DistSPECTRL: Distributing Specifications in Multi-Agent Reinforcement Learning Systems

Joe Eappen (🖂) and Suresh Jagannathan

Purdue University, West Lafayette IN, 47907, USA jeappen@purdue.edu suresh@cs.purdue.edu

Abstract. While notable progress has been made in specifying and learning objectives for general cyber-physical systems, applying these methods to distributed multi-agent systems still pose significant challenges. Among these are the need to (a) craft specification primitives that allow expression and interplay of both local and global objectives, (b) tame explosion in the state and action spaces to enable effective learning, and (c) minimize coordination frequency and the set of engaged participants for global objectives. To address these challenges, we propose a novel specification framework that allows natural composition of local and global objectives used to guide training of a multi-agent system. Our technique enables learning expressive policies that allow agents to operate in a coordination-free manner for local objectives, while using a decentralized communication protocol for enforcing global ones. Experimental results support our claim that sophisticated multi-agent distributed planning problems can be effectively realized using specification-guided learning. Code is provided at https://github.com/yokian/distspectrl.

Keywords: Multi-Agent Reinforcement Learning, Specification-Guided Learning

1 Introduction

Reinforcement Learning (RL) can be used to learn complex behaviors in many different problem settings. A main component of RL is providing feedback to an agent via a reward signal. This signal should encourage desired behaviors, and penalize undesirable ones, enabling the agent to eventually proceed through a sequence of tasks and is designed by the programmer beforehand. A commonly used technique to encode tasks in a reward signal is the sparse method of providing zero reward until a task is completed upon which a non-zero reward is given to the agent. Because this procedure has the significant shortcoming of delaying generating a useful feedback signal for a large portion of the agent-environment interaction process, a number of alternative techniques have been proposed [14,1].

Formulating a reward signal that reduces the sparsity of this feedback is known as *reward shaping*. Often, this is done manually but a more general, robust method would be to automatically shape a reward given a specification of desired behavior. SPECTRL [5] proposes a reward shaping mechanism for a set of temporal logic specifications on

a single-agent task that uses a compiled a finite-state automaton called a task monitor. Reward machines [19,2,18] are another objective-specifying method for RL problems that also define a finite automaton akin to the ones used in SPECTRL, with some subtle differences such as the lack of registers (used by the task monitor for memory).

Distributed multi-agent applications however, introduce new challenges in automating this reward shaping process. Agents have their own respective goals to fulfill as well as coordinated goals that must be performed in cooperation with other agents. While the expressiveness of the language in SPECTRL lends itself, with minor extensions, to specifying these kinds of goals, we require new compilation and execution algorithms to tackle inherent difficulties in multi-agent reinforcement learning (MARL); these include credit assignment of global objectives and the presence of large state and action spaces that grow as the number of agents increases. Learning algorithms for multi-agent problems have often encouraged distribution as a means of scaling in the presence of state and action space explosion. This is because purely centralized approaches have the disadvantages of not only requiring global knowledge of the system at all times but also induce frequent and costly synchronized agent control.

To address these issues, we develop a new specification-guided distributed multiagent reinforcement learning framework. Our approach has four main features. First, we introduce two classes of predicates (*viz.* local and global) to capture tasks in a multiagent world (Sec. 4). Second, we develop a new procedure for generating composite task monitors using these predicates and devise new techniques to distribute these monitors over all agents to address scalability and decentralization concerns (Sec. 5). Third, we efficiently solve the introduced problem of subtask synchronization (Sec. 6) among agents via synchronization states in the task monitors. Lastly, we describe a wide class of specification structures (Sec. 7) amenable to scaling in the number of agents and provide a means to perform such a scaling (Sec. 8).

By using these components in tandem, we provide the first solution to composing specifications and distributing them among agents in a scalable fashion within a multi-agent learning scenario supporting continuous state and action spaces. Before presenting details of our approach, we first provide necessary background information (Sec. 2) and formalize the problem (Sec. 3).

2 Background

Markov Decision Processes Reinforcement learning is a tool to solve Markov Decision Processes (MDPs). MDPs are tuples of the form $\langle S, D, A, P, R, T \rangle$ where $S \in \mathbb{R}^n$ is the state space, D is the initial state distribution, $A \in \mathbb{R}^m$ is the action space, $P: S \times A \times S \rightarrow [0,1]$ is the transition function, and T is the time horizon. A rollout $\zeta \in Z$ of length T is a sequence of states and actions $\zeta = (s_0, a_0, ..., a_{T-1}, s_T)$ where $s_i \in S$ and $a_i \in A$ are such that $s_{i+1} \sim P(s_i, a_i)$. $R: Z \rightarrow \mathbb{R}$ is a reward function used to score a rollout ζ .

Multi-Agent Reinforcement Learning A Markov game with $\mathcal{N} = \{1, \dots, N\}$ denoting the set of N agents is a tuple $\mathcal{M}_g = \langle \mathcal{N}, \{S^i\}_{i \in \mathcal{N}}, D, \{A^i\}_{i \in \mathcal{N}}, P, \{R^i\}_{i \in \mathcal{N}}, T \rangle$ where A^i, R^i define their agent-specific action spaces and reward functions. They are

3

a direct generalization of MDPs to the multi-agent scenario. Let $S_m = \{S^i\}_{i \in \mathcal{N}}$ and $A_m = \{A^i\}_{i \in \mathcal{N}}$, then $P : S_m \times A_m \times S_m \to [0, 1]$ is the transition function. A rollout $\zeta_m \in Z_m$ here corresponds to $\zeta_m = (\bar{s}_0, \bar{a}_0, ..., \bar{a}_{T-1}, \bar{s}_T)$ where $\bar{s} \in S_m$ and $\bar{a} \in A_m$. We also define an agent specific rollout $\zeta_m^i \in Z_m^i$, $\zeta_m^i = (s_0^i, a_0^i, ..., a_{T-1}^i, s_T^i)$ where $s^i \in S^i$ and $a^i \in A^i$. D is the initial state distribution over S_m .

Agents attempt to learn a policy $\pi^i : S^i \to \Delta(A^i)$ such that $\mathbb{E}\left[\sum_t R^i_t | \pi^i, \pi^{-i}\right]$ is maximized, where $\Delta(A^i)$ is a probability distribution over A^i and π^{-i} is the set of all policies apart from π^i . We use $\Pi = {\pi^i}_{i \in \mathcal{N}}$ to denote the set of all agent policies. For simplicity, we restrict our formulation to a *homogeneous* set of agents which operate over the same state $(S^i = S_A)$ and action space $(A^i = A_A)$.

SPECTRL Jothimurugan *et. al* [5] introduce a specification language for reinforcement learning problems built using temporal logic constraints and predicates. It is shown to be adept at handling complex compositions of task specifications through the use of a *task monitor* and well-defined monitor transition rules. Notably, one can encode Non-Markovian tasks into the MDP using the additional states of the automaton (task monitor) compiled from the given specification.

The atomic elements of this language are Boolean predicates b defined as functions of a state S with output $\llbracket b \rrbracket : S \to \mathbb{B}$. These elements have quantitative semantics $\llbracket b \rrbracket_q$ with the relation being $\llbracket b \rrbracket(s) = \text{True} \iff \llbracket b \rrbracket_q(s) > 0$. Specifications ϕ are Boolean functions of the state trajectory $\zeta = (s_1, s_2, ..., s_T)$. The specification language also includes composition functions for a specification ϕ and Boolean predicate b, with the language defined as

$$\phi ::= \texttt{achieve} \ b \mid \phi \texttt{ ensuring} \ b \mid \phi_1; \ \phi_2 \mid \phi_1 \texttt{ or } \phi_2$$

The description of these functions is as follows. achieve b is true when the trajectory satisfies b at least once. ϕ ensuring b is true when b is satisfied at all timesteps in the trajectory. ϕ_1 ; ϕ_2 is a sequential operator that is true when, in a given trajectory $\zeta = (s_1, s_2, ..., s_T), \exists k > 1$ such that $\phi_1(s_1, ..., s_k)$ is true and $\phi_2(s_{k+1}, ..., s_T)$ is true. In other words, ϕ_1 ; ϕ_2 represents an ordered sequential completion of specification ϕ_1 followed by ϕ_2 . Lastly, ϕ_1 or ϕ_2 is true when a trajectory satisfies either ϕ_1 or ϕ_2 .

Given a specification ϕ on a Markov Decision Process $\langle S, D, A, P, T \rangle$ (MDP) defined using SPECTRL, a task monitor $\langle Q, X, \Sigma, U, \Delta, q_0, v_0, F, \rho \rangle$ (a finite state automaton [20]) is compiled to record the completion status of tasks with monitor states Q; final monitor states F denote a satisfied trajectory. This is used to create an augmented version of the MDP $\langle \tilde{S}, \tilde{s_0}, \tilde{A}, \tilde{P}, \tilde{R}, \tilde{T} \rangle$ with an expanded state, action space and modified reward function. The task monitor provides a scoring function for trajectories in the augmented MDP to guide policy behavior.

While SPECTRL has been shown to work with trajectory-based algorithms for reinforcement learning [12], it is not immediately evident how to translate it to common RL algorithms such as DDPG [10] and PPO [16]. A simple solution would be to keep the episodic format with a trajectory $\zeta = (s_0, \dots, s_T)$ and assign the trajectory value of SPECTRL (a function of ζ) to the final state transition in the trajectory $s_{T-1} \rightarrow s_T$ and zero for all other states. Importantly, the maintains the trajectory ordering properties of SPECTRL in the episodic return $(\sum_{t=0}^{T} r_t)$.

3 Problem Statement

Directly appropriating SPECTRL for our use case of imposing specifications on multiagent problems poses significant scalability issues. Consider the case

 $\phi_a = \texttt{achieve}(\texttt{reach}(P)); \texttt{achieve}(\texttt{reach}(Q))$

where [[reach(P)]] = True when an agent reaches state P. To ease the illustration of our framework, we assume that all agents are homogeneous, *i.e.* $S^i = S_A, \forall i \in \mathcal{N}$. Now, the state space of the entire multi-agent system is $S = (S_A)^N$ for N agents (we omit m for perspicuity).

If the predicate reach was defined on the entire state S, it would yield a specification forcing synchronization between agents. On the other hand, if reach was defined on the agent state S_A , then it would create a *localized* specification where synchronization is not required. This would be akin to allowing individual agents to act independently of other agent behaviors.

However, using a centralized task monitor for the localized predicate would cause the number of monitor states to exponentially increase with the number of agents Nand subtasks K since the possible stages of task completion would be $\mathcal{O}(K^N)$.

To address this scalability issue, the benefits of task monitor distribution are apparent. In the case of ϕ_a above, assume reach is defined on the local state space S_A . If each agent had a separate task monitor stored locally to keep track of the task completion stages, the new number of monitor states is now reduced to $\mathcal{O}(NK)$.

Consider an example of robots in a warehouse. A few times a day, all robots must gather at a common point for damage inspection at the same time (akin to a global reach) to minimize the frequency of inspection (an associated cost). To ensure satisfaction of the entire specification, the reward given to an RL agent learning this objective should capture both the global and local tasks. For example, if the global reach task for the routine inspection is made local instead, the cost incurred may be larger than if it was a synchronized global objective.

Main Objective Given a specification ϕ on a system of N agents, we wish to find policies $\Pi = {\pi^1, \dots, \pi^N}$ to maximize the probability of satisfying ϕ for all agents. Formally, we seek

$$\varPi^* \in \mathop{\mathrm{arg\,max}}_{\pi^1, \cdots, \pi^N} \, \Pr_{\zeta_m \sim D_\varPi} \left[\llbracket \phi(\zeta_m) \rrbracket = \mathtt{True} \right]$$

where D_{Π} is the distribution of all system rollouts when all agents collectively follow policy set Π . We emphasize that ϕ acts on the entire rollout, $\phi : Z_m \to \{0, 1\}$ and not in an agent-specific manner, $\phi' : Z_m^i \to \{0, 1\}$. This discourages agents from attempting to simply satisfy their local objectives while preventing the system from achieving necessary global ones.

4 SPECTRL in a Multi-Agent World

Unlike the single agent case, multi-agent problems have two major classes of objectives. Agents have individual goals to fulfill as well as collective goals that require coordina-



Fig. 1: Example Composite Task Monitor for specification ϕ_{ex} (Sec. 4) with 4 task goals denoted by Q,P,P' and R where the agent starts at q_0 . Double circles represent *final* states while green circles represent *global* states. The diagram removes a state between q_1 and (q_3, q_4) as well as self-loops for ease of explanation.

tion and/or global system knowledge. These individual goals are often only dependent on the agent-specific state s^i while collective goals require full system knowledge \bar{s} .

Consequently, for a multi-agent problem, we see the need for two types of predicates *viz.* local and global. Local predicates are of the form $p_{lo} : S_A \to \mathbb{B}$ whereas global predicates have the form $p_{gl} : S \to \mathbb{B}$ where \mathbb{B} is the Boolean space. We introduce two simple extensions of reach [5] to demonstrate the capabilities of this distinction.

Local predicates are defined with respect to each agent and represent our individual goals. As an example, closely related to the problems observed in SPECTRL, we introduce the following local predicates for a state $s_a \in S_A$,

$$[\![reach_{lo}x]\!](s_a) = (||s_a - x||_{\infty} < 1)$$

which represents reaching near location x in terms of the L_{∞} norm. Now to enforce global restrictions, we introduce counterparts to these predicates that act on a global state $\bar{s} \in S$.

$$\llbracket \texttt{reach}_{gl}\bar{x} \rrbracket(\bar{s}) = (||\bar{s} - \bar{x}||_{\infty} < 1)$$

where we now have a set of locations $\bar{x} \in S$.

As in SPECTRL, each of these predicates b require quantitative semantics $[\![b]\!]_q$ to facilitate our reward shaping procedure. We define these semantics as follows:

- reach_{lo} has the same semantics as reach in [5] yet is defined on space S_A .

$$[\![\texttt{reach}_{lo}x]\!]_{a}(s_{a}) = 1 - d_{\infty}(s_{a}, x)$$

where $d_{\infty}(a, b)$ represents the L_{∞} distance between a and b with the usual extension to the case where b is a set.

- reach_{*ql*} is defined on the state space S as

$$[\![\texttt{reach}_{gl}\bar{x}]\!]_q(\bar{s}) = 1 - d_\infty(\bar{s},\bar{x})$$

We observe that the same composition rules can apply to these predicates and we thus attempt to solve RL systems described with these compositions. As shown in

5

Sec. 3, using a centralized SPECTRL compilation algorithm on the entire state space, even for simple sequences of tasks, leads to an explosion in monitor states. We, therefore, distribute task monitors over agents to handle scalability. Furthermore, we also need to change SPECTRL's compilation rules to handle mixed objective compositions such as¹

$$\phi = \operatorname{reach}_{lo}(P); \operatorname{reach}_{ql}(Q); \operatorname{reach}_{ql}(R)$$

To compile these specifications into a usable format, we utilize a *composite task monitor* as described in Sec. 5 and develop a new algorithm to achieve our goal. As an example, see Fig. 1 depicting a task monitor whose specification is:

 $\phi_{ex} = \operatorname{reach}_{al}(P') \operatorname{or} \operatorname{reach}_{lo}(Q); [\operatorname{reach}_{lo}(P) \operatorname{or} \operatorname{reach}_{al}(R)]$

Here, we have 4 task goals denoted by P, Q, R and P'. The agents all start at the root node q_0 . States q_2 , q_3 and q_4 are all final states in the task monitor while q_2 and q_4 are global monitor states. As shown in Sec. 6, q_0 and q_1 are a synchronization states. While it may seem that agents only require coordination at global states, it is also necessary for the agents to have the same task transition at these synchronization states as well.

5 Compilation Steps

Given a specification ϕ and the Markov game \mathcal{M}_g , we create a task monitor M that is distributed among agents by making agent-specific copies. This is used to create an augmented Markov game $\mathcal{M}'_g = \langle \mathcal{N}, \{\tilde{\mathcal{S}}_A\}_{i \in \mathcal{N}}, \tilde{D}, \{\tilde{A}_A\}_{i \in \mathcal{N}}, \tilde{P}, \{\tilde{R}^i\}_{i \in \mathcal{N}}, T \rangle$ on which the individual agent policies are trained.

Create Composite Task Monitor When the types of specifications are divided into two based on the domain, the solution can be modeled with a *composite task monitor* $M_{\phi} = \langle Q, \tilde{X}, \tilde{\Sigma}, \tilde{U}, \tilde{\Delta}, q_0, v_0, F, \rho \rangle$. As in SPECTRL, Q is a finite set of monitor states. $\tilde{X} = X_l \cup X_g$ is a finite set of registers that are partitioned into X_l for local predicates and X_g for global predicates. These registers are used to keep track of the degree of completion of the task at the current monitor state for local and global tasks respectively.

We describe below how to use the compiled composite task monitor to create an augmented Markov game \mathcal{M}'_g . Each $\tilde{\mathcal{S}}_A$ in \mathcal{M}'_g is an augmented state space with an augmented state being a tuple $(s_A, q, v) \in \mathcal{S}_A \times Q \times V$ where $V \in \mathbb{R}^X$ and $v \in V$ is a vector describing the register values.

 $\Delta = \Delta_l \cup \Delta_g$ houses the transitions of our task monitor. We require that: i) different transitions are allowed only under certain conditions defined by our states and register values; and, ii) furthermore, they must also provide rules on how to update the register values during each transition. To define these conditions for transition availability, we use $\tilde{\Sigma} = \Sigma_l \cup \Sigma_g$ where Σ_l is a set of predicates over $S_A \times V$ and Σ_g is a set of predicates over $S \times V$. Similarly, $\tilde{U} = U_l \cup U_g$ where U_l is a set of functions $u_l : S_A \times V \to V$ and U_g is a set of functions $u_g : S \times V \to V$. Now, we can define $\tilde{\Delta} \subseteq Q \times \tilde{\Sigma} \times V$

¹ We omit achieve in achieve(reach_{lo}(P)) and achieve(reach_{gl}(P)) from here on to reduce clutter; this specification is implied when we compose reach(P) with ; and or.

 $\tilde{U} \times Q$ to be a finite set of transitions that are non-deterministic. Transition $(q, \sigma, u, q') \in \tilde{\Delta}$ is an augmented transition either representing $(s^i, q, v) \xrightarrow{a^i \mid \Pi_{-i}} ((s^i)', q', u_l(s^i, v))$ or the form $(\bar{s}, q, v) \xrightarrow{a^i \mid \Pi_{-i}} (\bar{s}', q', u_g(s, v))$ depending on whether $\sigma \in \Sigma_l$ or $\sigma \in \Sigma_g$ respectively. Let $\delta_l \in \Delta_l$ represent the former (localized) and $\delta_g \in \Delta_g$ the latter (global) transition types. Here Π_{-i} denotes the policy set of all agents except agent *i*. Lastly, q_0 is the initial monitor state and v_0 is the initial register value (for all agents), $F \subset Q$ is the set of final monitor states, and $\rho : S \times F \times V \to \mathbb{R}$ is the reward function.

Copies of these composite task monitors M are distributed over agents \mathcal{N} to form the set $\{M^i\}_{i \in \mathcal{N}}$. These individually stored task monitors are used to let each agent $i \in \mathcal{N}$ keep track of its subtasks and the degree of completion of those subtasks by means of monitor state q^i and register value v^i .

Create Augmented Markov Game From our specification ϕ we create the augmented Markov game $\mathcal{M}'_g = \langle \mathcal{N}, \{\tilde{\mathcal{S}}_A\}_{i \in \mathcal{N}}, \tilde{D}, \{\tilde{A}_A\}_{i \in \mathcal{N}}, \tilde{P}, \{\tilde{R}^i\}_{i \in \mathcal{N}}, T \rangle$ using the compiled composite task monitor M. A set of policies \tilde{H}^* that maximizes rewards in \mathcal{M}'_g should maximize the chance of the specification ϕ being satisfied.

Each $S_A = S_A \times Q \times V$ and $D = (\{s_0\}_{i \in \mathcal{N}}, q_0, v_0)$. We use Δ to augment the transitions of P with monitor transition information. Since Δ may contain nondeterministic transitions, we require the policies \tilde{H} to decide which transition to choose. Thus $\tilde{A}_A = A_A \times A_{\phi}$ where $A_{\phi} = \Delta$ chooses among the set of available transitions at a monitor state q. Since monitors are distributed among all agents in \mathcal{N} , we denote the set of current monitor states as $\bar{q} = \{q^i\}_{i \in \mathcal{N}}$ and the set of register values as $\bar{v} =$ $\{v^i\}_{i \in \mathcal{N}}$. Now, each agent policy must output an *augmented action* $(a, \delta) \in \tilde{A}_A$ with the condition that $\delta_l = (q, \sigma_l, u_l, q')$ is possible in local augmented state $\tilde{s}_a = (s_a, q, v)$ if $\sigma_l(s_a, v)$ is True and $\delta_g = (q, \sigma_g, u_g, q')$ is possible in global augmented state $\tilde{s} =$ $(\bar{s}, \bar{q}, \bar{v})$ if $\sigma_g(\bar{s}, v)$ is True. We can write the augmented transition probability \tilde{P} as,

$$\hat{P}((\bar{s}, q, v), (a, (q, \sigma, u, q')), (\bar{s}', q', u(\bar{s}, v)))) = P(\bar{s}, a, \bar{s}')$$

for transitions $\delta_g \in \Delta_g$ with $(\sigma, u) = (\sigma_g, u_g)$ and transitions $\delta_l \in \Delta_l$ with $(\sigma, u) = (\sigma_l, u_l)$. Here, we let $u_l(\bar{s}, v) = u_l(s^i, v)$ for agent *i* since s^i is included in \bar{s} . An augmented rollout $\tilde{\zeta}_m$ where

$$\zeta_m = ((\bar{s}_0, \bar{q}_0, \bar{v}_0), \bar{a}_0, ..., \bar{a}_{T-1}, (\bar{s}_T, \bar{q}_T, \bar{v}_T))$$

is formed by these augmented transitions. To translate this trajectory back into the Markov game \mathcal{M}_q we can perform projection $\operatorname{proj}(\tilde{\zeta}_m) = (\bar{s}_0, \bar{a}_0, ..., \bar{a}_{T-1}, \bar{s}_T)$.

Determine Shaped Rewards Now that we have the augmented Markov game \mathcal{M}'_g and compiled our composite task monitor, we proceed to form our reward function that encourages the set of policies Π to satisfy our specification ϕ . We can perform shaping in a manner similar to SPECTRL's single-agent case on our distributed task monitor. Crucially, since reward shaping is done during the centralized training phase, we can assume we have access to the entire augmented rollout namely $\tilde{s}_t = (\bar{s}_t, \bar{q}_t, \bar{v}_t)$ at any

given $t \in [0, T]$. From the monitor reward function ρ , we can determine the weighting for a complete augmented rollout as

$$\tilde{R}^{i}(\tilde{\zeta}_{m}) = \begin{cases} \rho(\bar{s}_{T}, q_{T}^{i}, v_{T}^{i}), & \text{if } q_{T}^{i} \in F \\ -\infty & \text{otherwise} \end{cases}$$

Theorem 1. (Proof in Appendix Sec. F.) For any Markov game \mathcal{M}_g , specification ϕ and rollout ζ_m of \mathcal{M}_g , ζ_m satisfies ϕ if and only if there exists an augmented rollout $\tilde{\zeta}_m$ such that i) $\tilde{R}^i(\tilde{\zeta}_m) > 0 \forall i \in \mathcal{N}$ and ii) $\operatorname{proj}(\tilde{\zeta}_m) = \zeta_m$.

The \tilde{R}^i specified is $-\infty$ unless a trajectory reached a final state of the composite task monitor. To reduce the sparsity of this reward signal, we transform this into a shaped reward \tilde{R}^i_s that gives partial credit to completing subtasks in the composite task monitor.

Define for a non-final monitor state $q \in Q \setminus F$, function $\alpha : S \times Q \times V \to \mathbb{R}$.

$$\alpha(\bar{s},q,v) = \max_{(q,\sigma,u,q') \in \Delta, q \neq q'} \llbracket \sigma \rrbracket_q(\bar{s},v)$$

This represents how close an augmented state $\tilde{s} = (\bar{s}, q, v)$ is to transition to another state \tilde{s}' with a different monitor state. Intuitively, the larger α is, the higher the chance of moving deeper into the task monitor. In order to use this definition on all σ , we overload σ_l to also act on elements $\bar{s} = \{s^i\}_{i \in \mathcal{N}} \in S$ by yielding for agent *i*, the value $\sigma_l(\bar{s}) = \sigma_l(s^i)$.

Let C_l be a lower bound on the final reward at a final monitor state, and C_u being an upper bound on the absolute value of α over non-final monitor states. Also for $q \in Q$, let d_q be length of the longest path from q_0 to q in the graph M_{ϕ} (ignoring the self-loops in Δ) and $D = \max_{q \in Q} d_q$. For an augmented rollout $\tilde{\zeta}_m$ let $\tilde{s}_k = (\bar{s}_k, q_k^i, \bar{v})$ be the first augmented state in $\tilde{\zeta}_m$ such that $q_k^i = q_{k+1}^i = \cdots = q_T^i$. Then we have the shaped reward,

$$\tilde{R}_{s}^{i}(\tilde{\zeta}_{m}) = \begin{cases} \max_{k \leq j < T} \alpha(\bar{s}_{j}, q_{T}^{i}, v_{j}) + 2C_{u} \cdot (d_{q_{T}^{i}} - D) + C_{l} & \text{if } q_{T}^{i} \notin F \\ \tilde{R}^{i}(\tilde{\zeta}_{m}) & \text{otherwise} \end{cases}$$
(1)

Theorem 2. (Proof in Appendix Sec. F.) For two augmented rollouts $\tilde{\zeta}_m, \tilde{\zeta}'_m$, (i) if $\tilde{R}^i(\tilde{\zeta}_m) > \tilde{R}^i(\tilde{\zeta}'_m)$, then $\tilde{R}^i_s(\tilde{\zeta}_m) > \tilde{R}^i_s(\tilde{\zeta}'_m)$, and (ii) if $\tilde{\zeta}_m$ and $\tilde{\zeta}'_m$ end in distinct non-final monitor states q^i_T and $(q^i_T)'$ such that $d_{q^i_T} > d_{(q^i_T)'}$, then $\tilde{R}^i_s(\tilde{\zeta}_m) \ge \tilde{R}^i_s(\tilde{\zeta}'_m)$.

6 Sub-task Synchronization

Importance of Task Synchronization Consider the following example specification:

$$\phi'_{1a} = \operatorname{reach}_{lo}(P) \operatorname{or} \operatorname{reach}_{gl}(Q)$$

where P, Q are some goals. To ensure flexibility with respect to the possible acceptable rollouts within ϕ'_{1a} , the individual agent policies π^i are learnable and the task transition chosen is dependent on the agent-specific observations. This flexibility between agents however, adds an additional possible failure method in achieving a global specification - if even a single agent attempts to fulfill the global objective while the others decide to follow their local objectives, the specification would never be satisfied.



Fig. 2: Overview of the DistSPECTRL process for task synchronization. Branching in the task monitor diagram denotes potential non-deterministic choices between future tasks (such as in ϕ_{ex}). Left to right represents the order of policy actions over a trajectory. Green states represent the current monitor state of that agent.

Identifying Synchronization States As emphasized above, task synchronization is an important aspect of deploying these composite task monitors in the Markov game \mathcal{M}_g with specification ϕ . We show the existence of a subset of monitor states Sync $\in Q$ where in order to maintain task synchronization, agents simply require a consensus on which monitor transition $\delta = (q, \sigma, u, q')$ to take. If we use Q_g to symbolize the set of global monitor states, viz. all $q \in Q$ such that $\exists (q, \sigma_g, u_g, q') \in \Delta_g$, then we see that $Q_g \subseteq$ Sync. A valid choice for $q \in$ Sync with $q \notin Q_g$ is all branching states in the graph of M_{ϕ} with a set refinement presented in the Appendix (Sec. D).

During training, we enforce the condition that when an agent *i* has monitor state $q_t^i \in \text{Sync}$, it must wait for time $t_1 > t$ such that $q_{t_1}^j = q_t^i \ \forall j \in \mathcal{N}$ and then choose a common transition as the other agents. This is done during the centralized training phase by sharing the same transition between agents based on a majority vote.

7 Multi-Agent Specification Properties

Consider a specification ϕ and let $\mathcal{N} = \{1, \dots, N\}$ be the set of all agents with ζ_m being a trajectory sampled from the environment. $\phi(\zeta_m, n)$ is used to denote that the specification is satisfied on ζ_m for the set of agents $n \subseteq \mathcal{N}$ (i.e. $[\![\phi(\zeta_m, n)]\!] == \text{True}$).

MA-Distributive Many specifications pertaining to MA problems can be satisfied independent of the number of agents. At its core, we have the condition that a specification being satisfied with respect to a union of two disjoint sets of agents implies that it can be satisfied on both sets independently. Namely if $n_1, n_2 \subset \mathcal{N}$ with $n_1 \cap n_2 = \emptyset$ then an MA-Distributive specification satisfies the following condition:

$$\phi(\zeta_m, n_1 \cup n_2) \implies \phi(\zeta_m, n_1) \land \phi(\zeta_m, n_2)$$

MA-Decomposable Certain specifications satisfy a decomposibility property particular to multi-agent problems that can help in scaling with respect to the number of agents.

Say $\exists k \in \{1, \dots, N-1\}$ such that

$$\phi(\zeta_m, \mathcal{N}) \implies \phi_k(\zeta_m, \mathcal{N}) = \bigwedge_{j \in \{1, \dots, J\}} \phi(\zeta_m, n_j)$$

where

$$n_j \subset \mathcal{N} , \ k \le |n_j| < N , \ J = \lfloor \frac{N}{k} \rfloor , \ \bigcap_j n_j = \emptyset , \ \bigcup_j n_j = \mathcal{N}$$

with $\lfloor \rfloor$ representing the floor function. Each n_j is a set of at least k unique agents and $\{n_j\}_j$ forms a partition over \mathcal{N} .

We then call the specification ϕ *MA-Decomposable* with decomposibility factor k. Here ϕ_k can be thought of as a means to approximate the specification ϕ to smaller groups of agents within the set of agents \mathcal{N} . Provided we find a value of k, we can then use this as the basis of our MA-Dec scaling method to significantly improve training times for larger numbers of agents.

Theorem 3. (Proof in Appendix Sec. F.) All MA-Distributive specifications are also MA-Decomposable with decomposability factors $k \in \mathbb{Z}^+$, $1 \le k < N$.

Notably all compositions of $reach_{gl}$ and $reach_{lo}$ within our language are MA-Distributive and are thus MA-Decomposable with factor k = 2.² This is far from a general property however, as one can define specifications on N robots such as achieve("collect x fruits") where each robot can carry at most x/N fruits. In this case, no single subset of agents can satisfy the specification as the total capacity of fruits would be less than x and the specification is neither MA-Distributive nor MA-Decomposable.

8 Algorithm

Training Agents learn $\pi^i(s^i, v^i, q^i) = (a^i, \delta^i)$ on the augmented Markov game \mathcal{M}'_g where s^i, v^i, q^i are agent-specific state, register value and task monitor state respectively. Since training is centralized, all agent task monitors receive the same global state. Based on our discussion in Sec. 6, if an agent is in any given global monitor state, we wait for other agents to enter the same state, then do the arg max task transition for all agents in the same state. In addition, at the synchronization states (Sec. 6), we perform a similar process to select the task transition. These trained *augmented policies* are then projected into policies that can act in the original \mathcal{M}_q .

² It is satisfied with k = 1 as well but this is the trivial case where reach_{gl} and reach_{lo} are equivalent.

11



Fig. 3: Example MA-Dec Scaling Process with N = 10, k = 2, f = 2 on an MA-Decomposable spec ϕ with decomposability factor 2. At Stage 1, $g_1 = 2$ to start with 5 groups. Next $g_2 = fg_1 = 4$ which forms 2 groups. Finally at Stage 3, $g_3 = fg_2 = 8$ which forms one group (\mathcal{N}).

Scaling MA-Decomposable Specifications Our algorithm for scaling based off the MA-Decomposable property is shown in Alg. 1 (refer Appendix) and we name it MA-Dec scaling. Essentially, we approximate the spec. ϕ by first independently considering smaller groups within the larger set of agents \mathcal{N} and try to obtain a policy satisfying ϕ on these smaller groups. By progressively making the group sizes larger over stages and repeating the policy training process while continuing from the previous training stage's policy parameters, we form a curriculum that eases solving the original problem ϕ on all agents \mathcal{N} .

In Fig. 3 we demonstrate MA-Dec scaling for N = 10 agents on a spec. ϕ which is MA-Decomposable with decomposability factor 2. For this example we set the scaling parameters k = 2 and f = 2. Initially we have a min. group size $g_1 = 2$ and this is changed to $g_2 = 4$ and $g_3 = 8$ from setting the scaling factor. We increment the stage number every time all the groups of a stage have satisfied the entire specification ϕ w.r.t. their group. While separating training into stages, agents must be encouraged to move from stage i to stage i + 1. To ensure this, we need to scale rewards based on the stage. We chose a simple linear scaling where for stage number i and time step t, each agent receives reward $r_{i,t} = ic_k + CTM_{i,t}$ where $CTM_{i,t}$ is the original composite task monitor reward at stage i and $c_k \in \mathbb{R}$ is a constant. By bounding the reward terms such that rewards across stages are monotonically increasing $(r_{i,t} < r_{i+1,t'})$ we can find a suitable c_k to be $(2D + 1)C_u$ (refer Appendix Sec. B) where the terms are the same as in Eq. 1.

From setting the initial min. group size g_1 and scaling factor f, we get the total number of learning stages (T_s) as $T_s = \lfloor \log_f(N) - \log_f(k) \rfloor = \mathcal{O}(\log_f(N))$. We build the intuition behind why MA-Dec scaling is effective in the Appendix (Sec. B), by describing it as a form of curriculum learning.

Deployment Policies are constructed to proceed with only local information (s^i, v^i, q^i) . Since we cannot share the whole system state with the agent policies during deployment yet our composite task monitor requires access to this state at all times, we allow the following relaxations: 1) Global predicates $\sigma_g(\bar{s}, v)$ enabling task monitor transitions need global state and access it during deployment. 2) Global register updates $u_g(\bar{s}, v)$ are also a function of global state and access it during deployment.

In order to maintain task synchronization, agents use a consensus based communication method to decide task monitor transitions at global and synchronization states. If

agents choose different task transitions at these monitor states, the majority vote is used as done during training.

9 Experiment Setup

Our experiments aim to validate that the use of a distributed task monitor can achieve synchronization during the deployment of multiple agents on a range of specifications.

In addition, to emphasize the need for distribution of task monitors to alleviate the state space explosion caused by mixing local and global specifications, we include experiments with SPECTRL applied to a centralized controller.

Lastly, we provide results showing the efficacy of the MA-Dec scaling approach for larger numbers of agents when presented with a specification that satisfies the MA-Decomposability property (Sec. 7).

As a baseline comparison, we also choose to run our algorithm without giving policies access to the monitor state (**no_mon**). These are trained with the same shaped reward as DistSPECTRL. We also provide a Reward Machine baseline (**RM**) for ϕ_1 with continuous rewards since ϕ_1 is similar to the 'Rendezvous' specification in [13].

Environment Our first set of experiments are done on a 2D Navigation problem with N = 3 agents. The observations ($S \in \mathbb{R}^2$) used are coordinates within the space with the action space ($A \in \mathbb{R}^2$) providing the velocity of the agent.

The second set of experiments towards higher dimension 3D benchmarks, represent particle motion in a 3D space. We train multiple agents (N = 3) in the 3D space ($S \in \mathbb{R}^3$) with a 3D action space ($A \in \mathbb{R}^3$) to show the scaling potential of our framework.

The final set of experiments were on a modern discrete-action MARL benchmark built in Starcraft 2 [15] with N = 8 agents (the "8m" map). Each agent has 14 discrete actions with a state space $S \in \mathbb{R}^{80}$ representing a partial view of allies and enemies.

Algorithm Choices For the scaling experiments (Fig. 6) we used the 2D Navigation problem with horizon T = 500 and the scaling parameters³ k = 2 and f = 2. We also choose a version of PPO with a centralized Critic to train the augmented Markov Game noting that our framework is agnostic to the choice of training algorithm. The current stage is passed to the agents as an extra integer dimension. For other experiments we chose PPO with independent critics as our learning algorithm. Experiments were implemented using the RLLib toolkit [9].

Specifications (2D Navigation) The evaluated specifications are a mix of local and global objectives. The reach predicates have an error tolerance of 1 (the L_{∞} distance from the goal).

(i) $\phi_1 = \operatorname{reach}_{g1}(5,0)$; $\operatorname{reach}_{g1}(0,0)$, (ii) $\phi_2 = \phi_1$; $\operatorname{reach}_{g1}(3,0)$ (iii) $\phi_3 = \operatorname{reach}_{lo}(5,0)$; $\operatorname{reach}_{g1}(0,0)$; $\operatorname{reach}_{g1}(3,0)$ (iv) $\phi_4 = [\operatorname{reach}_{lo}(3,0) \circ \operatorname{reach}_{lo}(5,10)]$; ϕ_3

³ While we could start with k = 1, we set k = 2 to reduce the number of learning stages.

13

(SC2) ϕ_{sc} represents 'kiting' behaviour and is explained further in the Appendix (Sec. E). $\phi_{sc} = \phi_{sc_a}; \phi_{sc_a}; \phi_{sc_a}$ where $\phi_{sc_a} = away_from_enemy_{gl}; shooting_range_{lo};$ (3D Environment) $\phi_a = reach_{1o}(5,0,0); reach_{g1}(0,0,0); reach_{g1}(3,0,0)$ is the specification considered within X-Y-Z coordinates.



Fig. 4: Satisfaction percentages on specifications ϕ_1, ϕ_2, ϕ_3 and ϕ_4 with N = 3 agents. The shaded regions show the maximum and minimum achieved over 5 separate evaluation runs



Fig. 5: Specification satisfaction percentages (left) for the StarCraft 2 specification ϕ_{sc} with N = 8 agents and (right) for the 3D Navigation experiments on specification ϕ_a .



Fig. 6: Specification satisfaction percentages for N = 10 agents on ϕ_1 (left) and N = 6 agents on ϕ_3 (right) comparing the MA-Dec scaling (red) to centralized SPECTRL (blue) and vanilla DistSPECTRL (green) i.e. without scaling enhancements.

Spec.	DistSpectrl	. no_mon	Spectrl
ϕ_1	99.62	91.17	100.00
ϕ_2	99.05	00.00	97.38
ϕ_3	97.59	94.77	96.81
ϕ_4	97.31	00.00	90.78
ϕ_a	98.49	00.00	99.60
ϕ_{sc}	86.79	00.00	00.00

Table 1: Specification satisfaction percentages on convergence for Fig.4,5

Table	2:	Specification	satisfaction	percentages	on		
convergence for Fig.6, (Scaling to more Agents)							

_	U	U	0	0 /		
Spec. / # Agents MA-Dec DistSPECTRL SPECTRL						
	ϕ_3/N =6	94.09	0.00	0.00		
	$\phi_1/N = 6$ $\phi_1/N = 10$	97.83	80.67 72.30	98.96 99.28		
_	7 1/ - / 10	1 2 100	. 100			

10 Results

Handling Expressive Specifications The experiments in Fig. 4 demonstrate execution when the task monitor predicates have access to the the entire system state. This provides agents with information sufficient to calculate global predicates for task monitor transitions. The overall satisfaction percentage is reported with the value 0 being an incomplete task to 1.0 being the entire specification satisfied.

While SPECTRL has often been shown to be more effective [5,6] than many existing methods (e.g. **RM** case) for task specification, the further utility of the monitor state in enhancing coordination between agents is clearly evident in a distributed setting. The task monitor state is essential for coordination as our baseline **no_mon** is often unable to complete the entire task (even by exhaustively going through possible transitions) and global task completion requires enhanced levels of synchronization between agents.

From Table 1 we see that upon convergence of the learning algorithm, the agent is able to maintain a nearly 100% task completion rate for our tested specifications, a significant improvement in comparison to the **no_mon** case, showing the importance of the task monitor as part of a multi-agent policy.

Benefits of Distribution Over Centralization The centralized SPECTRL graphs (blue curves in Figs. 4, 5, 6) show that while distribution may not be necessary for certain specifications with few local portions (e.g. ϕ_2), concatenating them will quickly lead to learning difficulties with larger number of agents (Fig. 5, ϕ_{sc} and Fig. 6, ϕ_3). This difficulty is due in large part to state space explosion of the task monitor in these cases as is apparent by the significantly better performance of our distributed algorithm. We also remind the reader that a centralized algorithm is further disadvantageous in MARL settings due to the added synchronization cost between agents during deployment.

Scaling to Larger State Spaces The results in Fig. 5 show promise that the DistSPECTRL framework can be scaled up to larger dimension tasks as well. The 3D environment results exhibits similar behavior to the 2D case with the **no_mon** showing difficulty in progressing beyond the local tasks in the larger state space with sparser predicates. The ϕ_{sc} results also show promise in defining relevant predicates and achieving general specifications for modern MARL benchmarks.

Scaling to More Agents Table 2 and Fig. 6 demonstrate the benefits of MA-Dec scaling for larger N when presented with an MA-Decomposable specification. At smaller ranges of N as well as less complex combinations of mixed and global objectives, the effect of MA-Dec scaling is less pronounced. We observe that the stage based learning is crucial for a even simple mixed specification like ϕ_3 with as little as N = 6 agents.

11 Related Work

Multi-agent imitation learning [17,7,23] uses demonstrations of a task to specify desired behavior. However in many cases, directly being able to encode a specification by means of our framework is more straightforward and removes the need to have demonstrations beforehand. Given demonstrations, one may be able to infer the specification [21] and make refinements or compositions for use in our framework.

TLTL[8] is another scheme to incorporate temporal logic constraints into learning enabled controllers, although its insufficiency in handling non-Markovian specifications led us to choose SPECTRL as the basis for our methodology. Reward Machines (RMs) [19,2,18] are an automaton-based framework to encode different tasks into an MDP. While RMs can handle many non-Markovian reward structures, a major difference is that SPECTRL starts with a logical temporal logic specification and includes with the automaton the presence of memory (in the form of registers capable of storing realvalued information). Recent work [6] shows the relative advantages SPECTRL-based solutions may have over a range of continuous benchmarks.

Concurrent work has introduced the benefits of a temporal logic based approach to reward specification [4]. While experimental results are not yet displayed, the convergence guarantees of the given algorithm are promising. Since we use complex nonlinear function approximators (neural networks) in our work, such guarantees are harder to provide. Reward Machines have also been explored as a means of specifying behavior in multi-agent systems [13] albeit in discrete state-action systems that lend themselves to applying tabular RL methods such as Q-learning. One may extend this framework to continuous systems by means of function approximation but to the best of our knowledge, this has not been attempted yet. Similar to our synchronization state, the authors use a defined local event set to sync tasks between multiple agents and requires being aware of shared events visible to the other agents.

In the same spirit as our stage-based approach, transferring learning from smaller groups of agents to larger ones has also been explored [22]. Lastly, while we chose PPO to train the individual agents for its simplicity, our framework is agnostic to the RL algorithm used and can be made to work with other modern multi-agent RL setups [11,3] for greater coordination capabilities.

12 Conclusion

We have introduced a new specification language to help detail MARL tasks and describe how it can be used to compile a desired description of a distributed execution in order to achieve specified objectives. Our framework makes task synchronization realizable among agents through the use of: 1) Global predicates providing checks for task completion that are easily computed, well-defined and tractable; 2) A monitor state to keep track of task completion; and 3) Synchronization states to prevent objectives from diverging among agents.

Acknowledgements This work was supported in part by C-BRIC, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., Zaremba, W.: Hindsight experience replay. In: NIPS (2017)
- Camacho, A., Toro Icarte, R., Klassen, T.Q., Valenzano, R., McIlraith, S.A.: LTL and beyond: Formal languages for reward function specification in RL. In: IJCAI (2019)
- Foerster, J., Farquhar, G., Afouras, T., Nardelli, N., Whiteson, S.: Counterfactual multi-agent policy gradients. In: AAAI (2018)
- Hammond, L., Abate, A., Gutierrez, J., Wooldridge, M.: Multi-agent reinforcement learning with temporal logic specifications. In: AAMAS (2021)
- Jothimurugan, K., Alur, R., Bastani, O.: A composable specification language for reinforcement learning tasks. In: NeurIPS (2019)
- Jothimurugan, K., Bansal, S., Bastani, O., Alur, R.: Compositional reinforcement learning from logical specifications. In: NeurIPS (2021)
- 7. Le, H.M., Yue, Y., Carr, P., Lucey, P.: Coordinated MA imitation learning. In: ICML (2017)
- 8. Li, X., Vasile, C., Belta, C.: RL with temporal logic rewards. In: IROS (2017)
- Liang, E., Liaw, R., Nishihara, R., Moritz, P., Fox, R., Goldberg, K., Gonzalez, J., Jordan, M., Stoica, I.: RLlib: Abstractions for distributed reinforcement learning. In: ICML (2018)
- Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. In: ICLR (2016)
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., Mordatch, I.: Multi-agent actor-critic for mixed cooperative-competitive environments. In: NIPS (2017)
- Mania, H., Guy, A., Recht, B.: Simple random search of static linear policies is competitive for reinforcement learning. In: NeurIPS (2018)
- 13. Neary, C., Xu, Z., Wu, B., Topcu, U.: Reward machines for cooperative multi-agent reinforcement learning. In: AAMAS (2021)
- 14. Ng, A.Y., Harada, D., Russell, S.: Policy invariance under reward transformations: Theory and application to reward shaping. In: ICML (1999)
- Samvelyan, M., Rashid, T., Schroeder de Witt, C., Farquhar, G., Nardelli, N., Rudner, T.G.J., Hung, C.M., Torr, P.H.S., Foerster, J., Whiteson, S.: Starcraft challenge. In: AAMAS (2019)
- 16. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. CoRR (2017)
- Song, J., Ren, H., Sadigh, D., Ermon, S.: Multi-agent generative adversarial imitation learning. In: NeurIPS (2018)
- Toro Icarte, R., Klassen, T.Q., Valenzano, R., McIlraith, S.A.: Using reward machines for high-level task specification and decomposition in reinforcement learning. In: ICML (2018)
- Toro Icarte, R., Klassen, T.Q., Valenzano, R., McIlraith, S.A.: Reward machines: Exploiting reward function structure in reinforcement learning. arXiv preprint arXiv:2010.03950 (2020)
- 20. Vardi, M., Wolper, P.: Reasoning about infinite computations. Information and Comp. (1994)
- Vazquez-Chanlatte, M., Jha, S., Tiwari, A., Ho, M.K., Seshia, S.: Learning task specifications from demonstrations. In: NeurIPS (2018)
- 22. Wang, W., Yang, T., Liu, Y., Hao, J., Hao, X., Hu, Y., Chen, Y., Fan, C., Gao, Y.: From few to more: Large-scale dynamic multiagent curriculum learning. In: AAAI. No. 05 (2020)
- 23. Yu, L., Song, J., Ermon, S.: Multi-agent adversarial inverse RL. In: ICML (2019)

¹⁶ J. Eappen and S. Jagannathan