

MDPGT: Momentum-based Decentralized Policy Gradient Tracking

Zhanhong Jiang¹, Xian Yeow Lee², Sin Yong Tan², Kai Liang Tan²,
Aditya Balu², Young M. Lee¹, Chinmay Hegde³, Soumik Sarkar²

¹ Johnson Controls Inc., 507 East Michigan St, Milwaukee, WI 53202,

²Iowa State University, Ames, IA 50010,

³New York University, 6 MetroTech Center, Brooklyn, NY 11201,

{zhanhong.jiang, young.m.lee}@jci.com, {xylee, tsyong98, kailiang, baditya, soumiks}@iastate.edu, {chinmay.h}@nyu.edu

Abstract

We propose a novel policy gradient method for multi-agent reinforcement learning, which leverages two different variance-reduction techniques and does not require large batches over iterations. Specifically, we propose a momentum-based decentralized policy gradient tracking (MDPGT) where a new momentum-based variance reduction technique is used to approximate the local policy gradient surrogate with importance sampling, and an intermediate parameter is adopted to track two consecutive policy gradient surrogates. Moreover, MDPGT provably achieves best available sample complexity of $\mathcal{O}(N^{-1}\epsilon^{-3})$ for converging to an ϵ -stationary point of the global average of N local performance functions (possibly nonconcave). This outperforms the state-of-the-art sample complexity in decentralized model-free reinforcement learning and when initialized with a single trajectory, the sample complexity matches those obtained by the existing decentralized policy gradient methods. We further validate the theoretical claim for the Gaussian policy function. When the required error tolerance ϵ is small enough, MDPGT leads to a linear speed up, which has been previously established in decentralized stochastic optimization, but not for reinforcement learning. Lastly, we provide empirical results on a multi-agent reinforcement learning benchmark environment to support our theoretical findings.

Introduction

Multi-agent reinforcement learning (MARL) is an emerging topic which has been explored both in theoretical (Nguyen et al. 2014; Zhang et al. 2018; Qu et al. 2019; Zhang et al. 2021c) and empirical settings (Helou, Kalathil, and Xie 2020; Mukherjee, Bai, and Chakraborty 2020; Zhou et al. 2020). Several appealing applications of MARL can be seen in (Zhang, Yang, and Bašar 2019; Nguyen, Nguyen, and Nahavandi 2020) and relevant references therein.

While MARL can primarily be cast into two different categories, i.e., cooperative (Li, Chen, and Chen 2020; Wang et al. 2020; Li et al. 2020) and competitive (Chen et al. 2020), our focus is in the cooperative setting; see (Wei et al. 2021) for details on the competitive setting. Cooperative MARL is typically modeled as a networked multi-agent Markov decision process (MDP) (Zhang, Yang, and Basar 2018; Chu, Chinchali, and Katti 2020; Zhang et al. 2018) in which the

agents share a centralized reward function (Simões, Lau, and Reis 2020; Ackermann et al. 2019). However, in practice, this is not necessarily the case, and instead a more general yet challenging scenario is that agents have *heterogeneous* reward functions. Inherently, the ultimate goal in such a cooperative MARL setting is for agents to maximize the *global average* of local long-term returns. To address this problem, various algorithms have been proposed, including distributed-learning (Arslan and Yüksel 2016; Nguyen and Mukhopadhyay 2017) and distributed actor-critic (Li et al. 2020; Ryu, Shin, and Park 2020). More recent works have successfully showed finite-sample analysis for decentralized batch MARL (Zhang et al. 2021c) and leveraged advances in analysis of descent-ascent algorithms (Lu et al. 2021).

These preliminary attempts have facilitated the theoretical understanding of cooperative MARL by showing explicit sample complexity bounds, which match that of standard (vanilla) stochastic gradient descent (SGD). Additionally, recent works (Huang et al. 2020; Xu, Gao, and Gu 2019) in centralized RL have revealed that with simple *variance reduction* techniques, this sample complexity can be reduced to $\mathcal{O}(\epsilon^{-3})$ to reach an ϵ -stationary point (i.e., $\mathbb{E}[\|\nabla J(\mathbf{x})\|] \leq \epsilon$, where J is the return function and $\mathbf{x} \in \mathbb{R}^d$ is the decision variable to be optimized), which has been admitted as the best complexity in decentralized optimization (Das et al. 2020; Karimireddy et al. 2020). However, no similar matching bounds have yet been reported in the decentralized (cooperative MARL) setting. Hence, this motivates the question:

Can we achieve a sample complexity of $\mathcal{O}(\epsilon^{-3})$ in decentralized MARL via variance reduction?

In this paper, we answer this question affirmatively by proposing a variance-reduced policy gradient tracking approach, MDPGT (Algorithm 1), and analyzing it in Theorem 1. Additionally, we propose a variation (based on a different initialization) that enables state-of-the-art (SOTA) sample complexity for decentralized MARL (Zhang et al. 2021c; Lu et al. 2021). See Table 1 for SOTA comparisons. Specifically:

1. We propose MDPGT, in which we use a stochastic policy gradient surrogate, a convex combination of the vanilla stochastic policy gradient and an importance sampling-based stochastic recursive algorithm (SARAH) (Nguyen et al. 2017) for the local gradient update. Instead of directly applying the stochastic policy gradient surro-

gate in the parameter update, an intermediate parameter is adopted to track the difference between two consecutive stochastic policy gradient surrogates. For smooth nonconcave performance functions, we show that MDPGT with the mini-batch initialization can converge to an ϵ -stationary point in $\mathcal{O}(N^{-1}\epsilon^{-3})$ gradient-based updates which matches the best available known upper bounds (Huang et al. 2020).

2. We modify the initialization of the proposed algorithm MDPGT by using a single trajectory instead of a mini-batch of trajectories. Surprisingly, we find that *only one trajectory* results in a larger sampling complexity $\mathcal{O}(N^{-1}\epsilon^{-4})$, which, however, is the same as obtained by the SOTA (Zhang et al. 2021c; Lu et al. 2021) with a linear speed up when ϵ is sufficiently small. Additionally, our algorithm shows that when updating the policy parameter in MDPGT, the mini-batch size is $\mathcal{O}(1)$ instead of being ϵ -related (Xu, Gao, and Gu 2019; Qu et al. 2019), which can significantly improve practical efficiency.

3. To facilitate the theoretical understanding of MDPGT, we leverage a benchmark gridworld environment for numerical simulation and compare our proposed algorithm to a baseline decentralized policy gradient (DPG) and the momentum-based decentralized policy gradient (MDPG, described in Algorithm 2 in the supplementary materials), which is a new variant created in this work for the purpose of empirical comparison. We show that our theoretical claims are valid based on the experiments.

Related Works. Most previous decentralized MARL papers (Zhang et al. 2018; Suttle et al. 2020; Li et al. 2020; Chen et al. 2020; Bono et al. 2018) tend to focus on convergence to the optimal return. Exceptions include (Qu et al. 2019), where they proved non-asymptotic convergence rates with nonlinear function approximation using value propagation. This enables us to approximately derive the number of stochastic gradient evaluations. However, the algorithm involves the complex inner-outer structure and requires the size of the mini-batch to be \sqrt{K} , with K being the number of iterations, which may not be practically implementable. Zhang et al. (2021c) obtain $\mathcal{O}(\epsilon^{-4})$ for the cooperative setting by using gradient tracking (GT), which is a bias correction technique dedicated to decentralized optimization, but with several specifically imposed assumptions, such as stationary sample paths, which may not be realistic. Lu et al. (2021) also utilize GT but require dual parameter updates to achieve $\mathcal{O}(\epsilon^{-4})$; our approach is different and simpler. In this context, we mention that centralized counterparts of MARL (Huang et al. 2020; Xu, Gao, and Gu 2019; Papini et al. 2018) have achieved sample complexity of $\mathcal{O}(\epsilon^{-3})$. However, in both Xu, Gao, and Gu (2019) and Papini et al. (2018), the size of mini-batch is ϵ -related, which is more computationally sophisticated than those in both (Huang et al. 2020) and our proposed method. We provide additional discussion of related work in the supplementary materials.

Preliminaries

We first formulate MARL, followed by an overview of variance reduction techniques and decentralized policy gradients.

MARL Formulation

In this context, we consider a networked system involving multiple agents (say N) that *collaboratively* solve *dynamic* optimization problems. Specifically, the system can be quantified as a graph, i.e., $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, 2, \dots, N\}$ is the vertex set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set. Throughout the paper, we assume that \mathcal{G} is *static and undirected*, though in a few previous works (Lin et al. 2019; Suttle et al. 2020; Zhang et al. 2018) \mathcal{G} could be directed. The goal of this work is to provide the rigorous theoretical analysis for our decentralized MARL algorithm, with the property of the graph not being the main focus. When a pair of agents i and j can communicate with each other, we have $(i, j) \in \mathcal{E}$. We also define the neighborhood of a specific agent i , $Nb(i)$, such that $Nb(i) \triangleq \{j | j \in \mathcal{V}, (i, j) \in \mathcal{E} \text{ or } j = i\}$. Only agents in $Nb(i)$ are able to communicate with the agent i . We next present the definition of networked MARL on top of \mathcal{G} .

With multiple agents, the networked Markov decision process is thus characterized by a tuple $(\mathcal{S}, \{\mathcal{A}^i\}_{i \in \mathcal{V}}, \mathcal{P}, \{r^i\}_{i \in \mathcal{V}}, \mathcal{G}, \gamma)$, where \mathcal{S} indicates a global state space shared by all agents in \mathcal{V} with $|\mathcal{S}| < \infty$, \mathcal{A}^i signifies the action space specified for agent i , and $\gamma \in (0, 1]$ is the discount factor. Moreover, in the cooperative MARL setting, the environment is driven by the *joint* action space instead of individual action spaces. Thus, $\mathcal{A} \triangleq \prod_{i \in \mathcal{V}} \mathcal{A}^i$ is defined as the joint action space over all agents in \mathcal{V} . $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ represents the probability function to transition the current state to the next state. $\{r^i\}_{i \in \mathcal{V}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the local reward function of agent i and $r^i \in [-R, R]$ ($R > 0$). Additionally, states and actions are assumed to be *globally* observable, while the rewards are only locally observable. Such an assumption corresponds to the definition of the cooperative MARL setting and has been generic in previous works (Zhang et al. 2018; Zhang, Yang, and Basar 2018; Zhang et al. 2021c).

We next describe how agents behave in such an environment. Suppose that the current state of the environment is $s_k \in \mathcal{S}$, where k is the time step. Each agent i chooses its own action $a_k^i \in \mathcal{A}^i$, based on the local policy, $\pi^i : \mathcal{S} \times \mathcal{A}^i \rightarrow [0, 1]$. For a parameterized policy, we denote by $\pi_{\mathbf{x}^i}^i(s, a^i)$, which indicates the probability of agent i choosing action a^i given the current state s and $\mathbf{x}^i \in \mathbb{R}^{d_i}$ here is the policy parameter. Stacking all local policy parameter together yields $\mathbf{x} = [(\mathbf{x}^1)^\top, (\mathbf{x}^2)^\top, \dots, (\mathbf{x}^N)^\top]^\top \in \mathbb{R}^{\sum_{i \in \mathcal{V}} d_i}$. Hence, the joint policy can be denoted as $\pi_{\mathbf{x}} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, where $\pi_{\mathbf{x}}(s, a) \triangleq \prod_{i \in \mathcal{V}} \pi_{\mathbf{x}^i}^i(s, a^i)$ and $a \in \mathcal{A}$. In this context, the decisions are decentralized due to locally observable rewards, locally evaluated policies and locally executed actions. To simplify the notations, we drop the \mathbf{x}^i for $\pi_{\mathbf{x}^i}^i$ and \mathbf{x} for $\pi_{\mathbf{x}}$ respectively for local and joint policies throughout the rest of the paper. With the joint policy π and the state transition function \mathcal{P} , the environment evolves from s to s' with the probability $\mathcal{P}(s'|s, a)$. Another assumption imposed in this paper for the policy function is that for all $i \in \mathcal{V}, s \in \mathcal{S}, a^i \in \mathcal{A}^i, \pi^i(s, a^i)$ is continuously differentiable w.r.t. all $\mathbf{x}^i \in \mathbb{R}^{d_i}$. Such an assumption will assist in characterizing the smoothness of the objective function.

The goal for each agent is to learn a local policy π_*^i such

Table 1: Comparisons between existing and proposed approaches.

| Method | Complexity | Decentralized | Variance Reduction | Linear Speed Up | I.S. |
|------------------------------------|------------------------------|---------------|--------------------|-----------------|------|
| MBPG (Huang et al. 2020) | $\mathcal{O}(\epsilon^{-3})$ | ✗ | ✓ | ✗ | ✓ |
| Value Prop (Qu et al. 2019) | $\mathcal{O}(\epsilon^{-4})$ | ✓ | ✗ | ✗ | ✗ |
| DCPG (Zeng et al. 2020) | $\mathcal{O}(\epsilon^{-4})$ | ✓ | ✗ | ✗ | ✗ |
| Safe-Dec-PG (Lu et al. 2021) | $\mathcal{O}(\epsilon^{-4})$ | ✓ | ✓ | ✗ | ✗ |
| DFQI (Zhang et al. 2021c) | $\mathcal{O}(\epsilon^{-4})$ | ✓ | ✓ | ✗ | ✗ |
| Dec-TD (0) +GT (Lin and Ling 2021) | N/A | ✓ | ✓ | ✗ | ✗ |
| MDPGT (ours) | $\mathcal{O}(\epsilon^{-4})$ | ✓ | ✓ | ✓ | ✓ |
| MDPGT-MI (ours) | $\mathcal{O}(\epsilon^{-3})$ | ✓ | ✓ | ✓ | ✓ |

¹ **Complexity:** Sampling complexity for achieving $\mathbb{E}[\|\nabla J(\mathbf{x})\|] \leq \epsilon$. ² **Linear Speed Up:** If an algorithm has $\mathcal{O}(1/\sqrt{K})$ convergence, then its sampling complexity of attaining an $\mathcal{O}(\epsilon)$ accurate solution is ϵ^{-2} . Similarly, $\mathcal{O}(1/\sqrt{NK})$ corresponds to $N^{-1}\epsilon^{-2}$, which is N times faster than the former. Typically, K has to satisfy a certain condition. ³ MDPGT-MI is MDPGT with mini-batch initialization. We use this notation for conveniently classifying two different initialization approaches. In the rest of paper, we still adopt MDPGT to unify these two approaches. ⁴ I.S. denotes the utilization of importance sampling.

that the joint policy π_* is able to maximize the *global average* of expected cumulative discounted rewards, i.e.,

$$\pi_* = \operatorname{argmax}_{\mathbf{x} \in \mathbb{R}^d} J(\mathbf{x}) \triangleq \frac{1}{N} \sum_{i \in \mathcal{V}} \mathbb{E} \left[\sum_{h=0}^H \gamma^h r_h^i \right], \quad (1)$$

where H is the horizon and $d = \sum_{i \in \mathcal{V}} d_i$. Several works (Zhang et al. 2018; Qu et al. 2019; Zhang et al. 2021c; Lin et al. 2019; Suttle et al. 2020) have made their attempts to resolve this optimization problem, leading to different algorithms. Since each agent only has access to local information, a communication protocol needs to be introduced in the system, as done in decentralized optimization. With that, well-known centralized policy-based algorithms can be extended as MARL algorithms. Nevertheless, one issue that has not been sufficiently explored is *the inherent policy gradient variance, which could even be more significant in the MARL algorithms*. Consequently, this work propose novel MARL algorithms to investigate how to reduce the policy gradient variance during the optimization.

Variance Reduction and Bias Correction

In stochastic optimization, variance reduction techniques have been well studied and applied to either centralized or decentralized gradient descent type of algorithms, such as SVRG (Johnson and Zhang 2013), SARAH (Nguyen et al. 2017), SPIDER (Fang et al. 2018), Hybrid-SARAH (Tran-Dinh et al. 2019) and STORM (Cutkosky and Orabona 2019). In another line of work, the GT technique (Pu and Nedić 2020; Sun, Daneshmand, and Scutari 2019) was proposed specifically for consensus-based decentralized optimization techniques to improve the convergence rate by tracking and correcting each agent's locally aggregated gradients. In our work, we leverage both Hybrid-SARAH and GT to reduce the policy gradient variance and correct the policy gradient bias respectively in the MARL and achieve the best convergence rate. Hybrid-SARAH performs with a trade-off parameter to balance the effect between vanilla stochastic

gradient and SARAH. More detail on these techniques are elaborated in the supplementary materials.

So far, we are not aware of existing results that have successfully shown SARAH or Hybrid-SARAH type of variance reduction techniques well suited for *decentralized non-oblivious* learning problems, e.g., MARL. Consequently, the regular Hybrid-SARAH technique cannot be directly applied to MARL; we address this challenge in sections below.

Decentralized Policy Gradient

Given a time horizon H , we define a trajectory specifically for agent i , as $\tau^i \triangleq \{s_0, a_0^i, \dots, s_{H-1}, a_{H-1}^i\}$ under any stationary policy. By following the trajectory τ^i , a cumulative discounted reward is given as $\mathcal{R}_i(\tau^i) \triangleq \sum_{h=0}^H \gamma^h r_h^i$ such that an individual return can be obtained as:

$$J_i(\mathbf{x}^i) \triangleq \mathbb{E}_{\tau^i \sim p_i(\tau^i | \mathbf{x}^i)} [\mathcal{R}_i(\tau^i)] = \int \mathcal{R}_i(\tau^i) p_i(\tau^i | \mathbf{x}^i) d\tau^i, \quad (2)$$

where $p_i(\tau^i | \mathbf{x}^i)$ is the probability distribution over τ^i that is equivalent to the following expression given the initial distribution $\rho_0^i = \rho^i(s_0)$. Without loss of generalization, we can assume that the initial distribution is identical for all agents, namely $\rho(s_0)$. Then, we have,

$$p_i(\tau^i | \mathbf{x}^i) = \rho_0(s_0) \prod_{h=0}^{H-1} \mathcal{P}(s_{h+1} | s_h, a_h^i) \pi^i(a_h^i | s_h). \quad (3)$$

For each agent i , the goal is to find an optimal policy π_*^i to maximize the return $J_i(\mathbf{x}^i)$. As discussed above, the underlying dynamic distribution results in a non-oblivious learning problem, which is more significant in MARL. To resolve this issue, the decentralized policy gradient is a decent choice. As background knowledge of MARL, we next present how to arrive at the local stochastic policy gradient, which will help characterize the analysis for the proposed algorithms.

Computing the gradient of $J_i(\mathbf{x}^i)$ w.r.t \mathbf{x}^i yields the fol-

lowing formula:

$$\begin{aligned}\nabla J_i(\mathbf{x}^i) &= \int \mathcal{R}_i(\tau^i) \frac{\nabla p_i(\tau^i | \mathbf{x}^i)}{p_i(\tau^i | \mathbf{x}^i)} p_i(\tau^i | \mathbf{x}^i) d\tau^i \\ &= \mathbb{E}_{\tau^i \sim p_i(\tau^i | \mathbf{x}^i)} [\nabla \log p_i(\tau^i | \mathbf{x}^i) \mathcal{R}_i(\tau^i)]\end{aligned}\quad (4)$$

In practice, $p_i(\tau^i | \mathbf{x}^i)$ is typically unknown such that the accurate full policy gradient for agent i is difficult to obtain. Thus, similar to decentralized stochastic gradient descent (Jiang et al. 2017), we calculate the policy gradient by sampling a mini-batch of trajectories $\mathcal{B} = \{\tau_m^i\}_{m=1}^{|\mathcal{B}|}$ from the distribution $p_i(\tau^i | \mathbf{x}^i)$ such that

$$\hat{\nabla} J_i(\mathbf{x}^i) = \frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \nabla \log p_i(\tau_m^i | \mathbf{x}^i) \mathcal{R}_i(\tau_m^i). \quad (5)$$

In addition, combining Eq. 3, we can observe that $\nabla \log p_i(\tau_m^i | \mathbf{x}^i)$ is independent of the probability transition \mathcal{P} . Hence, Eq. 5 is written as

$$\begin{aligned}\hat{\nabla} J_i(\mathbf{x}^i) &= \frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \mathbf{g}_i(\tau_m^i | \mathbf{x}^i) \\ &= \frac{1}{|\mathcal{B}|} \sum_{m \in \mathcal{B}} \left(\sum_{h=0}^{H-1} \nabla_{\mathbf{x}^i} \log \pi^i(a_h^{i,m}, s_h^m) \right) \\ &\quad \left(\sum_{h=0}^{H-1} \gamma^h r_h^i(a_h^{i,m}, s_h^m) \right)\end{aligned}\quad (6)$$

In the above equation, $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$ is the unbiased estimate of $\nabla J_i(\mathbf{x}^i)$, i.e., $\mathbb{E}[\mathbf{g}_i(\tau^i | \mathbf{x}^i)] = \nabla J_i(\mathbf{x}^i)$. Some well-known policy gradient estimators can be obtained through Eq. 6, such as decentralized REINFORCE, which is the direct extension of its centralized version. We refer interested readers to (Huang et al. 2020) for more details.

Our Proposed Approach: MDPGT Hybrid Importance Sampling SARAH

In this subsection, we propose a hybrid importance sampling version of SARAH, termed HIS-SARAH, for decentralized policy gradient updates. First, we define the importance sampling weight (Metelli et al. 2020) as follows:

$$v(\tau | \mathbf{x}', \mathbf{x}) = \frac{p(\tau | \mathbf{x}')}{p(\tau | \mathbf{x})} = \prod_{h=0}^{H-1} \frac{\pi_{\mathbf{x}'}(a_h | s_h)}{\pi_{\mathbf{x}}(a_h | s_h)}. \quad (7)$$

As mentioned in the last section, due to the non-oblivious learning problem, $\mathbb{E}_{\tau \sim p(\tau | \mathbf{x})} [\mathbf{g}(\tau | \mathbf{x}) - \mathbf{g}(\tau | \mathbf{x}')] \neq \nabla J(\mathbf{x}) - \nabla J(\mathbf{x}')$. With Eq. 7 we have $\mathbb{E}_{\tau \sim p(\tau | \mathbf{x})} [\mathbf{g}(\tau | \mathbf{x}) - v(\tau | \mathbf{x}', \mathbf{x}) \mathbf{g}(\tau | \mathbf{x}')] = \nabla J(\mathbf{x}) - \nabla J(\mathbf{x}')$, which has been analyzed in (Huang et al. 2020) for centralized policy optimization methods and will be a key relationship in our proof. We denote by \mathbf{u}^i the stochastic policy gradient surrogate for agent i . Thus, applying Eq. 7 in a decentralized manner for Hybrid-SARAH (See Supplementary materials for definition) gives the following update law at a time step k :

$$\begin{aligned}\mathbf{u}_k^i &= \beta \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) + (1 - \beta) [\mathbf{u}_{k-1}^i + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) \\ &\quad - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)].\end{aligned}\quad (8)$$

The second term on the right hand side of Eq. 8 differ in the extra importance sampling weight compared to Eq. 13 in the supplementary materials. Intuitively, $v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i)$ resolves the non-stationarity in the MARL and retains the regular variance reduction property of HIS-SARAH as applied in supervised learning problems. Clearly, each \mathbf{u}_k^i is a *conditionally* biased estimator $\nabla J_i(\mathbf{x}_k^i)$, i.e., $\mathbb{E}[\mathbf{u}_k^i] \neq \nabla J_i(\mathbf{x}_k^i)$ typically. Nevertheless, it can be shown that $\mathbb{E}[\mathbf{u}_k^i] = \mathbb{E}[\nabla J_i(\mathbf{x}_k^i)]$, which implies that \mathbf{u}_k^i acts as a *surrogate* for the underlying exact full policy gradient. Therefore, \mathbf{u}_k^i will be called directly the stochastic policy gradient surrogate for the rest of analysis. With Eq. 8 in hand, we now are ready to present the algorithmic framework in the following.

Algorithmic Framework

We first present MDPGT (in Algorithm 1), which only takes a trajectory to initialize the policy gradient surrogate, leading to significant randomness due to the conditionally biased estimator property at the starting point of optimization, but still retaining the same sampling complexity as compared to the SOTA of MARL. To have a better initialization, we also present another way of initialization by sampling a mini-batch of trajectories from the distribution (in blue in Algorithm 1). Surprisingly, we will see that with a proper size of mini-batch initialization, the sampling complexity of our proposed approach complies with the best result in centralized RL, which improves the SOTA of MARL.

Algorithm 1: MDPGT

Result: $\tilde{\mathbf{x}}_K$ chosen uniformly random from $\{\mathbf{x}_k^i, i \in \mathcal{V}\}_{k=1}^K$
Input: $\mathbf{x}_0^i = \tilde{\mathbf{x}}_0 \in \mathbb{R}^d, \eta \in \mathbb{R}^+, \beta \in (0, 1), \mathbf{W} \in \mathbb{R}^{N \times N}, \mathbf{v}_0^i = \mathbf{0}_d, \mathbf{u}_{-1}^i = \mathbf{0}_d, K, \mathcal{B} \in \mathbb{Z}^+, k = 1$
 Initialize the local policy gradient surrogate by sampling a trajectory τ_0^i from $p_i(\tau^i | \mathbf{x}_0^i) : \mathbf{u}_0^i = \mathbf{g}_i(\tau_0^i | \mathbf{x}_0^i)$, or by sampling a *mini-batch* of trajectories $\{\tau_0^{i,m}\}_{m=1}^{|\mathcal{B}|}$ from $p_i(\tau^i | \mathbf{x}_0^i) : \mathbf{u}_0^i = \frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} \mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i)$
 Initialize the local policy gradient tracker: $\mathbf{v}_1^i = \sum_{j \in Nb(i)} \omega_{ij} \mathbf{v}_0^j + \mathbf{u}_0^i - \mathbf{u}_{-1}^i$
 Initialize the local estimate of the policy network parameter: $\mathbf{x}_1^i = \sum_{j \in Nb(i)} \omega_{ij}^i (\mathbf{x}_0^j + \eta \mathbf{v}_1^j)$
while $k < K$ **do**
for each agent **do**
 | Sample a trajectory τ_k^i from $p_i(\tau^i | \mathbf{x}_k^i)$ and compute the local policy gradient surrogate using Eq. 8
 | Update the local policy gradient tracker $\mathbf{v}_{k+1}^i = \sum_{j \in Nb(i)} \omega_{ij} \mathbf{v}_k^j + \mathbf{u}_k^i - \mathbf{u}_{k-1}^i$
 | Update the local estimate of the policy network parameters $\mathbf{x}_{k+1}^i = \sum_{j \in Nb(i)} \omega_{ij} (\mathbf{x}_k^j + \eta \mathbf{v}_{k+1}^j)$
end
 $k = k + 1$
end

A brief outline of Algorithm 1 is as follows. The initialization of the policy gradient surrogate \mathbf{u}_0^i can either be based on only a trajectory sampled from $p_i(\tau^i | \mathbf{x}_0^i)$ or a mini-batch.

Subsequently, the policy gradient tracker and network parameters are initialized based on \mathbf{u}_0^i . The core part of the algorithm consists of each individual update for \mathbf{u}_k^i , \mathbf{v}_k^i , and \mathbf{x}_k^i . By controlling the value of β in Eq. 8, MDPGT can degenerate to either vanilla decentralized policy gradient (with $\beta = 1$) or decentralized version of SRVR-PG (Xu, Gao, and Gu 2019) (with $\beta = 0$), both with the gradient tracking step. In our work, to emphasize the impact of the trade-off on the policy gradient surrogate, we keep $\beta \in (0, 1)$, which makes β act more closely as the momentum coefficient in accelerated SGD algorithms (Singh et al. 2020).

We emphasize that we are unaware of theoretical results for decentralized SRVR-PG. Hence, the proof techniques presented in this paper can also apply to this case. Another implication from Algorithm 1 is that at the beginning of each time step k , only one trajectory is required for computing the policy gradient, allowing for the batch size to be *independent* of ϵ , i.e., $\mathcal{O}(1)$, where we omit the number of agents N when considering the whole networked system.

Theoretical Convergence

In this section, we present an analysis of MDPGT. Most of the assumptions below are mild, and standard in the decentralized optimization and RL literature. Due to space limitations, we defer auxiliary lemmas and proofs to the supplementary materials.

Assumption 1. *Gradient and Hessian matrix of function $\log\pi^i(a^i|s)$ are bounded, i.e., there exist constants $C_g, C_h > 0$ such that $\|\nabla\log\pi^i(a^i|s)\| \leq C_g$ and $\|\nabla^2\log\pi^i(a^i|s)\| \leq C_h$, for all $i \in \mathcal{V}$.*

Note that we skip the subscript \mathbf{x}^i at π^i for the notation simplicity. In this context, we did not impose the bounded policy gradient assumption, though it can be derived based on the above assumption, which has been adopted in centralized RL algorithms (Zhang et al. 2021b; Huang et al. 2020; Xu, Gao, and Gu 2019). Additionally, it also helps derive the smoothness of $J_i(\mathbf{x}^i)$ that has typically been exposed as an assumption in decentralized learning/optimization literature.

Assumption 2. *The mixing matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is doubly stochastic such that $\lambda \triangleq \|\mathbf{W} - \mathbf{P}\| \in [0, 1)$, where λ signifies the second largest eigenvalue to measure the algebraic connectivity of the graph, and $\mathbf{P} = \frac{1}{N}\mathbf{1}^\top\mathbf{1}$ and $\mathbf{1}$ is a column vector with each entry being 1.*

Assumption 3. *Variance of importance sampling weight $v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)$ is bounded, i.e., there exists a constant $\mathcal{M} > 0$ such that $\mathbb{V}(v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)) \leq \mathcal{M}$, for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^{d_i}$ and $\tau^i \sim p_i(\tau^i|\mathbf{x}_1)$, for all $i \in \mathcal{V}$.*

Assumption 2 is generic in most previous works on decentralized optimization, though such a property has been relaxed in some works (Nedić and Olshevsky 2014). However, we have not been aware of any existing works in MARL doing such a relaxation and its investigation can be of independent interest. Assumption 3 is specifically introduced for importance sampling-based methods. Such an assumption is critical to construct the relationship between $\mathbb{V}(v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2))$ and $\|\mathbf{x}_1 - \mathbf{x}_2\|^2$, through which the impact of the variance of importance sampling on the convergence

can be explicitly quantified. Another typical assumption is for the bounded variance of stochastic gradient such that $\mathbb{E}[\|\mathbf{g}_i(\tau^i|\mathbf{x}^i) - \nabla J_i(\mathbf{x}^i)\|^2] \leq \sigma_i^2$. However, under MARL setting, such a result can be derived from Assumption 1 and we present the formal result in Lemma 1. In this context, we also have $\bar{\sigma}^2 = \frac{1}{N} \sum_{i=1}^N \sigma_i^2$, for all $i \in \mathcal{V}$. The explicit expression of $\bar{\sigma}^2$ is given in the supplementary materials.

Main Results

We present the main results to show specifically the convergence rates for MDPGT when it is initialized by a mini-batch of trajectories. We denote by $L > 0$ the *smoothness constant* and $G > 0$ the *upper bound* of $\|\mathbf{g}_i(\tau^i|\mathbf{x}^i)\|$ for all $i \in \mathcal{V}$. We further define a constant $C_v > 0$ such that $\mathbb{V}(v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)) \leq C_v^2 \|\mathbf{x}_1 - \mathbf{x}_2\|^2$. The explicit expressions of these constants are derived in lemmas in the supplementary materials. Note that in our work, G is not directly assumed, but instead derived based on Assumption 1.

Theorem 1. *Let Assumptions 1, 2 and 3 hold. Let the momentum coefficient $\beta = \frac{96L^2 + 96G^2C_v}{N}\eta^2$. If MDPGT is initialized by a mini-batch of trajectories with the size being \mathcal{B} and the step size satisfies the following condition*

$$0 < \eta \leq \min \left\{ \frac{(1 - \lambda^2)^2}{\lambda \sqrt{12844L^2 + 9792G^2C_v}}, \frac{\sqrt{N(1 - \lambda^2)}\lambda}{31\sqrt{L^2 + G^2C_v}}, \frac{1}{6\sqrt{6(L^2 + G^2C_v)}} \right\},$$

then the output $\tilde{\mathbf{x}}_K$ satisfies: for all $K \geq 2$:

$$\begin{aligned} \mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] &\leq \frac{4(J^* - J(\tilde{\mathbf{x}}_0))}{\eta K} + \frac{4\bar{\sigma}^2}{N|\mathcal{B}|\beta K} + \frac{8\beta\bar{\sigma}^2}{N} \\ &+ \frac{34\lambda^2}{KN(1 - \lambda^2)^3} \|\nabla \mathbf{J}(\tilde{\mathbf{x}}_0)\|^2 + \frac{68\lambda^2\bar{\sigma}^2}{(1 - \lambda^2)^3|\mathcal{B}|K} + \frac{204\lambda^2\beta^2\bar{\sigma}^2}{(1 - \lambda^2)^3}, \end{aligned} \quad (9)$$

where J^* is the upper bound of $J(\mathbf{x})$ and $\|\nabla \mathbf{J}(\tilde{\mathbf{x}}_0)\|^2 \triangleq \sum_{i=1}^N \|\nabla J_i(\tilde{\mathbf{x}}_0)\|^2$.

Theorem 1 depicts that when $K \rightarrow \infty$, MDPGT enables convergence to a *steady-state error* in a sublinear rate $\mathcal{O}(1/K)$ if η and β are selected properly, i.e.,

$$\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] \leq \frac{8\beta\bar{\sigma}^2}{N} + \frac{204\lambda^2\beta^2\bar{\sigma}^2}{(1 - \lambda^2)^3}. \quad (10)$$

By observing Eq. 10, the steady-state error is determined by the number of agents, the variance of stochastic policy gradient, and the spectral gap of the graph $1 - \lambda$. Increasing the number of agents leads to a small error bound. Though different network topologies imply different error bounds, the higher order term of β can reduce the impact of the spectral gap on the error bound. Another suggestion from Eq. 10 is that η and β can be reduced to make the steady-state error arbitrarily small, while in return this can affect the speed of convergence. Surprisingly, even though we have to adopt the bounded stochastic policy gradient derived from Assumption 1 for analysis, the error bound in Eq. 10 only

depends heavily on the variance, which is inherently consistent with most conclusions from decentralized optimization in literature without the bounded stochastic gradient assumption. While J^* is essentially correlated to the upper bound of reward R , in this context, we still adopt the implicit J^* for convenience. We next provide the analysis for the non-asymptotic behavior, defining appropriately η, β , and $|\mathcal{B}|$.

Corollary 1. *Let $\eta = \frac{N^{2/3}}{8LK^{1/3}}, \beta = \frac{DN^{1/3}}{64L^2K^{2/3}}, |\mathcal{B}| = \lceil \frac{K^{1/3}}{N^{2/3}} \rceil$ in Theorem 1. We have,*

$$\begin{aligned} \mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] &\leq \underbrace{\frac{256L^3D(J^* - J(\tilde{\mathbf{x}}_0)) + 2048L^4\bar{\sigma}^2 + D^2\bar{\sigma}^2}{8L^2D(NK)^{2/3}}}_{T_1} \\ &+ \underbrace{\frac{34\lambda^2}{KN(1-\lambda^2)^3}\|\nabla J(\tilde{\mathbf{x}}_0)\|^2 + \frac{\lambda^2\bar{\sigma}^2(51D^2 + 69632N^{2/3}L^4)}{1024(1-\lambda^2)^3K^{4/3}L^4}}_{T_2}, \end{aligned} \quad (11)$$

for all

$$K \geq \max \left\{ \frac{N^2D^{1.5}}{512L^3}, \frac{29791\sqrt{N}(L^2 + G^2C_v)^{1.5}}{512L^3\lambda^3(1-\lambda^2)^{1.5}}, \frac{(12844L^2 + 9792G^2C_v)^{1.5}N^2\lambda^3}{512L^3(1-\lambda^2)^6} \right\},$$

where $D = 96L^2 + 96G^2C_v$.

Remark 1. *An implication from Corollary 1 is that at the early stage of optimization, before T_1 in Eq. 11 dominates, the complexity is tightly related to the algebraic connectivity of the network topology in T_2 , which is measured by the spectral gap $1 - \lambda$. However, T_2 is in a large order of $1/K$. As the optimization moves towards the latter stage where T_1 dominates, the overall complexity is independent of the network topology.*

For the ease of exposition, with Corollary 1, when K is sufficiently large, it is an immediate consequence as $\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] \leq \mathcal{O}((NK)^{-2/3})$. Thus, for achieving $\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|] \leq \epsilon$, the following relationship is obtained:

$$\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|] = \sqrt{(\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|])^2} \leq \sqrt{\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2]} \leq \epsilon.$$

Combining the last two inequalities results in the ultimate sampling complexity, i.e., $\mathcal{O}(N^{-1}\epsilon^{-3})$, which exhibits *linear speed up*. More importantly, this is N times smaller than the sampling complexity of the centralized momentum-based policy gradient methods (Huang et al. 2020) that performs on a single node. However, we have known from Corollary 1 that typically K has to be large, which will in the linear speed up regime reduce η and β .

We also investigate a worse initialization with only a single trajectory sampled from $p_i(\tau^i|\mathbf{x}_0^i)$. However, without a mini-batch initialization, the eventual sampling complexity is $\mathcal{O}(N^{-1}\epsilon^{-4})$ (see Theorem 2 and Corollary 3). Though variance reduction techniques have not reduced the order of ϵ^{-1} , compared to the SOTA approaches, the linear speed up still enables the complexity to be N times smaller than that in (Xu, Gao, and Gu 2019; Huang et al. 2020). Additionally,

different from traditional decentralized learning problems, MARL has more significant variances in the optimization procedure due to the non-oblivious characteristic. Using just a single trajectory for each agent to initialize is can be a poor scheme, but the adopted variance reduction techniques can successfully maintain the SOTA sampling complexity in a decentralized setting. Please refer to the supplementary materials for formal results and proof.

Implication for Gaussian Policy. We study the sample complexity when the policy function $\pi^i(a^i|s)$ of each agent is explicitly a Gaussian distribution. For a bounded action space $\mathcal{A}^i \subset \mathbb{R}$, a Gaussian policy parameterized by \mathbf{x}_i is defined as

$$\pi^i(a^i|s) = \frac{1}{\sqrt{2\pi}} \exp \left(- \frac{((\mathbf{x}^i)^\top \phi_i(s) - a^i)^2}{2\xi^2} \right), \quad (12)$$

where ξ^2 is a constant standard deviation parameter and $\phi_i(s) : \mathcal{S} \rightarrow \mathbb{R}^{d_i}$ is a mapping from the state space to the feature space. Thus, the following formal result can be obtained. A more formal analysis and proof can be seen in the supplementary materials.

Corollary 2. *Let $\pi^i(a^i|s)$ be defined as a Gaussian distribution in Eq. 12 with $|a^i| \leq C_a$, where $C_a, C_f > 0$, and $\|\phi_i(s)\| \leq C_f$, and $\eta, \beta, |\mathcal{B}|$ be defined as in Corollary 1. The sampling complexity of attaining the accuracy $\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|] \leq \epsilon$ is $\mathcal{O}((1-\gamma)^{-4.5}N^{-1}\epsilon^{-3})$.*

Numerical Experiments and Results

To validate our proposed algorithm, we performed experiments on a multi-agent benchmark environment with a cooperative navigation task that has been commonly used as a benchmark in several previous works (Qu et al. 2019; Zhang et al. 2018; Lu et al. 2021). Our platform for cooperative navigation is derived off the particle environment introduced by (Lowe et al. 2017). In our modification, all agents are initialized at random locations with a specific goal in a 2-dimensional grid world. Each agent observes the combined position and velocity of itself and all other agents. The agent is capable of moving up, down, left or right with the objective of navigating to their respective goals. The reward function of each agent is defined as the negative euclidean distance of the agent to the goal. Additionally, a penalty of -1 is imposed whenever the agent collides with any other agents. All agent's policy is represented by a 3-layer dense neural network with 64 hidden units with *tanh* activation functions. The agents were trained for 50,000 episodes with a horizon of 50 steps and discount factor of 0.99. For the sake of brevity, we present numerical results in only one environment setting with five agents. Additional results with different number of agents and a simplified environment and computing infrastructure details are available in the supplementary materials¹.

Efficacy of MDPGT

Figure 1 illustrates the average training rewards obtained by the five agents in the cooperative navigation gridworld

¹Codes to reproduce these results are available at the following repository: <https://github.com/xylee95/MD-PGT>

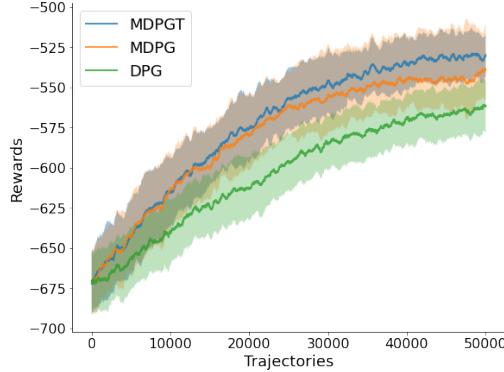


Figure 1: Average rewards obtain by MDPGT, MDPG and DPG in a cooperative navigation task for five agents. For MDPGT, $\beta = 0.5$ and results shown are averaged across five random seeds. The line plots denote the mean value and shaded regions denote the standard deviation of rewards.

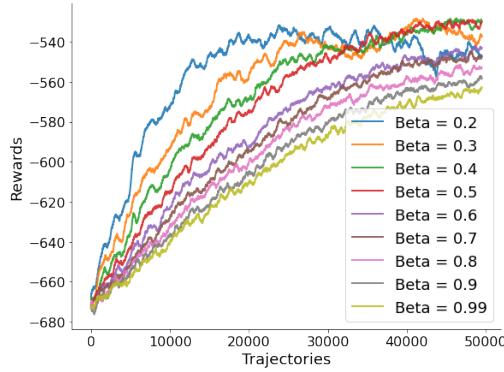


Figure 2: Ablation study illustrating the effect of various momentum coefficients, β on the performance of MDPGT for five agents in the cooperative navigation environment.

environment. As observed, both MDPG and MDPGT significantly outperforms the baseline, denoted as DPG. Comparing MDPG with MDPGT, we observe that while both algorithms initially have similar performance, MDPGT begins to outperform MDPG around the 15,000 iteration. Additionally, when we compare the standard deviation of the rewards, shown as shaded regions, we observe that standard deviation of MDPGT is also smaller than the standard deviation of MDPG. In summary, these results validate our theoretical findings that utilizing gradient tracking as bias correction technique does improve the performance of the algorithm. Additionally, the improvement in terms of sampling complexity over DPG is empirically evident through the result.

Effect of Momentum Coefficient

Next, we perform an additional ablation study to investigate the effect of the momentum coefficient β on the performance of MDPGT. As shown in Figure 2, we see that the choice of momentum coefficient does indeed have an effect on the performance. A β that is low can induce a faster convergence

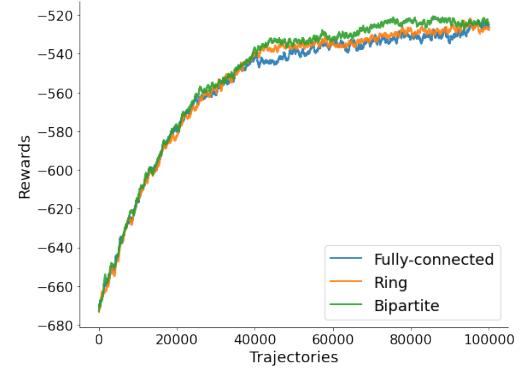


Figure 3: Experiment results for five agents in the cooperative navigation environment to compare the effects of different network topologies. $\beta = 0.5$ for all experiments shown.

rate, but at the cost of a higher fluctuations in rewards, as seen by $\beta = 0.2$ and 0.3. Conversely, a β value that is too high will cause the surrogate to degenerate into vanilla policy gradients and reflects a similar performance as DPG, which matches the implication by Eq. 10. Ultimately, we see that for this environment, $\beta = 0.4$ and 0.5 offers the perfect balance between convergence rate and stability/variance of the training. Hence, β can be viewed as hyper-parameter which can be tuned to trade off between optimizing for convergence versus training stability.

Effect of Different Topologies

Finally, we provide evidence which confirms the fact that our proposed method holds under various networks topologies. To test our hypothesis, we train five agents in the same cooperative navigation environment using three different network topologies: fully-connected, ring and bi-partite topology. As seen in Figure 3, the five agents achieves similar rewards despite communicating via different network topologies. This validates our claim in Remark 1.

Conclusions

This paper proposes a novel MARL algorithm that involves variance reduction techniques to reduce the sampling complexity of decentralized policy-based methods. Specifically we have developed the algorithmic framework and analyzed it in a principled manner. An importance sampling-based stochastic recursive momentum is presented as the policy gradient surrogate, which is taken as input to a policy gradient tracker. Through theoretical analysis, we have found that the proposed method can improve the sampling efficiency in the decentralized RL settings compared to the SOTA methods. To the best of our knowledge, this is the first time to achieve the best available rate, $\mathcal{O}(\epsilon^{-3})$, for generic (possibly non-concave) performance functions. Empirical results have shown the superiority of the proposed MDPGT over the baseline decentralized policy gradient methods. Future research directions include: 1) incorporating more complex decentralized environments in our experiments to reveal potentially novel and interesting results; 2) extending the proposed

method to model-based decentralized RL settings to improve further the sampling efficiency; 3) theoretically analyzing the robustness of the proposed method under adversarial attacks.

References

Ackermann, J.; Gabler, V.; Osa, T.; and Sugiyama, M. 2019. Reducing overestimation bias in multi-agent domains using double centralized critics. *arXiv preprint arXiv:1910.01465*.

Arslan, G.; and Yüksel, S. 2016. Decentralized Q-learning for stochastic teams and games. *IEEE Transactions on Automatic Control*, 62(4): 1545–1558.

Bono, G.; Dibangoye, J. S.; Matignon, L.; Pereyron, F.; and Simonin, O. 2018. Cooperative multi-agent policy gradient. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 459–476. Springer.

Chen, B.; Xu, M.; Liu, Z.; Li, L.; and Zhao, D. 2020. Delay-Aware Multi-Agent Reinforcement Learning for Cooperative and Competitive Environments. *arXiv e-prints*, arXiv–2005.

Chu, T.; Chinchali, S.; and Katti, S. 2020. Multi-agent reinforcement learning for networked system control. *arXiv preprint arXiv:2004.01339*.

Cortes, C.; Mansour, Y.; and Mohri, M. 2010. Learning Bounds for Importance Weighting. In *Nips*, volume 10, 442–450. Citeseer.

Cutkosky, A.; and Orabona, F. 2019. Momentum-based variance reduction in non-convex sgd. *arXiv preprint arXiv:1905.10018*.

Das, R.; Acharya, A.; Hashemi, A.; Sanghavi, S.; Dhillon, I. S.; and Topcu, U. 2020. Faster Non-Convex Federated Learning via Global and Local Momentum. *arXiv preprint arXiv:2012.04061*.

Fang, C.; Li, C. J.; Lin, Z.; and Zhang, T. 2018. Spider: Near-optimal non-convex optimization via stochastic path integrated differential estimator. *arXiv preprint arXiv:1807.01695*.

Helou, R. E.; Kalathil, D.; and Xie, L. 2020. Fully Decentralized Reinforcement Learning-based Control of Photovoltaics in Distribution Grids for Joint Provision of Real and Reactive Power. *arXiv preprint arXiv:2008.01231*.

Huang, F.; Gao, S.; Pei, J.; and Huang, H. 2020. Momentum-based policy gradient methods. In *International Conference on Machine Learning*, 4422–4433. PMLR.

Jiang, Z.; Balu, A.; Hegde, C.; and Sarkar, S. 2017. Collaborative deep learning in fixed topology networks. *arXiv preprint arXiv:1706.07880*.

Johnson, R.; and Zhang, T. 2013. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26: 315–323.

Karimireddy, S. P.; Jaggi, M.; Kale, S.; Mohri, M.; Reddi, S. J.; Stich, S. U.; and Suresh, A. T. 2020. Mime: Mimicking centralized stochastic algorithms in federated learning. *arXiv preprint arXiv:2008.03606*.

Li, M.-L.; Chen, S.; and Chen, J. 2020. Adaptive learning: A new decentralized reinforcement learning approach for cooperative multiagent systems. *IEEE Access*, 8: 99404–99421.

Li, W.; Jin, B.; Wang, X.; Yan, J.; and Zha, H. 2020. F2A2: Flexible Fully-decentralized Approximate Actor-critic for Cooperative Multi-agent Reinforcement Learning. *arXiv preprint arXiv:2004.11145*.

Lin, Q.; and Ling, Q. 2021. Decentralized TD (0) with Gradient Tracking. *IEEE Signal Processing Letters*.

Lin, Y.; Zhang, K.; Yang, Z.; Wang, Z.; Başar, T.; Sandhu, R.; and Liu, J. 2019. A communication-efficient multi-agent actor-critic algorithm for distributed reinforcement learning. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 5562–5567. IEEE.

Lowe, R.; Wu, Y.; Tamar, A.; Harb, J.; Abbeel, P.; and Mordatch, I. 2017. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. *Neural Information Processing Systems (NIPS)*.

Lu, S.; Zhang, K.; Chen, T.; Basar, T.; and Horesh, L. 2021. Decentralized Policy Gradient Descent Ascent for Safe Multi-Agent Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 8767–8775.

Metelli, A. M.; Papini, M.; Montali, N.; and Restelli, M. 2020. Importance Sampling Techniques for Policy Optimization. *Journal of Machine Learning Research*, 21(141): 1–75.

Mukherjee, S.; Bai, H.; and Chakrabortty, A. 2020. Model-free decentralized reinforcement learning control of distributed energy resources. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, 1–5. IEEE.

Nedić, A.; and Olshevsky, A. 2014. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3): 601–615.

Nguyen, D. T.; Yeoh, W.; Lau, H. C.; Zilberstein, S.; and Zhang, C. 2014. Decentralized multi-agent reinforcement learning in average-reward dynamic DCOPs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 28.

Nguyen, L. M.; Liu, J.; Scheinberg, K.; and Takáč, M. 2017. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, 2613–2621. PMLR.

Nguyen, T.; and Mukhopadhyay, S. 2017. Selectively decentralized Q-learning. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 328–333. IEEE.

Nguyen, T. T.; Nguyen, N. D.; and Nahavandi, S. 2020. Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE transactions on cybernetics*, 50(9): 3826–3839.

Papini, M.; Binaghi, D.; Canonaco, G.; Pirotta, M.; and Restelli, M. 2018. Stochastic variance-reduced policy gradient. In *International Conference on Machine Learning*, 4026–4035. PMLR.

Pirotta, M.; Restelli, M.; and Bascetta, L. 2013. Adaptive step-size for policy gradient methods. *Advances in Neural Information Processing Systems*, 26: 1394–1402.

Pu, S.; and Nedić, A. 2020. Distributed stochastic gradient tracking methods. *Mathematical Programming*, 1–49.

Qu, C.; Mannor, S.; Xu, H.; Qi, Y.; Song, L.; and Xiong, J. 2019. Value propagation for decentralized networked

deep multi-agent reinforcement learning. *arXiv preprint arXiv:1901.09326*.

Ryu, H.; Shin, H.; and Park, J. 2020. Multi-agent actor-critic with hierarchical graph attention network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7236–7243.

Simões, D.; Lau, N.; and Reis, L. P. 2020. Multi-agent actor centralized-critic with communication. *Neurocomputing*, 390: 40–56.

Singh, N.; Data, D.; George, J.; and Diggavi, S. 2020. SQuARM-SGD: Communication-Efficient Momentum SGD for Decentralized Optimization. *arXiv preprint arXiv:2005.07041*.

Sun, Y.; Daneshmand, A.; and Scutari, G. 2019. Convergence rate of distributed optimization algorithms based on gradient tracking. *arXiv preprint arXiv:1905.02637*.

Suttle, W.; Yang, Z.; Zhang, K.; Wang, Z.; Başar, T.; and Liu, J. 2020. A multi-agent off-policy actor-critic algorithm for distributed reinforcement learning. *IFAC-PapersOnLine*, 53(2): 1549–1554.

Sutton, R. S.; McAllester, D. A.; Singh, S. P.; Mansour, Y.; et al. 1999. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, volume 99, 1057–1063. Citeseer.

Tran-Dinh, Q.; Pham, N. H.; Phan, D. T.; and Nguyen, L. M. 2019. Hybrid stochastic gradient descent algorithms for stochastic nonconvex optimization. *arXiv preprint arXiv:1905.05920*.

Wang, X.; Wang, C.; Li, X.; Leung, V. C.; and Taleb, T. 2020. Federated deep reinforcement learning for internet of things with decentralized cooperative edge caching. *IEEE Internet of Things Journal*, 7(10): 9441–9455.

Wei, C.-Y.; Lee, C.-W.; Zhang, M.; and Luo, H. 2021. Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games. *arXiv preprint arXiv:2102.04540*.

Xin, R.; Khan, U. A.; and Kar, S. 2020. A near-optimal stochastic gradient method for decentralized non-convex finite-sum optimization. *arXiv preprint arXiv:2008.07428*.

Xin, R.; Khan, U. A.; and Kar, S. 2021. A hybrid variance-reduced method for decentralized stochastic non-convex optimization. *arXiv preprint arXiv:2102.06752*.

Xu, P.; Gao, F.; and Gu, Q. 2019. Sample efficient policy gradient methods with recursive variance reduction. *arXiv preprint arXiv:1909.08610*.

Zeng, S.; Anwar, A.; Doan, T.; Romberg, J.; and Raychowdhury, A. 2020. A decentralized policy gradient approach to multi-task reinforcement learning. *arXiv preprint arXiv:2006.04338*.

Zhang, J.; Bedi, A. S.; Wang, M.; and Koppel, A. 2021a. MARL with General Utilities via Decentralized Shadow Reward Actor-Critic. *arXiv preprint arXiv:2106.00543*.

Zhang, J.; Ni, C.; Yu, Z.; Szepesvari, C.; and Wang, M. 2021b. On the Convergence and Sample Efficiency of Variance-Reduced Policy Gradient Method. *arXiv preprint arXiv:2102.08607*.

Zhang, K.; Yang, Z.; and Başar, T. 2018. Networked multi-agent reinforcement learning in continuous spaces. In *2018 IEEE Conference on Decision and Control (CDC)*, 2771–2776. IEEE.

Zhang, K.; Yang, Z.; and Başar, T. 2019. Decentralized Multi-Agent Reinforcement Learning with Networked Agents: Recent Advances. *arXiv preprint arXiv:1912.03821*.

Zhang, K.; Yang, Z.; Liu, H.; Zhang, T.; and Başar, T. 2018. Fully decentralized multi-agent reinforcement learning with networked agents. In *International Conference on Machine Learning*, 5872–5881. PMLR.

Zhang, K.; Yang, Z.; Liu, H.; Zhang, T.; and Başar, T. 2021c. Finite-sample analysis for decentralized batch multi-agent reinforcement learning with networked agents. *IEEE Transactions on Automatic Control*.

Zhou, P.; Chen, X.; Liu, Z.; Braud, T.; Hui, P.; and Kangasharju, J. 2020. DRLE: Decentralized Reinforcement Learning at the Edge for Traffic Light Control. *arXiv preprint arXiv:2009.01502*.

Supplementary Materials

We present the additional analysis and detailed proof for the auxiliary lemmas and main theorems as well as additional empirical results.

Additional Related Works

The most recent work (Zhang et al. 2021b) in centralized RL introduced SIVR-PG, which employed the hidden convex nature of general utility function and leveraged techniques from composition optimization to attain the sampling complexity of $\tilde{\mathcal{O}}(\epsilon^{-2})$. Nevertheless, SIVR-PG requires an additional assumption that the unnormalized state-action occupancy measure in the utility function is convex and its corresponding inverse mapping exists, which may not rigorously hold for some RL settings. A concurrent work (Zhang et al. 2021a) recently that has been published on arXiv proposes a Decentralized Shadow Reward Actor Critic (DSAC) for general utilities, which can be treated as the extension to the decentralized setting of SIVR-PG. In their work, the authors develop the shadow reward that estimates the derivative of the local utility with respect to their occupancy measure and show that DSAC converges to ϵ -stationarity in $\mathcal{O}(\epsilon^{-2.5})$ or even faster $\mathcal{O}(\epsilon^{-2})$ with high probability. They further establish the global optimality by adopting the diminishing step size. Though DSAC reveals faster convergence rate, additional occupancy measure needs to be incorporated and updated in the algorithmic framework. Also, the shadow reward they have defined is essentially a derivative, which may vary significantly among different agents. This could be another source of variance during updates. In the analysis, we have observed that more assumptions are required to arrive at the theoretical claims, including strongly convex critic objective function. Different from most previous works, the constructed error bounds in their work are not based on expectation. The parameters, including step size, batch size, trajectory length, and even final iteration need to be carefully defined. In the empirical evaluation, the authors present a cooperative navigation with safety constraint and do not show any comparison with the decentralized baseline methods. We therefore are not going to compare our proposed method with theirs due to different emphasis.

Hybrid-SARAH and GT

Hybrid-SARAH. In this context, we denote by $\mathbf{g}(\zeta, \mathbf{x})$ the stochastic gradient, where ζ is the random seed. As discussed above, Hybrid-SARAH is the combination between vanilla stochastic gradient and SARAH such that the following expression can be obtained for a time step k :

$$\mathbf{u}_k = \beta \mathbf{g}(\zeta_k, \mathbf{x}_k) + (1 - \beta)(\mathbf{u}_{k-1} + \mathbf{g}(\zeta_k, \mathbf{x}_k) - \mathbf{g}(\zeta_k, \mathbf{x}_{k-1})), \quad (13)$$

where $\beta \in [0, 1]$ is the momentum parameter. One can observe from Eq. 13 that in the part of SARAH, the term $\mathbf{g}(\zeta_k, \mathbf{x}_{k-1})$ requires the computation of $\mathbf{g}(\zeta_k, \mathbf{x}_{k-1})$ and the copy of last iteration of parameters \mathbf{x}_{k-1} , which could increase the computational power and memory significantly if the dimension of \mathbf{x} is quite high, though this may not be an issue for a modern computer. Additionally, this form of update in SARAH enables to improve the convergence rates in different problems such as decentralized learning (Xin, Khan, and Kar 2020) and reinforcement learning (Huang et al. 2020; Xu, Gao, and Gu 2019). Therefore we will still adapt Hybrid-SARAH in our MARL algorithms.

GT. Another technique that has specifically been developed for decentralized optimization is GT, which is close to SARAH, but focusing on tracking and correcting each agent's locally aggregated gradients. However, in terms of formulation, GT typically appears in a form with a consensus or mixing step that can be expressed as follows:

$$\mathbf{v}_{k+1}^i = \sum_{j \in Nb(i)} \omega_{ij} \mathbf{v}_k^j + \mathbf{g}_i(\zeta_k^i, \mathbf{x}_k^i) - \mathbf{g}_i(\zeta_{k-1}^i, \mathbf{x}_{k-1}^i), \quad (14)$$

where $\omega_{ij} \in \mathbf{W} \in \mathbb{R}^{N \times N}$ is the probability of an edge in the communication network \mathcal{G} . \mathbf{W} is the mixing matrix to define the topology of the communication network and one assumption will be imposed for the matrix. An immediate observation from Eq. 14 is that the update requires the copy of $\mathbf{g}_i(\zeta_{k-1}^i, \mathbf{x}_{k-1}^i)$ from last iteration without the extra computation as in SARAH. The difference between GT and SARAH is not the focus of this work and instead we will combine both Hybrid-SARAH and GT together to reduce the stochastic gradient variance and correct the stochastic gradient bias in the MARL.

Additional Analysis for Algorithm 1

We discuss briefly how MDPGT enables to reduce the policy gradient variance. Let $\Delta_k^i \triangleq \mathbf{u}_k^i - \nabla J_i(\mathbf{x}_k^i)$. With simple mathematical manipulations, we can obtain

$$\begin{aligned} \mathbb{E}[\Delta_k^i] &= \mathbb{E}[(1 - \beta)\Delta_{k-1}^i + \beta(\underbrace{\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)}_{T_3}) + \\ &\quad (1 - \beta)(\underbrace{\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i) + \nabla J_i(\mathbf{x}_{k-1}^i)}_{T_4})]. \end{aligned} \quad (15)$$

We can observe from the last equation that $\mathbb{E}[\Delta_k^i] = (1 - \beta)\mathbb{E}[\Delta_{k-1}^i]$ due to $\mathbb{E}_{\tau_k^i \sim p_i(\tau^i | \mathbf{x}_k^i)}[T_3] = 0$ and $\mathbb{E}_{\tau_k^i \sim p_i(\tau^i | \mathbf{x}_k^i)}[T_4] = 0$. Applying Cauthy-Schwartz inequality, we further attain the upper bound for $\mathbb{E}[\|\Delta_k^i\|^2]$ as follows,

$$\begin{aligned}\mathbb{E}[\|\Delta_k^i\|^2] &\leq (1 - \beta)^2 \mathbb{E}[\|\Delta_{k-1}^i\|^2] + 2\beta^2 \mathbb{E}[\|T_3\|^2] \\ &\quad + 2(1 - \beta)^2 \mathbb{E}[\|T_4\|^2].\end{aligned}\tag{16}$$

In the sequel, we will establish the relationship between the importance sampling weight $v_i(\tau_k^i | \mathbf{x}_k^i, \mathbf{x}_{k-1}^i)$ and $\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2$ such that the following equality can be obtained

$$\begin{aligned}\mathcal{O}(\|T_4\|^2) &= \mathcal{O}(\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2) \\ &= \mathcal{O}(\|\mathbf{x}_k^i + \bar{\mathbf{x}}_k - \bar{\mathbf{x}}_k + \bar{\mathbf{x}}_{k-1} - \bar{\mathbf{x}}_{k-1} - \mathbf{x}_{k-1}^i\|^2).\end{aligned}\tag{17}$$

Eq. 17 is bounded above by the following

$$\mathcal{O}(3\|\mathbf{x}_k^i - \bar{\mathbf{x}}_k\|^2 + 3\|\mathbf{x}_{k-1}^i - \bar{\mathbf{x}}_{k-1}\|^2 + 3\eta^2\|\bar{\mathbf{u}}_{k-1}\|^2),$$

where $\bar{\mathbf{x}}$ is the ensemble average of all agents $i \in \mathcal{V}$ and the first two terms signify the consensus estimate, which can be controlled by β and η . The bound also follows from the fact that $\bar{\mathbf{u}}_{k-1} = \bar{\mathbf{v}}_k$ (see Lemma 4). Therefore, in light of Eq. 16, setting β and η properly is able to reduce the variance. This also motivates the adoption of policy gradient surrogate in the work.

Auxiliary Lemmas and Their Proof

We start with the results for the smoothness constant L and the upper bound of $\|\mathbf{g}_i(\tau^i | \mathbf{x}^i)\|$, G .

Lemma 1. *Let $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$ be defined in Eq. 6 for all $i \in \mathcal{V}$. Assumption 1 implies the following conclusions:*

- $\|\mathbf{g}_i(\tau^i | \mathbf{x}_1) - \mathbf{g}_i(\tau^i | \mathbf{x}_2)\| \leq L\|\mathbf{x}_1 - \mathbf{x}_2\|$, $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^{d_i}$, with $L = \frac{C_h R}{(1-\gamma)^2}$;
- $J(\mathbf{x})$ is L -smooth, i.e., $\|\nabla^2 J(\mathbf{x})\| \leq L$;
- $\|\mathbf{g}_i(\tau^i | \mathbf{x}^i)\| \leq G$, $\forall \mathbf{x}^i \in \mathbb{R}^{d_i}$, with $G = \frac{C_g R}{(1-\gamma)^2}$;
- Variance of stochastic policy gradient $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$ is bounded, i.e., $\mathbb{E}[\|\mathbf{g}_i(\tau^i | \mathbf{x}^i) - \nabla J_i(\mathbf{x}^i)\|^2] \leq \sigma_i^2$, $\bar{\sigma}^2 = \frac{1}{N} \sum_{i=1}^N \sigma_i^2$, for all $i \in \mathcal{V}$, where $\bar{\sigma}^2 = \frac{C_g^2 R^2}{(1-\gamma)^4}$.

Proof. Recall the definition of $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$ in Eq. 6 such that the PGT estimator (Sutton et al. 1999) is as follows:

$$\mathbf{g}_i(\tau^i | \mathbf{x}^i) = \sum_{h=0}^{H-1} \sum_{q=0}^h (\gamma^q r_q^i(a_q^i, s_q) - b_q) \nabla_{\mathbf{x}^i} \log \pi^i(a_h^i, s_h),$$

where b_q is a constant baseline and we specify the subscript \mathbf{x}^i to indicate the gradient w.r.t it. Later this is omitted unless specified appropriately. Further, we have

$$\begin{aligned}\|\nabla \mathbf{g}_i(\tau^i | \mathbf{x}^i)\| &= \left\| \sum_{h=0}^{H-1} \nabla^2 \log \pi^i(a_h^i, s_h) \left(\sum_{q=0}^h \gamma^q r_q^i(a_q^i, s_q) \right) \right\| \\ &\leq \left(\sum_{q=0}^{H-1} \|\nabla^2 \log \pi^i(a_q^i, s_q)\| \frac{R}{1-\gamma} \right) \\ &\leq \frac{C_h R}{(1-\gamma)^2}\end{aligned}\tag{18}$$

where $b_q = 0$. The above inequalities follow from that $\gamma \in (0, 1)$. When $b_q \neq 0$, we can easily scale it with γ^h and the above result still holds but with different constant coefficient. Thus, the first part of Lemma 1 can be proved.

As the PGT estimator is an unbiased estimator of the policy gradient $\nabla J_i(\mathbf{x}^i)$, we then have

$$\nabla J_i(\mathbf{x}^i) = \mathbb{E}_{\tau^i}[\mathbf{g}_i(\tau^i | \mathbf{x}^i)], \nabla^2 J_i(\mathbf{x}^i) = \mathbb{E}_{\tau^i}[\nabla \mathbf{g}_i(\tau^i | \mathbf{x}^i)].$$

Hence, the smoothness of $J_i(\mathbf{x}^i)$ can be directly implied from the Lipchitzness of $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$,

$$\|\nabla^2 J_i(\mathbf{x}^i)\| = \|\mathbb{E}_{\tau^i}[\nabla \mathbf{g}_i(\tau^i | \mathbf{x}^i)]\| \leq \|\nabla \mathbf{g}_i(\tau^i | \mathbf{x}^i)\| \leq \frac{C_h R}{(1-\gamma)^2}.$$

As the above relationship holds for each agent $i \in \mathcal{V}$, the second part of Lemma 1 can be obtained.

Similarly, taking the norm of $\mathbf{g}_i(\tau^i | \mathbf{x}^i)$ leads to

$$\|\mathbf{g}_i(\tau^i | \mathbf{x}^i)\| \leq \left\| \sum_{h=0}^{H-1} \nabla \log \pi^i(a_h^i, s_h) \frac{\gamma^h R(1 - \gamma^{H-h})}{1 - \gamma} \right\| \leq \frac{C_g R}{(1 - \gamma)^2}.$$

Based on the above inequality, it is immediately obtained that

$$\|\mathbf{g}_i(\tau^i | \mathbf{x}^i)\|^2 \leq \frac{C_g^2 R^2}{(1 - \gamma)^4}.$$

For a random variable X , we have that $\mathbb{E}[\|X - \mathbb{E}[X]\|^2] \leq \mathbb{E}[\|X\|^2]$ such that

$$\mathbb{E}[\|\mathbf{g}_i(\tau^i | \mathbf{x}^i) - \nabla J_i(\mathbf{x}^i)\|^2] \leq \frac{C_g^2 R^2}{(1 - \gamma)^4}.$$

Also, $\bar{\sigma}^2 = \frac{C_g^2 R^2}{(1 - \gamma)^4}$. The proof is complete. \square

Lemma 1 completes the definitions of L and G . One immediate observation for both L and G is that they are independent of the horizon H , which hence are tighter. In the following, we introduce a relation for the importance sampling weight. We first introduce the Rényi divergence between two distributions Q and Z as follows:

$$\mathcal{D}_\alpha(Q||Z) = \frac{1}{\alpha - 1} \log_2 \int_x Q(x) \left(\frac{Q(x)}{Z(x)} \right)^{\alpha-1} dx,$$

which is non-negative for all $\alpha > 0$. The exponentiated Rényi divergence is $\mathcal{D}_\alpha(Q||Z) = 2^{\mathcal{D}_\alpha(Q||Z)}$.

Lemma 2. (Lemma 1 in (Cortes, Mansour, and Mohri 2010)) Let $v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2) = \frac{p_i(\tau^i | \mathbf{x}_1)}{p_i(\tau^i | \mathbf{x}_2)}$ be the importance sampling weight for distributions $p_i(\tau^i | \mathbf{x}_1)$ and $p_i(\tau^i | \mathbf{x}_2)$. Then the expectation, second moment and variance of $v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2)$ satisfy the following results:

$$\mathbb{E}[v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2)] = 1, \mathbb{E}[v_i^2(\tau^i | \mathbf{x}_1, \mathbf{x}_2)] = \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)),$$

and

$$\mathbb{V}(v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2)) = \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) - 1.$$

With Lemma 2 in hand, we now are ready to state the relation between $\mathbb{V}(v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2))$ and $\|\mathbf{x}_1 - \mathbf{x}_2\|^2$.

Lemma 3. For any $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^{d_i}, \forall i \in \mathcal{V}$, let $v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2) = \frac{p_i(\tau^i | \mathbf{x}_1)}{p_i(\tau^i | \mathbf{x}_2)}$. With Assumptions 1 and 3, We have

$$\mathbb{V}[v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2)] \leq C_v \|\mathbf{x}_1 - \mathbf{x}_2\|^2, \quad (19)$$

where $C_v = H(2HC_g^2 + C_h)(\mathcal{M} + 1)$.

Proof. According to Lemma 2, the variance of the importance sampling weight is

$$\mathbb{V}[v_i(\tau^i | \mathbf{x}_1, \mathbf{x}_2)] = \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) - 1.$$

By definition, the following relationship can be obtained

$$\mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) = \int_{\tau^i} p_i(\tau^i | \mathbf{x}_1) \frac{p_i(\tau^i | \mathbf{x}_1)}{p_i(\tau^i | \mathbf{x}_2)} d\tau^i = \int_{\tau^i} p_i(\tau^i | \mathbf{x}_1)^2 p_i(\tau^i | \mathbf{x}_2)^{-1} d\tau^i.$$

Taking the gradient of $\mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2))$ w.r.t. \mathbf{x}_1 , we then have

$$\nabla_{\mathbf{x}_1} \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) = 2 \int_{\tau^i} p_i(\tau^i | \mathbf{x}_1) \nabla_{\mathbf{x}_1} p_i(\tau^i | \mathbf{x}_1) p_i(\tau^i | \mathbf{x}_2)^{-1} d\tau^i.$$

Letting $\mathbf{x}_1 = \mathbf{x}_2$ in the last equation yields

$$\nabla_{\mathbf{x}_1} \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) \Big|_{\mathbf{x}_1=\mathbf{x}_2} = 2 \int_{\tau^i} \nabla_{\mathbf{x}_1} p_i(\tau^i | \mathbf{x}_1) d\tau^i \Big|_{\mathbf{x}_1=\mathbf{x}_2} = 0.$$

Applying the mean value theorem w.r.t. \mathbf{x}_1 results in

$$\mathcal{D}_2(p_i(\tau^i | \mathbf{x}_1) || p_i(\tau^i | \mathbf{x}_2)) = 1 + \frac{1}{2} (\mathbf{x}_1 - \mathbf{x}_2)^\top \nabla_{\mathbf{x}_3}^2 \mathcal{D}_2(p_i(\tau^i | \mathbf{x}_3) || p_i(\tau^i | \mathbf{x}_2)) (\mathbf{x}_1 - \mathbf{x}_2), \quad (20)$$

where $\mathbf{x}_3 = t\mathbf{x}_1 + (1-t)\mathbf{x}_2, t \in [0, 1]$. The last inequality follows from the fact that $\mathcal{D}_2(p_i(\tau^i|\mathbf{x}_2)||p_i(\tau^i|\mathbf{x}_2)) = 1$. We have now obtained another expression for $\mathcal{D}_2(p_i(\tau^i|\mathbf{x}_1)||p_i(\tau^i|\mathbf{x}_2))$ and to bound it, we shall compute the Hessian matrix $\nabla_{\mathbf{x}_3}^2 \mathcal{D}_2(p_i(\tau^i|\mathbf{x}_3)||p_i(\tau^i|\mathbf{x}_2))$. Taking the gradient $\nabla_{\mathbf{x}_3} \mathcal{D}_2(p_i(\tau^i|\mathbf{x}_3)||p_i(\tau^i|\mathbf{x}_2))$ of w.r.t. \mathbf{x}_3 leads to

$$\begin{aligned} \nabla_{\mathbf{x}_3}^2 \mathcal{D}_2(p_i(\tau^i|\mathbf{x}_3)||p_i(\tau^i|\mathbf{x}_2)) &= 2 \int_{\tau^i} \nabla_{\mathbf{x}_3} \log p_i(\tau^i|\mathbf{x}_3) \nabla_{\mathbf{x}_3} \log p_i(\tau^i|\mathbf{x}_3)^\top \frac{p_i(\tau^i|\mathbf{x}_3)^2}{p_i(\tau^i|\mathbf{x}_2)} d\tau^i \\ &\quad + 2 \int_{\tau^i} \nabla_{\mathbf{x}_3}^2 p_i(\tau^i|\mathbf{x}_3) p_i(\tau^i|\mathbf{x}_3) p_i(\tau^i|\mathbf{x}_2)^{-1} d\tau^i. \end{aligned}$$

The above equation implies that to obtain the Hessian matrix, we need to attain $\nabla_{\mathbf{x}_3}^2 p_i(\tau^i|\mathbf{x}_3)$, which signifies the Hessian matrix of the trajectory distribution function. We next derive the the Hessian matrix of log-density function.

$$\nabla_{\mathbf{x}_3}^2 \log p_i(\tau^i|\mathbf{x}_3) = -p_i(\tau^i|\mathbf{x}_3)^2 \nabla_{\mathbf{x}_3} p_i(\tau^i|\mathbf{x}_3) \nabla_{\mathbf{x}_3} p_i(\tau^i|\mathbf{x}_3)^\top + p_i(\tau^i|\mathbf{x}_3)^{-1} \nabla_{\mathbf{x}_3}^2 p_i(\tau^i|\mathbf{x}_3).$$

Combining the last two equations yields

$$\begin{aligned} \|\nabla_{\mathbf{x}_3}^2 \mathcal{D}_2(p_i(\tau^i|\mathbf{x}_3)||p_i(\tau^i|\mathbf{x}_2))\| &= \left\| 4 \int_{\tau^i} \nabla_{\mathbf{x}_3} \log p_i(\tau^i|\mathbf{x}_3) \nabla_{\mathbf{x}_3} p_i(\tau^i|\mathbf{x}_3)^\top \frac{p_i(\tau^i|\mathbf{x}_3)^2}{p_i(\tau^i|\mathbf{x}_2)} d\tau^i \right. \\ &\quad \left. + 2 \int_{\tau^i} \nabla_{\mathbf{x}_3}^2 \log p_i(\tau^i|\mathbf{x}_3) \frac{p_i(\tau^i|\mathbf{x}_3)^2}{p_i(\tau^i|\mathbf{x}_2)} d\tau^i \right\| \\ &\leq \int_{\tau^i} \frac{p_i(\tau^i|\mathbf{x}_3)^2}{p_i(\tau^i|\mathbf{x}_2)} (4\|\nabla_{\mathbf{x}_3} \log p_i(\tau^i|\mathbf{x}_3)\|^2 + 2\|\nabla_{\mathbf{x}_3}^2 \log p_i(\tau^i|\mathbf{x}_3)\|) d\tau^i \\ &\leq (4H^2 C_g^2 + 2HC_h) \mathbb{E}[v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)^2] \\ &\leq 2H(2HC_g^2 + C_h)(\mathcal{M} + 1), \end{aligned}$$

where the second inequality follows from Assumption 1 and Lemma 2 and the last inequality follows from Assumption 3. Therefore, substituting the above result into Eq. 20 attain the following desirable result

$$\mathbb{V}[v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)] = \mathcal{D}_2(p_i(\tau^i|\mathbf{x}_1)||p_i(\tau^i|\mathbf{x}_2)) - 1 \leq C_v \|\mathbf{x}_1 - \mathbf{x}_2\|^2.$$

□

Lemma 3 is critical in the proof as it will translate HIS-SARAH into a squared norm that can be bounded above. We next present the outline of how to prove Theorem 1 and apply the same proof techniques to show Theorem 2. We first define some notations for the convenience of proof. In a generalized case, the dimension for each \mathbf{x}^i is not necessarily the same, while in this work we assume that $d_1 = d_2 = \dots = d_N = \frac{d}{N}$ for the ease of exposition and in this context, d/N is assumed to be an integer. Recalling the update laws for \mathbf{v}_k^i and \mathbf{x}_k^i in Algorithm 1 in a matrix form, we have:

$$\mathbf{v}_{k+1} = \underline{\mathbf{W}} \mathbf{v}_k + \mathbf{u}_k - \mathbf{u}_{k-1} \quad (21a)$$

$$\mathbf{x}_{k+1} = \underline{\mathbf{W}}(\mathbf{x}_k + \eta \mathbf{v}_{k+1}) \quad (21b)$$

where $\underline{\mathbf{W}} \triangleq \mathbf{W} \otimes \mathbf{I}_{d/N}$ and $\mathbf{x}_k, \mathbf{v}_k, \mathbf{u}_k$ are square-integrable random vectors in \mathbb{R}^d that concatenate the local estimates of the solution $\{\mathbf{x}_k^i\}_{i=1}^N$, gradient trackers $\{\mathbf{v}_k^i\}_{i=1}^N$, and the stochastic policy gradient surrogates $\{\mathbf{u}_k^i\}_{i=1}^N$. Additionally, we denote $\nabla \mathbf{J}(\mathbf{x}_k) = [\nabla J_1(\mathbf{x}_k^1)^\top, \dots, \nabla J_N(\mathbf{x}_k^N)^\top]^\top$ and define the ensemble averages as follows.

$$\bar{\mathbf{x}}_k \triangleq \frac{1}{N} (\mathbf{1}_N^\top \otimes \mathbf{I}_{d/N}) \mathbf{x}_k, \bar{\mathbf{v}}_k \triangleq \frac{1}{N} (\mathbf{1}_N^\top \otimes \mathbf{I}_{d/N}) \mathbf{v}_k,$$

$$\bar{\mathbf{u}}_k \triangleq \frac{1}{N} (\mathbf{1}_N^\top \otimes \mathbf{I}_{d/N}) \mathbf{u}_k, \bar{\nabla \mathbf{J}}(\mathbf{x}_k) \triangleq \frac{1}{N} (\mathbf{1}_N^\top \otimes \mathbf{I}_{d/N}) \nabla \mathbf{f}(\mathbf{x}_k).$$

With these definitions, we present some known results for gradient tracking type of algorithms in decentralized optimization and refer interested readers to (Xin, Khan, and Kar 2021) for the detail of proof.

Lemma 4. *The following relationships hold for MDPGT.*

- Define $\Lambda = \frac{1}{N} (\mathbf{1}_N^\top \otimes \mathbf{I}_{d/N})$. Thus, $\|\underline{\mathbf{W}} \mathbf{x} - \Lambda \mathbf{x}\| \leq \lambda \|\mathbf{x} - \Lambda \mathbf{x}\|$.
- $\bar{\mathbf{v}}_{k+1} = \bar{\mathbf{u}}_k$, for all $k \geq 0$.
- $\|\bar{\nabla \mathbf{J}}(\mathbf{x}_k) - \nabla J(\bar{\mathbf{x}}_k)\|^2 \leq \frac{L^2}{N} \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2$, for all $k \geq 0$.

The first relationship in Lemma 4 has been well-known due to a doubly stochastic mixing matrix \mathbf{W} and the second one is an immediate consequence of the update law for the gradient tracker. The third relationship is due to the second conclusion from Lemma 1. Multiplying the update for \mathbf{x}_{k+1}^i by Λ yields the following equality

$$\bar{\mathbf{x}}_{k+1} = \bar{\mathbf{x}}_k + \eta \bar{\mathbf{v}}_{k+1} = \bar{\mathbf{x}}_k + \eta \bar{\mathbf{u}}_k, \quad \forall k \geq 0. \quad (22)$$

We next establish a key technical lemma that sheds light on the convergence in terms of the second moment of $\nabla J(\bar{\mathbf{x}}_k)$.

Lemma 5. *Let $\bar{\mathbf{x}}_k$ be generated by Eq. 22. If the step size $\eta \in (0, \frac{1}{2L}]$, then for all $K \geq 0$, we have:*

$$\begin{aligned} \sum_{k=0}^K \|\nabla J(\bar{\mathbf{x}}_k)\|^2 &\leq \frac{2(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{2} \sum_{k=0}^K \|\bar{\mathbf{u}}_k\|^2 \\ &+ 2 \sum_{k=0}^K \|\bar{\mathbf{u}}_k - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2 + \frac{2L^2}{N} \sum_{k=0}^K \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2. \end{aligned} \quad (23)$$

Proof. Based on the second conclusion of Lemma 1, we have the following relationship:

$$\begin{aligned} J(\bar{\mathbf{x}}_{k+1}) &\geq J(\bar{\mathbf{x}}_k) + \langle \nabla J(\bar{\mathbf{x}}_k), (\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{x}}_k) \rangle - \frac{L}{2} \|\bar{\mathbf{x}}_{k+1} - \bar{\mathbf{x}}_k\|^2 \\ &\geq J(\bar{\mathbf{x}}_k) + \eta \langle \nabla J(\bar{\mathbf{x}}_k), \bar{\mathbf{u}}_k \rangle - \frac{\eta^2 L}{2} \|\bar{\mathbf{u}}_k\|^2. \end{aligned} \quad (24)$$

According to the basic inequality $\langle \mathbf{a}, \mathbf{b} \rangle = \frac{1}{2}(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - \|\mathbf{a} - \mathbf{b}\|^2)$, the last inequality, the following can be obtained

$$\begin{aligned} J(\bar{\mathbf{x}}_{k+1}) &\geq J(\bar{\mathbf{x}}_k) + \frac{\eta}{2} \|\nabla J(\bar{\mathbf{x}}_k)\|^2 + \left(\frac{\eta}{2} - \frac{\eta^2 L}{2} \right) \|\bar{\mathbf{u}}_k\|^2 - \frac{\eta}{2} \|\bar{\mathbf{u}}_k - \nabla J(\bar{\mathbf{x}}_k)\|^2 \\ &\geq J(\bar{\mathbf{x}}_k) + \frac{\eta}{2} \|\nabla J(\bar{\mathbf{x}}_k)\|^2 + \left(\frac{\eta}{2} - \frac{\eta^2 L}{2} \right) \|\bar{\mathbf{u}}_k\|^2 - \frac{\eta}{2} \|\bar{\mathbf{u}}_k - \nabla J(\bar{\mathbf{x}}_k) + \bar{\nabla \mathbf{J}}(\mathbf{x}_k) - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2 \\ &\geq J(\bar{\mathbf{x}}_k) + \frac{\eta}{2} \|\nabla J(\bar{\mathbf{x}}_k)\|^2 + \frac{\eta}{4} \|\bar{\mathbf{u}}_k\|^2 - \eta \|\bar{\mathbf{u}}_k - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2 - \frac{\eta L^2}{N} \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2, \end{aligned} \quad (25)$$

where the third inequality is due to $\|\mathbf{a} + \mathbf{b}\|^2 \leq 2(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2)$, the smoothness of J and $\eta \in (0, \frac{1}{2L}]$. Rewriting Eq. 25 yields the following relationship

$$\|\nabla J(\bar{\mathbf{x}}_k)\|^2 \leq \frac{2(J(\bar{\mathbf{x}}_{k+1}) - J(\bar{\mathbf{x}}_k))}{\eta} - \frac{1}{2} \|\bar{\mathbf{u}}_k\|^2 + 2 \|\bar{\mathbf{u}}_k - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2 + \frac{2L^2}{N} \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2. \quad (26)$$

Summing up the last equation from 0 to K completes the proof. \square

According to Lemma 5, it is clearly observed that in order to arrive at the main results presented in the previous section, the upper bound of the right hand side in the last inequality needs to be explicitly quantified. One may argue that in the main theorems it is $\tilde{\mathbf{x}}_K$ instead of $\bar{\mathbf{x}}_K$. Dividing Eq. 23 by $\frac{1}{K+1}$ allows us to approximate $\tilde{\mathbf{x}}_K$ using the ensemble average. Additionally, a sufficiently large K implies all agents have converged to the (near-) optimal solution, suggesting that the gap between $\tilde{\mathbf{x}}_K$ and $\bar{\mathbf{x}}_K$ is quite close to 0. Though in the final result, the convergence rate is specifically for $\tilde{\mathbf{x}}_K$, Lemma 5 helps facilitate the understanding of convergence for MDPGT and will be critical to derive the final error bounds in the main theorems. Thus, in the sequel, the analysis is dedicated to finding the upper bounds for $\sum_{k=0}^K \|\bar{\mathbf{u}}_k - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2$ and $\sum_{k=0}^K \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2$. To start the analysis, we present two relationships that will help significantly characterize the recursions of gradient variances, consensus, and gradient tracking errors.

Lemma 6. *Let $\{e_k\}_{k \geq 0}$, $\{b_k\}_{k \geq 0}$, and $\{c_k\}_{k \geq 0}$ be nonnegative sequences such that $e_k \leq q e_{k-1} + q b_{k-1} + c_k + C, \forall k \geq 1$, where $q \in (0, 1)$ and $C \geq 0$. Then, we have for all $K \geq 1$,*

$$\sum_{k=1}^K e_k \leq \frac{e_0}{1-q} + \frac{1}{1-q} \sum_{k=0}^{K-1} b_k + \frac{1}{1-q} \sum_{k=0}^K c_k + \frac{CK}{1-q}. \quad (27)$$

Similarly, if $e_{k+1} \leq q e_k + b_{k-1} + C, \forall k \geq 1$, then, we have for all $K \geq 2$,

$$\sum_{k=1}^K e_k \leq \frac{e_1}{1-q} + \frac{1}{1-q} \sum_{k=0}^{K-2} b_k + \frac{CK}{1-q}. \quad (28)$$

Please follow the proof of Lemma 6 in (Xin, Khan, and Kar 2021) for more detail. We next derive the relevant recursions for the gradient variances $\mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2]$ and $\mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2]$. Essentially the former can be treated intuitively as the average of the latter up to a factor of $\frac{1}{N^2}$, but they are proved separately.

Lemma 7. *Let $\bar{\mathbf{u}}_k$ and \mathbf{x}_k be generated by MDPGT. Then for all $k \geq 1$, we have*

$$\begin{aligned} \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2] &\leq (1-\beta)^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})\|^2] \\ &+ \frac{12(L^2 + G^2 C_v) \eta^2 (1-\beta)^2}{N} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] + \frac{2\beta^2 \bar{\sigma}^2}{N} \\ &+ \frac{12(L^2 + G^2 C_v) (1-\beta)^2}{N^2} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 \\ &+ \|\mathbf{x}_{k-1} - \Lambda \mathbf{x}_{k-1}\|^2], \end{aligned} \quad (29)$$

and

$$\begin{aligned} \mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2] &\leq (1-\beta)^2 \mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2] \\ &+ 12N(L^2 + G^2 C_v) \eta^2 (1-\beta)^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] + 2\beta^2 \bar{\sigma}^2 N \\ &+ 12(L^2 + G^2 C_v) (1-\beta)^2 \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 \\ &+ \|\mathbf{x}_{k-1} - \Lambda \mathbf{x}_{k-1}\|^2]. \end{aligned} \quad (30)$$

Proof. Recalling

$$\mathbf{u}_k^i = \beta \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) + (1-\beta)[\mathbf{u}_{k-1}^i + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)], \quad (31)$$

we then have

$$\begin{aligned} \mathbf{u}_k^i - \nabla J_i(\mathbf{x}_k^i) &= \beta \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) + (1-\beta)[\mathbf{u}_{k-1}^i + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)] \\ &- \beta \nabla J_i(\mathbf{x}_k^i) - (1-\beta) \nabla J_i(\mathbf{x}_k^i) \\ &= \beta(\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)) + (1-\beta)[\nabla J_i(\mathbf{x}_{k-1}^i) + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) \\ &- v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)] + (1-\beta)(\mathbf{u}_k^i - \nabla J_i(\mathbf{x}_{k-1}^i)). \end{aligned} \quad (32)$$

As $\mathbb{E}[\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)] = \mathbf{0}_{d/N}$ and $\mathbb{E}[\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)] = \mathbf{0}_{d/N}$. Applying Eq. 32 from 1 to N and taking the ensemble average results in

$$\begin{aligned} \bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k) &= (1-\beta)(\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})) + \beta \frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)) \\ &+ (1-\beta) \frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)) \end{aligned} \quad (33)$$

Let

$$\mathbf{n}_k = \frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i))$$

and

$$\mathbf{z}_k = \frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)).$$

We have

$$\begin{aligned} \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2] &= (1-\beta)^2 \|\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})\|^2 + \mathbb{E}[\|\beta \mathbf{n}_k + (1-\beta) \mathbf{z}_k\|^2] \\ &+ 2\mathbb{E}[\langle (1-\beta)(\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})), \beta \mathbf{n}_k + (1-\beta) \mathbf{z}_k \rangle] \\ &= (1-\beta)^2 \|\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})\|^2 + \mathbb{E}[\|\beta \mathbf{n}_k + (1-\beta) \mathbf{z}_k\|^2] \\ &\leq (1-\beta)^2 \|\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})\|^2 + 2\beta^2 \mathbb{E}[\|\mathbf{n}_k\|^2] + 2(1-\beta)^2 \mathbb{E}[\|\mathbf{z}_k\|^2]. \end{aligned} \quad (34)$$

The second equality follows from that $\mathbb{E}[\langle (1-\beta)(\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{J}(\mathbf{x}_{k-1})), \beta \mathbf{n}_k + (1-\beta) \mathbf{z}_k \rangle] = 0$ as $\mathbb{E}[\mathbf{n}_k] = \mathbf{0}_{d/N}$ and $\mathbb{E}[\mathbf{z}_k] = \mathbf{0}_{d/N}$. We next bound the terms $\mathbb{E}[\|\mathbf{n}_k\|^2]$ and $\mathbb{E}[\|\mathbf{z}_k\|^2]$. As for $k \geq 1$,

$$\begin{aligned} \mathbb{E}[\|\mathbf{n}_k\|^2] &= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] + \frac{1}{N^2} \sum_{i \neq j} \mathbb{E}[\langle \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i), \mathbf{g}_j(\tau_k^j | \mathbf{x}_k^j) - \nabla J_j(\mathbf{x}_k^j) \rangle] \\ &= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \leq \frac{\bar{\sigma}^2}{N}. \end{aligned} \quad (35)$$

The second equality follows from that

$$\begin{aligned}
& \mathbb{E}[\langle \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i), \mathbf{g}_j(\tau_k^j | \mathbf{x}_k^j) - \nabla J_j(\mathbf{x}_k^j) \rangle] \\
&= \mathbb{E}[\langle \mathbb{E}[\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i)] - \nabla J_i(\mathbf{x}_k^i), \mathbf{g}_j(\tau_k^j | \mathbf{x}_k^j) - \nabla J_j(\mathbf{x}_k^j) \rangle] \\
&= 0.
\end{aligned} \tag{36}$$

For the term $\mathbb{E}[\|\mathbf{z}_k\|^2]$, we have for all $k \geq 1$,

$$\begin{aligned}
\mathbb{E}[\|\mathbf{z}_k\|^2] &= \mathbb{E}[\|\frac{1}{N} \sum_{i=1}^N (\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i))\|^2] \\
&= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \\
&\quad + \frac{1}{N^2} \sum_{i \neq j} \mathbb{E}[\langle \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i), \\
&\quad \mathbf{g}_j(\tau_k^j | \mathbf{x}_k^j) - v_j(\tau_k^j | \mathbf{x}_{k-1}^j, \mathbf{x}_k^j) \mathbf{g}_j(\tau_k^j | \mathbf{x}_{k-1}^j) + \nabla J_j(\mathbf{x}_{k-1}^j) - \nabla J_j(\mathbf{x}_k^j) \rangle] \\
&= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \nabla J_i(\mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \\
&\leq \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2].
\end{aligned} \tag{37}$$

The third equality follows from the same argument we have for the $\mathbb{E}[\langle \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i), \mathbf{g}_j(\tau_k^j | \mathbf{x}_k^j) - \nabla J_j(\mathbf{x}_k^j) \rangle]$. The last inequality follows from the fact that

$$\mathbb{E}[\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)] = \nabla J_i(\mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_{k-1}^i)$$

and the variance decomposition, i.e., for any vector $\mathbf{a} \in \mathbb{R}^p$,

$$\mathbb{E}[\|\mathbf{a} - \mathbb{E}[\mathbf{a}]\|^2] = \mathbb{E}[\|\mathbf{a}\|^2] - \|\mathbb{E}[\mathbf{a}]\|^2.$$

Hence, we have now

$$\begin{aligned}
\mathbb{E}[\|\mathbf{z}_k\|^2] &\leq \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2] \\
&= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) + \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2] \\
&= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[2\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2 + 2\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2] \\
&\leq \frac{2L^2}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + \frac{2}{N^2} \sum_{i=1}^N \mathbb{E}[\|1 - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\|^2 \|\mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2] \\
&\leq \frac{2L^2}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + \frac{2G^2}{N^2} \mathbb{E}[\|1 - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\|^2] \\
&= \frac{2L^2}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + \frac{2G^2}{N^2} \sum_{i=1}^N \mathbb{V}(v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i)) \\
&\leq \frac{2L^2}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + \frac{2G^2 C_v}{N^2} \sum_{i=1}^N \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] \\
&= \frac{2L^2 + 2G^2 C_v}{N^2} \mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] \\
&= \frac{2L^2 + 2G^2 C_v}{N^2} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k + \Lambda \mathbf{x}_k - \Lambda \mathbf{x}_{k-1} + \Lambda \mathbf{x}_{k-1} - \mathbf{x}_{k-1}\|^2] \\
&\leq \frac{6L^2 + 6G^2 C_v}{N} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] + \frac{6L^2 + 6G^2 C_v}{N^2} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 + \|\Lambda \mathbf{x}_{k-1} - \mathbf{x}_{k-1}\|^2].
\end{aligned} \tag{38}$$

The second inequality follows from the Cauthy-Schwartz inequality. The third inequality follows from Lemma 1 while the fourth inequality follows from Lemma 3. Hence, the following relationship can be obtained

$$\begin{aligned}
\mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{f}(\mathbf{x}_k)\|^2] &\leq (1 - \beta)^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1} - \bar{\nabla} \mathbf{f}(\mathbf{x}_{k-1})\|^2] + \frac{12L^2 + 12G^2 C_v}{N} \eta^2 (1 - \beta)^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\
&\quad + \frac{2\beta^2 \bar{\sigma}^2}{N} + \frac{12L^2 + 12G^2 C_v}{N^2} (1 - \beta)^2 \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 + \|\Lambda \mathbf{x}_{k-1} - \mathbf{x}_{k-1}\|^2],
\end{aligned} \tag{39}$$

which completes the first part of proof for Lemma 7.

Recalling Eq. 32, we have

$$\begin{aligned}
\mathbf{u}_k^i - \nabla J_i(\mathbf{x}_k^i) &= \beta(\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)) + (1 - \beta)[\nabla J_i(\mathbf{x}_{k-1}^i) + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) \\
&\quad - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)] + (1 - \beta)(\mathbf{u}_k^i - \nabla J_i(\mathbf{x}_{k-1}^i)).
\end{aligned} \tag{40}$$

Based on the analysis above, it is immediately obtained that the expectations of the first and second terms are all zero. Following the similar proof in Eq. 34, we have for all $k \geq 1$,

$$\begin{aligned}
\mathbb{E}[\|\mathbf{u}_k^i - \nabla J_i(\mathbf{x}_k^i)\|^2] &\leq (1 - \beta)^2 \|\mathbf{u}_{k-1}^i - \nabla J_i(\mathbf{x}_{k-1}^i)\|^2 + 2\beta^2 \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \\
&\quad + 2(1 - \beta)^2 \mathbb{E}[\|\nabla J_i(\mathbf{x}_{k-1}^i) + \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \\
&\leq (1 - \beta)^2 \|\mathbf{u}_{k-1}^i - \nabla J_i(\mathbf{x}_{k-1}^i)\|^2 + 2\beta^2 \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \nabla J_i(\mathbf{x}_k^i)\|^2] \\
&\quad + 2(1 - \beta)^2 \mathbb{E}[\|\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)\|^2]
\end{aligned} \tag{41}$$

According to Lemma 1 and Lemma 3, the following relationship can be attained

$$\begin{aligned}
\mathbb{E}[\|\mathbf{u}_k^i - \nabla J_i(\mathbf{x}_k^i)\|^2] &\leq (1 - \beta)^2 \|\mathbf{u}_{k-1}^i - \nabla J_i(\mathbf{x}_{k-1}^i)\|^2 + 2\beta^2 \sigma_i^2 + 4(1 - \beta)^2 (L^2 + C_v G^2) \mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] \\
&\leq (1 - \beta)^2 \|\mathbf{u}_{k-1}^i - \nabla J_i(\mathbf{x}_{k-1}^i)\|^2 + 2\beta^2 \sigma_i^2 \\
&\quad + 12(1 - \beta)^2 (L^2 + C_v G^2) (\mathbb{E}[\|\mathbf{x}_k^i - \bar{\mathbf{x}}_k\|^2] + \|\bar{\mathbf{x}}_k - \bar{\mathbf{x}}_{k-1}\|^2 + \|\bar{\mathbf{x}}_{k-1} - \mathbf{x}_{k-1}^i\|^2) \\
&= (1 - \beta)^2 \|\mathbf{u}_{k-1}^i - \nabla J_i(\mathbf{x}_{k-1}^i)\|^2 + 2\beta^2 \sigma_i^2 + 12(1 - \beta)^2 (L^2 + C_v G^2) \eta^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\
&\quad + 12(1 - \beta)^2 (L^2 + C_v G^2) (\mathbb{E}[\|\mathbf{x}_k^i - \bar{\mathbf{x}}_k\|^2] + \|\mathbf{x}_{k-1}^i - \bar{\mathbf{x}}_{k-1}\|^2).
\end{aligned} \tag{42}$$

Applying Eq. 42 over i from 1 to N completes the second part of the proof for Lemma 7. \square

With Lemma 6 and Lemma 7 in hand, we now are ready to present the upper bounds for both $\sum_{k=1}^K \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2]$ and $\sum_{k=1}^K \mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2]$ in the following lemma.

Lemma 8. *Let $\bar{\mathbf{u}}_k$ and \mathbf{x}_k be generated by MDPGT initialized with a mini-batch of trajectories \mathcal{B} . Then for any $\beta \in (0, 1)$, $\forall K \geq 1$, we have*

$$\begin{aligned} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2] &\leq \frac{\bar{\sigma}^2}{|\mathcal{B}|N\beta} + \frac{12(L^2 + G^2C_v)\eta^2}{N\beta} \times \\ &\sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{24(L^2 + G^2C_v)}{N^2\beta} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] \\ &+ \frac{2\beta\bar{\sigma}^2K}{N}, \end{aligned} \quad (43)$$

and

$$\begin{aligned} \sum_{k=0}^K \mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2] &\leq \frac{N\bar{\sigma}^2}{|\mathcal{B}|\beta} + \frac{12N(L^2 + G^2C_v)\eta^2}{\beta} \times \\ &\sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{24(L^2 + G^2C_v)}{\beta} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] \\ &+ 2N\beta\bar{\sigma}^2K. \end{aligned} \quad (44)$$

Proof. We apply the conclusions from Lemma 6 to Lemma 7. We first show the upper error bound for $\sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2]$. Substituting Eq. 29 into Eq. 27 leads to the following inequality

$$\begin{aligned} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k - \bar{\nabla} \mathbf{J}(\mathbf{x}_k)\|^2] &\leq \frac{1}{1 - (1 - \beta)^2} \mathbb{E}[\|\bar{\mathbf{u}}_0 - \bar{\nabla} \mathbf{J}(\mathbf{x}_0)\|^2] + \frac{12(L^2 + G^2C_v)\eta^2(1 - \beta)^2}{N(1 - (1 - \beta)^2)} \sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\ &+ \frac{2\beta^2\bar{\sigma}^2K}{N(1 - (1 - \beta)^2)} + \frac{12(L^2 + G^2C_v)(1 - \beta)^2}{N^2(1 - (1 - \beta)^2)} \mathbb{E} \sum_{k=0}^{K-1} [\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 + \|\mathbf{x}_{k-1} - \Lambda \mathbf{x}_{k-1}\|^2] \\ &\leq \frac{1}{1 - (1 - \beta)^2} \mathbb{E}[\|\bar{\mathbf{u}}_0 - \bar{\nabla} \mathbf{J}(\mathbf{x}_0)\|^2] + \frac{12(L^2 + G^2C_v)\eta^2(1 - \beta)^2}{N(1 - (1 - \beta)^2)} \sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\ &+ \frac{2\beta^2\bar{\sigma}^2K}{N(1 - (1 - \beta)^2)} + \frac{24(L^2 + G^2C_v)(1 - \beta)^2}{N^2(1 - (1 - \beta)^2)} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2]. \end{aligned} \quad (45)$$

We now process the first term on the right hand side of the last inequality.

$$\mathbb{E}[\|\bar{\mathbf{u}}_0 - \bar{\nabla} \mathbf{J}(\mathbf{x}_0)\|^2] = \mathbb{E}\left[\left\|\frac{1}{N} \sum_{i=1}^N \frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} (\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i))^2\right\|^2\right].$$

It is immediately obtained that

$$\mathbb{E}[\|\bar{\mathbf{u}}_0 - \bar{\nabla} \mathbf{J}(\mathbf{x}_0)\|^2] = \frac{1}{N^2|\mathcal{B}|^2} \sum_{i=1}^N \sum_{m=1}^{|\mathcal{B}|} \mathbb{E}[\|\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i)\|^2] \leq \frac{\bar{\sigma}^2}{N|\mathcal{B}|}.$$

Observing that $\frac{1}{1 - (1 - \beta)^2} \leq \frac{1}{\beta}$, we can obtain the first conclusion in Lemma 8. Likewise, if we apply Eq. 27 to Eq. 30, the following relationship can be attained

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2] &\leq \frac{\mathbb{E}[\|\mathbf{u}_0 - \nabla \mathbf{J}(\mathbf{x}_0)\|^2]}{\beta} + \frac{12N(L^2 + G^2C_v)\eta^2}{\beta} \sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\ &+ \frac{24(L^2 + G^2C_v)}{\beta} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] + 2N\beta\bar{\sigma}^2K. \end{aligned} \quad (46)$$

As

$$\begin{aligned}\mathbb{E}[\|\mathbf{u}_0 - \nabla \mathbf{J}(\mathbf{x}_0)\|^2] &= \sum_{i=1}^N \mathbb{E}[\|\frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} (\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i))\|^2] \\ &= \frac{1}{|\mathcal{B}|^2} \sum_{i=1}^N \sum_{m=1}^{|\mathcal{B}|} \mathbb{E}[\|\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i)\|^2] \leq \frac{N\bar{\sigma}^2}{|\mathcal{B}|}.\end{aligned}\tag{47}$$

Substituting Eq. 47 into Eq. 46 completes the proof for the second conclusion in Lemma 8. \square

It is clearly observed from Eq. 43 that the error bound is dependent of the consensus error. To bound the consensus, we first present a lemma to bound the gradient tracking errors. As the initialization approach can be to use either a single trajectory or a mini-batch of trajectories, the lemma also consists of the explicit bound for the initial gradient tracking error. Before that, we present a fact that establishes the relationship between $\|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2$ and $\|\mathbf{x}_{k+1} - \Lambda \mathbf{x}_{k+1}\|^2$.

Fact 1. For all $k \geq 0$, we have

$$\begin{aligned}\|\mathbf{x}_{k+1} - \Lambda \mathbf{x}_{k+1}\|^2 &\leq \frac{1+\lambda^2}{2} \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 \\ &\quad + \frac{2\eta^2\lambda^2}{1-\lambda^2} \|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2,\end{aligned}\tag{48}$$

and

$$\begin{aligned}\|\mathbf{x}_{k+1} - \Lambda \mathbf{x}_{k+1}\|^2 &\leq 2\lambda^2 \|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2 \\ &\quad + 2\eta^2\lambda^2 \|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2.\end{aligned}\tag{49}$$

This fact can be easily shown true by using the update laws of MDPGT and the Young's inequality.

Lemma 9. Let \mathbf{v}_k be generated by MDPGT initialized with a mini-batch of trajectories \mathcal{B} . We have the following relationships:

$$\mathbb{E}[\|\mathbf{v}_1 - \Lambda \mathbf{v}_1\|^2] \leq \frac{N\bar{\sigma}^2}{|\mathcal{B}|} + \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2,\tag{50}$$

and if $0 < \eta \leq \frac{1-\lambda^2}{8\lambda\sqrt{9L^2+3G^2C_v}}$, then for all $k \leq 1$,

$$\begin{aligned}\mathbb{E}[\|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2] &\leq \frac{3+\lambda^2}{4} \mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] \\ &\quad + \frac{48L^2 + 24(L^2 + G^2C_v)}{1-\lambda^2} N\eta^2 \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] + 4N\beta^2\bar{\sigma}^2 \\ &\quad + \frac{114L^2 + 72(L^2 + G^2C_v)}{1-\lambda^2} \mathbb{E}[\|\mathbf{x}_{k-1} - \Lambda \mathbf{x}_{k-1}\|^2] \\ &\quad + \frac{10\beta^2}{1-\lambda^2} \mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2].\end{aligned}\tag{51}$$

Proof. According to the update law, we have

$$\mathbf{v}_1 - \Lambda \mathbf{v}_1 = \mathbf{W} \mathbf{v}_0 + \mathbf{u}_0 - \mathbf{u}_{-1} - \Lambda \mathbf{W} \mathbf{v}_0 - \Lambda \mathbf{u}_0 + \Lambda \mathbf{u}_{-1} = (\mathbf{I} - \Lambda) \mathbf{u}_0.\tag{52}$$

It follows from $\mathbf{v}_0 = \mathbf{0}$ and $\mathbf{u}_{-1} = \mathbf{0}$. Hence, we have

$$\mathbb{E}[\|\mathbf{v}_1 - \Lambda \mathbf{v}_1\|^2] = \mathbb{E}[\|(\mathbf{I} - \Lambda) \mathbf{u}_0\|^2] \leq \mathbb{E}[\|\mathbf{u}_0 - \nabla \mathbf{J}(\mathbf{x}_0) + \nabla \mathbf{J}(\mathbf{x}_0)\|^2]\tag{53}$$

Since $\mathbf{u}_0^i = \frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} \mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i)$, we then obtain

$$\begin{aligned}\mathbb{E}[\|\mathbf{u}_0 - \nabla \mathbf{J}(\mathbf{x}_0) + \nabla \mathbf{J}(\mathbf{x}_0)\|^2] &= \sum_{i=1}^N \mathbb{E}[\|\mathbf{u}_0^i - \nabla J_i(\mathbf{x}_0^i)\|^2] + \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 \\ &= \sum_{i=1}^N \mathbb{E}[\|\frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} (\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i))\|^2] + \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 \\ &= \frac{1}{|\mathcal{B}|^2} \sum_{i=1}^N \sum_{m=1}^{|\mathcal{B}|} \mathbb{E}[\|\mathbf{g}_i(\tau_0^{i,m} | \mathbf{x}_0^i) - \nabla J_i(\mathbf{x}_0^i)\|^2] + \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 \\ &\leq \frac{N\bar{\sigma}^2}{|\mathcal{B}|} + \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2,\end{aligned}\tag{54}$$

which completes the first part of the proof.

We now proceed to prove the second part. As

$$\begin{aligned}\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1} &= \underline{\mathbf{W}} \mathbf{v}_k + \mathbf{u}_k - \mathbf{u}_{k-1} - \Lambda(\underline{\mathbf{W}} \mathbf{v}_k + \mathbf{u}_k - \mathbf{u}_{k-1}) \\ &= (\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k + (\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1}),\end{aligned}\tag{55}$$

conducting the squared norm results in the following relationship

$$\begin{aligned}\|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2 &= \|(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k + (\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1})\|^2 \\ &= \|(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k\|^2 + 2\langle(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k, (\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1})\rangle \\ &\quad + \|(\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1})\|^2 \\ &\leq \lambda^2 \|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2 + 2\langle(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k, (\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1})\rangle + \|\mathbf{u}_k - \mathbf{u}_{k-1}\|^2.\end{aligned}\tag{56}$$

We next investigate the second term on the right hand side of the last inequality. Based on the update for the policy gradient surrogate, we have

$$\mathbf{u}_k^i = \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) + (1 - \beta) \mathbf{u}_{k-1}^i - (1 - \beta) v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\mathbf{x}_{k-1}^i).$$

Thus,

$$\begin{aligned}\mathbf{u}_k^i - \mathbf{u}_{k-1}^i &= \mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \beta \mathbf{u}_{k-1}^i - (1 - \beta) v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\mathbf{x}_{k-1}^i) \\ &= (1 - \beta)(\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - v_i(\tau_k^i | \mathbf{x}_{k-1}^i, \mathbf{x}_k^i) \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)) + \beta(\mathbf{g}_i(\tau_k^i | \mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i)) \\ &\quad + \beta(\mathbf{g}_i(\tau_k^i | \mathbf{x}_{k-1}^i) - \nabla J(\mathbf{x}_{k-1}^i) + \nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i).\end{aligned}\tag{57}$$

Hence, taking the expectation on both sides of the last equality leads to

$$\mathbb{E}[\mathbf{u}_k - \mathbf{u}_{k-1}] = (1 - \beta)(\nabla \mathbf{J}(\mathbf{x}_k) - \nabla \mathbf{J}(\mathbf{x}_{k-1})) + \beta \mathbb{E}[\mathbf{g}(\tau_k | \mathbf{x}_k) - \mathbf{g}(\tau_k | \mathbf{x}_{k-1})] - \beta(\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})).\tag{58}$$

We now investigate the term of $2\langle(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k, (\mathbf{I} - \Lambda)(\mathbf{u}_k - \mathbf{u}_{k-1})\rangle$ and denote it as B_k . Consequently, the following relationship is attained

$$\begin{aligned}\mathbb{E}[B_k] &= 2\langle(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k, (\mathbf{I} - \Lambda) \mathbb{E}[(\mathbf{u}_k - \mathbf{u}_{k-1})]\rangle \\ &= 2\langle(\underline{\mathbf{W}} - \Lambda) \mathbf{v}_k, (\mathbf{I} - \Lambda)[(1 - \beta)(\nabla \mathbf{J}(\mathbf{x}_k) - \nabla \mathbf{J}(\mathbf{x}_{k-1})) + \beta \mathbb{E}[\mathbf{g}(\tau_k | \mathbf{x}_k) - \mathbf{g}(\tau_k | \mathbf{x}_{k-1})] \\ &\quad - \beta(\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1}))]\rangle \\ &\leq 2\lambda \|\mathbf{v}_k - \Lambda \mathbf{v}_k\| \|(1 - \beta)(\nabla \mathbf{J}(\mathbf{x}_k) - \nabla \mathbf{J}(\mathbf{x}_{k-1})) + \beta \mathbb{E}[\mathbf{g}(\tau_k | \mathbf{x}_k) - \mathbf{g}(\tau_k | \mathbf{x}_{k-1})] \\ &\quad - \beta(\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1}))\| \\ &\leq \frac{1 - \lambda^2}{2} \|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2 + \frac{6\lambda^3(1 - \beta)^2}{1 - \lambda^2} \|\nabla \mathbf{J}(\mathbf{x}_k) - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2 \\ &\quad + \frac{6\lambda^3\beta^2}{1 - \lambda^2} \mathbb{E}[\|\mathbf{g}(\tau_k | \mathbf{x}_k) - \mathbf{g}(\tau_k | \mathbf{x}_{k-1})\|^2] + \frac{6\lambda^3\beta^2}{1 - \lambda^2} \|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2 \\ &\leq \frac{1 - \lambda^2}{2} \|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2 + (\beta^2 + (1 - \beta)^2) \frac{6\lambda^3 L^2}{1 - \lambda^2} \mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] + \frac{6\lambda^3\beta^2}{1 - \lambda^2} \|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2.\end{aligned}\tag{59}$$

The first inequality is due to Cauchy-Schwartz inequality, while the second inequality follows from basic inequalities $2ab \leq ea^2 + b^2/e$, with $e = \frac{1 - \lambda^2}{2\lambda}$ for all $a, b \in \mathbb{R}$ and $\|\mathbf{a} + \mathbf{b} + \mathbf{c}\|^2 \leq 3\|\mathbf{a}\|^2 + 3\|\mathbf{b}\|^2 + 3\|\mathbf{c}\|^2, \forall \mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^d$. The last inequality is

due to the smoothness property. Using Eq. 57 yields the upper bound of the second moment of $\mathbf{u}_k^i - \mathbf{u}_{k-1}^i$,

$$\begin{aligned}
\mathbb{E}[\|\mathbf{u}_k^i - \mathbf{u}_{k-1}^i\|^2] &\leq 4(1-\beta)^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - v_i(\tau_k^i|\mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] + 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - \nabla J(\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2] \\
&= 4(1-\beta)^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) + \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - v_i(\tau_k^i|\mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] + 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - \nabla J(\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2] \\
&\leq 4(1-\beta)^2\mathbb{E}[2\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2 + 2\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - v_i(\tau_k^i|\mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] + 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - \nabla J(\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2] \\
&= 8(1-\beta)^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] \\
&+ 8(1-\beta)^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - v_i(\tau_k^i|\mathbf{x}_{k-1}^i, \mathbf{x}_k^i)\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_k^i) - \mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i)\|^2] + 4\beta^2\mathbb{E}[\|\mathbf{g}_i(\tau_k^i|\mathbf{x}_{k-1}^i) - \nabla J(\mathbf{x}_{k-1}^i)\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2] \\
&\leq 8(1-\beta)^2L^2\mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + 8(1-\beta)^2G^2C_v\mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] \\
&+ 4\beta^2L^2\mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] + 4\beta^2\sigma_i^2 + 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2] \\
&= [8(1-\beta)^2L^2 + 8(1-\beta)^2G^2C_v + 4\beta^2L^2]\mathbb{E}[\|\mathbf{x}_k^i - \mathbf{x}_{k-1}^i\|^2] \\
&+ 4\beta^2\sigma_i^2 + 4\beta^2\mathbb{E}[\|\nabla J(\mathbf{x}_{k-1}^i) - \mathbf{u}_{k-1}^i\|^2].
\end{aligned} \tag{60}$$

The last inequality is based on the smoothness property, Lemma 1 and Lemma 2. We then have that

$$\begin{aligned}
\mathbb{E}[\|\mathbf{u}_k - \mathbf{u}_{k-1}\|^2] &\leq [8(1-\beta)^2L^2 + 8(1-\beta)^2G^2C_v + 4\beta^2L^2]\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] + \\
&4\beta^2N\bar{\sigma}^2 + 4\beta^2\mathbb{E}[\|\nabla \mathbf{J}(\mathbf{x}_{k-1} - \mathbf{u}_{k-1})\|^2].
\end{aligned} \tag{61}$$

Taking the expectation on both sides of Eq. 56 and substituting Eq. 59 and Eq. 61 into it, the following relationship can be obtained

$$\begin{aligned}
\mathbb{E}[\|\mathbf{v}_{k+1} - \Lambda \mathbf{v}_{k+1}\|^2] &\leq \lambda^2\mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] + \frac{1-\lambda^2}{2}\mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] \\
&+ (\beta^2 + (1-\beta)^2)\frac{6\lambda^3L^2}{1-\lambda^2}\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] \\
&+ \frac{6\lambda^3\beta^2}{1-\lambda^2}\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2] + 4\beta^2N\bar{\sigma}^2 \\
&+ [8(1-\beta)^2L^2 + 8(1-\beta)^2G^2C_v + 4\beta^2L^2]\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] \\
&+ 4\beta^2\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2] \\
&= \frac{1+\lambda^2}{2}\lambda^2\mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] \\
&+ \left[(\beta^2 + (1-\beta)^2)\frac{6\lambda^3L^2}{1-\lambda^2} + 8(1-\beta)^2L^2 + 8(1-\beta)^2G^2C_v + 4\beta^2L^2\right]\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] \\
&+ \left(4\beta^2 + \frac{6\lambda^3\beta^2}{1-\lambda^2}\right)\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2] + 4\beta^2N\bar{\sigma}^2 \\
&\leq \frac{1+\lambda^2}{2}\lambda^2\mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] + \left(\frac{12\lambda^3L^2}{1-\lambda^2} + 8(L^2 + G^2C_v) + 4L^2\right)\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] \\
&+ 4N\beta^2\bar{\sigma}^2 + \frac{10\beta^2}{1-\lambda^2}\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2] \\
&\leq \frac{1+\lambda^2}{2}\lambda^2\mathbb{E}[\|\mathbf{v}_k - \Lambda \mathbf{v}_k\|^2] + \frac{16L^2 + 8(L^2 + G^2C_v)}{1-\lambda^2}\mathbb{E}[\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2] + 4\beta^2N\bar{\sigma}^2 \\
&+ \frac{10\beta^2}{1-\lambda^2}\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla \mathbf{J}(\mathbf{x}_{k-1})\|^2].
\end{aligned} \tag{62}$$

Since

$$\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 \leq 3\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2 + 3N\eta^2\|\bar{\mathbf{u}}_{k-1}\|^2 + 3\|\mathbf{x}_{k-1} - \Lambda\mathbf{x}_{k-1}\|^2,$$

combining Fact 1, the following is obtained

$$\|\mathbf{x}_k - \mathbf{x}_{k-1}\|^2 \leq 6\lambda^2\eta^4\|\mathbf{v}_k - \Lambda\mathbf{v}_k\|^2 + 3N\eta^2\|\bar{\mathbf{u}}_{k-1}\|^2 + 9\|\mathbf{x}_{k-1} - \Lambda\mathbf{x}_{k-1}\|^2. \quad (63)$$

Combining the last two inequalities, one can attain

$$\begin{aligned} \mathbb{E}[\|\mathbf{v}_{k+1} - \Lambda\mathbf{v}_{k+1}\|^2] &\leq \left(\frac{1+\lambda^2}{2} + \frac{96L^2 + 48(L^2 + G^2C_v)}{1-\lambda^2}\lambda^2\eta^2 \right) \mathbb{E}[\|\mathbf{v}_k - \Lambda\mathbf{v}_k\|^2] \\ &+ \frac{48L^2 + 24(L^2 + G^2C_v)}{1-\lambda^2}N\eta^2\mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] + \frac{144L^2 + 72(L^2 + G^2C_v)}{1-\lambda^2}\mathbb{E}[\|\mathbf{x}_{k-1} - \Lambda\mathbf{x}_{k-1}\|^2] \\ &+ 4\beta^2N\bar{\sigma}^2 + \frac{10\beta^2}{1-\lambda^2}\mathbb{E}[\|\mathbf{u}_{k-1} - \nabla\mathbf{J}(\mathbf{x}_{k-1})\|^2]. \end{aligned} \quad (64)$$

Combining the fact that $0 < \eta \leq \frac{1-\lambda^2}{8\lambda\sqrt{9L^2+3G^2C_v}}$ yields the desirable result. \square

According to Lemma 8 and Lemma 9, the remaining step to obtain the explicitly accurate error bound is to inaugurate the correlation between $\sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2]$ and $\sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2]$. Thus, the following lemma is constructed for this purpose.

Lemma 10. *Let \mathbf{x}_k be generated by MDPGT initialized with a mini-batch of trajectories \mathcal{B} . If $0 < \eta \leq \frac{(1-\lambda^2)^2}{\lambda\sqrt{14592L^2+9984G^2C_v}}$ and $\beta \in (0, 1)$, then for $K \geq 2$, we have*

$$\begin{aligned} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] &\leq \left[\frac{1436L^2 + 4608(L^2 + G^2C_v)}{(1-\lambda^2)^4} \right] \\ &+ \lambda^2N\eta^4 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{32\lambda^2N\bar{\sigma}^2\eta^2}{(1-\lambda^2)^3|\mathcal{B}|} \left(1 + \frac{10\beta}{1-\lambda^2} \right) \\ &+ \frac{32\lambda^2\eta^2}{(1-\lambda^2)^3} \|\nabla\mathbf{J}(\mathbf{x}_0)\|^2 + \frac{128\lambda^2N\beta^2K\bar{\sigma}^2\eta^2}{(1-\lambda^2)^3} \left(1 + \frac{5\beta}{1-\lambda^2} \right). \end{aligned} \quad (65)$$

Proof. Applying Eq. 27 to the fact that

$$\|\mathbf{x}_{k+1} - \Lambda\mathbf{x}_{k+1}\|^2 \leq \frac{1+\lambda^2}{2}\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2 + \frac{2\eta^2\lambda^2}{1-\lambda^2}\|\mathbf{v}_{k+1} - \Lambda\mathbf{v}_{k+1}\|^2$$

leads to

$$\sum_{k=0}^K \|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2 \leq \frac{4\lambda^2\eta^2}{(1-\lambda^2)^2} \sum_{k=1}^K \|\mathbf{v}_k - \Lambda\mathbf{v}_k\|^2.$$

Similarly, applying Eq. 28 to Eq. 51 results in

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}[\|\mathbf{v}_k - \Lambda\mathbf{v}_k\|^2] &\leq \frac{4}{1-\lambda^2}\mathbb{E}[\|\mathbf{v}_1 - \Lambda\mathbf{v}_1\|^2] + \frac{192L^2 + 96(L^2 + G^2C_v)}{1-\lambda^2}N\eta^2 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\ &+ \frac{576L^2 + 288(L^2 + G^2C_v)}{(1-\lambda^2)^2} \sum_{k=0}^{K-2} \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] + \frac{40\beta^2}{(1-\lambda^2)^2} \sum_{k=0}^{K-2} \mathbb{E}[\|\mathbf{u}_k - \nabla\mathbf{J}(\mathbf{x}_k)\|^2] \\ &+ \frac{16\beta^2N\bar{\sigma}^2K}{1-\lambda^2} \\ &\leq \frac{4N\bar{\sigma}^2}{(1-\lambda^2)|\mathcal{B}|} + \frac{4\|\nabla\mathbf{J}(\mathbf{x}_0)\|^2}{1-\lambda^2} + \frac{192L^2 + 96(L^2 + G^2C_v)}{1-\lambda^2}N\eta^2 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_{k-1}\|^2] \\ &+ \frac{576L^2 + 288(L^2 + G^2C_v)}{(1-\lambda^2)^2} \sum_{k=0}^{K-2} \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] + \frac{40\beta^2}{(1-\lambda^2)^2} \sum_{k=0}^{K-2} \mathbb{E}[\|\mathbf{u}_k - \nabla\mathbf{J}(\mathbf{x}_k)\|^2] \\ &+ \frac{16\beta^2N\bar{\sigma}^2K}{1-\lambda^2}. \end{aligned} \quad (66)$$

The last inequality follows from Lemma 9. As

$$\begin{aligned}
& \frac{40\beta^2}{(1-\lambda^2)^2} \sum_{k=0}^{K-2} \mathbb{E}[\|\mathbf{u}_k - \nabla \mathbf{J}(\mathbf{x}_k)\|^2] \leq \frac{40N\bar{\sigma}^2\beta}{(1-\lambda^2)^2|\mathcal{B}|} + \frac{480\beta N(L^2 + G^2C_v)\eta^2}{(1-\lambda^2)^2} \\
& \quad \sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{960\beta(L^2 + G^2C_v)}{(1-\lambda^2)^2} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda(\mathbf{x}_k)\|^2] \\
& \quad + \frac{80N\beta^3\bar{\sigma}^2K}{(1-\lambda^2)^2},
\end{aligned} \tag{67}$$

we then have

$$\begin{aligned}
\sum_{k=1}^K \mathbb{E}[\|\mathbf{v}_k - \Lambda\mathbf{v}_k\|^2] & \leq \frac{192L^2 + 576(L^2 + G^2C_v)}{(1-\lambda^2)^2} N\eta^2 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\
& \quad + \frac{576L^2 + 1248(L^2 + G^2C_v)}{(1-\lambda^2)^2} \sum_{k=0}^{K-1} \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] \\
& \quad \frac{4N\bar{\sigma}^2}{(1-\lambda^2)|\mathcal{B}|} \left(1 + \frac{10\beta}{1-\lambda^2}\right) + \frac{16\beta^2N\bar{\sigma}^2K}{1-\lambda^2} \left(1 + \frac{5\beta}{1-\lambda^2}\right) + \frac{4\|\nabla \mathbf{J}(\mathbf{x}_0)\|^2}{1-\lambda^2}.
\end{aligned} \tag{68}$$

Thus, we have

$$\begin{aligned}
\sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] & \leq \frac{768L^2 + 2304(L^2 + G^2C_v)\lambda^2N\eta^4}{(1-\lambda^2)^4} \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\
& \quad + \frac{2304L^2 + 4992(L^2 + G^2C_v)}{(1-\lambda^2)^4} \lambda^2\eta^2 \sum_{k=0}^{K-1} \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] + \frac{16N\bar{\sigma}^2\lambda^2\eta^2}{(1-\lambda^2)^3|\mathcal{B}|} \left(1 + \frac{10\beta}{1-\lambda^2}\right) \\
& \quad + \frac{64\beta^2N\bar{\sigma}^2K\eta^2\lambda^2}{(1-\lambda^2)^3} \left(1 + \frac{5\beta}{1-\lambda^2}\right) + \frac{16\lambda^2\eta^2\|\nabla \mathbf{J}(\mathbf{x}_0)\|^2}{(1-\lambda^2)^3}.
\end{aligned} \tag{69}$$

With simple mathematical manipulations, based on the conditions $0 < \eta \leq \frac{(1-\lambda^2)^2}{\lambda\sqrt{14592L^2+9984G^2C_v}}$ and $\beta \in (0, 1)$, it is obtained that

$$1 - \frac{2304L^2 + 4992(L^2 + G^2C_v)}{(1-\lambda^2)^4} \lambda^2\eta^2 \leq \frac{1}{2}.$$

Hence the proof is completed by adopting this inequality. \square

Hence, with Lemma 10, it suffices to show Theorem 1, whose proof is presented next. While for Theorem 2, the same proof ideas can be applied and the corresponding results are obtained by setting $|\mathcal{B}| = 1$. Thus, we are not going to repeat statements for the auxiliary lemmas, instead using the conclusions from the lemmas and adjusting slightly the constants in the error bounds.

Proof of Theorem 1

Proof. Recall the conclusion of Lemma 5, we have

$$\sum_{k=0}^K \|\nabla J(\bar{\mathbf{x}}_k)\|^2 \leq \frac{2(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{2} \sum_{k=0}^K \|\bar{\mathbf{u}}_k\|^2 + 2 \sum_{k=0}^K \|\bar{\mathbf{u}}_k - \bar{\nabla \mathbf{J}}(\mathbf{x}_k)\|^2 + \frac{2L^2}{N} \sum_{k=0}^K \|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2.$$

Substituting Eq. 43 into the last inequality yields

$$\begin{aligned}
\sum_{k=0}^K \|\nabla J(\bar{\mathbf{x}}_k)\|^2 & \leq \frac{2(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{2L^2}{N} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] \\
& \quad + \frac{2\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{24(L^2 + G^2C_v)\eta^2}{N\beta} \sum_{k=0}^{K-1} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{48(L^2 + G^2C_v)}{N^2\beta} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] + \frac{4\beta\bar{\sigma}^2K}{N} \\
& = \frac{2(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{4} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{2}{N} \left(L^2 + \frac{24(L^2 + G^2C_v)}{N\beta} \right) \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda\mathbf{x}_k\|^2] \\
& \quad + \frac{2\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{4\beta\bar{\sigma}^2K}{N} - \left(\frac{1}{4} - \frac{24(L^2 + G^2C_v)\eta^2}{N\beta} \right) \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2].
\end{aligned} \tag{70}$$

To get rid of the last term on the right hand side of the above inequality, the following relationship can be obtained

$$\frac{1}{4} - \frac{24(L^2 + G^2 C_v) \eta^2}{N \beta} \geq 0 \Rightarrow \frac{96(L^2 + G^2 C_v) \eta^2}{N} \leq \beta < 1 \Rightarrow 0 < \eta < \frac{1}{6\sqrt{6(L^2 + G^2 C_v)}}.$$

It is easily to verified the above relationship as we have that $\beta = \frac{96(L^2 + G^2 C_v) \eta^2}{N}$. Thus,

$$\begin{aligned} \sum_{k=0}^K \|\nabla J(\bar{\mathbf{x}}_k)\|^2 &\leq \frac{2(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{4} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{2\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{4\beta\bar{\sigma}^2 K}{N} \\ &+ \frac{2}{N} \left(L^2 + \frac{24(L^2 + G^2 C_v)}{N\beta} \right) \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2]. \end{aligned} \quad (71)$$

As

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \sum_{k=0}^K \mathbb{E}[\|\nabla J(\mathbf{x}_k^i)\|^2] &\leq \frac{2}{N} \sum_{i=1}^N \sum_{k=0}^K \mathbb{E}[\|\nabla J(\mathbf{x}_k^i) - \nabla J(\bar{\mathbf{x}}_k)\|^2 + \|\nabla J(\bar{\mathbf{x}}_k)\|^2] \\ &\leq \frac{2L^2}{N} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] + 2 \sum_{k=0}^K \mathbb{E}[\|\nabla J(\bar{\mathbf{x}}_k)\|^2], \end{aligned} \quad (72)$$

we have

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \sum_{k=0}^K \mathbb{E}[\|\nabla J(\mathbf{x}_k^i)\|^2] &\leq \frac{2L^2}{N} \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] + \frac{4(J^* - J(\bar{\mathbf{x}}_0))}{\eta} \\ &- \frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{4\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{8\beta\bar{\sigma}^2 K}{N} + \frac{4}{N} \left(L^2 + \frac{24(L^2 + G^2 C_v)}{N\beta} \right) \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] \\ &= \frac{6}{N} \left(L^2 + \frac{16(L^2 + G^2 C_v)}{N\beta} \right) \sum_{k=0}^K \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] + \frac{4(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\ &+ \frac{4\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{8\beta\bar{\sigma}^2 K}{N}. \end{aligned} \quad (73)$$

As $\frac{6}{N} \left(L^2 + \frac{16(L^2 + G^2 C_v)}{N\beta} \right) = \frac{6}{N} \left(L^2 + \frac{1}{6\eta^2 L^2} \right)$ and $\eta^2 L^2 \leq \eta^2 (L^2 + G^2 C_v) \leq \frac{1}{96}$ based on $\eta < \frac{1}{6\sqrt{6(L^2 + G^2 C_v)}}$, we have

$$\frac{6}{N} \left(L^2 + \frac{16(L^2 + G^2 C_v)}{N\beta} \right) < \frac{17}{16N\eta^2}.$$

Thus, the following relationship can be attained

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \sum_{k=0}^K \mathbb{E}[\|\nabla J(\mathbf{x}_k^i)\|^2] &\leq \frac{17}{16\eta^2 N} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] + \frac{4(J^* - J(\bar{\mathbf{x}}_0))}{\eta} - \frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\ &+ \frac{4\bar{\sigma}^2}{N\beta|\mathcal{B}|} + \frac{8\beta\bar{\sigma}^2 K}{N}. \end{aligned} \quad (74)$$

We investigate the combined term $-\frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{17}{16\eta^2 N} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2]$. Since

$$\begin{aligned}
& -\frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{17}{16\eta^2 N} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] \leq -\frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\
& + \frac{17}{16\eta^2 N} \left[\left(\frac{1436L^2 + 4608(L^2 + G^2C_v)}{(1-\lambda^2)^4} \right) \lambda^2 N \eta^4 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{32\lambda^2 N \bar{\sigma}^2 \eta^2}{(1-\lambda^2)^3 |\mathcal{B}|} \right. \\
& \left. \left(1 + \frac{10\beta}{1-\lambda^2} \right) + \frac{32\lambda^2 \eta^2}{(1-\lambda^2)^3} \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 + \frac{128\lambda^2 N \beta^2 K \bar{\sigma}^2 \eta^2}{(1-\lambda^2)^3} \left(1 + \frac{5\beta}{1-\lambda^2} \right) \right] \\
& = -\frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{1526L^2 + 4896(L^2 + G^2C_v)}{(1-\lambda^2)^4} \lambda^2 \eta^2 \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\
& + \frac{34\lambda^2}{N(1-\lambda^2)^3} \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 + \frac{34\lambda^2 \bar{\sigma}^2}{(1-\lambda^2)^3 |\mathcal{B}|} \left(1 + \frac{10\beta}{1-\lambda^2} \right) + \frac{136\lambda^2 \beta^2 K \bar{\sigma}^2}{(1-\lambda^2)^3} \left(1 + \frac{5\beta}{1-\lambda^2} \right) \\
& = -\frac{1}{2} \left(1 - \frac{3052L^2 + 9792(L^2 + G^2C_v)}{(1-\lambda^2)^4} \lambda^2 \eta^2 \right) \sum_{k=0}^{K-2} \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] \\
& + \frac{34\lambda^2}{N(1-\lambda^2)^3} \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 + \frac{34\lambda^2 \bar{\sigma}^2}{(1-\lambda^2)^3 |\mathcal{B}|} \left(1 + \frac{10\beta}{1-\lambda^2} \right) + \frac{136\lambda^2 \beta^2 K \bar{\sigma}^2}{(1-\lambda^2)^3} \left(1 + \frac{5\beta}{1-\lambda^2} \right).
\end{aligned} \tag{75}$$

With the condition that $\eta \leq \frac{(1-\lambda^2)^2}{\lambda \sqrt{12844L^2 + 9792G^2C_v}}$, it is immediately that

$$1 - \frac{3052L^2 + 9792(L^2 + G^2C_v)}{(1-\lambda^2)^4} \lambda^2 \eta^2 \geq 0$$

such that the first term on the right hand side of Eq. 75 can be removed. Moreover, since $\eta \leq \frac{\sqrt{N(1-\lambda^2)}\lambda}{31\sqrt{L^2 + G^2C_v}}$, which leads to $\beta \leq \frac{(1-\lambda^2)\lambda^2}{10}$ such that

$$1 + \frac{10\beta}{1-\lambda^2} < 2, \quad 1 + \frac{5\beta}{1-\lambda^2} < \frac{3}{2}.$$

By now we can conclude that

$$\begin{aligned}
& -\frac{1}{2} \sum_{k=0}^K \mathbb{E}[\|\bar{\mathbf{u}}_k\|^2] + \frac{17}{16\eta^2 N} \mathbb{E}[\|\mathbf{x}_k - \Lambda \mathbf{x}_k\|^2] \leq \frac{34\lambda^2}{N(1-\lambda^2)^3} \|\nabla \mathbf{J}(\mathbf{x}_0)\|^2 + \frac{68\lambda^2 \bar{\sigma}^2}{(1-\lambda^2)^3 |\mathcal{B}|} \\
& + \frac{204\lambda^2 \beta^2 K \bar{\sigma}^2}{(1-\lambda^2)^3}.
\end{aligned} \tag{76}$$

Substituting Eq. 76 into Eq. 74 and dividing both sides by $\frac{1}{K+1}$, and with $\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] = \frac{1}{N(K+1)} \sum_{i=1}^N \sum_{k=0}^K \mathbb{E}[\|\nabla J(\mathbf{x}_k^i)\|^2]$ completes the proof. \square

Proof of Corollary 1

Proof. Substituting η, β , and $|\mathcal{B}|$ into the Eq. 9 and conducting some simple mathematical manipulations can easily attain the desirable result. \square

Analysis for MDPGT with Single Trajectory Initialization

Theorem 2. *Let Assumptions 1,2 and 3 hold. Let the momentum coefficient $\beta = \frac{96L^2 + 96G^2C_v}{N} \eta^2$. If MDPGT is initialized by a single trajectory and the step size satisfies the following condition*

$$0 < \eta \leq \min \left\{ \frac{(1-\lambda^2)^2}{\lambda \sqrt{12844L^2 + 9792G^2C_v}}, \frac{\sqrt{N(1-\lambda^2)}\lambda}{31\sqrt{L^2 + G^2C_v}}, \frac{1}{6\sqrt{6(L^2 + G^2C_v)}} \right\}, \tag{77}$$

then the output $\tilde{\mathbf{x}}_K$ satisfies: for all $K \geq 2$:

$$\begin{aligned}\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] &\leq \frac{4(J^* - J(\bar{\mathbf{x}}_0))}{\eta K} + \frac{4\bar{\sigma}^2}{N\beta K} + \frac{8\beta\bar{\sigma}^2}{N} \\ &+ \frac{34\lambda^2}{KN(1-\lambda^2)^3} \|\nabla \mathbf{J}(\bar{\mathbf{x}}_0)\|^2 + \frac{68\lambda^2\bar{\sigma}^2}{(1-\lambda^2)^3 K} \\ &+ \frac{204\lambda^2\beta^2\bar{\sigma}^2}{(1-\lambda^2)^3},\end{aligned}\tag{78}$$

where J^* is the upper bound of $J(\mathbf{x})$ and $\|\nabla \mathbf{J}(\bar{\mathbf{x}}_0)\|^2 \triangleq \sum_{i=1}^N \|\nabla J_i(\bar{\mathbf{x}}_0)\|^2$.

Proof. As the initialization is single trajectory, $|\mathcal{B}| = 1$. Following the similar proof techniques as already shown in Theorem 1 yields the desirable result. The proof process is not repeated in this context. \square

Similarly, in view of Theorem 2, the asymptotic behavior is the same as in Theorem 1, having the same steady-state error. Due to an infinite time horizon, regardless of the initialization strategy, diverse agents are able to learn effectively in a collaborative manner. However, in terms of the non-asymptotic property, the single trajectory initialization strategy makes a difference in the sampling complexity, which is reflected by the following result.

Corollary 3. Let $\eta = \frac{N^{3/4}}{8LK^{1/4}}$, $\beta = \frac{DN^{1/2}}{64L^2K^{1/2}}$, in Theorem 2. We have

$$\begin{aligned}\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] &\leq \frac{32L(J^* - J(\bar{\mathbf{x}}_0))}{(NK)^{3/4}} + \frac{2048L^4\bar{\sigma}^2 + D^2\bar{\sigma}^2N}{8L^2DN^{3/2}K^{1/2}} \\ &+ \frac{\lambda^2}{(1-\lambda^2)^3 K} \left(\frac{34\|\nabla \mathbf{J}(\bar{\mathbf{x}}_0)\|^2}{N} + 68\bar{\sigma}^2 + \frac{51\bar{\sigma}^2D^2N}{1024L^4} \right),\end{aligned}\tag{79}$$

for all

$$K \geq \max \left\{ \frac{N^3D^2}{4096L^4}, \frac{923521N(L^2 + G^2C_v)^2}{4096L^4(1-\lambda^2)^2\lambda^4}, \frac{(12844L^2 + 9792G^2C_v)^2\lambda^4N^3}{4096L^4(1-\lambda^2)^8} \right\},\tag{80}$$

where $D = 96L^2 + 96G^2C_v$.

Proof. Substituting η, β into the Eq. 78 and conducting some simple mathematical manipulations can easily attain the desirable result. \square

Corollary 3 implies that with only a single trajectory initialization for MDPGT, when K is sufficiently large, the mean-squared convergence rate is

$$\mathbb{E}[\|\nabla J(\tilde{\mathbf{x}}_K)\|^2] \leq \mathcal{O}\left(\frac{1}{(NK)^{1/2}}\right),$$

which is slower than that obtained in Corollary 1. With similar mathematical manipulation, the eventual sampling complexity is $\mathcal{O}(N^{-1}\epsilon^{-4})$. Though variance reduction techniques has not reduced the order of ϵ^{-1} , compared to the SOTA approaches, the linear speed up still enables the complexity to be N times smaller than that in (Xu, Gao, and Gu 2019; Huang et al. 2020). Additionally, different from traditional decentralized learning problems, MARL has more significant variances in the optimization procedure due to the non-oblivious characteristic. Using just a single trajectory for each agent to initialize is a quite poor scheme, but the adopted variance reduction techniques can successfully maintain the SOTA sampling complexity in a decentralized setting.

Implication for Gaussian Policy

In this section, we study the sample complexity when the policy function $\pi^i(a^i|s)$ of each agent is explicitly a Gaussian distribution. For a bounded action space $\mathcal{A}^i \subset \mathbb{R}$, a Gaussian policy parameterized by \mathbf{x}_i is defined as

$$\pi^i(a^i|s) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{((\mathbf{x}^i)^\top \phi_i(s) - a^i)^2}{2\xi^2}\right),\tag{81}$$

where ξ^2 is a constant standard deviation parameter and $\phi_i(s) : \mathcal{S} \rightarrow \mathbb{R}^{d_i}$ is mapping from the state space to the feature space. Note that the standard deviation parameter can be varying in terms of different agents, while for simplicity in this context, we assume that it is the same for each agent. Next, we verify the assumptions on the Gaussian policy. We first impose a

mild assumption that the action space and feature space are bounded, i.e., there exist constants $C_a > 0$ and $C_f > 0$ such that for all $i \in \mathcal{V}$, $|a^i| \leq C_a$, $\forall a^i \in \mathcal{A}^i$ and $\|\phi_i(s)\| \leq C_f$, $\forall s \in \mathcal{S}$. Then we can show that Assumption 1 hold. Due to the limit of space, we defer the derivation to the supplementary materials. While Assumption 3 does not always hold for all Gaussian distributions, but based on a result from (Cortes, Mansour, and Mohri 2010), it has been shown that for two Gaussian distributions, $\pi_{\mathbf{x}_1}^i(a^i|s) = \mathcal{N}(\mu_1, \xi_1^2)$ and $\pi_{\mathbf{x}_2}^i(a^i|s) = \mathcal{N}(\mu_2, \xi_2^2)$, if $\xi_2 > \frac{\sqrt{2}}{2}\xi_1$, then the variance of the importance sampling weight $v_i(\tau^i|\mathbf{x}_1, \mathbf{x}_2)$ is bounded. Due to a constant standard deviation parameter defined for the Gaussian distributions, it is easily verified that for any time step $k \geq 0$, $\mathbb{V}(v_i(\tau_k^i|\mathbf{x}_{k-1}^i, \mathbf{x}_k^i))$ should be bounded by some constant $\mathcal{M} > 0$. We recall Corollary 1 that can apply to any general policy models. Therefore, based on the above discussion, it also applies to the Gaussian policy function scenario. We present the result in the following corollary to explicitly state the relationship between γ and ϵ in the sampling complexity. We first present an established result that bounds the variance of policy gradient estimator $\mathbf{g}_i(\tau^i|\mathbf{x}^i)$ in Eq. 6 when the policy function π^i follows a Gaussian distribution.

Lemma 11. (Lemma 5.5 in (Pirotta, Restelli, and Bascetta 2013)) *Given a Gaussian policy $\pi_{\mathbf{x}^i}^i = \mathcal{N}((\mathbf{x}^i)^\top \phi_i(s), \xi^2)$, for all $i \in \mathcal{V}$, if $|r_i(s, a^i)| \leq R$ and $\|\phi^i(s)\| \leq C_f$ for all $s \in \mathcal{S}$, $a^i \in \mathcal{A}^i$ and $R > 0$, $C_f > 0$ are constants, then the variance of policy gradient estimator $\mathbb{V}(\mathbf{g}_i(\tau^i|\mathbf{x}^i))$ can be bounded as*

$$\mathbb{V}(\mathbf{g}_i(\tau^i|\mathbf{x}^i)) \leq \frac{R^2 C_f^2}{(1-\gamma)^2 \xi^2} \left(\frac{1-\gamma^{2H}}{1-\gamma^2} - H\gamma^{2H} - 2\gamma^H \frac{1-\gamma^H}{1-\gamma} \right). \quad (82)$$

In this context, we have defined a constant standard deviation parameter for the policy of each agent to simplify the analysis. However, one can still define separately this parameter for each agent, i.e., ξ_i . Then in the upper bound of last inequality, $\xi^2 = \min\{\xi_1^2, \xi_2^2, \dots, \xi_N^2\}$. We show the proof for Corollary 2 in the following.

Proof. According to Corollary 1, when K is sufficiently large, T_1 dominates the convergence such that we investigate it. According to Lemma 11, we can immediately obtain that

$$\bar{\sigma}^2 \leq \frac{R^2 C_f^2}{(1-\gamma)^2 \xi^2} \left(\frac{1-\gamma^{2H}}{1-\gamma^2} - H\gamma^{2H} - 2\gamma^H \frac{1-\gamma^H}{1-\gamma} \right) = \mathcal{O}\left(\frac{1}{(1-\gamma)^3}\right).$$

As $L = \frac{C_h R}{(1-\gamma)^2}$ and $G = \frac{C_g R}{(1-\gamma)^2}$, we have $D = 96L^2 + 96G^2C_v = \mathcal{O}\left(\frac{1}{(1-\gamma)^4}\right)$. Thus,

$$\frac{256L^3D(J^* - J(\bar{\mathbf{x}}_0)) + 2048L^4\bar{\sigma}^2 + D^2\bar{\sigma}^2}{8L^2D(NK)^{2/3}} \leq \mathcal{O}\left(\frac{1}{(1-\gamma)^3(NK)^{2/3}}\right).$$

We can easily obtain that the sampling complexity is $\mathcal{O}\left(\frac{1}{(1-\gamma)^{4.5}N\epsilon^3}\right)$ when the policy is parameterized by a Gaussian distribution. \square

Remark 2. *When the policy function is a Gaussian distribution for each agent, Corollary 2 shows that the sampling complexity is inherently in the order of $\mathcal{O}(\epsilon^{-3})$, which matches the generalized result. Moreover, the linear speed up is still retained. Additionally, we also show that the complexity does not rely on the horizon. The dependence on $(1-\gamma)^{-4.5}$ stems from the variance of the policy gradient estimator, which has been known in the centralized counterpart (Xu, Gao, and Gu 2019). When the initialization is with a single trajectory, it can be inferred from Corollary 3 that the sampling complexity is $\mathcal{O}((1-\gamma)^{-6}N^{-1}\epsilon^{-4})$.*

Algorithmic Framework for MDPG

We next present a new decentralized algorithm adapted from the centralized MBPG (Huang et al. 2020), which serves as a new baseline for empirical comparison. Different from MBPG, the updates are simplified with constant step size and β value. This simplification is to enable a fair comparison among all algorithms.

Algorithm 2: Momentum-based Decentralized Policy Gradient (MDPG)

Result: $\tilde{\mathbf{x}}_K$ chosen uniformly random from $\{\mathbf{x}_k^i, i \in \mathcal{V}\}_{k=1}^K$

Input: $\mathbf{x}_1^i = \bar{\mathbf{x}}_1 \in \mathbb{R}^d, \eta \in \mathbb{R}^+, \beta \in (0, 1), \mathbf{W} \in \mathbb{R}^{N \times N}, K, \mathcal{B} \in \mathbb{Z}^+, k = 1$

while $k < K$ **do**

- for** each agent **do**

 - if $k = 1$: compute the local policy gradient surrogate by sampling a trajectory τ_1^i from $p_i(\tau^i | \mathbf{x}_1^i) : \mathbf{u}_1^i = \mathbf{g}_i(\tau_1^i | \mathbf{x}_1^i)$, or by sampling a *mini-batch* of trajectories $\{\tau_1^{i,m}\}_{m=1}^{|\mathcal{B}|}$ from $p_i(\tau^i | \mathbf{x}_1^i) : \mathbf{u}_1^i = \frac{1}{|\mathcal{B}|} \sum_{m=1}^{|\mathcal{B}|} \mathbf{g}_i(\tau_1^{i,m} | \mathbf{x}_1^i)$
 - if $k > 1$: sample a trajectory τ_k^i from $p_i(\tau^i | \mathbf{x}_k^i)$ and compute the local policy gradient surrogate using Eq. 8
 - Update the local estimate of the policy network parameters $\mathbf{x}_{k+1}^i = \sum_{j \in Nb(i)} \omega_{ij} (\mathbf{x}_k^j + \eta \mathbf{u}_k^j)$

- end**

$k = k + 1$

end

Computing Resource Details

All numerical experiments presented in the main article and supplementary section were performed on a computing cluster with 2 x 16 Core Intel Xeon CPU E5-2620 v4 and a memory of 256 GB. All codes were implemented using PyTorch version 1.8.1.

Additional Experimental Results

In the following section, we include additional results for comparing MDPGT, MDPG and DPG in a gridworld environment with two and three agents. Additionally, we also provide results for five, ten, twenty and thirty agents in a simplified lineworld environment.

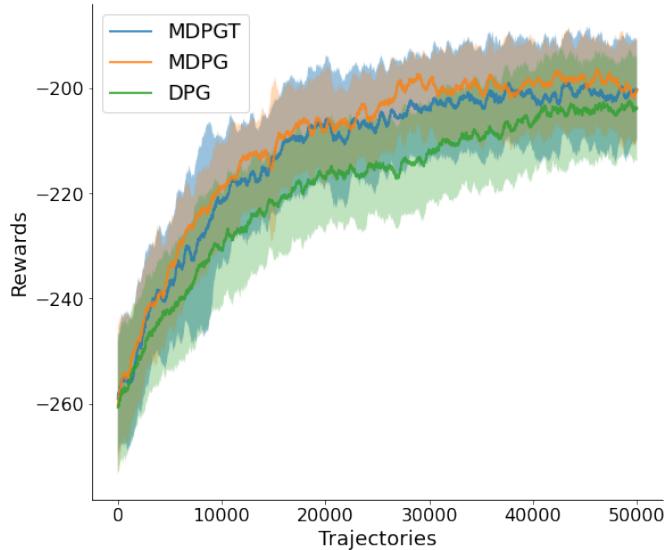


Figure 4: Experimental results comparing MDPGT, MDPG and DPG in the gridworld environment with 2 agents using fully-connected topology.

Gridworld Environment Figures 4, 5, and 6 illustrate the average reward obtained by the three algorithms, MDPGT, MDPG and DPG in with two, three and ten agents respectively using a fully-connected network topology with $\beta = 0.5$. It should be noted that in Figure 6 the number of trajectories is only 25,000 such that MDPGT and MDPG perform similarly. However, when the number is 50,000, it suffices to show that the result should resemble the plot in Figure 1, with even a larger gap between MDPGT and MDPG. In all scenarios, we observed that both MDPGT and MDPG still outperforms DPG. Interestingly, with fewer agents, the performance of MDPGT and MDPG are similar, in contrast to the five agent scenario presented in the main article. Fewer agents may reduce the impact of linear speed up on the error bound, which is w.r.t $\mathcal{O}(N^{-1})$. By observing carefully, when the number of agents becomes larger, MDPGT is more advantageous in a complex graph than MDPG as the extra tracking step contributes to correcting the gradient bias caused by diverse agents. This interesting finding will be validated in the following simplified lineworld environment where more agents are incorporated.

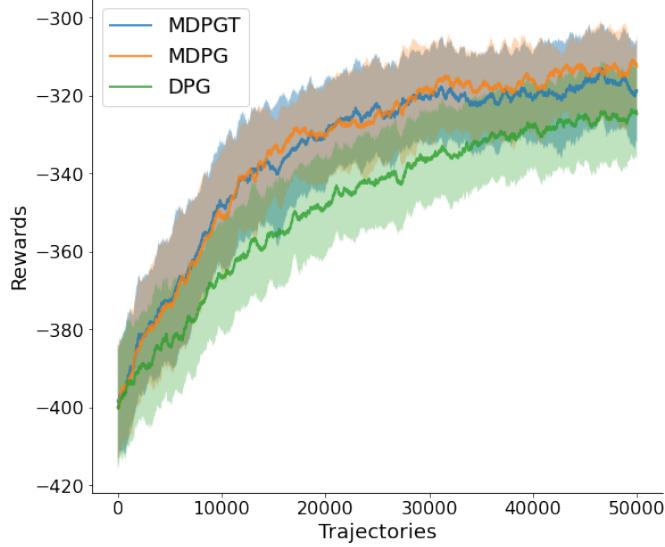


Figure 5: Experimental results comparing MDPGT, MDPG and DPG in the gridworld environment with 3 agents using fully-connected topology.

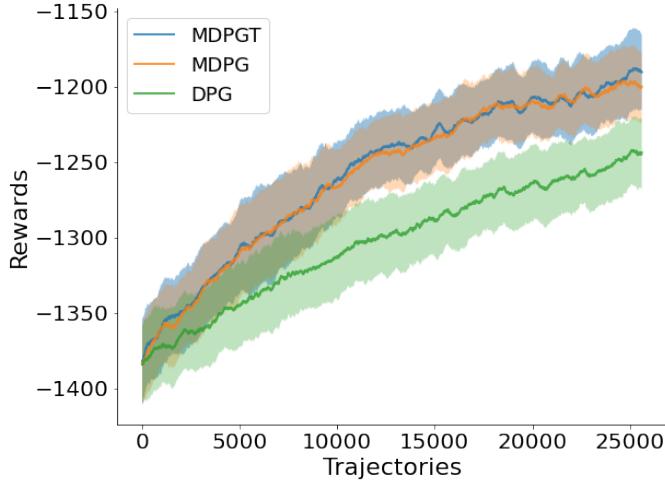


Figure 6: Experimental results comparing MDPGT, MDPG and DPG in the gridworld environment with 10 agents using fully-connected topology.

Lineworld Environment As an initial proof of concept and to demonstrate our algorithm’s ability to scale to a larger number of agents, we created a simplification of the gridworld environment called the Lineworld environment. In this environment, all agents are randomly initialized with a 1-D coordinate. The goal of all the agents are to cooperatively arrive at the coordinate **0**. In this setting, the agents are not allowed to collide and have a smaller action set of either moving up, down or stay stationary. The reward function we used for this environment is also defined as the Euclidean distance of individual agent to its respective goal. All agent’s policy is represented by a 3-layer neural network with 64 hidden units with *tanh* activation functions. The agents were trained for 10,000 episodes with a horizon of 500 steps and discount factor of 0.99. A learning rate of 3E-4 was used in all our experiments and we average the results conducted over 5 random seeds.

Figures 7 and 8 shows the rewards of our proposed methods and the baseline DPG in the Lineworld environment with five, ten, twenty and thirty agents respectively. In summary, the findings from the performance in the Lineworld environment strengthens the conclusion attained from the Gridworld setting that MDPGT outperforms MDPG in a more sophisticated graphs involving more agents. With five agents, we observed that MDPGT performs better than DPG but worse than MDPG. However, as the number of agent increases, we see that the performance of MDPGT improves to be on-par or slightly better than MDPG and significantly better than DPG. Thus, a practical guideline from an application point-of-view is that one can choose MDPG for simple graphs

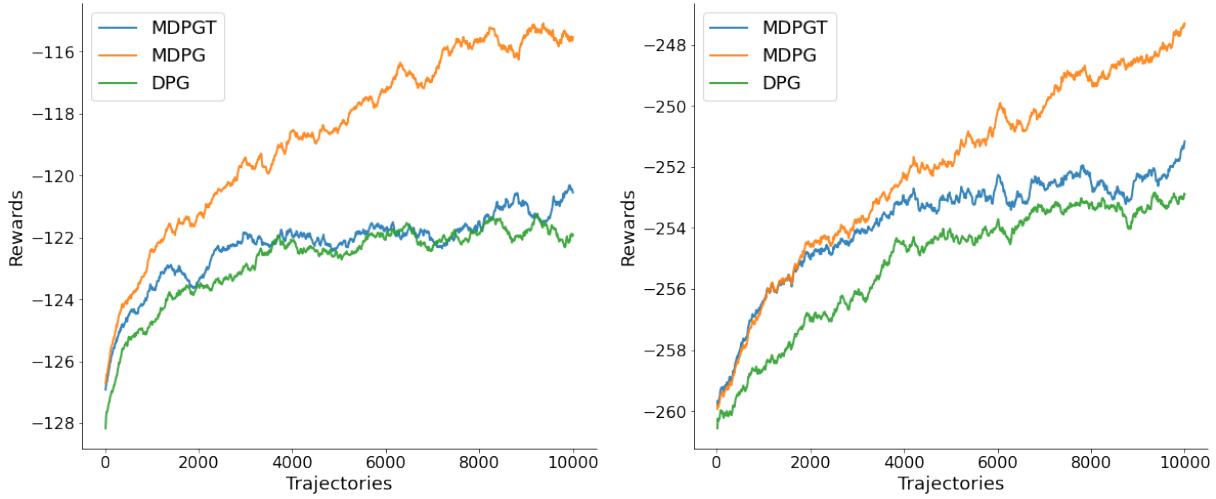


Figure 7: Experimental results comparing MDPGT, MDPG and DPG in the lineworld environment with 5 agents (left) and 10 agents (right) using fully-connected topology.

and expect MDPGT to perform much better in complex graphs.

Minibatch Initialization To validate the theoretical findings obtained in this work for the difference between two different versions of MDPGT, we implement several experiments with 5 agents using fully-connected topology with different mini-batch initializations. In this context, mini-batch size equal to 1 corresponds to MDPGT with only single trajectory initialization. From Figure 9, one can observe that the reward curves due to different mini-batch initializations are depicted in three phases. Before approximately 20,000 trajectories, regardless of whichever initialization, MDPGT yields similar performance. After 20,000 trajectories, the difference becomes appealing, which is reflected by the outperforming of MDPGT with mini-batch size equal to 4. This supports the claim of Corollary 1 that K has to satisfy a certain condition to start showing the non-asymptotic behavior. Additionally, when the number of trajectories is larger than 60,000, another interesting phenomenon is that all reward curves start converging together again, which can be explained by using Theorem 1 as MDPGT enables the asymptotic results once K is sufficiently larger than a certain number. We have pointed out in the analysis that they all converge to the same steady-state error defined in Eq. 10. Intuitively, practitioners can set a sufficiently large number of trajectories to get the same optimal solution regardless of what initialization is taken. While different *approximate* solutions can also be obtained when different specific mini-batch sizes for initialization are defined. One may argue how to set the mini-batch size for a specific problem. For example, we can observe from Figure 9 that to get better approximate solutions, the mini-batch size could be selected in between 4 and 16. This is also suggested from Corollary 1 that the mini-batch size is w.r.t both the number of agents and the number of trajectories. Thus, either mini-batch size equal to 64 or 256 results in slightly worse result than only one trajectory initialization.

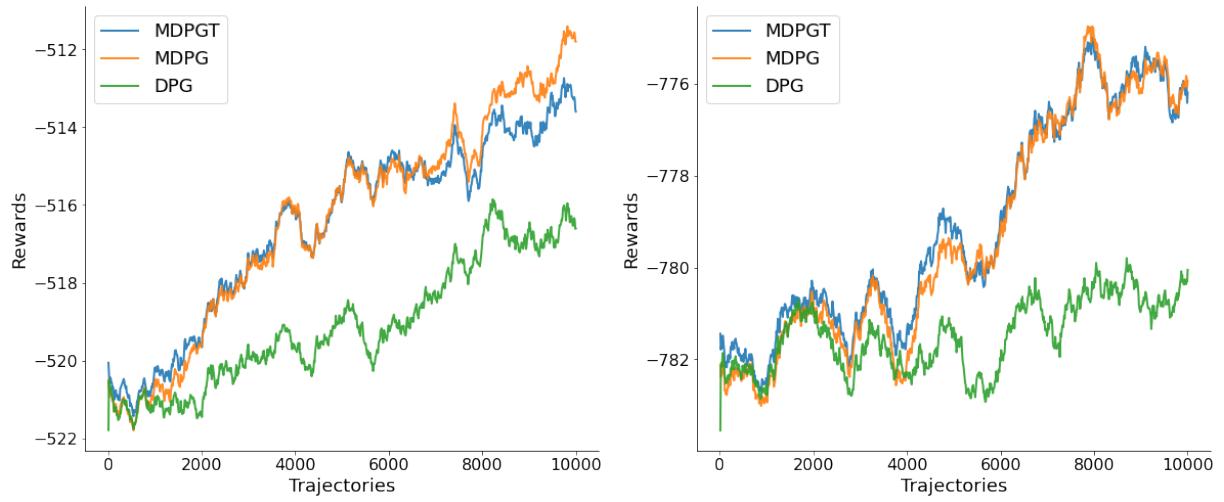


Figure 8: Experimental results comparing MDPGT, MDPG and DPG in the lineworld environment with 20 agents (left) and 30 agents (right) using fully-connected topology.

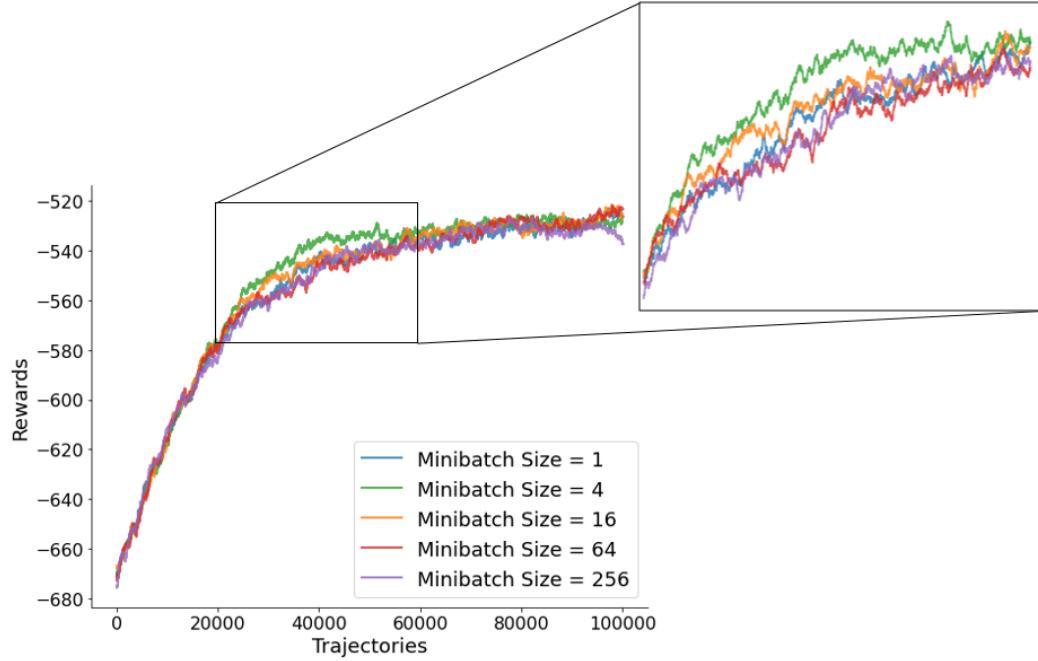


Figure 9: Experimental results comparing MDPGT in the gridworld environment with 5 agents and $\beta = 0.5$ using fully-connected topology with different mini-batch initializations.