

Discrete-Time Mean Field Control with Environment States

Kai Cui, Anam Tahir, Mark Sinzger and Heinz Koepl

Abstract—Multi-agent reinforcement learning methods have shown remarkable potential in solving complex multi-agent problems but mostly lack theoretical guarantees. Recently, mean field control and mean field games have been established as a tractable solution for large-scale multi-agent problems with many agents. In this work, driven by a motivating scheduling problem, we consider a discrete-time mean field control model with common environment states. We rigorously establish approximate optimality as the number of agents grows in the finite agent case and find that a dynamic programming principle holds, resulting in the existence of an optimal stationary policy. As exact solutions are difficult in general due to the resulting continuous action space of the limiting mean field Markov decision process, we apply established deep reinforcement learning methods to solve the associated mean field control problem. The performance of the learned mean field control policy is compared to typical multi-agent reinforcement learning approaches and is found to converge to the mean field performance for sufficiently many agents, verifying the obtained theoretical results and reaching competitive solutions.

I. INTRODUCTION

Reinforcement Learning (RL) has proven to be a very successful approach for solving sequential decision-making problems [1]. Today it has numerous applications e.g. in robotics [2], strategic games [3] or communication networks [4]. Many such applications are modelled as special cases of Markov games, which has led to empirical success in the multi-agent RL (MARL) domain.

However, MARL problems quickly become intractable for large numbers of agents and proposed solutions offer few rigorous guarantees [5]. An increasingly popular approach in resolving this curse of dimensionality are mean field approximation. The main idea is to convert a many-agent system with N indistinguishable and interchangeable agents into a problem where one representative agent interacts with e.g. the empirical state distribution – the mean field – of the other agents. Since the N -agent model is reduced to a single agent and a mean field, this lends the problem tractability with theoretical guarantees for sufficiently large N .

The framework of mean field games (MFG) was first introduced in [6] and [7] for stochastic differential games and has since been extended to discrete-time [8], [9]. It provides a framework for analyzing many-agent competitive problems, for which learning-based solutions have become increasingly popular [10], [11], [12]. Mean field theory applied to the cooperative setting is known as mean field control (MFC), where one assumes that many agents cooperate to achieve

The authors are with the Department of Electrical Engineering, Technische Universität Darmstadt, 64287 Darmstadt, Germany. Contact: {kai.cui, anam.tahir, mark.sinzger, heinz.koepl}@bcs.tu-darmstadt.de




Fig. 1. Overview of the multi-agent system as a probabilistic graphical model using plate notation [23], where circles and diamonds indicate stochastic and deterministic nodes respectively. Each agent i chooses an action $u_t^{i,N}$ conditional on the environment state $x_t^{0,N}$ and local agent state $x_t^{i,N}$, influencing the next environment state $x_{t+1}^{0,N}$ only via their empirical distribution G_t^N . Agent states are assumed i.i.d. for simplicity of analysis.

Pareto optima [13], [14]. MFC has various applications e.g. in smart heating [15] or portfolio management [16].

The dimensions of the MFC problem are independent of the specific number of agents, making it more tractable. However, solving the MFC problem has the challenge of time-inconsistency due to the non-Markovian nature of the problem [13], [17], [18]. A recent way of handling this inherent time-inconsistency problem is to use an enlarged state-action space [19], [20], [21], [22]. We similarly apply this technique by lifting up the state-action space into its probability measure space, since it will enable usage of dynamic programming and established reinforcement learning methods.

In this work we extend the theory of discrete-time MFC by considering additional environment states. An advantage of discrete-time models is applicability of a plethora of reinforcement learning solutions. Our model can be considered a special case of the MFC equivalent of major-minor mean field games [24], [25] with trivial major agent policy, which to the best of our knowledge has not been formulated yet. We expect that our results can be generalized, similar to approaches e.g. in [9] for the competitive mean field game, although for deterministic mean fields.

The main contributions of this paper are: (i) We propose a new discrete-time MFC formulation that transforms large-scale multi-agent control problems with common environment states into a simple Markov decision process (MDP) with lifted state-action space; (ii) we rigorously show approximate optimality for sufficiently large systems as well as existence of an optimal stationary policy through a dynamic programming principle, and (iii) associated with this standard discrete-time MDP with continuous action space, we verify our theoretical findings empirically using modern reinforcement learning techniques. As a result, we outperform existing baselines for the many-agent case and obtain a methodology to solve large multi-agent control problems such as the following.




Fig. 2. Overview of the queuing system. Many schedulers (middle) obtain packets at a fixed rate (left) that must be assigned to one of the accessible queues (right) such that total packet drops are minimized.

II. SCHEDULING SCENARIO

While the concept of mean field limits has been used in queuing systems before, it has mostly been used for the state of the buffer fillings of queues or the number of servers/queues [26], [27]. In this work we use mean fields to represent the state of a large amount of schedulers while modeling the queues exactly. See also Figure 2 for a visualization of the problem. Note that in principle, our model could be used for any similar resource allocation problem such as allocation of many firefighters to houses on fire.

Consider a queuing system with N agents called schedulers, $\{s_1, \dots, s_N\}$, and M parallel servers, each with its own finite FIFO queue. Denote the queue filling by $b_i \in \{0, \dots, B_i\}$, $i = 1, \dots, M$ where B_i is the maximum buffer space for the i -th queue. At any time step t , the state $x_t^{i,N} \in \mathcal{X}$ of a scheduler is the set of queues it has access to. The agent state space \mathcal{X} therefore consists of all combinations of queue access where every agent has access to at least one of the queues. The environment state is the current buffer filling $x^0 = [b_1, \dots, b_M]$, where b_j is the buffer filling of queue j .

In discrete-time, the number of job arrivals to be assigned at each time step t is Poisson distributed with rate $\lambda\Delta T$ and the number of serviced jobs for each server is Poisson distributed with rate $\beta\Delta T$, where $\Delta T > 0$ can be considered the time span between each synchronization of schedulers. As an approximation, we assume that all queue departures in a time slot happen before the new arrivals, and newly arrived jobs thus cannot be serviced in the same time slot.

We split the total number of job packets which arrive in some time step ΔT uniformly at random amongst the schedulers. The jobs assigned to each scheduler need to be sent out immediately. Each scheduler decides which of the accessible queues it sends its arrived jobs to during each time step. If a job is mapped to a full buffer, it is lost and a penalty c_d is incurred. The goal of the system is therefore to minimize the number of job drops. At each step of the decision making, we assume that the state of the environment x^0 and their own accessible queues are known to the schedulers.

We can model the dynamics of the environment state dependent on the empirical state-action distribution of all schedulers: Consider agents choosing some choice of queues as their action, where inaccessible queues are treated as randomly picking a destination. In that case, to assign a packet

to its destination queue, it is clearly sufficient to consider the empirical distribution: Sampling from the empirical distribution, using the sampled action and, if inaccessible, resampling an accessible queue provides the desired behavior.

III. MEAN FIELD CONTROL

In this section, we formulate a N -agent model that in the limit of $N \rightarrow \infty$ results in a more tractable MFC problem. Importantly, we will then show approximate optimality and a dynamic programming principle for the MFC problem, allowing for application of reinforcement learning.

Notation. Let \mathcal{A} be a finite set. We equip \mathcal{A} with the discrete metric and denote the set of real-valued functions on \mathcal{A} by $\mathbb{R}^{\mathcal{A}}$. For $f \in \mathbb{R}^{\mathcal{A}}$ let $\|f\|_{\infty} = \max_{a \in \mathcal{A}} f(a)$. Denote by $|\mathcal{A}|$ the cardinality of \mathcal{A} . Denote by $\mathcal{P}(\mathcal{A}) = \{p \in \mathbb{R}^{\mathcal{A}}: p(a) \geq 0, \sum_{a \in \mathcal{A}} p(a) = 1\}$ the space of probability simplices, equivalent to the probability measures on \mathcal{A} . Equip $\mathcal{P}(\mathcal{A})$ with the l_1 -norm $\|\mu - \nu\|_1 = \sum_{a \in \mathcal{A}} |\mu(a) - \nu(a)|$. For readability, we uncurry occurrences of multiple parentheses, e.g. $\pi_t(x_t^0)(x) \equiv \pi_t(x_t^0, x)$. Define $\mu(f) := \sum_{a \in \mathcal{A}} f(a)\mu(a)$ for any $\mu \in \mathcal{P}(\mathcal{A})$, $f: \mathcal{A} \rightarrow \mathbb{R}$.

A. Finite Agent Model

Let \mathcal{X}, \mathcal{U} be a finite state and action space respectively. Let \mathcal{X}^0 be a finite environment state space. For any $N \in \mathbb{N}$, at each time $t = 0, 1, \dots$, the states and actions of agent $i = 1, \dots, N$ are random variables denoted by $x_t^{i,N} \in \mathcal{X}$ and $u_t^{i,N} \in \mathcal{U}$. Analogously, the environment state is a random variable denoted by $x_t^{0,N} \in \mathcal{X}^0$. Define the empirical state-action distribution $\mathbb{G}_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{(x_t^{i,N}, u_t^{i,N})} \in \mathcal{P}(\mathcal{X} \times \mathcal{U})$. For each agent i , we consider locally Markovian policies $\pi^i = \{\pi_t^i\}_{t \geq 0} \in \Pi_N$ from the space of admissible Markov policies Π_N where $\pi_t^i: \mathcal{X}^0 \times \mathcal{X} \rightarrow \mathcal{P}(\mathcal{U})$. Further, we define the policy profile $\pi = (\pi^1, \dots, \pi^N) \in \Pi_N^N$.

Acting only on local and environment information may seem like a strong restriction. However, other agent states are uninformative under continuity assumptions as $N \rightarrow \infty$ as the interaction between agents will be restricted to the increasingly deterministic empirical state-action distribution.

Let $\mu_0 \in \mathcal{P}(\mathcal{X})$ be the initial agent state distribution, $\mu_0^0 \in \mathcal{P}(\mathcal{X}^0)$ the initial environment state distribution and $P^0: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathcal{P}(\mathcal{X}^0)$ a transition kernel. The random variables shall follow $x_0^{0,N} \sim \mu_0^0$ and subsequently

$$x_t^{i,N} \sim \mu_0, \quad (1)$$

$$u_t^{i,N} \sim \pi_t^i(x_t^{0,N}, x_t^{i,N}), \quad (2)$$

$$x_{t+1}^{0,N} \sim P^0(x_t^{0,N}, \mathbb{G}_t^N), \quad (3)$$

where for simplicity of further analysis the agent states are always sampled according to μ_0 .

Remark. While this is a strong dynamics assumption, our formulation is nonetheless sufficient for the scheduling problem. In principle, any results should similarly hold under appropriate assumptions for nontrivial agent state dynamics by considering mean field and environment state together. As this will significantly complicate analysis, an according extension of theoretical results is left to future works.

Let us introduce another notation. First, define the space of decision rules $\mathcal{H} := \{h: \mathcal{X} \rightarrow \mathcal{P}(\mathcal{U})\}$. Then a one-step policy profile $\mathbf{h} = (h^1, \dots, h^N) \in \mathcal{H}^N$ is an N -fold decision rule. Our major example of a one-step policy profile is $(\pi_t^1(x^0), \dots, \pi_t^N(x^0))$ for fixed $t \geq 0$, fixed $x^0 \in \mathcal{X}^0$ and potentially different policies for the N agents. For given agent state distribution μ_0 and a one-step policy profile $\mathbf{h} \in \mathcal{H}^N$ let $x^{i,N} \sim \mu_0, u^{i,N} \sim h^j(x^{i,N})$, s.t. $(x^{i,N}, u^{i,N})_{i=1, \dots, N}$ are independent. Then, consider a random measure $\mathbb{G}_{\mathbf{h}}^N \in \mathcal{P}(\mathcal{X} \times \mathcal{U})$ or equivalently its random probability mass function $\mathbb{G}_{\mathbf{h}}^N: \mathcal{X} \times \mathcal{U} \rightarrow [0, 1]$,

$$\mathbb{G}_{\mathbf{h}}^N(x, u) := \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{x, u}(x^{i,N}, u^{i,N}). \quad (4)$$

Define $\mathcal{G}^N(\mu_0, \mathbf{h})$ as the distribution of $\mathbb{G}_{\mathbf{h}}^N$, so $\mathcal{G}^N(\mu_0, \mathbf{h})$ is a distribution over the set $\mathcal{P}(\mathcal{X} \times \mathcal{U})$ and $\mathbb{G}_{\mathbf{h}}^N \sim \mathcal{G}^N(\mu_0, \mathbf{h})$. Consider the primary example $\mathbf{h} = (\pi_t^1(x^0), \dots, \pi_t^N(x^0))$. In contrast to the empirical distribution $\mathbb{G}_{\mathbf{h}}^N$ that depends on a random $x_t^{0,N}$, the random probability mass function $\mathbb{G}_{\mathbf{h}}^N$ has $x_t^{0,N} = x^0$ fixed. By $\mathbb{E}[\mathbb{G}_{\mathbf{h}}^N]$ we denote in the following the entry-wise expectation $\{\mathbb{E}[\mathbb{G}_{\mathbf{h}}^N(x, u)]\}_{(x, u) \in \mathcal{X} \times \mathcal{U}}$.

Let $\gamma \in (0, 1)$ be the discount factor and $r: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ a reward function. The goal is to maximise the discounted accumulated reward

$$J^N(\boldsymbol{\pi}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t^{0,N}, \mathbb{G}_t^N) \right] \quad (5)$$

which generalizes optimizing an average per-agent reward

$$J^N(\boldsymbol{\pi}) = \sum_{i=1}^N \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t \tilde{r}(x_t^{i,N}, x_t^{0,N}, \mathbb{G}_t^N) \right] \quad (6)$$

for some shared $\tilde{r}: \mathcal{X} \times \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ through $r(x_t^{0,N}, \mathbb{G}_t^N) \equiv \sum_{x \in \mathcal{X}} \tilde{r}(x, x_t^{0,N}, \mathbb{G}_t^N) \sum_{u \in \mathcal{U}} \mathbb{G}_t^N(x, u)$.

As the optimality concept in this work, we therefore define approximate Pareto optimality.

Definition 1 (Pareto optimality). *For $\epsilon > 0$, $\boldsymbol{\pi}^\epsilon \in \Pi_N^N$ is ϵ -Pareto optimal if and only if*

$$J^N(\boldsymbol{\pi}^\epsilon) \geq \sup_{\boldsymbol{\pi}} J^N(\boldsymbol{\pi}) - \epsilon. \quad (7)$$

A visualization of this model can be found in Figure 1.

B. Mean Field Model

As $N \rightarrow \infty$, we formally obtain the following mean field MDP, which will be rigorously justified in the sequel. At each time $t = 0, 1, \dots$, the environment state is a random variable denoted by $x_t^0 \in \mathcal{X}^0$. We consider Markovian upper-level policies $\boldsymbol{\pi} = \{\pi_t\}_{t \geq 0} \in \Pi$ from the space of such policies Π where $\pi_t: \mathcal{X}^0 \rightarrow \mathcal{H}$. We equip both \mathcal{H} and Π with the supremum metric. As mentioned, the population state distribution is fixed to $\mu_0 \in \mathcal{P}(\mathcal{X})$ at all times. The random state-action distribution is therefore given by

$$\mathbb{G}_t := \mathbb{G}(\mu_0, \pi_t(x_t^0)) \quad (8)$$

where $\mathbb{G}: \mathcal{P}(\mathcal{X}) \times \mathcal{H} \rightarrow \mathcal{P}(\mathcal{X} \times \mathcal{U})$ is defined by

$$\mathbb{G}(\mu, h)(x, u) := h(x, u)\mu(x) \quad (9)$$

for any $x \in \mathcal{X}, u \in \mathcal{U}$. The random environment state variables therefore follow $x_0^0 \sim \mu_0^0$ and subsequently

$$x_{t+1}^0 \sim P^0(x_t^0, \mathbb{G}_t). \quad (10)$$

Analogously, the objective becomes

$$J(\boldsymbol{\pi}) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t^0, \mathbb{G}_t) \right]. \quad (11)$$

We require the following simple continuity assumption to obtain meaningful results in the limit as $N \rightarrow \infty$.

Assumption 1 (Continuity of r and P^0). *The functions r and P^0 are continuous, i.e. for all $x^0 \in \mathcal{X}^0$ and $\mathbb{G}_n \rightarrow \mathbb{G} \in \mathcal{P}(\mathcal{X} \times \mathcal{U})$ we have*

$$r(x^0, \mathbb{G}_n) \rightarrow r(x^0, \mathbb{G}), \quad P^0(x^0, \mathbb{G}_n) \rightarrow P^0(x^0, \mathbb{G}). \quad (12)$$

By compactness of $\mathcal{P}(\mathcal{X} \times \mathcal{U})$, we have boundedness.

Proposition 1. *Under Assumption 1, r is bounded by some R , i.e. for any $x^0 \in \mathcal{X}^0, \mathbb{G} \in \mathcal{P}(\mathcal{X} \times \mathcal{U})$ we have*

$$|r(x^0, \mathbb{G})| \leq R. \quad (13)$$

Our first goal will be to show that as $N \rightarrow \infty$, the optimal solution to the MFC is approximately Pareto optimal in the finite N case. This will motivate solving the MFC problem.

IV. APPROXIMATE OPTIMALITY

We first show the following lemma on uniform convergence in probability of empirical state-action distributions to their state-action-wise average for fixed one-step policy profiles.

Lemma 1. *Let $x^0 \in \mathcal{X}^0$ and $\mathbf{h} \in \mathcal{H}^N$ be an arbitrary one-step policy profile. Let $\mathbb{G}^N \sim \mathcal{G}(\mu_0, \mathbf{h})$. Then*

- (i) $\mathbb{E} [\|\mathbb{G}^N - \mathbb{E}[\mathbb{G}^N]\|_1^2] \leq \frac{|\mathcal{X}|^2 |\mathcal{U}|^2}{4N}$
- (ii) $\mathbb{P} (\|\mathbb{G}^N - \mathbb{E}[\mathbb{G}^N]\|_1 \geq \epsilon) \leq \frac{|\mathcal{X}|^2 |\mathcal{U}|^2}{4\epsilon^2 N}$

Proof. By Chebyshev's inequality, (i) implies (ii). It remains to prove (i). Let $x^{i,N} \sim \mu_0$ be i.i.d. and $u^{i,N} \sim \pi^i(x^0, x^{i,N})$, s.t. $(x^{i,N}, u^{i,N})_{i=1, \dots, N}$ are independent. Then by the sub-additivity of $\mathbb{E}[(\cdot)^2]^{\frac{1}{2}}$, we have

$$\begin{aligned} & \mathbb{E} [\|\mathbb{G}^N - \mathbb{E}[\mathbb{G}^N]\|_1^2]^{\frac{1}{2}} \\ &= \mathbb{E} \left[\left(\sum_{x \in \mathcal{X}, u \in \mathcal{U}} \left| \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{x, u}(x^{i,N}, u^{i,N}) - \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{1}_{x, u}(x^{i,N}, u^{i,N}) \right] \right| \right)^2 \right]^{\frac{1}{2}} \\ &\leq \sum_{x \in \mathcal{X}, u \in \mathcal{U}} \left(\mathbb{V} \left[\frac{1}{N} \sum_{i=1}^N \mathbf{1}_{x, u}(x^{i,N}, u^{i,N}) \right] \right)^{\frac{1}{2}} \\ &= \sum_{x \in \mathcal{X}, u \in \mathcal{U}} \left(\frac{1}{N^2} \sum_{i=1}^N \mathbb{V} [\mathbf{1}_{x, u}(x^{i,N}, u^{i,N})] \right)^{\frac{1}{2}} \\ &\leq \sum_{x \in \mathcal{X}, u \in \mathcal{U}} \left(\frac{1}{N^2} \sum_{i=1}^N \frac{1}{4} \right)^{\frac{1}{2}} = \frac{|\mathcal{X}| |\mathcal{U}|}{2\sqrt{N}} \end{aligned}$$

using the trivial variance bound $\frac{1}{4}$ for indicator functions. \square

To achieve approximate optimality of mean field solutions in the N -agent case, we first define how to obtain an N -agent policy $\pi^N \in \Pi_N^N$ from a mean field policy $\hat{\pi} \in \Pi$ by

$$\pi^N(\hat{\pi}) = (\pi^1, \dots, \pi^N) \text{ with } \pi_t^i(x^0, x) = \hat{\pi}_t(x^0)(x)$$

for all $i = 1, \dots, N$, i.e. all agents with state $x \in \mathcal{X}$ will follow the action distribution $\hat{\pi}_t(x^0)(x)$ at times $t \geq 0$.

Theorem 1. *Under Assumption 1, we have uniform convergence of the N -agent objective to the mean field objective as $N \rightarrow \infty$, i.e.*

$$\lim_{N \rightarrow \infty} \sup_{\pi \in \Pi} |J^N(\pi^N(\pi)) - J(\pi)| = 0. \quad (14)$$

Proof. We have by definition

$$\sup_{\pi \in \Pi} |J^N(\pi^N(\pi)) - J(\pi)| \quad (15)$$

$$= \sup_{\pi \in \Pi} \left| \sum_{t=0}^{\infty} \gamma^t \mathbb{E} \left[r(x_t^{0,N}, \mathbb{G}_t^N) - r(x_t^0, \mathbb{G}_t) \right] \right| \quad (16)$$

$$\leq \sum_{t=0}^{\infty} \gamma^t \sup_{\pi \in \Pi} \left| \mathbb{E} \left[r(x_t^{0,N}, \mathbb{G}_t^N) - r(x_t^0, \mathbb{G}_t) \right] \right|. \quad (17)$$

To obtain the desired result, we first show for any $t \geq 0$ that $\sup_{\pi \in \Pi} \|\mathcal{L}(x_t^{0,N}) - \mathcal{L}(x_t^0)\|_1 \rightarrow 0$ implies $\mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N) \rightarrow \mathcal{L}(x_t^0, \mathbb{G}_t)$ weakly uniformly over all $\pi \in \Pi$. Note that $\sup_{\pi \in \Pi} \|\mathcal{L}(x_t^{0,N}) - \mathcal{L}(x_t^0)\|_1 \rightarrow 0$ by definition implies

$$\sup_{\pi \in \Pi} \left| \mathcal{L}(x_t^{0,N})(x^0) - \mathcal{L}(x_t^0)(x^0) \right| \rightarrow 0$$

for any $x^0 \in \mathcal{X}^0$. For the joint law, consider any $f: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$, continuous and bounded by $|f| \leq F$. Then

$$\begin{aligned} & \sup_{\pi \in \Pi} \left| \mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N)(f) - \mathcal{L}(x_t^0, \mathbb{G}_t)(f) \right| \\ &= \sup_{\pi \in \Pi} \left| \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \right] - \mathbb{E} \left[f(x_t^0, \mathbb{G}_t) \right] \right| \\ &\leq \sup_{\pi \in \Pi} \sum_{x^0 \in \mathcal{X}^0} \left| \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \mathcal{L}(x_t^{0,N})(x) \right. \\ &\quad \left. - f(x^0, \mathbb{G}(\mu_0, \pi_t(x^0))) \mathcal{L}(x_t^0)(x^0) \right| \\ &\leq \sum_{x^0 \in \mathcal{X}^0} \sup_{\pi \in \Pi} \left| \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \right| \\ &\quad \cdot \sup_{\pi \in \Pi} \left| \mathcal{L}(x_t^{0,N})(x^0) - \mathcal{L}(x_t^0)(x^0) \right| \\ &+ \sum_{x^0 \in \mathcal{X}^0} \sup_{\pi \in \Pi} \left| f(x^0, \mathbb{G}(\mu_0, \pi_t(x^0))) \right. \\ &\quad \left. - \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \right| \\ &\quad \cdot \sup_{\pi \in \Pi} \mathcal{L}(x_t^0)(x^0), \end{aligned}$$

where the first sum goes to zero by assumption and boundedness of f . For the second term, consider arbitrary fixed $x^0 \in \mathcal{X}^0$. Write \mathbb{G}_π short for $\mathbb{G}(\mu_0, \pi_t(x^0))$ and introduce $\mathbb{G}_\pi^N \sim \mathcal{G}(\mu_0, (\pi_t(x^0), \dots, \pi_t(x^0)))$ for all $N, \pi \in \Pi$. So in

contrast to \mathbb{G}_t^N that depends on a random $x_t^{0,N}$, the random probability mass function \mathbb{G}_π^N has $x_t^{0,N} = x^0$ fixed. Then

$$\begin{aligned} & f(x^0, \mathbb{G}(\mu_0, \pi_t(x^0))) - \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \\ &= f(x^0, \mathbb{G}_\pi) - \mathbb{E} \left[f(x^0, \mathbb{G}_\pi^N) \right] \end{aligned}$$

We observe that for any $(x, u) \in \mathcal{X} \times \mathcal{U}$

$$\mathbb{E}[\mathbb{G}_\pi^N(x, u)] = \mathbb{G}_\pi(x, u). \quad (18)$$

For this purpose, let $x^{i,N} \sim \mu_0$ be i.i.d. and $u^{i,N} \sim \pi_t(x^0, x^{i,N})$, s.t. $(x^{i,N}, u^{i,N})_{i=1, \dots, N}$ are independent. Then for any $(x, u) \in \mathcal{X} \times \mathcal{U}$ we have

$$\begin{aligned} \mathbb{E}[\mathbb{G}_\pi^N(x, u)] &= \frac{1}{N} \sum_{i=1}^N \mathbb{E} [\mathbf{1}_{x,u}(x^{i,N}, u^{i,N})] \\ &= \frac{1}{N} \sum_{i=1}^N \mu_0(x) \pi_t(x^0, x, u) \\ &= \mathbb{G}_\pi(x, u). \end{aligned}$$

Let $\epsilon > 0$ arbitrary. By compactness of $\mathcal{P}(\mathcal{X} \times \mathcal{U})$, the function $f(x^0, \cdot): \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ is uniformly continuous. Consequently, there exists $\delta > 0$ such that for all $\pi \in \Pi$

$$\begin{aligned} \|\mathbb{G}_\pi - \mathbb{G}_\pi^N\|_1 &< \delta \\ \implies |f(x^0, \mathbb{G}_\pi) - f(x^0, \mathbb{G}_\pi^N)| &< \frac{\epsilon}{2}. \end{aligned}$$

By Lemma 1 (ii) and (18) there exists $N' \in \mathbb{N}$ such that for $N > N'$ and for all $\pi \in \Pi$ we have

$$\mathbb{P}(\|\mathbb{G}_\pi - \mathbb{G}_\pi^N\|_1 \geq \delta) \leq \frac{\epsilon}{4F}.$$

As a result, we have

$$\begin{aligned} & \mathbb{E} [|f(x^0, \mathbb{G}_\pi) - f(x^0, \mathbb{G}_\pi^N)|] \\ &\leq \mathbb{P}(|f(x^0, \mathbb{G}_\pi) - f(x^0, \mathbb{G}_\pi^N)| \geq \frac{\epsilon}{2}) \cdot 2F + 1 \cdot \frac{\epsilon}{2} \\ &\leq \mathbb{P}(\|\mathbb{G}_\pi - \mathbb{G}_\pi^N\|_1 \geq \delta) \cdot 2F + \frac{\epsilon}{2} \\ &\leq \frac{\epsilon}{4F} \cdot 2F + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

Since ϵ was arbitrary, and no choices depended on $\pi \in \Pi$, we have the desired convergence of the second term

$$\begin{aligned} & \lim_{N \rightarrow \infty} \sup_{\pi \in \Pi} \left| f(x, \mathbb{G}(\mu_0, \pi_t(x))) \right. \\ &\quad \left. - \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x \right] \right| = 0. \end{aligned}$$

We can now show $\mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N) \rightarrow \mathcal{L}(x_t^0, \mathbb{G}_t)$ weakly uniformly over all $\pi \in \Pi$ by induction over all t , which by Assumption 1 will imply

$$\sup_{\pi \in \Pi} \left| \mathbb{E} \left[r(x_t^{0,N}, \mathbb{G}_t^N) - r(x_t^0, \mathbb{G}_t) \right] \right| \rightarrow 0 \quad (19)$$

for all $t \geq 0$ and hence the desired statement by the dominated convergence theorem applied to (17).

At $t = 0$, we trivially have $\mathcal{L}(x_t^{0,N}) = \mu_0^0 = \mathcal{L}(x_t^0)$ and therefore $\mathcal{L}(x_0^{0,N}, \mathbb{G}_0^N) \rightarrow \mathcal{L}(x_0^0, \mathbb{G}_0)$ uniformly by the

prequel. Assume that the induction assumption holds at time t , then at time $t + 1$ we have

$$\begin{aligned} & \|\mathcal{L}(x_{t+1}^{0,N}) - \mathcal{L}(x_{t+1}^0)\|_1 \\ &= \sum_{x^0 \in \mathcal{X}^0} \left| \mathcal{L}(x_{t+1}^{0,N})(x^0) - \mathcal{L}(x_{t+1}^0)(x^0) \right| \\ &= \sum_{x^0 \in \mathcal{X}^0} \left| \mathbb{E} \left[P^0(x^0 \mid x_t^{0,N}, \mathbb{G}_t^N) \right] - \mathbb{E} \left[P^0(x^0 \mid x_t^0, \mathbb{G}_t) \right] \right| \\ &\rightarrow 0 \end{aligned}$$

uniformly by Assumption 1 and induction assumption. \square

To extend to optimality over arbitrary asymmetric policy tuples, we show that the performance of policy tuples is close to the averaged policy as $N \rightarrow \infty$.

Theorem 2. *Under Assumption 1, as $N \rightarrow \infty$ we have similar performance of any policy tuple $\pi = (\pi^1, \dots, \pi^N) \in \Pi_N^N$ and its average policy $\hat{\pi}(\pi) \in \Pi$ defined by $\hat{\pi}_t(x^0)(a \mid x) = \frac{1}{N} \sum_{i=1}^N \pi_t^i(a \mid x^0, x)$ in the N -agent case, i.e. with shorthand $\hat{\pi} = \hat{\pi}(\pi)$ we have*

$$\lim_{N \rightarrow \infty} \sup_{\pi \in \Pi_N^N} \left| J^N(\pi^1, \dots, \pi^N) - J^N(\pi^N(\hat{\pi})) \right| = 0. \quad (20)$$

Proof. Let $\pi \in \Pi_N^N$ arbitrary. Again, we have by definition

$$\begin{aligned} & \sup_{\pi \in \Pi_N^N} \left| J^N(\pi^1, \dots, \pi^N) - J^N(\pi^N(\hat{\pi})) \right| \quad (21) \\ & \leq \sum_{t=0}^{\infty} \gamma^t \sup_{\pi \in \Pi_N^N} \left| \mathbb{E} \left[r(x_t^{0,N}, \mathbb{G}_t^N) - r(x_t^{0,N}, \hat{\mathbb{G}}_t^N) \right] \right| \quad (22) \end{aligned}$$

by introducing random variables $\hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N, \hat{x}_t^{i,N}, \hat{u}_t^{i,N}, i = 1, \dots, N$ induced by instead applying the averaged policy tuple $\pi^N(\hat{\pi})$ in (2). By dominated convergence, it is sufficient to show term-wise convergence to zero in (22).

Fix $t \geq 0$. As in the proof of Theorem 1, we show that $\sup_{\pi \in \Pi} \|\mathcal{L}(x_t^{0,N}) - \mathcal{L}(\hat{x}_t^{0,N})\|_1 \rightarrow 0$ implies $\sup_{\pi \in \Pi_N^N} \left| \mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N)(f) - \mathcal{L}(\hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N)(f) \right| \rightarrow 0$ for any $f: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ continuous and bounded, since

$$\begin{aligned} & \sup_{\pi \in \Pi_N^N} \left| \mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N)(f) - \mathcal{L}(\hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N)(f) \right| \\ &= \sup_{\pi \in \Pi_N^N} \left| \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \right] - \mathbb{E} \left[f(\hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N) \right] \right| \\ &\leq \sum_{x^0 \in \mathcal{X}^0} \sup_{\pi \in \Pi_N^N} \left| \mathbb{E} \left[f(x_t^{0,N}, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \right. \\ &\quad \cdot \sup_{\pi \in \Pi_N^N} \left| \mathcal{L}(x_t^{0,N})(x^0) - \mathcal{L}(\hat{x}_t^{0,N})(x^0) \right| \\ &\quad + \sum_{x^0 \in \mathcal{X}^0} \sup_{\pi \in \Pi_N^N} \left| \mathbb{E} \left[f(x^0, \hat{\mathbb{G}}_t^N) \mid \hat{x}_t^{0,N} = x^0 \right] \right. \\ &\quad \left. - \mathbb{E} \left[f(x^0, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \right| \\ &\quad \cdot \sup_{\pi \in \Pi_N^N} \mathcal{L}(\hat{x}_t^{0,N})(x^0) \end{aligned}$$

where the first sum goes to zero by assumption and boundedness of f . For the second term, consider arbitrary

fixed $x^0 \in \mathcal{X}^0, \pi \in \Pi_N^N$. Then introduce random variables $\mathbb{G}_\pi^N \sim \mathcal{G}^N(\mu_0, (\pi_t^1(x^0), \dots, \pi_t^N(x^0)))$ and $\mathbb{G}_{\hat{\pi}}^N \sim \mathcal{G}^N(\mu_0, (\hat{\pi}_t(x^0), \dots, \hat{\pi}_t(x^0)))$ for every $N \in \mathbb{N}$ and $\pi \in \Pi_N^N$. Then we have

$$\begin{aligned} & \mathbb{E} \left[f(x^0, \hat{\mathbb{G}}_t^N) \mid \hat{x}_t^{0,N} = x^0 \right] - \mathbb{E} \left[f(x^0, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \\ &= \mathbb{E} \left[f(x^0, \mathbb{G}_{\hat{\pi}}^N) \right] - \mathbb{E} \left[f(x^0, \mathbb{G}_\pi^N) \right]. \end{aligned}$$

We observe that for any $(x, u) \in \mathcal{X} \times \mathcal{U}$:

$$\mathbb{E}[\mathbb{G}_{\hat{\pi}}^N(x, u)] = \mathbb{E}[\mathbb{G}_\pi^N(x, u)]. \quad (23)$$

For this purpose, let $x^{i,N} \sim \mu_0$ and $u^{i,N} \sim \pi_t^i(x^0, x^{i,N})$ as well as $\hat{x}^{i,N} \sim \mu_0$ and $\hat{u}^{i,N} \sim \hat{\pi}_t(x^0, x^{i,N})$, s.t. $(x^{i,N}, u^{i,N})_{i=1, \dots, N}$ and $(\hat{x}^{i,N}, \hat{u}^{i,N})_{i=1, \dots, N}$ are independent, respectively. Then for any $(x, u) \in \mathcal{X} \times \mathcal{U}$:

$$\begin{aligned} \mathbb{E}[\mathbb{G}_{\hat{\pi}}^N(x, u)] &= \frac{1}{N} \sum_{i=1}^N \mathbb{P}(\hat{x}^{i,N} = x, \hat{u}^{i,N} = u) \\ &= \frac{1}{N} \sum_{i=1}^N \mu_0(x) \hat{\pi}(x, u) = \frac{1}{N} \sum_{i=1}^N \mu_0(x) \frac{1}{N} \sum_{j=1}^N \pi^j(x, u) \\ &= \frac{1}{N} \sum_{j=1}^N \mu_0(x) \pi^j(x, u) = \frac{1}{N} \sum_{j=1}^N \mathbb{P}(x^{i,N} = x, u^{i,N} = u) \\ &= \mathbb{E}[\mathbb{G}_\pi^N(x, u)]. \end{aligned}$$

Then by (23), sub-additivity of $\mathbb{E}[(\cdot)^2]^{\frac{1}{2}}$ and Lemma 1 (i),

$$\begin{aligned} & \mathbb{E}[\|\mathbb{G}_{\hat{\pi}}^N - \mathbb{G}_\pi^N\|_1^2]^{\frac{1}{2}} \\ & \leq \mathbb{E} \left[\left(\|\mathbb{G}_{\hat{\pi}}^N - \mathbb{E}[\mathbb{G}_{\hat{\pi}}^N]\|_1 + \|\mathbb{G}_\pi^N - \mathbb{E}[\mathbb{G}_\pi^N]\|_1 \right)^2 \right]^{\frac{1}{2}} \\ & \leq \mathbb{E}[\|\mathbb{G}_{\hat{\pi}}^N - \mathbb{E}[\mathbb{G}_{\hat{\pi}}^N]\|_1^2]^{\frac{1}{2}} + \mathbb{E}[\|\mathbb{G}_\pi^N - \mathbb{E}[\mathbb{G}_\pi^N]\|_1^2]^{\frac{1}{2}} \\ & \leq 2 \frac{|\mathcal{X}| \cdot |\mathcal{U}|}{\sqrt{4N}} = \frac{|\mathcal{X}| \cdot |\mathcal{U}|}{\sqrt{N}}. \end{aligned}$$

Chebyshev's inequality implies

$$\mathbb{P}(\|\mathbb{G}_{\hat{\pi}}^N - \mathbb{G}_\pi^N\|_1 \geq \epsilon) \leq \frac{|\mathcal{X}|^2 |\mathcal{U}|^2}{\epsilon^2 N} \quad (24)$$

independent of $\pi \in \Pi_N^N$.

Then analogously to the proof of Theorem 1,

$$\begin{aligned} & \sup_{\pi \in \Pi_N^N} \left| \mathbb{E} \left[f(x^0, \hat{\mathbb{G}}_t^N) \mid \hat{x}_t^{0,N} = x^0 \right] \right. \\ & \quad \left. - \mathbb{E} \left[f(x^0, \mathbb{G}_t^N) \mid x_t^{0,N} = x^0 \right] \right| \rightarrow 0 \end{aligned}$$

can be concluded, showing the desired implication.

We now show by induction over all t that for any $t \geq 0$, and any $f: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ continuous and bounded, $\sup_{\pi \in \Pi_N^N} \left| \mathcal{L}(x_t^{0,N}, \mathbb{G}_t^N)(f) - \mathcal{L}(\hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N)(f) \right| \rightarrow 0$ which by Assumption 1 will again imply that (22) goes to zero.

At $t = 0$, we have $\mathcal{L}(x_0^{0,N}) = \mu_0^0 = \mathcal{L}(\hat{x}_0^{0,N})$, implying $\sup_{\pi \in \Pi_N^N} \left| \mathcal{L}(x_0^{0,N}, \mathbb{G}_0^N)(f) - \mathcal{L}(\hat{x}_0^{0,N}, \hat{\mathbb{G}}_0^N)(f) \right| \rightarrow 0$ for any $f: \mathcal{X}^0 \times \mathcal{P}(\mathcal{X} \times \mathcal{U}) \rightarrow \mathbb{R}$ by the prequel. Assuming the induction assumption holds at time t , then at time $t + 1$

$$\|\mathcal{L}(x_{t+1}^{0,N}) - \mathcal{L}(\hat{x}_{t+1}^{0,N})\|_1$$

$$= \sum_{x^0 \in \mathcal{X}^0} \left| \mathbb{E} \left[P^0(x^0 | x_t^{0,N}, \mathbb{G}_t^N) - P^0(x^0 | \hat{x}_t^{0,N}, \hat{\mathbb{G}}_t^N) \right] \right| \\ \rightarrow 0$$

uniformly by induction assumption and continuity and boundedness of P^0 , which implies the desired statement. \square

Corollary 1. *Under Assumption 1, for any $\epsilon > 0$ there exists $N(\epsilon)$ such that for all $N > N(\epsilon)$ a policy π^* optimal in the MFC MDP – that is, $J(\pi^*) = \sup_{\pi \in \Pi} J(\pi)$ – is ϵ -Pareto optimal in the N -agent case, i.e.*

$$J^N(\pi^N(\pi^*)) \geq \sup_{\pi} J^N(\pi) - \epsilon. \quad (25)$$

Proof. By Theorem 1 and Theorem 2, there exists $N' \in \mathbb{N}$ such that for average policy $\hat{\pi}$ of π and all $N > N'$ we have

$$\begin{aligned} & \sup_{\pi} (J^N(\pi) - J^N(\pi^N(\pi^*))) \\ & \leq \sup_{\pi} (J^N(\pi) - J^N(\pi^N(\hat{\pi}))) \\ & \quad + \sup_{\pi} (J^N(\pi^N(\hat{\pi})) - J(\hat{\pi})) \\ & \quad + \sup_{\pi} (J(\hat{\pi}) - J(\pi^*)) \\ & \quad + \sup_{\pi} (J(\pi^*) - J^N(\pi^N(\pi^*))) \\ & < \frac{\epsilon}{3} + \frac{\epsilon}{3} + 0 + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Reordering terms gives the desired inequality. \square

V. DYNAMIC PROGRAMMING PRINCIPLE

The following dynamic programming principle for the MFC MDP is a standard result, for which the MDP state will be only the environment state, see e.g. [20], [22].

Define action-value function $Q: \mathcal{X}^0 \times \mathcal{H} \rightarrow \mathbb{R}$,

$$Q(x^0, h) := \sup_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t^0, \mathbb{G}(\mu_0, \pi_t(x_t^0))) \mid x_0^0 = x^0, \pi_0(x^0) = h \right]. \quad (26)$$

Note that by boundedness of r , we trivially have

$$|Q| \leq \frac{R}{1-\gamma}.$$

As we have an MDP with finite state space \mathcal{X}^0 , the following Bellman equation will hold, see [28].

Theorem 3. *The Bellman equation*

$$Q(x^0, h) = r(x^0, \mathbb{G}(\mu_0, h)) + \gamma \mathbb{E}_{\tilde{x}^0 \sim P^0(x^0, \mathbb{G}(\mu_0, h))} \left[\sup_{\tilde{h} \in \mathcal{H}} Q(\tilde{x}^0, \tilde{h}) \right] \quad (27)$$

holds for all $x^0 \in \mathcal{X}^0, h \in \mathcal{H}$.

In the following, we obtain existence of an optimal stationary policy by compactness of \mathcal{H} and continuity of Q , which shall be inherited from the continuity of r and P^0 .

Lemma 2. *The unique function that satisfies the Bellman equation is given by Q . Further, if there exists $h_{x^0} \in$*

arg max $_{h \in \mathcal{H}} Q(x^0, h)$ for any $x^0 \in \mathcal{X}^0$, then the policy π^ with $\pi_t^*(x^0) = h_{x^0}$ is an optimal stationary policy.*

Proof. For uniqueness, define the space of $\frac{R}{1-\gamma}$ -bounded functions $\mathcal{Q} := \{f: \mathcal{X}^0 \times \mathcal{H} \rightarrow [-\frac{R}{1-\gamma}, \frac{R}{1-\gamma}]\}$ and the Bellman operator $B: \mathcal{Q} \rightarrow \mathcal{Q}$ defined by

$$(BQ)(x^0, h) = r(x^0, \mathbb{G}(\mu_0, h)) + \gamma \mathbb{E}_{\tilde{x}^0 \sim P^0(x^0, \mathbb{G}(\mu_0, h))} \left[\sup_{\tilde{h} \in \mathcal{H}} Q(\tilde{x}^0, \tilde{h}) \right]. \quad (28)$$

We show that \mathcal{Q} is a complete metric space under the supremum norm. Let $(Q_n)_{n \in \mathbb{N}}$ be a Cauchy sequence of functions $Q_n \in \mathcal{Q}$. Then by definition, for any $\epsilon > 0$ there exists $n' \in \mathbb{N}$ such that for all $n, m > n'$ we have

$$\begin{aligned} & \|Q_n - Q_m\|_{\infty} < \epsilon \\ \implies & \forall x^0 \in \mathcal{X}^0, h \in \mathcal{H}: |Q_n(x^0, h) - Q_m(x^0, h)| < \epsilon \end{aligned}$$

such that for all $x^0 \in \mathcal{X}^0, h \in \mathcal{H}$ there exists a value $c_{x^0, h} \in [-\frac{R}{1-\gamma}, \frac{R}{1-\gamma}]$ for which $Q_n(x^0, h) \rightarrow c_{x^0, h}$. Define the function $Q' \in \mathcal{Q}$ by $Q'(x^0, h) = c_{x^0, h}$, then we have

$$\begin{aligned} & |Q_n(x^0, h) - Q'(x^0, h)| \\ & = \lim_{m \rightarrow \infty} |Q_n(x^0, h) - Q_m(x^0, h)| < \epsilon \end{aligned}$$

for all $x^0 \in \mathcal{X}^0, h \in \mathcal{H}, n > n'$, and hence $Q_n \rightarrow Q' \in \mathcal{Q}$ as $n \rightarrow \infty$. This implies completeness of $(\mathcal{Q}, \|\cdot\|_{\infty})$.

We now show that B is a contraction under the supremum norm, i.e.

$$\|BQ_1 - BQ_2\|_{\infty} \leq C \|Q_1 - Q_2\|_{\infty}$$

for some $C < 1$. Define the shorthand $\tilde{x}^0 \sim P^0(x^0, \mathbb{G}(\mu_0, h))$. We have

$$\begin{aligned} & \|BQ_1 - BQ_2\|_{\infty} \\ & = \sup_{x^0 \in \mathcal{X}^0, h \in \mathcal{H}} |BQ_1(x^0, h) - BQ_2(x^0, h)| \\ & \leq \sup_{x^0 \in \mathcal{X}^0, h \in \mathcal{H}} \gamma \mathbb{E}_{\tilde{x}^0} \left[\left| \sup_{\tilde{h} \in \mathcal{H}} Q_1(\tilde{x}^0, \tilde{h}) - \sup_{\tilde{h} \in \mathcal{H}} Q_2(\tilde{x}^0, \tilde{h}) \right| \right] \\ & \leq \sup_{x^0 \in \mathcal{X}^0, h \in \mathcal{H}} \gamma \|Q_1 - Q_2\|_{\infty} \end{aligned}$$

with $\gamma < 1$. Therefore, by Banach fixed point theorem, B has the unique fixed point Q .

For optimality, define the policy action-value function Q^{π} for $\pi \in \Pi$ as the fixed point of $B^{\pi}: \mathcal{Q} \rightarrow \mathcal{Q}$ defined by

$$(B^{\pi}Q)(x^0, h) = r(x^0, \mathbb{G}(\mu_0, h)) + \gamma \mathbb{E}_{\tilde{x}^0} [Q(\tilde{x}^0, \pi(\tilde{x}^0))].$$

From this, we immediately have

$$\begin{aligned} Q^{\pi^*}(x^0, h) & = r(x^0, \mathbb{G}(\mu_0, h)) + \gamma \mathbb{E}_{\tilde{x}^0} [Q(\tilde{x}^0, \pi^*(\tilde{x}^0))] \\ & = r(x^0, \mathbb{G}(\mu_0, h)) + \gamma \mathbb{E}_{\tilde{x}^0} \left[\sup_{\tilde{h} \in \mathcal{H}} Q(\tilde{x}^0, \tilde{h}) \right] \\ & = Q(x^0, h) \end{aligned}$$

which implies that π^* is optimal, see also [28]. \square

Lemma 3. *The action-value function Q is continuous.*

Proof. We will show as $x_n^0 \rightarrow x^0 \in \mathcal{X}^0$ and $h_n \rightarrow h \in \mathcal{H}$,

$$Q(x_n^0, h_n) \rightarrow Q(x^0, h).$$

By the Bellman equation, we immediately have

$$\begin{aligned} & |Q(x_n^0, h_n) - Q(x^0, h)| \\ & \leq |r(x_n^0, \mathbb{G}(\mu_0, h_n)) - r(x^0, \mathbb{G}(\mu_0, h))| \\ & \quad + \left| \gamma \sum_{\tilde{x}^0 \in \mathcal{X}^0} (P^0(\tilde{x}^0 | x_n^0, \mathbb{G}(\mu_0, h_n)) \right. \\ & \quad \quad \left. - P^0(\tilde{x}^0 | x^0, \mathbb{G}(\mu_0, h))) \sup_{\tilde{h} \in \mathcal{H}} Q(\tilde{x}^0, \tilde{h}) \right| \\ & \leq |r(x_n^0, \mathbb{G}(\mu_0, h_n)) - r(x^0, \mathbb{G}(\mu_0, h))| \\ & \quad + \frac{\gamma R}{1 - \gamma} \sum_{\tilde{x}^0 \in \mathcal{X}^0} |P^0(\tilde{x}^0 | x_n^0, \mathbb{G}(\mu_0, h_n)) \\ & \quad \quad - P^0(\tilde{x}^0 | x^0, \mathbb{G}(\mu_0, h))| \rightarrow 0 \end{aligned}$$

since r, P^0, \mathbb{G} are continuous and Q is bounded. \square

Corollary 2. *There exists an optimal stationary policy $\pi^*: \mathcal{X}^0 \rightarrow \mathcal{H}$ such that $Q^{\pi^*} = Q$.*

Proof. By Lemma 3, Q is continuous. Furthermore, \mathcal{H} is compact. By the extreme value theorem, there exists $h_{x^0} \in \arg \max_{h \in \mathcal{H}} Q(x^0, h)$ for any $x^0 \in \mathcal{X}^0$. By Lemma 2, there exists an optimal stationary policy π^* . \square

Since \mathcal{H} is continuous, general exact solutions are difficult. Instead, we apply reinforcement learning with stochastic policies to find an optimal stationary policy.

VI. EXPERIMENTS

We compare the empirical performance of the mean field solution in the aforementioned scheduling problem. Since there exist few theoretical guarantees for tractable multi-agent reinforcement learning methods [5], we compare our approach (MF) to empirically effective independent learning (IL) [29], i.e. applying single-agent RL to each separate agent (NA), as well as the well-known Join-Shortest-Queue (JSQ) algorithm [26], where agents choose the shortest queue accessible and otherwise randomly. To make independent learning more tractable, we also share policy parameters between all agents using parameter sharing (PS) [30] and train each policy via the PPO algorithm [31] using the RLlib 1.2.0 Pytorch implementation [32] for 400,000 time steps in the N -agent cases and 2 million time steps in the MF case, which is sufficient for convergence of MF and N -agent policies up to $N = 4$, after which N -agent training becomes unstable under the shared hyperparameters in Table I and continues to fail even with more time steps.

For policies and critics, we use separate feedforward networks with two hidden layers of 256 nodes and tanh activations. In the mean field case the policy outputs parameters μ, σ of a diagonal Gaussian distribution over actions, which are sampled and clipped between 0 and 1. We normalize each of these output values such that they give the probability of assigning to an accessible queue given some agent state, i.e.

TABLE I
PARAMETER AND HYPERPARAMETER SETTINGS USED IN THE EXPERIMENTS OF THIS WORK.

Symbol	Function	Value
c_d	Packet drop penalty	1
M	Number of queues	2
B_i	Queue buffer sizes	5
ΔT	Time step size	0.5 s
λ	Packet arrival rate	$(3M - 1) \text{ s}^{-1}$
β	Queue servicing rate	3 s^{-1}
γ	Discount factor	0.99
PPO		
l_r	Learning rate	$5e - 5$
λ_{PPO}	GAE coefficient	0.2
β_{PPO}	Initial KL coefficient	0.2
d_{target}	KL target	0.01
ϵ	Clip parameter	0.3
B	Training batch size	4000
B_m	SGD mini-batch size	128
k	SGD iterations per batch	30

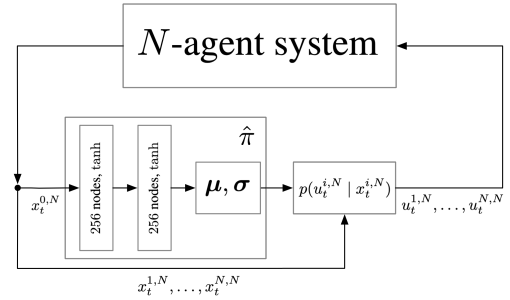


Fig. 3. Overview of mean field control application in N -agent systems: Conditional on the environment state $x_t^{0,N}$, the upper-level mean field policy $\hat{\pi}$ outputs a sampled, shared lower-level policy for all agents i , from which random actions $u_t^{i,N}$ are sampled conditional on local agent states $x_t^{i,N}$.

a shared lower-level decision rule $h \in \mathcal{H}$ for all agents. A visualization of this process can be found in Figure 3.

Note that we use stochastic policies as required by stochastic policy gradient methods, though we can easily obtain a deterministic policy if necessary by simply using the mean parameter of the Gaussian distribution. In the N -agent case, we output queue assignment probabilities for each of the agents via a standard softmax final layer. Invalid assignments to queues that are not accessible by an agent are treated as randomly sampling one from all accessible queues.

As can be seen in Figure 4 for μ_0 given such that the probability of access to both queues is 0.6 and otherwise uniformly random, the mean field solution reaches its mean field performance in the N -agent case as N grows large. This validates our theoretical findings empirically. Our solution further appears to outperform NA and PS for sufficiently many agents, as IL approaches increasingly fail to learn due to the credit assignment problem.

Moreover, our best learned policy is close to JSQ and competitive with slight irregularities at $b_0 = 0$. Observe in Figure 4 that the MF policy gives an interpretable solution. As a queue becomes more filled, the optimal solution will be more likely to avoid assignment of packets to that queue.




Fig. 4. (a): Cumulative reward average over 500 runs with 95% confidence interval achieved against number of agents N . The dotted line indicates cumulative reward of MF in the MFC MDP. NA and PS are trained separately for each N , while the MF policy is trained only once and used for all N . As N grows, the MF policy performance becomes increasingly close to the MFC MDP and competitive with JSQ, while NA and PS begin to fail learning due to the credit assignment problem. (b): MF policy probabilities of assigning to queue 1 against buffer fillings b_0, b_1 for agents with access to both queues, averaged over 500 samples.

VII. CONCLUSION

In this work, we have formulated a discrete-time mean field control model with common environment states motivated by a scheduling problem. We have rigorously shown approximate optimality as $N \rightarrow \infty$ and applied reinforcement learning to solve the MFC MDP. Empirically, we obtain competitive results for sufficiently many agents and validate our theoretical results. For future work, it could be interesting to consider partial observability of the system for schedulers, or methods to scale to large numbers of queues. Potential extensions are manifold and include dynamic agent states, major-minor systems, partial observability and general non-finite spaces.

ACKNOWLEDGMENT

This work has been co-funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY center, the European Research Council (ERC) within the Consolidator Grant CONSYN (grant agreement no. 773196) and the German Research Foundation (DFG) as part of sub-project C3 within the Collaborative Research Center (CRC) 1053 – MAKI.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [3] N. Brown and T. Sandholm, “Superhuman ai for multiplayer poker,” *Science*, vol. 365, no. 6456, pp. 885–890, 2019.
- [4] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y.-C. Liang, and D. I. Kim, “Applications of deep reinforcement learning in communications and networking: A survey,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3133–3174, 2019.
- [5] K. Zhang, Z. Yang, and T. Başar, “Multi-agent reinforcement learning: A selective overview of theories and algorithms,” *Handbook of Reinforcement Learning and Control*, pp. 321–384, 2021.
- [6] M. Huang, R. P. Malhamé, P. E. Caines *et al.*, “Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle,” *Communications in Information & Systems*, vol. 6, no. 3, pp. 221–252, 2006.
- [7] J.-M. Lasry and P.-L. Lions, “Mean field games,” *Japanese journal of mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [8] D. A. Gomes, J. Mohr, and R. R. Souza, “Discrete time, finite state space mean field games,” *Journal de mathématiques pures et appliquées*, vol. 93, no. 3, pp. 308–328, 2010.

- [9] N. Saldi, T. Basar, and M. Raginsky, “Markov–nash equilibria in mean-field games with discounted cost,” *SIAM Journal on Control and Optimization*, vol. 56, no. 6, pp. 4256–4287, 2018.
- [10] D. Mguni, J. Jennings, and E. M. de Cote, “Decentralised learning in systems with many, many strategic agents,” *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [11] X. Guo, A. Hu, R. Xu, and J. Zhang, “Learning mean-field games,” in *Advances in Neural Information Processing Systems*, 2019, pp. 4966–4976.
- [12] K. Cui and H. Koepl, “Approximately solving mean field games via entropy-regularized deep reinforcement learning,” in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 1909–1917.
- [13] D. Andersson and B. Djehiche, “A maximum principle for sdes of mean-field type,” *Applied Mathematics & Optimization*, vol. 63, no. 3, pp. 341–356, 2011.
- [14] A. Bensoussan, J. Frehse, P. Yam *et al.*, *Mean field games and mean field type control theory*. Springer, 2013, vol. 101.
- [15] A. C. Kizilkale and R. P. Malhamé, “Collective target tracking mean field control for electric space heaters,” in *22nd Mediterranean Conference on Control and Automation*. IEEE, 2014, pp. 829–834.
- [16] B. Djehiche and H. Tembine, “Risk-sensitive mean-field type control under partial observation,” in *Stochastics of Environmental and Financial Economics*. Springer, Cham, 2016, pp. 243–263.
- [17] B. Djehiche, H. Tembine, and R. Tempone, “A stochastic maximum principle for risk-sensitive mean-field type control,” *IEEE Transactions on Automatic Control*, vol. 60, no. 10, pp. 2640–2649, 2015.
- [18] M. F. Djete, D. Possamaï, and X. Tan, “Mckean-vlasov optimal control: the dynamic programming principle,” *arXiv preprint arXiv:1907.08860*, 2019.
- [19] H. Pham and X. Wei, “Bellman equation and viscosity solutions for mean-field stochastic control problem,” *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 24, no. 1, pp. 437–461, 2018.
- [20] M. Motte and H. Pham, “Mean-field markov decision processes with common noise and open-loop controls,” *arXiv preprint arXiv:1912.07883*, 2019.
- [21] H. Gu, X. Guo, X. Wei, and R. Xu, “Dynamic programming principles for learning mfcs,” *arXiv preprint arXiv:1911.07314*, 2019.
- [22] —, “Mean-field controls with q-learning for cooperative marl: Convergence and complexity analysis,” *arXiv preprint arXiv:2002.04131*, 2020.
- [23] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [24] M. Nourian and P. E. Caines, “ ϵ -nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents,” *SIAM Journal on Control and Optimization*, vol. 51, no. 4, pp. 3302–3331, 2013.
- [25] P. E. Caines and A. C. Kizilkale, “ ϵ -nash equilibria for partially observed lqg mean field games with a major player,” *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3225–3234, 2016.
- [26] A. Mukhopadhyay and R. R. Mazumdar, “Analysis of randomized join-the-shortest-queue (jsq) schemes in large heterogeneous processor-sharing systems,” *IEEE Transactions on Control of Network Systems*, vol. 3, no. 2, pp. 116–126, 2016.
- [27] W. R. KhudaBukhsh, S. Kar, B. Alt, A. Rizk, and H. Koepl, “Generalized cost-based job scheduling in very large heterogeneous cluster systems,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 11, pp. 2594–2604, 2020.
- [28] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [29] M. Tan, “Multi-agent reinforcement learning: Independent vs. cooperative agents,” in *Proceedings of the tenth international conference on machine learning*, 1993, pp. 330–337.
- [30] J. K. Gupta, M. Egorov, and M. Kochenderfer, “Cooperative multi-agent control using deep reinforcement learning,” in *International Conference on Autonomous Agents and Multiagent Systems*. Springer, 2017, pp. 66–83.
- [31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [32] E. Liang, R. Liaw, R. Nishihara, P. Moritz, R. Fox, K. Goldberg, J. Gonzalez, M. Jordan, and I. Stoica, “Rllib: Abstractions for distributed reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 3053–3062.