too restrictive for RL algorithms.

Distributed and Multi-agent RL. Some recent works have analyzed distributed and multi-agent RL algorithms in the presence of Markovian noise in various settings such as decentralized stochastic approximation (Doan et al., 2019; Sun et al., 2020; Wai, 2020; Zeng et al., 2020), TD learning with linear function approximation (Wang et al., 2020a), and off-policy TD in actor-critic algorithms (Chen et al., 2021e;f). However, all these works consider decentralized settings, where the agents communicate with their neighbors after every local update. On the other hand, we consider a federated setting, with each agent performing multiple local updates between successive communication rounds, thereby resulting in communication savings. In (Shen et al., 2020), a parallel implementation of asynchronous advantage actor-critic (A3C) algorithm (which does not have local updates) has been proposed under both i.i.d. and Markov sampling. However, the authors prove a linear speedup only for the i.i.d. case, and an almost linear speedup is observed experimentally for the Markovian case.

3. Preliminaries: Single Node Setting

We model our RL setting with a Markov Decision Process (MDP) with 5 tuples $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where \mathcal{S} and \mathcal{A} are finite sets of states and actions, \mathcal{P} is the set of transition probabilities, \mathcal{R} is the reward function, and $\gamma \in (0, 1)$ denotes the discount factor. At each time step t , the system is in some state S_t , and the agent takes some action A_t according to a policy $\pi(\cdot|S_t)$ in hand, which results in reward $\mathcal{R}(S_t, A_t)$ for the agent. In the next time step, the system transitions to a new state S_{t+1} according to the state transition probability $\mathcal{P}(\cdot|S_t, A_t)$. This series of states and actions $(S_t, A_t)_{t \geq 0}$ constructs a Markov chain, which is the source of the Markovian noise in RL. Throughout this paper we assume that this Markov chain is irreducible and aperiodic (also known as ergodic). It is known that this Markov

chain asymptotically converges to a steady state, and we denote its stationary distribution with μ^π .

To measure the long-term reward achieved by following policy π , we define the value function

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(S_t, A_t) \mid S_0 = s, A_t \sim \pi(\cdot \mid S_t) \right]. \quad (1)$$

Equation (1) is the tabular representation of the value function. Sometimes, however, the size of the state space $|\mathcal{S}|$ is large, and storing $V^\pi(s)$ for all $s \in \mathcal{S}$ is computationally infeasible. Hence, a low dimensional vector $\mathbf{v}^\pi \in \mathbb{R}^d$, where $d \ll |\mathcal{S}|$, can be used to approximate the value function as $V^\pi(s) \approx \phi(s)^\top \mathbf{v}^\pi$ (Tsitsiklis & Van Roy, 1997). Here $\phi(s) \in \mathbb{R}^d$ is a given feature vector corresponding to the state s . Using a low-dimensional vector \mathbf{v}^π to approximate a high-dimensional vector $(V^\pi(s))_{s \in \mathcal{S}}$ is referred to as the function approximation paradigm in RL. For each (s, a) pair, we also define the Q -function, $Q^\pi(s, a) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(S_t, A_t) \mid S_0 = s, A_0 = a]$, which will be employed in Q -learning.

3.1. Temporal Difference Learning

An intermediate goal in RL is to estimate the value function (either $(V^\pi(s))_{s \in \mathcal{S}}$ or \mathbf{v}^π) corresponding to a particular policy π using data collected from the environment. This task is denoted as *policy evaluation* and one of the commonly-used approaches to accomplish this is *Temporal Difference (TD)-learning* (Sutton, 1988). TD-learning is an iterative algorithm where the elements of a d (or $|\mathcal{S}|$, in the tabular setting) dimensional vector is updated until it converges to \mathbf{v}^π (or V^π). This evaluated value function can be employed in different RL algorithms such as actor-critic (Konda & Tsitsiklis, 2000). In the on-policy function approximation setting, the update of the n -step TD-learning is as follows

Sample $A_{t+n} \sim \pi(\cdot \mid S_{t+n}), S_{t+n+1} \sim \mathcal{P}(\cdot \mid S_{t+n}, A_{t+n})$

$$\begin{aligned} \text{update } \mathbf{v}_{t+1} = & \mathbf{v}_t + \alpha \phi(S_t) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathcal{R}(S_l, A_l) \\ & + \gamma \phi(S_{l+1})^\top \mathbf{v}_t - \phi(S_l)^\top \mathbf{v}_t), \end{aligned} \quad (2)$$

where α is the step size. Note that in this setting, the evaluating policy and the sampling policy coincide. In contrast, in the off-policy setting these two policies can in general differ, and we need to account for this difference while running the algorithm. We will further expand on TD-learning and its variants in Sections 4.1 and 5.1.1.

3.2. Control Problem and Q -learning

Assuming some initial distribution ξ on the state space, the average value function corresponding to policy π is

defined as $V^\pi(\xi) = \mathbb{E}_{s \sim \xi}[V^\pi(s)]$. This scalar quantity is a metric of average long-term rewards achieved by the agent, when it starts from distribution ξ and follows policy π . The ultimate goal of the agent is to obtain an optimal policy π^* which results in the maximum long-term rewards, i.e. $\pi^* \in \arg \max_{\pi} V^\pi(\xi)$. Throughout the paper, we denote the parameters corresponding to the optimal policy with $*$, e.g., $V^{\pi^*}(\xi) \equiv V^*(\xi)$. The task of obtaining the optimal policy in RL is denoted as the *control problem*.

Q -learning (Watkins & Dayan, 1992) is one of the most widely used algorithms in RL to solve the control problem. At each time step t , Q -learning preserves a $|\mathcal{S}| \cdot |\mathcal{A}|$ dimensional table Q_t , and updates it table as $Q_{t+1}(s, a) = Q_t(S_t, A_t) + \alpha(\mathcal{R}(S_t, A_t) + \gamma \max_a Q_t(S_{t+1}, a) - Q_t(S_t, A_t))$, if $(s, a) = (S_t, A_t)$ and $Q_{t+1}(s, a) = Q_t(s, a)$ otherwise. The $|\mathcal{S}| \cdot |\mathcal{A}|$ elements of the vector Q_t are updated iteratively until it converges to Q^* , corresponding to an optimal policy. Using Q^* , one can obtain an optimal policy via greedy selection.

3.3. Stochastic Approximation and Finite Sample Bounds

Both TD-learning and Q -learning can be seen as variants of stochastic approximation (Chen et al., 2020b; 2019b;a; 2021d; Tsitsiklis, 1994). While generic stochastic approximation algorithms are studied under i.i.d. noise (Even-Dar & Mansour, 2004; Shah & Xie, 2018; Wainwright, 2019; Liu et al., 2015; Dalal et al., 2018), to apply them for studying RL we need to understand stochastic approximation under Markovian noise (Tsitsiklis, 1994; Qu & Wierman, 2020; Srikant & Ying, 2019; Chen et al., 2021c) which is significantly more challenging.

For a generic stochastic approximation (i.i.d. or Markovian noise) with constant step size α , parameter vector \mathbf{x}_T , and convergent point \mathbf{x}^* , it can be shown that the algorithm have the following convergence behaviour

$$\mathbb{E}[\|\mathbf{x}_T - \mathbf{x}^*\|^2] \leq \mathcal{C}_1(1 - \mathcal{C}_0\alpha)^T + \mathcal{C}_2\alpha, \quad (3)$$

where $\mathcal{C}_0, \mathcal{C}_1$, and \mathcal{C}_2 are some problem dependent positive constants (Look at Appendix A for a discussion on a lower bound on the convergence of general stochastic approximation). The first term is denoted as the bias and the second term is called the variance. According to this bound, \mathbf{x}_T geometrically converges to a ball around \mathbf{x}^* with radius proportional to $\mathcal{C}_2\alpha$. Notice that we can always reduce the variance term by reducing the step size α , but this will lead to slower convergence in the bias term. In particular, in order to get $\mathbb{E}[\|\mathbf{x}_T - \mathbf{x}^*\|^2] \leq \epsilon$, it is easy to see that we need $T \geq \mathcal{O}(\frac{\mathcal{C}_2}{\epsilon} \log \frac{1}{\epsilon})$ sample complexity. Now suppose the constant \mathcal{C}_2 is large. In this case, the variance term in the bound in (3) is large, and the sample complexity, which is proportional to \mathcal{C}_2 will be poor. Notice that by the dis-

cussion in Appendix A, this bound is tight and cannot be improved.

This is where the FL can be employed in order to control the variance term by generating more data. For instance, in federated TD-learning, multiple agents work together to evaluate the value function simultaneously. Due to this collaboration, the agents can estimate the true value function with a lower variance. The same holds for estimating Q^* in Q -learning.

4. Federated On-policy RL

4.1. On-policy TD-learning with linear function approximation

In this section we describe the TD-learning with linear function approximation and online data samples in the single node setting. In this problem, we consider a full rank feature matrix $\Phi \in |\mathcal{S}| \times d$, and we denote s -th row of this matrix with $\phi(s)$, $s \in \mathcal{S}$. The goal is to find $\mathbf{v}^\pi \in \mathbb{R}^d$ which solves the following fixed point equation:

$$\Phi \mathbf{v}^\pi = \Pi_\Phi((\mathcal{T}^\pi)^n \Phi \mathbf{v}^\pi). \quad (4)$$

In equation (4), $\Pi_\Phi(\cdot)$ is the projection with respect to the weighted 2-norm, i.e., $\Pi_\Phi(\mathbf{V}) = \arg \min_{\mathbf{v} \in \mathbb{R}^d} \|\Phi \mathbf{v} - \mathbf{V}\|_\pi$. Here $\|\mathbf{V}\|_\pi = \sqrt{\mathbf{V}^\top \boldsymbol{\mu}^\pi \mathbf{V}}$ and $\boldsymbol{\mu}^\pi$ is a diagonal matrix with diagonal entries corresponding to μ^π . In equation (4), $(\mathcal{T}^\pi)^n$ denotes the n -step Bellman operator (Tsitsiklis & Van Roy, 1997). It is known (Tsitsiklis & Van Roy, 1997) that equation (4) has a unique solution \mathbf{v}^π , and $\Phi \mathbf{v}^\pi$ is ‘‘close’’ to the true value function V^π . n -step TD-learning algorithm, which was shown in (2), is an iterative algorithm to obtain this unique fixed point using samples from the environment. Note that in this algorithm states and actions are sampled over a single trajectory, and hence the noise in updating \mathbf{v}_t is Markovian. Furthermore, since the policy which samples the actions and the evaluating policy are both π , this algorithm is on-policy. As described in (Tsitsiklis & Van Roy, 1997; Bertsekas & Tsitsiklis, 1996a), the TD-learning algorithm can be studied under the umbrella of linear stochastic approximation with Markovian noise. More recently, the authors in (Bhandari et al., 2018; Srikant & Ying, 2019) have shown that the update parameter of TD-learning \mathbf{v}_t converges to \mathbf{v}^π in the form $\mathbb{E}[\|\mathbf{v}_t - \mathbf{v}^\pi\|_2^2] \leq \mathcal{O}((1 - C_0\alpha)^t + \alpha)$. In the next section we show how FL can improve this result.

4.2. Federated TD-learning with linear function approximation

The federated version of on-policy n -step TD-learning with linear function approximation is shown in Algorithm 1. In this algorithm we consider N agents which collaboratively work together to evaluate \mathbf{v}^π . For each agent i , $i = 1, 2, \dots, N$, we initialize their corresponding pa-

rameters $\mathbf{v}_0^i = \mathbf{0}$. Furthermore, each agent i samples its initial state S_0^i from some given distribution ξ . In the next time steps, each agent follows a single Markovian trajectory generated by policy π , independently from other agents. At each time t , the parameter of each agent i is updated using this independently generated trajectory as $\mathbf{v}_{t+1}^i = \mathbf{v}_t^i + \alpha \phi(S_t^i) E_{t,n}^i$. Finally, in order to ensure convergence to a global optimum, every K time steps all the agents send their parameters to a central server. The central server evaluates the average of these parameters and returns this average to each of the agents. Each agent then continues their update procedure using this average.

Notice that the averaging step is essential to ensure synchronization among the agents. Smaller K results in more frequent synchronization, and hence better convergence guarantees. However, setting smaller K is equivalent to more number of communications between the single agents and the central server, which incurs higher cost. Hence, an intermediate value for K has to be chosen to strike a balance between the communication cost and the accuracy. At the end, the algorithm samples a time step $\hat{T} \sim q_T^c$, where

$$q_T^c(t) = \frac{c^{-t}}{\sum_{t'=0}^{T-1} c^{-t'}} \quad \text{for } t = 0, 1, \dots, T-1 \quad (5)$$

and $c > 1$ is some constant. Since we have $q_T^c(t) \geq 0$ and $\sum_{t=0}^{T-1} q_T^c(t) = 1$, it is clear that $q_T^c(\cdot)$ is a probability distribution over the time interval $[0, T-1]$. In Theorem 4.1 we characterize the convergence of this algorithm as a function of α , N , and K . Throughout the paper, $\tilde{O}(\cdot)$ ignores the logarithmic terms.

Algorithm 1 Federated n -step TD (On-policy, Function Approx.)

- 1: **Input:** Policy π, ξ
 - 2: **Initialization:** $\mathbf{v}_0^i = \mathbf{0}$ and $S_0^i \sim \xi$ and $\{A_l^i, S_{l+1}^i\} \sim \pi$ for $0 \leq l \leq n-1$ and all i
 - 3: **for** $t = 0$ to $T-1$ **do**
 - 4: **for** $i = 1, \dots, N$ **do**
 - 5: Sample $A_{t+n}^i \sim \pi(\cdot | S_{t+n}^i), S_{t+n+1}^i \sim \mathcal{P}(\cdot | S_{t+n}^i, A_{t+n}^i)$
 - 6: $e_{t,l}^i = \mathcal{R}(S_l^i, A_l^i) + \gamma \phi(S_{l+1}^i)^\top \mathbf{v}_t^i - \phi(S_l^i)^\top \mathbf{v}_t^i$ for $l = t, \dots, t+n-1$
 - 7: $E_{t,n}^i = \sum_{l=t}^{t+n-1} \gamma^{l-t} e_{t,l}^i$
 - 8: $\mathbf{v}_{t+1}^i = \mathbf{v}_t^i + \alpha \phi(S_t^i) E_{t,n}^i$
 - 9: **end for**
 - 10: **if** $t+1 \bmod K = 0$ **then**
 - 11: $\mathbf{v}_{t+1}^j \leftarrow \frac{1}{N} \sum_{j=1}^N \mathbf{v}_{t+1}^j, \forall i \in [N]$
 - 12: **end if**
 - 13: **end for**
 - 14: **Sample** $\hat{T} \sim q_T^{c,DDL}$
 - 15: **Return:** $\frac{1}{N} \sum_{i=1}^N \mathbf{v}_{\hat{T}}^i$
-

Theorem 4.1. Let $\mathbf{v}_{\hat{T}} = \frac{1}{N} \sum_{i=1}^N \mathbf{v}_{\hat{T}}^i$ denote the average of the parameters across agents at the random time \hat{T} . For small enough step size α , and $T \geq \mathcal{O}(\log \frac{1}{\alpha})$, there exist constant $c_{TDL} \in (0, 1)$ (see Section C.1 for precise statement), such that we have

$$\begin{aligned} \mathbb{E}[\|\mathbf{v}_{\hat{T}} - \mathbf{v}^\pi\|_2^2] \leq & C_1^{TDL} \frac{1}{\alpha} (1 - \alpha C_0^{TDL})^T + C_2^{TDL} \frac{\alpha \tau_\alpha^2}{N} \\ & + C_3^{TDL} (K - 1) \alpha^2 \tau_\alpha + C_4^{TDL} \alpha^3 \tau_\alpha^2, \end{aligned}$$

where C_i^{TDL} , $i = 0, 1, 2, 3, 4$ are problem dependent constants, and $\tau_\alpha = \mathcal{O}(\log(1/\alpha))$. By choosing $\alpha = \mathcal{O}(\frac{\log(NT)}{T})$ and $K = T/N$, we achieve $\mathbb{E}[\|\mathbf{v}_{\hat{T}} - \mathbf{v}^\pi\|_2^2] \leq \epsilon$ within $T = \tilde{\mathcal{O}}(\frac{1}{N\epsilon})$ iterations.

For brevity purposes, here we did not show the exact dependence of the constants C_i^{TDL} , $i = 0, 1, 2, 3$ on the problem dependent constants. For a discussion on the detailed expression look at Section C.1 in the appendix.

Theorem 4.1 shows that federated TD-learning with linear function approximation enjoys a linear speedup with respect to the number of agents. Compared to the convergence bound of general stochastic approximation in (3), the bound in Theorem 4.1 has three differences. Firstly, the variance term which is proportional to the step size α is divided with the number of agents N . This will allow us to control the variance (and hence improve sample complexity) by employing more number of agents. Secondly, we have an extra term which is zero with perfect synchronization $K = 1$. Although this term is not divided with N , but it is proportional to α^2 , which is one order higher than the variance term in (3). Finally, the last term is of the order $\tilde{\mathcal{O}}(\alpha^3)$, which can be handled by choosing small enough step size.

Furthermore, according to the choice of K in Theorem 4.1, after T iterations, the communication cost of federated TD is $T/K = N$. However, by employing federated TD-learning in the naive setting where all the agents communicate with the central server at every time step, the communication cost will be $\mathcal{O}(T)$. Hence, we observe that by carefully tuning the hyper parameters of federated TD, we can significantly reduce the communication cost of the overall algorithm, while not losing performance in terms of the sample complexity.

Finally, federated TD-learning Algorithm 1 preserves the privacy of the agents. In particular, since the single agents only require to share their parameters \mathbf{v}_{t+1}^i , the central server will not be exposed to the state-action-reward trajectory generated by each agent. This can be essential in some applications where privacy is an issue (Mothukuri et al., 2021; Truex et al., 2019). Examples of such applications include autonomous driving (Liang et al., 2019; Zhao et al., 2021), Internet of Things (IoT) (Nguyen et al., 2021;

Ren et al., 2019; Wang et al., 2020b), and cloud robotics (Liu et al., 2019; Xu et al., 2021).

Remark. In algorithm 1, the randomness in choosing \hat{T} is independent of all the other randomness in the problem. Hence, in a practical setting, one can sample \hat{T} ahead of time, before running the algorithm, and stop the algorithm at time step \hat{T} and output $\mathbf{v}_{\hat{T}}$. By this method, we require only a single data point to be saved, which results in the memory complexity of $\mathcal{O}(1)$ for the algorithm.

5. Federated Off-Policy RL

On-policy TD-learning requires online sampling from the environment, which might be costly (e.g. robotics (Gu et al., 2017; Levine et al., 2020)), high risk (e.g. self-driving cars (Yurtsever et al., 2020; Maddern et al., 2017)), or unethical (e.g. in clinical trials (Gottesman et al., 2019; Liu et al., 2018; Gottesman et al., 2020)). Off-policy training in RL refers to the paradigm where we use data collected by a fixed behaviour policy to run the algorithm. When employed in federated setting, off-policy RL has privacy advantages as well (Foerster et al., 2016; Qi et al., 2021; Zhuo et al., 2019). In particular, suppose each single agent attains a unique sampling policy, and they do not wish to reveal these policies to the central server. In off-policy FL, agents only transmit sampled data, and hence the sampling policies remain private to each agent.

In Section 5.1 we will discuss off-policy TD-learning and in Section 5.2 we will discuss Q -learning, which is an off-policy control algorithm. For the off-policy algorithms, we only study the tabular setting. Notice that it has been observed that the combination of off-policy sampling and function approximation in RL (also known as deadly triad (Sutton & Barto, 2018)) can result in instability or even divergence (Baird, 1995). Recently there has been some work to overcome deadly triad (Chen et al., 2021b). Extension of our work to function approximation in the off-policy setting is a future research direction.

5.1. Federated Off-Policy TD-learning

In the following, we first discuss single-node off-policy TD-learning, and then we generalize it to the federated setting.

5.1.1. OFF-POLICY TD-LEARNING

In off-policy TD-learning the goal is to evaluate the value function $\mathbf{V}^\pi = (V^\pi(s))_{s \in \mathcal{S}}$ corresponding to the policy π using data sampled from some fixed behaviour policy π_b . In this setting, the evaluating policy π and the sampling policy π_b can be arbitrarily different, and we need to account for this difference while performing the evaluation. Although π and π_b can be different, notice that the value function \mathbf{V}^π does not depend on π_b . In order to account for this

difference, we introduce the notion of *importance sampling* as $\mathfrak{J}^b(s, a) = \frac{\pi(a|s)}{\pi^b(a|s)}$ which is employed in the off-policy TD-learning.

Recently, several works studied the finite-time convergence of off-policy TD-learning. In particular, the authors in (Khadadian et al., 2021; Chen et al., 2021c; 2020b; 2021d) show that, similar to on-policy TD, off-policy TD-learning can be studied under the umbrella of stochastic approximation. Hence, this algorithm enjoys similar convergence behaviour as (3).

5.1.2. FEDERATED OFF-POLICY TD-LEARNING

The federated version of n -step off-policy TD-learning is shown in Algorithm 2. In this algorithm, each agent i attains a unique (and private) sampling policy π^i and follows an independent trajectory generated by this policy. Furthermore, at each time step t , each agent i attains a $|\mathcal{S}|$ -dimensional vector \mathbf{V}_t^i and updates this vector using the samples generated by π^i . In order to account for the off-policy sampling, each agent utilizes $\mathfrak{J}^{(i)}(S^i, A^i) = \frac{\pi(A^i|S^i)}{\pi^i(A^i|S^i)}$ in the update of their algorithm. We further define $\mathfrak{J}_{\max} = \max_{s,a,i} \mathfrak{J}^{(i)}(s, a)$, which is a measure of discrepancy between the evaluating policy π and sampling policy π^i of all the agents.

In order to ensure synchronization, all the agents transmit their parameter vectors to the central server every K time steps. The central server returns the average of these vectors to each agent and each agent follows this averaged vector afterwards. Notice that in federated off-policy TD-learning Algorithm 2, each agent share neither their sampled trajectory of state-action-rewards, nor their sampling policy with the central server. This provides two levels of privacy for the single agents. At the end, the algorithm samples a time step $\hat{T} \sim q_T^{cTD}$, where the distribution q_T^c is defined in (5) and $c_{TD} = 1 - \frac{\alpha\varphi_{TD}}{2}$, where $\varphi_{TD} = 1 - \frac{0.5e^{1/4}(2-\mu_{\min}(1-\gamma^{n+1}))}{\sqrt{e-1+\left(\frac{2-\mu_{\min}(1-\gamma^{n+1})}{2-2\mu_{\min}(1-\gamma^{n+1})}\right)^2}}$. Here, we de-

note $\mu_{\min} = \min_{s,i} \mu^{\pi^i}(s)$. The constant c_{TD} is carefully chosen to ensure the convergence of Algorithm 2. Furthermore, for small enough step size α , it can be shown that $0 < c_{TD} < 1$.

Theorem 5.1 states the convergence of this Algorithm.

Theorem 5.1. *Consider the federated n -step off-policy TD-learning Algorithm 2. Denote $\mathbf{V}_{\hat{T}} = \frac{1}{N} \sum_{i=1}^N \mathbf{V}_{\hat{T}}^i$. For small enough step size α and large enough T , we have*

$$\mathbb{E}[\|\mathbf{V}_{\hat{T}} - \mathbf{V}^\pi\|_\infty^2] \leq \mathcal{C}_1^{TD_T} \frac{1}{\alpha} \mathcal{C}_{TD}^T + \mathcal{C}_2^{TD_T} \frac{\alpha \tau_\alpha^2}{N} + \mathcal{C}_3^{TD_T} (K-1) \alpha^2 \tau_\alpha,$$

$$\text{where } \mathcal{C}_1^{TD_T} = \bar{\mathcal{C}}_1^{TD_T} \cdot \frac{\mathfrak{J}_{\max}^{4n-2}}{\mu_{\min}(1-\gamma)^3}, \quad \mathcal{C}_2^{TD_T} =$$

Algorithm 2 Federated n -step TD (Off-policy Tabular Setting)

- 1: **Input:** Policy π, ξ
- 2: **Initialization:** $\mathbf{V}_0^i = \mathbf{0}$ and $S_0^i \sim \xi$ and $\{A_l^i, S_{l+1}^i\} \sim \pi^i$ for $0 \leq l \leq n-1$ and all i
- 3: **for** $t = 0$ to $T-1$ **do**
- 4: **for** $i = 1, \dots, N$ **do**
- 5: Sample $A_{t+n}^i \sim \pi^i(\cdot | S_{t+n}^i), S_{t+n+1}^i \sim \mathcal{P}(\cdot | S_{t+n}^i, A_{t+n}^i)$
- 6: $e_{t,l}^i = \mathcal{R}(S_t^i, A_t^i) + \gamma \mathbf{V}_t^i(S_{l+1}^i) - \mathbf{V}_t^i(S_t^i)$
- 7: Update $\mathbf{V}_{t+1}^i(s) = \mathbf{V}_t^i(s) + \alpha \sum_{l=t}^{t+n-1} \gamma^{l-t} [\prod_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i)] e_{t,l}^i$ if $s = S_t^i$ and $\mathbf{V}_{t+1}^i(s) = \mathbf{V}_t^i(s)$ otherwise.
- 8: **end for**
- 9: **if** $t+1 \bmod K = 0$ **then**
- 10: $\mathbf{V}_{t+1}^i \leftarrow \frac{1}{N} \sum_{j=1}^N \mathbf{V}_{t+1}^j, \forall i \in [N]$
- 11: **end if**
- 12: **end for**
- 13: **Sample** $\hat{T} \sim q_T^{cTD}$
- 14: **Return:** $\frac{1}{N} \sum_{i=1}^N \mathbf{V}_{\hat{T}}^i$

$\bar{\mathcal{C}}_2^{TD_T} \cdot \frac{\mathfrak{J}_{\max}^{3n-1} |\mathcal{S}| \log^2(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^4}, \mathcal{C}_3^{TD_T} = \bar{\mathcal{C}}_3^{TD_T} \cdot \frac{\mathfrak{J}_{\max}^{7n-3} |\mathcal{S}|^2 \log^2(|\mathcal{S}|)}{\mu_{\min}^4 (1-\gamma)^8}$, and $\bar{\mathcal{C}}_i^{TD_T}, i = 1, 2, 3$ are universal problem independent constants. In addition, choosing $\alpha = \frac{8 \log(NT)}{T \varphi_{TD}}$ and $K = T/N$, we have $\mathbb{E}[\|\mathbf{V}_{\hat{T}} - \mathbf{V}^\pi\|_\infty^2] \leq \epsilon$ after $T = \tilde{\mathcal{O}}\left(\frac{1}{N\epsilon} \cdot \frac{\mathfrak{J}_{\max}^{7n-3} |\mathcal{S}|^2 \log^2(|\mathcal{S}|)}{\mu_{\min}^5 (1-\gamma)^9}\right)$ iterations.

The proof is given in Section C.2 in the appendix.

Note that similar to on-policy TD-learning Algorithm 1, off-policy TD-learning also enjoys a linear speedup while maintaining a low communication cost. In addition, this algorithm preserves the privacy of the agents by holding both the data and the sampling policy private.

5.2. Federated Q -learning

So far we have discussed policy evaluation problem with on and off-policy samples. Next we aim at solving the control problem by employing the celebrated Q -learning algorithm (Watkins & Dayan, 1992; Tsitsiklis, 1994). In the next section we will explain the Q -learning algorithm. Further, in Section 5.2.2 we will provide a federated version of Q -learning along with its convergence result.

5.2.1. Q -LEARNING

The goal of Q -learning is to evaluate Q^* , which is the unique Q -function corresponding to the optimal policy. Knowing Q^* , one can obtain an optimal policy through a greedy selection (Puterman, 2014), and hence resolve the control problem.

Suppose $\{S_t, A_t\}_{t \geq 0}$ is generated by a fixed behaviour policy π_b . At each time step t , Q -learning preserves a $|\mathcal{S}| \cdot |\mathcal{A}|$ table Q_t and updates the elements of this table as shown in Section 3.2. By assuming π_b to be an ergodic policy, the asymptotic convergence of Q_t to Q^* has been established in (Bertsekas & Tsitsiklis, 1996b). Furthermore, it can be shown that Q -learning is a special case of stochastic approximation and enjoys a convergence bound similar to (3) (Beck & Srikant, 2012a; Li et al., 2020; Qu & Wierman, 2020; Chen et al., 2021c).

Two points worth mentioning about the Q -learning algorithm. Firstly, Q -learning is an off-policy algorithm in the sense that only samples from a fixed ergodic policy is needed to perform the algorithm. Secondly, as opposed to the TD-learning, the update of the Q -learning is *non-linear*. This imposes a sharp contrast between the analysis of Q -learning and TD-learning (Chen et al., 2019a).

5.2.2. FEDERATED Q -LEARNING

Algorithm 3 provides the federated version of Q -learning. We characterize its convergence in the following theorem.

Algorithm 3 Federated Q -learning

- 1: **Input:** Sampling policy π_b^i for $i = 1, 2, \dots, N$, initial distribution ξ
 - 2: **Initialization:** $Q_0^i = \mathbf{0}$ and $S_0^i \sim \xi$ for all i
 - 3: **for** $t = 0$ to $T - 1$ **do**
 - 4: **for** $i = 1, \dots, N$ **do**
 - 5: Sample $A_t^i \sim \pi_b^i(\cdot | S_t^i), S_{t+1}^i \sim \mathcal{P}(\cdot | S_t^i, A_t^i)$
 - 6: Update $Q_{t+1}^i(s, a) = Q_t^i(S_t^i, A_t^i) + \alpha [\mathcal{R}(S_t^i, A_t^i) + \gamma \max_a Q_t^i(S_{t+1}^i, a) - Q_t^i(S_t^i, A_t^i)]$, if $(s, a) = (S_t^i, A_t^i)$ and $Q_{t+1}^i(s, a) = Q_t^i(s, a)$ otherwise.
 - 7: **end for**
 - 8: **if** $t + 1 \bmod K = 0$ **then**
 - 9: $Q_{t+1}^i \leftarrow \frac{1}{N} \sum_{j=1}^N Q_{t+1}^j, \forall i \in [N]$
 - 10: **end if**
 - 11: **end for**
 - 12: **Sample:** $\hat{T} \sim q_T^{c_Q}$
 - 13: **Return:** $\frac{1}{N} \sum_{i=1}^N Q_{\hat{T}}^i$
-

Theorem 5.2. Consider the federated Q -learning Algorithm 3 with $c_Q = 1 - \frac{\alpha \varphi_Q}{2} \in (0, 1)$, where $\varphi_Q = 1 - \frac{0.5e^{1/4}(2 - \mu_{\min}(1 - \gamma))}{\sqrt{e-1} + \left(\frac{2 - \mu_{\min}(1 - \gamma)}{2 - 2\mu_{\min}(1 - \gamma)}\right)^2}$ and we denote $\mu_{\min} = \min_{s,a,i} \mu^{\pi^i}(s) \pi^i(a|s)$. Denote $Q_{\hat{T}} = \frac{1}{N} \sum_{i=1}^N Q_{\hat{T}}^i$. For small enough step size α and large enough T , we have

$$\mathbb{E}[\|Q_{\hat{T}} - Q^*\|_\infty^2] \leq \mathcal{C}_1^Q \frac{1}{\alpha} c_Q^T + \mathcal{C}_2^Q \frac{\alpha \tau_\alpha^2}{N} + \mathcal{C}_3^Q (K - 1) \alpha^2 \tau_\alpha,$$

where $\mathcal{C}_1^Q = \bar{\mathcal{C}}_1^Q \cdot \frac{1}{\mu_{\min}(1 - \gamma)^3}$, $\mathcal{C}_2^Q = \bar{\mathcal{C}}_2^Q \cdot \frac{|\mathcal{S}| \log^2(|\mathcal{S}|)}{\mu_{\min}^2(1 - \gamma)^4}$, $\mathcal{C}_3^Q = \bar{\mathcal{C}}_3^Q \cdot \frac{|\mathcal{S}|^2 \log^2(|\mathcal{S}|)}{\mu_{\min}^4(1 - \gamma)^8}$, and $\bar{\mathcal{C}}_i^Q, i = 1, 2, 3$ are universal problem

independent constants. In addition, choosing $\alpha = \frac{8 \log(NT)}{T \varphi_Q}$ and $K = T/N$, we have $\mathbb{E}[\|Q_{\hat{T}} - Q^*\|_\infty^2] \leq \epsilon$ within $T = \tilde{\mathcal{O}}\left(\frac{1}{N\epsilon} \cdot \frac{|\mathcal{S}|^2 \log^2(|\mathcal{S}|)}{\mu_{\min}^5(1 - \gamma)^9}\right)$ iterations.

According to Theorem 5.2, federated Q -learning Algorithm 3, similar to federated off-policy TD-learning, enjoys linear speedup, communication efficiency as well as privacy guarantees. We would like to emphasize that the update of Q -learning is non-linear. Hence the result of Theorem 5.2 cannot be derived from Theorems 4.1 and 5.1.

6. Generalized Federated Stochastic Approximation

In this section we study the convergence of a general federated stochastic approximation for contractive operators, FedSAM, which is presented in Algorithm 4. In this algorithm there are N agents $i = 1, 2, \dots, N$. At each time step $t \geq 0$, each agent i maintains the parameter $\theta_t^i \in \mathbb{R}^d$. At time $t = 0$, all agents initialize their parameters with $\theta_0^i = \theta_0$. Next, at time $t \geq 0$, each agent i updates its parameter as $\theta_{t+1}^i = \theta_t^i + \alpha (\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i) - \theta_t^i + \mathbf{b}^i(\mathbf{y}_t^i))$. Here α denotes the step size, and \mathbf{y}_t^i is a noise which is Markovian along the time t , but is independent across the agents i . This notion is defined more concretely in Assumption 6.4. We note that functions $\mathbf{G}^i(\cdot, \cdot)$ and $\mathbf{b}^i(\cdot)$ are allowed to be dependent on the agent i . This allows us to employ the convergence bound of FedSAM in order to derive the convergence bound of off-policy TD-learning with different behaviour policies across agents. In order to avoid divergence, every K time steps we synchronize the parameters of all the agents as $\theta_t^i \leftarrow \theta_t \triangleq \frac{1}{N} \sum_{j=1}^N \theta_t^j$, for all $i \in [N]$. Note that although smaller K corresponds to more frequent synchronization and hence more ‘‘accurate’’ updates, at the same time it results in a higher communication cost, which is not desirable. Hence, in order to determine the optimal choice of synchronization period, it is essential to characterize the dependence of the convergence on K . This is one of the results which we will derive in Theorem B.1. Finally, the algorithm samples $\hat{T} \sim q_T^c$, where $q_T^c(t) = \frac{e^{-t}}{\sum_{t'=0}^{T-1} e^{-t'}}$ and outputs $\theta_{\hat{T}}$. This sampling scheme is essential for the convergence of overall algorithm. We further make some assumptions regarding the underlying process.

First, we assume that the expectation of $\mathbf{G}^i(\theta, \mathbf{y}_t^i)$ geometrically converges to some function $\bar{\mathbf{G}}^i(\theta)$ and the expectation of $\mathbf{b}^i(\mathbf{y}_t^i)$ geometrically converges to $\mathbf{0}$. In particular, we have the following assumption.

Assumption 6.1. For every agent i , there exist a function $\bar{\mathbf{G}}^i(\theta)$ such that we have

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\mathbf{G}^i(\theta, \mathbf{y}_t^i)] &= \bar{\mathbf{G}}^i(\theta) \\ \lim_{t \rightarrow \infty} \mathbb{E}[\mathbf{b}^i(\mathbf{y}_t^i)] &= \mathbf{0}. \end{aligned} \quad (6)$$

Algorithm 4 Federated Stochastic Approximation with Markovian Noise (FedSAM)

```

1: Input:  $c_{FSAM}, T, \theta_0, K, \alpha$ 
2:  $\theta_0^i = \theta_0$  for all  $i = 1, \dots, N$ .
3: for  $t = 0$  to  $T - 1$  do
4:    $\theta_{t+1}^i = \theta_t^i + \alpha (\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i) - \theta_t^i + \mathbf{b}^i(\mathbf{y}_t^i)), \forall i \in [N]$ 
5:   if  $t + 1 \bmod K = 0$  then
6:      $\theta_{t+1}^i = \theta_{t+1}^i \triangleq \frac{1}{N} \sum_{j=1}^N \theta_{t+1}^j, \forall i \in [N]$ 
7:   end if
8: end for
9: Sample  $\hat{T} \sim q_T^{c_{FSAM}}(\cdot)$ .
10: Return:  $\frac{1}{N} \sum_{i=1}^N \theta_{\hat{T}}^i$ 
    
```

Furthermore, there exists $m_1, m_2 \geq 0$ and $\rho \in [0, 1)$, such that for every $i = 1, 2, \dots, N$,

$$\begin{aligned} \|\bar{\mathbf{G}}^i(\theta) - \mathbb{E}[\bar{\mathbf{G}}^i(\theta, \mathbf{y}_t^i)]\|_c &\leq m_1 \|\theta\|_c \rho^t \\ \|\mathbb{E}[\bar{\mathbf{b}}^i(\mathbf{y}_t^i)]\|_c &\leq m_2 \rho^t, \end{aligned} \quad (7)$$

where $\|\cdot\|_c$ is a given norm.

Next, we assume a contraction property on the expected operator $\bar{\mathbf{G}}^i(\theta)$.

Assumption 6.2. We assume all expected operators $\bar{\mathbf{G}}^i(\theta)$ are contraction mappings with respect to $\|\cdot\|_c$ with contraction factor $\gamma_c \in (0, 1)$. That is, for all $i = 1, 2, \dots, N$,

$$\|\bar{\mathbf{G}}^i(\theta_1) - \bar{\mathbf{G}}^i(\theta_2)\|_c \leq \gamma_c \|\theta_1 - \theta_2\|_c, \quad \forall \theta_1, \theta_2 \in \mathbb{R}^d.$$

Next, we consider some Lipschitz and boundedness properties on $\mathbf{G}^i(\cdot, \cdot)$ and $\mathbf{b}^i(\cdot)$.

Assumption 6.3. For all $i = 1, \dots, N$, there exist constants A_1, A_2 and B such that

1. $\|\mathbf{G}^i(\theta_1, \mathbf{y}^i) - \mathbf{G}^i(\theta_2, \mathbf{y}^i)\|_c \leq A_1 \|\theta_1 - \theta_2\|_c$, for all $\theta_1, \theta_2, \mathbf{y}^i$.
2. $\|\mathbf{G}^i(\theta, \mathbf{y}^i)\|_c \leq A_2 \|\theta\|_c$ for all θ, \mathbf{y}^i .
3. $\|\mathbf{b}^i(\mathbf{y}^i)\|_c \leq B$ for all \mathbf{y}^i .

Remark. By Assumption 6.2 and due to the Banach fixed point theorem, $\bar{\mathbf{G}}^i(\cdot)$ has a unique fixed point for all $i = 1, 2, \dots, N$. Furthermore, by Assumption 6.3, we have $\mathbf{G}^i(\mathbf{0}, \mathbf{y}) = \mathbf{0}$. Hence the point $\mathbf{0}$ is the unique fixed point of $\bar{\mathbf{G}}^i(\cdot)$.

Finally, we impose an assumption on the random data \mathbf{y}_t^i .

Assumption 6.4. We assume that the Markovian noise \mathbf{y}_t^i (Markovian with respect to time t) is independent across agents i . In other words, for all measurable functions $f(\cdot)$ and $g(\cdot)$, we assume the following

$$\mathbb{E}_{t-r}[f(\mathbf{y}_t^i) \times g(\mathbf{y}_t^j)] = \mathbb{E}_{t-r}[f(\mathbf{y}_t^i)] \times \mathbb{E}_{t-r}[g(\mathbf{y}_t^j)],$$

for all $r \leq t, i \neq j$.

Theorem 6.1 states the convergence of Algorithm 4.

Theorem 6.1. Consider the federated stochastic approximation Algorithm 4 with $c_{FSAM} = 1 - \frac{\alpha\varphi_2}{2} \in (0, 1)$ (φ_2 is defined in Equation (14) in the appendix), and synchronization frequency K . Denote $\theta_t = \frac{1}{N} \sum_{i=1}^N \theta_t^i$, and consider $\theta_{\hat{T}}$ as the output of this algorithm after T iterations. Assume $\tau_\alpha = \lceil 2 \log_\rho \frac{1}{\alpha} \rceil$. For $T \geq \max\{K + \tau_\alpha, 2\tau_\alpha\}$ and small enough step size α , we have

$$\begin{aligned} \mathbb{E}[\|\theta_{\hat{T}}\|_c^2] &\leq \mathcal{C}_1 \frac{1}{\alpha} c_{FSAM}^{T-2\tau_\alpha+1} + \mathcal{C}_2 \frac{\alpha\tau_\alpha^2}{N} + \mathcal{C}_3 (K-1)\alpha^2\tau_\alpha \\ &\quad + \mathcal{C}_4 \alpha^3 \tau_\alpha^2, \end{aligned} \quad (8)$$

where $\mathcal{C}_i, i = 1, 2, 3, 4$ are some constants which are specified precisely in Appendix B, and are independent of K, α, N . Choosing $\alpha = \frac{8 \log(NT)}{T\varphi_2}$ and $K = T/N$, we get $T = \tilde{\mathcal{O}}(\frac{1}{N\epsilon})$ sample complexity for achieving $\mathbb{E}[\|\theta_{\hat{T}}\|_c^2] \leq \epsilon$.

Theorem 6.1 establishes the convergence of $\theta_{\hat{T}}$ to zero in the expected mean-squared sense. The first term in (8) converges geometrically to zero as T grows. The second term is proportional to α similar as (3). However, the number of agents N in the denominator ensures *linear speedup*, meaning that for small enough α (such that α/N is the dominant term), the sample complexity of each individual agent, relative to a centralized system, is reduced by a factor of N . The third term has quadratic dependence on α , and is zero when we have perfect synchronization, i.e. $K = 1$. The last term is proportional to α^3 , and has the weakest dependency on the step size α . For $K > 1$ we can merge the last two terms by upper bounding $\alpha^3 \leq \alpha^2$. The current upper bound, however, is tighter since with $K = 1$ (i.e. perfect synchronization) we have no term in the order α^2 . Note that similar bounds (sans the last α^3 term) have been established for the simpler i.i.d. noise case in the federated setting (Khaled et al., 2020; Koloskova et al., 2020). Consequently, we achieve the same sample complexity results for the more general federated setting with Markov noise.

Remark. The bound in Theorem 6.1 holds only after $T > \max\{K + \tau_\alpha, 2\tau_\alpha\}$ and for all synchronization periods $K \geq 1$. At $K = 1$ the third term in the bound goes away, and we will be left only with the first order term, which is linearly decreasing with respect to the number of agents N , and the third order term $\tilde{\mathcal{O}}(\alpha^3)$. The last term, however, is not tight and can be further improved to be of the order $\mathcal{O}(\alpha^j), j > 3$. However, for that we need to assume larger τ_α , which means the bound only hold after a longer waiting time. In particular, by choosing $\tau_\alpha = \lceil r \log_\rho \alpha \rceil$, we can get $\tilde{\mathcal{O}}(\alpha^{2r-1})$ for the last term (see the proof of Lemma B.2).

Acknowledgment

This work was partially supported by NSF awards CCF-1944993, CCF-2045694, CNS-2112471, CMMI-2112533, EPCN-2144316, and an award from Raytheon technologies.

References

- Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., and Zhang, L. Solving Rubik’s cube with a robot hand. *Preprint arXiv:1910.07113*, 2019.
- Baird, L. Residual algorithms: Reinforcement learning with function approximation. In *Machine Learning Proceedings 1995*, pp. 30–37. Elsevier, 1995.
- Banach, S. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fund. math.*, 3(1):133–181, 1922.
- Beck, A. *First-order methods in optimization*, volume 25. SIAM, 2017.
- Beck, C. L. and Srikant, R. Error bounds for constant step-size Q -learning. *Syst. Control. Lett.*, 61:1203–1208, 2012a.
- Beck, C. L. and Srikant, R. Error bounds for constant step-size Q -learning. *Systems & control letters*, 61(12):1203–1208, 2012b.
- Beck, C. L. and Srikant, R. Improved upper bounds on the expected error in constant step-size Q -learning. In *2013 American Control Conference*, pp. 1926–1931. IEEE, 2013.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-dynamic programming*. Athena Scientific, 1996a.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-dynamic programming*. Athena Scientific, 1996b.
- Bertsekas, D. P., Bertsekas, D. P., Bertsekas, D. P., and Bertsekas, D. P. *Dynamic programming and optimal control*, volume 2. Athena scientific Belmont, MA, 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. PMLR, 2018.
- Borkar, V. S. *Stochastic approximation: a dynamical systems viewpoint*, volume 48. Springer, 2009.
- Borkar, V. S. and Meyn, S. P. The ODE method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38(2):447–469, 2000.
- Chen, T. and Giannakis, G. B. Bandit convex optimization for scalable and dynamic iot management. *IEEE Internet of Things Journal*, 6(1):1276–1286, 2018.
- Chen, T., Zhang, K., Giannakis, G. B., and Basar, T. Communication-efficient policy gradient methods for distributed reinforcement learning. *IEEE Transactions on Control of Network Systems*, 2021a.
- Chen, Z., Zhang, S., Doan, T. T., Maguluri, S. T., and Clarke, J.-P. Finite-sample analysis of nonlinear stochastic approximation with applications in reinforcement learning. *Under review by Automatica, Preprint arXiv:1905.11425*, 2019a.
- Chen, Z., Zhang, S., Doan, T. T., Maguluri, S. T., and Clarke, J.-P. Performance of Q -learning with linear function approximation: Stability and finite-time analysis. In *OptRL Workshop at NeurIPS 2019*, 2019b.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. Finite-sample analysis of stochastic approximation using smooth convex envelopes. *Under Review at Mathematics of Operations Research, Preprint arXiv:2002.00874*, 2020a.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. Finite-sample analysis of stochastic approximation using smooth convex envelopes. In *Advances in Neural Information Processing Systems*, 2020b. URL <https://proceedings.neurips.cc/paper/2020/file/5d44ee6f2c3f71b73125876103c8f6c4-Paper.pdf>.
- Chen, Z., Khodadadian, S., and Maguluri, S. T. Finite-Sample Analysis of Off-Policy Natural Actor-Critic with Linear Function Approximation. *Preprint arXiv:2105.12540*, 2021b. Submitted to NeurIPS 2021.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. A Lyapunov Theory for Finite-Sample Guarantees of Asynchronous Q -Learning and TD-Learning Variants. *Under review by JMLR, Preprint arXiv:2102.01567*, 2021c.
- Chen, Z., Maguluri, S. T., Shakkottai, S., and Shanmugam, K. Finite-Sample Analysis of Off-Policy TD-Learning via Generalized Bellman Operators. *Preprint arXiv:2106.12729*, 2021d.
- Chen, Z., Zhou, Y., and Chen, R. Multi-agent off-policy td learning: Finite-time analysis with near-optimal sample complexity and communication complexity. *arXiv preprint arXiv:2103.13147*, 2021e.
- Chen, Z., Zhou, Y., Chen, R., and Zou, S. Sample and communication-efficient decentralized actor-critic algorithms with finite-time analysis. *arXiv preprint arXiv:2109.03699*, 2021f.

- Dalal, G., Szörényi, B., Thoppe, G., and Mannor, S. Finite sample analysis for TD(0) with function approximation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Doan, T., Maguluri, S., and Romberg, J. Finite-Time Analysis of Distributed TD(0) with Linear Function Approximation on Multi-Agent Reinforcement Learning. In *International Conference on Machine Learning*, pp. 1626–1635, 2019.
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I., and Abbeel, P. RI^2 : Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. In *International Conference on Machine Learning*, pp. 1406–1415, 2018.
- Even-Dar, E. and Mansour, Y. Learning Rates for Q -Learning. *J. Mach. Learn. Res.*, 5:1–25, 2004. ISSN 1532-4435.
- Foerster, J. N., Assael, Y. M., De Freitas, N., and Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. *arXiv preprint arXiv:1605.06676*, 2016.
- Gottesman, O., Johansson, F., Komorowski, M., Faisal, A., Sontag, D., Doshi-Velez, F., and Celi, L. A. Guidelines for reinforcement learning in healthcare. *Nature medicine*, 25(1):16–18, 2019.
- Gottesman, O., Futoma, J., Liu, Y., Parbhoo, S., Celi, L., Brunskill, E., and Doshi-Velez, F. Interpretable off-policy evaluation in reinforcement learning by highlighting influential transitions. In *International Conference on Machine Learning*, pp. 3658–3667. PMLR, 2020.
- Gu, S., Holly, E., Lillicrap, T., and Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 3389–3396. IEEE, 2017.
- Hu, B. and Syed, U. A. Characterizing the exact behaviors of temporal difference learning algorithms using markov jump linear system theory. *arXiv preprint arXiv:1906.06781*, 2019.
- Islam, R., Henderson, P., Gomrokchi, M., and Precup, D. Reproducibility of benchmarked deep reinforcement learning tasks for continuous control. *arXiv preprint arXiv:1708.04133*, 2017.
- Jaakkola, T., Jordan, M. I., and Singh, S. P. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in neural information processing systems*, pp. 703–710, 1994.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. Advances and open problems in federated learning. *Preprint arXiv:1912.04977*, 2019.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., and Levine, S. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *Preprint arXiv:1806.10293*, 2018.
- 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., Chen, Z., and Maguluri, S. T. Finite-Sample Analysis of Off-Policy Natural Actor-Critic Algorithm. In *International Conference on Machine Learning*, 2021.
- Kober, J., Bagnell, J. A., and Peters, J. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- 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.
- Konda, V. R. and Tsitsiklis, J. N. Actor-critic algorithms. In *Advances in neural information processing systems*, pp. 1008–1014, 2000.
- 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.
- Lakshminarayanan, C. and Szepesvari, C. Linear stochastic approximation: How far does constant step-size and iterate averaging go? In *International Conference on Artificial Intelligence and Statistics*, pp. 1347–1355, 2018.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *Preprint arXiv:2005.01643*, 2020.
- Li, G., Wei, Y., Chi, Y., Gu, Y., and Chen, Y. Sample Complexity of Asynchronous Q -Learning: Sharper Analysis and Variance Reduction. *Advances in neural information processing systems*, 2020.

- Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019.
- Liang, X., Liu, Y., Chen, T., Liu, M., and Yang, Q. Federated transfer reinforcement learning for autonomous driving. *arXiv preprint arXiv:1910.06001*, 2019.
- 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, B., Liu, J., Ghavamzadeh, M., Mahadevan, S., and Petrik, M. Finite-sample analysis of proximal gradient TD algorithms. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pp. 504–513, 2015.
- Liu, B., Wang, L., and Liu, M. Lifelong federated reinforcement learning: a learning architecture for navigation in cloud robotic systems. *IEEE Robotics and Automation Letters*, 4(4):4555–4562, 2019.
- Liu, R. and Olshesky, A. Distributed td (0) with almost no communication. *arXiv preprint arXiv:2104.07855*, 2021.
- Liu, Y., Gottesman, O., Raghu, A., Komorowski, M., Faisal, A. A., Doshi-Velez, F., and Brunskill, E. Representation Balancing MDPs for Off-policy Policy Evaluation. *Advances in Neural Information Processing Systems*, 31: 2644–2653, 2018.
- Maddern, W., Pascoe, G., Linegar, C., and Newman, P. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017.
- Maei, H. R. Convergent actor-critic algorithms under off-policy training and function approximation. *Preprint arXiv:1802.07842*, 2018.
- 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.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937, 2016.
- Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., and Srivastava, G. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.
- Nadiger, C., Kumar, A., and Abdelhak, S. Federated reinforcement learning for fast personalization. In *2019 IEEE Second International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*, pp. 123–127. IEEE, 2019.
- Nair, A., Srinivasan, P., Blackwell, S., Alcicek, C., Fearon, R., De Maria, A., Panneershelvam, V., Suleyman, M., Beattie, C., Petersen, S., et al. Massively parallel methods for deep reinforcement learning. *arXiv preprint arXiv:1507.04296*, 2015.
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., and Poor, H. V. Federated learning for internet of things: A comprehensive survey. *arXiv preprint arXiv:2104.07914*, 2021.
- Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Qi, J., Zhou, Q., Lei, L., and Zheng, K. Federated reinforcement learning: Techniques, applications, and open challenges. *arXiv preprint arXiv:2108.11887*, 2021.
- Qu, G. and Wierman, A. Finite-Time Analysis of Asynchronous Stochastic Approximation and Q -Learning. In *Conference on Learning Theory*, pp. 3185–3205. PMLR, 2020.
- Qu, Z., Lin, K., Kalagnanam, J., Li, Z., Zhou, J., and Zhou, Z. Federated learning’s blessing: Fedavg has linear speedup. *arXiv preprint arXiv:2007.05690*, 2020.
- Rakhlin, A., Shamir, O., and Sridharan, K. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning*, pp. 1571–1578, 2012.
- Recht, B., Re, C., Wright, S., and Niu, F. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/218a0aefd1d1a4be65601cc6ddc1520e-Paper.pdf>.
- Ren, J., Wang, H., Hou, T., Zheng, S., and Tang, C. Federated learning-based computation offloading optimization in edge computing-supported internet of things. *IEEE Access*, 7:69194–69201, 2019.
- Shah, D. and Xie, Q. Q -learning with nearest neighbors. In *Advances in Neural Information Processing Systems*, pp. 3111–3121, 2018.

- Shalev-Shwartz, S., Shammah, S., and Shashua, A. Safe, multi-agent, reinforcement learning for autonomous driving. *Preprint arXiv:1610.03295*, 2016.
- Shalev-Shwartz, S. et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- Shao, K., Tang, Z., Zhu, Y., Li, N., and Zhao, D. A survey of deep reinforcement learning in video games. *arXiv preprint arXiv:1912.10944*, 2019.
- Shen, H., Zhang, K., Hong, M., and Chen, T. Asynchronous advantage actor critic: Non-asymptotic analysis and linear speedup. *arXiv preprint arXiv:2012.15511*, 2020.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484, 2016.
- Spiridonoff, A., Olshevsky, A., and Paschalidis, I. C. Communication-efficient sgd: From local sgd to one-shot averaging. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- 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.
- Stich, S. U. Local sgd converges fast and communicates little. In *International Conference on Learning Representations*, 2018.
- Sun, J., Wang, G., Giannakis, G. B., Yang, Q., and Yang, Z. Finite-time analysis of decentralized temporal-difference learning with linear function approximation. In *International Conference on Artificial Intelligence and Statistics*, pp. 4485–4495. PMLR, 2020.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Tadić, V. On the convergence of temporal-difference learning with linear function approximation. *Machine learning*, 42(3):241–267, 2001.
- Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., and Zhou, Y. A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, pp. 1–11, 2019.
- Tsitsiklis, J. N. Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3):185–202, 1994.
- Tsitsiklis, J. N. and Van Roy, B. Analysis of temporal-difference learning with function approximation. In *Advances in neural information processing systems*, pp. 1075–1081, 1997.
- Wai, H.-T. On the convergence of consensus algorithms with markovian noise and gradient bias. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 4897–4902. IEEE, 2020.
- Wainwright, M. J. Stochastic approximation with cone-contractive operators: Sharp ℓ_∞ -bounds for Q-learning. *Preprint arXiv:1905.06265*, 2019.
- Wang, G., Lu, S., Giannakis, G., Tesauro, G., and Sun, J. Decentralized TD Tracking with Linear Function Approximation and its Finite-Time Analysis. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 13762–13772. Curran Associates, Inc., 2020a. URL <https://proceedings.neurips.cc/paper/2020/file/9ec51f6eb240fb631a35864e13737bca-Paper.pdf>.
- Wang, J. and Joshi, G. Cooperative sgd: A unified framework for the design and analysis of local-update sgd algorithms. *Journal of Machine Learning Research*, 22(213):1–50, 2021.
- Wang, X., Wang, C., Li, X., Leung, V. C., and Taleb, T. Federated deep reinforcement learning for internet of things with decentralized cooperative edge caching. *IEEE Internet of Things Journal*, 7(10):9441–9455, 2020b.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Xu, M., Peng, J., Gupta, B., Kang, J., Xiong, Z., Li, Z., and Abd El-Latif, A. A. Multi-agent federated reinforcement learning for secure incentive mechanism in intelligent cyber-physical systems. *IEEE Internet of Things Journal*, 2021.
- Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., and Yu, H. Federated learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 13(3):1–207, 2019.
- Yun, W. J., Park, S., Kim, J., Shin, M., Jung, S., Mohaisen, A., and Kim, J.-H. Cooperative multi-agent deep reinforcement learning for reliable surveillance via autonomous multi-uav control. *IEEE Transactions on Industrial Informatics*, 2022.

- Yurtsever, E., Lambert, J., Carballo, A., and Takeda, K. A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access*, 8:58443–58469, 2020.
- Zeng, S., Doan, T. T., and Romberg, J. Finite-time analysis of decentralized stochastic approximation with applications in multi-agent and multi-task learning. *arXiv preprint arXiv:2010.15088*, 2020.
- Zhang, S., Liu, B., Yao, H., and Whiteson, S. Provably convergent two-timescale off-policy actor-critic with function approximation. In *International Conference on Machine Learning*, pp. 11204–11213. PMLR, 2020.
- Zhang, S. Q., Lin, J., and Zhang, Q. A multi-agent reinforcement learning approach for efficient client selection in federated learning. *arXiv preprint arXiv:2201.02932*, 2022.
- Zhao, L., Ran, Y., Wang, H., Wang, J., and Luo, J. Towards cooperative caching for vehicular networks with multi-level federated reinforcement learning. In *ICC 2021-IEEE International Conference on Communications*, pp. 1–6. IEEE, 2021.
- Zhuo, H. H., Feng, W., Lin, Y., Xu, Q., and Yang, Q. Federated deep reinforcement learning. *arXiv preprint arXiv:1901.08277*, 2019.

Appendices

The appendices are organized as follows. In section A we discuss the lower bound on the convergence of general stochastic approximation. In Section B we derive the convergence bound of FedSAM algorithm. Next, we employ the results in Section B to derive the convergence bounds of federated TD-learning in Section C and federated Q -learning in Section D.

A. Lower Bound on the Convergence of General Stochastic Approximation

In this section we discuss the convergence of general stochastic approximation. In this discussion we provide a simple stochastic approximation with iid noise which can give insight into the general convergence bound in (3). In particular, we show that the convergence bound in (3) is tight and cannot be improved.

Consider a one dimensional random variable X with zero mean $\mathbb{E}[X] = 0$ and bounded variance $\mathbb{E}[X^2] = \sigma^2$. Consider the following update

$$x_{t+1} = x_t + \alpha(X_t - x_t), \quad t \geq 0, \quad (9)$$

where we start with some fixed deterministic x_0 and X_t is a iid sample of the random variable X . It is easy to see that the update (9) is a special case of the update of the general stochastic approximation with the fixed point $x^* = 0$.

By expanding the update (9), we have

$$x_t = (1 - \alpha)^t x_0 + \alpha \sum_{k=0}^{t-1} (1 - \alpha)^{t-k-1} X_k.$$

Hence, we have

$$x_t^2 = (1 - \alpha)^{2t} x_0^2 + \left(\alpha \sum_{k=0}^{t-1} (1 - \alpha)^{t-k-1} X_k \right)^2 + 2\alpha(1 - \alpha)^t x_0 \sum_{k=0}^{t-1} (1 - \alpha)^{t-k-1} X_k.$$

Taking expectation on both sides, and using the zero mean property of X_k , we have

$$\begin{aligned} \mathbb{E}[x_t^2] &= (1 - \alpha)^{2t} x_0^2 + \alpha^2 \mathbb{E} \left(\sum_{k=0}^{t-1} (1 - \alpha)^{t-k-1} X_k \right)^2 \\ &= (1 - \alpha)^{2t} x_0^2 + \alpha^2 \mathbb{E} \left(\sum_{k=0}^{t-1} (1 - \alpha)^{2(t-k-1)} X_k^2 + \sum_{k,k'=0, k \neq k'}^{t-1} (1 - \alpha)^{2t-k-k'-2} X_k X_{k'} \right) \\ &= (1 - \alpha)^{2t} x_0^2 + \alpha^2 \left(\sum_{k=0}^{t-1} (1 - \alpha)^{2(t-k-1)} \sigma^2 + \sum_{k,k'=0, k \neq k'}^{t-1} (1 - \alpha)^{2t-k-k'-2} \mathbb{E}[X_k] \mathbb{E}[X_{k'}] \right) \quad (\text{iid property}) \\ &= (1 - \alpha)^{2t} x_0^2 + \alpha^2 \left(\sum_{k=0}^{t-1} (1 - \alpha)^{2(t-k-1)} \sigma^2 \right) \quad (\text{zero mean}) \\ &= (1 - \alpha)^{2t} x_0^2 + \alpha^2 \sigma^2 \frac{1 - (1 - \alpha)^{2t}}{1 - (1 - \alpha)^2} \\ &= x_0^2 (1 - \alpha)^{2t} + \sigma^2 \frac{1 - (1 - \alpha)^{2t}}{2 - \alpha} \alpha \\ &= \underbrace{\left(x_0^2 - \frac{\alpha \sigma^2}{2 - \alpha} \right)}_{T_1: \text{bias}} (1 - \alpha)^{2t} + \underbrace{\frac{\sigma^2}{2 - \alpha} \alpha}_{T_2: \text{variance}}. \quad (10) \end{aligned}$$

It is clear that (10) has the same form as the bound in (3) with T_1 as the geometric term which converges to zero as $t \rightarrow \infty$, and T_2 term proportional to the step size α . In addition, note that in the above derivation, we did not use any inequality, and hence the bounds in (10) as well as (3) are tight.

B. Analysis of Federated Stochastic Approximation

First, we restate Theorem 6.1 with explicit expressions of the different constants.

Theorem B.1. Consider the FedSAM Algorithm 4 with $c_{FSAM} = 1 - \frac{\alpha\varphi_2}{2}$ (φ_2 is defined in (14)), and suppose Assumptions 6.1, 6.2, 6.3, 6.4 are satisfied. Consider small enough step size α which satisfies the assumptions in (21), (23), (32), (35). Furthermore, denote $\tau_\alpha = \lceil 2 \log_\rho \alpha \rceil$, and take large enough T such that $T > \max\{K + \tau_\alpha, 2\tau_\alpha\}$. Then, the output of the FedSAM Algorithm 4, $\hat{\theta}_T \triangleq \frac{1}{N} \sum_{i=1}^N \theta_T^i$, satisfies

$$\mathbb{E}[\|\hat{\theta}_T\|_c^2] \leq C_1 \frac{1}{\alpha} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau_\alpha+1} + C_2 \frac{\alpha\tau_\alpha^2}{N} + C_3(K-1)\alpha^2\tau_\alpha + C_4\alpha^3\tau_\alpha^2, \quad (11)$$

where $C_1 = 16u_{cm}^2 M_0(\log_\rho \frac{1}{e} + \frac{1}{\varphi_2})$, $M_0 = \frac{1}{l_{cm}^2} \left(\frac{1}{C_1^2} \left(B + (A_2 + 1) \left(\|\theta_0\|_c + \frac{B}{2C_1}\right)\right)^2 + \|\theta_0\|_c^2\right)$, $C_2 = 8u_{cm}^2 (C_8 + \frac{1}{2} + C_{12})/\varphi_2$, and $C_3 = 80\tilde{B}^2 C_{17} u_{cm}^2 \left(1 + \frac{4m\rho}{B(1-\rho)}\right)/\varphi_2$ and $C_4 = 8u_{cm}^2 (C_7 + C_{11} + 0.5C_3^2 C_9^2 + C_3 C_{10} + 2C_1 C_3 C_{10} + 3A_1 C_3 + C_{13})/\varphi_2$. Here the constants $u_{cm}, l_{cm}, \varphi_2, C_1, C_3, C_7, C_8, C_9, C_{10}, C_{11}, C_{12}, C_{13}, C_{17}, \tilde{B}$ are problem dependent constants which are defined in the following proposition and lemmas.

Next, we characterize the sample complexity of the FedSAM Algorithm 4, where we establish a linear speedup in the convergence of the algorithm.

Corollary B.1.1. Consider FedSAM Algorithm 4 with fixed number of iterations T and step size $\alpha = \frac{8 \log(NT)}{\varphi_2 T}$. Suppose T is large enough, such that α satisfies the requirements of the step size in Theorem B.1 and $T > 4\tau_\alpha$. Also take $K = T/N$. Then we have $\mathbb{E}[\|\hat{\theta}_T\|_c^2] \leq \epsilon$ after $T = \mathcal{O}\left(\frac{1}{N\epsilon}\right)$ iterations.

Corollary B.1.1 establishes the sample and communication complexity of FedSAM Algorithm 4. The $\mathcal{O}(1/(N\epsilon))$ sample complexity shows the linear speedup with increasing N . Another aspect of the cost is the number of communications required between the agents and the central server. According to Corollary B.1.1, we need $T/N = \tilde{\mathcal{O}}(N)$ rounds of communications in order to reach an ϵ -optimal solution. Hence, even in the presence of Markov noise, the required number of communications is independent of the desired final accuracy ϵ , and grows linearly with the number of agents. Our result generalizes the existing result achieved for the simpler i.i.d. noise case in (Khaled et al., 2020; Spiridonoff et al., 2021).

In the following sections, we discuss the proof of Theorem B.1. In Section B.1, we introduce some notations and preliminary results to facilitate our Lyapunov-function based analysis. Next, in Section B.2, we state some primary propositions, which are then used to prove Theorem B.1. In Sections B.3, B.4, B.5, we prove the aforementioned propositions. Along the way, we state several intermediate results, which are stated and proved in Sections B.6 and B.7, respectively.

Throughout the appendix we have several sets of constants. The constants $C_i, i = 1, \dots, 17$ are problem dependent constants which we define recursively. The final constants which appears in the resulting bound in Theorem B.1 are shown as $C_i, i = 1, \dots, 4$. Finally, the constant c is used in the sampling of the time step \hat{T} .

B.1. Preliminaries

We define the following notations:

- $\theta_t \triangleq \frac{1}{N} \sum_{i=1}^N \theta_t^i$: virtual sequence of average (across agents) parameter.
- $\Theta_t = \{\theta_t^1, \dots, \theta_t^N\}$: set of local parameters at individual nodes.
- $\mathbf{Y}_t = \{\mathbf{y}_t^1, \dots, \mathbf{y}_t^N\}$: Markov chains at individual nodes.
- μ^i : the stationary distribution of \mathbf{y}_t^i as $t \rightarrow \infty$.
- $\mathbf{G}(\Theta_t, \mathbf{Y}_t) \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i)$: average of the noisy local operators at the individual local parameters.
- $\mathbf{G}(\theta_t, \mathbf{Y}_t) \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{G}^i(\theta_t, \mathbf{y}_t^i)$: average of the noisy local operators at the average parameter.
- $\mathbf{b}(\mathbf{Y}_t) \triangleq \frac{1}{N} \sum_{i=1}^N \mathbf{b}^i(\mathbf{y}_t^i)$: average of Markovian noise.

- $\bar{\mathbf{G}}(\Theta_t) \triangleq \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{G}}^i(\theta_t^i)$: average of expected operators evaluated at the local parameter.
- $\bar{\mathbf{G}}(\theta_t) \triangleq \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{G}}^i(\theta_t)$: average of expected operator evaluated at the average parameter.
- $\theta^* = \mathbf{0}$ is the unique fixed point satisfying $\theta^* = \bar{\mathbf{G}}^i(\theta^*)$ for all $i = 1, \dots, N$. Note that this follows directly from the assumptions. In particular, by Assumption 6.2, $\bar{\mathbf{G}}^i(\cdot)$ is a contraction, and hence by Banach fixed point theorem (Banach, 1922), there exist a unique fixed point of this operator. Furthermore, by Assumption 6.3, we have

$$\|\bar{\mathbf{G}}^i(\mathbf{0})\|_c = \|\mathbb{E}_{\mathbf{y} \sim \mu^i}[\mathbf{G}^i(\mathbf{0}, \mathbf{y})]\|_c \leq \mathbb{E}_{\mathbf{y} \sim \mu^i} \|\mathbf{G}^i(\mathbf{0}, \mathbf{y})\|_c \leq A_2 \mathbb{E}_{\mathbf{y} \sim \mu^i} [\|\mathbf{0}\|_c] = 0.$$

Hence $\bar{\mathbf{G}}^i(\mathbf{0}) = \mathbf{0}$ and $\mathbf{0}$ is the unique fixed point of the operator $\bar{\mathbf{G}}^i(\cdot)$.

- $\Delta_t^i = \|\theta_t^i - \theta_t^i\|_c$, $\Delta_t = \frac{1}{N} \sum_{i=1}^N \Delta_t^i$, $\Omega_t = \frac{1}{N} \sum_{i=1}^N (\Delta_t^i)^2$: measures of synchronization error.

Throughout this proof we assume $\|\cdot\|_c$ as some given norm. $\mathbb{E}_t[\cdot] \triangleq \mathbb{E}[\cdot | \mathcal{F}_t]$, where \mathcal{F}_t is the sigma-algebra generated by $\{\theta_r^i\}_{r=1, \dots, t}^{i=1, \dots, N}$. Unless specified otherwise, $\|\cdot\|$ denotes the Euclidean norm.

Generalized Moreau Envelope: Consider the norm $\|\cdot\|_c$ which appears in Assumptions 6.1-6.3. Square of this norm need not be smooth. Inspired by (Chen et al., 2020a), we use the Generalized Moreau Envelope as a Lyapunov function for the analysis of the convergence of Algorithm 4. The Generalized Moreau Envelope of $f(\cdot)$ with respect to $g(\cdot)$, for $\psi > 0$, is defined as

$$M_f^{\psi, g}(\mathbf{x}) = \min_{\mathbf{u} \in \mathbb{R}^d} \left\{ f(\mathbf{u}) + \frac{1}{\psi} g(\mathbf{x} - \mathbf{u}) \right\}. \quad (12)$$

Let $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_c^2$ and $g(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_s^2$, which is L -smooth with respect to $\|\cdot\|_s$ norm. For this choice of f, g , $M_f^{\psi, g}(\cdot)$ is essentially a smooth approximation to f , which is henceforth denoted with the simpler notation $M(\cdot)$. Also, due to the equivalence of norms, there exist $l_{cs}, u_{cs} > 0$ such that

$$l_{cs} \|\cdot\|_s \leq \|\cdot\|_c \leq u_{cs} \|\cdot\|_s. \quad (13)$$

We next summarize the properties of $M(\cdot)$ in the following proposition, which were established in (Chen et al., 2020a).

Proposition B.1 ((Chen et al., 2020b)). *The function $M(\cdot)$ satisfies the following properties.*

- (1) $M(\cdot)$ is convex, and $\frac{L}{\psi}$ -smooth with respect to $\|\cdot\|_s$. That is, $M(\mathbf{y}) \leq M(\mathbf{x}) + \langle \nabla M(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2\psi} \|\mathbf{x} - \mathbf{y}\|_s^2$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$.
- (2) There exists a norm, denoted by $\|\cdot\|_m$, such that $M(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_m^2$.
- (3) Let $\ell_{cm} = (1 + \psi \ell_{cs}^2)^{1/2}$ and $u_{cm} = (1 + \psi u_{cs}^2)^{1/2}$. Then it holds that $\ell_{cm} \|\cdot\|_m \leq \|\cdot\|_c \leq u_{cm} \|\cdot\|_m$.

By Proposition B.1, we can use $M(\cdot)$ as a smooth surrogate for $\frac{1}{2} \|\cdot\|_c^2$. Furthermore, we denote

$$\varphi_1 = \frac{1 + \psi u_{cs}^2}{1 + \psi \ell_{cs}^2}, \quad \varphi_2 = 1 - \gamma_c \varphi_1^{1/2}, \quad \text{and} \quad \varphi_3 = \frac{114L(1 + \psi u_{cs}^2)}{\psi \ell_{cs}^2}. \quad (14)$$

Note that by choosing $\psi > 0$ small enough, we can ensure $\varphi_2 \in (0, 1)$.

B.2. Proof of Theorem B.1

In this section, first we state three key results (Propositions B.2, B.3, B.4). These are then used to prove Theorem B.1. The first step of the proof is to characterize the one-step drift of the Lyapunov function $M(\cdot)$, with the parameters generated by the FedSAM Algorithm 4, which is formally stated in the following proposition.

Proposition B.2 (One-step drift - I). *Consider the update of the FedSAM Algorithm 4. Suppose the assumptions 6.1, 6.2, 6.3, and 6.4 are satisfied. Consider $\tau = \lceil 2 \log_\rho \alpha \rceil$ and $t \geq 2\tau$, we have*

$$\mathbb{E}_{t-2\tau} [M(\theta_{t+1})]$$

$$\leq \left(1 - \alpha \left[2\varphi_2 - \zeta_1^2 - \zeta_2^2 \frac{2LA_2}{\psi l_{cs}^2} - \zeta_3^2 \frac{L(A_1+1)}{\psi l_{cs}^2} - \zeta_4^2\right] + \alpha^2 \frac{6Lu_{cm}^2}{l_{cs}^2\psi} \left[m_1 + \frac{(A_2+1)^2}{2}\right]\right) \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)]$$

$$+ \alpha^2 \underbrace{\left(\frac{m_2 L}{\psi l_{cs}^2}\right)}_{C_3} \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c] \quad (15)$$

$$+ \underbrace{\frac{L}{\psi l_{cs}^2} \left(\frac{L}{2\psi l_{cs}^2} + \alpha A_2 \left(1 + \frac{u_{cm}^2}{\zeta_2^2}\right) + \alpha(A_1+1) \left(\frac{3u_{cm}^2}{2\zeta_3^2} + 2\right) + 3m_1\alpha^2\right)}_{C_4} \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2] \quad (16)$$

$$+ \alpha \underbrace{\frac{L}{\psi l_{cs}^2} \left(\left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2}\right)(A_1+1) + A_2 + \frac{Lu_{cm}^2}{2l_{cs}^2\zeta_4^2\psi} + \frac{3A_1^2\alpha^2}{2}\right)}_{C_5} \mathbb{E}_{t-2\tau} [\Omega_t] \quad (17)$$

$$+ \alpha \underbrace{\frac{L}{2l_{cs}^2\psi} \left(\left(\frac{3u_{cm}^2}{\zeta_3^2} + 3\right)(A_1+1) + m_1\alpha\right)}_{C_6} \mathbb{E}_{t-2\tau} [\Omega_{t-\tau}] \quad (18)$$

$$+ \underbrace{\frac{m_2^2 u_{cm}^2 L^2}{2\zeta_1^2 l_{cs}^4 \psi^2}}_{C_7} \alpha \alpha^4 \quad (19)$$

$$+ \underbrace{\frac{1}{2} \left(1 + \frac{3L}{\psi l_{cs}^2}\right)}_{C_8} \alpha^2 \mathbb{E}_{t-2\tau} [\|\mathbf{b}(\mathbf{Y}_t)\|_c^2], \quad (20)$$

where $\zeta_1, \zeta_2, \zeta_3,$ and ζ_4 are arbitrary positive constants.

Proof. The proof of Proposition B.2 is presented in full detail in Section B.3. \square

Before discussing the bound in its full generality, we discuss a few special cases.

- Perfect synchronization ($K = 1$) with i.i.d. noise: since $\Omega_t = 0$ for all t , $\tau = 0$ (independence across time), and $C_7 = 0$ (see Lemmas B.2 and B.5 in Section B.6), the terms (15), (16), (17), (18), (19) will not appear in the bound, which is the form we get for centralized systems with i.i.d. noise (Rakhlin et al., 2012).
- Infrequent synchronization ($K > 1$) with i.i.d. noise: $\tau = 0$, and $C_7 = 0$, the terms (15), (16), (18), (19) will not appear in the bound, which is the form we get for federated stochastic optimization with i.i.d. noise (Khaled et al., 2020).
- Perfect synchronization ($K = 1$) with Markov noise: since $\Omega_t = 0$ for all t , the bound in Proposition B.2 generalizes the results in (Srikant & Ying, 2019; Chen et al., 2021c).

Next, we substitute the bound on $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c$ (Lemma B.6) and $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2$ (Lemma B.7) to further bound the one-step drift. Establishing a tight bound for these two quantities are essential to ensure linear speedup.

Proposition B.3 (One-step drift - II). *Consider the update of the FedSAM Algorithm 4. Suppose Assumptions 6.1, 6.2, 6.3, and 6.4 are satisfied. Define $C_{15}(\tau) = \left(\frac{6m_1 Lu_{cm}^2}{l_{cs}^2\psi} + \frac{6L(A_2+1)^2 u_{cm}^2}{2\psi l_{cs}^2} + 144(A_2^2 + 1)C_4 u_{cm}^2\right) \tau^2$. For*

$$\alpha \leq \min \left\{ \frac{1}{360\tau(A_2^2 + 1)}, \frac{\varphi_2}{2C_{15}(\tau)} \right\}, \quad (21)$$

$\tau = \lceil 2 \log_\rho \alpha \rceil$, and $t \geq 2\tau$, we have

$$\mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_{t+1})] \leq (1 - \alpha\varphi_2) \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + C_{14}(\tau)\alpha^4 + C_{16}(\tau)\frac{\alpha^2}{N} + \alpha C_{17} \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k],$$

where $C_{14}(\tau) = \left(C_7 + C_{11} + 0.5C_3^2C_9^2 + C_3C_{10} + 2C_1C_3C_{10} + 3A_1C_3 + C_{13} + C_8\frac{u_{cD}^2}{l_{cD}^2}2m_2^2\alpha^2 \right) \tau^2$, $C_{16}(\tau) = \left(C_8\frac{u_{cD}^2B^2}{l_{cD}^2} + \frac{1}{2} + C_{12} \right) \tau^2$ and $C_{17} = (3A_1C_3 + 8A_1^2C_4 + C_5 + C_6)$. Here we define $C_9 = \frac{8u_{cD}B}{l_{cD}}$, $C_{10} = \frac{8m_2u_{cD}}{l_{cD}(1-\rho)}$, $C_{11} = \frac{8C_1^2C_3^2u_{cm}^2}{C_6^2}$, $C_{12} = \frac{8C_4u_{cD}^2B^2}{l_{cD}^2}$, $C_{13} = \frac{14C_4u_{cD}^2m_2^2}{l_{cD}^2(1-\rho^2)}$.

Proof. The proof of Proposition B.3 is presented in full detail in Section B.4. \square

Conditional expectation $\mathbb{E}_{t-2\tau}[\cdot]$. The conditional expectation $\mathbb{E}_{t-2\tau}[\cdot]$ used in Proposition B.3 is essential when dealing with Markovian noise. The idea of using conditional expectation to deal with Markovian noise is not novel per se. In the previous work (Bhandari et al., 2018; Srikant & Ying, 2019; Chen et al., 2021c), conditioning on $t - \tau$ is sufficient to establish the convergence results. Due to the mixing property (Assumption 6.1), the Markov chain geometrically converges to its stationary distribution. Therefore, choosing “large enough” τ , and conditioning on $t - \tau$, one can ensure that the Markov chain at time t is “almost in steady state.” However, in federated setting, conditioning on $t - \tau$ results in bounds that are too loose. In particular, consider the differences $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c$ and $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2$ in (15) and (16) respectively. In the centralized setting, as in (Bhandari et al., 2018; Srikant & Ying, 2019; Chen et al., 2021c), these terms can be bounded deterministically to yield $\asymp \alpha^2$ bound. However, in the federated setting, this crude bound does not result in linear speedup in N . In this work, to achieve a finer bound on $\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c$, we go τ steps further back in time. This ensures that the difference behaves almost like the difference of average of i.i.d. random variables, resulting in a tighter bound (see Lemma B.6). By exploiting the conditional expectation $\mathbb{E}_{t-2\tau}[\cdot]$, we derive a refined analysis to bound this term as $\mathcal{O}(\alpha^2/N + \alpha^4)$, which guarantees a linear speedup (see Lemmas B.6 and B.7).

Taking total expectation in Proposition B.3 (using tower property), we get

$$\mathbb{E}[M(\boldsymbol{\theta}_{t+1})] \leq (1 - \alpha\varphi_2) \mathbb{E}[M(\boldsymbol{\theta}_t)] + C_{14}(\tau)\alpha^4 + C_{16}(\tau)\frac{\alpha^2}{N} + \alpha C_{17} \sum_{k=t-\tau}^t \mathbb{E}[\Omega_k] \quad (22)$$

To understand the bound in Proposition B.3, consider the case of $K = 1$ (i.e. full synchronization). In this case we have $\Omega_i = 0$ for all i , and the bound in Proposition B.3 simplifies to $\mathbb{E}[M(\boldsymbol{\theta}_{t+1})] \leq (1 - \alpha\varphi_2) \mathbb{E}[M(\boldsymbol{\theta}_t)] + C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha^2}{N}$. This recursion is sufficient to achieve linear speedup. However, the bound in Proposition B.3 also include terms which are proportional to the error due to synchronization. In order to ensure convergence along with linear speedup, we need to further upper bound this term with terms which are of the order $\mathcal{O}(\alpha^3(K - 1))$ and $M(\boldsymbol{\theta}_i)$. The following lemma is the next important contribution of the paper where we establish such a bound for weighted sum of the synchronization error. Notice that the weights $\{w_t\}$ are carefully chosen to ensure the best rate of convergence for the overall algorithm.

Proposition B.4 (Synchronization Error). *Suppose $T > K + \tau$. For α such that*

$$\begin{aligned} \alpha^2 &\leq \min \left\{ \frac{1}{\varphi_2^2}, \frac{\ln(5/4)}{2(1 + \tilde{A}_1)(K - 1)^2} \right\}, \\ \alpha^2(\log_\rho(\alpha) + 1) &\leq \frac{v}{2K^2}, \\ \alpha^2(\log_\rho(\alpha) + 1)^3 \exp(\alpha\varphi_2(2\log_\rho\alpha + 1)) &\leq \frac{v}{4K} \exp\left(-\varphi_2\sqrt{\frac{\ln(5/4)}{2(1 + \tilde{A}_1)}}\right), \end{aligned} \quad (23)$$

where $v = \frac{\varphi_2}{80\tilde{A}_2C_{17}u_{cm}^2}$, the weighted consensus error satisfies

$$\frac{2C_{17}}{\varphi_2W_T} \sum_{t=2\tau}^T w_t \left[\sum_{\ell=t-\tau}^t \mathbb{E}\Omega_\ell \right] \leq \alpha^2\tilde{B}^2\frac{10C_{17}}{\varphi_2} \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) (\tau + 1)(K - 1) + \frac{1}{2W_T} \sum_{t=0}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t). \quad (24)$$

Proof. The proof is presented in Section B.5. \square

Finally, by incorporating the results in Propositions B.2, B.3, and B.4, we can establish the convergence of FedSAM in Theorem B.1.

Proof of Theorem B.1. Assume $w_0 = 1$, and consider the weights w_t generated by the recursion $w_t = w_{t-1} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{-1}$. Multiplying both sides of (22) by $\frac{2w_t}{\alpha\varphi_2}$, and rearranging the terms, we get

$$\begin{aligned} w_t \mathbb{E}M(\boldsymbol{\theta}_t) &\leq \frac{2w_t}{\alpha\varphi_2} \left(1 - \frac{\alpha\varphi_2}{2}\right) \mathbb{E}M(\boldsymbol{\theta}_t) - \frac{2w_t}{\alpha\varphi_2} \mathbb{E}M(\boldsymbol{\theta}_{t+1}) \\ &\quad + \frac{2w_t}{\alpha\varphi_2} \left(C_{14}(\tau)\alpha^4 + C_{16}(\tau)\frac{\alpha^2}{N} \right) + \frac{2w_t\alpha C_{17}}{\alpha\varphi_2} \sum_{\ell=t-\tau}^t \mathbb{E}[\Omega_\ell] \\ &= \frac{2}{\alpha\varphi_2} [w_{t-1} \mathbb{E}M(\boldsymbol{\theta}_t) - w_t \mathbb{E}M(\boldsymbol{\theta}_{t+1})] + \frac{2w_t}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) + \frac{2w_t C_{17}}{\varphi_2} \sum_{\ell=t-\tau}^t \mathbb{E}[\Omega_\ell], \end{aligned} \quad (25)$$

where we use $w_{t-1} = w_t \left(1 - \frac{\alpha\varphi_2}{2}\right)$. Summing (25) over $t = 2\tau$ to T (define $W_T = \sum_{t=2\tau}^T w_t$), we get

$$\begin{aligned} &\frac{1}{W_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \\ &\leq \frac{2}{\alpha\varphi_2 W_T} [w_{2\tau-1} \mathbb{E}M(\boldsymbol{\theta}_{2\tau}) - w_T \mathbb{E}M(\boldsymbol{\theta}_{T+1})] + \frac{2}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) \frac{1}{W_T} \sum_{t=2\tau}^T w_t \\ &\quad + \frac{1}{W_T} \sum_{t=2\tau}^T \frac{2w_t C_{17}}{\varphi_2} \left[\sum_{\ell=t-\tau}^t \Omega_\ell \right] \\ &\leq \frac{2w_{2\tau-1}}{\alpha\varphi_2 w_T} \mathbb{E}M(\boldsymbol{\theta}_{2\tau}) + \frac{2}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) + \frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t \Omega_\ell \quad (\because M(\boldsymbol{\theta}) \geq 0, W_T \geq w_T) \\ &= \frac{2}{\alpha\varphi_2} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} \mathbb{E}M(\boldsymbol{\theta}_{2\tau}) + \frac{2}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) + \frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \left[\sum_{\ell=t-\tau}^t \Omega_\ell \right]. \end{aligned} \quad (26)$$

Substituting the bound on $\frac{1}{W_T} \sum_{t=2\tau}^T w_t \left[\sum_{\ell=t-\tau}^t \Omega_\ell \right]$ from Proposition B.4 into (26), we get

$$\begin{aligned} \frac{1}{W_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) &\leq \frac{2}{\alpha\varphi_2} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} \mathbb{E}M(\boldsymbol{\theta}_{2\tau}) + \frac{2}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) \\ &\quad + \alpha^2 \tilde{B}^2 \frac{10C_{17}}{\varphi_2} \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) (\tau+1)(K-1) + \frac{1}{2W_T} \sum_{t=0}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \\ \Rightarrow \frac{1}{W_T} \sum_{t=2\tau}^T w_t \mathbb{E}[M(\boldsymbol{\theta}_t)] &\leq \frac{4}{\alpha\varphi_2} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} M_0 + \frac{4}{\varphi_2} \left(C_{14}(\tau)\alpha^3 + C_{16}(\tau)\frac{\alpha}{N} \right) \\ &\quad + \alpha^2 \tilde{B}^2 \frac{20C_{17}}{\varphi_2} \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) (\tau+1)(K-1) + 2\tau \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} M_0, \end{aligned} \quad (27)$$

where M_0 is a problem dependent constant and is defined in Lemma B.6. To simplify (27), we define $\bar{C}_{18}(\tau) = \tilde{B}^2 \frac{40C_{17}}{\varphi_2} \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) \tau$, $\bar{C}_{19}(\tau) = \frac{4}{\varphi_2} C_{16}(\tau)$. We have

$$\frac{1}{W_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \leq \left(\frac{4}{\alpha\varphi_2} + 2\tau \right) \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} M_0 + \frac{4C_{14}(\tau)}{\varphi_2} \alpha^3 + \bar{C}_{18}(\tau)(K-1)\alpha^2 + \bar{C}_{19}(\tau)\frac{\alpha}{N}. \quad (28)$$

Furthermore, define $\tilde{W}_T = \sum_{t=0}^T w_t \geq W_T$. By definition of \hat{T} , we have $\mathbb{E}[M(\boldsymbol{\theta}_{\hat{T}})] = \frac{1}{\tilde{W}_T} \sum_{t=0}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t)$, and hence

$$\mathbb{E}[M(\boldsymbol{\theta}_{\hat{T}})] = \frac{1}{\tilde{W}_T} \sum_{t=0}^{2\tau-1} w_t \mathbb{E}M(\boldsymbol{\theta}_t) + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t)$$

$$\begin{aligned}
 &\leq \frac{M_0}{\tilde{W}_T} \sum_{t=0}^{2\tau-1} w_t + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) && \text{(Lemma B.6)} \\
 &\leq \frac{2\tau M_0 w_{2\tau-1}}{\tilde{W}_T} + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) && (w_t \leq w_{t+1} \text{ for all } t \geq 0) \\
 &\leq \frac{2\tau M_0 w_{2\tau-1}}{w_T} + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \\
 &= 2\tau M_0 \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \\
 &\leq 2\tau M_0 \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} + \frac{1}{\tilde{W}_T} \sum_{t=2\tau}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t) \\
 &\leq \left(4\tau + \frac{4}{\alpha\varphi_2}\right) M_0 \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} + \frac{4C_{14}(\tau)}{\varphi_2} \alpha^3 + \bar{C}_{18}(\tau)(K-1)\alpha^2 + \bar{C}_{19}(\tau) \frac{\alpha}{N} && \text{(by (28))} \\
 &= \bar{C}_{20}(\tau) \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} + \frac{4C_{14}(\tau)}{\varphi_2} \alpha^3 + \bar{C}_{18}(\tau)(K-1)\alpha^2 + \bar{C}_{19}(\tau) \frac{\alpha}{N},
 \end{aligned}$$

where $\bar{C}_{20}(\alpha, \tau) = \left(4\tau M_0 + \frac{4M_0}{\alpha\varphi_2}\right)$. Furthermore, by Proposition B.1, we have $M(\boldsymbol{\theta}_{\hat{T}}) = \frac{1}{2} \|\boldsymbol{\theta}_{\hat{T}}\|_m^2 \geq \frac{1}{2u_{cm}^2} \|\boldsymbol{\theta}_{\hat{T}}\|_c^2$, and hence

$$\mathbb{E}[\|\boldsymbol{\theta}_{\hat{T}}\|_c^2] \leq \bar{C}_1(\alpha, \tau) \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} + \mathcal{C}_2(\tau) \frac{\alpha}{N} + \mathcal{C}_3(\tau)(K-1)\alpha^2 + \mathcal{C}_4(\tau)\alpha^3 \quad (29)$$

where $\bar{C}_1(\alpha, \tau) = 2u_{cm}^2 \bar{C}_{20}(\alpha, \tau)$, $\mathcal{C}_2(\tau) = 2u_{cm}^2 \bar{C}_{19}(\tau)$, $\mathcal{C}_3(\tau) = 2u_{cm}^2 \bar{C}_{18}(\tau)$ and $\mathcal{C}_4(\tau) = 2u_{cm}^2 \cdot \frac{4C_{14}(\tau)}{\varphi_2}$.

Finally, note that by definition of τ , we have $\tau = \lceil 2 \log_\rho \alpha \rceil \leq 1 + 2 \log_\rho \alpha = 1 + 2(\ln \alpha)(\log_\rho e) = 1 + (2 \log_\rho \frac{1}{e}) \ln \frac{1}{\alpha} \leq 1 + (2 \log_\rho \frac{1}{e}) \frac{1}{\alpha}$. Hence, we have $\bar{C}_1(\alpha, \tau) \leq 2u_{cm}^2 M_0 (4 + \frac{8}{\alpha} \log_\rho \frac{1}{e} + \frac{4}{\alpha\varphi_2}) \leq 16u_{cm}^2 M_0 (\log_\rho \frac{1}{e} + \frac{1}{\varphi_2}) \frac{1}{\alpha} = \mathcal{C}_1 \cdot \frac{1}{\alpha}$, where $\mathcal{C}_1 = 16u_{cm}^2 M_0 (\log_\rho \frac{1}{e} + \frac{1}{\varphi_2})$. Furthermore, we have $\mathcal{C}_2(\tau) = 2u_{cm}^2 \bar{C}_{19}(\tau) = 2u_{cm}^2 \frac{4}{\varphi_2} C_{16}(\tau) = \frac{8u_{cm}^2 (C_8 + \frac{1}{2} + C_{12})}{\varphi_2} \tau^2 \equiv \frac{8u_{cm}^2 (C_8 + \frac{1}{2} + C_{12})}{\varphi_2} \tau_\alpha^2 = \mathcal{C}_2 \tau_\alpha^2$, where we denote $\tau_\alpha \equiv \tau$ to emphasize the dependence of τ on α , and $\mathcal{C}_2 = \frac{8u_{cm}^2 (C_8 + \frac{1}{2} + C_{12})}{\varphi_2}$. Note that we have $\tau_\alpha = \mathcal{O}(\log(1/\alpha))$.

In addition, $\mathcal{C}_3(\tau) = 2u_{cm}^2 \bar{C}_{18}(\tau) = \frac{80\bar{B}^2 C_{17} u_{cm}^2 (1 + \frac{4m_2\rho}{B(1-\rho)})}{\varphi_2} \cdot \tau \equiv \mathcal{C}_3 \tau_\alpha$, where $\mathcal{C}_3 = \frac{80\bar{B}^2 C_{17} u_{cm}^2 (1 + \frac{4m_2\rho}{B(1-\rho)})}{\varphi_2}$. And lastly, $\mathcal{C}_4(\tau) = \frac{8u_{cm}^2 (C_7 + C_{11} + 0.5C_3^2 C_9^2 + C_3 C_{10} + 2C_1 C_3 C_{10} + 3A_1 C_3 + C_{13})}{\varphi_2} \tau^2 \equiv \mathcal{C}_4 \tau_\alpha^2$, where $\mathcal{C}_4 = 8u_{cm}^2 (C_7 + C_{11} + 0.5C_3^2 C_9^2 + C_3 C_{10} + 2C_1 C_3 C_{10} + 3A_1 C_3 + C_{13}) / \varphi_2$. \square

Next we will state the proof of Corollary B.1.1.

Proof of Corollary B.1.1. By this choice of step size, for large enough T , α will be small enough and can satisfy the requirements of the step size in (21), (23), (32), (35). Furthermore, the first term in (29) will be

$$\begin{aligned}
 \mathcal{C}_1 \frac{1}{\alpha} \left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-2\tau+1} &\leq \mathcal{C}_1 \frac{1}{\alpha} e^{-\frac{4\log(NT)}{T}(T-2\tau+1)} = \mathcal{C}_1 \frac{\varphi_2 T}{8\log(NT)} e^{-\log((NT)^4)(1 + \frac{1-2\tau}{T})} \\
 &= \frac{\mathcal{C}_1 \varphi_2 T}{8\log(NT)} \left(\frac{1}{N^4 T^4}\right)^{(1 + \frac{1-2\tau}{T})} \\
 &\leq \frac{\mathcal{C}_1 \varphi_2 T}{8\log(NT)} \left(\frac{1}{N^4 T^4}\right)^{0.5} && \text{(Assumption on } T) \\
 &= \frac{\mathcal{C}_1 \varphi_2 T}{8\log(NT)} \frac{1}{N^2 T^2}
 \end{aligned}$$

$$= \tilde{\mathcal{O}}\left(\frac{\mathcal{C}_1 \varphi_2}{NT}\right).$$

Furthermore, for the second term we have

$$\mathcal{C}_2 \frac{\alpha \tau_\alpha^2}{N} = \tilde{\mathcal{O}}\left(\frac{\mathcal{C}_2 / \varphi_2}{NT}\right).$$

Finally, for the third and the fourth terms we have

$$\mathcal{C}_3(K-1)\alpha^2 \tau_\alpha + \mathcal{C}_4 \alpha^3 \tau_\alpha^2 = \tilde{\mathcal{O}}\left(\left(\mathcal{C}_3 / \varphi_2^2 + \mathcal{C}_4 / \varphi_2^3\right) \frac{T}{N} \times \frac{\log^2(NT)}{T^2}\right) = \tilde{\mathcal{O}}\left(\frac{(\mathcal{C}_3 / \varphi_2^2 + \mathcal{C}_4 / \varphi_2^3)}{NT}\right).$$

Upper bounding (29) with ϵ , we get $\tilde{\mathcal{O}}\left(\frac{\mathcal{C}_1 \varphi_2 + \mathcal{C}_2 / \varphi_2 + \mathcal{C}_3 / \varphi_2^2 + \mathcal{C}_4 / \varphi_2^3}{NT}\right) \leq \epsilon$. Hence, we need to have $T = \tilde{\mathcal{O}}\left(\frac{\mathcal{C}_1 \varphi_2 + \mathcal{C}_2 / \varphi_2 + \mathcal{C}_3 / \varphi_2^2 + \mathcal{C}_4 / \varphi_2^3}{N\epsilon}\right)$ number of iterations to get to a ball around the optimum with radius ϵ . \square

B.3. Proof of Proposition B.2

The update of the virtual parameter sequence $\{\boldsymbol{\theta}_t\}$ can be written as follows

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha (\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)). \quad (30)$$

Using $\frac{p-1}{\psi}$ -smoothness of $M(\cdot)$ (Proposition B.1), we get

$$\begin{aligned} M(\boldsymbol{\theta}_{t+1}) &\leq M(\boldsymbol{\theta}_t) + \langle \nabla M(\boldsymbol{\theta}_t), \boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t \rangle + \frac{L}{2\psi} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t\|_s^2 && \text{(Smoothness of } M(\cdot)\text{)} \\ &= M(\boldsymbol{\theta}_t) + \alpha \langle \nabla M(\boldsymbol{\theta}_t), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t) \rangle + \frac{L\alpha^2}{2\psi} \|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)\|_s^2 \\ &= M(\boldsymbol{\theta}_t) + \alpha \underbrace{\langle \nabla M(\boldsymbol{\theta}_t), \bar{\mathbf{G}}(\boldsymbol{\theta}_t) - \boldsymbol{\theta}_t \rangle}_{T_1: \text{ Expected update}} + \alpha \underbrace{\langle \nabla M(\boldsymbol{\theta}_t), \mathbf{b}(\mathbf{Y}_t) \rangle}_{T_2: \text{ Error due to Markovian noise } \mathbf{b}(\mathbf{Y}_t)} \\ &\quad + \alpha \underbrace{\langle \nabla M(\boldsymbol{\theta}_t), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) \rangle}_{T_3: \text{ Error due to Markovian noise } Y_k} + \alpha \underbrace{\langle \nabla M(\boldsymbol{\theta}_t), \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) - \bar{\mathbf{G}}(\boldsymbol{\theta}_t) \rangle}_{T_4: \text{ Error due local updates}} \\ &\quad + \frac{L\alpha^2}{2\psi} \underbrace{\|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)\|_s^2}_{T_5: \text{ Error due to noise and discretization}}. \end{aligned} \quad (31)$$

The inequality in (31) characterizes one step drift of the Lyapunov function $M(\boldsymbol{\theta}_t)$. The term T_1 is responsible for negative drift of the overall recursion. T_2 and T_3 appear due to the presence of the Markovian noises $\mathbf{b}^i(\mathbf{y}_t^i)$ and $\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)$ in the update of Algorithm 4. T_4 appears due to the mismatch between the parameters of the agents $\boldsymbol{\theta}_t^i, i = 1, \dots, N$. Finally, T_5 appears due to discretization error in the smoothness upper bound. Next, we state bounds on T_1, T_2, T_3, T_4, T_5 in the following intermediate lemmas.

Lemma B.1. *For all $\boldsymbol{\theta} \in \mathbb{R}^d$, the operator $\bar{\mathbf{G}}(\boldsymbol{\theta})$ satisfies the following*

$$T_1 = \langle \nabla M(\boldsymbol{\theta}), \bar{\mathbf{G}}(\boldsymbol{\theta}) - \boldsymbol{\theta} \rangle \leq -2\varphi_2 M(\boldsymbol{\theta}).$$

Lemma B.1 guarantees the negative drift in the one step recursion analysis of Proposition B.2. It follows from the Moreau envelope construction (Chen et al., 2021c) and the contraction property of the operators $\bar{\mathbf{G}}^i(\cdot), i = 1, \dots, N$ (Assumption 6.2).

Lemma B.2. *Consider the iteration t of the Algorithm 4, and consider $\tau = \lceil 2 \log_\rho \alpha \rceil$. We have*

$$\mathbb{E}_{t-\tau}[T_2] = \mathbb{E}_{t-\tau} \langle \nabla M(\boldsymbol{\theta}_t), \mathbf{b}(\mathbf{Y}_t) \rangle$$

$$\begin{aligned} &\leq \frac{L^2}{2\alpha\psi^2l_{cs}^4}\mathbb{E}_{t-\tau}\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\alpha}{2}\mathbb{E}_{t-\tau}[\|\mathbf{b}(\mathbf{Y}_t)\|_c^2] \\ &\quad + \frac{\alpha m_2 L}{\psi l_{cs}^2}\mathbb{E}_{t-\tau}\|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_c + \frac{1}{2}\left(\frac{Lm_2\alpha^2 u_{cm}}{\zeta_1 l_{cs}^2 \psi}\right)^2 + \zeta_1^2 \mathbb{E}_{t-\tau}[M(\boldsymbol{\theta}_t)], \end{aligned}$$

where ζ_1 is an arbitrary positive constant.

In the i.i.d. noise setting, $\mathbb{E}[T_2] = 0$. In Markov noise setting, going back τ steps (which introduces $\boldsymbol{\theta}_{t-\tau}$) enables us to use Markov chain mixing property (Assumption 6.1).

Lemma B.3. For any $t \geq 0$, denote $T_3 = \langle \nabla M(\boldsymbol{\theta}_t), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) \rangle$. For any $\tau < t$, we have

$$\begin{aligned} \mathbb{E}_{t-\tau}[T_3] &\leq \left(\zeta_2 \frac{2LA_2}{\psi l_{cs}^2} + \zeta_3 \frac{L(A_1+1)}{\psi l_{cs}^2} + \frac{6\alpha m_1 L u_{cm}^2}{l_{cs}^2 \psi} \right) \mathbb{E}_{t-\tau}[M(\boldsymbol{\theta}_t)] \\ &\quad + \left(\frac{2LA_2}{\psi l_{cs}^2} \left(\frac{1}{2} + \frac{u_{cm}^2}{2\zeta_2^2} \right) + \frac{L(A_1+1)}{\psi l_{cs}^2} \left(\frac{3u_{cm}^2}{2\zeta_3^2} + 2 \right) + \frac{3m_1 L \alpha}{l_{cs}^2 \psi} \right) \mathbb{E}_{t-\tau}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2] \\ &\quad + \left(\left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) \cdot \frac{L(A_1+1)}{\psi l_{cs}^2} + \frac{LA_2}{\psi l_{cs}^2} \right) \mathbb{E}_{t-\tau}[\Omega_t] \\ &\quad + \left(\left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) \frac{L(A_1+1)}{\psi l_{cs}^2} + \frac{m_1 L \alpha}{2l_{cs}^2 \psi} \right) \mathbb{E}_{t-\tau}[\Omega_{t-\tau}], \end{aligned}$$

where ζ_2 and ζ_3 are arbitrary positive constants.

Lemma B.4. For any $t \geq 0$, we have

$$T_4 = \langle \nabla M(\boldsymbol{\theta}_t), \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) - \bar{\mathbf{G}}(\boldsymbol{\theta}_t) \rangle \leq \zeta_4^2 M(\boldsymbol{\theta}_t) + \frac{L^2 u_{cm}^2}{2l_{cs}^4 \zeta_4^2 \psi^2} \Omega_t,$$

where ζ_4 is an arbitrary positive constant.

Lemma B.5. For any $0 \leq \tau < t$, we have

$$T_5 = \|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)\|_s^2 \leq \frac{6(A_2+1)^2 u_{cm}^2}{l_{cs}^2} M(\boldsymbol{\theta}_t) + \frac{3A_1^2}{l_{cs}^2} \Omega_t + \frac{3}{l_{cs}^2} \|\mathbf{b}(\mathbf{Y}_t)\|_c^2.$$

Substituting the bounds in Lemmas B.1, B.2, B.3, B.4, B.5, and taking expectation, we get the final bound in Proposition B.2.

B.4. Proof of Proposition B.3

First we state the following two intermediate lemmas, which are proved in Section B.7.

Lemma B.6. Suppose $\tau = \lceil 2 \log_\rho \alpha \rceil$ and

$$\alpha\tau \leq \min \left\{ \frac{1}{12\sqrt{A_2^2+1}}, \frac{1}{8(A_2+1)} \right\}. \quad (32)$$

For any $0 \leq t \leq 2\tau$ we have the following

$$M(\boldsymbol{\theta}_t) \leq \frac{1}{l_{cm}^2} \left(\frac{1}{C_1^2} \left(B + (A_2+1) \left(\|\boldsymbol{\theta}_0\|_c + \frac{B}{2C_1} \right) \right)^2 + \|\boldsymbol{\theta}_0\|_c^2 \right) \equiv M_0 \quad (33)$$

Furthermore, for any $t \geq 2\tau$, we have the following

$$\begin{aligned} \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c] &\leq 4\alpha\tau C_1 \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_t\|_c] + 8\alpha\tau \frac{u_{cD}}{l_{cD}} \frac{B}{\sqrt{N}} + \frac{u_{cD}}{l_{cD}} \frac{8m_2}{1-\rho} \alpha^2 (1+2C_1\tau) \\ &\quad + 6A_1\alpha \sum_{i=t-\tau}^t \mathbb{E}_{t-2\tau}[\Delta_i]. \end{aligned} \quad (34)$$

Lemma B.7. Suppose $\tau = \lceil 2 \log_\rho \alpha \rceil$ and

$$\alpha \leq \min \left\{ \frac{1}{C_1}, \frac{1}{4\tau C_2}, \frac{1}{40\tau C_1^2} \right\}, \quad (35)$$

where $C_1 = 3\sqrt{A_2^2 + 1}$ and $C_2 = 3C_1 + 8$. We have the following

$$\begin{aligned} \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2] &\leq 8\tau^2 \alpha^2 C_1^2 \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_t\|_c^2 + 8 \frac{u_{cD}^2 B^2}{l_{cD}^2} \alpha^2 \tau^2 \\ &\quad + 14 \frac{u_{cD}^2 m_2^2 \alpha^4 \tau}{l_{cD}^2 (1-\rho^2)} + 8\alpha^2 A_1^2 \tau \sum_{i=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-i}^2]. \end{aligned}$$

In Lemma B.6, we define $C_9 \triangleq \frac{8u_{cD}B}{l_{cD}}$, $C_{10} \triangleq \frac{8m_2 u_{cD}}{l_{cD}(1-\rho)}$. Hence, we can bound the term in (15) as

$$\begin{aligned} &\alpha^2 C_3 \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c] \\ &\leq \alpha^2 C_3 \left(4\alpha\tau C_1 \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t\|_c] + C_9 \frac{\alpha\tau}{\sqrt{N}} + C_{10} \alpha^2 (1 + 2C_1\tau) + 6A_1 \alpha \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Delta_k] \right) \\ &\leq \alpha^2 C_3 \left(4C_1 u_{cm} \alpha\tau \mathbb{E}_{t-2\tau} [\sqrt{2M(\boldsymbol{\theta}_t)}] + C_9 \frac{\alpha\tau}{\sqrt{N}} + C_{10} \alpha^2 (1 + 2C_1\tau) + 6A_1 \alpha \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Delta_k] \right) \\ &\hspace{20em} \text{(Proposition B.1)} \\ &= \frac{1}{\zeta_6} 4C_1 C_3 u_{cm} \alpha^{5/2} \tau \cdot \sqrt{\alpha} \zeta_6 \sqrt{2\mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)]} \\ &\quad + \alpha^2 C_3 \left(C_9 \frac{\alpha\tau}{\sqrt{N}} + C_{10} \alpha^2 (1 + 2C_1\tau) + 6A_1 \alpha \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Delta_k] \right) \quad (\sqrt{\cdot} \text{ is concave, } \zeta_6 > 0) \\ &\leq \frac{8C_1^2 C_3^2 u_{cm}^2}{\zeta_6^2} \alpha^5 \tau^2 + \alpha \zeta_6^2 \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + C_3 C_9 \frac{\alpha^3 \tau}{\sqrt{N}} + C_3 C_{10} \alpha^4 (1 + 2C_1\tau) \\ &\quad + 6A_1 C_3 \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} \left[\frac{1}{2} \alpha^5 + \frac{1}{2} \alpha \Delta_k^2 \right] \quad \text{(Young's inequality)} \\ &\leq \underbrace{\frac{8C_1^2 C_3^2 u_{cm}^2}{\zeta_6^2}}_{C_{11}} \alpha^5 \tau^2 + \alpha \zeta_6^2 \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + C_3 C_9 \frac{\alpha^3 \tau}{\sqrt{N}} + \alpha^4 [C_3 C_{10} (1 + 2C_1\tau) + 3A_1 C_3 \alpha (\tau + 1)] \\ &\quad + 3\alpha A_1 C_3 \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k] \quad \text{(By (58))} \\ &\leq C_{11} \alpha^5 \tau^2 + \alpha \zeta_6^2 \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + \frac{1}{2} C_3^2 C_9^2 \alpha^4 \tau + \frac{1}{2} \frac{\alpha^2 \tau}{N} + \alpha^4 [C_3 C_{10} (1 + 2C_1\tau) + 3A_1 C_3 \alpha (\tau + 1)] \\ &\quad + 3\alpha A_1 C_3 \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k]. \quad (36) \end{aligned}$$

Furthermore, using Lemma B.7, (58) and Proposition B.1, the term in (16) can be bounded as follows:

$$\begin{aligned} C_4 \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2] &\leq 16\tau^2 \alpha^2 C_1^2 C_4 u_{cm}^2 \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + 8\alpha^2 A_1^2 C_4 \tau \sum_{k=0}^{\tau} \mathbb{E}_{t-2\tau} [\Omega_{t-k}] \\ &\quad + \underbrace{\frac{8C_4 u_{cD}^2 B^2}{l_{cD}^2}}_{C_{12}} \frac{\alpha^2}{N} \tau^2 + \underbrace{\frac{14C_4 u_{cD}^2 m_2^2}{l_{cD}^2 (1-\rho^2)}}_{C_{13}} \alpha^4 \tau. \quad (37) \end{aligned}$$

Inserting the upper bounds in (36) and (37) in the upper bound in Proposition B.2, we have

$$\begin{aligned}
 & \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_{t+1})] \\
 & \leq \left(1 - \alpha \left[2\varphi_2 - \zeta_1^2 - \zeta_2^2 \frac{2LA_2}{\psi l_{cs}^2} - \zeta_3^2 \frac{L(A_1+1)}{\psi l_{cs}^2} - \zeta_4^2 - \zeta_6^2 \right] \right. \\
 & \quad \left. + \alpha^2 \frac{6Lu_{cm}^2}{l_{cs}^2 \psi} \left[m_1 + \frac{(A_2+1)^2}{2} \right] + \alpha^2 16\tau^2 C_1^2 C_4 u_{cm}^2 \right) \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] \\
 & \quad + \alpha C_5 \mathbb{E}_{t-2\tau} [\Omega_t] + \alpha C_6 \mathbb{E}_{t-2\tau} [\Omega_{t-\tau}] \\
 & \quad + \underbrace{\left(C_7 + C_{11} \alpha \tau^2 + 0.5 C_3^2 C_9^2 \tau + C_3 C_{10} + 2C_1 C_3 C_{10} \tau + 3A_1 C_3 \alpha (\tau + 1) + C_{13} \tau + C_8 \frac{u_{cD}^2}{l_{cD}^2} 2m_2^2 \alpha^2 \right)}_{C'_{14}(\tau)} \alpha^4 \\
 & \quad + \left(C_8 \frac{u_{cD}^2 B^2}{l_{cD}^2} + \frac{\tau}{2} + C_{12} \tau^2 \right) \frac{\alpha^2}{N} + (3\alpha A_1 C_3 + 8\alpha^2 A_1^2 C_4 \tau) \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k]. \tag{38}
 \end{aligned}$$

We define $C'_{14}(\tau) \leq C_{14}(\tau) \triangleq (C_7 + C_{11} + 0.5 C_3^2 C_9^2 + C_3 C_{10} + 2C_1 C_3 C_{10} + 3A_1 C_3 + C_{13} + C_8 \frac{u_{cD}^2}{l_{cD}^2} 2m_2^2 \alpha^2) \tau^2$. Also, we choose $\zeta_1 = \zeta_4 = \zeta_6 = \sqrt{\varphi_2/10}$, $\zeta_2 = \sqrt{\frac{\varphi_2}{10} \cdot \frac{\psi l_{cs}^2}{2LA_2}}$, $\zeta_3 = \sqrt{\frac{\varphi_2}{10} \cdot \frac{\psi l_{cs}^2}{L(A_1+1)}}$, and denote

$$C_{15}(\tau) = \left(\frac{6m_1 Lu_{cm}^2}{l_{cs}^2 \psi} + \frac{6L(A_2+1)^2 u_{cm}^2}{2\psi l_{cs}^2} + 144(A_2^2 + 1)C_4 u_{cm}^2 \right) \tau^2. \tag{39}$$

This yields

$$\begin{aligned}
 & \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_{t+1})] \leq \left(1 - \frac{3}{2} \alpha \varphi_2 + \alpha^2 C_{15}(\tau) \right) \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + \left(C_8 \frac{u_{cD}^2 B^2}{l_{cD}^2} + \frac{\tau}{2} + C_{12} \tau^2 \right) \frac{\alpha^2}{N} \\
 & \quad + C_{14}(\tau) \alpha^4 + (3\alpha A_1 C_3 + 8\alpha^2 A_1^2 C_4 \tau + \alpha C_5 + \alpha C_6) \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k] \\
 & \leq \left(1 - \frac{3}{2} \alpha \varphi_2 + \alpha^2 C_{15}(\tau) \right) \mathbb{E}_{t-2\tau} [M(\boldsymbol{\theta}_t)] + C_{14}(\tau) \alpha^4 + C_{16}(\tau) \frac{\alpha^2}{N} + \alpha C_{17} \sum_{k=t-\tau}^t \mathbb{E}_{t-2\tau} [\Omega_k], \tag{40}
 \end{aligned}$$

where $C_{16}(\tau) = \left(C_8 \frac{u_{cD}^2 B^2}{l_{cD}^2} + \frac{1}{2} + C_{12} \right) \tau^2$ and $C_{17} = (3A_1 C_3 + 8A_1^2 C_4 + C_5 + C_6)$. Due to $\alpha \leq \frac{\varphi_2}{2C_{15}(\tau)}$, we have $1 - \frac{3}{2} \alpha \varphi_2 + \alpha^2 C_{15}(\tau) \leq 1 - \alpha \varphi_2$. This completes the proof.

B.5. Proof of Proposition B.4

First we state the following lemma, which characterizes a bound on the expectation of the synchronization error Ω_t .

Lemma B.8. *Suppose Assumptions 6.1, 6.3 holds and the step size α satisfies $\alpha \leq \sqrt{\frac{\ln(5/4)}{2(1+\tilde{A}_1)(K-1)^2}}$. Then, for $sK \leq t \leq (s+1)K - 1$, where $s = \lfloor t/K \rfloor$, the network consensus error $\Omega_t \triangleq \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_t^i\|_c^2$ satisfies*

$$\mathbb{E} \Omega_t \leq 5\alpha^2 (t - sK) \tilde{B}^2 \left(1 + \frac{4m_2 \rho}{B(1-\rho)} \right) + 5\alpha^2 (t - sK) \tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E} \|\boldsymbol{\theta}_{t'}\|_c^2, \tag{41}$$

where $\tilde{A}_1 = \frac{2A_1^2 u_{c2}^2}{l_{c2}^2}$, $\tilde{A}_2 = \frac{2A_2^2 u_{c2}^2}{l_{c2}^2}$, $\tilde{B} = \frac{u_{c2}}{l_{c2}} B$. Here, A_1, A_2, B are the constants defined in Assumption 6.3, and l_{c2}, u_{c2} are constants involved in the equivalence of the norms: $l_{c2} \|\cdot\|_2 \leq \|\cdot\|_c \leq u_{c2} \|\cdot\|_2$.

Due to the assumption on the step size, the bound in Lemma B.8 hold. Substituting the bound on Ω_ℓ from Lemma B.8, we get

$$\frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \left[\sum_{\ell=t-\tau}^t \mathbb{E} \Omega_\ell \right]$$

$$\leq \frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t \left[5\alpha^2(\ell - s_\ell K) \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) + 5\alpha^2(\ell - s_\ell K) \tilde{A}_2 \sum_{t'=s_\ell K}^{\ell-1} \mathbb{E} \|\boldsymbol{\theta}_{t'}\|_c^2 \right]. \quad (42)$$

where $s_\ell = \lfloor \ell/K \rfloor$. Hence, $s_\ell K$ denotes the last time instant before ℓ when synchronization happened. The first term in (42) can be upper bounded as follows.

$$\begin{aligned} & \frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t 5\alpha^2(\ell - s_\ell K) \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \\ &= \alpha^2 \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \frac{10C_{17}}{\varphi_2} \frac{1}{W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \\ &\leq \alpha^2 \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \frac{10C_{17}}{\varphi_2} \frac{1}{W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t (K-1) \quad (\ell - s_\ell K \leq K-1) \\ &= \alpha^2 \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \frac{10C_{17}}{\varphi_2} \left[\frac{1}{W_T} \sum_{t=2\tau}^T w_t \right] (\tau+1)(K-1) \\ &\leq \alpha^2 \tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \frac{10C_{17}}{\varphi_2} (\tau+1)(K-1). \end{aligned} \quad (43)$$

Next, we compute the second term in (42).

$$\begin{aligned} & \frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t 5\alpha^2(\ell - s_\ell K) \tilde{A}_2 \sum_{\ell'=s_\ell K}^{\ell-1} \|\boldsymbol{\theta}_{\ell'}\|_c^2 \\ &\leq \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \quad (\text{Proposition B.1}) \\ &\leq \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2 W_T} \left[\underbrace{\sum_{t=2\tau}^K w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'})}_{\mathcal{I}_1} + \underbrace{\sum_{t=K+1}^T w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'})}_{\mathcal{I}_2} \right], \end{aligned} \quad (44)$$

where if $K < 2\tau$, $\mathcal{I}_1 = 0$. Next, we bound $\mathcal{I}_1, \mathcal{I}_2$ separately.

$$\begin{aligned} \mathcal{I}_1 &= \sum_{t=2\tau}^K w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \\ &\leq (K-1) \sum_{t=2\tau}^K w_t \sum_{\ell=t-\tau}^t \sum_{\ell'=0}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \quad (\text{for } \ell < K, s_\ell = 0; \text{ for } \ell = K, \ell - s_\ell K = 0) \\ &\leq (K-1)(\tau+1) \sum_{t=2\tau}^K w_t \sum_{\ell'=0}^{t-1} M(\boldsymbol{\theta}_{\ell'}) \\ &\leq (K-1)(K-2\tau+1)(\tau+1) w_K \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t), \end{aligned} \quad (45)$$

where, (45) follows since $w_{t-1} \leq w_t, \forall t$. Next, to bound \mathcal{I}_2 in (44), we again split it into two terms.

$$\mathcal{I}_2 = \underbrace{\sum_{t=K+1}^{K+\tau} w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'})}_{\mathcal{I}_3} + \underbrace{\sum_{t=K+\tau+1}^T w_t \sum_{\ell=t-\tau}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'})}_{\mathcal{I}_4}.$$

First, we bound \mathcal{I}_3 .

$$\begin{aligned}
 \mathcal{I}_3 &\leq \sum_{t=K+1}^{K+\tau} w_t \left[\sum_{\ell=t-\tau}^K (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) + \sum_{\ell=K+1}^t (\ell - s_\ell K) \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \right] \\
 &\leq \sum_{t=K+1}^{K+\tau} w_t \left[\sum_{\ell=t-\tau}^K \ell \sum_{\ell'=0}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) + \sum_{\ell=K+1}^t (\ell - K) \sum_{\ell'=K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \right] \quad (s_\ell = 0 \text{ for } \ell < K, s_\ell \geq 1 \text{ for } K \leq \ell \leq K + \tau) \\
 &\leq \sum_{t=K+1}^{K+\tau} w_t \left[K(K+1-t+\tau) \sum_{\ell'=0}^{K-1} M(\boldsymbol{\theta}_{\ell'}) + (t-K)^2 \sum_{\ell'=K}^{t-1} M(\boldsymbol{\theta}_{\ell'}) \right] \\
 &\leq w_{K+\tau} K \tau^2 \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t) + \tau^3 w_{K+\tau} \sum_{t=K}^{K+\tau-1} M(\boldsymbol{\theta}_t). \tag{46}
 \end{aligned}$$

Next, we bound \mathcal{I}_4 , assuming $t_0 K + \tau \leq T < (t_0 + 1)K + \tau$, where t_0 is a non-negative integer.

$$\begin{aligned}
 \mathcal{I}_4 &= \sum_{t=K+\tau+1}^T w_t \sum_{\ell=t-\tau}^t \underbrace{(\ell - s_\ell K)}_{\leq K-1} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \\
 &\leq (K-1) \sum_{t=K+\tau+1}^T w_t \sum_{\ell=t-\tau}^t \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \\
 &= (K-1) \left[w_{K+\tau+1} \sum_{\ell=K+1}^{K+\tau+1} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) + \cdots + w_{K+\tau+K-1} \sum_{\ell=2K-1}^{2K+\tau-1} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \right] \\
 &\quad + (K-1) \left[w_{2K+\tau} \sum_{\ell=2K}^{2K+\tau} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) + \cdots + w_{3K+\tau-1} \sum_{\ell=3K-1}^{3K+\tau-1} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \right] \\
 &\quad + \cdots + (K-1) \left[w_{t_0 K + \tau} \sum_{\ell=t_0 K}^{t_0 K + \tau} \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) + \cdots + w_T \sum_{\ell=T-\tau}^T \sum_{\ell'=s_\ell K}^{\ell-1} M(\boldsymbol{\theta}_{\ell'}) \right] \\
 &\leq (K-1)(\tau+1) \left[w_{K+\tau+1} \sum_{\ell=K}^{K+\tau} M(\boldsymbol{\theta}_\ell) + \cdots + w_{2K+\tau-1} \sum_{\ell=K}^{2K+\tau-2} M(\boldsymbol{\theta}_\ell) \right] \\
 &\quad + (K-1)(\tau+1) \left[w_{2K+\tau} \sum_{\ell=2K}^{2K+\tau-1} M(\boldsymbol{\theta}_\ell) + \cdots + w_{3K+\tau-1} \sum_{\ell=2K}^{3K+\tau-2} M(\boldsymbol{\theta}_\ell) \right] \\
 &\quad + \cdots + (K-1)(\tau+1) \left[w_{t_0 K + \tau} \sum_{\ell=t_0 K}^{t_0 K + \tau - 1} M(\boldsymbol{\theta}_\ell) + \cdots + w_T \sum_{\ell=t_0 K}^{T-1} M(\boldsymbol{\theta}_\ell) \right] \\
 &\leq (K-1)K(\tau+1) \left[w_{2K+\tau-1} \sum_{\ell=K}^{2K+\tau-2} M(\boldsymbol{\theta}_\ell) + w_{3K+\tau-1} \sum_{\ell'=2K}^{3K+\tau-2} M(\boldsymbol{\theta}_{\ell'}) + \cdots + w_T \sum_{\ell=t_0 K}^{T-1} M(\boldsymbol{\theta}_\ell) \right] \\
 &\leq (K-1)K(\tau+1) \left[w_{2K+\tau-1} \left(\sum_{\ell=K}^{2K-1} M(\boldsymbol{\theta}_\ell) + \sum_{\ell=2K}^{2K+\tau-2} M(\boldsymbol{\theta}_\ell) \right) \right. \\
 &\quad \left. + w_{3K+\tau-1} \left(\sum_{\ell'=2K}^{3K-1} M(\boldsymbol{\theta}_{\ell'}) + \sum_{\ell'=3K}^{3K+\tau-2} M(\boldsymbol{\theta}_{\ell'}) \right) \right. \\
 &\quad \left. + \cdots + w_{t_0 K + \tau - 1} \left(\sum_{\ell'=(t_0-1)K}^{t_0 K - 1} M(\boldsymbol{\theta}_{\ell'}) + \sum_{\ell'=t_0 K}^{t_0 K + \tau - 1} M(\boldsymbol{\theta}_{\ell'}) \right) + w_T \sum_{\ell=t_0 K}^{T-1} M(\boldsymbol{\theta}_\ell) \right] \\
 &\leq (K-1)K(\tau+1) \left\lceil \frac{(\tau-1)}{K} \right\rceil \left[\sum_{\ell=K}^{2K-1} w_{\ell+K+\tau-1} M(\boldsymbol{\theta}_\ell) + \sum_{\ell=2K}^{3K-1} w_{\ell+K+\tau-1} M(\boldsymbol{\theta}_\ell) + \cdots + w_T \sum_{\ell=t_0 K}^{T-1} M(\boldsymbol{\theta}_\ell) \right]
 \end{aligned}$$

$$\begin{aligned}
 &\leq (K-1)\tau(\tau+1) \left[\sum_{\ell=K}^{T-K-\tau} w_{\ell+K+\tau-1} M(\boldsymbol{\theta}_\ell) + w_T \sum_{\ell=T-K-\tau+1}^{T-1} M(\boldsymbol{\theta}_\ell) \right] \\
 &= (K-1)\tau(\tau+1) \left[\sum_{\ell=K}^{T-K-\tau} \frac{w_\ell}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} M(\boldsymbol{\theta}_\ell) + \sum_{\ell=T-K-\tau+1}^{T-1} \frac{w_\ell}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{T-\ell}} M(\boldsymbol{\theta}_\ell) \right] \\
 &\hspace{20em} (w_t = w_{t+1} \left(1 - \frac{\alpha\varphi_2}{2}\right)) \\
 &\leq \frac{(K-1)\tau(\tau+1)}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} \sum_{t=K}^{T-1} w_t M(\boldsymbol{\theta}_t). \tag{47}
 \end{aligned}$$

Using the bounds on $\mathcal{I}_3, \mathcal{I}_4$ from (46) and (47) respectively, we can bound \mathcal{I}_2 .

$$\mathcal{I}_2 \leq w_{K+\tau} K \tau^2 \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t) + \tau^3 w_{K+\tau} \sum_{t=K}^{K+\tau-1} M(\boldsymbol{\theta}_t) + \frac{(K-1)\tau(\tau+1)}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} \sum_{t=K}^{T-1} w_t M(\boldsymbol{\theta}_t). \tag{48}$$

Substituting the bounds on $\mathcal{I}_1, \mathcal{I}_2$ from (45), (48) respectively, into (44), we get

$$\begin{aligned}
 &\frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t 5\alpha^2 (\ell - s_\ell K) \tilde{A}_2 \sum_{\ell'=s_\ell K}^{\ell-1} \|\boldsymbol{\theta}_{\ell'}\|_c^2 \\
 &\leq \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2 W_T} \left[(K-1)(K-2\tau+1)(\tau+1)w_K \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t) + K\tau^2 w_{K+\tau} \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t) \right] \\
 &\quad + \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2 W_T} \left[\tau^3 w_{K+\tau} \sum_{t=K}^{K+\tau-1} M(\boldsymbol{\theta}_t) + \frac{(K-1)\tau(\tau+1)}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} \sum_{t=K}^{T-1} w_t M(\boldsymbol{\theta}_t) \right] \tag{49}
 \end{aligned}$$

We analyze the terms in (49) separately. First, for the terms with $\sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t)$,

$$\begin{aligned}
 &\frac{1}{W_T} \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} K(\tau+1) [(K-2\tau+1)w_K + \tau w_{K+\tau}] \sum_{t=0}^{K-1} M(\boldsymbol{\theta}_t) \\
 &= \frac{1}{W_T} \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} K(\tau+1) \sum_{t=0}^{K-1} \left[(K-2\tau+1) \frac{w_t}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K-t}} + \tau \frac{w_t}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-t}} \right] M(\boldsymbol{\theta}_t) \\
 &= \frac{1}{W_T} \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} K(\tau+1) \sum_{t=0}^{K-1} \left[(K-2\tau+1) + \frac{\tau}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} \right] \frac{w_t}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K-t}} M(\boldsymbol{\theta}_t) \\
 &\leq \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} K(\tau+1) \left[(K-2\tau+1) + \frac{\tau}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} \right] \frac{1}{W_T} \sum_{t=0}^{K-1} w_t M(\boldsymbol{\theta}_t) \\
 &\leq \frac{1}{2W_T} \sum_{t=0}^{K-1} w_t M(\boldsymbol{\theta}_t), \tag{50}
 \end{aligned}$$

where (50) holds since we choose α small enough such that

$$\alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} \left[(K-2\tau+1) + \frac{\tau}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} \right] K(\tau+1) \leq \frac{1}{2}.$$

To get this, we use the inequality $1 - x \geq \exp\left(-\frac{x}{1-x}\right)$ for $x < 1$, $\frac{\alpha\varphi_2}{2} \leq \frac{1}{2}$ and $\tau < 2\log_\rho \alpha + 1$, we get $\frac{1}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} \leq \exp(\alpha\varphi_2(2\log_\rho \alpha + 1))$. For (50) to hold, it is sufficient that

$$\alpha^2 \leq \frac{\varphi_2}{80C_{17}u_{cm}^2 \tilde{A}_2 K 2(\log_\rho(\alpha) + 1)} \min \left\{ \frac{1}{K}, \frac{1}{2(\log_\rho(\alpha) + 1) \exp(\alpha\varphi_2(2\log_\rho \alpha + 1))} \right\} \tag{51}$$

Next, for the remaining terms in the third line of (49), by the assumption on the step size, we have

$$\begin{aligned} \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} \frac{\tau^3}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} &\leq \frac{1}{4}, \\ \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} \frac{(\tau+1)\tau(K-1)}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} &\leq \frac{1}{4}, \end{aligned} \quad (52)$$

and hence we get

$$\begin{aligned} &\frac{1}{W_T} \alpha^2 \tilde{A}_2 \frac{20C_{17}u_{cm}^2}{\varphi_2} \left[\frac{\tau^3 w_K}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^\tau} \sum_{t=K}^{K+\tau-1} M(\boldsymbol{\theta}_t) + \frac{(\tau+1)\tau(K-1)}{\left(1 - \frac{\alpha\varphi_2}{2}\right)^{K+\tau-1}} \sum_{t=K}^T w_t M(\boldsymbol{\theta}_t) \right] \\ &\leq \frac{1}{2W_T} \sum_{t=K}^T w_t M(\boldsymbol{\theta}_t). \end{aligned} \quad (53)$$

Substituting (50), (53) in (49), we get

$$\frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \sum_{\ell=t-\tau}^t 5\alpha^2(\ell - s_\ell K) \tilde{A}_2 \sum_{t=s_\ell K}^{\ell-1} \|\boldsymbol{\theta}_t\|_c^2 \leq \frac{1}{2W_T} \sum_{t=0}^T w_t M(\boldsymbol{\theta}_t). \quad (54)$$

Finally, substituting the bounds in (43), (54) into (42), we get

$$\begin{aligned} &\frac{2C_{17}}{\varphi_2 W_T} \sum_{t=2\tau}^T w_t \left[\sum_{\ell=t-\tau}^t \mathbb{E}\Omega_\ell \right] \\ &\leq \alpha^2 \tilde{B}^2 \frac{10C_{17}}{\varphi_2} \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) (\tau+1)(K-1) + \frac{1}{2W_T} \sum_{t=0}^T w_t \mathbb{E}M(\boldsymbol{\theta}_t). \end{aligned} \quad (55)$$

B.6. Auxiliary Lemmas

The following lemma is of central importance in proving linear speedup of FedSAM.

Lemma B.9. *Let l_{cD} and u_{cD} be constants that satisfy $l_{cD}\|\cdot\|_D \leq \|\cdot\|_c \leq u_{cD}\|\cdot\|_D$, where $\|\cdot\|_D = \sqrt{x^\top D x}$ for some positive definite matrix $D \succ 0$. Note that for any $D \succ 0$, these constants always exist due to norm equivalence. Furthermore, in case the norm $\|x\|_c$ is defined in the form $\sqrt{x^\top D x}$ for some $D \succ 0$, we take $l_{cD} = u_{cD} = 1$. We have*

$$\mathbb{E}_{t-r}[\|\mathbf{b}(\mathbf{Y}_t)\|_c] \leq \frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} + 2m_2\rho^r \right] \quad (56)$$

$$\mathbb{E}_{t-r}[\|\mathbf{b}(\mathbf{Y}_t)\|_c^2] \leq \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} + 2m_2^2\rho^{2r} \right]. \quad (57)$$

Lemma B.9 is essential in characterizing the linear speedup in Theorem B.1. This lemma characterizes the bound on the conditional expectation of $\|\mathbf{b}(\mathbf{Y}_t)\|_c$ and $\|\mathbf{b}(\mathbf{Y}_t)\|_c^2$, conditioned on r time steps before. In order to understand this lemma, consider the bound in (57). For the sake of understanding, suppose the noise \mathbf{Y}_t is i.i.d. In this case we will end up with the first term which is proportional to $1/N$. This is precisely the linear reduction of the variance of sum of N i.i.d. random variables. Furthermore, in order to extend the i.i.d. noise setting to the more general Markovian noise, we need to pay an extra price by adding the exponentially decreasing term to the first variance term.

Lemma B.10. *The following hold*

$$\begin{aligned} \|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t)\|_c^2 &\leq A_1^2 \Delta_t^2 \leq A_1^2 \Omega_t \\ \Delta_t^2 &\leq \Omega_t. \end{aligned} \quad (58)$$

Lemma B.11. *For the generalized Moreau Envelope defined in (12), it holds that*

$$\begin{aligned} \left\| \nabla M_f^{\psi,g}(x) \right\|_m^* &= \|x\|_m, \\ \left\langle \nabla M_f^{\psi,g}(x), x \right\rangle &\geq 2M_f^{\psi,g}(x). \end{aligned}$$

B.7. Proof of Lemmas

Proof of Lemma B.1. Using Cauchy-Schwarz inequality, we have

$$\langle \nabla M(\boldsymbol{\theta}), \bar{\mathbf{G}}(\boldsymbol{\theta}) - \boldsymbol{\theta} \rangle \leq \underbrace{\|\nabla M(\boldsymbol{\theta})\|_m^* \cdot \|\bar{\mathbf{G}}(\boldsymbol{\theta})\|_m}_{T_{11}} - \underbrace{\langle \nabla M(\boldsymbol{\theta}), \boldsymbol{\theta} \rangle}_{T_{12}}, \quad (59)$$

where $\|\cdot\|_m^*$ denotes the dual of the norm $\|\cdot\|_m$. Furthermore, by Proposition B.1 we have

$$\begin{aligned} T_{11} &\stackrel{(a)}{\leq} \|\boldsymbol{\theta}\|_m \cdot \|\bar{\mathbf{G}}(\boldsymbol{\theta})\|_m && \text{(Lemma B.11)} \\ &\leq \|\boldsymbol{\theta}\|_m \cdot l_{cm}^{-1} \left\| \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{G}}^i(\boldsymbol{\theta}) \right\|_c && \text{(By (13))} \\ &\leq \|\boldsymbol{\theta}\|_m \cdot l_{cm}^{-1} \frac{1}{N} \sum_{i=1}^N \|\bar{\mathbf{G}}^i(\boldsymbol{\theta})\|_c && \text{(triangle inequality)} \\ &\leq \|\boldsymbol{\theta}\|_m \cdot \frac{u_{cm} \gamma_c}{l_{cm}} \|\boldsymbol{\theta}\|_m && \text{(Assumption 6.2, Proposition B.1)} \\ &= \frac{2u_{cm} \gamma_c}{l_{cm}} M(\boldsymbol{\theta}). && (60) \end{aligned}$$

Furthermore, by the convexity of the $\|\cdot\|_m$ norm (Lemma B.11), we have

$$T_{12} \geq \|\boldsymbol{\theta}\|_m^2 = 2M(\boldsymbol{\theta}). \quad (61)$$

Hence, using (60), (61) in (59) we get

$$\langle \nabla M(\boldsymbol{\theta}), \bar{\mathbf{G}}(\boldsymbol{\theta}) - \boldsymbol{\theta} \rangle \leq -2 \left(1 - \frac{u_{cm} \gamma_c}{l_{cm}} \right) M(\boldsymbol{\theta}) = -2\varphi_2 M(\boldsymbol{\theta}), \quad (62)$$

where φ_2 is defined in (14). \square

Proof of Lemma B.2. Given some $\tau < t$, $T_2 = \langle \nabla M(\boldsymbol{\theta}_t), \mathbf{b}(\mathbf{Y}_t) \rangle$ can be written as follows:

$$\begin{aligned} T_2 &= \langle \nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{b}(\mathbf{Y}_t) \rangle + \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{b}(\mathbf{Y}_t) \rangle \\ &\leq \|\nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \cdot \|\mathbf{b}(\mathbf{Y}_t)\|_s + \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{b}(\mathbf{Y}_t) \rangle && \text{(Cauchy-Schwarz)} \\ &\leq \frac{L}{\psi l_{cs}} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c \cdot \frac{1}{l_{cs}} \|\mathbf{b}(\mathbf{Y}_t)\|_c + \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{b}(\mathbf{Y}_t) \rangle && \text{(Proposition B.1)} \\ &\leq \frac{1}{2\alpha} \left(\frac{L}{\psi l_{cs}^2} \right)^2 \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\alpha}{2} \|\mathbf{b}(\mathbf{Y}_t)\|_c^2 + \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{b}(\mathbf{Y}_t) \rangle. && (63) \end{aligned}$$

Taking expectation on both sides, we have

$$\begin{aligned} \mathbb{E}_{t-\tau}[T_2] &= \frac{L^2}{2\alpha\psi^2 l_{cs}^4} \mathbb{E}_{t-\tau} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\alpha}{2} \mathbb{E}_{t-\tau} \|\mathbf{b}(\mathbf{Y}_t)\|_c^2 + \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbb{E}_{t-\tau}[\mathbf{b}(\mathbf{Y}_t)] \rangle \\ &\leq \frac{L^2}{2\alpha\psi^2 l_{cs}^4} \mathbb{E}_{t-\tau} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\alpha}{2} \mathbb{E}_{t-\tau} \|\mathbf{b}(\mathbf{Y}_t)\|_c^2 + \underbrace{\|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \cdot \|\mathbb{E}_{t-\tau}[\mathbf{b}(\mathbf{Y}_t)]\|_s}_{T_{21}}. && (64) \end{aligned}$$

For T_{21} , we have

$$\begin{aligned} T_{21} &\leq l_{cs}^{-1} [\|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \|\mathbb{E}_{t-\tau}[\mathbf{b}(\mathbf{Y}_t)]\|_c] \\ &\leq l_{cs}^{-1} [\|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* m_2 \rho^\tau] && \text{(Assumption 6.1)} \\ &\leq \frac{m_2 \alpha^2}{l_{cs}} [\|\nabla M(\boldsymbol{\theta}_{t-\tau}) - \nabla M(\boldsymbol{\theta}_t)\|_s^* + \|\nabla M(\boldsymbol{\theta}_t)\|_s^*] && \text{(assumption on } \tau) \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{m_2 \alpha^2}{l_{cs}} \left[\frac{L}{\psi} \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_s + \|\nabla M(\boldsymbol{\theta}_t) - \nabla M(\mathbf{0})\|_s^* \right] && \text{(Proposition B.1)} \\
 &\leq \frac{m_2 \alpha^2}{l_{cs}^2} \frac{L}{\psi} [\|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_c + \|\boldsymbol{\theta}_t\|_c] && \text{(By (13))} \\
 &\leq \frac{m_2 \alpha}{l_{cs}^2} \frac{L}{\psi} \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_c + \frac{m_2 \alpha^2 u_{cm}}{\zeta_1 l_{cs}^2} \frac{L}{\psi} \cdot \zeta_1 \|\boldsymbol{\theta}_t\|_m && \text{(Proposition B.1; } \zeta_1 > 0) \\
 &\leq \frac{m_2 \alpha}{l_{cs}^2} \frac{L}{\psi} \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_c + \frac{1}{2} \left(\frac{m_2 \alpha^2 u_{cm}}{\zeta_1 l_{cs}^2} \frac{L}{\psi} \right)^2 + \zeta_1^2 M(\boldsymbol{\theta}_t). && (65)
 \end{aligned}$$

Substituting (65) in (64), we get

$$\begin{aligned}
 \mathbb{E}_{t-\tau}[T_2] &\leq \frac{L^2}{2\alpha\psi^2 l_{cs}^4} \mathbb{E}_{t-\tau} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\alpha}{2} \mathbb{E}_{t-\tau} \|b(\mathbf{Y}_t)\|_c^2 + \frac{\alpha m_2 L}{\psi l_{cs}^2} \mathbb{E}_{t-\tau} \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_t\|_c \\
 &\quad + \frac{1}{2} \left(\frac{m_2 \alpha^2 u_{cm}}{\zeta_1 l_{cs}^2} \frac{L}{\psi} \right)^2 + \zeta_1^2 \mathbb{E}_{t-\tau}[M(\boldsymbol{\theta}_t)].
 \end{aligned}$$

□

Proof of Lemma B.3.

$$\begin{aligned}
 T_3 &= \langle \nabla M(\boldsymbol{\theta}_t), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) \rangle \\
 &= \underbrace{\langle \nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) \rangle}_{T_{31}} \\
 &\quad + \underbrace{\langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \mathbf{G}(\boldsymbol{\Theta}_{t-\tau}, \mathbf{Y}_t) + \bar{\mathbf{G}}(\boldsymbol{\Theta}_{t-\tau}) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) \rangle}_{T_{32}} \\
 &\quad + \underbrace{\langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{G}(\boldsymbol{\Theta}_{t-\tau}, \mathbf{Y}_t) - \bar{\mathbf{G}}(\boldsymbol{\Theta}_{t-\tau}) \rangle}_{T_{33}}.
 \end{aligned} \tag{66}$$

Next, we bound all three terms individually.

I. Bound on T_{31} :

$$\begin{aligned}
 T_{31} &= \frac{1}{N} \sum_{i=1}^N \langle \nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i) \rangle \\
 &\leq \frac{1}{N} \sum_{i=1}^N \underbrace{\|\nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^*}_{T_{311}} \cdot \underbrace{\|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i)\|_s}_{T_{312}}.
 \end{aligned} \tag{67}$$

For T_{311} , we have

$$T_{311} = \|\nabla M(\boldsymbol{\theta}_t) - \nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \leq \frac{L}{\psi l_{cs}} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c, \tag{68}$$

where the inequality follows from Proposition B.1 and (13). For T_{312} , we have,

$$\begin{aligned}
 T_{312} &= \|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbb{E}_{\mathbf{y} \sim \mu^i} \mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y})\|_s \\
 &\leq l_{cs}^{-1} [\|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)\|_c + \mathbb{E}_{\mathbf{y} \sim \mu^i} \|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y})\|_c] && \text{(Triangle and Jensen's inequality)} \\
 &\leq l_{cs}^{-1} [A_2 \|\boldsymbol{\theta}_t^i\|_c + A_2 \|\boldsymbol{\theta}_t^i\|_c] && \text{(Assumption 6.3)} \\
 &\leq \frac{2A_2}{l_{cs}} [\|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c + \|\boldsymbol{\theta}_t\|_c] && \text{(triangle inequality)} \\
 &= \frac{2A_2}{l_{cs}} [\Delta_t^i + \|\boldsymbol{\theta}_t\|_c].
 \end{aligned} \tag{69}$$

Substituting (68) and (69) in (67), we get

$$\begin{aligned}
 T_{31} &\leq \frac{2LA_2}{\psi l_{cs}^2} \left[\frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c \cdot [\Delta_t^i + \|\boldsymbol{\theta}_t\|_c] \right] \\
 &\leq \frac{2LA_2}{\psi l_{cs}^2} \left[\frac{1}{N} \sum_{i=1}^N [\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c \cdot \Delta_t^i] + \frac{u_{cm}}{\zeta_2} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c \cdot \zeta_2 \|\boldsymbol{\theta}_t\|_m \right] && \text{(Proposition B.1)} \\
 &\leq \frac{2LA_2}{\psi l_{cs}^2} \left[\frac{1}{N} \sum_{i=1}^N \left[\frac{1}{2} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{1}{2} (\Delta_t^i)^2 \right] + \frac{u_{cm}^2}{2\zeta_2^2} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{\zeta_2^2}{2} \|\boldsymbol{\theta}_t\|_m^2 \right] && \text{(Young's inequality)} \\
 &= \frac{2LA_2}{\psi l_{cs}^2} \left[\left(\frac{1}{2} + \frac{u_{cm}^2}{2\zeta_2^2} \right) \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \zeta_2^2 M(\boldsymbol{\theta}_t) + \frac{1}{2} \Omega_t \right]. && (70)
 \end{aligned}$$

II. Bound on T_{32} :

$$\begin{aligned}
 T_{32} &= \frac{1}{N} \sum_{i=1}^N \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i) + \bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i) \rangle \\
 &\leq \frac{1}{N} \sum_{i=1}^N \underbrace{\|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^*}_{T_{321}} \cdot \underbrace{\|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i) + \bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i)\|_s}_{T_{322}}, && (71)
 \end{aligned}$$

where the last inequality is by Hölder's inequality. For T_{321} , we have

$$\begin{aligned}
 T_{321} &= \|\nabla M(\boldsymbol{\theta}_{t-\tau}) - \nabla M(\mathbf{0})\|_s^* \leq \frac{L}{\psi} \|\boldsymbol{\theta}_{t-\tau} - \mathbf{0}\|_s && \text{(Proposition B.1)} \\
 &\leq \frac{L}{\psi l_{cs}} [\|\boldsymbol{\theta}_t\|_c + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c]. && (72)
 \end{aligned}$$

For T_{322} , we have

$$\begin{aligned}
 T_{322} &= \|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i) + \bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i)\|_s \\
 &\leq \frac{1}{l_{cs}} [\|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i)\|_c + \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i)\|_c] && \text{(triangle inequality)} \\
 &\leq \frac{1}{l_{cs}} [A_1 \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_{t-\tau}^i\|_c + \gamma_c \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_{t-\tau}^i\|_c] && \text{(Assumptions 6.2 and 6.3)} \\
 &\leq \frac{A_1 + 1}{l_{cs}} [\|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_{t-\tau}^i\|_c] && (\gamma_c \leq 1; \text{triangle inequality)} \\
 &= \frac{A_1 + 1}{l_{cs}} [\Delta_t^i + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \Delta_{t-\tau}^i]. && (73)
 \end{aligned}$$

Substituting (72) and (73) in (71), we get

$$\begin{aligned}
 T_{32} &\leq \frac{L(A_1 + 1)}{\psi l_{cs}^2} \left[\frac{1}{N} \sum_{i=1}^N \zeta_3 \|\boldsymbol{\theta}_t\|_c \cdot \frac{1}{\zeta_3} [\Delta_t^i + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \Delta_{t-\tau}^i] \right. \\
 &\quad \left. + \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c \cdot [\Delta_t^i + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \Delta_{t-\tau}^i] \right] \\
 &\leq \frac{L(A_1 + 1)}{\psi l_{cs}^2} \left[\frac{1}{2} \zeta_3^2 \|\boldsymbol{\theta}_t\|_m^2 + \frac{1}{2N} \sum_{i=1}^N \frac{u_{cm}^2}{\zeta_3^2} [\Delta_t^i + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \Delta_{t-\tau}^i]^2 \right. \\
 &\quad \left. + \frac{1}{2} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{1}{2N} \sum_{i=1}^N [\Delta_t^i + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \Delta_{t-\tau}^i]^2 \right] && \text{(Young's inequality)}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{L(A_1 + 1)}{\psi l_{cs}^2} \left[\zeta_3^2 M(\boldsymbol{\theta}_t) + \left(\frac{3u_{cm}^2}{2\zeta_3^2} + 2 \right) \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) \frac{1}{N} \sum_{i=1}^N [(\Delta_t^i)^2 + (\Delta_{t-\tau}^i)^2] \right] \\
 &\leq \frac{L(A_1 + 1)}{\psi l_{cs}^2} \left[\zeta_3^2 M(\boldsymbol{\theta}_t) + \left(\frac{3u_{cm}^2}{2\zeta_3^2} + 2 \right) \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) [\Omega_t + \Omega_{t-\tau}] \right]. \tag{74}
 \end{aligned}$$

III. Bound on T_{33} : Taking expectation on both sides of T_{33} , we have

$$\begin{aligned}
 \mathbb{E}_{t-\tau}[T_{33}] &= \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbb{E}_{t-\tau}[\mathbf{G}(\boldsymbol{\Theta}_{t-\tau}, \mathbf{Y}_t)] - \bar{\mathbf{G}}(\boldsymbol{\Theta}_{t-\tau}) \rangle \\
 &= \frac{1}{N} \sum_{i=1}^N \langle \nabla M(\boldsymbol{\theta}_{t-\tau}), \mathbb{E}_{t-\tau}[\mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i)] - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i) \rangle \\
 &\leq \frac{1}{l_{cs}} \frac{1}{N} \sum_{i=1}^N \|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \cdot \|\mathbb{E}_{t-\tau}[\mathbf{G}^i(\boldsymbol{\theta}_{t-\tau}^i, \mathbf{y}_t^i)] - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_{t-\tau}^i)\|_c \tag{By (13)} \\
 &\leq \frac{1}{l_{cs}} \frac{1}{N} \sum_{i=1}^N \|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \cdot m_1 \|\boldsymbol{\theta}_{t-\tau}^i\|_c \rho^\tau \tag{Assumption 6.1} \\
 &\leq \frac{m_1 \alpha}{l_{cs}} \frac{1}{N} \sum_{i=1}^N \|\nabla M(\boldsymbol{\theta}_{t-\tau})\|_s^* \|\boldsymbol{\theta}_{t-\tau}^i\|_c \tag{assumption on τ } \\
 &\leq \frac{m_1 L \alpha}{l_{cs}^2 \psi} \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_{t-\tau}\|_c \|\boldsymbol{\theta}_{t-\tau}^i\|_c \tag{By Proposition B.1 and (13)} \\
 &\leq \frac{m_1 L \alpha}{l_{cs}^2 \psi} \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_{t-\tau}\|_c [\|\boldsymbol{\theta}_{t-\tau}\|_c + \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_{t-\tau}^i\|_c] \tag{triangle inequality} \\
 &\leq \frac{m_1 L \alpha}{l_{cs}^2 \psi} \left[\|\boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{1}{N} \sum_{i=1}^N \left[\frac{1}{2} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{1}{2} \|\boldsymbol{\theta}_{t-\tau} - \boldsymbol{\theta}_{t-\tau}^i\|_c^2 \right] \right] \tag{Young's inequality} \\
 &\leq \frac{m_1 L \alpha}{l_{cs}^2 \psi} \mathbb{E}_{t-\tau} \left[3 \left(\|\boldsymbol{\theta}_t\|_c^2 + \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 \right) + \frac{1}{2} \Omega_{t-\tau} \right] \tag{((a+b)^2 \leq 2a^2 + 2b^2)} \\
 &= \frac{m_1 L \alpha}{l_{cs}^2 \psi} \mathbb{E}_{t-\tau} \left[6u_{cm}^2 M(\boldsymbol{\theta}_t) + 3 \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + \frac{1}{2} \Omega_{t-\tau} \right]. \tag{75}
 \end{aligned}$$

Substituting the bounds on T_{31} , T_{32} , T_{33} from (70), (74), (75), in (66), we get the result. \square

Proof of Lemma B.4. Denote $T_4 = \langle \nabla M(\boldsymbol{\theta}_t), \bar{\mathbf{G}}(\boldsymbol{\Theta}_t) - \bar{\mathbf{G}}(\boldsymbol{\theta}_t) \rangle$. By the Cauchy–Schwarz inequality, we have

$$T_4 \leq \underbrace{\|\nabla M(\boldsymbol{\theta}_t)\|_s^*}_{T_{41}} \cdot \underbrace{\|\bar{\mathbf{G}}(\boldsymbol{\Theta}_t) - \bar{\mathbf{G}}(\boldsymbol{\theta}_t)\|_s}_{T_{42}}. \tag{76}$$

For T_{41} we have

$$\begin{aligned}
 \|\nabla M(\boldsymbol{\theta}_t)\|_s^* &= \|\nabla M(\boldsymbol{\theta}_t) - \nabla M(\mathbf{0})\|_s^* \\
 &\leq \frac{L}{\psi} \|\boldsymbol{\theta}_t\|_s \tag{Proposition B.1} \\
 &\leq \frac{L}{l_{cs} \psi} \|\boldsymbol{\theta}_t\|_c \tag{By (13)} \\
 &\leq \frac{L u_{cm}}{l_{cs} \psi} \|\boldsymbol{\theta}_t\|_m \tag{Proposition B.1} \\
 &= \frac{L u_{cm}}{l_{cs} \psi} \sqrt{2M(\boldsymbol{\theta}_t)}. \tag{77}
 \end{aligned}$$

For T_{42} we have

$$\begin{aligned} \|\bar{\mathbf{G}}(\boldsymbol{\Theta}_t) - \bar{\mathbf{G}}(\boldsymbol{\theta}_t)\|_s &= \left\| \frac{1}{N} \sum_{i=1}^N (\bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t)) \right\|_s \leq \frac{1}{N} \sum_{i=1}^N \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t)\|_s && \text{(Jensen's inequality)} \\ &\leq \frac{1}{l_{cs}N} \sum_{i=1}^N \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_t^i) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_t)\|_c \leq \frac{\gamma_c}{l_{cs}N} \sum_{i=1}^N \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c, && (78) \end{aligned}$$

where (78) follows from Assumption 6.2. Combining (77) and (78), for an arbitrary $\zeta_4 > 0$, we get

$$\begin{aligned} T_4 &\leq \frac{Lu_{cm}}{l_{cs}\psi} \sqrt{2M(\boldsymbol{\theta}_t)} \cdot \frac{\gamma_c}{l_{cs}N} \sum_{i=1}^N \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c \\ &\leq \frac{1}{N} \sum_{i=1}^N \zeta_4 \sqrt{2M(\boldsymbol{\theta}_t)} \cdot \frac{1}{\zeta_4} \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c \cdot \frac{Lu_{cm}}{l_{cs}^2\psi} && (\gamma_c \leq 1) \\ &\leq \frac{1}{N} \sum_{i=1}^N \left[\zeta_4^2 M(\boldsymbol{\theta}_t) + \frac{1}{2\zeta_4^2} \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c^2 \cdot \frac{L^2 u_{cm}^2}{l_{cs}^4 \psi^2} \right] && \text{(Young's inequality)} \\ &= \zeta_4^2 M(\boldsymbol{\theta}_t) + \frac{L^2 u_{cm}^2}{2l_{cs}^4 \zeta_4^2 \psi^2} \Omega_t. && (79) \end{aligned}$$

□

Proof of Lemma B.5. Denote $T_5 = \|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)\|_s^2$. We have

$$\begin{aligned} T_5 &= \|\mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t) + \mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t + \mathbf{b}(\mathbf{Y}_t)\|_s^2 \\ &\leq 3 \underbrace{\|\mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t) - \boldsymbol{\theta}_t\|_s^2}_{T_{51}} + 3 \underbrace{\|\mathbf{G}(\boldsymbol{\Theta}_t, \mathbf{Y}_t) - \mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t)\|_s^2}_{T_{52}} + 3 \|\mathbf{b}(\mathbf{Y}_t)\|_s^2. && (80) \end{aligned}$$

For T_{51} we have

$$\begin{aligned} T_{51} &\leq l_{cs}^{-2} (\|\mathbf{G}(\boldsymbol{\theta}_t, \mathbf{Y}_t)\|_c + \|\boldsymbol{\theta}_t\|_c)^2 && \text{(By (13))} \\ &\leq l_{cs}^{-2} (A_2 \|\boldsymbol{\theta}_t\|_c + \|\boldsymbol{\theta}_t\|_c)^2 && \text{(Assumption 6.3)} \\ &= \frac{2(A_2 + 1)^2 u_{cm}^2}{l_{cs}^2} M(\boldsymbol{\theta}_t). && (81) \end{aligned}$$

For T_{52} we have

$$\begin{aligned} T_{52} &= \left\| \frac{1}{N} \sum_{i=1}^N (\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_t, \mathbf{y}_t^i)) \right\|_s^2 \leq \frac{1}{N} \sum_{i=1}^N \|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_t, \mathbf{y}_t^i)\|_s^2 && (\|\cdot\|_s^2 \text{ is convex}) \\ &\leq \frac{1}{N l_{cs}^2} \sum_{i=1}^N \|\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\boldsymbol{\theta}_t, \mathbf{y}_t^i)\|_c^2 && \text{(By (13))} \\ &\leq \frac{A_1^2}{l_{cs}^2} \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_t^i - \boldsymbol{\theta}_t\|_c^2 && \text{(Assumption 6.3)} \\ &= \frac{A_1^2}{l_{cs}^2} \Omega_t. && (82) \end{aligned}$$

Using the bounds in (81), (82), we get

$$T_5 \leq \frac{6(A_2 + 1)^2 u_{cm}^2}{l_{cs}^2} M(\boldsymbol{\theta}_t) + \frac{3A_1^2}{l_{cs}^2} \Omega_t + 3 \|\mathbf{b}(\mathbf{Y}_t)\|_s^2, \quad (83)$$

where the last inequality is by the assumption on τ .

□

Proof of Lemma B.6. Define $\theta_l = \frac{1}{N} \sum_{i=1}^N \|\theta_l^i\|_c$. By the update rule of Algorithm 4, if $l + 1 \bmod K \neq 0$, we have

$$\begin{aligned}
 \theta_{l+1} &= \frac{1}{N} \sum_{i=1}^N \|\theta_{l+1}^i\|_c = \frac{1}{N} \sum_{i=1}^N \|\theta_l^i + \alpha(\mathbf{G}^i(\theta_l^i, \mathbf{y}_l^i) - \theta_l^i + \mathbf{b}^i(\mathbf{y}_l^i))\|_c \\
 &\leq \theta_l + \frac{\alpha}{N} \sum_{i=1}^N [\|\mathbf{G}^i(\theta_l^i, \mathbf{y}_l^i)\|_c + \|\theta_l^i\|_c + \|\mathbf{b}^i(\mathbf{y}_l^i)\|_c] && \text{(triangle inequality)} \\
 &\leq \theta_l + \alpha \frac{1}{N} \sum_{i=1}^N [(A_2 + 1)\|\theta_l^i\|_c + B] && \text{(Assumption 6.3)} \\
 &= (1 + \alpha(A_2 + 1))\theta_l + \alpha B. && (84)
 \end{aligned}$$

Furthermore, if $l + 1 \bmod K = 0$, we have $\theta_{l+1}^j = \frac{1}{N} \sum_{j=1}^N (\theta_l^j + \alpha(\mathbf{G}^j(\theta_l^j, \mathbf{y}_l^j) - \theta_l^j + \mathbf{b}^j(\mathbf{y}_l^j)))$, and hence

$$\begin{aligned}
 \theta_{l+1} &= \frac{1}{N} \sum_{i=1}^N \|\theta_{l+1}^i\|_c = \frac{1}{N} \sum_{i=1}^N \left\| \frac{1}{N} \sum_{j=1}^N (\theta_l^j + \alpha(\mathbf{G}^j(\theta_l^j, \mathbf{y}_l^j) - \theta_l^j + \mathbf{b}^j(\mathbf{y}_l^j))) \right\|_c \\
 &\leq \theta_l + \frac{\alpha}{N} \sum_{i=1}^N [\|\mathbf{G}^i(\theta_l^i, \mathbf{y}_l^i)\|_c + \|\theta_l^i\|_c + \|\mathbf{b}^i(\mathbf{y}_l^i)\|_c], && \text{(triangle inequality)}
 \end{aligned}$$

and the same bound as in (84) holds. By recursive application of (84), we get

$$\begin{aligned}
 \theta_{l+1} &\leq (1 + \alpha(A_2 + 1))^{l+1} \theta_0 + \alpha B \sum_{\ell=0}^l (1 + \alpha(A_2 + 1))^\ell \\
 &= (1 + \alpha(A_2 + 1))^{l+1} \theta_0 + \alpha B \frac{(1 + \alpha(A_2 + 1))^{l+1} - 1}{\alpha(A_2 + 1)}. && (85)
 \end{aligned}$$

Notice that for $x \leq \frac{\log 2}{\tau}$, we have $(1 + x)^{\tau+1} \leq 1 + 2x(\tau + 1)$. If $0 \leq l \leq 2\tau - 1$, by the assumption on α , we have $(1 + \alpha(A_2 + 1))^{l+1} \leq (1 + \alpha(A_2 + 1))^{2\tau} \leq 1 + 4\alpha(A_2 + 1)\tau \leq 2$. Hence, we have

$$\theta_l \leq 2\theta_0 + \alpha B \frac{4\alpha(A_2 + 1)\tau}{\alpha(A_2 + 1)} = 2\theta_0 + 4\alpha B\tau. \quad (86)$$

Furthermore, we have

$$\begin{aligned}
 \|\theta_{l+1} - \theta_l\|_c &= \alpha \|\mathbf{G}(\Theta_l, \mathbf{y}_l) - \theta_l + \mathbf{b}(\mathbf{y}_l)\|_c \\
 &\leq \alpha \|\mathbf{G}(\Theta_l, \mathbf{y}_l) - \theta_l\|_c + \alpha \|\mathbf{b}(\mathbf{y}_l)\|_c && \text{(triangle inequality)} \\
 &= \alpha \left\| \frac{1}{N} \sum_{i=1}^N (\mathbf{G}^i(\theta_l^i, \mathbf{y}_l^i) - \theta_l^i) \right\|_c + \alpha \|\mathbf{b}(\mathbf{y}_l)\|_c \\
 &\leq \alpha \frac{1}{N} \sum_{i=1}^N \|\mathbf{G}^i(\theta_l^i, \mathbf{y}_l^i) - \theta_l^i\|_c + \alpha B && \text{(convexity of norm)} \\
 &\leq \alpha \frac{1}{N} \sum_{i=1}^N (A_2 + 1) \|\theta_l^i\|_c + \alpha B && \text{(Assumption 6.3)} \\
 &= \alpha(A_2 + 1)\theta_l + \alpha B. && (87)
 \end{aligned}$$

Suppose $0 \leq t \leq 2\tau$. We have

$$\|\theta_t - \theta_0\|_c \leq \left[\sum_{k=0}^{t-1} \|\theta_{k+1} - \theta_k\|_c \right] \quad \text{(triangle inequality)}$$

$$\leq \left[\sum_{k=0}^{t-1} \alpha(A_2 + 1)\theta_k + \alpha B \right] \quad (\text{By (87)})$$

$$\leq B\alpha t + \alpha(A_2 + 1) \sum_{k=0}^{t-1} (2\theta_0 + 4B\alpha\tau) \quad (\text{By (86)})$$

$$\leq 2\alpha\tau(B + (A_2 + 1)(2\theta_0 + 4B\alpha\tau)) \quad (t \leq 2\tau)$$

$$\leq \frac{1}{C_1} \left(B + (A_2 + 1) \left(\theta_0 + \frac{B}{2C_1} \right) \right). \quad (88)$$

Furthermore, by Proposition B.1, we have $M(\boldsymbol{\theta}_t) = \frac{1}{2}\|\boldsymbol{\theta}_t\|_m^2$, and hence for any $0 \leq t \leq 2\tau$ we have

$$\begin{aligned} M(\boldsymbol{\theta}_t) &\leq \frac{1}{2l_{cm}^2} \|\boldsymbol{\theta}_t\|_c^2 \\ &= \frac{1}{2l_{cm}^2} (\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_0 + \boldsymbol{\theta}_0\|_c)^2 \\ &\leq \frac{1}{2l_{cm}^2} (\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_0\|_c + \|\boldsymbol{\theta}_0\|_c)^2 \quad (\text{triangle inequality}) \\ &\leq \frac{1}{l_{cm}^2} (\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_0\|_c^2 + \|\boldsymbol{\theta}_0\|_c^2) \\ &\leq \frac{1}{l_{cm}^2} \left(\frac{1}{C_1^2} \left(B + (A_2 + 1) \left(\|\boldsymbol{\theta}_0\|_c + \frac{B}{2C_1} \right) \right)^2 + \|\boldsymbol{\theta}_0\|_c^2 \right), \end{aligned}$$

which proves the first claim.

Next we prove the second claim. By the update rule in 30, we have

$$\begin{aligned} \|\boldsymbol{\theta}_{l+1} - \boldsymbol{\theta}_l\|_c^2 &= \alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \boldsymbol{\theta}_l + \mathbf{b}(\mathbf{Y}_l)\|_c^2 \\ &\leq 3\alpha^2 \|\mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l) - \boldsymbol{\theta}_l\|_c^2 + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &\leq 6\alpha^2 (\|\mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 + \|\boldsymbol{\theta}_l\|_c^2) + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &= 6\alpha^2 \left(\left\| \frac{1}{N} \sum_{i=1}^N \mathbf{G}^i(\boldsymbol{\theta}_l, \mathbf{y}_l^i) \right\|_c^2 + \|\boldsymbol{\theta}_l\|_c^2 \right) + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &\leq 6\alpha^2 \left(\left(\frac{1}{N} \sum_{i=1}^N \|\mathbf{G}^i(\boldsymbol{\theta}_l, \mathbf{y}_l^i)\|_c \right)^2 + \|\boldsymbol{\theta}_l\|_c^2 \right) + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &\quad (\text{convexity of norm}) \\ &\leq 6\alpha^2 \left(\left(\frac{1}{N} \sum_{i=1}^N A_2 \|\boldsymbol{\theta}_l\|_c \right)^2 + \|\boldsymbol{\theta}_l\|_c^2 \right) + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &\quad (\text{Assumption 6.3}) \\ &= 6\alpha^2 (A_2^2 + 1) \|\boldsymbol{\theta}_l\|_c^2 + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 \|\mathbf{G}(\boldsymbol{\Theta}_l, \mathbf{Y}_l) - \mathbf{G}(\boldsymbol{\theta}_l, \mathbf{Y}_l)\|_c^2 \\ &\leq 6\alpha^2 (A_2^2 + 1) \|\boldsymbol{\theta}_l\|_c^2 + 3\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 3\alpha^2 A_1^2 \Delta_l^2, \quad (89) \end{aligned}$$

where (89) follows from Lemma B.10. Taking square root on both sides, we get

$$\|\boldsymbol{\theta}_{l+1} - \boldsymbol{\theta}_l\|_c \leq 3\alpha \sqrt{A_2^2 + 1} \|\boldsymbol{\theta}_l\|_c + 2\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c + 2\alpha A_1 \Delta_l. \quad (90)$$

Combining the above inequality with the fact that $\|\boldsymbol{\theta}_{l+1}\|_c - \|\boldsymbol{\theta}_l\|_c \leq \|\boldsymbol{\theta}_{l+1} - \boldsymbol{\theta}_l\|_c$, we get

$$\|\boldsymbol{\theta}_{l+1}\|_c \leq (1 + 3\alpha \sqrt{A_2^2 + 1}) \|\boldsymbol{\theta}_l\|_c + 2\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c + 2\alpha A_1 \Delta_l$$

$$=(1 + \alpha C_1)\|\boldsymbol{\theta}_l\|_c + 2\alpha\|\mathbf{b}(\mathbf{Y}_l)\|_c + 2\alpha A_1 \Delta_l, \quad (91)$$

where we denote $C_1 = 3\sqrt{A_2^2 + 1}$. Assuming $t - \tau \leq l \leq t$, and taking expectation on both sides, we have

$$\begin{aligned} \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{l+1}\|_c] &\leq (1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] + 2\alpha\mathbb{E}_{t-2\tau}[\|\mathbf{b}(\mathbf{Y}_l)\|_c] + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l] \\ &\leq (1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] + 2\alpha \frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} + 2m_2\rho^{l-t+2\tau} \right] + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l] \quad (\text{Lemma B.9}) \\ &\leq (1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] + 2\alpha \frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} + 2m_2\alpha\rho^{l-t+\tau} \right] + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l] \quad (\text{Assumption on } \tau) \\ &= (1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] + \alpha c_t(l) + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l], \end{aligned} \quad (92)$$

where $c_t(l) = 2\frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} + 2m_2\alpha\rho^{l-t+\tau} \right]$. By applying this inequality recursively, we have

$$\begin{aligned} \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{l+1}\|_c] &\leq (1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] + \alpha c_t(l) + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l] \\ &\leq (1 + \alpha C_1) \left[(1 + \alpha C_1)\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{l-1}\|_c] + \alpha c_t(l-1) + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_{l-1}] \right] + \alpha c_t(l) + 2\alpha A_1 \mathbb{E}_{t-2\tau}[\Delta_l] \\ &\leq (1 + \alpha C_1)^{l+1-t+\tau} \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{t-\tau}\|_c] + \alpha \sum_{k=t-\tau}^l (1 + \alpha C_1)^{l-k} c_t(k) + 2\alpha A_1 \mathbb{E}_{t-2\tau} \left[\sum_{k=t-\tau}^l (1 + \alpha C_1)^{l-k} \Delta_k \right] \\ &\leq (1 + \alpha C_1)^{\tau+1} \mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{t-\tau}\|_c] + \underbrace{\alpha \sum_{k=t-\tau}^t (1 + \alpha C_1)^{t-k} c_t(k)}_{T_1} + 2\alpha A_1 \mathbb{E}_{t-2\tau} \left[\underbrace{\sum_{k=t-\tau}^t (1 + \alpha C_1)^{t-k} \Delta_k}_{T_2} \right]. \end{aligned} \quad (93)$$

We study T_1 and T_2 in (93) separately. For T_1 we have

$$\begin{aligned} T_1 &= 2\frac{u_{cD}}{l_{cD}} \sum_{k=0}^{\tau} (1 + \alpha C_1)^{\tau-k} \left[\frac{B}{\sqrt{N}} + 2m_2\alpha\rho^k \right] \\ &= 2\frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} \frac{(1 + \alpha C_1)^{\tau+1} - 1}{\alpha C_1} + 2m_2\alpha(1 + \alpha C_1)^\tau \sum_{k=0}^{\tau} \left(\frac{\rho}{1 + \alpha C_1} \right)^k \right] \\ &\leq 2\frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} \frac{(1 + \alpha C_1)^{\tau+1} - 1}{\alpha C_1} + 2m_2\alpha(1 + \alpha C_1)^\tau \sum_{k=0}^{\tau} \rho^k \right] \quad (\alpha > 0) \\ &\leq 2\frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} \frac{(1 + \alpha C_1)^{\tau+1} - 1}{\alpha C_1} + 2m_2\alpha(1 + \alpha C_1)^\tau \frac{1}{1 - \rho} \right]. \end{aligned} \quad (94)$$

Notice that for $x \leq \frac{\log 2}{\tau}$, we have $(1 + x)^{\tau+1} \leq 1 + 2x(\tau + 1)$. By the assumption on α , we have $(1 + \alpha C_1)^{\tau+1} \leq 1 + 2\alpha C_1(\tau + 1) \leq 1 + 4\alpha\tau C_1 \leq 2$ and $(1 + \alpha C_1)^\tau \leq 1 + 2\alpha C_1\tau \leq 1 + 1/2 \leq 2$. Hence, we have

$$T_1 \leq 2\frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} 2(\tau + 1) + \frac{4m_2\alpha}{1 - \rho} \right]. \quad (95)$$

Furthermore, for the term T_2 we have

$$\begin{aligned} T_2 &= \sum_{k=0}^{\tau} (1 + \alpha C_1)^{\tau-k} \Delta_{t-\tau+k} \leq \sum_{k=0}^{\tau} (1 + \alpha C_1)^\tau \Delta_{t-\tau+k} \quad (\text{due to } \alpha > 0) \\ &\leq \sum_{k=0}^{\tau} (1 + 2\alpha C_1\tau) \Delta_{t-\tau+k} \leq 2 \sum_{k=0}^{\tau} \Delta_{t-k}. \end{aligned} \quad (96)$$

Substituting (95), (96) in (93), for every $t - \tau \leq l \leq t$, we get

$$\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_l\|_c] \leq 2\mathbb{E}_{t-2\tau}[\|\boldsymbol{\theta}_{t-\tau}\|_c] + \alpha \frac{2u_{cD}}{l_{cD}} \left[\frac{4B}{\sqrt{N}}\tau + \frac{4m_2\alpha}{1 - \rho} \right] + 4A_1\alpha \sum_{k=0}^{\tau} \Delta_{t-k}. \quad (97)$$

But we have

$$\begin{aligned}
 \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c &\leq \mathbb{E}_{t-2\tau} \left[\sum_{i=t-\tau}^{t-1} \|\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i\|_c \right] && \text{(triangle inequality)} \\
 &\leq \sum_{i=t-\tau}^{t-1} \mathbb{E}_{t-2\tau} [\alpha C_1 \|\boldsymbol{\theta}_i\|_c + 2\alpha \|\mathbf{b}(\mathbf{Y}_i)\|_c + 2\alpha A_1 \Delta_i] && \text{(by (90))} \\
 &\leq \sum_{i=t-\tau}^{t-1} \alpha C_1 \left[2\mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_{t-\tau}\|_c] + \alpha \frac{2u_{cD}}{l_{cD}} \left[\frac{4B}{\sqrt{N}}\tau + \frac{4\alpha m_2}{1-\rho} \right] + 4A_1 \alpha \sum_{j=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-j}] \right] && \text{(by (97))} \\
 &\quad + 2\alpha \sum_{i=t-\tau}^{t-1} \frac{u_{cD}}{l_{cD}} \left[\frac{B}{\sqrt{N}} + 2m_2 \rho^{i-t+2\tau} \right] && \text{(Lemma B.9)} \\
 &\quad + 2\alpha A_1 \sum_{i=t-\tau}^{t-1} \mathbb{E}_{t-2\tau} [\Delta_i] \\
 &\leq \alpha \tau C_1 \left[2\mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_{t-\tau}\|_c] + \alpha \frac{2u_{cD}}{l_{cD}} \left[\frac{4B}{\sqrt{N}}\tau + \frac{4\alpha m_2}{1-\rho} \right] + 4A_1 \alpha \sum_{j=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-j}] \right] \\
 &\quad + 2\alpha \tau \frac{u_{cD}}{l_{cD}} \frac{B}{\sqrt{N}} + 4\alpha^2 \frac{u_{cD}}{l_{cD}} m_2 \frac{1}{1-\rho} + 2\alpha A_1 \sum_{i=t-\tau}^{t-1} \mathbb{E}_{t-2\tau} [\Delta_i] && \text{(assumption on } \tau) \\
 &\leq 2\alpha \tau C_1 \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_{t-\tau}\|_c] + \frac{u_{cD}}{l_{cD}} \frac{B}{\sqrt{N}} [8C_1 \alpha^2 \tau^2 + 2\alpha \tau] + \frac{u_{cD}}{l_{cD}} \left[\frac{8m_2 C_1 \alpha^2 \tau}{1-\rho} + \frac{4m_2 \alpha^2}{1-\rho} \right] \\
 &\quad + (4A_1 \alpha^2 \tau C_1 + 2\alpha A_1) \sum_{i=t-\tau}^t \mathbb{E}_{t-2\tau} [\Delta_i] && (\Delta_i \geq 0) \\
 &\leq 2\alpha \tau C_1 \mathbb{E}_{t-2\tau} [\|\boldsymbol{\theta}_{t-\tau}\|_c] + 4\alpha \tau \frac{u_{cD}}{l_{cD}} \frac{B}{\sqrt{N}} + \frac{u_{cD}}{l_{cD}} \frac{4m_2}{1-\rho} \alpha^2 (1 + 2C_1 \tau) + 3A_1 \alpha \sum_{i=t-\tau}^t \mathbb{E}_{t-2\tau} [\Delta_i]. && (98)
 \end{aligned}$$

Furthermore, by triangle inequality, we have $\|\boldsymbol{\theta}_{t-\tau}\|_c \leq \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \|\boldsymbol{\theta}_t\|_c$. By assumption on α , we have $2\alpha \tau C_1 \|\boldsymbol{\theta}_{t-\tau}\|_c \leq 2\alpha \tau C_1 \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + 2\alpha \tau C_1 \|\boldsymbol{\theta}_t\|_c \leq 0.5\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + 2\alpha \tau C_1 \|\boldsymbol{\theta}_t\|_c$. By taking expectation on both sides, and substituting it in (98), we get (34). \square

Proof of Lemma B.7. From (91) we have

$$\begin{aligned}
 \|\boldsymbol{\theta}_{l+1}\|_c^2 &\leq (1 + \alpha C_1)^2 \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 4\alpha^2 A_1^2 \Delta_l^2 \\
 &\quad + \underbrace{4\alpha(1 + \alpha C_1) \|\boldsymbol{\theta}_l\|_c \|\mathbf{b}(\mathbf{Y}_l)\|_c}_{T_1} + \underbrace{4\alpha A_1 (1 + \alpha C_1) \|\boldsymbol{\theta}_l\|_c \Delta_l}_{T_2} + \underbrace{8\alpha^2 A_1 \Delta_l \|\mathbf{b}(\mathbf{Y}_l)\|_c}_{T_3}. && (99)
 \end{aligned}$$

For T_1 we have

$$\begin{aligned}
 T_1 &= 2\sqrt{\alpha(1 + \alpha C_1)} \|\boldsymbol{\theta}_l\|_c \cdot 2\sqrt{\alpha(1 + \alpha C_1)} \|\mathbf{b}(\mathbf{Y}_l)\|_c \\
 &\leq 2\alpha(1 + \alpha C_1) \|\boldsymbol{\theta}_l\|_c^2 + 2\alpha(1 + \alpha C_1) \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 && (ab \leq \frac{1}{2}a^2 + \frac{1}{2}b^2) \\
 &\leq 4\alpha \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c^2, && (100)
 \end{aligned}$$

where the last inequality is by the assumption on α . Analogously for T_2 we have

$$\begin{aligned}
 T_2 &= 2\sqrt{\alpha(1 + \alpha C_1)} \|\boldsymbol{\theta}_l\|_c \cdot 2A_1 \sqrt{\alpha(1 + \alpha C_1)} \Delta_l \\
 &\leq 4\alpha \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha A_1^2 \Delta_l^2. && (101)
 \end{aligned}$$

For T_3 we have

$$\begin{aligned} T_3 &= 2\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c \cdot 4\alpha A_1 \Delta_l \\ &\leq 2\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 8\alpha^2 A_1^2 \Delta_l^2. \end{aligned} \quad (102)$$

Combining the bounds in (100), (101), and (102), and noting that $(1 + \alpha C_1)^2 \leq 1 + 3\alpha C_1$, from (99) we have

$$\begin{aligned} \|\boldsymbol{\theta}_{l+1}\|_c^2 &\leq (1 + 3\alpha C_1) \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 4\alpha^2 A_1^2 \Delta_l^2 \\ &\quad + 4\alpha \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 4\alpha \|\boldsymbol{\theta}_l\|_c^2 + 4\alpha A_1^2 \Delta_l^2 + 2\alpha^2 \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 8\alpha^2 A_1^2 \Delta_l^2 \\ &= (1 + \alpha(3C_1 + 8)) \|\boldsymbol{\theta}_l\|_c^2 + (6\alpha^2 + 4\alpha) \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + A_1^2 (12\alpha^2 + 4\alpha) \Delta_l^2 \\ &\leq (1 + \alpha C_2) \|\boldsymbol{\theta}_l\|_c^2 + 10\alpha \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 16\alpha A_1^2 \Delta_l^2. \end{aligned} \quad (103)$$

where $C_2 = 3C_1 + 8$. Taking expectation on both sides, we have

$$\begin{aligned} \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{l+1}\|_c^2 &\leq (1 + \alpha C_2) \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_l\|_c^2 + 10\alpha \mathbb{E}_{t-2\tau} \|\mathbf{b}(\mathbf{Y}_l)\|_c^2 + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} [\Delta_l^2] \\ &\leq (1 + \alpha C_2) \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_l\|_c^2 + 10\alpha \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} + 2m_2^2 \rho^{2(l-t+2\tau)} \right] + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} [\Delta_l^2] \quad (\text{Lemma B.9}) \\ &= (1 + \alpha C_2) \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_l\|_c^2 + \alpha \bar{c}_t(l) + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} [\Delta_l^2], \end{aligned} \quad (104)$$

where $\bar{c}_t(l) = 10 \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} + 2m_2^2 \alpha^2 \rho^{2(l-t+\tau)} \right]$. Hence, for any $t - \tau \leq l \leq t$, we have

$$\begin{aligned} \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{l+1}\|_c^2 &\leq (1 + \alpha C_2) \left[(1 + \alpha C_2) \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{l-1}\|_c^2 + \alpha \bar{c}_t(l-1) + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} [\Delta_{l-1}^2] \right] \\ &\quad + \alpha \bar{c}_t(l) + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} [\Delta_l^2] \\ &\leq (1 + \alpha C_2)^{l+1-t+\tau} \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + \alpha \sum_{i=t-\tau}^l (1 + \alpha C_2)^{l-i} \bar{c}_t(i) + 16\alpha A_1^2 \mathbb{E}_{t-2\tau} \left[\sum_{i=t-\tau}^l (1 + \alpha C_2)^{l-i} \Delta_i^2 \right] \\ &\leq (1 + \alpha C_2)^{\tau+1} \mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + \underbrace{\alpha \sum_{i=t-\tau}^t (1 + \alpha C_2)^{t-i} \bar{c}_t(i)}_{T_4} + \underbrace{16\alpha A_1^2 \mathbb{E}_{t-2\tau} \left[\sum_{i=t-\tau}^t (1 + \alpha C_2)^{t-i} \Delta_i^2 \right]}_{T_5}. \end{aligned} \quad (105)$$

For the term T_4 we have

$$\begin{aligned} T_4 &= 10 \frac{u_{cD}^2}{l_{cD}^2} \sum_{i=0}^{\tau} (1 + \alpha C_2)^{\tau-i} \left[\frac{B^2}{N} + 2m_2^2 \alpha^2 \rho^{2i} \right] \\ &= 10 \frac{u_{cD}^2}{l_{cD}^2} \frac{B^2}{N} \frac{(1 + \alpha C_2)^{\tau+1} - 1}{\alpha C_2} + 20 \frac{u_{cD}^2}{l_{cD}^2} m_2^2 \alpha^2 (1 + \alpha C_2)^{\tau} \sum_{i=0}^{\tau} \left(\frac{\rho^2}{1 + \alpha C_2} \right)^i \\ &\leq 10 \frac{u_{cD}^2}{l_{cD}^2} \frac{B^2}{N} \frac{(1 + \alpha C_2)^{\tau+1} - 1}{\alpha C_2} + 20 \frac{u_{cD}^2}{l_{cD}^2} m_2^2 \alpha^2 (1 + \alpha C_2)^{\tau} \sum_{i=0}^{\tau} \rho^{2i} \quad (\alpha C_2 \geq 0) \\ &\leq 10 \frac{u_{cD}^2}{l_{cD}^2} \frac{B^2}{N} \frac{(1 + \alpha C_2)^{\tau+1} - 1}{\alpha C_2} + 20 \frac{u_{cD}^2}{l_{cD}^2} m_2^2 \alpha^2 (1 + \alpha C_2)^{\tau} \frac{1}{1 - \rho^2}. \end{aligned} \quad (106)$$

By the same argument as in Lemma B.6, and by the assumption on α , we have $(1 + \alpha C_2)^{\tau+1} \leq 1 + 2\alpha C_2(\tau + 1) \leq 1 + 4\alpha\tau C_2 \leq 2$ and $(1 + \alpha C_2)^{\tau} \leq 1 + 2\alpha C_2\tau \leq 1 + 1/2 \leq 2$. Hence, we have

$$T_4 \leq 40 \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} \tau + \frac{m_2^2 \alpha^2}{1 - \rho^2} \right]. \quad (107)$$

Furthermore, for T_5 , we have

$$T_5 = \sum_{i=0}^{\tau} (1 + \alpha C_2)^{\tau-i} \mathbb{E}_{t-2\tau} [\Delta_{t-\tau+i}^2]$$

$$\begin{aligned}
 &\leq \sum_{i=0}^{\tau} (1 + \alpha C_2)^\tau \mathbb{E}_{t-2\tau} [\Delta_{t-\tau+i}^2] && (\alpha C_2 > 0) \\
 &\leq \sum_{i=0}^{\tau} (1 + 2\alpha C_2 \tau) \mathbb{E}_{t-2\tau} [\Delta_{t-\tau+i}^2] && (\text{assumption on } \alpha) \\
 &\leq 2 \sum_{i=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-i}^2]. && (108)
 \end{aligned}$$

Combining the bounds on T_4 (107) and T_5 (108) in (105) we have for any $t - \tau \leq l \leq t$,

$$\mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_l\|_c^2 \leq 2\mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + 40 \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} \alpha \tau + m_2^2 \alpha^3 \frac{1}{1 - \rho^2} \right] + 32\alpha A_1^2 \sum_{i=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-i}^2]. \quad (109)$$

Furthermore, we have

$$\begin{aligned}
 \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 &\leq \left(\sum_{i=t-\tau}^{t-1} \|\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i\|_c \right)^2 && (\text{triangle inequality}) \\
 &\leq \tau \sum_{i=t-\tau}^{t-1} \|\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i\|_c^2 \\
 &\leq \tau \sum_{i=t-\tau}^{t-1} [\alpha^2 C_1^2 \|\boldsymbol{\theta}_i\|_c^2 + 3\alpha^2 \|\mathbf{b}(\mathbf{y}_i)\|_c^2 + 3\alpha^2 A_1^2 \Delta_i^2]. && (\text{from (89)})
 \end{aligned}$$

Taking expectation on both sides, and using the bounds in Lemma B.9 and (109), we get

$$\begin{aligned}
 &\mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 \\
 &\leq \tau \alpha^2 C_1^2 \sum_{i=t-\tau}^{t-1} \left[2\mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + 40 \frac{u_{cD}^2}{l_{cD}^2} \left(\frac{B^2}{N} \alpha \tau + \frac{m_2^2 \alpha^3}{1 - \rho^2} \right) + 32\alpha A_1^2 \sum_{i=0}^{\tau} \mathbb{E}_{t-2\tau} [\Delta_{t-i}^2] \right] \\
 &\quad + 3\alpha^2 \tau \sum_{i=t-\tau}^{t-1} \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{B^2}{N} + 2m_2^2 \rho^{2(i-t+2\tau)} \right] + 3\alpha^2 A_1^2 \tau \sum_{i=t-\tau}^{t-1} \mathbb{E}_{t-2\tau} [\Delta_i^2] \\
 &\leq \tau^2 \alpha^2 C_1^2 \left[2\mathbb{E}_{t-2\tau} \|\boldsymbol{\theta}_{t-\tau}\|_c^2 + 40 \frac{u_{cD}^2}{l_{cD}^2} \left(\frac{B^2}{N} \alpha \tau + \frac{m_2^2 \alpha^3}{1 - \rho^2} \right) \right] + 3\alpha^2 \tau \frac{u_{cD}^2}{l_{cD}^2} \left[\frac{\tau B^2}{N} + \frac{2m_2^2 \alpha^2}{1 - \rho^2} \right] && (\rho^\tau \leq \alpha) \\
 &\quad + \alpha^2 A_1^2 \tau (3 + 32\alpha \tau C_1^2) \sum_{i=0}^{\tau} \mathbb{E}_{t-2\tau} \Delta_{t-i}^2. && (110)
 \end{aligned}$$

Furthermore, by triangle inequality, we have $\|\boldsymbol{\theta}_{t-\tau}\|_c \leq \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \|\boldsymbol{\theta}_t\|_c$. Squaring both sides, we have $\|\boldsymbol{\theta}_{t-\tau}\|_c^2 \leq (\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c + \|\boldsymbol{\theta}_t\|_c)^2 \leq 2\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + 2\|\boldsymbol{\theta}_t\|_c^2$. By assumption on α , we have $2\tau^2 \alpha^2 C_1^2 \|\boldsymbol{\theta}_{t-\tau}\|_c^2 \leq 4\tau^2 \alpha^2 C_1^2 \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + 4\tau^2 \alpha^2 C_1^2 \|\boldsymbol{\theta}_t\|_c^2 \leq 0.5\|\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-\tau}\|_c^2 + 4\tau^2 \alpha^2 C_1^2 \|\boldsymbol{\theta}_t\|_c^2$. \square

Proof of Lemma B.8. For $sK + 1 \leq t \leq (s+1)K - 1$, where $s = \lfloor t/K \rfloor$,

$$\begin{aligned}
 \Omega_t &\triangleq \frac{1}{N} \sum_{i=1}^N (\Delta_i^t)^2 = \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{\theta}_t - \boldsymbol{\theta}_i^t\|_c^2 \\
 &= \frac{1}{N} \sum_{i=1}^N \left\| \left(\boldsymbol{\theta}_{sK}^i - \boldsymbol{\theta}_{sK} \right) + \alpha \sum_{t'=sK}^{t-1} \left[\left(\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \boldsymbol{\theta}_{t'} + \mathbf{b}(\mathbf{y}_{t'}) \right) - \left(\mathbf{G}(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \boldsymbol{\theta}_{t'} + \mathbf{b}(\mathbf{y}_{t'}) \right) \right] \right\|_c^2 \\
 &\leq \frac{2\alpha^2}{N} \sum_{i=1}^N \left[\left\| \sum_{t'=sK}^{t-1} [\mathbf{b}(\mathbf{y}_{t'}^i) - \mathbf{b}(\mathbf{y}_{t'})] \right\|_c^2 + \left\| \sum_{t'=sK}^{t-1} \left[\left(\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \boldsymbol{\theta}_{t'} \right) - \left(\mathbf{G}(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \boldsymbol{\theta}_{t'} \right) \right] \right\|_c^2 \right] \\
 &\quad (\because \boldsymbol{\theta}_{sK}^i = \boldsymbol{\theta}_{sK} \text{ and Young's inequality})
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{2\alpha^2 u_{c2}^2}{N} \sum_{i=1}^N \left[\sum_{t'=sK}^{t-1} \|\mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'})\|_2^2 + \sum_{\substack{t', t''=sK \\ t' \neq t''}}^{t-1} \langle \mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'}), \mathbf{b}^i(\mathbf{y}_{t''}) - \mathbf{b}(\mathbf{y}_{t''}) \rangle \right] \\
 &\quad + \frac{4\alpha^2}{N} (t-sK) \sum_{t'=sK}^{t-1} \sum_{i=1}^N \left[u_{c2}^2 \|\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \mathbf{G}(\boldsymbol{\Theta}_{t'}, \mathbf{y}_{t'})\|_2^2 + \|\boldsymbol{\theta}_{t'}^i - \boldsymbol{\theta}_{t'}\|_c^2 \right] \\
 &\hspace{15em} \text{(Triangle inequality, and } \|\cdot\|_c \leq u_{c2} \|\cdot\|_2) \\
 &\leq \frac{2\alpha^2 u_{c2}^2}{N} \sum_{i=1}^N \left[\sum_{t'=sK}^{t-1} \frac{1}{l_{c2}^2} \|\mathbf{b}^i(\mathbf{y}_{t'})\|_c^2 + \sum_{\substack{t', t''=sK \\ t' \neq t''}}^{t-1} \langle \mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'}), \mathbf{b}^i(\mathbf{y}_{t''}) - \mathbf{b}(\mathbf{y}_{t''}) \rangle \right] \\
 &\quad + \frac{4\alpha^2}{N} (t-sK) \sum_{t'=sK}^{t-1} \sum_{i=1}^N \left[\frac{u_{c2}^2}{l_{c2}^2} \|\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'})\|_c^2 + \|\boldsymbol{\theta}_{t'}^i - \boldsymbol{\theta}_{t'}\|_c^2 \right]. \quad (l_{c2} \|\cdot\|_2 \leq \|\cdot\|_c \text{ and } \text{Var}(X) \leq E[X^2])
 \end{aligned}$$

Taking expectation

$$\begin{aligned}
 \mathbb{E}[\Omega_t] &\leq \frac{2\alpha^2 u_{c2}^2}{l_{c2}^2} (t-sK) B^2 + \frac{4\alpha^2 u_{c2}^2}{N} \sum_{i=1}^N \sum_{\substack{t', t''=sK \\ t' < t''}}^{t-1} \mathbb{E} \langle \mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'}), \mathbb{E}_{t'} [\mathbf{b}^i(\mathbf{y}_{t''}) - \mathbf{b}(\mathbf{y}_{t''})] \rangle \\
 &\quad + \frac{4\alpha^2}{N} (t-sK) \sum_{t'=sK}^{t-1} \sum_{i=1}^N \mathbb{E} \left[\frac{2u_{c2}^2}{l_{c2}^2} \left\{ \|\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'}) - \mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'})\|_c^2 + \|\mathbf{G}^i(\boldsymbol{\theta}_{t'}, \mathbf{y}_{t'})\|_c^2 \right\} + \|\boldsymbol{\theta}_{t'}^i - \boldsymbol{\theta}_{t'}\|_c^2 \right] \\
 &\hspace{15em} \text{(using Assumption 6.3)} \\
 &\leq \frac{2\alpha^2 u_{c2}^2}{l_{c2}^2} (t-sK) B^2 + \frac{4\alpha^2 u_{c2}^2}{l_{c2} N} \sum_{i=1}^N \sum_{\substack{t', t''=sK \\ t' < t''}}^{t-1} \mathbb{E} [\|\mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'})\|_2 \|\mathbb{E}_{t'} [\mathbf{b}^i(\mathbf{y}_{t''}) - \mathbf{b}(\mathbf{y}_{t''})]\|_c] \\
 &\hspace{15em} \text{(using } l_{c2} \|\cdot\|_2 \leq \|\cdot\|_c) \\
 &\quad + \frac{4\alpha^2 (t-sK)}{N} \sum_{t'=sK}^{t-1} \sum_{i=1}^N \mathbb{E} \left[\frac{2u_{c2}^2}{l_{c2}^2} \left\{ A_1^2 \|\boldsymbol{\theta}_{t'}^i - \boldsymbol{\theta}_{t'}\|_c^2 + A_2^2 \|\boldsymbol{\theta}_{t'}\|_c^2 \right\} + \|\boldsymbol{\theta}_{t'}^i - \boldsymbol{\theta}_{t'}\|_c^2 \right] \quad \text{(Assumption 6.3)} \\
 &\stackrel{(a)}{\leq} \frac{2\alpha^2 u_{c2}^2}{l_{c2}^2} (t-sK) B^2 + \frac{4\alpha^2 u_{c2}^2}{l_{c2} N} \frac{2m_2 B}{l_{c2}} \sum_{i=1}^N \sum_{t'=sK}^{t-1} \sum_{\substack{t''=sK \\ t' < t''}}^{t-1} \mathbb{E} [\rho^{t''-t'}] \\
 &\quad + 4\alpha^2 (t-sK) \sum_{t'=sK}^{t-1} \left[\left(1 + \frac{2A_1^2 u_{c2}^2}{l_{c2}^2} \right) \Omega_{t'} + \frac{2A_2^2 u_{c2}^2}{l_{c2}^2} \|\boldsymbol{\theta}_{t'}\|_c^2 \right] \\
 &\stackrel{(b)}{\leq} \frac{2\alpha^2 u_{c2}^2}{l_{c2}^2} (t-sK) \left[B^2 + \frac{4m_2 B \rho}{1-\rho} \right] + 4\alpha^2 (t-sK) \sum_{t'=sK}^{t-1} \left[\left(1 + \frac{2A_1^2 u_{c2}^2}{l_{c2}^2} \right) \Omega_{t'} + \frac{2A_2^2 u_{c2}^2}{l_{c2}^2} \|\boldsymbol{\theta}_{t'}\|_c^2 \right], \quad (111)
 \end{aligned}$$

where (a) follows since

$$\begin{aligned}
 &\mathbb{E} [\|\mathbf{b}^i(\mathbf{y}_{t'}) - \mathbf{b}(\mathbf{y}_{t'})\|_2 \|\mathbb{E}_{t'} [\mathbf{b}^i(\mathbf{y}_{t''}) - \mathbf{b}(\mathbf{y}_{t''})]\|_c] \\
 &\leq \mathbb{E} \left[\|\mathbf{b}^i(\mathbf{y}_{t'})\|_2 \cdot \left\{ \|\mathbb{E}_{t'} [\mathbf{b}^i(\mathbf{y}_{t''})]\|_c + \left\| \mathbb{E}_{t'} \left[\frac{1}{N} \sum_{j=1}^N \mathbf{b}(\mathbf{y}_{t''}^j) \right] \right\|_c \right\} \right] \quad (\text{Var}(X) \leq E[X^2], \text{ and Triangle inequality}) \\
 &\leq \mathbb{E} \left[\frac{1}{l_{c2}} \|\mathbf{b}^i(\mathbf{y}_{t'})\|_c \cdot \left\{ \|\mathbb{E}_{t'} [\mathbf{b}^i(\mathbf{y}_{t''})]\|_c + \frac{1}{N} \sum_{j=1}^N \left\| \mathbb{E}_{t'} [\mathbf{b}(\mathbf{y}_{t''}^j)] \right\|_c \right\} \right] \quad \text{(Jensen's inequality)} \\
 &\leq \frac{B}{l_{c2}} \cdot 2m_2 \mathbb{E} [\rho^{t''-t'}]. \quad \text{(Assumption 6.1, 6.3)}
 \end{aligned}$$

Finally, the inequality in (b) follows since

$$\begin{aligned} \sum_{t'=sK}^{t-1} \sum_{\substack{t''=sK \\ t' < t''}}^{t-1} \mathbb{E} \left[\rho^{t''-t'} \right] &= \sum_{t'=sK}^{t-1} \sum_{t''=t'+1}^{t-1} \rho^{t''-t'} = \sum_{t'=sK}^{t-1} \frac{\rho - \rho^{t-t'}}{1 - \rho} \\ &= \frac{\rho(t - sK)}{1 - \rho} - \frac{\rho(1 - \rho^{t-sK})}{1 - \rho} \\ &\leq \frac{\rho(t - sK)}{1 - \rho}. \end{aligned}$$

For simplicity, we define $\tilde{A}_1 = \frac{2A_1^2 u_{c2}^2}{l_{c2}^2}$, $\tilde{A}_2 = \frac{2A_2^2 u_{c2}^2}{l_{c2}^2}$, $\tilde{B} = \frac{u_{c2}}{l_{c2}} B$. Hence, (111) simplifies to

$$\mathbb{E}\Omega_t \leq 2\alpha^2(t - sK) \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) + 2 \sum_{t'=sK}^{t-1} \left((1 + \tilde{A}_1)\mathbb{E}\Omega_{t'} + \tilde{A}_2\mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right) \right]. \quad (112)$$

Recursively applying (112), going back 2 steps, we see

$$\begin{aligned} \mathbb{E}\Omega_t &\leq 2\alpha^2\tilde{B}^2(t - sK) \left[1 + \frac{4m_2\rho}{B(1-\rho)} \right] + 4\alpha^2(t - sK)\tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 + 4\alpha^2(t - sK)(1 + \tilde{A}_1) \sum_{t'=sK}^{t-2} \mathbb{E}\Omega_{t'} \\ &\quad + 4\alpha^2(t - sK)(1 + \tilde{A}_1) 2\alpha^2(t - sK - 1) \underbrace{\left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) + 2 \sum_{t'=sK}^{t-2} \left((1 + \tilde{A}_1)\mathbb{E}\Omega_{t'} + \tilde{A}_2\mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right) \right]}_{\geq \Omega_{t-1}} \\ &= 2\alpha^2\tilde{B}^2(t - sK) \left[1 + \frac{4m_2\rho}{B(1-\rho)} \right] \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - 1 - sK) \right] \\ &\quad + 4\alpha^2(t - sK)\tilde{A}_2 \left[\sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 + 4\alpha^2(1 + \tilde{A}_1)(t - 1 - sK) \sum_{t'=sK}^{t-2} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right] \\ &\quad + 4\alpha^2(t - sK)(1 + \tilde{A}_1) \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - 1 - sK) \right] \sum_{t'=sK}^{t-2} \mathbb{E}\Omega_{t'} \\ &\leq 4\alpha^2(t - sK) \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - 1 - sK) \right] \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) + \tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right. \\ &\quad \left. + (1 + \tilde{A}_1) \sum_{t'=sK}^{t-2} \mathbb{E}\Omega_{t'} \right]. \quad (113) \end{aligned}$$

To derive the bound for going back, in general, j steps (such that $t - j \geq sK$), we use an induction argument. Suppose for going back $k (< j)$ steps, the bound is

$$\begin{aligned} \mathbb{E}\Omega_t &\leq 4\alpha^2(t - sK)\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)} \right) \prod_{\ell=1}^{k-1} \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - \ell - sK) \right] \\ &\quad + 4\alpha^2(t - sK)\tilde{A}_2 \prod_{\ell=1}^{k-1} \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - \ell - sK) \right] \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \\ &\quad + 4\alpha^2(t - sK)(1 + \tilde{A}_1) \prod_{\ell=1}^{k-1} \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - \ell - sK) \right] \sum_{t'=sK}^{t-k} \mathbb{E}\Omega_{t'}. \quad (114) \end{aligned}$$

We derive the bound for $k + 1$ steps. For this, we further bound the last term in (114).

$$4\alpha^2(t - sK)(1 + \tilde{A}_1) \prod_{\ell=1}^{k-1} \left[1 + 4\alpha^2(1 + \tilde{A}_1)(t - \ell - sK) \right] \left[\sum_{t'=sK}^{t-k-1} \mathbb{E}\Omega_{t'} + \mathbb{E}\Omega_{t-k} \right]$$

$$\begin{aligned}
 &\leq 4\alpha^2(t-sK)(1+\tilde{A}_1) \prod_{\ell=1}^{k-1} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] \sum_{t'=sK}^{t-k-1} \mathbb{E}\Omega_{t'} \\
 &\quad + 4\alpha^2(t-sK)(1+\tilde{A}_1) \prod_{\ell=1}^{k-1} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] \\
 &\quad \times 4\alpha^2(t-k-sK) \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) + \sum_{t'=sK}^{t-k-1} \left((1+\tilde{A}_1)\mathbb{E}\Omega_{t'} + \tilde{A}_2\mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right) \right] \quad (\text{using (112)}) \\
 &\leq 4\alpha^2(t-sK)(1+\tilde{A}_1) \prod_{\ell=1}^{k-1} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] \sum_{t'=sK}^{t-k-1} \mathbb{E}\Omega_{t'} \\
 &\quad + 4\alpha^2(t-sK) \prod_{\ell=1}^{k-1} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] \\
 &\quad \times 4\alpha^2(1+\tilde{A}_1)(t-k-sK) \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) + \sum_{t'=sK}^{t-k-1} \left((1+\tilde{A}_1)\mathbb{E}\Omega_{t'} + \tilde{A}_2\mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 \right) \right]. \quad (115)
 \end{aligned}$$

Substituting (115) into (114), we see that the induction hypothesis in (114) holds. We can go back as far as the last instant of synchronization, $j \leq t-sK$. For $j = t-sK$, we get

$$\begin{aligned}
 \mathbb{E}\Omega_t &\leq 4\alpha^2(t-sK) \prod_{\ell=1}^{t-1-sK} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] \\
 &\quad \times \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) + \tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 + (1+\tilde{A})\mathbb{E}\Omega_{sK} \right]. \quad (116)
 \end{aligned}$$

Next, using $1+x \leq e^x$ for $x \geq 0$, we get

$$\begin{aligned}
 \prod_{\ell=1}^{t-1-sK} \left[1+4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right] &\leq \exp\left(\sum_{\ell=1}^{t-1-sK} 4\alpha^2(1+\tilde{A}_1)(t-\ell-sK)\right) \\
 &\leq \exp\left(2\alpha^2(1+\tilde{A}_1)(t-sK)^2\right) \\
 &\leq \frac{5}{4}, \quad (117)
 \end{aligned}$$

if α is small enough such that $2\alpha^2(1+\tilde{A}_1)(K-1)^2 \leq \ln \frac{5}{4}$, which holds true by the assumption on the step size. Using (117) in (116), we get

$$\begin{aligned}
 \mathbb{E}\Omega_t &\leq 5\alpha^2(t-sK) \left[\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) + \tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2 + (1+\tilde{A})\underbrace{\mathbb{E}\Omega_{sK}}_{=0} \right] \\
 &= 5\alpha^2(t-sK)\tilde{B}^2 \left(1 + \frac{4m_2\rho}{B(1-\rho)}\right) + 5\alpha^2(t-sK)\tilde{A}_2 \sum_{t'=sK}^{t-1} \mathbb{E}\|\boldsymbol{\theta}_{t'}\|_c^2, \quad (118)
 \end{aligned}$$

which concludes the proof. \square

Proof of Lemma B.9. We have

$$\begin{aligned}
 \mathbb{E}_{t-r}[\|\mathbf{b}(\mathbf{Y}_t)\|_c] &\leq u_{cD}\mathbb{E}_{t-r}[\|\mathbf{b}(\mathbf{Y}_t)\|_D] \\
 &= u_{cD}\mathbb{E}_{t-r} \left[\sqrt{\mathbf{b}(\mathbf{Y}_t)^\top D \mathbf{b}(\mathbf{Y}_t)} \right] \\
 &\leq u_{cD} \sqrt{\mathbb{E}_{t-r}[\mathbf{b}(\mathbf{Y}_t)^\top D \mathbf{b}(\mathbf{Y}_t)]} \quad (\text{concavity of square root})
 \end{aligned}$$

$$\begin{aligned}
 &= u_{cD} \sqrt{\mathbb{E}_{t-r} \left[\left(\frac{1}{N} \sum_{i=1}^N \mathbf{b}^i(\mathbf{y}_t^i) \right)^\top D \left(\frac{1}{N} \sum_{i=1}^N \mathbf{b}^i(\mathbf{y}_t^i) \right) \right]} \\
 &= u_{cD} \sqrt{\mathbb{E}_{t-r} \left[\frac{1}{N^2} \sum_{i=1}^N \mathbf{b}^i(\mathbf{y}_t^i)^\top D \mathbf{b}^i(\mathbf{y}_t^i) + \frac{2}{N^2} \sum_{i<j} \mathbf{b}^i(\mathbf{y}_t^i)^\top D \mathbf{b}^j(\mathbf{y}_t^j) \right]} \\
 &= u_{cD} \sqrt{\frac{1}{N^2} \sum_{i=1}^N \mathbb{E}_{t-r} \left[\|\mathbf{b}^i(\mathbf{y}_t^i)\|_D^2 \right] + \frac{2}{N^2} \sum_{i<j} \mathbb{E}_{t-r} \left[\mathbf{b}^i(\mathbf{y}_t^i) \right]^\top D \mathbb{E}_{t-r} \left[\mathbf{b}^j(\mathbf{y}_t^j) \right]} \quad (\text{Assumption 6.4}) \\
 &\leq u_{cD} \left[\underbrace{\sqrt{\frac{1}{N^2} \sum_{i=1}^N \mathbb{E}_{t-r} \left[\|\mathbf{b}^i(\mathbf{y}_t^i)\|_D^2 \right]}}_{T_1} + \underbrace{\sqrt{\frac{2}{N^2} \sum_{i<j} \mathbb{E}_{t-r} \left[\mathbf{b}^i(\mathbf{y}_t^i) \right]^\top D \mathbb{E}_{t-r} \left[\mathbf{b}^j(\mathbf{y}_t^j) \right]}}_{T_2} \right]. \quad (119)
 \end{aligned}$$

For the term T_1 we have

$$\begin{aligned}
 T_1 &\leq \frac{1}{N} \sqrt{\sum_{i=1}^N \mathbb{E}_{t-r} \left[l_{cD}^{-2} \|\mathbf{b}^i(\mathbf{y}_t^i)\|_c^2 \right]} \leq \frac{1}{N l_{cD}} \sqrt{\sum_{i=1}^N \mathbb{E}_{t-r} [B^2]} \quad (\text{Assumption 6.3}) \\
 &= \frac{B}{l_{cD}} \frac{1}{\sqrt{N}}. \quad (120)
 \end{aligned}$$

For the term T_2 we have

$$\begin{aligned}
 T_2 &\leq \frac{2}{N} \sqrt{\sum_{i<j} \|\mathbb{E}_{t-r} \left[\mathbf{b}^i(\mathbf{y}_t^i) \right]\|_D \cdot \|\mathbb{E}_{t-r} \left[\mathbf{b}^j(\mathbf{y}_t^j) \right]\|_D} \quad (\text{Cauchy-Schwarz}) \\
 &\leq \frac{2}{N} \sqrt{\sum_{i<j} l_{cD}^{-2} \|\mathbb{E}_{t-r} \left[\mathbf{b}^i(\mathbf{y}_t^i) \right]\|_c \cdot \|\mathbb{E}_{t-r} \left[\mathbf{b}^j(\mathbf{y}_t^j) \right]\|_c} \\
 &\leq \frac{2}{N} \sqrt{\sum_{i<j} l_{cD}^{-2} m_2 \rho^r \cdot m_2 \rho^r} \quad (\text{Assumption 6.1}) \\
 &\leq \frac{2m_2 \rho^r}{l_{cD}}. \quad (121)
 \end{aligned}$$

Substituting (120) and (121) in (119), we get the result in (56).

The proof of (57) follows analogously. \square

Proof of Lemma B.10. By definition, we have

$$\begin{aligned}
 \|\mathbf{G}(\Theta_t, \mathbf{Y}_t) - \mathbf{G}(\theta_t, \mathbf{Y}_t)\|_c^2 &= \left\| \frac{1}{N} \sum_{i=1}^N (\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\theta_t, \mathbf{y}_t^i)) \right\|_c^2 \\
 &\leq \left(\frac{1}{N} \sum_{i=1}^N \|\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i) - \mathbf{G}^i(\theta_t, \mathbf{y}_t^i)\|_c \right)^2 \quad (\text{convexity of norm}) \\
 &\leq \left(\frac{1}{N} \sum_{i=1}^N A_1 \|\theta_t^i - \theta_t\|_c \right)^2 \quad (\text{Assumption 6.3}) \\
 &= (A_1 \Delta_t)^2 \\
 &\leq \frac{1}{N} \sum_{i=1}^N A_1^2 \|\theta_t^i - \theta_t\|_c^2 \quad (\text{convexity of square})
 \end{aligned}$$

$$= A_1^2 \Omega_t. \quad (\text{By definition of } \Omega_t)$$

Furthermore, by the convexity of $(\cdot)^2$, we have

$$\Delta_t^2 = \left(\frac{1}{N} \sum_{i=1}^N \Delta_t^i \right)^2 \leq \frac{1}{N} \sum_{i=1}^N (\Delta_t^i)^2 = \Omega_t.$$

□

Proof of Lemma B.11. Since $M_f^{\psi,g}(\cdot)$ is convex, and there exists a norm, $\|\cdot\|_m$, such that $M_f^{\psi,g}(x) = \frac{1}{2} \|x\|_m^2$ (see Proposition B.1), using the chain rule of subdifferential calculus,

$$\nabla M_f^{\psi,g}(x) = \|x\|_m u_x,$$

where $u_x \in \partial \|x\|_m$ is a subgradient of $\|x\|_m$ at x . Hence,

$$\left\| \nabla M_f^{\psi,g}(x) \right\|_m^* = \|x\|_m \|u_x\|_m^*,$$

where $\|\cdot\|_m^*$ is the dual norm of $\|\cdot\|_m$. Since $\|\cdot\|_m$ is convex and, as a function of x , is 1-Lipschitz w.r.t. $\|\cdot\|_m$, we have $\|u_x\|_m^* \leq 1$ (see Lemma 2.6 in (Shalev-Shwartz et al., 2012)).

Further, by convexity of $\|\cdot\|_m$ norm, $\|0\|_m \geq \|x\|_m + \langle u_x, -x \rangle$. Therefore,

$$\langle \nabla M_f^{\psi,g}(x), x \rangle = \|x\|_m \langle u_x, x \rangle \geq \|x\|_m^2 = 2M_f^{\psi,g}(x).$$

□

C. Federated TD-learning

C.1. On-policy Function Approximation

Proposition C.1. *On-policy TD-learning with linear function approximation Algorithm 1 satisfies the following:*

1. $\theta_t^i = \mathbf{v}_t^i - \mathbf{v}^\pi$
2. $S_t = (S_t^1, \dots, S_t^N)$ and $A_t = (A_t^1, \dots, A_t^N)$
3. $\mathbf{y}_t^i = (S_t^i, A_t^i, \dots, S_{t+n-1}^i, A_{t+n-1}^i, S_{t+n}^i)$ and $\mathbf{Y}_t = (S_t, A_t, \dots, S_{t+n-1}, A_{t+n-1}, S_{t+n})$
4. μ^π : Stationary distribution of the policy π .

Furthermore, choose some arbitrary positive constant $\beta > 0$. The corresponding $\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i)$ and $\mathbf{b}^i(\mathbf{y}_t^i)$ in Algorithm 1 for On-policy TD-learning with linear function approximation is as follows

1. $\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i) = \theta_t^i + \frac{1}{\beta} \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\gamma \phi(S_{l+1}^i)^\top \theta_t^i - \phi(S_l^i)^\top \theta_t^i)$
2. $\mathbf{b}^i(\mathbf{y}_t^i) = \frac{1}{\beta} \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathcal{R}(S_l^i, A_l^i) + \gamma \phi(S_{l+1}^i)^\top \mathbf{v}^\pi - \phi(S_l^i)^\top \mathbf{v}^\pi)$

where \mathbf{v}^π solves the projected bellman equation $\Phi \mathbf{v}^\pi = \Pi_\pi((\mathcal{T}^\pi)^n \Phi \mathbf{v}^\pi)$. Furthermore, the corresponding step size α in Algorithm 4 is $\alpha \times \beta$.

Lemma C.1. *Consider the federated on-policy TD-learning Algorithm 1 as a special case of FedSAM Algorithm 4 (see Proposition C.1). Suppose the trajectory $\{S_t^i\}_{t=0,1,\dots}$ converges geometrically fast to its stationary distribution as follows $d_{TV}(P(S_t^i = \cdot | S_0^i) || \mu^i(\cdot)) \leq \bar{m} \bar{\rho}^t$ for all $i = 1, 2, \dots, N$. The corresponding $\bar{\mathbf{G}}^i(\theta)$ in Assumption 6.1 for the federated TD-learning is as follows*

$$\bar{\mathbf{G}}^i(\boldsymbol{\theta}) = \boldsymbol{\theta} + \frac{1}{\beta} \Phi^\top \boldsymbol{\mu}^\pi (\gamma^n (P^\pi)^n \Phi \boldsymbol{\theta} - \Phi \boldsymbol{\theta}), \quad (122)$$

where $\beta > 0$ is an arbitrary constant introduced in Proposition C.1. In addition, (6) holds. Furthermore, for $t \geq n + 1$, we have $m = \max\{\frac{2A_2\bar{m}}{\bar{\rho}^n}, 2B\bar{m}\}$, where A_2 and B are specified in Lemma C.3 and $\bar{\rho} = \rho$.

Lemma C.2. Consider the federated on-policy TD-learning 1 as a special case of FedSAM (as specified in Proposition C.1). Consider the $|\mathcal{S}| \times |\mathcal{S}|$ matrix $\mathbf{U} = \Phi^\top \boldsymbol{\mu}^\pi (\gamma^n (P^\pi)^n - \mathbf{I}) \Phi$ with eigenvalues $\{\lambda_1, \dots, \lambda_{|\mathcal{S}|}\}$. Define $\lambda_{\max} = \max_i |\lambda_i|$ and $\delta = -\max_i \Re[\lambda_i] > 0$, where $\Re[\cdot]$ evaluates the real part. By choosing β large enough in the linear function (122), there exist a weighted 2-norm $\|\boldsymbol{\theta}\|_\Lambda = \sqrt{\boldsymbol{\theta}^\top \Lambda \boldsymbol{\theta}}$, such that $\bar{\mathbf{G}}^i(\boldsymbol{\theta})$ is a contraction with respect to this norm, i.e., $\|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_1) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_2)\|_\Lambda \leq \gamma_c \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\Lambda$ for $\gamma_c = 1 - \frac{\delta^2}{8\lambda_{\max}^2}$.

Lemma C.3. Consider the federated on-policy TD-learning Algorithm 1 as a special case of FedSAM (as specified in Proposition C.1). There exist some constants A_1 , A_2 , and B such that the properties of Assumption 6.3 are satisfied.

Lemma C.4. Consider the federated on-policy TD-learning Algorithm 1 as a special case of FedSAM (as specified in Proposition C.1). Assumption 6.4 holds for this algorithm.

C.1.1. PROOFS

Proof of Proposition C.1. Items 1-4 are by definition. Subtracting \mathbf{v}^π from both sides of the update of the TD-learning, we have

$$\begin{aligned} \underbrace{\mathbf{v}_{t+1}^i - \mathbf{v}^\pi}_{\boldsymbol{\theta}_{t+1}^i} &= \underbrace{\mathbf{v}_t^i - \mathbf{v}^\pi}_{\boldsymbol{\theta}_t^i} + \alpha \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathcal{R}(S_l^i, A_l^i) + \gamma \phi(S_{l+1}^i)^\top \mathbf{v}_t^i - \phi(S_l^i)^\top \mathbf{v}_t^i) \\ &= \boldsymbol{\theta}_t^i + \alpha \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\underbrace{\mathcal{R}(S_l^i, A_l^i) + \gamma \phi(S_{l+1}^i)^\top (\mathbf{v}_t^i - \mathbf{v}^\pi + \mathbf{v}^\pi)}_{\boldsymbol{\theta}_t^i} - \phi(S_l^i)^\top (\mathbf{v}_t^i - \mathbf{v}^\pi + \mathbf{v}^\pi) \right) \\ &= \boldsymbol{\theta}_t^i + \alpha \beta \underbrace{\left(\boldsymbol{\theta}_t^i + \frac{1}{\beta} \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\gamma \phi(S_{l+1}^i)^\top \boldsymbol{\theta}_t^i - \phi(S_l^i)^\top \boldsymbol{\theta}_t^i) \right)}_{\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)} \\ &\quad - \underbrace{\boldsymbol{\theta}_t^i + \frac{1}{\beta} \phi(S_t^i) \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathcal{R}(S_l^i, A_l^i) + \gamma \phi(S_{l+1}^i)^\top \mathbf{v}^\pi - \phi(S_l^i)^\top \mathbf{v}^\pi)}_{\mathbf{b}^i(\mathbf{y}_t^i)}. \end{aligned}$$

which proves items 1 and 2. Furthermore, for the synchronization part of TD-learning, we have

$$\begin{aligned} \mathbf{v}_t^i &\leftarrow \frac{1}{N} \sum_{j=1}^N \mathbf{v}_t^j \\ \implies \underbrace{\mathbf{v}_t^i - \mathbf{v}^\pi}_{\boldsymbol{\theta}_t^i} &\leftarrow \frac{1}{N} \sum_{j=1}^N \underbrace{(\mathbf{v}_t^j - \mathbf{v}^\pi)}_{\boldsymbol{\theta}_t^j}, \end{aligned}$$

which is equivalent to the synchronization step in FedSAM Algorithm 4. Notice that here we used the fact that all agents have the same fixed point \mathbf{v}^π .

□

Proof of Lemma C.1. It is easy to observe that

$$\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i) = \boldsymbol{\theta} + \frac{1}{\beta} \phi(S_t^i) (\gamma^n \phi(S_{t+n}^i)^\top \boldsymbol{\theta} - \phi(S_t^i)^\top \boldsymbol{\theta}).$$

Taking expectation with respect to the stationary distribution, we have

$$\begin{aligned}
 \bar{\mathbf{G}}^i(\boldsymbol{\theta}) &= \mathbb{E}_{S_t^i \sim \mu^\pi} \left[\boldsymbol{\theta} + \frac{1}{\beta} \phi(S_t^i) (\gamma^n \phi(S_{t+n}^i)^\top \boldsymbol{\theta} - \phi(S_t^i)^\top \boldsymbol{\theta}) \right] \\
 &= \mathbb{E}_{S_t^i \sim \mu^\pi} \left[\mathbb{E} \left[\boldsymbol{\theta} + \frac{1}{\beta} \phi(S_t^i) (\gamma^n \phi(S_{t+n}^i)^\top \boldsymbol{\theta} - \phi(S_t^i)^\top \boldsymbol{\theta}) \mid S_t^i \right] \right] && \text{(tower property of expectation)} \\
 &= \mathbb{E}_{S_t^i \sim \mu^\pi} \left[\boldsymbol{\theta} + \frac{1}{\beta} \phi(S_t^i) (\gamma^n \mathbb{E}[(\Phi \boldsymbol{\theta})(S_{t+n}^i) \mid S_t^i] - (\Phi \boldsymbol{\theta})(S_t^i)) \right] \\
 &= \mathbb{E}_{S_t^i \sim \mu^\pi} \left[\boldsymbol{\theta} + \frac{1}{\beta} \phi(S_t^i) (\gamma^n ((P^\pi)^n \Phi \boldsymbol{\theta})(S_t^i) - (\Phi \boldsymbol{\theta})(S_t^i)) \right] \\
 &= \boldsymbol{\theta} + \frac{1}{\beta} \Phi^\top \boldsymbol{\mu}^\pi (\gamma^n (P^\pi)^n \Phi \boldsymbol{\theta} - \Phi \boldsymbol{\theta}).
 \end{aligned}$$

where P^π is the transition probability matrix corresponding to the policy π , and $\boldsymbol{\mu}^\pi$ is a diagonal matrix with diagonal entries corresponding to elements of μ^π .

As explained in (Tsiriklis & Van Roy, 1997), the projection operator Π_π is a linear operator and can be written as $\Pi_\pi = \Phi(\Phi^\top \boldsymbol{\mu}^\pi \Phi)^{-1} \Phi^\top \boldsymbol{\mu}^\pi$, where $\boldsymbol{\mu}^\pi$ is a diagonal matrix with diagonal entries corresponding to the stationary distribution of the policy π . Hence, the fixed point equation is as follows $\Phi \mathbf{v}^\pi = \Phi(\Phi^\top \boldsymbol{\mu}^\pi \Phi)^{-1} \Phi^\top \boldsymbol{\mu}^\pi ((\mathcal{T}^\pi)^n \Phi \mathbf{v}^\pi)$. Since Φ is a full column matrix, we can eliminate it from both sides of the equality, and further multiply both sides with $\Phi^\top \boldsymbol{\mu}^\pi \Phi$. We have $\Phi^\top \boldsymbol{\mu}^\pi \Phi \mathbf{v}^\pi = \Phi^\top \boldsymbol{\mu}^\pi ((\mathcal{T}^\pi)^n \Phi \mathbf{v}^\pi)$, and hence $\Phi^\top \boldsymbol{\mu}^\pi ((\mathcal{T}^\pi)^n \Phi \mathbf{v}^\pi - \Phi \mathbf{v}^\pi) = \mathbf{0}$, which is equivalent to $\mathbb{E}_{S_0^i \sim \mu^\pi} [\phi^\top(S_0^i) \sum_{l=0}^{n-1} (\mathcal{R}(S_l^i, A_l^i) + \gamma (\Phi \mathbf{v}^\pi)(S_{l+1}^i) - (\Phi \mathbf{v}^\pi)(S_l^i))] = \mathbf{0}$, which means

$$\mathbb{E}_{\mathbf{y}^i \sim \mu^\pi} \mathbf{b}^i(\mathbf{y}^i) = \mathbf{0}, \quad (123)$$

and proves (6).

Moreover, we have

$$\begin{aligned}
 \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}) - \mathbb{E}[\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)]\|_c &= \left\| \mathbb{E}_{\mathbf{y}_t^i \sim \mu^\pi} [\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)] - \mathbb{E}[\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)] \right\|_c \\
 &= \left\| \sum_{\mathbf{y}_t^i} (\mu^\pi(\mathbf{y}_t^i) - P(\mathbf{y}_t^i = \mathbf{y}_t^i \mid \mathbf{y}_0^i)) \mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i) \right\|_c \\
 &\leq \sum_{\mathbf{y}_t^i} |\mu^\pi(\mathbf{y}_t^i) - P(\mathbf{y}_t^i = \mathbf{y}_t^i \mid \mathbf{y}_0^i)| \cdot \|\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)\|_c && (\|ax\|_c = |a|\|x\|_c) \\
 &\leq \sum_{\mathbf{y}_t^i} |\mu^\pi(\mathbf{y}_t^i) - P(\mathbf{y}_t^i = \mathbf{y}_t^i \mid \mathbf{y}_0^i)| \cdot A_2 \|\boldsymbol{\theta}\|_c. && \text{(Assumption 6.3)}
 \end{aligned}$$

For brevity, we denote $P(S_t^i = s_t^i) = P(s_t^i)$. We have

$$\begin{aligned}
 &\sum_{\mathbf{y}_t^i} |\mu^\pi(\mathbf{y}_t^i) - P(\mathbf{y}_t^i = \mathbf{y}_t^i \mid \mathbf{y}_0^i)| \\
 &= \sum_{s_t^i, a_t^i, \dots, s_{t+n}^i} |\mu^\pi(s_t^i, a_t^i, \dots, s_{t+n}^i) - P(s_t^i, a_t^i, \dots, s_{t+n}^i \mid \mathbf{y}_0^i)| \\
 &= \sum_{s_t^i, a_t^i, \dots, s_{t+n}^i} \left| \mu^\pi(s_t^i) \pi(a_t^i \mid s_t^i) \mathcal{P}(s_{t+1}^i \mid s_t^i, a_t^i) \dots \mathcal{P}(s_{t+n}^i \mid s_{t+n-1}^i, a_{t+n-1}^i) \right. \\
 &\quad \left. - P(s_t^i \mid S_n^i) \pi(a_t^i \mid s_t^i) \mathcal{P}(s_{t+1}^i \mid s_t^i, a_t^i) \dots \mathcal{P}(s_{t+n}^i \mid s_{t+n-1}^i, a_{t+n-1}^i) \right| && (t \geq n+1) \\
 &= \sum_{s_t^i, a_t^i, \dots, s_{t+n}^i} \left| \mu^\pi(s_t^i) - P(s_t^i \mid S_n^i) \right| \pi(a_t^i \mid s_t^i) \mathcal{P}(s_{t+1}^i \mid s_t^i, a_t^i) \dots \mathcal{P}(s_{t+n}^i \mid s_{t+n-1}^i, a_{t+n-1}^i)
 \end{aligned}$$

$$\begin{aligned}
 &= \sum_{s_t^i} \left| \mu^\pi(s_t^i) - P(s_t^i | S_n^i) \right| \\
 &= 2d_{TV}(P(S_t^i = \cdot | S_n^i) \| \mu^\pi(\cdot)) \\
 &\leq 2\bar{m}\bar{\rho}^{t-n} \\
 &= (2\bar{m}\bar{\rho}^{-n})\bar{\rho}^t.
 \end{aligned}$$

$$\begin{aligned}
 \|\mathbb{E}[\mathbf{b}^i(\mathbf{y}_t^i)]\|_c &= \left\| \mathbb{E}[\mathbf{b}^i(\mathbf{y}_t^i)] - \mathbb{E}_{\mathbf{y}_t^i \sim \mu^\pi}[\mathbf{b}^i(\mathbf{y}_t^i)] \right\|_c && \text{(By (123))} \\
 &= \left\| \sum_{y_t^i} (\mu^\pi(y_t^i) - P(\mathbf{y}_t^i = y_t^i | \mathbf{y}_0^i)) \mathbf{b}^i(\mathbf{y}_t^i) \right\|_c \\
 &\leq \sum_{y_t^i} |\mu^\pi(y_t^i) - P(\mathbf{y}_t^i = y_t^i | \mathbf{y}_0^i)| \|\mathbf{b}^i(\mathbf{y}_t^i)\|_c \\
 &\leq \sum_{y_t^i} |\mu^\pi(y_t^i) - P(\mathbf{y}_t^i = y_t^i | \mathbf{y}_0^i)| B && \text{(Assumption 6.3)} \\
 &\leq 2B\bar{m}\bar{\rho}^t
 \end{aligned}$$

□

Proof of Lemma C.2. Consider the $|\mathcal{S}| \times |\mathcal{S}|$ matrix $\mathbf{U} = \Phi^\top \boldsymbol{\mu}^\pi (\gamma^n (P^\pi)^n - \mathbf{I}) \Phi$ with eigenvalues $\{\lambda_1, \dots, \lambda_{|\mathcal{S}|}\}$. As shown in (Tsitsiklis & Van Roy, 1997), since Φ is a full rank matrix, the real part of λ_i is strictly negative for all $i = 1, \dots, |\mathcal{S}|$. Furthermore, define $\lambda_{\max} = \max_i |\lambda_i|$ and $\delta = -\max_i \Re[\lambda_i] > 0$, where $\Re[\cdot]$ evaluates the real part. Consider the matrix $\mathbf{U}' = \mathbf{I} + \frac{1}{2\lambda_{\max}^2/\delta} \mathbf{U}$. It is easy to show that the eigenvalues of \mathbf{U}' are $\{1 + \frac{\lambda_1}{2\lambda_{\max}^2/\delta}, \dots, 1 + \frac{\lambda_{|\mathcal{S}|}}{2\lambda_{\max}^2/\delta}\}$. For an arbitrary i , the norm of the i 'th eigenvalue satisfies

$$\begin{aligned}
 \left| 1 + \frac{\lambda_i}{2\lambda_{\max}^2/\delta} \right|^2 &= \left(1 + \frac{\Re[\lambda_i]}{2\lambda_{\max}^2/\delta} \right)^2 + \left(\frac{\Im[\lambda_i]}{2\lambda_{\max}^2/\delta} \right)^2 \\
 &\leq \left(1 + \frac{\Re[\lambda_i]}{2\lambda_{\max}^2/\delta} \right)^2 + \left(\frac{\Im[\lambda_i]}{2\lambda_{\max}^2/\delta} \right)^2 && (\Re[\lambda_i] < 0) \\
 &\leq \left(1 + \frac{-\delta}{2\lambda_{\max}^2/\delta} \right)^2 + \left(\frac{\lambda_{\max}}{2\lambda_{\max}^2/\delta} \right)^2 \\
 &= 1 - \frac{\delta^2}{4\lambda_{\max}^2}.
 \end{aligned}$$

Hence, all the eigenvalues of \mathbf{U}' are in the unit circle. By (Bertsekas et al., 1995), Page 46 footnote, we can find a weighted 2-norm as $\|\boldsymbol{\theta}\|_\Lambda = \sqrt{\boldsymbol{\theta}^\top \Lambda \boldsymbol{\theta}}$ such that \mathbf{U}' is contraction with respect to this norm with some contraction factor γ_c . In particular, there exist a choice of Λ such that we have $\gamma_c = 1 - \frac{\delta^2}{8\lambda_{\max}^2}$.

□

Proof of Lemma C.3. The existence of A_1 and A_2 immediately follows after observing that $\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)$ is a linear function of $\boldsymbol{\theta}_t^i$. Furthermore, the result on B follows due to \mathbf{v}^π being bounded as shown in (Chen et al., 2021b). □

Proof of Lemma C.4. For the sake of brevity, we write $S_t^i = s_t^i$ simply as s_t^i , and similarly for other random variables. We have

$$\mathbb{E}_{t-r}[f(\mathbf{y}_t^i) \times g(\mathbf{y}_t^j)]$$

$$\begin{aligned}
 &= \sum_{\substack{s_t^i, a_t^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i, s_t^j, a_t^j, \dots, s_{t+n-1}^j, a_{t+n-1}^j, s_{t+n}^j | \mathcal{F}_{t-r}}}} P(s_t^i, a_t^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i, s_t^j, a_t^j, \dots, s_{t+n-1}^j, a_{t+n-1}^j, s_{t+n}^j | \mathcal{F}_{t-r}) f(\mathbf{y}_t^i) \times g(\mathbf{y}_t^j) \\
 &= \sum_{\substack{s_{t-r}^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i, s_{t-r}^j, \dots, s_{t+n-1}^j, a_{t+n-1}^j, s_{t+n}^j | \mathcal{F}_{t-r}}} P(s_{t-r}^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i, s_{t-r}^j, \dots, s_{t+n-1}^j, a_{t+n-1}^j, s_{t+n}^j | \mathcal{F}_{t-r}) f(\mathbf{y}_t^i) \times g(\mathbf{y}_t^j) \\
 &= \sum_{\substack{s_{t-r}^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i | \mathcal{F}_{t-r}}} P(s_{t-r}^i, \dots, s_{t+n-1}^i, a_{t+n-1}^i, s_{t+n}^i | \mathcal{F}_{t-r}) P(s_{t-r}^j, \dots, s_{t+n-1}^j, a_{t+n-1}^j, s_{t+n}^j | \mathcal{F}_{t-r}) f(\mathbf{y}_t^i) g(\mathbf{y}_t^j) \\
 &= \mathbb{E}_{t-r}[f(\mathbf{y}_t^i)] \times \mathbb{E}_{t-r}[g(\mathbf{y}_t^j)].
 \end{aligned}$$

□

Proof of Theorem 4.1. By Proposition C.1 and Lemmas C.1, C.2, C.3, and C.4, it is clear that the federated TD-learning with linear function approximation Algorithm 1 satisfies all the Assumptions 6.1, 6.2, 6.3, and 6.4 on the FedSAM Algorithm 4. Furthermore, by the proof of Theorem B.1, we have $w_t = (1 - \frac{\alpha\varphi_2}{2})^{-t}$, and the constant c_{TDL} in the sampling distribution $q_T^{c_{TDL}}$ in Algorithm 1 is $c_{TDL} = (1 - \frac{\alpha\varphi_2}{2})^{-1}$. Furthermore, by choosing the step size α small enough, we can satisfy the requirements in (21), (23), (32), (35). By choosing K large enough, we can satisfy $K > \tau_\alpha$. Hence, the result of Theorem B.1 holds for this algorithm with some $c_{TDL} > 1$. Also, it is easy to see that $(1 - \frac{\alpha\varphi_2}{2})^{-\tau_\alpha} = \mathcal{O}(1)$, which is a constant that can be absorbed in \mathcal{C}_1^{TDL} . Finally, for the sample complexity result, we simply employ Corollary B.1.1.

Next, we derive the constant c_{TDL} . Since $\|\cdot\|_c = \|\cdot\|_\Lambda$, which is smooth, we choose $g(\cdot) = \frac{1}{2}\|\cdot\|_\Lambda^2$. By taking $\psi = 1$, we have $l_{cs} = u_{cs} = 1$. Therefore, we have $\varphi_1 = 1$, and $\varphi_2 = 1 - \gamma_c$, and $c_{TDL} = \left(1 - \frac{\alpha(1-\gamma_c)}{2}\right)^{-1}$, where γ_c is defined in Lemma C.2.

□

C.2. Off-policy Tabular Setting

In this subsection, we verify that the Off-policy federated TD-learning Algorithm 2 satisfies the properties of the FedSAM Algorithm 4. In the following, V^π is the solution to the Bellman equation (124).

$$V^\pi(s) = \sum_a \pi(a|s) \left[\mathcal{R}(s, a) + \gamma \sum_{s'} \mathcal{P}(s'|s, a) V^\pi(s') \right] \quad (124)$$

Note that V^π is independent of the sampling policy of the agent. Furthermore, we take $\|\cdot\|_c = \|\cdot\|_\infty$.

Proposition C.2. *Off-policy n -step federated TD-learning is equivalent to the FedSAM Algorithm 4 with the following parameters.*

1. $\boldsymbol{\theta}_t^i = \mathbf{V}_t^i - \mathbf{V}^\pi$
2. $S_t = (S_t^1, \dots, S_t^N)$ and $A_t = (A_t^1, \dots, A_t^N)$
3. $\mathbf{y}_t^i = (S_t^i, A_t^i, \dots, S_{t+n-1}^i, A_{t+n-1}^i, S_{t+n}^i)$ and $\mathbf{Y}_t = (S_t, A_t, \dots, S_{t+n-1}, A_{t+n-1}, S_{t+n})$
4. μ^i : Stationary distribution of the sampling policy of the i -th agent.
5. $\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)_s = \boldsymbol{\theta}_t^i(s) + \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}_t^i(S_{l+1}^i) - \boldsymbol{\theta}_t^i(S_l^i)] \right)$
6. $\mathbf{b}^i(\mathbf{y}_t^i)_s = \mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)]$

Lemma C.5. *Consider the federated off-policy TD-learning Algorithm 2 as a special case of FedSAM (as specified in Proposition C.2). Suppose the trajectory $\{S_t^i\}_{t=0,1,\dots}$ converges geometrically fast to its stationary distribution as follows*

$d_{TV}(P(S_t^i = \cdot | S_0^i) | | \mu^i(\cdot)) \leq \bar{m} \bar{\rho}^t$ for all $i = 1, 2, \dots, N$. The corresponding $\bar{\mathbf{G}}^i(\boldsymbol{\theta})$ in Assumption 6.1 for the federated TD-learning is as follows

$$\bar{\mathbf{G}}^i(\boldsymbol{\theta})_s = \boldsymbol{\theta}(s) + [\gamma^n \boldsymbol{\mu}^i(P^\pi)^n \boldsymbol{\theta} - \boldsymbol{\mu}^i \boldsymbol{\theta}](s).$$

Furthermore, for $t \geq n + 1$, we have $m_1 = \frac{2A_2 \bar{m}}{\bar{\rho}^n}$, where A_2 is the constant specified in Assumption 6.3, $m_2 = 0$, and $\bar{\rho} = \rho$.

Lemma C.6. Consider the federated off-policy TD-learning 2 as a special case of FedSAM (as specified in Proposition C.2). The corresponding contraction factor γ_c in Assumption 6.2 for this algorithm is $\gamma_c = 1 - \mu_{\min}(1 - \gamma^{n+1})$, where $\mu_{\min} = \min_{s,i} \mu^i(s)$.

Lemma C.7. Consider the federated off-policy TD-learning 2 as a special case of FedSAM (as specified in Proposition C.2). The constants A_1 , A_2 , and B in Assumption 6.3 can be chosen as follows: $A_1 = A_2 = 1 + (1 + \gamma) \begin{cases} n & \text{if } \gamma \mathfrak{J}_{\max} = 1 \\ \frac{1 - (\gamma \mathfrak{J}_{\max})^n}{1 - \gamma \mathfrak{J}_{\max}} & \text{o.w.} \end{cases}$, and $B = \frac{2\mathfrak{J}_{\max}}{1 - \gamma} \begin{cases} n & \text{if } \gamma \mathfrak{J}_{\max} = 1 \\ \frac{1 - (\gamma \mathfrak{J}_{\max})^n}{1 - \gamma \mathfrak{J}_{\max}} & \text{o.w.} \end{cases}$, where $\mathfrak{J}_{\max} = \max_{s^i, a^i, i} \mathfrak{J}^{(i)}(s^i, a^i)$.

Lemma C.8. Consider the federated off-policy TD-learning 2 as a special case of FedSAM (as specified in Proposition C.2). Assumption 6.4 holds for this algorithm.

C.2.1. PROOFS

Proof of Proposition C.2. Items 1-4 are by definition. Furthermore, by the update of the TD-learning, and subtracting \mathbf{V}^π from both sides, we have

$$\begin{aligned} \underbrace{\mathbf{V}_{t+1}^i - \mathbf{V}^\pi(s)}_{\boldsymbol{\theta}_{t+1}^i(s)} &= \underbrace{\mathbf{V}_t^i - \mathbf{V}^\pi(s)}_{\boldsymbol{\theta}_t^i(s)} \\ &+ \alpha \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left[\prod_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right] (\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}_t^i(S_{l+1}^i) - \mathbf{V}_t^i(S_l^i)) \right) \\ &= \boldsymbol{\theta}_t^i(s) \\ &+ \alpha \mathbb{1}_{\{s=S_t^i\}} \left\{ \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right) \left[\mathcal{R}(S_l^i, A_l^i) + \gamma \left(\underbrace{\mathbf{V}_t^i(S_{l+1}^i) - \mathbf{V}^\pi(S_{l+1}^i)}_{\boldsymbol{\theta}_t^i(S_{l+1}^i)} + \mathbf{V}^\pi(S_{l+1}^i) \right) \right. \right. \\ &\quad \left. \left. - \left(\underbrace{\mathbf{V}_t^i(S_l^i) - \mathbf{V}^\pi(S_l^i)}_{\boldsymbol{\theta}_t^i(S_l^i)} + \mathbf{V}^\pi(S_l^i) \right) \right] \right\} \\ &= \boldsymbol{\theta}_t^i(s) + \alpha \left\{ \underbrace{\boldsymbol{\theta}_t^i(s) + \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}_t^i(S_{l+1}^i) - \boldsymbol{\theta}_t^i(S_l^i)] \right)}_{\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)_s} \right\} \\ &\quad \left. - \underbrace{\mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)]}_{\mathbf{b}^i(\mathbf{y}_t^i)_s} \right\}, \end{aligned}$$

which proves items 5 and 6. Furthermore, for the synchronization part of TD-learning, if $t \bmod K = 0$,

$$\begin{aligned} \mathbf{V}_t^i &\leftarrow \frac{1}{N} \sum_{j=1}^N \mathbf{V}_t^j \\ \implies \underbrace{\mathbf{V}_t^i - \mathbf{V}^\pi}_{\boldsymbol{\theta}_t^i} &\leftarrow \frac{1}{N} \sum_{j=1}^N \underbrace{(\mathbf{V}_t^j - \mathbf{V}^\pi)}_{\boldsymbol{\theta}_t^j}, \end{aligned}$$

which is equivalent to the synchronization step in FedSAM Algorithm 4. Notice that here we used the fact that all agents have the same fixed point \mathbf{V}^π . \square

Proof of Lemma C.5. By taking expectation of $\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)_s$, we have

$$\begin{aligned}\bar{\mathbf{G}}^i(\boldsymbol{\theta})_s &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\boldsymbol{\theta}(s) + \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\prod_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}(S_{l+1}^i) - \boldsymbol{\theta}(S_l^i)] \right) \right] \\ &= \boldsymbol{\theta}(s) + \underbrace{\sum_{l=t}^{t+n-1} \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} \left(\prod_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}(S_{l+1}^i) - \boldsymbol{\theta}(S_l^i)] \right]}_{T_l}.\end{aligned}$$

Denote $\mathbb{E}_k^i[\cdot] = \mathbb{E}[\cdot | \{S_r^i, A_r^i\}_{r \leq k-1}, S_k^i]$. For T_l , we have

$$\begin{aligned}T_l &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{E}_l^i \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} \left(\prod_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}(S_{l+1}^i) - \boldsymbol{\theta}(S_l^i)] \right] \right] \\ &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} \left(\prod_{j=t}^{l-1} \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) \left[\gamma \mathbb{E}_l^i \left[\mathcal{J}^{(i)}(S_l^i, A_l^i) \boldsymbol{\theta}(S_{l+1}^i) \right] - \boldsymbol{\theta}(S_l^i) \right] \right].\end{aligned}$$

Here,

$$\begin{aligned}\mathbb{E}_l^i \left[\mathcal{J}^{(i)}(S_l^i, A_l^i) \boldsymbol{\theta}(S_{l+1}^i) \right] &= \sum_{s,a} P(S_{l+1}^i = s, A_l^i = a | S_l^i) \mathcal{J}^{(i)}(S_l^i, A_l^i = a) \boldsymbol{\theta}(s) \\ &= \sum_{s,a} P(S_{l+1}^i = s, A_l^i = a | S_l^i) \frac{\pi(a | S_l^i)}{\pi^i(a | S_l^i)} \boldsymbol{\theta}(s) \\ &= \sum_{s,a} \pi^i(a | S_l^i) \mathcal{P}(s | S_l^i, a) \frac{\pi(a | S_l^i)}{\pi^i(a | S_l^i)} \boldsymbol{\theta}(s) \\ &= \sum_{s,a} \mathcal{P}(s | S_l^i, a) \pi(a | S_l^i) \boldsymbol{\theta}(s) \equiv [P^\pi \boldsymbol{\theta}](S_l^i),\end{aligned}$$

where $[P^\pi]_{s_0, s_1} = \sum_a \mathcal{P}(s_1 | s_0, a) \pi(a | s_0)$. Hence, we have

$$\begin{aligned}T_l &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} \prod_{j=t}^{l-2} \mathcal{J}^{(i)}(S_j^i, A_j^i) \mathbb{E}_{l-1}^i \left[\mathcal{J}^{(i)}(S_{l-1}^i, A_{l-1}^i) [\gamma (P^\pi \boldsymbol{\theta})(S_l^i) - \boldsymbol{\theta}(S_l^i)] \right] \right] \\ &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} \prod_{j=t}^{l-2} \mathcal{J}^{(i)}(S_j^i, A_j^i) [\gamma ((P^\pi)^2 \boldsymbol{\theta})(S_{l-1}^i) - (P^\pi \boldsymbol{\theta})(S_{l-1}^i)] \right] \\ &= \dots \\ &= \mathbb{E}_{S_t^i \sim \mu^i} \left[\mathbb{1}_{\{s=S_t^i\}} \gamma^{l-t} [\gamma ((P^\pi)^{l-t+1} \boldsymbol{\theta})(S_t^i) - ((P^\pi)^{l-t} \boldsymbol{\theta})(S_t^i)] \right] \\ &= \boldsymbol{\mu}^i(s) \gamma^{l-t} [\gamma ((P^\pi)^{l-t+1} \boldsymbol{\theta})(s) - ((P^\pi)^{l-t} \boldsymbol{\theta})(s)] \\ &= \gamma^{l-t} [\gamma (\boldsymbol{\mu}^i (P^\pi)^{l-t+1} \boldsymbol{\theta})(s) - (\boldsymbol{\mu}^i (P^\pi)^{l-t} \boldsymbol{\theta})(s)] \\ &= [\gamma^{l-t+1} \boldsymbol{\mu}^i (P^\pi)^{l-t+1} \boldsymbol{\theta} - \gamma^{l-t} \boldsymbol{\mu}^i (P^\pi)^{l-t} \boldsymbol{\theta}](s),\end{aligned}$$

where we denote $\boldsymbol{\mu}^i$ as diagonal matrix with diagonal entries corresponding to the stationary distribution μ^i . Hence, in total we have

$$\begin{aligned}\bar{\mathbf{G}}^i(\boldsymbol{\theta})_s &= \boldsymbol{\theta}(s) + \sum_{l=t}^{t+n-1} [\gamma^{l-t+1} \boldsymbol{\mu}^i (P^\pi)^{l-t+1} \boldsymbol{\theta} - \gamma^{l-t} \boldsymbol{\mu}^i (P^\pi)^{l-t} \boldsymbol{\theta}](s) \\ &= \boldsymbol{\theta}(s) + [\gamma^n \boldsymbol{\mu}^i (P^\pi)^n \boldsymbol{\theta} - \boldsymbol{\mu}^i \boldsymbol{\theta}](s).\end{aligned}$$

Furthermore, by the same argument as in the proof of Lemma C.1, we have

$$\|\bar{\mathbf{G}}^i(\boldsymbol{\theta}) - \mathbb{E}[\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)]\|_c \leq \sum_{y_t^i} |\mu^i(y_t^i) - P(\mathbf{y}_t^i = y_t^i | \mathbf{y}_0^i)| \cdot A_2 \|\boldsymbol{\theta}\|_c,$$

and

$$\sum_{y_t^i} |\mu^i(y_t^i) - P(\mathbf{y}_t^i = y_t^i | \mathbf{y}_0^i)| \leq (2\bar{m}\bar{\rho}^{-n})\bar{\rho}^t,$$

which proves that $m_1 = 2\bar{m}A_2\bar{\rho}^{-n}$ constant. In addition, we have

$$\begin{aligned} \|\mathbb{E}[\mathbf{b}^i(\mathbf{y}_t^i)]\|_c &= \left\| \mathbb{E} \left[\mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)] \right] \right\|_c \\ &= \left\| \mathbb{E} \left[\mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \mathbb{E}_l \left[\Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)] \right] \right] \right\|_c \\ &= \left\| \mathbb{E} \left[\mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \Pi_{j=t}^{l-1} \mathcal{J}^{(i)}(S_j^i, A_j^i) \underbrace{\mathbb{E}_l \left[\mathcal{J}^{(i)}(S_l^i, A_l^i) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)] \right]}_T \right] \right\|_c. \end{aligned}$$

For the term T , we have

$$\begin{aligned} T &= \sum_a \pi^i(a|S_l^i) \cdot \frac{\pi(a|S_l^i)}{\pi^i(a|S_l^i)} \left[\mathcal{R}(S_l^i, a) + \gamma \sum_{s'} \mathcal{P}(s'|S_l^i, a) \mathbf{V}^\pi(s') - \mathbf{V}^\pi(S_l^i) \right] \\ &= \sum_a \pi(a|S_l^i) \left[\mathcal{R}(S_l^i, a) + \gamma \sum_{s'} \mathcal{P}(s'|S_l^i, a) \mathbf{V}^\pi(s') - \mathbf{V}^\pi(S_l^i) \right] \\ &= 0, \end{aligned}$$

which shows that $m_2 = 0$. □

Proof of Lemma C.6.

$$\begin{aligned} \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_1) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_2)\|_c &= \|\boldsymbol{\theta}_1 + [\gamma^{n+1} \boldsymbol{\mu}^i(P^\pi)^{n+1} \boldsymbol{\theta}_1 - \boldsymbol{\mu}^i \boldsymbol{\theta}_1] - (\boldsymbol{\theta}_2 + [\gamma^{n+1} \boldsymbol{\mu}^i(P^\pi)^{n+1} \boldsymbol{\theta}_2 - \boldsymbol{\mu}^i \boldsymbol{\theta}_2])\|_\infty \\ &= \|(I - \boldsymbol{\mu}^i(I - \gamma^{n+1}(P^\pi)^{n+1}))(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2)\|_\infty \\ &\leq \|I - \boldsymbol{\mu}^i(I - \gamma^{n+1}(P^\pi)^{n+1})\|_\infty \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty. \quad (\text{definition of matrix norm}) \end{aligned}$$

Since the elements of the matrix $I - \boldsymbol{\mu}^i(I - \gamma^{n+1}(P^\pi)^{n+1})$ is all positive, we have $\|I - \boldsymbol{\mu}^i(I - \gamma^{n+1}(P^\pi)^{n+1})\|_\infty = \|(I - \boldsymbol{\mu}^i(I - \gamma^{n+1}(P^\pi)^{n+1}))\mathbf{1}\|_\infty = \|\mathbf{1} - \boldsymbol{\mu}^i(\mathbf{1} - \gamma^{n+1}\mathbf{1})\|_\infty = 1 - \mu_{\min}^i(1 - \gamma^{n+1}) \leq 1 - \mu_{\min}(1 - \gamma^{n+1})$. □

Proof of Lemma C.7.

$$\begin{aligned} &\|\mathbf{G}^i(\boldsymbol{\theta}_1, \mathbf{y}) - \mathbf{G}^i(\boldsymbol{\theta}_2, \mathbf{y})\|_c \\ &= \max_s \left| \boldsymbol{\theta}_1(s) - \boldsymbol{\theta}_2(s) + \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}_1(S_{l+1}^i) - \boldsymbol{\theta}_1(S_l^i)] \right) \right. \\ &\quad \left. - \mathbb{1}_{\{s=S_t^i\}} \left(\sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) \right) [\gamma \boldsymbol{\theta}_2(S_{l+1}^i) - \boldsymbol{\theta}_2(S_l^i)] \right) \right| \quad (\because c = \infty) \\ &\leq \max_s \left[|\boldsymbol{\theta}_1(s) - \boldsymbol{\theta}_2(s)| + \left| \sum_{l=t}^{t+n-1} \gamma^{l-t} \Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) [\gamma (\boldsymbol{\theta}_1(S_{l+1}^i) - \boldsymbol{\theta}_2(S_{l+1}^i)) - (\boldsymbol{\theta}_1(S_l^i) - \boldsymbol{\theta}_2(S_l^i))] \right| \right] \\ &\quad (\text{triangle inequality}) \\ &\leq \max_s \left[|\boldsymbol{\theta}_1(s) - \boldsymbol{\theta}_2(s)| + \sum_{l=t}^{t+n-1} \gamma^{l-t} \Pi_{j=t}^l \mathcal{J}^{(i)}(S_j^i, A_j^i) [\gamma |\boldsymbol{\theta}_1(S_{l+1}^i) - \boldsymbol{\theta}_2(S_{l+1}^i)| + |\boldsymbol{\theta}_1(S_l^i) - \boldsymbol{\theta}_2(S_l^i)|] \right] \\ &\quad (\text{triangle inequality}) \end{aligned}$$

$$\begin{aligned}
 &\leq \max_s \left[|\boldsymbol{\theta}_1(s) - \boldsymbol{\theta}_2(s)| + \sum_{l=t}^{t+n-1} \gamma^{l-t} \Pi_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) [\gamma \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty] \right] && \text{(definition of } \|\cdot\|_\infty) \\
 &\leq \max_s \left[|\boldsymbol{\theta}_1(s) - \boldsymbol{\theta}_2(s)| + \sum_{l=t}^{t+n-1} \gamma^{l-t} \mathfrak{J}_{\max}^{l-t+1} [\gamma \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty] \right] && \text{(definition of } \mathfrak{J}_{\max}) \\
 &= \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + [\gamma \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty] \mathfrak{J}_{\max} \sum_{l=0}^{n-1} (\gamma \mathfrak{J}_{\max})^l \\
 &= \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + [\gamma \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty] \mathfrak{J}_{\max} \begin{cases} n & \text{if } \gamma \mathfrak{J}_{\max} = 1 \\ \frac{1 - (\gamma \mathfrak{J}_{\max})^n}{1 - \gamma \mathfrak{J}_{\max}} & \text{o.w.} \end{cases}
 \end{aligned}$$

Furthermore, we have

$$\begin{aligned}
 \|\mathbf{b}^i(\mathbf{y}_t^i)\|_c &= \max_{S_t^i, A_t^i, \dots, S_{t+n}^i} \left| \mathbb{1}_{\{s=S_t^i\}} \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\Pi_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right) [\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)] \right| \\
 &\leq \max_{S_t^i, A_t^i, \dots, S_{t+n}^i} \sum_{l=t}^{t+n-1} \gamma^{l-t} \left(\Pi_{j=t}^l \mathfrak{J}^{(i)}(S_j^i, A_j^i) \right) |\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)| && \text{(triangle inequality)} \\
 &\leq \max_{S_t^i, A_t^i, \dots, S_{t+n}^i} \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathfrak{J}_{\max}^{l-t+1}) |\mathcal{R}(S_l^i, A_l^i) + \gamma \mathbf{V}^\pi(S_{l+1}^i) - \mathbf{V}^\pi(S_l^i)| \\
 &\leq \max_{S_t^i, A_t^i, \dots, S_{t+n}^i} \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathfrak{J}_{\max}^{l-t+1}) [|\mathcal{R}(S_l^i, A_l^i)| + \gamma |\mathbf{V}^\pi(S_{l+1}^i)| + |\mathbf{V}^\pi(S_l^i)|] && \text{(triangle inequality)} \\
 &\leq \max_{S_t^i, A_t^i, \dots, S_{t+n}^i} \sum_{l=t}^{t+n-1} \gamma^{l-t} (\mathfrak{J}_{\max}^{l-t+1}) \left[1 + \frac{\gamma}{1-\gamma} + \frac{1}{1-\gamma} \right] \\
 &= \frac{2\mathfrak{J}_{\max}}{1-\gamma} \sum_{l=t}^{t+n-1} (\gamma \mathfrak{J}_{\max})^{l-t} \\
 &= \frac{2\mathfrak{J}_{\max}}{1-\gamma} \begin{cases} n & \text{if } \gamma \mathfrak{J}_{\max} = 1 \\ \frac{1 - (\gamma \mathfrak{J}_{\max})^n}{1 - \gamma \mathfrak{J}_{\max}} & \text{o.w.} \end{cases}
 \end{aligned}$$

□

Proof of Lemma C.8. The proof follows similar to Lemma C.4. □

Proof of Theorem 5.1. By Proposition C.2 and Lemmas C.5, C.6, C.7, and C.8, it is clear that the federated off-policy TD-learning Algorithm 2 satisfies all the Assumptions 6.1, 6.2, 6.3, and 6.4 of the FedSAM Algorithm 4. Furthermore, by the proof of Theorem B.1, we have $w_t = (1 - \frac{\alpha \varphi_2}{2})^{-t}$, and the constant c in the sampling distribution q_t^c in Algorithm 1 is $c = (1 - \frac{\alpha \varphi_2}{2})^{-1}$. In equation (125) we evaluate the exact value of w_t .

Furthermore, by choosing step size α small enough, we can satisfy the requirements in (21), (23), (32), (35). By choosing K large enough, we can satisfy $K > \tau$, and by choosing T large enough we can satisfy $T > K + \tau$. Hence, the result of Theorem B.1 holds for this algorithm.

Next, we derive the constants involved in Theorem B.1 step by step. After deriving the constants $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$, and \mathcal{C}_4 in Theorem B.1, we can directly get the constants \mathcal{C}_i^{TD} for $i = 1, 2, 3, 4$.

In this analysis we only consider the terms involving $|\mathcal{S}|, |\mathcal{A}|, \frac{1}{1-\gamma}, \mathfrak{J}_{\max}$, and μ_{\min} . Since $\|\cdot\|_c = \|\cdot\|_\infty$, we choose $g(\cdot) = \frac{1}{2} \|\cdot\|_p^2$, i.e. the p -norm with $p = 2 \log(|\mathcal{S}|)$. It is known that $g(\cdot)$ is $(p-1)$ smooth with respect to $\|\cdot\|_p$ norm (Beck, 2017), and hence $L = \Theta(\log(|\mathcal{S}|))$. Hence, we have $l_{cs} = |\mathcal{S}|^{-1/p} = \frac{1}{\sqrt{e}} = \Theta(1)$ and $u_{cs} = 1$. Therefore, we have $\varphi_1 = \frac{1+\psi u_{cs}^2}{1+\psi l_{cs}^2} = \frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}} \leq 1 + \psi$. By choosing $\psi = (\frac{1+\gamma_c}{2\gamma_c})^2 - 1 = \frac{1+2\gamma_c-3\gamma_c^2}{4\gamma_c^2} \geq (1-\gamma_c) = \mu_{\min}(1-\gamma^{n+1}) =$

$\Omega(\mu_{\min}(1-\gamma))$, which is $\psi = \mathcal{O}(1)$, we have $\varphi_1 = \frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}} = \sqrt{e} \frac{(\frac{1+\gamma_c}{2\gamma_c})^2}{\sqrt{e} + (\frac{1+\gamma_c}{2\gamma_c})^2 - 1} = \mathcal{O}(1)$, and

$$\begin{aligned} \varphi_2 &= 1 - \gamma_c \sqrt{\frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}}} = 1 - \gamma_c \sqrt{\frac{(\frac{1+\gamma_c}{2\gamma_c})^2}{1 + \frac{(\frac{1+\gamma_c}{2\gamma_c})^2 - 1}{\sqrt{e}}}} = 1 - \frac{0.5(1+\gamma_c)e^{1/4}}{\sqrt{\sqrt{e}-1 + \left(\frac{1+\gamma_c}{2\gamma_c}\right)^2}} = 1 - \frac{0.5e^{1/4}(2 - \mu_{\min}(1-\gamma^{n+1}))}{\sqrt{\sqrt{e}-1 + \left(\frac{2-\mu_{\min}(1-\gamma^{n+1})}{2-2\mu_{\min}(1-\gamma^{n+1})}\right)^2}} \\ &> 1 - \gamma_c \sqrt{1+\psi} = 1 - \gamma_c \frac{1+\gamma_c}{2\gamma_c} = \frac{1-\gamma_c}{2} = 0.5\mu_{\min}(1-\gamma^{n+1}) = \Omega(\mu_{\min}(1-\gamma)) \end{aligned}$$

$$\varphi_3 = \frac{L(1+\psi u_{cs}^2)}{\psi l_{cs}^2} = \mathcal{O}\left(\frac{\log(|\mathcal{S}|)(1+\psi)}{\psi}\right) \leq \mathcal{O}\left(\frac{\log(|\mathcal{S}|)}{1-\gamma_c}\right) = \mathcal{O}\left(\frac{\log(|\mathcal{S}|)}{\mu_{\min}(1-\gamma)}\right).$$

Using φ_2 , we have

$$w_t = \left(1 - \frac{\alpha\varphi_2}{2}\right)^{-t} = \left(1 - \alpha/2 + \frac{0.25\alpha e^{1/4}(2 - \mu_{\min}(1-\gamma^{n+1}))}{\sqrt{\sqrt{e}-1 + \left(\frac{2-\mu_{\min}(1-\gamma^{n+1})}{2-2\mu_{\min}(1-\gamma^{n+1})}\right)^2}}\right)^{-t}. \quad (125)$$

Further, we have

$$\begin{aligned} l_{cm} &= (1 + \psi l_{cs}^2)^{1/2} = \Theta(1) \\ u_{cm} &= (1 + \psi u_{cs}^2)^{1/2} = \Theta(1) \end{aligned}$$

Since TV-divergence is upper bounded with 1, we have $\bar{m} = \mathcal{O}(1)$. By Lemma C.7, we have

$$A_1 = A_2 = 1 + (1+\gamma) \begin{cases} n & \text{if } \gamma\mathfrak{J}_{\max} = 1 \\ \frac{1-(\gamma\mathfrak{J}_{\max})^n}{1-\gamma\mathfrak{J}_{\max}} & \text{o.w.} \end{cases} = \mathcal{O}(\mathfrak{J}_{\max}^{n-1})$$

and $A_1 = A_2 = \Omega(1)$,

$$B = \frac{2\mathfrak{J}_{\max}}{1-\gamma} \begin{cases} n & \text{if } \gamma\mathfrak{J}_{\max} = 1 \\ \frac{1-(\gamma\mathfrak{J}_{\max})^n}{1-\gamma\mathfrak{J}_{\max}} & \text{o.w.} \end{cases} = \mathcal{O}\left(\frac{\mathfrak{J}_{\max}^n}{1-\gamma}\right),$$

and $B = \Omega(1)$. Hence $m_1 = \frac{2A_2\bar{m}}{\rho^n} = \mathcal{O}(\mathfrak{J}_{\max}^{n-1})$. Also, we have $m_2 = 0$.

We choose the D -norm in Lemma B.9 as the 2-norm $\|\cdot\|_2$. Hence, by primary norm equivalence, we have $l_{cD} = \frac{1}{\sqrt{|\mathcal{S}|}}$, and $u_{cD} = 1$, and hence $\frac{u_{cD}}{l_{cD}} = \sqrt{|\mathcal{S}|}$.

We can evaluate the rest of the constants as follows

$$\begin{aligned} \zeta_1 &= \zeta_4 = \zeta_6 = \sqrt{\varphi_2/10} = \Omega(\sqrt{\mu_{\min}(1-\gamma)}) \\ \zeta_2 &= \sqrt{\frac{\varphi_2}{10} \cdot \frac{\psi l_{cs}^2}{2LA_2}} = \Omega\left(\sqrt{\mu_{\min}(1-\gamma) \cdot \frac{\mu_{\min}(1-\gamma)}{\log(|\mathcal{S}|)}}\right) = \Omega\left(\frac{\mu_{\min}(1-\gamma)}{\sqrt{\log(|\mathcal{S}|)}}\right), \end{aligned}$$

and similarly

$$\zeta_3 = \sqrt{\frac{\varphi_2}{10} \cdot \frac{\psi l_{cs}^2}{L(A_1+1)}} = \Omega\left(\frac{\mu_{\min}(1-\gamma)}{\sqrt{\log(|\mathcal{S}|)}}\right).$$

$$C_1 = 3\sqrt{A_2^2 + 1} = \mathcal{O}(\mathfrak{J}_{\max}^{n-1}),$$

and

$$C_1 = \Omega(1),$$

$$C_2 = 3C_1 + 8 = \mathcal{O}(\mathfrak{J}_{\max}^{n-1}),$$

$$C_3 = \frac{m_2 L}{\psi l_{cs}^2} = 0,$$

$$\begin{aligned} C_4 &= \left(\frac{L^2}{2\psi^2 l_{cs}^4} + \frac{2\alpha L A_2}{\psi l_{cs}^2} \left(\frac{1}{2} + \frac{u_{cm}^2}{2\zeta_3^2} \right) + \frac{\alpha L (A_1 + 1)}{\psi l_{cs}^2} \left(\frac{3u_{cm}^2}{2\zeta_3^2} + 2 \right) + \frac{3m_1 L \alpha^2}{l_{cs}^2 \psi} \right) \\ &= \mathcal{O} \left(\frac{\log^2(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^2} + \frac{\log(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min} (1-\gamma)} \cdot \frac{\log(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^2} + \frac{\log(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min} (1-\gamma)} \cdot \frac{\log(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^2} + \frac{\mathfrak{J}_{\max}^{n-1} \log(|\mathcal{S}|)}{\mu_{\min} (1-\gamma)} \right) \\ &= \mathcal{O} \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min}^2 (1-\gamma)^2} \right), \end{aligned}$$

$$\begin{aligned} C_5 &= \left(\left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) \frac{L(A_1 + 1)}{\psi l_{cs}^2} + \frac{L A_2}{\psi l_{cs}^2} + \frac{L^2 u_{cm}^2}{2l_{cs}^4 \zeta_4^2 \psi^2} + \frac{3A_1^2 L \alpha^2}{2\psi l_{cs}^2} \right) \\ &= \mathcal{O} \left(\frac{\log(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^2} \cdot \frac{\log(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min} (1-\gamma)} + \frac{\log(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min} (1-\gamma)} + \frac{\log^2(|\mathcal{S}|)}{\mu_{\min}^3 (1-\gamma)^3} + \frac{\mathfrak{J}_{\max}^{2n-2} \log(|\mathcal{S}|)}{\mu_{\min} (1-\gamma)} \right) \\ &= \mathcal{O} \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{2n-2}}{\mu_{\min}^3 (1-\gamma)^3} \right), \end{aligned}$$

$$\begin{aligned} C_6 &= \left(\left(\frac{3u_{cm}^2}{2\zeta_3^2} + \frac{3}{2} \right) \frac{L(A_1 + 1)}{\psi l_{cs}^2} + \frac{m_1 L \alpha}{2l_{cs}^2 \psi} \right) = \mathcal{O} \left(\frac{\log(|\mathcal{S}|)}{\mu_{\min}^2 (1-\gamma)^2} \cdot \frac{\log(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min} (1-\gamma)} + \frac{\mathfrak{J}_{\max}^{n-1} \log(|\mathcal{S}|)}{\mu_{\min} (1-\gamma)} \right) \\ &= \mathcal{O} \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min}^3 (1-\gamma)^3} \right), \end{aligned}$$

$$C_7 = \frac{m_2^2 u_{cm}^2 L^2}{2\zeta_1^2 l_{cs}^4 \psi^2} \alpha = 0,$$

$$C_8 = \left(\frac{1}{2} + \frac{3L}{2\psi l_{cs}^2} \right) = \mathcal{O} \left(\frac{\log(|\mathcal{S}|)}{\mu_{\min} (1-\gamma)} \right),$$

$$C_9 = \frac{8u_{cD} B}{l_{cD}} = \mathcal{O} \left(\frac{\sqrt{|\mathcal{S}|} \mathfrak{J}_{\max}^n}{1-\gamma} \right),$$

$$C_{10} = \frac{8m_2 u_{cD}}{l_{cD} (1-\rho)} = 0,$$

$$C_{11} = \frac{8C_1^2 C_3^2 u_{cm}^2}{\zeta_6^2} = 0,$$

$$C_{12} = \frac{8C_4 u_{cD}^2 B^2}{l_{cD}^2} = \mathcal{O} \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min}^2 (1-\gamma)^2} \cdot |\mathcal{S}| \cdot \frac{\mathfrak{J}_{\max}^{2n}}{(1-\gamma)^2} \right) = \mathcal{O} \left(\frac{|\mathcal{S}| \log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{3n-1}}{\mu_{\min}^2 (1-\gamma)^4} \right),$$

$$C_{13} = \frac{14C_4 u_{cD}^2 m_2^2}{l_{cD}^2 (1-\rho^2)} = 0.$$

$$C_{14}(\tau) = \left(C_7 + C_{11} + 0.5C_3^2C_9^2 + C_3C_{10} + 2C_1C_3C_{10} + 3A_1C_3 + C_{13} + C_8 \frac{u_{cD}^2}{l_{cD}^2} 2m_2^2\alpha^2 \right) \tau^2 = 0,$$

$$\begin{aligned} C_{16}(\tau) &= \left(C_8 \frac{u_{cD}^2 B^2}{l_{cD}^2} + \frac{1}{2} + C_{12} \right) \tau^2 = \mathcal{O} \left(\frac{|\mathcal{S}| \log(|\mathcal{S}|) \mathfrak{J}_{\max}^{2n}}{(1-\gamma)^3 \mu_{\min}} + 1 + \left(\frac{|\mathcal{S}| \log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{3n-1}}{\mu_{\min}^2 (1-\gamma)^4} \right) \right) \tau^2 \\ &= \mathcal{O} \left(\frac{|\mathcal{S}| \log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{3n-1}}{(1-\gamma)^4 \mu_{\min}^2} \right) \tau^2, \end{aligned}$$

$$\begin{aligned} C_{17} &= (3A_1C_3 + 8A_1^2C_4 + C_5 + C_6) \\ &= \mathcal{O} \left(0 + \mathfrak{J}_{\max}^{2n-2} \cdot \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min}^2 (1-\gamma)^2} \right) + \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{2n-2}}{\mu_{\min}^3 (1-\gamma)^3} \right) + \left(\frac{\log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{n-1}}{\mu_{\min}^3 (1-\gamma)^3} \right) \right) \\ &= \mathcal{O} \left(\frac{\mathfrak{J}_{\max}^{3n-3} \log^2(|\mathcal{S}|)}{\mu_{\min}^3 (1-\gamma)^3} \right). \end{aligned}$$

Similar to $\|\cdot\|_D = \|\cdot\|_2$, we have $l_{c2} = \frac{1}{\sqrt{|\mathcal{S}|}}$ and $u_{c2} = 1$.

$$\tilde{A}_1 = \frac{2A_1^2 u_{c2}^2}{l_{c2}^2} = \mathcal{O}(\mathfrak{J}_{\max}^{2n-2} |\mathcal{S}|),$$

$$\tilde{A}_2 = \frac{2A_2^2 u_{c2}^2}{l_{c2}^2} = \mathcal{O}(\mathfrak{J}_{\max}^{2n-2} |\mathcal{S}|),$$

$$\tilde{B} = \frac{u_{c2}^2 B^2}{l_{c2}^2} = \mathcal{O} \left(\frac{|\mathcal{S}| \mathfrak{J}_{\max}^{2n}}{(1-\gamma)^2} \right),$$

$$\begin{aligned} M_0 &= \frac{1}{l_{cm}^2} \left(\frac{1}{C_1^2} \left(B + (A_2 + 1) \left(\|\theta_0\|_c + \frac{B}{2C_1} \right) \right)^2 + \|\theta_0\|_c^2 \right) = \mathcal{O} \left(\left(\frac{\mathfrak{J}_{\max}^n}{1-\gamma} + (\mathfrak{J}_{\max}^{n-1}) \cdot \frac{\mathfrak{J}_{\max}^n}{1-\gamma} \right)^2 + 1 \right) \\ &= \mathcal{O} \left(\frac{\mathfrak{J}_{\max}^{4n-2}}{(1-\gamma)^2} \right) \end{aligned}$$

$$C_1 = 16u_{cm}^2 M_0 \left(\log_\rho \frac{1}{e} + \frac{1}{\varphi_2} \right) = \mathcal{O} \left(\frac{\mathfrak{J}_{\max}^{4n-2}}{(1-\gamma)^2} \cdot \frac{1}{\mu_{\min} (1-\gamma)} \right) = \mathcal{O} \left(\frac{\mathfrak{J}_{\max}^{4n-2}}{(1-\gamma)^3 \mu_{\min}} \right)$$

$$\begin{aligned} C_2 &= \frac{8u_{cm}^2 (C_8 + \frac{1}{2} + C_{12})}{\varphi_2} = \mathcal{O} \left(\frac{1}{\mu_{\min} (1-\gamma)} \left(\frac{\log(|\mathcal{S}|)}{\mu_{\min} (1-\gamma)} + 1 + \frac{|\mathcal{S}| \log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{3n-1}}{\mu_{\min}^2 (1-\gamma)^4} \right) \right) \\ &= \mathcal{O} \left(\frac{|\mathcal{S}| \log^2(|\mathcal{S}|) \mathfrak{J}_{\max}^{3n-1}}{\mu_{\min}^2 (1-\gamma)^4} \right) \end{aligned}$$

$$\begin{aligned} C_3 &= \frac{80\tilde{B}^2 C_{17} u_{cm}^2 \left(1 + \frac{4m_2 \rho}{B(1-\rho)} \right)}{\varphi_2} = \mathcal{O} \left(\frac{1}{\mu_{\min} (1-\gamma)} \cdot \frac{|\mathcal{S}|^2 \mathfrak{J}_{\max}^{4n}}{(1-\gamma)^4} \cdot \frac{\mathfrak{J}_{\max}^{3n-3} \log^2(|\mathcal{S}|)}{\mu_{\min}^3 (1-\gamma)^3} \right) \\ &= \mathcal{O} \left(\frac{\mathfrak{J}_{\max}^{7n-3} |\mathcal{S}|^2 \log^2(|\mathcal{S}|)}{\mu_{\min}^4 (1-\gamma)^8} \right), \end{aligned}$$

$$C_4 = 8u_{cm}^2 (C_7 + C_{11} + 0.5C_3^2C_9^2 + C_3C_{10} + 2C_1C_3C_{10} + 3A_1C_3 + C_{13}) / \varphi_2 = 0.$$

Finally, for the sample complexity result, we simply employ Corollary B.1.1.

□

D. Federated Q-learning

In this section, we verify that the federated Q-learning algorithm 3 satisfies the properties of the FedSAM Algorithm 4. In the following, Q^* is the solution to the Bellman optimality equation (126)

$$Q^*(s, a) = \mathcal{R}(s, a) + \gamma \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s, a)} \left[\max_{a'} Q^*(S', a') \right]. \quad (126)$$

Note that Q^* is independent of the sampling policy of the agent. Furthermore, $\|\cdot\|_c = \|\cdot\|_\infty$.

Proposition D.1. *Federated Q-learning algorithm 3 is equivalent to the FedSAM Algorithm 4 with the following parameters.*

1. $\theta_t^i = Q_t^i - Q^*$
2. $S_t = (S_t^1, \dots, S_t^N)$ and $A_t = (A_t^1, \dots, A_t^N)$
3. $\mathbf{y}_t^i = (S_t^i, A_t^i, S_{t+1}^i, A_{t+1}^i)$ and $\mathbf{Y}_t = (S_t, A_t, S_{t+1}, A_{t+1})$
4. μ^i : Stationary distribution of the sampling policy of the i 'th agent.
5. $\mathbf{G}^i(\theta_t^i, \mathbf{y}_t^i)_{(s, a)} = \theta_t^i(s, a) + \mathbb{1}_{\{S_t^i=s, A_t^i=a\}} \times [\gamma \max_{a'} (\theta_t^i + Q^*(S_{t+1}^i, a')) - \theta_t^i(S_t^i, A_t^i) - \gamma \max_{a'} Q^*(S_{t+1}^i, a')]$
6. $\mathbf{b}^i(\mathbf{y}_t^i)_{(s, a)} = \mathbb{1}_{\{S_t^i=s, A_t^i=a\}} [\mathcal{R}(S_t^i, A_t^i) + \gamma \max_{a'} Q^*(S_{t+1}^i, a') - Q^*(S_t^i, A_t^i)]$

where $\mathbb{1}_A$ is the indicator function corresponding to set A , such that $\mathbb{1}_A = 1$ if A is true, and 0 otherwise.

Lemma D.1. *Consider the federated Q-learning Algorithm 3 as a special case of FedSAM (as specified in Proposition D.1). Suppose the trajectory $\{S_t^i, A_t^i\}_{t=0,1,\dots}$ converges geometrically fast to its stationary distribution as follows $d_{TV}(P(S_t^i = \cdot, A_t^i = \cdot | S_0^i, A_0^i) | | \mu^i(\cdot, \cdot)) \leq \bar{m} \bar{\rho}^t$ for all $i = 1, 2, \dots, N$. The corresponding $\bar{\mathbf{G}}^i(\theta)$ in Assumption 6.1 for the federated Q-learning is as follows*

$$\bar{\mathbf{G}}^i(\theta)_{(s, a)} = \theta(s, a) + \mu^i(s, a) \times \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s, a)} \left[\gamma \max_{a'} (\theta + Q^*(S', a')) - \theta(s, a) - \gamma \max_{a'} Q^*(S', a') \right].$$

Furthermore, we have $m_1 = 2A_2\bar{m}$, where A_2 is specified in Lemma D.3, $m_2 = 0$, and $\bar{\rho} = \rho$.

Lemma D.2. *Consider the federated Q-learning as a special case of FedSAM (as specified in Proposition D.1). The corresponding contraction factor γ_c in Assumption 6.2 for this algorithm is $\gamma_c = (1 - (1 - \gamma)\mu_{\min})$, where $\mu_{\min} = \min_{s, a, i} \mu^i(s, a)$*

Lemma D.3. *Consider the federated Q-learning as a special case of FedSAM (as specified in Proposition D.1). The constants A_1 , A_2 , and B in Assumption 6.2 are as follows: $A_1 = A_2 = 2$ and $B = \frac{2}{1-\gamma}$.*

Lemma D.4. *Consider the federated Q-learning as a special case of FedSAM (as specified in Proposition D.1). Assumption 6.4 holds for this algorithm.*

D.1. Proofs

Proof of Proposition D.1. Items 1-4 are by definition. Furthermore, by the update of the Q-learning, and subtracting Q^* from both sides, we have

$$\begin{aligned} & \underbrace{Q_{t+1}^i(s, a) - Q^*(s, a)}_{\theta_{t+1}^i(s, a)} = \underbrace{Q_t^i(s, a) - Q^*(s, a)}_{\theta_t^i(s, a)} \\ & + \alpha \mathbb{1}_{\{(s, a) = (S_t^i, A_t^i)\}} \left(\mathcal{R}(S_t^i, A_t^i) + \gamma \max_a Q_t^i(S_{t+1}^i, a) - Q_t^i(S_t^i, A_t^i) \right) \\ & = \theta_t^i(s, a) \\ & + \alpha \mathbb{1}_{\{(s, a) = (S_t^i, A_t^i)\}} \left(\mathcal{R}(S_t^i, A_t^i) + \gamma \max_a \underbrace{Q_t^i(S_{t+1}^i, a) - Q^*(S_{t+1}^i, a)}_{\theta_t^i(S_{t+1}^i, a)} + Q^*(S_{t+1}^i, a) \right) \end{aligned}$$

$$\begin{aligned}
 & - \left(\underbrace{Q_t^i(S_t^i, A_t^i) - Q^*(S_t^i, A_t^i)}_{\boldsymbol{\theta}_t^i(S_t^i, A_t^i)} + Q^*(S_t^i, A_t^i) \right) \quad (\text{addition and subtraction}) \\
 & = \boldsymbol{\theta}_t^i(s, a) \\
 & + \alpha \left(\underbrace{\boldsymbol{\theta}_t^i(s, a) + \mathbb{1}_{\{S_t^i=s, A_t^i=a\}} \times [\gamma \max_{a'} (\boldsymbol{\theta}_t^i + Q^*(S_{t+1}^i, a')) - \boldsymbol{\theta}_t^i(S_t^i, A_t^i) - \gamma \max_{a'} Q^*(S_{t+1}^i, a')]}_{\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)_{(s,a)}} - \boldsymbol{\theta}_t^i(s, a) \right. \\
 & \quad \left. + \mathbb{1}_{\{S_t^i=s, A_t^i=a\}} \underbrace{\left(\mathcal{R}(S_t^i, A_t^i) + \gamma \max_{a'} Q^*(S_{t+1}^i, a') - Q^*(S_t^i, A_t^i) \right)}_{\mathbf{b}^i(\mathbf{y}_t^i)_{(s,a)}} \right),
 \end{aligned}$$

which proves items 5 and 6. Furthermore, for the synchronization part of Q -learning, we have

$$\begin{aligned}
 Q_t^i & \leftarrow \frac{1}{N} \sum_{j=1}^N Q_t^j \\
 \implies \underbrace{Q_t^i - Q^*}_{\boldsymbol{\theta}_t^i} & \leftarrow \frac{1}{N} \sum_{j=1}^N \underbrace{(Q_t^j - Q^*)}_{\boldsymbol{\theta}_t^j},
 \end{aligned}$$

which is equivalent to the synchronization step in FedSAM Algorithm 4. Notice that here we used the fact that all agents have the same fixed point Q^* . \square

Proof of Lemma D.1. $\bar{\mathbf{G}}(\boldsymbol{\theta})_{(s,a)}$ can be found by simply taking expectation of $\mathbf{G}^i(\boldsymbol{\theta}_t^i, \mathbf{y}_t^i)_{(s,a)}$, defined in Proposition D.1, with respect to the stationary distribution μ^i . Furthermore, we have

$$\begin{aligned}
 & \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}) - \mathbb{E}[\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)]\|_c \\
 & = \left\| \mathbb{E}_{\mathbf{y}_t^i \sim \mu^i} [\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)] - \mathbb{E}[\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)] \right\|_c \\
 & = \left\| \sum_{s,a,s',a'} (\mu^i(s, a, s', a') - P(S_t^i = s, A_t^i = a, S_{t+1}^i = s', A_{t+1}^i = a' | S_0^i, A_0^i)) \mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i) \right\|_c \\
 & = \sum_{s,a,s',a'} |\mu^i(s, a, s', a') - P(S_t^i = s, A_t^i = a, S_{t+1}^i = s', A_{t+1}^i = a' | S_0^i, A_0^i)| \cdot \|\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y}_t^i)\|_c \quad (\|ax\|_c = |a|\|x\|_c) \\
 & \leq \sum_{s,a,s',a'} |\mu^i(s, a, s', a') - P(S_t^i = s, A_t^i = a, S_{t+1}^i = s', A_{t+1}^i = a' | S_0^i, A_0^i)| \cdot A_2 \|\boldsymbol{\theta}\|_c \quad (\text{Assumption D.3}) \\
 & = \sum_{s,a,s',a'} \left| \mu^i(s, a) \mathcal{P}(s'|s, a) \pi^i(a'|s') - P(S_t^i = s, A_t^i = a | S_0^i, A_0^i) \mathcal{P}(s'|s, a) \pi^i(a'|s') \right| \cdot A_2 \|\boldsymbol{\theta}\|_c \\
 & \quad (\text{definition of transition probability}) \\
 & = \sum_{s,a} \left| \mu^i(s, a) - P(S_t^i = s, A_t^i = a | S_0^i, A_0^i) \right| \cdot A_2 \|\boldsymbol{\theta}\|_c \\
 & = 2d_{TV}(\mu^i(\cdot, \cdot), P(S_t^i = \cdot, A_t^i = \cdot | S_0^i, A_0^i)) \cdot A_2 \|\boldsymbol{\theta}\|_c \\
 & \leq 2A_2 \|\boldsymbol{\theta}\|_c \bar{m} \bar{\rho}^t.
 \end{aligned}$$

In addition, we have

$$\begin{aligned}
 \|\mathbb{E}[\mathbf{b}^i(\mathbf{y}_t^i)]\|_c & = \max_{s,a} \left| P(S_t^i = s, A_t^i = a | S_0^i, A_0^i) \left(\mathcal{R}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot | s, a)} [\max_{a'} Q^*(S', a')] - Q^*(s, a) \right) \right| \\
 & = 0. \quad (\text{Bellman optimality equation (126)})
 \end{aligned}$$

\square

Proof of Lemma D.2.

$$\begin{aligned}
 & \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_1) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_2)\|_c = \max_{s,a} |\bar{\mathbf{G}}^i(\boldsymbol{\theta}_1)_{(s,a)} - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_2)_{(s,a)}| \\
 &= \max_{s,a} \left| \boldsymbol{\theta}_1(s,a) + \mu^i(s,a) \times \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left[\gamma \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \boldsymbol{\theta}_1(s,a) - \gamma \max_{a'} Q^*(S', a') \right] \right. \\
 & \quad \left. - \left(\boldsymbol{\theta}_2(s,a) + \mu^i(s,a) \times \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left[\gamma \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) - \boldsymbol{\theta}_2(s,a) - \gamma \max_{a'} Q^*(S', a') \right] \right) \right| \\
 &= \max_{s,a} \left| (1 - \mu^i(s,a))(\boldsymbol{\theta}_1(s,a) - \boldsymbol{\theta}_2(s,a)) \right. \\
 & \quad \left. + \gamma \mu^i(s,a) \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left[\max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right] \right| \\
 &\leq \max_{s,a} \left[\left| (1 - \mu^i(s,a))(\boldsymbol{\theta}_1(s,a) - \boldsymbol{\theta}_2(s,a)) \right| \right. \\
 & \quad \left. + \gamma \left| \mu^i(s,a) \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left[\max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right] \right| \right] \quad (\text{triangle inequality}) \\
 &\leq \max_{s,a} \left[(1 - \mu^i(s,a)) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \right. \\
 & \quad \left. + \gamma \mu^i(s,a) \left[\left| \mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left[\max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right] \right| \right] \right] \quad (\text{definition of } \|\cdot\|_\infty) \\
 &\leq \max_{s,a} \left[(1 - \mu^i(s,a)) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \right. \\
 & \quad \left. + \gamma \mu^i(s,a) \left[\mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left| \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right| \right] \right] \quad (\mu^i(s,a) \geq 0) \\
 &\leq \max_{s,a} \left[(1 - \mu^i(s,a)) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \right. \\
 & \quad \left. + \gamma \mu^i(s,a) \left[\mathbb{E}_{S' \sim \mathcal{P}(\cdot|s,a)} \left| \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right| \right] \right] \quad (\text{Jensen's inequality})
 \end{aligned}$$

Next, we note that for any functions $f(\cdot)$ and $g(\cdot)$, we have

$$\left| (\max_x f(x)) - (\max_x g(x)) \right| \leq \max_x |f(x) - g(x)|. \quad (127)$$

The reason is as follows. We have $\max_x f(x) = \max_x f(x) - g(x) + g(x) \leq (\max_x f(x) - g(x)) + (\max_x g(x))$. Hence, $(\max_x f(x)) - (\max_x g(x)) \leq \max_x f(x) - g(x) \leq \max_x |f(x) - g(x)|$. Now suppose $\max_x f(x) \geq \max_x g(x)$. Then we can apply absolute value to the left hand side of the inequality, and we get the bound. By a similar argument for the case $\max_x f(x) \leq \max_x g(x)$, we get the bound in (127). Hence, we have $|\max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a'))| \leq \max_{a'} |\boldsymbol{\theta}_1(S', a') - \boldsymbol{\theta}_2(S', a')| \leq \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty$. As a result, we have

$$\begin{aligned}
 \|\bar{\mathbf{G}}^i(\boldsymbol{\theta}_1) - \bar{\mathbf{G}}^i(\boldsymbol{\theta}_2)\|_c &\leq \max_{s,a} \left[(1 - \mu^i(s,a)) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \gamma \mu^i(s,a) [\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty] \right] \\
 &= \max_{s,a} \left(1 - (1 - \gamma) \mu^i(s,a) \right) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \\
 &\leq (1 - (1 - \gamma) \mu_{\min}^i) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \\
 &\leq (1 - (1 - \gamma) \mu_{\min}) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \\
 &= (1 - (1 - \gamma) \mu_{\min}) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_c.
 \end{aligned}$$

□

Proof of Lemma D.3. First, for A_1 , we have

$$\|\mathbf{G}^i(\boldsymbol{\theta}_1, \mathbf{y}) - \mathbf{G}^i(\boldsymbol{\theta}_2, \mathbf{y})\|_c$$

$$\begin{aligned}
 &= \max_{s,a} \left| \boldsymbol{\theta}_1(s, a) + \mathbb{1}_{\{S=s, A=a\}} \times [\gamma \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \boldsymbol{\theta}_1(S, A) - \gamma \max_{a'} Q^*(S', a')] \right. \\
 &\quad \left. - \left(\boldsymbol{\theta}_2(s, a) + \mathbb{1}_{\{S=s, A=a\}} \times [\gamma \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) - \boldsymbol{\theta}_2(S, A) - \gamma \max_{a'} Q^*(S', a')] \right) \right| \\
 &\leq \max_{s,a} \left[|(1 - \mathbb{1}_{\{S=s, A=a\}})(\boldsymbol{\theta}_1(s, a) - \boldsymbol{\theta}_2(s, a))| \right. \\
 &\quad \left. + \gamma \mathbb{1}_{\{S=s, A=a\}} \left| \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right| \right] \quad (\text{triangle inequality}) \\
 &\leq \max_{s,a} \left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \gamma \mathbb{1}_{\{S=s, A=a\}} \left| \max_{a'} (\boldsymbol{\theta}_1 + Q^*(S', a')) - \max_{a'} (\boldsymbol{\theta}_2 + Q^*(S', a')) \right| \right] \quad (\text{definition of } \|\cdot\|_\infty) \\
 &\leq \max_{s,a} \left[\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty + \gamma \mathbb{1}_{\{S=s, A=a\}} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \right] \quad (\text{By (127)}) \\
 &\leq 2\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_\infty \\
 &= 2\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_c
 \end{aligned}$$

Second, for A_2 , we have

$$\begin{aligned}
 \|\mathbf{G}^i(\boldsymbol{\theta}, \mathbf{y})\|_c &= \max_{s,a} \left| \boldsymbol{\theta}(s, a) + \mathbb{1}_{\{S=s, A=a\}} \times [\gamma \max_{a'} (\boldsymbol{\theta} + Q^*(S', a')) - \boldsymbol{\theta}(S, A) - \gamma \max_{a'} Q^*(S', a')] \right| \\
 &\leq \max_{s,a} \left[|(1 - \mathbb{1}_{\{S=s, A=a\}})\boldsymbol{\theta}(a, s)| + \gamma \mathbb{1}_{\{S=s, A=a\}} \left| \max_{a'} (\boldsymbol{\theta} + Q^*(S', a')) - \max_{a'} Q^*(S', a') \right| \right] \quad (\text{triangle inequality}) \\
 &\leq \max_{s,a} \left[\|\boldsymbol{\theta}\|_\infty + \gamma \mathbb{1}_{\{S=s, A=a\}} \left| \max_{a'} (\boldsymbol{\theta} + Q^*(S', a')) - \max_{a'} Q^*(S', a') \right| \right] \quad (\text{definition of } \|\cdot\|_\infty) \\
 &\leq \max_{s,a} \left[\|\boldsymbol{\theta}\|_\infty + \gamma \max_{a'} |\boldsymbol{\theta}(S', a')| \right] \quad (\text{By (127)}) \\
 &\leq 2\|\boldsymbol{\theta}\|_\infty \\
 &= 2\|\boldsymbol{\theta}\|_c.
 \end{aligned}$$

Lastly, for B , we have

$$\begin{aligned}
 \|b^i(\mathbf{y}^i)\|_c &= \max_{s,a} \left| \mathbb{1}_{\{S=s, A=a\}} \left[\mathcal{R}(S, A) + \gamma \max_{a'} Q^*(S', a') - Q^*(S, A) \right] \right| \\
 &\leq \max_{s,a} \mathbb{1}_{\{S=s, A=a\}} \left[|\mathcal{R}(S, A)| + \gamma \max_{a'} |Q^*(S', a')| + |Q^*(S, A)| \right] \quad (\text{triangle inequality}) \\
 &\leq \max_{s,a} \left[1 + \gamma \frac{1}{1-\gamma} + \frac{1}{1-\gamma} \right] \\
 &= \frac{2}{1-\gamma}.
 \end{aligned}$$

□

Proof of Lemma D.4. The proof follows similar to Lemma C.4. □

Proof of Theorem 5.2. By Proposition D.1 and Lemmas D.1, D.2, D.3, and D.4, it is clear that the federated Q Algorithm 3 satisfies all the Assumptions 6.1, 6.2, 6.3, and 6.4 of the FedSAM Algorithm 4. Furthermore, by the proof of Theorem B.1, we have $w_t = (1 - \frac{\alpha\varphi_2}{2})^{-t}$, and the constant c in the sampling distribution q_T^c in Algorithm 1 is $c = (1 - \frac{\alpha\varphi_2}{2})^{-1}$. In equation (128) we evaluate the exact value of w_t .

Furthermore, by choosing step size α small enough, we can satisfy the requirements in (21), (23), (32), (35). By choosing K large enough, we can satisfy $K > \tau$, and by choosing T large enough we can satisfy $T > K + \tau$. Hence, the result of Theorem B.1 holds for this algorithm.

Next, we derive the constants involved in Theorem B.1 step by step. In this analysis we only consider the terms involving $|\mathcal{S}|$, $|\mathcal{A}|$, $\frac{1}{1-\gamma}$, \mathcal{J}_{\max} , and μ_{\min} . Since $\|\cdot\|_c = \|\cdot\|_\infty$, we choose $g(\cdot) = \frac{1}{2}\|\cdot\|_p^2$, i.e. the p -norm with $p = 2\log(|\mathcal{S}|)$.

It is known that $g(\cdot)$ is $(p-1)$ smooth with respect to $\|\cdot\|_p$ norm (Beck, 2017), and hence $L = \Theta(\log(|\mathcal{S}|))$. Hence, we have $l_{cs} = |\mathcal{S}|^{-1/p} = \frac{1}{\sqrt{e}} = \Theta(1)$ and $u_{cs} = 1$. Therefore, we have $\varphi_1 = \frac{1+\psi u_{cs}^2}{1+\psi l_{cs}^2} = \frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}} \leq 1+\psi$. By choosing $\psi = (\frac{1+\gamma_c}{2\gamma_c})^2 - 1 = \frac{1+2\gamma_c-3\gamma_c^2}{4\gamma_c^2} \geq (1-\gamma_c) = \mu_{\min}(1-\gamma) = \Omega(\mu_{\min}(1-\gamma))$, which is $\psi = \mathcal{O}(1)$, we have $\varphi_1 = \frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}} = \sqrt{e} \frac{(\frac{1+\gamma_c}{2\gamma_c})^2}{\sqrt{e} + (\frac{1+\gamma_c}{2\gamma_c})^2 - 1} = \mathcal{O}(1)$, and

$$\begin{aligned} \varphi_2 &= 1 - \gamma_c \sqrt{\frac{1+\psi}{1+\frac{\psi}{\sqrt{e}}}} = 1 - \gamma_c \sqrt{\frac{(\frac{1+\gamma_c}{2\gamma_c})^2}{1+\frac{(\frac{1+\gamma_c}{2\gamma_c})^2-1}{\sqrt{e}}}} = 1 - \frac{0.5(1+\gamma_c)e^{1/4}}{\sqrt{\sqrt{e}-1 + \left(\frac{1+\gamma_c}{2\gamma_c}\right)^2}} = 1 - \frac{0.5e^{1/4}(2-\mu_{\min}(1-\gamma))}{\sqrt{\sqrt{e}-1 + \left(\frac{2-\mu_{\min}(1-\gamma)}{2-2\mu_{\min}(1-\gamma)}\right)^2}} \\ &> 1 - \gamma_c \sqrt{1+\psi} = 1 - \gamma_c \frac{1+\gamma_c}{2\gamma_c} = \frac{1-\gamma_c}{2} = 0.5\mu_{\min}(1-\gamma) = \Omega(\mu_{\min}(1-\gamma)) \end{aligned}$$

$$\varphi_3 = \frac{L(1+\psi u_{cs}^2)}{\psi l_{cs}^2} = \mathcal{O}\left(\frac{\log(|\mathcal{S}|)(1+\psi)}{\psi}\right) \leq \mathcal{O}\left(\frac{\log(|\mathcal{S}|)}{1-\gamma_c}\right) = \mathcal{O}\left(\frac{\log(|\mathcal{S}|)}{\mu_{\min}(1-\gamma)}\right).$$

Using φ_2 , we have

$$w_t = \left(1 - \frac{\alpha\varphi_2}{2}\right)^{-t} = \left(1 - \alpha/2 + \frac{0.25\alpha e^{1/4}(2-\mu_{\min}(1-\gamma))}{\sqrt{\sqrt{e}-1 + \left(\frac{2-\mu_{\min}(1-\gamma)}{2-2\mu_{\min}(1-\gamma)}\right)^2}}\right)^{-t}. \quad (128)$$

Further, we have

$$\begin{aligned} l_{cm} &= (1+\psi l_{cs}^2)^{1/2} = \Theta(1) \\ u_{cm} &= (1+\psi u_{cs}^2)^{1/2} = \Theta(1) \end{aligned}$$

Since TV-divergence is upper bounded with 1, we have $\bar{m} = \mathcal{O}(1)$. By Lemma D.3, we have

$$A_1 = A_2 = 2 = \mathcal{O}(1)$$

and $A_1 = A_2 = \Omega(1)$,

$$B = \frac{2}{1-\gamma} = \mathcal{O}\left(\frac{1}{1-\gamma}\right).$$

Hence $m_1 = 2A_2\bar{m} = \mathcal{O}(1)$. Also, we have $m_2 = 0$.

We choose the D -norm in Lemma B.9 as the 2-norm $\|\cdot\|_2$. Hence, by primary norm equivalence, we have $l_{cD} = \frac{1}{\sqrt{|\mathcal{S}|}}$, and $u_{cD} = 1$, and hence $\frac{u_{cD}}{l_{cD}} = \sqrt{|\mathcal{S}|}$. The rest of the proof is similar to the proof of Theorem 5.1 where \mathfrak{J}_{\max} is substituted with 1. The sample complexity can also be derived using Corollary B.1.1.

□