# Multi-Agent Reinforcement Learning in Stochastic Networked Systems

Anonymous Author(s) Affiliation Address email

### Abstract

We study multi-agent reinforcement learning (MARL) in a stochastic network of 1 agents. The objective is to find localized policies that maximize the (discounted) 2 global reward. In general, scalability is a challenge in this setting because the size 3 of the global state/action space can be exponential in the number of agents. Scalable 4 algorithms are only known in cases where dependencies are static, fixed and local, 5 e.g., between neighbors in a fixed, time-invariant underlying graph. In this work, 6 we propose a Scalable Actor Critic framework that applies in settings where the 7 dependencies can be non-local and stochastic, and provide a finite-time error bound 8 that shows how the convergence rate depends on the speed of information spread 9 in the network. Additionally, as a byproduct of our analysis, we obtain novel 10 11 finite-time convergence results for a general stochastic approximation scheme and for temporal difference learning with state aggregation, which apply beyond the 12 setting of MARL in networked systems. 13

### 14 **1 Introduction**

<sup>15</sup> Multi-Agent Reinforcement Learning (MARL) has achieved impressive performance in a wide <sup>16</sup> array of applications including multi-player game play [41, 31], multi-robot systems [13], and <sup>17</sup> autonomous driving [25]. In comparison to single-agent reinforcement learning (RL), MARL poses <sup>18</sup> many challenges, chief of which is scalability [56]. Even if each agent's local state/action spaces are <sup>19</sup> small, the size of the global state/action space can be large, potentially exponentially large in the <sup>20</sup> number of agents, which renders many RL algorithms such as *Q*-learning not applicable.

A promising approach for addressing the scalability challenge that has received attention in recent 21 years is to exploit application-specific structures, e.g., [18, 35, 38]. A particularly important example 22 of such a structure is a networked structure, e.g., applications in multi-agent networked systems 23 such as social networks [7, 27], communication networks [59, 50], queueing networks [34], and 24 smart transportation networks [58]. In these networked systems, it is often possible to exploit *static*, 25 local dependency structures [16, 17, 1, 32], e.g., the fact that agents only interact with a fixed set of 26 neighboring agents throughout the game. This sort of dependency structure often leads to scalable, 27 distributed algorithms for optimization and control 16 1. 32, and has proven effective for designing 28 scalable and distributed MARL algorithms, e.g. [35, 38]. 29 However, many real-world networked systems have inherently *time-varying*, non-local dependencies. 30

For example, in the context of wireless networks, each node can send packets to other nodes within a fixed transmission range. However, the interference range, in which other nodes can interfere the transmission, can be larger than the transmission range [52]. As a result, due to potential collisions, the local reward of each node not only depends on its own local state/action, but also depends on the actions of other nodes within the interference range, which may be more than one-hop away. In

<sup>36</sup> addition, a node may be able to observe other nodes' local states before picking its local action 33.

37 Things become even more complex when mobility and stochastic network conditions are considered.

<sup>38</sup> These lead to dependencies that are both stochastic and non-local. Although one can always fix and

localize the dependence model, this leads to considerably reduced performance. Beyond wireless
 networks, similar stochastic and non-local dependencies exists in epidemics [30], social networks

41 [7, 27], and smart transportation networks [58].

A challenging open question in MARL is to understand how to obtain algorithms that are scalable in 42 settings where the dependencies are stochastic and non-local. Prior work considers exclusively static 43 and local dependencies, e.g., [35, 38]. It is clear that hardness results apply when the dependencies 44 are too general [24]. Further, results in the static, local setting to this point rely on the concept of 45 exponential decay [35, 16], meaning the agents' impact on each other decays exponentially in their 46 graph distance. This property relies on the fact that the dependencies are purely local and static, and 47 it is not clear whether it can still be exploited when the interactions are more general. This motivates 48 an important open question: Is it possible to design scalable algorithms for stochastic, non-local 49 networked MARL? 50

**Contributions.** In this paper, we introduce a class of stochastic, non-local dependency structures 51 where every agent is allowed to depend on a random subset of agents. In this context, we propose 52 and analyze a Scalable Actor Critic (SAC) algorithm that provably learns a near-optimal local policy 53 in a scalable manner (Theorem D.2). This result represents the *first* provably scalable method for 54 stochastic networked MARL. Key to our approach is that the class of dependencies we consider leads 55 to a  $\mu$ -decay property (Definition 4.1). This property generalizes the exponential decay property 56 underlying recent results such as [35, 16], which does not apply to stochastic non-local dependencies, 57 and enables the design of an efficient and scalable algorithm for settings with stochastic, non-local 58 dependencies. Our analysis of the algorithm reveals an important trade-off: as deeper interactions 59 appear more frequently, the "information" can spread more quickly from one part of the network to 60 another, which leads to the efficiency of the proposed method to degrade. This is to be expected, 61 as when the agents are allowed to interact globally, the problem becomes a single-agent tabular 62 Q-learning problem with an exponentially large state space, which is known to be intractable since 63 the sample complexity is polynomial in the size of the state/action space [12, 24]. 64

The key technical result underlying our analysis of the Scalable Actor Critic algorithm is a finite-time 65 analysis of a general stochastic approximation scheme featuring infinite-norm contraction and state 66 aggregation (Theorem 2.1). We apply this result to networked MARL using the local neighborhood of 67 each agent to provide state aggregation (SA). This result also applies beyond MARL. Specifically, we 68 show that it yields finite-time bounds on Temporal Difference (TD)/Q learning with state aggregation 69 (Theorem 3.1). To the best of our knowledge the resulting bound is the first finite-time bound on 70 asynchronous Q-learning with state aggregation. Additionally, it yields a novel analysis for TD-71 learning with state aggregation (the first error bound in the infinity norm) that sheds new insight 72 into how the error depends on the quality of state abstraction. These two results are important 73 contributions in their own right. Due to space constraints, we discuss asynchronous Q-learning with 74 state aggregation in Appendix C.4. 75

**Related literature.** The prior work that is most related to our paper is [38], which also studies 76 MARL in a networked setting. The key difference is that we allow the dependency structure among 77 agents to be non-local and stochastic, while 38 requires the dependency structure to be local and 78 static. The generality of setting means techniques from [38] do not apply and adds considerable 79 complexity to the proof in two aspects. First, instead of analyzing the algorithm directly like [38]. 80 we derive a finite-time error bound for TD learning with state aggregation (Section 2 and 3), and 81 then establish its connection with the algorithm (Section 4.3). Second, we need a more general decay 82 property (Definition 4.1) than the exponential one used in [38]. Defining and establishing this general 83 decay property for the non-local and stochastic setting is highly non-trivial (Section 4.1). 84

More broadly, MARL has received considerable attention in recent years, see 56 for a survey. The 85 86 line of work most relevant to the current paper focuses on cooperative MARL. In the cooperative setting, each agent can decide its local actions but share a common global state with other agents. 87 The objective is to maximize a global reward by working cooperatively. Notable examples of this 88 approach include [6] 10] and the references therein. In contrast, we study a situation where each 89 agent has its own state that it acts upon. Despite the differences, like our situation, cooperative 90 MARL problems still face scalability issues since the joint-action space is exponentially large. A 91 variety of methods have been proposed to deal with this, including independent learners [8, 29], 92

where each agent employs a single-agent RL policy. Function approximation is another approach
that can significantly reduce the space/computational complexity. One can use linear functions
[57] or neural networks [28] in the approximation. A limitation of these approaches is the lack of
theoretical guarantees on the approximation error. In contrast, our technique not only reduces the
space/computational complexity significantly, but also has theoretical guarantees on the performance
loss in settings with stochastic and non-local dependencies.

The mean-field approach [44, 55, 19] provides another way to address the scalability issue, but under very different settings compared to ours. Specifically, the mean-field approach typically assumes homogeneous agents with identical local state/action space and policies, and each agent depends on other agents through their population or "mean" behavior. In contrast, our approach considers a local-interaction model, where there is an underlying graph and each agent depends on neighboring agents in the graph. Further, our approach allows heterogeneous agents, which means that the local state/action spaces and policies can differ among the agents.

Another related line of work uses centralized training with decentralized execution, e.g., [28, 15], where there is a centralized coordinator that can communicate with all the agents and keep track of their experiences and policies. In contrast, our work only requires distributed training, where we constrain the scale of communication in training within the  $\kappa$ -hop neighborhood of each agent.

More broadly, this paper contributes to a growing literature that uses exponential decay to derive scalable algorithms for learning in networked systems. The specific form of exponential decay that we generalize is related to the idea of "correlation decay" studied in [16, 17], though their focus is on solving static combinatorial optimization problems whereas ours is on learning policies in dynamic environments. Most related to the current paper is [38], which shows an exponential decay property in a restricted networked MARL model with purely local dependencies. In contrast, we show a more general  $\mu$ -decay property holds for a general form of stochastic, non-local dependencies.

The technical work in this paper contributes to the analysis of stochastic approximation (SA), which 117 has received considerable attention over the past decade [53, 43, 11, 54]. Our work is most related 118 to [37], which uses an asynchronous nonlinear SA to study the finite-time convergence rate for 119 asynchronous Q-learning on a single trajectory. Beyond [37], there are many other works that use 120 SA schemes to study TD learning and Q-learning, e.g. [43, 51, 20]. The finite-time error bound for 121 TD learning with state aggregation in our work is most related to the asymptotic convergence limit 122 given in [48] and the application of SA scheme to asynchronous Q-learning in [37]. Beyond these 123 papers, other related work in the broader area of RL with state aggregation includes [26, 23, 22, 9, 42]. 124 We add to this literature with a novel finite-time convergence bound for a general SA with state 125 aggregation. This result, in turn, yields the first finite-time error bound in the infinity norm for both 126 TD learning with state aggregation and Q-learning with state aggregation. 127

### **128 2** Stochastic Approximation

In this section, we present the key technical innovation underlying our results on MARL: a new
finite-time analysis of a general asynchronous stochastic approximation (SA) scheme. This analysis
underlies our approach for MARL in networked systems (presented in Section 4). Further, this SA
scheme is of interest more broadly, e.g., to the settings of TD learning with state aggregation (Section
and asynchronous *Q*-learning with state aggregation (Appendix C.4).

Consider a finite-state Markov chain whose state space is given by  $\mathcal{N} = \{1, 2, \dots, n\}$ . Let  $\{i_t\}_{t=0}^{\infty}$ be the sequence of states visited by this Markov chain. Our focus is generalizing the following asynchronous stochastic approximation (SA) scheme, which is studied in [47] 40, 51]: Let parameter  $x \in \mathbb{R}^{\mathcal{N}}$ , and  $F : \mathbb{R}^{\mathcal{N}} \to \mathbb{R}^{\mathcal{N}}$  be a  $\gamma$ -contraction in the infinity norm. The update rule of the SA scheme is given by

$$\begin{aligned} x_{i_t}(t+1) &= x_{i_t}(t) + \alpha_t (F_{i_t}(x(t)) - x_{i_t}(t) + w(t)), \\ x_j(t+1) &= x_j(t) \text{ for } j \neq i_t, j \in \mathcal{N}, \end{aligned} \tag{1}$$

where w(t) is a noise sequence. It is shown in [37] that parameter x(t) converges to the unique fixed point of F at the rate of  $O(1/\sqrt{t})$ .

While general, in many cases, including networked MARL, we do not wish to calculate an entry for every state in  $\mathcal{N}$  in parameter x, but instead, wish to calculate "aggregated entries." Specifically, at each time step, after  $i_t$  is generated, we use a surjection h to decide which dimension of parameter xshould be updated. This technique, referred to as state aggregation, is one of the easiest-to-deploy schemes for state space compression in the RL literature [21, 42]. In the generalized SA scheme, our

objective is to specify the convergence point as well as obtain a finite-time error bound.

Formally, to define the generalization of (1), let  $\mathcal{N} = \{1, \dots, n\}$  be the state space of  $\{i_t\}$  and  $\mathcal{M} = \{1, \dots, m\}, (m \leq n)$  be the *abstract* state space. The surjection  $h : \mathcal{N} \to \mathcal{M}$  is used to convert every state in  $\mathcal{N}$  to its abstraction in  $\mathcal{M}$ . Given parameter  $x \in \mathbb{R}^{\mathcal{M}}$  and function  $F : \mathbb{R}^{\mathcal{N}} \to \mathbb{R}^{\mathcal{N}}$ , we consider the generalized SA scheme that updates  $x(t) \in \mathbb{R}^{\mathcal{M}}$  starting from  $x(0) = \mathbf{0}$ ,

$$x_{h(i_t)}(t+1) = x_{h(i_t)}(t) + \alpha_t \big( F_{i_t}(\Phi x(t)) - x_{h(i_t)}(t) + w(t) \big), x_j(t+1) = x_j(t) \text{ for } j \neq h(i_t), j \in \mathcal{M},$$
(2)

where the feature matrix  $\Phi \in \mathbb{R}^{\mathcal{N} \times \mathcal{M}}$  is defined as

$$\Phi_{ij} = \begin{cases} 1 & \text{if } h(i) = j \\ 0 & \text{otherwise} \end{cases}, \forall i \in \mathcal{N}, j \in \mathcal{M}.$$
(3)

In order to state our main result characterizing the convergence of (2), we must first state a few definitions and assumptions. To begin, we define the weighted infinity norm as in [37], except that we extend its definition so as to define the contraction of function F. The reason we use the weighted infinity norm as opposed to the standard infinity norm is that its generality can be used in certain settings for undiscounted RL, as shown in [47, 2].

**Definition 2.1** (Weighted Infinity Norm). *Fix a positive vector*  $v \in \mathbb{R}^{\mathcal{M}}$ . *For*  $x \in \mathbb{R}^{\mathcal{M}}$ , we define  $\|x\|_{v} := \sup_{i \in \mathcal{M}} \frac{|x_{i}|}{v_{i}}$ . *For*  $x \in \mathbb{R}^{\mathcal{N}}$ , we define  $\|x\|_{v} := \sup_{i \in \mathcal{N}} \frac{|x_{i}|}{v_{h(i)}}$ .

Next, we state our assumption on the mixing rate of the Markov chain  $\{i_t\}$ , which is common in the literature [49, 43]. It holds for any finite-state Markov chain which is aperiodic and irreducible [5].

Assumption 2.1 (Stationary Distribution and Geometric Mixing Rate).  $\{i_t\}$  is an aperiodic and irreducible Markov chain on state space  $\mathcal{N}$  with stationary distribution  $d = (d_1, d_2, \dots, d_n)$ . Let  $d'_j = \sum_{i \in h^{-1}(j)} d_i$  and  $\sigma' = \inf_{j \in \mathcal{M}} d'_j$ . There exists positive constants  $K_1, K_2$  which satisfy that sup\_ $S \subseteq \mathcal{N} \left| \sum_{i \in S} d_i - \sum_{i \in S} \mathbb{P}(i_t = i \mid i_0 = j) \right| \leq K_1 \exp(-t/K_2), \forall j \in \mathcal{N}, \forall t \geq 0$  and  $K_2 \geq 1$ .

Our next assumption ensures contraction of F. It is also standard, e.g., [47, 51, 37], and ensures that F has a unique fixed point  $y^*$ .

**Assumption 2.2** (Contraction). Operator F is a  $\gamma$  contraction in  $\|\cdot\|_v$ , i.e., for any  $x, y \in \mathbb{R}^N$ , we have  $\|F(x) - F(y)\|_v \leq \gamma \|x - y\|_v$ . Further, there exists some constant C > 0 such that for any  $x \in \mathbb{R}^N$ , we have  $\|F(x)\|_v \leq \gamma \|x\|_v + C$ .

In Assumption 2.2, notice that the first sentence directly implies the second with  $C = (1 + \gamma) ||y^*||_v$ , where  $y^* \in \mathbb{R}^N$  is the unique fixed point of F. Further, while Assumption 2.2 implies that F has a unique fixed point  $y^*$ , we do not expect our stochastic approximation scheme to converge to it. Instead, we show that the convergence is to the unique  $x^*$  that solves

$$\Pi F(\Phi x^*) = x^*, \text{ where } \Pi := (\Phi^\top D \Phi)^{-1} \Phi^\top D.$$
(4)

Here  $D = diag(d_1, d_2, \dots, d_n)$  denotes the steady-state probabilities for the process  $\{i_t\}$ . Note that  $x^*$  is well-defined because the operator  $\Pi F(\Phi \cdot)$ , which defines a mapping from  $\mathbb{R}^{\mathcal{M}}$  to  $\mathbb{R}^{\mathcal{M}}$ , is also a contraction in  $\|\cdot\|_{v}$ . We state and prove this as Proposition B.1 in Appendix B.1.

Our last assumption is on the noise sequence w(t). It is also standard, e.g., [40, 37].

Assumption 2.3 (Martingale Difference Sequence).  $w_t$  is  $\mathcal{F}_{t+1}$  measurable and satisfies  $\mathbb{E}w(t) \mid \mathcal{F}_t = 0$ . Further,  $|w(t)| \leq \bar{w}$  almost surely for constant  $\bar{w}$ .

181 We are now ready to state our finite-time convergence result for stochastic approximation.

**Theorem 2.1.** Suppose Assumptions 2.1 2.2 2.3 hold. Further, assume there exists constant  $\bar{x} \ge \|x^*\|_v$  such that  $\forall t, \|x(t)\|_v \le \bar{x}$  almost surely. Let the step size be  $\alpha_t = \frac{H}{t+t_0}$  with  $t_0 =$ 

<sup>&</sup>lt;sup>1</sup>The assumption on  $\bar{x}$  follows from Assumptions 2.2 and 2.3 See Proposition B.2 in Appendix B.3

max(4H, 2K<sub>2</sub> log T), and  $H \ge \frac{2}{\sigma'(1-\gamma)}$ . Let  $x^*$  be the unique solution of equation  $\Pi F(\Phi x^*) = x^*$ , and define constants  $C_1 := 2\bar{x} + C + \frac{w}{\underline{v}}, C_2 := 4\bar{x} + 2C + \frac{\bar{w}}{\underline{v}}, C_3 := 2K_1(2\bar{x}+C)(1+2K_2+4H).$ Then, with probability at least  $1 - \delta$ ,

$$||x(T) - x^*||_v \le \frac{C_a}{\sqrt{T+t_0}} + \frac{C_a'}{T+t_0} = \tilde{O}\left(\frac{1}{\sqrt{T}}\right),$$

187 where the constants are given by  $C_a = \frac{4HC_2}{1-\gamma}\sqrt{K_2\log T\left(\log\left(\frac{4mK_2T}{\delta}\right) + \log\log T\right)}$  and  $C'_a = 4\max\left\{\frac{48K_2C_1H\log T + \sigma'C_3}{(1-\gamma)\sigma'}, \frac{2\bar{x}(2K_2\log T + t_0)}{1-\gamma}\right\}.$ 

A proof of Theorem 2.1 can be found in Appendix B.2. Compared with Theorem 4 in [37], Theorem 189 2.1 holds for a more general SA scheme where state aggregation is used to reduce the dimension 190 of the parameter x. The proof technique used in [37] does not apply to our setting because our 191 stationary point  $x^*$  has a more complex form (3). To do the generalization, we need to use a different 192 error decomposition method compared to [37] that leverages the stationary distribution D rather than 193 the distribution of  $i_t$  condition on  $i_{t-\tau}$  (see Appendix B.2 for details). Because of this generality, 194 Theorem 2.1 requires a stronger but standard assumption on the mixing rate of the Markov chain 195  $\{i_t\}.$ 196

### 197 **3** State Aggregation

Before applying our analysis of SA (Theorem 2.1) in the network setting, we first illustrate its 198 importance via a simpler application to the cases of TD-learning and Q-learning with state aggregation. 199 Understanding state aggregation methods is a foundational goal of analysis in the RL literature and it 200 has been studied in many previous works, e.g., [26, 23, 22, 9, 42]. Further, the result is extremely 201 useful in the analysis in networked MARL that follows since the  $\mu$ -decay property we introduce 202 (Definition 4.1) provides a natural state aggregation in the network setting (see Corollary 4.4). Due 203 to space constraints, in this section we only introduce the results on TD-learning; the results on 204 Q-learning are given in Appendix C.4. 205

In TD learning with state aggregation [42, 48], given the sequence of states visited by the Markov chain is  $\{i_t\}$ , the update rule of TD(0) is given by

$$\theta_{h(i_t)}(t+1) = \theta_{h(i_t)}(t) + \alpha_t \big( r_t + \gamma \theta_{h(i_{t+1})}(t) - \theta_{h(i_t)}(t) \big), \theta_j(t+1) = \theta_j(t) \text{ for } j \neq h(i_t), j \in \mathcal{M},$$
(5)

where  $h : \mathcal{N} \to \mathcal{M}$  is a surjection that maps each state in  $\mathcal{N}$  to an abstract state in  $\mathcal{M}$  and  $r_t$  is the reward at time step t such that  $\mathbb{E}[r_t] = r(i_t, i_{t+1})$ .

Taking F as the Bellman Policy Operator, i.e., the *i*'th dimension of function F is given by

$$F_i(V) = \mathbb{E}_{i' \sim \mathbb{P}(\cdot|i)}[r(i,i') + \gamma V_{i'}], \forall i \in \mathcal{N}, V \in \mathbb{R}^{\mathcal{N}}$$

The value function (vector)  $V^*$  is defined as  $V_i^* = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(i_t, i_{t+1}) \mid i_0 = i], i \in \mathcal{N}$  [48]. By defining the feature matrix  $\Phi$  as (3) and the noise sequence as

$$w(t) = r_t + \gamma \theta_{h(i_{t+1})}(t) - \mathbb{E}_{i' \sim \mathbb{P}(\cdot|i_t)}[r(i_t, i') + \gamma \theta_{h(i')}(t)],$$

we can rewrite the update rule of TD(0) in (5) in the form of an SA scheme (2). Therefore, we can apply Theorem 2.1 to obtain a finite-time error bound for TD learning with state aggregation. A proof of Theorem 3.1 can be found in Appendix C.2.

**Theorem 3.1.** Let Assumption 2.1 hold for the Markov chain  $\{i_t\}$  and let the stage reward  $r_t$  be upper bounded by  $\bar{r}$  almost surely. Assume that if h(i) = h(i') for  $i, i' \in \mathcal{N}$ , we have  $|V_i^* - V_{i'}^*| \leq \zeta$ for a constant  $\zeta$ . Consider TD(0) with the step size  $\alpha_t = \frac{H}{t+t_0}$ , where  $t_0 = \max(4H, 2K_2 \log T)$  and  $H \geq \frac{2}{\sigma'(1-\gamma)}$ . Define constant  $C_4 := 4K_1(1+2K_2+4H)$ . Then, with probability at least  $1-\delta$ ,

$$\|\Phi \cdot \theta(T) - V^*\|_{\infty} \le \frac{C_a}{\sqrt{T+t_0}} + \frac{C'_a}{T+t_0} + \frac{\zeta}{1-\gamma},$$

where the constants are given by  $C_a = \frac{40H\bar{r}}{(1-\gamma)^2}\sqrt{K_2\log T\left(\log\left(\frac{4mK_2T}{\delta}\right) + \log\log T\right)}$  and  $C'_a = \frac{8\bar{r}}{(1-\gamma)^2}\max\left\{\frac{144K_2H\log T}{\sigma'} + C_4, 2K_2\log T + t_0\right\}.$ 

The most related prior results to Theorem 3.1 are [43, 4]. In contrast to these, Theorem 3.1 considers the infinity norm, which is more natural for measuring error when using state aggregation. Further, our analysis is different and extends to the case of *Q*-learning with state aggregation (see Appendix C.4), where we obtain the first finite-time error bound. Moreover, unlike [4], our TD-learning algorithm does not require a projection step.

### 227 4 Networked MARL

We now present our main results, which apply the results in the previous sections to a stochastic 228 networked MARL setting. We consider a network of agents that are associated with an underlying 229 undirected graph  $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ , where  $\mathcal{N} = \{1, 2, \cdots, n\}$  denotes the set of agents and  $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ 230 denotes the set of edges. The distance  $d_{\mathcal{G}}(i, j)$  between two agents i and j is defined as the number 231 of edges on the shortest path that connects them on graph  $\mathcal{G}$ . Each agent is associated with its local 232 state  $s_i \in S_i$  and local action  $a_i \in A_i$  where  $S_i$  and  $A_i$  are finite sets. The global state/action is 233 defined as the combination of all local states/actions, i.e.,  $s = (s_1, \dots, s_n) \in S := S_1 \times \dots \times S_n$ , 234 and  $a = (a_1, \dots, a_n) \in \mathcal{A} := \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ . We use  $N_i^{\kappa}$  to denote the  $\kappa$ -hop neighborhood of agent i on  $\mathcal{G}$ , i.e.,  $N_i^{\kappa} := \{j \in \mathcal{N} \mid d_{\mathcal{G}}(i, j) \leq \kappa\}$ . Let  $f(\kappa) := \sup_i |N_i^{\kappa}|$ . For a subset  $M \subseteq \mathcal{N}$ , we 235 236 use  $s_M/a_M$  to denote the tuple formed by the states/actions of agents in M. 237

Before we define the transitions and rewards, we first define the notion of active link sets, which are 238 directed graphs on the agents  $\mathcal{N}$  and they characterize the interaction structure among the agents. 239 More specifically, an active link set is a set of directed edges that contains all self-loops, i.e., a subset 240 of  $\mathcal{N} \times \mathcal{N}$  and a super set of  $\{(i, i) \mid i \in \mathcal{N}\}$ . Generally speaking,  $(j, i) \in L$  means agent j can affect 241 agent i in the active link set L. Given an active link set L, we also use  $N_i(L) := \{j \in \mathcal{N} \mid (j,i) \in L\}$ 242 to denote the set of all agents (include itself) who can affect agent i in the active link set L. In this 243 paper, we consider a pair of active link sets  $(L_t^s, L_t^r)$  that is independently drawn from some joint distribution  $\mathcal{D}$  at each time step  $t_t^{\mathcal{D}}$  where the distribution  $\mathcal{D}$  will be defined using the underlying 244 245 graph  $\mathcal{G}$  later in Section 4.1. The role of  $L_t^s/L_t^r$  is that they define the dependence structure of state 246 transition/reward at time t, which we detail below. 247

Transitions. At time t, given the current state, action s(t), a(t) and the active link set  $L_t^s$ , the next individual state  $s_i(t+1)$  is independently generated and only depends on the state/action of the agents in  $N_i(L_t^s)$ . In other words, we have,

$$P(s(t+1)|s(t), a(t), L_t^s) = \prod_{i \in \mathcal{N}} P_i(s_i(t+1)|s_{N_i(L_t^s)}(t), a_{N_i(L_t^s)}(t), L_t^s).$$
(6)

*Rewards.* Each agent is associated with a local reward function  $r_i$ . At time t, it is a function of  $L_t^r$ and the state/action of agents in  $N_i(L_t^r)$ :  $r_i(L_t^r, s_{N_i(L_t^r)}(t), a_{N_i(L_t^r)}(t))$ . The global reward r(t) is defined to be the summation of the local rewards  $r_i(t)$ .

*Policy.* Each agent follows a localized policy that depends on its  $\beta$ -hop neighborhood, where  $\beta \ge 0$ is a fixed integer. Specifically, at time step t, given the global state s(t), agent i adopts a local policy

 $\zeta_i$  parameterized by  $\theta_i$  to decide the distribution of  $a_i(t)$  based on the states of agents in  $N_i^{\beta}$ .

Our objective is for all the agents to *cooperatively* maximize the discounted global reward, i.e.,  $J(\theta) = \mathbb{E}_{s \sim \pi_0} \left[ \sum_{t=0}^{\infty} \gamma^t r(s(t), a(t)) \mid s(0) = s \right]$ , where  $\pi_0$  is a given distribution on the initial global state, and we recall r(s(t), a(t)) is the global stage reward defined as the sum of all local

 $_{260}$  rewards at time t.

*Examples.* To highlight the applicability of the general model, we include two examples of networked systems that feature the dependence structure captured by our model in Appendix A: a wireless communication example and an example of controlling a process that spreads over a network.

Note that a limitation of our setting is that the dependence structure we consider is stationary, in the sense that dependencies are sampled i.i.d. from the distribution  $\mathcal{D}$ . It is important to consider more general time-varying forms (e.g. Markovian) in future research.

*Background.* Before moving on, we review a few key concepts in RL which will be useful in the rest of the section. We use  $\pi_t^{\theta}$  to denote the distribution of s(t) under policy  $\theta$  given that

<sup>&</sup>lt;sup>2</sup>Here, correlations between  $L_t^s$  and  $L_t^r$  are possible

- $s(0) \sim \pi_0$ . A well-known result [46] is that the gradient of the objective  $\nabla J(\theta)$  can be computed by  $\frac{1}{1-\gamma} \mathbb{E}_{s \sim \pi^{\theta}, a \sim \zeta^{\theta}(\cdot|s)} Q^{\theta}(s, a) \nabla \log \zeta^{\theta}(a \mid s)$ , where distribution  $\pi^{\theta}(s) = (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \pi_t^{\theta}(s)$ 269 270
- is the discounted state visitation distribution. Evaluating the Q-function  $Q^{\theta}(s, a)$  plays a key
- 271 role in approximating  $\nabla J(\theta)$ . Local Q-function for agent i is the discounted local reward, i.e. 272
- $Q_i^{\theta}(s,a) = \mathbb{E}_{\zeta^{\theta}} \left[ \sum_{t=0}^{\infty} \gamma^t r_i(t) \mid s(0) = s, a(0) = a \right], \text{ where we use } r_i(t) \text{ to denote the local reward of agent } i \text{ at time step } t. \text{ Using local } Q\text{-functions, we can decompose the global } Q\text{-function as } Q^{\theta}(s,a) = \frac{1}{n} \sum_{i=1}^{n} Q_i^{\theta}(s,a), \text{ which allows each node to evaluate its local } Q\text{-function separately.}$ 273
- 274 275
- A key challenge in our MARL setting is that directly estimating the Q-functions is not scalable since 276 the size of the Q-functions is exponentially large in the number of agents. Therefore, in Section 4.1 277 we study structural properties of the Q-functions resulting from the dependence structure in the 278 transition (6), which enables us to design a scalable RL algorithm in Section 4.2279

#### 4.1 μ-decay Property 280

One of the core challenges for MARL is that the size of the Q function is exponentially large in the 281 number of agents. The key to our algorithm and its analysis is the identification of a novel structural 282 decay property for the Q-function, which says that the local Q-function of each agent i is mainly 283 decided by the states of the agents who are near *i*. This property is critical for the design of scalable 284 algorithms because it enables the agents to reduce the dimension of the Q-function by truncating 285 its dependence of the states and actions of far away agents. Recently, exponential decay has been 286 shown to hold in networked MARL when the network is static [38] 36], which is exploited to design a 287 scalable RL algorithm. However, in stochastic network settings it is too much to hope for exponential 288 decay in general 14, and so we introduce the more general notion of  $\mu$ -decay here, where  $\mu$  is a 289 function that converges to 0 as  $\kappa$  tends to infinity. The case of exponential decay that has been studied 290 previously corresponds to  $\mu(\kappa) = \gamma^{\kappa}/(1-\gamma)$ . The formal definition of  $\mu$ -decay is given below, 291

where for simplicity, we use  $i \xrightarrow{L} j$  to denote  $(i, j) \in L$  and denote  $N_{-i}^{\kappa} := \mathcal{N} \setminus N_i^{\kappa}$ . 292

**Definition 4.1.** For a function  $\mu : \mathbb{N} \to \mathbb{R}^+$  that satisfies  $\lim_{\kappa \to +\infty} \mu(\kappa) = 0$ , the  $\mu$ -decay property 293 holds if for any policy  $\theta$  and any  $i \in \mathcal{N}$ , the local Q function  $Q_i^{\theta}$  satisfies  $|Q_i^{\theta}(s,a) - Q_i^{\theta}(s',a')| \leq \mu(\kappa)$  for any (s,a), (s',a') that are identical within  $N_i^{\kappa}$ , i.e.  $s_{N_i^{\kappa}} = s'_{N_i^{\kappa}}, a_{N_i^{\kappa}} = a'_{N_i^{\kappa}}$ . 294 295

Intuitively, if the  $\mu$ -decay property holds and  $\mu(\kappa)$  decays quickly as  $\kappa$  increases, we can approximately decompose the global Q function as  $Q^{\theta}(s, a) = \frac{1}{n} \sum_{i=1}^{n} Q_{i}^{\theta}(s, a) \approx \frac{1}{n} \sum_{i=1}^{n} \hat{Q}_{i}^{\theta}(s_{N_{i}^{\kappa}}, a_{N_{i}^{\kappa}})$ , 296 297 where  $\hat{Q}_i$  only depends on the states and actions within the  $\kappa$ -hop neighborhood of agent *i*. Before 298 our work, [45] empirically showed that such a value decomposition allows efficient training of 299 MARL. Under the assumption that such decomposition exists, 45 propose an approach to learn this 300 decomposition. In contrast, as we prove in this section, the  $\mu$  decay property holds provably and 301 therefore, the global Q function can be directly decomposed in the networked MARL model and that 302 the error of such decomposition is provably small. 303

Our first result is Theorem 4.1 which shows the relationship between the random active link sets and 304 the  $\mu$ -decay property. The proof of Theorem 4.1 is deferred to Appendix D.1 305

**Theorem 4.1.** Define  $L^a$  as the static active link set that contains all pairs (i, j) whose graph 306 distance on  $\mathcal{G}$  is less than or equal to  $\beta$ , which is the dependency of local policy. Let random variable 307  $X_i(\kappa)$  denote the smallest  $t \in \mathbb{N}$  such that there exists a chain of agents 308

$$j_0^a \xrightarrow{L_0^s} j_1^s \xrightarrow{L^a} j_1^a \xrightarrow{L_1^s} \cdots \xrightarrow{L_{t-1}^s} j_t^s \xrightarrow{L^a} j_t^a,$$

that satisfies  $j_0^a \in N_{-i}^{\kappa}$  and  $j_t^a \xrightarrow{L_t^r} i$ . The  $\mu$ -decay property holds for  $\mu(\kappa) = \frac{1}{1-\gamma} \mathbb{E}[\gamma^{X_i(\kappa)}]$ . 309

- To make the  $\mu$ -decay result more concrete, we provide several scenarios that yield different upper 310
- bounds on the term  $\mathbb{E}[\gamma^{X_i(\kappa)}]$ . In the first scenario, we study the case where long range links do 311
- not exist in Corollary 4.2. In this case, we obtain an exponential decay property that generalizes the 312
- result in [38]. A proof is in Appendix D.2 313
- **Corollary 4.2** (Exponential Decay). Consider a distribution  $\mathcal{D}$  of active link sets that satisfies 314

$$P_{(L^s,L^r)\sim\mathcal{D}}\{(i,j)\in L^s\}=0, \text{ for all } i,j\in\mathcal{N} \text{ s.t. } d_{\mathcal{G}}(i,j)\geq\alpha_1, \\ P_{(L^s,L^r)\sim\mathcal{D}}\{(i,j)\in L^r\}=0, \text{ for all } i,j\in\mathcal{N} \text{ s.t. } d_{\mathcal{G}}(i,j)\geq\alpha_2.$$

### Algorithm 1 Scalable Actor Critic

- 1: for  $m = 0, 1, 2, \cdots$  do
- Sample initial global state  $s(0) \sim \pi_0$ . 2:
- Each node *i* takes action  $a_i(0) \sim \zeta_i^{\theta_i(m)}(\cdot \mid s_{N_i^{\theta}}(0))$  to obtain the global state s(1). 3:
- Each node *i* records  $s_{N_i^{\kappa}}(0), a_{N_i^{\kappa}}(0), r_i(0)$  and initialize  $\hat{Q}_i^0$  to be all zero vector. 4:
- for  $t = 1, \cdots, T$  do 5:
- Each node *i* takes action  $a_i(t) \sim \zeta_i^{\theta_i(m)}(\cdot \mid s_{N_i^{\beta}}(t))$  to obtain the global state s(t+1). 6:
- Each node *i* update the local estimation  $\hat{Q}_i$  with step size  $\alpha_{t-1} = \frac{H}{t-1+t_0}$ , 7:

$$\begin{aligned} \hat{Q}_{i}^{t} \big( s_{N_{i}^{\kappa}}(t-1), a_{N_{i}^{\kappa}}(t-1) \big) &= \\ (1-\alpha_{t-1}) \hat{Q}_{i}^{t-1} \big( s_{N_{i}^{\kappa}}(t-1), a_{N_{i}^{\kappa}}(t-1) \big) + \alpha_{t-1} \Big( r_{i}(t) + \gamma \hat{Q}_{i}^{t-1} \big( s_{N_{i}^{\kappa}}(t), a_{N_{i}^{\kappa}}(t) \big) \Big), \\ \hat{Q}_{i}^{t} \big( s_{N_{i}^{\kappa}}, a_{N_{i}^{\kappa}} \big) &= \hat{Q}_{i}^{t-1} \big( s_{N_{i}^{\kappa}}, a_{N_{i}^{\kappa}} \big) \text{ for } \big( s_{N_{i}^{\kappa}}, a_{N_{i}^{\kappa}} \big) \neq \big( s_{N_{i}^{\kappa}}(t-1), a_{N_{i}^{\kappa}}(t-1) \big). \end{aligned}$$

- 8: Each node *i* approximate  $\nabla_{\theta_i} J(\theta)$  by  $\hat{g}_{i}(m) = \sum_{t=0}^{T} \gamma^{t} \frac{1}{n} \sum_{j \in N_{i}^{\kappa}} \hat{Q}_{j}^{T} \left( s_{N_{j}^{\kappa}}(t), a_{N_{j}^{\kappa}}(t) \right) \nabla_{\theta_{i}} \log \zeta_{i}^{\theta_{i}(m)} \left( a_{i}(t) \mid s_{N_{i}^{\beta}}(t) \right).$ Each node *i* conducts gradient ascent by  $\theta_i(m+1) = \theta_i(m) + \eta_m \hat{q}_i(m)$ . 9:
- Then,  $\mathbb{E}[\gamma^{X_i(\kappa)}] < C\rho^{\kappa}$ , where  $\rho = \gamma^{1/(\alpha_1+\beta)}, C = \gamma^{-\alpha_2/(\alpha_1+\beta)}$ . 315

1.

In the second scenario, long range active links can occur, but with exponentially small probability 316 with respect to their distance. In this case, we can obtain a near-exponential decay property where 317  $\mu(\kappa) = O(\rho^{\kappa/\log \kappa}))$  for some  $\rho \in (0, 1)$ . A proof can be found in Appendix D.3 318

**Theorem 4.3** (Near-Exponential Decay). Suppose the distribution  $\mathcal{D}$  of active link sets satisfies 319

 $P_{(L^s, L^r) \sim \mathcal{D}}\{(i, j) \in L^s \cup L^r\} < c\lambda^{d_{\mathcal{G}}(i, j)}, \text{ for all } i, j \in \mathcal{N},$ 

where  $c \ge 1, 1 > \lambda > 0$  are constants. If the largest size of the  $\kappa$  neighborhood in the underlying 320 graph  $\mathcal{G}$  can be bounded by a polynomial of  $\kappa$ , i.e., there exists some constants  $c_0 \geq 1, n_0 \in \mathbb{N}$  such that  $|\{j \in \mathcal{N} \mid d_{\mathcal{G}}(i,j) = \kappa\}| \leq c_0(\kappa+1)^{n_0}$  holds for all i, then  $\mathbb{E}[\gamma^{X_i(\kappa-1)}] \leq C\rho^{\kappa/(1+\ln(\kappa+1))}$ 321 322 for some positive constant C and decay rate  $\rho < 1$ .<sup>3</sup> 323

It is interesting to compare the result above with models of the so-called "small world phenomena" in 324 social networks, e.g., [14]. In these models, a link (i, j) occurs with probability  $1/poly(d_{\mathcal{G}}(i, j))$ , as 325 opposed to the exponential dependence in Lemma 4.3. In this case, one can see function  $\mu$  is lower 326 bounded by  $1/poly(\kappa)$ , which leads us to conjecture that  $\mu$  is also upper bounded by  $O(1/poly(\kappa))$ . 327 Thus, when information spreads "slowly" it helps a localized algorithm to learn efficiently. 328

#### 4.2 A Scalable Actor Critic Algorithm 329

Motivated by the  $\mu$ -decay property of the Q-functions, we design a novel Scalable Actor Critic 330 algorithm (Algorithm [I]) for networked MARL problem, which exploits the  $\mu$ -decay result in the 331 previous section. The Critic part (from line 2 to line 7) uses the local trajectory  $\{(s_{N_{i}^{\mu}}, a_{N_{i}^{\mu}}, r_{i})\}$  to 332 evaluate the local Q-functions under parameter  $\theta(m)$ . Intuitively, the  $\mu$ -decay property guarantees 333 that we can achieve good approximation error even when  $\kappa$  is not large. The Actor part (from line 334 8 to line 9 computes the estimated partial derivative using the estimated local Q-functions, and 335 uses the partial derivative to update local parameter  $\theta_i$ . The step size sequence  $\{\eta_m\}$  will be defined 336 in Theorem D.2. Compared with the Scalable Actor Critic algorithm proposed in [38], Algorithm 337 1 extends the policy dependency structure considered. No longer is the dependency completely 338 local; it now extends to all agents within the  $\beta$ -hop neighborhood. Interestingly, the time-varying 339 dependencies do not add complexity into the algorithm (though the analysis is more complex). 340

Algorithm is highly scalable. Each agent *i* needs only to query and store the information within its 341  $\kappa$ -hop neighborhood during the learning process. The parameter  $\kappa$  can be set to balance accuracy and 342 complexity. Specifically, as  $\kappa$  increases, the error bound becomes tighter at the expense of increasing 343 computation, communication, and space complexity. 344

<sup>&</sup>lt;sup>3</sup>The explicit expression of C and  $\rho$  can be found in Appendix D.3

### 345 4.3 Convergence

We now present our main result, a finite-time error bound for the Scalable Actor Critic algorithm (Algorithm 1) that holds under general (non-local) dependencies. To that end, we first describe the assumption needed in our result. It focuses on the Markov chain formed by the global state-action pair (*s*, *a*) under a fixed policy parameter  $\theta$  and is standard for finite-time convergence results in RL, e.g., [43, 5, 37].

Assumption 4.1. Under any fixed policy  $\theta$ ,  $\{z(t) := (s(t), a(t))\}$  is an aperiodic and irreducible Markov chain on state space  $\mathcal{Z} := \mathcal{S} \times \mathcal{A}$  with a unique stationary distribution  $d^{\theta} = (d_z^{\theta}, z \in \mathcal{Z})$ , which satisfies  $d_z^{\theta} > 0, \forall z \in \mathcal{Z}$ . Define  $d^{\theta}(z') = \sum_{z \in \mathcal{Z}: z_{N_i^{\kappa}} = z'} d^{\theta}(z)$ and  $\sigma'(\kappa) := \inf_{z' \in \mathcal{Z}_{N_i^{\kappa}}} d^{\theta}(z')$ . There exists positive constants  $K_1, K_2$  such that  $K_2 \geq 1$  and  $\forall z' \in \mathcal{Z}, \forall t \geq 0, \sup_{\mathcal{K} \subseteq \mathcal{Z}} \left| \sum_{z \in \mathcal{K}} d_z^{\theta} - \sum_{z \in \mathcal{K}} \mathbb{P}(z(t) = z \mid z(0) = z') \right| \leq K_1 e^{-t/K_2}$ .

Recall that in TD learning with state aggregation (Section  $\overline{3}$ ), we defined a surjection h that maps a 356 state to an abstract state. To have a good approximate equivalence, we need to find a good h, i.e., 357 if two states are mapped to the same abstract state, their value functions are required to be close 358 (Theorem 3.1). In the context of networked MARL, the  $\mu$  decay property (Definition 4.1) provides a 359 natural mapping h for state aggregation. To see this, for each agent i, let h map the global state/action 360 to the local states/actions in agent i's  $\kappa$ -hop neighborhood, i.e.,  $h(s, a) = (s_{N_{\kappa}^{\kappa}}, a_{N_{\kappa}^{\kappa}})$ . The  $\mu$ -decay 361 property guarantees that if h(s, a) = h(s', a'), the difference in their Q-functions is upper bounded 362 by  $\mu(\kappa)$ , which is vanishing as  $\kappa$  increases. This idea leads to the following corollary by applying 363 Theorem 3.1 to the networked MARL system. 364

**Corollary 4.4.** Suppose Assumption [4.1] and  $\mu$ -decay property (Definition [4.1]) hold. Let the step size be  $\alpha_t = \frac{H}{t+t_0}$  with  $t_0 = \max(4H, 2K_2 \log T)$ , and  $H \ge \frac{2}{(1-\gamma)\sigma'(\kappa)}$ . Define constant  $C_3$ as in Theorem [3.1] Then, inside outer loop iteration m, for each  $i \in \mathcal{N}$ , with probability at least  $1 - \delta$ , we have  $\sup_{(s,a)\in\mathcal{S}\times\mathcal{A}} \left| Q_i^{\theta(m)}(s,a) - \hat{Q}^T(s_{N_i^{\kappa}}, a_{N_i^{\kappa}}) \right| \le \frac{C_a}{\sqrt{T+t_0}} + \frac{C_a}{T+t_0} + \frac{\mu(\kappa)}{1-\gamma}$ ,

where the constants are given by 
$$C_a = \frac{40H}{(1-\gamma)^2} \sqrt{K_2 \log T \left( \log \left( \frac{4f(\kappa)K_2T}{\delta} \right) + \log \log T \right)}$$
 and  
 $C'_a = \frac{8}{(1-\gamma)^2} \max \{ \frac{144K_2H \log T}{\sigma'(\kappa)} + C_3, 2K_2 \log T + t_0 \}.$ 

The most related result in the literature to the above is Theorem 7 in [38]. In comparison, Corollary 4.4 applies for more general, potentially non-local, dependencies and, also, improves the constant term by a factor of  $1/(1 - \gamma)$ .

To analyze the Actor part of Algorithm 1, we make the following additional boundedness and Lipschitz continuity assumptions on the gradients. These are standard assumptions in the literature.

**Assumption 4.2.** For any  $i, a_i, s_{N_i^{\beta}}$  and  $\theta_i$ , we assume  $\left\| \nabla_{\theta_i} \log \zeta_i^{\theta_i}(a_i \mid s_{N_i^{\beta}}) \right\| \le W_i$ . Then, for any  $L_t^a$ ,  $\left\| \nabla_{\theta} \log \zeta^{\theta}(a \mid s) \right\| \le W := \sqrt{\sum_{i=1}^n W_i^2}$ . We further assume  $\nabla J(\theta)$  is W'-Lipschitz in  $\theta$ .

Intuitively, since the quality of the estimated policy gradient depends on the quality of the estimation of *Q*-functions, if every agent *i* has learned a good approximation of its local *Q*-function in the Critic part of Algorithm 1, the policy gradient can be approximated well. Therefore, the Actor part can obtain a good approximation of a stationary point of the objective function. We state the sample complexity result in Theorem 4.5 and defer the detailed bounds and a proof to Appendix D.4.

**Theorem 4.5.** Under Assumption 4.2 to reach an  $O(\epsilon)$ -approximate stationary point with probability at least  $1 - \delta$ , we need to choose  $\kappa$  such that  $\mu(\kappa) = O(W^{-2}(1 - \gamma)^4 \epsilon)$ . The number of required iterations of the outer loop should satisfy  $M = \tilde{\Omega}(\epsilon^{-2}poly(W, W', \frac{1}{1-\gamma}))$  and the number of required iterations of the inner loop is  $T = \tilde{\Omega}(\epsilon^{-2}poly(W, \frac{1}{\sigma'(\kappa)}, K_2, \frac{1}{1-\gamma}, \log f(\kappa), \log(1/\delta)))$ .

Note that W scales with the number of agents n. Thus, Theorem 4.5 shows that the complexity of our algorithm scales with the largest state-action space size of any  $\kappa$ -hop neighborhood and the number of agents n, which avoids the exponential blowup in n when the graph is sparse and achieves scalable RL for networked agents even under stochastic, non-local settings.

### 391 **References**

- [1] B. Bamieh, F. Paganini, and M. A. Dahleh. Distributed control of spatially invariant systems.
   *IEEE Transactions on automatic control*, 47(7):1091–1107, 2002.
- [2] D. P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. II*. Athena Scientific, 3rd edition, 2007.
- [3] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*, volume 5. Athena Scientific
   Belmont, MA, 1996.
- J. Bhandari, D. Russo, and R. Singal. A finite time analysis of temporal difference learning with
   linear function approximation. In *Conference On Learning Theory*, pages 1691–1692. PMLR,
   2018.
- [5] P. Bremaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues.* Texts in
   Applied Mathematics. Springer New York, 2013.
- [6] L. Bu, R. Babu, B. De Schutter, et al. A comprehensive survey of multiagent reinforcement
   learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- [7] D. Chakrabarti, Y. Wang, C. Wang, J. Leskovec, and C. Faloutsos. Epidemic thresholds in real
   networks. *ACM Transactions on Information and System Security (TISSEC)*, 10(4):1, 2008.
- [8] C. Claus and C. Boutilier. The dynamics of reinforcement learning in cooperative multiagent
   systems. AAAI/IAAI, 1998:746–752, 1998.
- [9] C. Dann, N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. On oracle efficient pac rl with rich observations. In *Advances in Neural Information Processing Systems*,
   pages 1422–1432, 2018.
- [10] T. Doan, S. Maguluri, and J. Romberg. Finite-time analysis of distributed TD(0) with linear function approximation on multi-agent reinforcement learning. In K. Chaudhuri and
  R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1626–1635, Long Beach,
  California, USA, 09–15 Jun 2019. PMLR.
- <sup>418</sup> [11] T. T. Doan. Finite-time analysis and restarting scheme for linear two-time-scale stochastic <sup>419</sup> approximation, 2019.
- [12] K. Dong, Y. Wang, X. Chen, and L. Wang. Q-learning with ucb exploration is sample efficient
   for infinite-horizon mdp. *arXiv preprint arXiv:1901.09311*, 2019.
- Y. Duan, X. Chen, R. Houthooft, J. Schulman, and P. Abbeel. Benchmarking deep reinforcement
   learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.
- <sup>425</sup> [14] D. Easley, J. Kleinberg, et al. Networks, crowds, and markets: Reasoning about a highly <sup>426</sup> connected world. *Significance*, 9:43–44, 2012.
- [15] J. Foerster, G. Farquhar, T. Afouras, N. Nardelli, and S. Whiteson. Counterfactual multi-agent
   policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [16] D. Gamarnik. Correlation decay method for decision, optimization, and inference in large-scale
   networks. In *Theory Driven by Influential Applications*, pages 108–121. INFORMS, 2013.
- [17] D. Gamarnik, D. A. Goldberg, and T. Weber. Correlation decay in random decision networks.
   *Mathematics of Operations Research*, 39(2):229–261, 2014.
- [18] H. Gu, X. Guo, X. Wei, and R. Xu. Q-learning for mean-field controls, 2020.
- 435 [19] H. Gu, X. Guo, X. Wei, and R. Xu. Q-learning for mean-field controls. *arXiv preprint* 436 *arXiv:2002.04131*, 2020.

- [20] D. hwan Lee and N. He. A unified switching system perspective and o.d.e. analysis of q-learning
   algorithms. *ArXiv*, abs/1912.02270, 2019.
- [21] N. Jiang. Notes on state abstractions. http://nanjiang.web.engr.illinois.edu/
   files/cs598/note4.pdf, 2018.
- [22] N. Jiang, A. Kulesza, and S. Singh. Abstraction selection in model-based reinforcement learning.
   In *International Conference on Machine Learning*, pages 179–188, 2015.
- [23] N. K. Jong and P. Stone. State abstraction discovery from irrelevant state variables. In *IJCAI*,
   volume 8, pages 752–757, 2005.
- [24] T. Lattimore and M. Hutter. Pac bounds for discounted mdps. In N. H. Bshouty, G. Stoltz,
   N. Vayatis, and T. Zeugmann, editors, *Algorithmic Learning Theory*, pages 320–334, Berlin,
   Heidelberg, 2012. Springer Berlin Heidelberg.
- [25] D. Li, D. Zhao, Q. Zhang, and Y. Chen. Reinforcement learning and deep learning based lateral
   control for autonomous driving [application notes]. *IEEE Computational Intelligence Magazine*,
   14(2):83–98, 2019.
- [26] L. Li, T. J. Walsh, and M. L. Littman. Towards a unified theory of state abstraction for MDPs.
   In *ISAIM*, 2006.
- [27] M. Llas, P. M. Gleiser, J. M. López, and A. Díaz-Guilera. Nonequilibrium phase transition
   in a model for the propagation of innovations among economic agents. *Physical Review E*,
   68(6):066101, 2003.
- [28] R. Lowe, Y. Wu, A. Tamar, J. Harb, O. P. Abbeel, and I. Mordatch. Multi-agent actor-critic for
   mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pages 6379–6390, 2017.
- [29] L. Matignon, G. J. Laurent, and N. Le Fort-Piat. Independent reinforcement learners in coopera tive markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 27(1):1–31, 2012.
- [30] W. Mei, S. Mohagheghi, S. Zampieri, and F. Bullo. On the dynamics of deterministic epidemic
   propagation over networks. *Annual Reviews in Control*, 44:116–128, 2017.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves,
   M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep rein forcement learning. *Nature*, 518(7540):529, 2015.
- <sup>467</sup> [32] N. Motee and A. Jadbabaie. Optimal control of spatially distributed systems. *IEEE Transactions* <sup>468</sup> *on Automatic Control*, 53(7):1616–1629, 2008.
- [33] M. J. Neely. Optimal backpressure routing for wireless networks with multi-receiver diversity.
   In 2006 40th Annual Conference on Information Sciences and Systems, pages 18–25, 2006.
- [34] C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of optimal queuing network control.
   *Mathematics of Operations Research*, 24(2):293–305, 1999.
- [35] G. Qu and N. Li. Exploiting fast decaying and locality in multi-agent mdp with tree dependence
   structure. In 2019 IEEE 58th Conference on Decision and Control (CDC), pages 6479–6486.
   IEEE, 2019.
- [36] G. Qu, Y. Lin, A. Wierman, and N. Li. Scalable multi-agent reinforcement learning for networked systems with average reward. *Advances in Neural Information Processing Systems*, 33, 2020.
- [37] G. Qu and A. Wierman. Finite-time analysis of asynchronous stochastic approximation and
   q-learning. In *Conference on Learning Theory*, pages 3185–3205. PMLR, 2020.
- [38] G. Qu, A. Wierman, and N. Li. Scalable reinforcement learning of localized policies for multi-agent networked systems. In *Learning for Dynamics and Control*, pages 256–266. PMLR, 2020.

- [39] N. A. Ruhi, C. Thrampoulidis, and B. Hassibi. Improved bounds on the epidemic threshold of
   exact sis models on complex networks. In 2016 IEEE 55th Conference on Decision and Control
   (CDC), pages 3560–3565. IEEE, 2016.
- [40] D. Shah and Q. Xie. Q-learning with nearest neighbors. In Advances in Neural Information
   Processing Systems, pages 3111–3121, 2018.
- [41] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser,
   I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural
   networks and tree search. *nature*, 529(7587):484, 2016.
- [42] S. P. Singh, T. Jaakkola, and M. I. Jordan. Reinforcement learning with soft state aggregation.
   In Advances in neural information processing systems, pages 361–368, 1995.
- [43] R. Srikant and L. Ying. Finite-time error bounds for linear stochastic approximation andtd
   learning. pages 2803–2830, 2019.
- [44] J. Subramanian and A. Mahajan. Reinforcement learning in stationary mean-field games.
   In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 251–259, 2019.
- [45] P. Sunehag, G. Lever, A. Gruslys, W. M. Czarnecki, V. F. Zambaldi, M. Jaderberg, M. Lanctot,
   N. Sonnerat, J. Z. Leibo, K. Tuyls, et al. Value-decomposition networks for cooperative
   multi-agent learning based on team reward. In *AAMAS*, pages 2085–2087, 2018.
- [46] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforce ment learning with function approximation. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, NIPS'99, page 1057–1063, Cambridge, MA, USA,
   1999. MIT Press.
- [47] J. N. Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine learning*, 16(3):185–202, 1994.
- [48] J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690, 1997.
- [49] J. N. Tsitsiklis and B. Van Roy. Analysis of temporal-difference learning with function approximation. In *Advances in neural information processing systems*, pages 1075–1081, 1997.
- [50] W. Vogels, R. van Renesse, and K. Birman. The power of epidemics: Robust communication
   for large-scale distributed systems. *SIGCOMM Comput. Commun. Rev.*, 33(1):131–135, Jan.
   2003.
- 515 [51] M. J. Wainwright. Stochastic approximation with cone-contractive operators: Sharp  $\ell_i nfty$ -516 bounds for Q-learning. *arXiv preprint arXiv:1905.06265*, 2019.
- [52] S. Wang, V. Venkateswaran, and X. Zhang. Fundamental analysis of full-duplex gains in wireless networks. *IEEE/ACM Transactions on Networking*, 25(3):1401–1416, 2017.
- [53] Y. Wu, W. Zhang, P. Xu, and Q. Gu. A finite time analysis of two time-scale actor critic methods,
   2020.
- [54] T. Xu, S. Zou, and Y. Liang. Two time-scale off-policy td learning: Non-asymptotic analysis
   over markovian samples. In *Advances in Neural Information Processing Systems 32*, pages
   10634–10644. Curran Associates, Inc., 2019.
- Y. Yang, R. Luo, M. Li, M. Zhou, W. Zhang, and J. Wang. Mean field multi-agent reinforcement
   learning. In *International Conference on Machine Learning*, pages 5571–5580. PMLR, 2018.
- [56] K. Zhang, Z. Yang, and T. Başar. Multi-agent reinforcement learning: A selective overview of
   theories and algorithms. *arXiv preprint arXiv:1911.10635*, 2019.
- [57] K. Zhang, Z. Yang, H. Liu, T. Zhang, and T. Başar. Fully decentralized multi-agent reinforcement
   learning with networked agents. *arXiv preprint arXiv:1802.08757*, 2018.

- <sup>530</sup> [58] R. Zhang and M. Pavone. Control of robotic mobility-on-demand systems: a queueingtheoretical perspective. *The International Journal of Robotics Research*, 35(1-3):186–203,
- 532 2016.
- [59] A. Zocca. Temporal starvation in multi-channel csma networks: an analytical framework.
   *Queueing Systems*, 91(3-4):241–263, 2019.

## 535 Checklist

536	1. For all authors
537 538 539	(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Our main results in Sections 2.3 and 4 are consistent with our abstract and introduction.
540 541	(b) Did you describe the limitations of your work? [Yes] We discuss the limitation of our networked MARL setting after 'Examples' in Section 4.
542 543 544	(c) Did you discuss any potential negative societal impacts of your work? [N/A] The goal of this work is to advance theoretical foundations of MARL algorithms. We do not see any negative societal impacts of this work.
545 546	(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We did not find any aspect of our work that may lead to ethics violations.
547	2. If you are including theoretical results
548 549	(a) Did you state the full set of assumptions of all theoretical results? [Yes] In the statement of each theorem, we refer to all the assumptions needed clearly.
550 551	(b) Did you include complete proofs of all theoretical results? [Yes] See Appendices for complete proofs of our results.
552	3. If you ran experiments
553 554	(a) Did you include the code, data, and instructions needed to reproduce the main experi- mental results (either in the supplemental material or as a URL)? [N/A]
555 556	(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
557 558	(c) Did you report error bars (e.g., with respect to the random seed after running experi- ments multiple times)? [N/A]
559 560	(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
561	4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets
562	(a) If your work uses existing assets, did you cite the creators? [N/A]
563	(b) Did you mention the license of the assets? [N/A]
564 565	(c) Did you include any new assets either in the supplemental material or as a URL? $[N/A]$
566 567	(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
568 569	(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
570	5. If you used crowdsourcing or conducted research with human subjects
571 572	<ul> <li>(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]</li> </ul>
573 574	(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
575 576	(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]