
SIMULATION-BASED OPTIMISTIC POLICY ITERATION FOR MULTI-AGENT MDPs WITH KULLBACK-LEIBLER CONTROL COST

Khaled Nakhleh

Electrical and Computer Engineering Department
Texas A&M University
College Station, TX
{khaled.jamal}@tamu.edu

Ceyhun Eksin

Industrial and System Engineering Department
Electrical and Computer Engineering Department
Texas A&M University
College Station, TX
eksinc@tamu.edu

Sabit Ekin

Engineering Technology and Industrial Distribution Department
Electrical and Computer Engineering Department
Texas A&M University
College Station, TX
sabitekin@tamu.edu

Keywords reinforcement learning · optimistic policy iteration · KL control · multi-agent MDPs · Markov games

ABSTRACT

This paper proposes an agent-based optimistic policy iteration (OPI) scheme for learning stationary optimal stochastic policies in multi-agent Markov Decision Processes (MDPs), in which agents incur a Kullback-Leibler (KL) divergence cost for their control efforts and an additional cost for the joint state. The proposed scheme consists of a greedy policy improvement step followed by an m -step temporal difference (TD) policy evaluation step. We use the separable structure of the instantaneous cost to show that the policy improvement step follows a Boltzmann distribution that depends on the current value function estimate and the uncontrolled transition probabilities. This allows agents to compute the improved joint policy independently. We show that both the synchronous (entire state space evaluation) and asynchronous (a uniformly sampled set of substates) versions of the OPI scheme with finite policy evaluation rollout converge to the optimal value function and an optimal joint policy asymptotically. Simulation results on a multi-agent MDP with KL control cost variant of the Stag-Hare game validates our scheme's performance in terms of minimizing the cost return.

1 Introduction

Consider the two-agent MDP version [12] of the Stag-Hare game [21] where two hunters move on a gridworld to hunt hares or a stag (see Figure (1) for an illustration). Each hunter can only determine their own next position (local state) in this world by moving to a neighboring grid, but their costs are determined based on their joint position (state), and whether they hunt a stag or a hare. The hunters can individually hunt a hare, but need to coordinate together to be able to hunt a stag. Considering the returns are the same amongst the hunters, this setting is an example of a multi-agent Markov decision process (MDP), in which each agent has control of its local state but there is a common value function (identical interest) that depends on the joint state.

We assume the instantaneous costs are decomposed into two parts: an identical cost term that only depends on the joint state, and another term that captures the cost of control. The cost of control is measured by the KL divergence between the uncontrolled transition probabilities and transition probabilities collectively chosen by the agents. The KL cost represents the control cost the agents are willing to pay in order to modify the uncontrolled transition probability




Figure 1: A multi-agent MDP with KL control cost: Two hunters hunting either hares or a stag on a 5×5 gridworld. See Section (6) for details.

function. We coin the class of MDPs we consider as the multi-agent MDPs with Kullback-Leibler (KL) control cost which stem from the linearly solvable MDPs [22, 23] and the adversarial linearly solvable Markov games [8].

In the proposed Kullback-Leibler controlled optimistic policy iteration (KLC-OPI) scheme, each agent straddles between policy improvement and policy evaluation steps as is the case with the single agent OPI [16]. Optimistic policy iteration generalizes value iteration and policy iteration methods by considering a finite m -step temporal difference (TD) rollout in the policy evaluation step, where we obtain the value iteration method when $m = 1$ and the policy iteration method when $m \rightarrow \infty$. In the policy improvement step, the agents update their policies greedily considering their current policy valuations. The updated policies are then evaluated using the finite m -step TD rollout with sampled trajectories, where agents concurrently generate trajectories for each joint state in the joint state space, and compute the discounted returns. The m -step TD rollout allows for a less noisy unbiased estimation of the policy compared to value iteration, while performing a finite rollout remains practically viable in contrast to policy iteration.

The instantaneous cost decomposition along with the infinite horizon discounted sum of the costs render a close form for the greedy optimal policies, that follow a Boltzmann distribution using a Cole-Hopf transformation of the value function (Lemma 4.1). The resultant controlled transition probabilities selects next states inversely proportional to the exponent of the states' cost returns. That is, the KL control cost minimization solution results in a stochastic policy that only depends on the uncontrolled transition probability function and the current value function estimate, which eliminates the combinatorial search over the state-action space. Indeed, we do not need to restrict the action space to be finite, and allow the individual action spaces to be continuous. In addition, each agent can sample their next sub-state from its marginal controlled policy having computed the closed-form joint policy.

For simulation-based policy iteration schemes, a synchronous policy evaluation step requires simulating a trajectory for each joint state in the state space. We also consider the asynchronous implementation of the KLC-OPI where a subset of the joint state space is evaluated per iteration, potentially eliminating the need to run trajectories for each joint state. We show that the asynchronous version of the KLC-OPI, namely ASYNC-KLC-OPI, can be related to the synchronous version of the scheme. We analyze the asymptotic convergence of the two schemes, KLC-OPI and ASYNC-KLC-OPI, and prove given standard assumptions on the learning rate and the initialized value functions, that the schemes' iterates asymptotically converge to the optimal value function for all agents and to an optimal joint policy (Theorem 4.4 and Corollary 5.2, respectively). Finally, simulation results on the aforementioned multi-agent MDP variant of the Stag-Hare game show that the scheme is able to learn a joint policy that minimizes the cost return (Section (6)).

2 Related Work and Contributions

2.1 Simulation-Based Optimistic Policy Iteration

Optimistic, also known as simulation-based modified, policy iteration methods are preferred in settings with large state spaces due to their fast convergence [20]. However, the theoretical underpinnings of this practical success remain unclear. The pioneering theoretical guarantee for OPIs with Monte-Carlo estimation established convergence assuming

infinitely long trajectories that start from each state at every iteration [24]. Indeed, we follow the same proof structure here but consider finitely long trajectories. Similar, OPI schemes were considered for the stochastic shortest path problem in [6, 14]. More recent efforts focused on showing convergence of finite-step rollouts during the policy evaluation step after applying multi-step greedy or lookahead policies [9, 27, 28]. In particular, prior the asymptotic error bound was strengthened in [28] using stochastic approximation techniques for the setting where policy evaluation is done using only a single Monte-Carlo trajectory in each iteration. In contrast to these works, we prove asymptotic convergence to the optimal value function and an optimal joint policy without requiring a lookahead operation, i.e., a greedy improvement, while also performing a finite rollout in each iteration (Theorem 4.4). Our asymptotic convergence proof technique is applicable to any single-agent MDP with deterministic policies, not just MDPs with KL control costs.

2.2 Decentralized Learning in Multi-Agent MDPs and Markov Games

The KLC-OPI scheme considers multiple agents that perform the policy improvement step independently of other agents using an agent-local estimate of the value function (Lemma 4.1). Thus, the proposed scheme falls under the framework of multi-agent reinforcement learning (MARL) [29], which has seen growing interest in the context of learning in identical interest or potential games [13, 11, 25], zero-sum Markov games [17, 18, 4], and multi-agent systems with KL control cost [5]. The considered multi-agent MDP framework is also equivalent to identical-interest Markov games or Markov cooperative games considered in [13, 7], where the costs/rewards are identical for all the agents.

Work on Markov games consider either policy gradients, in which agents consider parameterized policies updated using gradients computed through episodic returns [13, 1], or a combination of standard learning protocols, e.g., best-response, fictitious, with a standard reinforcement learning algorithm, e.g., Q-learning [18]. Policy gradient methods are applicable to continuous state and action spaces' environments, but they suffer from large variance and the convergence rate being sensitive to the choice of parameters.

Recently, localized policy iteration methods for networked multi-agent systems [30] and zero-sum games [2] are shown to converge near globally optimal policy. These schemes are based on state space partitioning with policy dependent mapping that gives a uniform sup-norm contraction property which enables convergence of the algorithms to the optimal value functions.

The proposed KLC-OPI scheme is a novel agent-based learning scheme that is shown to carry over the convergence properties of OPI designed for single-agent MDPs to the multi-agent MDP with KL control cost setting. The KL cost structure allows for continuous action spaces through the close-form solution to the policy improvement step.

3 Multi-Agent MDPs With KL Control Cost

We consider an infinite-horizon discounted n -agent MDP given as the tuple $\Gamma := \{N, \mathcal{S}, \{\mathcal{A}\}_{i=1}^n, P, \rho, C, \gamma\}$ with a finite number of players $N := \{1, \dots, n\}$, finite joint state space \mathcal{S} , and a continuous action space $\mathcal{A} = \mathcal{A}_{i=1} \times \mathcal{A}_{i=2} \times \dots \times \mathcal{A}_{i=n}$. The transition probability function $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ determines the transition probability to the next joint state $s_{t+1} \in \mathcal{S}$ given the joint state $s_t \in \mathcal{S}$ and joint action profile \mathbf{a}_t at time step $t \in \mathbb{N}^+$. The initially sampled joint state is chosen from a prior ρ , for which we assume $\rho(s) > 0$ for all joint states $s \in \mathcal{S}$.

The intrinsic joint state cost function is defined as $C : \mathcal{S} \rightarrow \mathbb{R}$.

As done in [22] for single-agent linearly solvable MDPs, in multi-agent MDPs with KL control cost, the one-step cost function is composed of two terms: the intrinsic joint state cost function C , and the control cost, measured using the Kullback-Leibler (KL) divergence between the controlled and uncontrolled transition probability function. Finally, $\gamma \in [0, 1)$ is the discount factor.

3.1 Stochastic Joint and Marginal Policies

In the KL control setting, agents pick a joint policy by re-weighting P with a continuous-valued action profile $\mathbf{a} \in \mathcal{A}$ that directly specifies the transition probability from s to $s' \in \mathcal{S}$. Hence, the agents avoid the combinatorial search over the state-action space for an action profile that is then applied to the transition probability function P .

Assumption 3.1. The joint state $s_t \in \mathcal{S}$ at time step t is composed of n sub-states where the sub-state $s_{i,t} \in \mathcal{S}_i$ can only be controlled by agent i at time step t . The joint state is then $s_t = [s_{1,t}, s_{2,t}, \dots, s_{n,t}]$.

Given Assumption 3.1, the joint state space can be written as $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \dots \times \mathcal{S}_n$. Moreover, and similar to product games [10], the multi-agent MDP with KL control cost transition structure is derived by taking the product of n Markov transition structures. Agent $i \in N$ only controls their sub-state transitions through a probability transition

function $P_i : \mathcal{S} \times \mathcal{A}_i \rightarrow \Delta(\mathcal{S}_i)$ such that $P = \times_{i \in \mathcal{N}} P_i$. Agent i obtains its uncontrolled transition probability function $P_{i,0} : \mathcal{S} \rightarrow \Delta(\mathcal{S}_i)$ by applying action $a_i = 0$ such that the transition probability for any $s'_i \in \mathcal{S}_i$ and $s \in \mathcal{S}$

$$P_{i,0}(s'_i|s) = P_i(s'_i|s, a_i = 0). \quad (1)$$

Assumption 3.2. Given Γ and a joint state $s \in \mathcal{S}$, the uncontrolled transition probabilities of agents' sub-states are equal, i.e., $P_{i,0}(s'_i|s) = P_{j,0}(s'_j|s)$ for all $s'_i = s'_j$ sub-state pairs and for all agent pairs $i, j \in \mathcal{N}$.

Given Assumptions 3.1 and 3.2, the uncontrolled transition probability function $P_0 : \mathcal{S} \rightarrow \Delta(\mathcal{S})$ can be written for any state $s \in \mathcal{S}$

$$P_0(s'|s) = \prod_{i=1}^n P_{i,0}(s'_i|s). \quad (2)$$

In this sense, when no re-weighting is made of P , i.e. $\mathbf{a} = \mathbf{0}$, we obtain the uncontrolled transition probability function P_0 . Note that, given Assumption 3.2, agent $i \in \mathcal{N}$ only requires the knowledge of $P_{i,0}$ for every joint state $s \in \mathcal{S}$ in order to compute P_0 .

Let $\pi_{P_0} \in \Pi$ be the joint stochastic policy, or simply the joint policy, from the class of model-dependent stochastic policies Π given the uncontrolled transition probability function P_0 . The joint policy is equal to a re-weighting of P with a continuous-valued action profile such that for any $s, s' \in \mathcal{S}$, the weighted transition probability $P(s'|s, \mathbf{a}) = \pi_{P_0}(s'|s)$. Hence, the joint policy given P_0 is the mapping $\pi_{P_0} : \mathcal{S} \rightarrow \Delta(\mathcal{S})$.

Assumption 3.3. Given any joint state pair $s, s' \in \mathcal{S}$, the joint policy $\pi_{P_0}(s'|s) = 0$ when $P_0(s'|s) = 0$.

Finally, the controlled sub-state transition is done using a stochastic marginal policy $\pi_{i,P_0} \in \Pi_i : \mathcal{S} \rightarrow \Delta(\mathcal{S}_i)$ that is derived from π_{P_0} for all $i \in \mathcal{N}$.

3.2 Decomposition of the One-Step Cost Function

Similar to single-agent linearly solvable MDPs [22, 23], the control cost is captured by the KL divergence between π_{P_0} and the uncontrolled transition probability function P_0 , i.e.,

$$D_{KL}(\pi_{P_0}(\cdot|s) \| P_0(\cdot|s)) := \sum_{s' \in \mathcal{S}} \pi_{P_0}(s'|s) \ln \left(\frac{\pi_{P_0}(s'|s)}{P_0(s'|s)} \right) = \mathbb{E}_{s' \sim \pi_{P_0}(\cdot|s)} \left[\ln \left(\frac{\pi_{P_0}(\cdot|s)}{P_0(\cdot|s)} \right) \right]. \quad (3)$$

The one-step cost function is then written as $q(s, \pi_{P_0}) = C(s) + D_{KL}(\pi_{P_0}(\cdot|s) \| P_0(\cdot|s))$. We then define the value function under π_{P_0} to be state-wise

$$V^{\pi_{P_0}}(s) := \mathbb{E}_{s' \sim \pi_{P_0}(\cdot|s)} \left[\sum_{t=0}^{\infty} \gamma^t \left[C(s_t) + D_{KL}(\pi_{P_0}(\cdot|s_t) \| P_0(\cdot|s_t)) \right] \middle| s_{t=0} = s \right]. \quad (4)$$

With q_{max} being the maximum one-step cost, the value function is bounded by $q_{max}/(1 - \gamma)$. Note that if the joint policy π_{P_0} is equal to P_0 , then the KL control cost is zero and the next joint state s' is sampled according to $s' \sim P_0(\cdot|s)$, with $s' = [s'_1, s'_2, \dots, s'_n]$. Given Assumption 3.3, prohibitive transitions to another joint state in a single transition and the D_{KL} control cost boundedness conditions are met.

3.3 The KL Evaluation and KL Optimal Bellman Operators

We define the KL evaluation Bellman operator $\mathcal{T}^{\pi_{P_0}}$ for π_{P_0} applied state-wise to the value function as

$$(\mathcal{T}^{\pi_{P_0}} V)(s) = C(s) + D_{KL}(\pi_{P_0}(\cdot|s) \| P_0(\cdot|s)) + \gamma \sum_{s'} \pi_{P_0}(s'|s) V(s'). \quad (5)$$

The KL *optimal* Bellman operator, denoted as \mathcal{T} , is the Bellman operator given an optimal joint policy, i.e., $\mathcal{T}V := \min_{\pi_{P_0} \in \Pi} \mathcal{T}^{\pi_{P_0}} V$. Then, the optimal value function V^* satisfies $V^* = \mathcal{T}V^*$. Moreover, we denote the obtained optimal joint policy as $\pi_{P_0}^*$. Both the KL evaluation and the KL optimal Bellman operators have the three properties: monotonicity, distributivity, and γ -contraction. For any two arbitrary value functions $V_1, V_2 \in \mathbb{R}^{|\mathcal{S}|}$ where $V_1 \leq V_2$, $\mathcal{T}^{\pi_{P_0}}$ exhibits monotonicity such that $\mathcal{T}^{\pi_{P_0}} V_1 \leq \mathcal{T}^{\pi_{P_0}} V_2$. In addition, for any constant $c_1 \in \mathbb{R}$, and an all-ones vector $\bar{\mathbf{e}}$, the operator $\mathcal{T}^{\pi_{P_0}}$ has the distributivity property such that $\mathcal{T}^{\pi_{P_0}}(V + c_1 \cdot \bar{\mathbf{e}}) = \mathcal{T}^{\pi_{P_0}} V + \gamma \cdot c_1 \cdot \bar{\mathbf{e}}$. A single application of the Bellman operator gives γ -contractions in the L_∞ -Norm with $\|\mathcal{T}^{\pi_{P_0}} V_1 - \mathcal{T}^{\pi_{P_0}} V_2\|_\infty \leq \gamma \|V_1 - V_2\|_\infty$.

4 Multi-Agent Simulation-Based Kullback-Leibler Controlled Optimistic Policy Iteration KLC-OPI

4.1 The Scheme Description

The scheme is run by each agent $i \in \mathcal{N}$ on iterations $k = 0, 1, 2, \dots, K$ independently of other agents. At each iteration k and for agent i , the scheme stores the value function estimate $V_{i,k}$ used in evaluating the joint policy. Agent i performs a greedy policy improvement step to obtain its independently calculated joint policy $\pi_{P_0,k+1}^{(i)}$ as

$$\mathcal{G}(V_{i,k}) := \arg \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \mathcal{T}^{\pi_{P_0,k+1}^{(i)}} V_{i,k} \subseteq \mathcal{S} \rightarrow \Delta(\mathcal{S}). \quad (6)$$

Given that the joint state can be decomposed into an agent $i \in \mathcal{N}$ sub-state s_i as stated in Assumption 3.1, agent i only controls their transitions to a sub-state $s'_i \in \mathcal{S}_i$ using the marginal policy

$$\pi_{i,P_0,k+1}^{(i)}(s'_i|s) := \sum_{j \in \mathcal{N} \setminus \{i\}} \sum_{s'_j \in \mathcal{S}_j} \pi_{P_0,k+1}^{(i)}(s'_i, s'_j|s). \quad (7)$$

The agents collectively evaluate their value function estimates by generating $|\mathcal{S}|$ synchronous and coupled m -step TD sampled trajectories using their marginal policies (7) with each trajectory starting at a joint state $s \in \mathcal{S}$. The term *synchronous* denotes evaluating the value function estimates for all joint states in the state space, while the term *coupled* means that the joint state transitions are done simultaneously by all agents. The simulation-based policy evaluation step results in the per joint state cost function $q_t(s_t, \pi_{P_0,k+1}^{(i)})$ for a timestep t using the m -step TD trajectories. Each agent uses the discounted sum of the returns to obtain an estimate of $(\mathcal{T}^{\pi_{P_0,k+1}^{(i)}})^m V_{i,k}$ which we represent using the noise term $\epsilon_{m,k} \in \mathbb{R}^{|\mathcal{S}|}$ as follows,

$$(\mathcal{T}^{\pi_{P_0,k+1}^{(i)}})^m V_{i,k} + \epsilon_{m,k} = \sum_{t=0}^{m-1} \gamma^t q_t(s_t, \pi_{P_0,k+1}^{(i)}) + \gamma^m V_{i,k}(s_{t=m}). \quad (8)$$

Agent i 's value function estimate is updated using the noisy returns and results in the estimate $V_{i,k+1}$ at the end of iteration k . We summarize the simulation-based KLC-OPI scheme run by each agent $i \in \mathcal{N}$ in the following:

$$\text{KLC-OPI} \begin{cases} \pi_{P_0,k+1}^{(i)} &= \mathcal{G}(V_{i,k}), \\ V_{i,k+1} &= (I - A_k)V_{i,k} + A_k((\mathcal{T}^{\pi_{P_0,k+1}^{(i)}})^m V_{i,k} + \epsilon_{m,k}), \end{cases} \quad (9)$$

where I is the $|\mathcal{S}| \times |\mathcal{S}|$ identity matrix, A_k is a $|\mathcal{S}| \times |\mathcal{S}|$ diagonal matrix with joint state learning rates $\alpha_k(s) \in \mathbb{R}^+$ as its elements, and the rollout value is $m \in \mathbb{N}^+$. Note that $m = 1$ corresponds to value iteration with KL control cost KLC-VI, and letting $m \rightarrow \infty$ gives a policy iteration with KL control cost KLC-PI rendition of the scheme. In addition, the learning rates' matrix A_k is the same for all agents.

Similar to [Appendix 1.2 in [23]] for the single-agent linearly solvable MDP case, we show that the optimal joint policy (6) follows a Boltzmann distribution by applying the Cole-Hopf transformation to the Bellman equation with $\mathcal{T}^{\pi_{P_0,k+1}^{(i)}}$.

Lemma 4.1. *At iteration k and for any $i \in \mathcal{N}$, the joint policy $\pi_{P_0,k+1}^{(i)}$ given P_0 that minimizes the discounted return of the policy evaluation step in (9) follows a Boltzmann distribution*

$$\pi_{P_0,k+1}^{(i)}(s'|s) = \frac{P_0(s'|s)(Z_{i,k}(s'))^\gamma}{\sum_{s' \in \mathcal{S}} P_0(s'|s)(Z_{i,k}(s'))^\gamma}, \quad (10)$$

where $Z_{i,k}$ for agent i is the Cole-Hopf transformation of the value function such that $Z_{i,k}(s) = e^{-V_{i,k}(s)}$ state-wise in iteration k .

Proof. We have state-wise

$$\begin{aligned}
 \mathcal{T}V_{i,k}(s) &= \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \{C(s) + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} \left[\ln \frac{\pi_{P_0,k+1}^{(i)}(\cdot|s)}{P_0(\cdot|s)} \right] + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} [\gamma V_{i,k}(\cdot)]\} \\
 &= \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \{C(s) + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} \left[\ln \frac{\pi_{P_0,k+1}^{(i)}(\cdot|s)}{P_0(\cdot|s)} \right] + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} \left[\ln \frac{1}{(Z_{i,k}(\cdot))^\gamma} \right]\} \\
 &= \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \{C(s) + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} \left[\ln \frac{\pi_{P_0,k+1}^{(i)}(\cdot|s)}{P_0(\cdot|s)(Z_{i,k}(\cdot))^\gamma} \right]\}. \tag{11}
 \end{aligned}$$

Define the constant

$$d_{i,P_0,k}(s; \gamma) := \sum_{s' \in \mathcal{S}} P_0(s'|s)(Z_{i,k}(s'))^\gamma.$$

The value function under the joint policy that achieves the minimum discounted cost becomes

$$\begin{aligned}
 \mathcal{T}V_{i,k}(s) &= \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \{C(s) + \mathbb{E}_{s' \sim \pi_{P_0,k+1}^{(i)}(\cdot|s)} \left[\ln \frac{\pi_{P_0,k+1}^{(i)}(\cdot|s)}{\frac{P_0(\cdot|s)(Z_{i,k}(\cdot))^\gamma}{d_{i,P_0,k}(s; \gamma)}} \right] - \ln d_{i,P_0,k}(s; \gamma)\} \\
 &= \min_{\pi_{P_0,k+1}^{(i)} \in \Pi} \{C(s) - \ln d_{i,P_0,k}(s; \gamma) + D_{KL}(\pi_{P_0,k+1}^{(i)}(\cdot|s) \parallel \frac{P_0(\cdot|s)(Z_{i,k}(\cdot))^\gamma}{d_{i,P_0,k}(s; \gamma)})\}. \tag{12}
 \end{aligned}$$

Note that the first two terms ($C(s)$ and $d_{i,P_0,k}(s; \gamma)$) do not depend on $\pi_{P_0,k+1}^{(i)}$. Thus, the minimum is achieved when the last term, the KL cost, in (12) is equal to zero. The optimal joint policy calculated by each agent $i \in \mathcal{N}$ is then given by (10). \square

Remark 4.2. We let the initial value estimates be the same for all agents, $V_{i,0} = V_{j,0}$ for all $i, j \in \mathcal{N}$. Given the initialization, the value function estimate $V_{i,k}$ for agent i and thus the joint policy $\pi_{P_0,k}^{(i)}$ computed using (10) is identical to other agents' estimates for all the iterations.

Remark 4.3. The noise value $\epsilon_{m,k}$ at iteration k in (8) captures two sources of error: the error from simulating only a single trajectory per joint state instead of averaging over an infinite number of trajectories and the estimation error term incurred from the estimated value function of the last visited joint state in the simulated trajectory.

4.2 Asymptotic Convergence

Our main result is that each agent's value function estimates $V_{i,k}$ and individually calculated joint policy $\pi_{P_0,k}^{(i)}$ converge to the optimal value function and an optimal joint policy respectively under the KLC-OPI iterations.

Theorem 4.4. *Assume the initial value function estimate is such that $\mathcal{T}V_{i,0} - V_{i,0} \leq \mathbf{0}$ for each agent $i \in \mathcal{N}$ and $\alpha_k(s) = O(\frac{1}{k})$ state-wise. The value function estimate $V_{i,k}$ and joint policy $\pi_{P_0,k}^{(i)}$ iterates of $i \in \mathcal{N}$ for the KLC-OPI scheme in (9) asymptotically converge to V^* and to $\pi_{P_0}^*$, respectively.*

The convergence to optimal value function builds on two main assumptions. First one is that the initialized value function estimates $i \in \mathcal{N}$ satisfy $\mathcal{T}V_{i,k=0} - V_{i,k=0} \leq \mathbf{0}$. The assumption is not restrictive since the initial value function estimate can be set to a large value state-wise such that a single application of the KL optimal Bellman operator will guarantee a lower value function estimate. Our second assumption on the step size is standard in stochastic approximation and ensures that the steps are square summable but not summable [24, 28].

Similar to optimistic policy iteration with a deterministic policy [Proposition 1 in [19]], applying the operator $\mathcal{T}^{\pi_{P_0,k+1}^{(i)}}$ $m > 1$ times in a single iteration of KLC-OPI does not guarantee a contraction in any norm nor monotonic improvement of the value function estimate, unlike exact policy iteration schemes where the monotonicity property is preserved [3, 15]. Hence, the KLC-OPI simulation-based policy evaluation step may result in a worse performing value function estimate, i.e., a larger discounted cost return, compared to the previous iteration's estimate.

Given the non-contracting, and in general without additional assumptions, non-monotonic improvement challenges of optimistic policy iteration schemes' updates, the updated value functions' performance has to be bounded using an analysis technique that does not utilize the contraction property. To prove asymptotic convergence to the optimal

value function, we show that the scheme's iterative updates are upper and lower bounded by two standard stochastic approximation processes that we later show both converge to the optimal value function for each agent.

We begin by defining the filtration that captures the algorithm's history up to and before the noise $\epsilon_{m,k}$ at the end of iteration k is realized,

$$\mathcal{F}_{m,k} := \{(\epsilon_{m,k'})_{k' \leq k-1}\}. \quad (13)$$

Given the filtration, it follows that the noise term has zero mean, $\mathbb{E}[\epsilon_{m,k} | \mathcal{F}_{m,k}] = \mathbf{0}$, and bounded variance $\mathbb{E}[\|\epsilon_{m,k}\|_\infty | \mathcal{F}_{m,k}] \leq \frac{q_{\max} \bar{\mathbf{e}}}{1-\gamma}$ state-wise.

Next, we state the upper bound on the policy evaluation step performance in a single iteration which is a restatement of Lemma 2.c in [24] with our notation.

Lemma 4.5 (Lemma 2.c in [24]). *For the policy evaluation step at iteration k in the KLC-OPI scheme, and given the estimate of agent $i \in \mathcal{N}$ value function $V_{i,k}$, we have $(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^m V_{i,k} \leq V_{i,k} + \frac{r_{i,k} \bar{\mathbf{e}}}{1-\gamma}$, where $r_{i,k} := \max_{s \in \mathcal{S}} [\mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k}(s) - V_{i,k}(s)]$ and $\bar{\mathbf{e}}$ is an all-ones vector.*

We now prove that the KLC-OPI policy evaluation step output in (9) is upper bounded by a standard stochastic approximation process which results in the asymptotic improvement property $\limsup_{k \rightarrow \infty} \mathcal{T} V_{i,k} - V_{i,k} \leq \mathbf{0}$ for each agent $i \in \mathcal{N}$.

Proposition 4.6. *If $\mathcal{T} V_{i,k=0} - V_{i,k=0} \leq \mathbf{0}$ and the learning rate state-wise is $\alpha_k(s) = \mathcal{O}(\frac{1}{k})$ such that $0 \leq \alpha_k(s) \leq 1$ for any iteration k , then the policy evaluation step for any $m \in \mathbb{N}^+$ in (9) gives*

$$\limsup_{k \rightarrow \infty} \mathcal{T} V_{i,k} - V_{i,k} \leq \mathbf{0}. \quad (14)$$

Proof. Note that $\mathcal{T}^{\pi_{p_0,k+1}^{(i)}}$ is an affine transformation such that $\mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} = q_i^{\pi_{p_0,k+1}^{(i)}} + \gamma \pi_{p_0,k+1}^{(i)} V_{i,k}$. We then have the following for $k = 1$, using the policy evaluation update rule in (9),

$$\begin{aligned} \mathcal{T}^{\pi_{p_0,1}^{(i)}} V_{i,1} &= q_i^{\pi_{p_0,1}^{(i)}} + \gamma \pi_{p_0,1}^{(i)} \left[(I - A_0) V_{i,0} + A_0 \left((\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} + \epsilon_{m,0} \right) \right] \\ &= q_i^{\pi_{p_0,1}^{(i)}} + \gamma \pi_{p_0,1}^{(i)} V_{i,0} - A_0 \gamma \pi_{p_0,1}^{(i)} V_{i,0} + A_0 \gamma \pi_{p_0,1}^{(i)} (\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} + A_0 \gamma \pi_{p_0,1}^{(i)} \epsilon_{m,0}, \end{aligned} \quad (15)$$

which we rewrite as follows using the fact that $(\mathcal{T}^{\pi_{p_0,1}^{(i)}})^{m+1} V_{i,0} = \mathcal{T}^{\pi_{p_0,1}^{(i)}} (\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} = q_i^{\pi_{p_0,1}^{(i)}} + \gamma \pi_{p_0,1}^{(i)} (\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0}$, and by reorganizing the terms,

$$\mathcal{T}^{\pi_{p_0,1}^{(i)}} V_{i,1} = (I - A_0) (\mathcal{T}^{\pi_{p_0,1}^{(i)}} V_{i,0}) + A_0 \left((\mathcal{T}^{\pi_{p_0,1}^{(i)}})^{m+1} V_{i,0} + \gamma \pi_{p_0,1}^{(i)} \epsilon_{m,0} \right). \quad (16)$$

We know that the optimal Bellman operator is going to be better, i.e., $\mathcal{T} V_{i,1} \leq \mathcal{T}^{\pi_{p_0,1}^{(i)}} V_{i,1}$. Let $y_{i,k} := \mathcal{T} V_{i,k} - V_{i,k}$. Subtracting $V_{i,1}$ from both sides of (16), and using the fact that $\mathcal{T} V_{i,k} = \mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k}$ due to the greedy step in (9), we have

$$\begin{aligned} y_{i,1} &\leq (I - A_0) (\mathcal{T} V_{i,0}) + A_0 \left((\mathcal{T}^{\pi_{p_0,1}^{(i)}})^{m+1} V_{i,0} + \gamma \pi_{p_0,1}^{(i)} \epsilon_{m,0} \right) - (I - A_0) V_{i,0} - A_0 \left((\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} + \epsilon_{m,0} \right) \\ &= (I - A_0) y_{i,0} + A_0 \left[(\gamma \pi_{p_0,1}^{(i)} - I) \epsilon_{m,0} \right] + A_0 \left((\mathcal{T}^{\pi_{p_0,1}^{(i)}})^{m+1} V_{i,0} - (\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} \right). \end{aligned} \quad (17)$$

Since we assume that $\mathcal{T} V_{i,0} \leq V_{i,0}$, i.e., $y_{i,0} \leq \mathbf{0}$, then applying the operator m times on both sides gives $(\mathcal{T}^{\pi_{p_0,1}^{(i)}})^{m+1} V_{i,0} - (\mathcal{T}^{\pi_{p_0,1}^{(i)}})^m V_{i,0} \leq \mathbf{0}$. Thus we obtain the upper bound

$$y_{i,1} \leq A_0 \left[(\gamma \pi_{p_0,1}^{(i)} - I) \epsilon_{m,0} \right]. \quad (18)$$

Define $w_k := (\gamma \pi_{p_0,k+1}^{(i)} - I) \epsilon_{m,k}$ where w_k satisfies the same properties as $\epsilon_{m,k}$, i.e., zero-mean and bounded variance, then $y_{i,1} \leq A_0 w_0 = U_1$ where U_1 is the upper bound on $y_{i,1}$. Similarly, we have the following upper bound for $k = 2$

$$y_{i,2} \leq (I - A_1) y_{i,1} + A_1 w_1 + A_1 \left((\mathcal{T}^{\pi_{p_0,2}^{(i)}})^{m+1} V_{i,1} - (\mathcal{T}^{\pi_{p_0,2}^{(i)}})^m V_{i,1} \right). \quad (19)$$

The upper bound on $y_{i,1} \leq U_1$ implies $\mathcal{T}^{\pi_{p_0,2}^{(i)}} V_{i,1} \leq V_{i,1} + U_1$. When we apply the operator $\mathcal{T}^{\pi_{p_0,k+1}^{(i)}}$ to both sides of the inequality for $k = 2$, we obtain the following using the distributivity property $\mathcal{T}^{\pi_{p_0,k+1}^{(i)}}(V_{i,k} + c \cdot \bar{\mathbf{e}}) = \mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} + \gamma \cdot c \cdot \bar{\mathbf{e}}$ for some constant $c \in \mathbb{R}$,

$$\begin{aligned} (\mathcal{T}^{\pi_{p_0,2}^{(i)}})^2 V_{i,1} &\leq (\mathcal{T}^{\pi_{p_0,2}^{(i)}})(V_{i,1} + U_1) \\ &= (\mathcal{T}^{\pi_{p_0,2}^{(i)}}) V_{i,1} + \gamma U_1. \end{aligned} \quad (20)$$

Applying $\mathcal{T}^{\pi_{p_0,2}^{(i)}}$ m times gives $(\mathcal{T}^{\pi_{p_0,2}^{(i)}})^{m+1} V_{i,1} \leq (\mathcal{T}^{\pi_{p_0,2}^{(i)}})^m V_{i,1} + \gamma^m U_1$. Then, we can replace the last term in (19) with $\gamma^m U_1$ to get the following bound

$$\begin{aligned} y_{i,2} &\leq (I - A_1)U_1 + A_1 w_1 + A_1 \gamma^m U_1 \\ &= (I + A_1(\gamma^m I - I))U_1 + A_1 w_1. \end{aligned} \quad (21)$$

We let $U_2 := (I + A_1(\gamma^m I - I))U_1 + A_1 w_1$. Repeating the above steps for $k = 3, 4, \dots, K$ we obtain the upper bound for any k

$$\begin{aligned} (\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^{m+1} V_{i,k} - (\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^m V_{i,k} &\leq \gamma^m U_k \\ &= \gamma^m [(I + A_{k-1}(\gamma^m I - I))U_{k-1} + A_{k-1} w_{k-1}], \end{aligned} \quad (22)$$

with $U_0 = \mathbf{0}$. Since the noise is zero-mean with bounded variance, $\gamma \leq 1$, and $\alpha_k(s) = \mathcal{O}(1/k)$ with $0 \leq \alpha_k(s) \leq 1$ state-wise for any k , we have that

$$\lim_{k \rightarrow \infty} \alpha_k [(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^{m+1} V_{i,k} - (\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^m V_{i,k}] \leq \gamma^m \lim_{k \rightarrow \infty} \frac{U_k}{k} = \mathbf{0}. \quad (23)$$

Define a noise function as $f(w_k, \gamma^m, U_k) := w_k + \gamma^m U_k$, then we have for any $i \in \mathcal{N}$ and k

$$y_{i,k+1} \leq (I - A_k)y_{i,k} + A_k f(w_k, \gamma^m, U_k). \quad (24)$$

Define another sequence $X_{k+1} = (I - A_k)X_k + A_k f(w_k, \gamma^m, U_k)$ with $X_0 = y_{i,0}$ for any $i \in \mathcal{N}$. We have that $y_{i,k} \leq X_k$. Given (23) and $f(w_k, \gamma^m, U_k)$ is a zero-mean noise function, the sequence $\{X_k\}_{0,1,\dots}$ is a standard stochastic sequence that converges to zero in the limit. Finally, we have $\lim_{k \rightarrow \infty} X_k = \mathbf{0}$ and since $y_{i,k} \leq X_k$, we obtain $\limsup_{k \rightarrow \infty} y_{i,k} \leq \mathbf{0}$. \square

Establishing the asymptotic policy improvement property of (9) is the key contribution of our proof approach. This result is similar to the property obtained in [24] for showing convergence of OPI with infinitely long single trajectories. However, here we consider finitely long trajectories which is in agreement with practical implementation. In the proof, we overcome the error introduced by finite-trajectories by upper bounding the noisy and finite rollout with a term that only depends on the discount factor, learning rates, the joint policy from the previous iteration, and the realized noise up to iteration k . Moreover, we achieve this property using greedy policy improvements (6), i.e., without the need to perform a lookahead operation for each iteration as is done in [28, 26].

We are ready to provide the proof of Theorem 4.4.

Proof (Theorem 4.4). Given the asymptotic improvement guarantee in (14), for every $\delta > 0$, there exists an iteration $k(\delta)$ such that $r_{i,k} \leq \delta$ for $k \geq k(\delta)$ where we recall that $r_{i,k} = \max_{s \in S} [\mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} - V_{i,k}]$. Using the result from Lemma 4.5, we have $(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^{m-1} V_{i,k} \leq V_{i,k} + \frac{r_{i,k} \cdot \bar{\mathbf{e}}}{1-\gamma}$. Then applying the operator on both sides gives $(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^m V_{i,k} \leq \mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} + \frac{\gamma \cdot r_{i,k} \cdot \bar{\mathbf{e}}}{1-\gamma} \leq \mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} + \frac{\gamma \cdot \delta \cdot \bar{\mathbf{e}}}{1-\gamma}$ for $k \geq k(\delta)$.

We then write the update rule from (9) for $k \geq k(\delta)$ as

$$\begin{aligned} V_{i,k+1} &= (I - A_k)V_{i,k} + A_k \left((\mathcal{T}^{\pi_{p_0,k+1}^{(i)}})^m V_{i,k} + \epsilon_{m,k} \right) \\ &\leq (I - A_k)V_{i,k} + A_k \left(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}} V_{i,k} + \frac{\gamma \cdot \delta \cdot \bar{\mathbf{e}}}{1-\gamma} + \epsilon_{m,k} \right). \end{aligned} \quad (25)$$

For every positive integer \bar{k} , define the random variable sequence

$$W_{i,k+1,\delta}^{(\bar{k})} = (I - A_k)W_{i,k,\delta}^{(\bar{k})} + A_k \left(\mathcal{T}^{\pi_{p_0,k+1}^{(i)}} W_{i,k,\delta}^{(\bar{k})} + \frac{\gamma \cdot \delta \cdot \bar{\mathbf{e}}}{1-\gamma} + \epsilon_{m,k} \right), \quad (26)$$

for all $k \geq \bar{k}$. In addition, let $A^{(\bar{k})}$ be the event that $k(\delta) = \bar{k}$. Then for any \bar{k} , any sample path in $A^{(\bar{k})}$, and for all $k \geq \bar{k}$, we have $V_{i,k+1} \leq W_{i,k+1,\delta}^{(\bar{k})}$. Next, define the mappings $\mathcal{H} : \mathbb{R}^{|\mathcal{S}|} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ and $\mathcal{H}^{\pi_{P_0,k+1}^{(i)}} : \mathbb{R}^{|\mathcal{S}|} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ such that $\mathcal{H}^{\pi_{P_0,k+1}^{(i)}} V_{i,k} = \mathcal{T}^{\pi_{P_0,k+1}^{(i)}} V_{i,k} + \frac{\gamma\delta\bar{\epsilon}}{1-\gamma}$. Rewriting the sequence in Equation (26) using $\mathcal{H}^{\pi_{P_0,k+1}^{(i)}}$ gives

$$W_{i,k+1,\delta}^{(\bar{k})} = (I - A_k)W_{i,k,\delta}^{(\bar{k})} + A_k\left(\mathcal{H}^{\pi_{P_0,k+1}^{(i)}} W_{i,k,\delta}^{(\bar{k})} + \epsilon_{m,k}\right). \quad (27)$$

Since a single application of $\mathcal{T}^{\pi_{P_0,k+1}^{(i)}}$ is a contraction mapping, a single application of $\mathcal{H}^{\pi_{P_0,k+1}^{(i)}}$ is also a contraction mapping with a unique fixed point $W_\delta^* = V^* + \frac{\gamma\delta\bar{\epsilon}}{(1-\gamma)^2}$. Given that $\epsilon_{m,k}$ is a zero-mean noise, the sequence in Equation (27) converges almost surely to W_δ^* for sample paths in $A^{(\bar{k})}$. Since the union of the events $A^{(\bar{k})}$ is the entire sample space, we have $\limsup_{k \rightarrow \infty} V_{i,k} \leq \limsup_{k \rightarrow \infty} W_{i,k,\delta}^{(\bar{k})} = W_\delta^*$. Given that δ can be chosen arbitrarily close to zero, then $\limsup_{k \rightarrow \infty} V_{i,k} \leq V^*$.

Similarly, since $(\mathcal{T}^{\pi_{P_0,k+1}^{(i)}})^m V_{i,k} \geq V^*$, we have

$$\begin{aligned} V_{i,k+1} &= (I - A_k)V_{i,k} + A_k\left((\mathcal{T}^{\pi_{P_0,k+1}^{(i)}})^m V_{i,k} + \epsilon_{m,k}\right) \\ &\geq (I - A_k)V_{i,k} + A_k(V^* + \epsilon_{m,k}). \end{aligned} \quad (28)$$

Define another random variable sequence $B_{i,k+1} = (I - A_k)B_{i,k} + A_k(V^* + \epsilon_{m,k})$ with $B_{i,k=0} = V_{i,k=0}$. The sequence $B_{i,k}$ converges asymptotically to V^* and we have $\liminf_{k \rightarrow \infty} V_{i,k} \geq V^*$. Given the asymptotic upper and lower convergence bounds to V^* , the update rule in (9) also converges to the optimal value function V^* for each $i \in \mathcal{N}$ and we obtain an optimal joint policy $\pi_{P_0}^*$ using (10). \square

Note that we do not use any structural feature of multi-agent MDPs with KL control cost in the proof, and thus the convergence result for Theorem 4.4 holds for general MDPs.




Figure 2: ASYNC-KLC-OPI performance on the multi-agent MDP with KL control cost variant of the Stag-Hare game averaged over 10 simulation runs.

5 The Asynchronous KLC-OPI Scheme

For simulation-based policy iteration schemes, it can be computationally expensive to run a synchronous policy evaluation step with $|\mathcal{S}|$ m -step TD sampled trajectories especially if the joint state space is large. For this reason, we focus on the asynchronous policy evaluation step of value function estimates using $1 \leq D \leq |\mathcal{S}|$ joint states in each iteration k . In a single iteration, a larger D value means that a larger subspace of the joint state space is evaluated at the expense of additional sampling per iteration.

In order to prove the convergence of the asynchronous version of the scheme, we assume that the initial state distribution ρ is a uniform distribution. Let $\mathcal{D}_k \subseteq \mathcal{S}$ be the set of joint states that are evaluated at iteration k under the current joint policy $\pi_{P_0,k+1}^{(i)}$ for any $i \in \mathcal{S}$. The set \mathcal{D}_k of each iteration satisfies the following assumption.

Assumption 5.1. The cardinality of the sampled joint states' set $D = |\mathcal{D}_k|$ is a constant with the set having unique joint state elements in each iteration $k = 0, 1, 2, \dots, K$.

Assumption 5.1 ensures that no joint state is evaluated more than once per iteration, and that the number of evaluated joint states is fixed in the proposed asynchronous policy evaluation step.

Similar to the synchronous policy evaluation step in (9), each joint state $s \in \mathcal{D}_k$ is evaluated using its independently simulated m -step TD trajectory and its corresponding value function estimate is updated at the end of the current iteration. For joint states that are not in the set \mathcal{D}_k , their value function estimates are kept the same. The asynchronous KLC-OPI scheme is then

$$\text{ASYNC-KLC-OPI} \begin{cases} \pi_{P_0, k+1}^{(i)} &= \mathcal{G}(V_{i,k}), \\ V_{i, k+1}(s) &= (1 - \alpha_k(s))V_{i,k}(s) + \alpha_k(s) \left((\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k}(s) + \epsilon_{m,k}(s) \right) & \forall s \in \mathcal{D}_k, \\ V_{i, k+1}(s) &= V_{i,k}(s) & \forall s \notin \mathcal{D}_k. \end{cases} \quad (29)$$

Note that the policy improvement step taken by each agent is with respect to all joint state elements of the current value function estimate regardless of whether a joint state's value estimate was evaluated or not. We can re-write the policy evaluation step in (29) as follows

$$V_{i, k+1}(s) = \mathcal{I}_{s \in \mathcal{D}_k} \left[(1 - \alpha_k(s))V_{i,k}(s) + \alpha_k(s) \left[(\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k}(s) + \epsilon_{m,k}(s) \right] \right] + \mathcal{I}_{s \notin \mathcal{D}_k} \left[V_{i,k}(s) \right], \quad (30)$$

where $\mathcal{I}_{s \in \mathcal{D}_k}$ is the indicator function that specifies if the joint state s is evaluated in iteration k or not. In the following, we establish the same asymptotic convergence guarantee as in Theorem 4.4 for the iterates in ASYNC-KLC-OPI.

Corollary 5.2. Assume that the initial value function is such that $\mathcal{T}V_{i,0} - V_{i,0} \leq \mathbf{0}$ for each agent $i \in \mathcal{N}$, $\alpha_k(s) = \mathcal{O}(\frac{1}{k})$ state-wise, and that the initial joint state distribution ρ is a uniform distribution, then the ASYNC-KLC-OPI scheme in (29) value function $V_{i,k}$ and joint policy $\pi_{P_0, k}^{(i)}$ iterates asymptotically converge to V^* and to $\pi_{P_0}^*$ for all $i \in \mathcal{N}$, respectively.

Proof. We can write the policy evaluation update rule of the asynchronous scheme in (29) similar fashion to the synchronous case (9),

$$V_{i, k+1} = (I - A_k H) V_{i,k} + A_k H \left((\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k} + g_{m,k} \right), \quad (31)$$

with H being the $|\mathcal{S}| \times |\mathcal{S}|$ diagonal joint state policy evaluation probability matrix, i.e. the matrix H element $h(s) > 0$ state-wise is the stationary probability that the joint state s is sampled according to. The added noise function $g_{m,k}$ is given by

$$g_{m,k} := \epsilon_{m,k} + \left((H^{-1} \mathcal{X}_k - I) \cdot (-V_{i,k} + (\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k} + \epsilon_{m,k}) \right), \quad (32)$$

with \mathcal{X}_k being a $|\mathcal{S}| \times |\mathcal{S}|$ diagonal matrix with Bernoulli random variable elements $x_k(s)$ state-wise at iteration k . The random variable $x_k(s) = 1$ if the joint state s is evaluated during iteration k , and it is zero otherwise. Since we assume that the initial state distribution ρ is a uniform distribution in the asynchronous case, the expected value of the Bernoulli random variables is $\mathbb{E}[x_k(s) | \mathcal{F}_{m,k}] = h(s)$ state-wise. We then have $H^{-1} \cdot \mathbb{E}[\mathcal{X}_k | \mathcal{F}_{m,k}] = H^{-1} \cdot H = I$, which means that $g_{m,k}$ is also a zero-mean noise function such that $\mathbb{E}[g_{m,k} | \mathcal{F}_{m,k}] = \mathbf{0}$. Similar to the noise function $\epsilon_{m,k}$, the added noise function $g_{m,k}$ variance is also bounded since the value function estimate and $(\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k}$ in Equation (32) are bounded by the value $\frac{q_{\max}}{1-\gamma}$ state-wise.

Given Assumption 5.1, the set's cardinality $|\mathcal{D}_k|$ is a constant throughout the iterations, then the diagonal matrix H has equal and constant elements throughout the iterations $k = 0, 1, 2, \dots, K$. In other words, the matrix elements $h(s) = h(s')$ for all $(s, s') \in \mathcal{S}$ pairs. Define $\mathcal{A}_k = A_k \cdot H$ as the product of the diagonal learning rate matrix A_k and the diagonal joint state policy evaluation probability matrix H . We can then write the policy evaluation step as

$$V_{i, k+1} = (I - \mathcal{A}_k) V_{i,k} + \mathcal{A}_k \left((\mathcal{T}^{\pi_{P_0, k+1}^{(i)}})^m V_{i,k} + g_{m,k} \right). \quad (33)$$

Since H has constant elements throughout the iterations, then \mathcal{A}_k is a matrix that commutes with the joint policy $\pi_{P_0, k+1}^{(i)}$ in each iteration such that $\mathcal{A}_k \cdot \pi_{P_0, k+1}^{(i)} = \pi_{P_0, k+1}^{(i)} \cdot \mathcal{A}_k$. Since $\mathcal{A}_k \cdot \pi_{P_0, k+1}^{(i)}$ satisfy the commutative property and the added noise function $g_{m,k}$ also has zero mean with bounded variance, we can directly use the same steps as in Proposition 4.6 to show that $\limsup_{k \rightarrow \infty} \mathcal{T}V_{i,k} - V_{i,k} \leq \mathbf{0}$ for each agent $i \in \mathcal{N}$. The convergence to the optimal value function V^* for each agent and to an optimal joint policy $\pi_{P_0}^*$ follows the same steps as in Theorem 4.4. \square

6 Simulations

6.1 A Multi-Agent MDP: Stag-Hare

We consider a two-agent MDP with KL control cost variant [12] of the Stag-Hare game [21]. There are $|\mathcal{N}| = 2$ hunters who hunt on a 5×5 gridworld and can only move to adjacent grids in a single timestep.

Hunter i 's sub-state is their grid location $s_i = 0, 1, \dots, 24$ such that there are $|\mathcal{S}_i| = 25$ possible grid locations. For the two hunters, this gives a total of $|\mathcal{S}| = |\mathcal{S}_1| \times |\mathcal{S}_2| = 625$ joint states. The uncontrolled transition probability $P_{i,0}$ forces agent $i \in \mathcal{N}$ to remain in their current sub-state w.p. 0.9, and transition to one of the b adjacent sub-states w.p. $0.1/b$.

The gridworld has four hares and one stag at sub-state locations $s_h \in \{0, 4, 20, 24\}$ and $s_s \in \{12\}$, respectively. When one of the hunters reaches a hare's location, both hunters obtain a negative intrinsic joint state cost of -2 added to the KL control cost per timestep. If both hunters cooperate and move together to the stag's location, they would obtain a lower intrinsic joint state cost of -10 added to the KL control cost per timestep. The intrinsic joint state cost incurred for each agent $i \in \mathcal{N}$ in a single transition is

$$C(s) = -2 \cdot \sum_{i=1}^n \mathcal{I}\{s_i \in s_h\} - 10 \cdot \mathcal{I}\left\{\left[\sum_{i=1}^n \mathcal{I}\{s_i \in s_s\}\right] > 1\right\}, \quad (34)$$

where $\mathcal{I}\{\cdot\}$ being the indicator function.

6.2 ASYNC-KLC-OPI Simulation Results

We set the discount factor to be $\gamma = 0.95$, and the fixed episode time horizon m to be the average of a geometrically distributed random variable with distribution $\text{geom}(1-\gamma)$ such that $m = 20$. We test the scheme using $D = [20, 40, 60, 80]$ joint states for $K = 3000$ iterations and show averaged results over 10 simulation runs in Figure (2).

In Figure 2(a), ASYNC-KLC-OPI converges faster to a minimum cost return as the number of sampled joint states D in a single iteration K increases. With a larger D value, the scheme evaluates a larger subset of the joint state space in one iteration, and as a result it obtains a better joint policy for the next iteration compared with a smaller D value. However, the average runtime per iteration increases as D increases and is 6.66, 14.57, 16.18, and 20.75 seconds for the selected D values, respectively.

Figure 2(b) compares the obtained stochastic policy for $D = 80$ joint states against the optimal deterministic joint policy. The optimal deterministic joint policy executes grid transitions that direct the agents to the stag's sub-state in the lowest number of time steps. It can be seen from the figure that the stochastic joint policy gives a slightly better performance for the joint states $[20, 4]$, $[5, 12]$ and $[18, 14]$ in terms of the undiscounted cost return. For the joint state $[11, 13]$ where each hunter starts in a cell neighboring the stag, the two policies achieve a similar performance. The cost return difference is due to the increased KL cost when using the optimal deterministic joint policy. Agents using the optimal deterministic joint policy transition to a selected joint state w.p. 1 which results in a higher KL cost compared with a stochastic joint policy.

For Figure 2(c), we plot the L_∞ -Norm for the difference between iteration k value function estimates and the value function from iteration $K = 3000$. As shown, the convergence rate to the iteration $K = 3000$ value function estimate is faster with a larger D value.

7 Conclusion

We presented a synchronous and an asynchronous simulation-based optimistic policy iteration schemes for multi-agent MDPs with KL control costs that are run independently by each agent. The separation between control costs and joint state costs rendered the optimal joint policy to have a close-form solution in the form of a Boltzmann distribution that depends on the current value function estimate and uncontrolled transition probabilities. Given standard assumptions on the learning rates and the initial value function estimate, we showed the asymptotic convergence of both schemes to the optimal value function and an optimal joint policy. The convergence result applies to any simulation-based OPI scheme with finite and noisy rollout returns on any MDP. For different number of sampled joint states in an iteration, simulation results on a multi-agent MDP variant of the Stag-Hare game showed that the asynchronous scheme converges to a minimum cost return for the agents, with better performance than the optimal deterministic joint policy.

References

- [1] Sarper Aydin and Ceyhun Eksin. Policy gradient play over time-varying networks in markov potential games. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1997–2002. IEEE, 2023.
- [2] Dimitri Bertsekas. Distributed asynchronous policy iteration for sequential zero-sum games and minimax control. *arXiv preprint arXiv:2107.10406*, 2021.
- [3] Dimitri Bertsekas and John N Tsitsiklis. *Neuro-dynamic programming*. Athena Scientific, 1996.
- [4] Sarnaduti Brahma, Yitao Bai, Duy Anh Do, and Thinh T. Doan. Convergence rates of asynchronous policy iteration for zero-sum markov games under stochastic and optimistic settings. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 3493–3498, 2022.
- [5] Neil Cammardella, Ana Bušić, and Sean P. Meyn. Kullback–leibler-quadratic optimal control. *SIAM Journal on Control and Optimization*, 61(5):3234–3258, 2023.
- [6] Yuanlong Chen. On the convergence of optimistic policy iteration for stochastic shortest path problem. *arXiv preprint arXiv:1808.08763*, 2018.
- [7] Dongsheng Ding, Chen-Yu Wei, Kaiqing Zhang, and Mihailo Jovanovic. Independent policy gradient for large-scale Markov potential games: Sharper rates, function approximation, and game-agnostic convergence. In *International Conference on Machine Learning*, pages 5166–5220. PMLR, 2022.
- [8] Krishnamurthy Dvijotham and Emo Todorov. Linearly solvable markov games. In *2012 American Control Conference (ACC)*, pages 1845–1850. IEEE, 2012.
- [9] Yonathan Efroni, Gal Dalal, Bruno Scherrer, and Shie Mannor. How to combine tree-search methods in reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3494–3501, 2019.
- [10] János Flesch, Gijs Schoenmakers, and Koos Vrieze. Stochastic games on a product state space. *Mathematics of Operations Research*, 33(2):403–420, 2008.
- [11] Xin Guo, Xinyu Li, Chinmay Maheshwari, Shankar Sastry, and Manxi Wu. Markov α -potential games: Equilibrium approximation and regret analysis. *arXiv preprint arXiv:2305.12553*, 2023.
- [12] Hilbert J Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87:159–182, 2012.
- [13] Stefanos Leonardos, Will Overman, Ioannis Panageas, and Georgios Piliouras. Global convergence of multi-agent policy gradient in Markov potential games. *arXiv preprint arXiv:2106.01969*, 2021.
- [14] Jun Liu. On the convergence of reinforcement learning with monte carlo exploring starts. *Automatica*, 129:109693, 2021.
- [15] Rémi Munos. Error bounds for approximate policy iteration. In *ICML*, volume 3, pages 560–567. Citeseer, 2003.
- [16] Martin L. Puterman and Moon Chirl Shin. Modified policy iteration algorithms for discounted markov decision problems. *Management Science*, 24(11):1127–1137, 1978.
- [17] Muhammed Sayin, Kaiqing Zhang, David Leslie, Tamer Basar, and Asuman Ozdaglar. Decentralized q-learning in zero-sum Markov games. *Advances in Neural Information Processing Systems*, 34:18320–18334, 2021.
- [18] Muhammed O Sayin, Francesca Parise, and Asuman Ozdaglar. Fictitious play in zero-sum stochastic games. *SIAM Journal on Control and Optimization*, 60(4):2095–2114, 2022.
- [19] Bruno Scherrer, Mohammad Ghavamzadeh, Victor Gabillon, Boris Lesner, and Matthieu Geist. Approximate modified policy iteration and its application to the game of tetris. *Journal of Machine Learning Research*, 16(49):1629–1676, 2015.
- [20] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [21] Brian Skyrms. *The stag hunt and the evolution of social structure*. Cambridge University Press, 2004.
- [22] Emanuel Todorov. Linearly-solvable markov decision problems. *Advances in neural information processing systems*, 19, 2006.
- [23] Emanuel Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009.
- [24] John N. Tsitsiklis. On the convergence of optimistic policy iteration. *J. Mach. Learn. Res.*, 3(null):59–72, mar 2003.

- [25] Onur Unlu and Muhammed O Sayin. Episodic logit-q dynamics for efficient learning in stochastic teams. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1985–1990. IEEE, 2023.
- [26] Anna Winnicki, Joseph Lubars, Michael Livesay, and R Srikant. The role of lookahead and approximate policy evaluation in reinforcement learning with linear value function approximation. *Operations Research*, 2024.
- [27] Anna Winnicki and R Srikant. Reinforcement learning with unbiased policy evaluation and linear function approximation. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 801–806. IEEE, 2022.
- [28] Anna Winnicki and R Srikant. On the convergence of policy iteration-based reinforcement learning with monte carlo policy evaluation. In *International Conference on Artificial Intelligence and Statistics*, pages 9852–9878. PMLR, 2023.
- [29] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of Reinforcement Learning and Control*, pages 321–384, 2021.
- [30] Yizhou Zhang, Guannan Qu, Pan Xu, Yiheng Lin, Zaiwei Chen, and Adam Wierman. Global convergence of localized policy iteration in networked multi-agent reinforcement learning. *Proc. ACM Meas. Anal. Comput. Syst.*, 7(1), mar 2023.