


Compositional Shielding and Reinforcement Learning for Multi-Agent Systems

Asger Horn Brorholt 
Aalborg University
Aalborg, Denmark
asgerhb@cs.aau.dk

Kim Guldstrand Larsen 
Aalborg University
Aalborg, Denmark
kgl@cs.aau.dk

Christian Schilling 
Aalborg University
Aalborg, Denmark
christianms@cs.aau.dk

ABSTRACT

Deep reinforcement learning has emerged as a powerful tool for obtaining high-performance policies. However, the safety of these policies has been a long-standing issue. One promising paradigm to guarantee safety is a *shield*, which “shields” a policy from making unsafe actions. However, computing a shield scales exponentially in the number of state variables. This is a particular concern in multi-agent systems with many agents. In this work, we propose a novel approach for multi-agent shielding. We address scalability by computing individual shields for each agent. The challenge is that typical safety specifications are global properties, but the shields of individual agents only ensure local properties. Our key to overcome this challenge is to apply assume-guarantee reasoning. Specifically, we present a sound proof rule that decomposes a (global, complex) safety specification into (local, simple) obligations for the shields of the individual agents. Moreover, we show that applying the shields during reinforcement learning significantly improves the quality of the policies obtained for a given training budget. We demonstrate the effectiveness and scalability of our multi-agent shielding framework in two case studies, reducing the computation time from hours to seconds and achieving fast learning convergence.

KEYWORDS

Multi-agent reinforcement learning; Shielding; Safety; Assume-guarantee reasoning

ACM Reference Format:

Asger Horn Brorholt , Kim Guldstrand Larsen , and Christian Schilling . 2025. Compositional Shielding and Reinforcement Learning for Multi-Agent Systems. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Reinforcement learning (RL) [29, 34], and in particular deep RL, has demonstrated success in automatically learning high-performance policies for complex systems [4, 25]. However, learned policies lack guarantees, which prevents applications in safety-critical domains.

An attractive algorithmic paradigm to provably safe RL is *shielding* [3]. In this paradigm, one constructs a *shield*, which is a non-deterministic policy that only allows safe actions. The shield acts

as a guardrail for the RL agent to enforce safety both during learning (of a concrete policy) and operation. This way, one obtains a *safe-by-design shielded policy with high performance*.

Shield synthesis automatically computes a shield from a safety specification and a model of the system, but scales exponentially in the number of state variables. This is a particular concern in multi-agent (MA) systems, which typically consist of many variables. Shielding of MA systems will be our focus in this work.

Existing approaches to MA shielding address scalability by computing individual shields for each agent. Yet, these shields are either not truly safe or not truly independent; rather, they require online communication among all agents, which is often unrealistic.

In this paper, we present the first MA shielding approach that is truly compositional, does not require online communication, and provides absolute safety guarantees. Concretely, we assume that agents observe a subset of all system variables (i.e., operate in a projection of the global state space). We show how to tractably synthesize individual shields in low-dimensional projections. The challenge we need to overcome is that a straightforward generalization of the classical shield synthesis to the MA setting for truly independent shields often fails. The reason is that the projection removes the potential to coordinate between the agents, but often some form of coordination is required.

To address the need for coordination, we get inspiration from *compositional reasoning*, which is a powerful approach, allowing to scale up the analysis of distributed systems. The underlying principle is to construct a correctness proof of multi-component systems by smaller, “local” proofs for each individual component. In particular, *assume-guarantee reasoning* for concurrent programs was popularized in seminal works [5, 21, 26, 28, 33]. By writing $\langle A \rangle C \langle G \rangle$ for “assuming A , component C will guarantee G ,” the standard (acyclic) assume-guarantee rule for finite state machines with handshake synchronization looks as follows [16]:

$$\frac{\langle \top \rangle C_1 \langle G_1 \rangle, \langle G_1 \rangle C_2 \langle G_2 \rangle, \dots, \langle G_{n-2} \rangle C_{n-1} \langle G_{n-1} \rangle, \langle G_{n-1} \rangle C_n \langle \phi \rangle}{\langle \top \rangle C_1 \parallel C_2 \parallel \dots \parallel C_n \langle \phi \rangle}$$

By this chain of assume-guarantee pairs, it is clear that, together, the components ensure safety property ϕ .

In this work, we adapt the above rule to multi-agent shielding. Instead of one shield for the whole system, we synthesize an individual shield for each agent, which together we call a *distributed shield*. Thus, we arrive at n shield synthesis problems (corresponding to the rule’s premise), but each of which is efficient. In our case studies, this reduces the synthesis time from hours to seconds. The *guarantees* G_i allow the individual shields to coordinate on responsibilities at synthesis time. Yet, distributed shields do not require communication when deployed. Altogether, this allows us to *synthesize safe shields* in a compositional and scalable way.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).




Figure 1: Car platoon example for $n = 10$ cars.

The crucial challenge is that, in the classical setting, the components C_i are fixed. In our synthesis setting, the components C_i are our agents, which are *not* fixed at the time of the shield synthesis. In this work, we assume that the guarantees G_i are given, which allows us to derive corresponding individual agent shields via standard shield synthesis.

Motivating Example. A multi-agent car platoon with adaptive cruise controls consists of n cars, numbered from back to front [22] (Figure 1). The cars 1 to $n - 1$ are each controlled by an agent, while (front) car n is driven by the environment. The state variables are the car velocities v_i and distances d_i between cars i and $i + 1$. For $i < n$, car i follows car $i + 1$, observing the variables (v_i, v_{i+1}, d_i) . With a decision period of 1 second, cars act by choosing an acceleration from $\{-2, 0, 2\}$ [m/s²]. Velocities are capped between $[-10, 20]$ m/s.

When two cars have distance 0, they enter an uncontrollable “damaged” state where both cars get to a standstill. The global safety property is to maintain a safe but bounded distance between all cars, i.e., the set of safe states is $\phi = \{s \mid \bigwedge_i 0 < d_i < 200\}$.

As a first attempt, we design the agents’ individual safety properties to only maintain a safe distance to the car in front, i.e., $\phi_i = \{s \mid 0 < d_i < 200\}$ for all i . However, safe agent shields for cars $i > 1$ do not exist for this property: car i cannot prevent a crash with car $i - 1$ (behind), and in the “damaged” (halting) state car i cannot guarantee to avoid crashing with car $i + 1$. Note that making the distance d_{i-1} observable for car i does not help.

To overcome this seemingly impossible situation, we will allow car i to *assume* that the (unobservable) car $i - 1$ *guarantees* to never crash into car i . This guarantee will be provided by the shield of car $i - 1$ and eliminates the critical behavior preventing a local shield for car i . In that way, we iteratively obtain local shields for all agents. Note that this coincides with human driver reasoning.

Beside synthesis of a distributed shield, we also study learning policies for shielded agents. In general, multi-agent reinforcement learning (MARL) [38] is complex due to high-dimensional state and action spaces, which impede convergence to optimal policies.

Here, we identify a class of systems where learning the agents in a *cascading* way is both effective and efficient. Concretely, if we assign an index to each agent, and each agent only depends on agents with lower index, we can learn policies in a sequential order. This leads to a low-dimensional space for the learning algorithm, which leads to fast convergence. While in general suboptimal, we show that this approach still leads to Pareto-optimal results.

In summary, this paper makes the following main contributions:

- We propose *distributed shielding*, the first MA shielding approach with absolute safety guarantees and scalability, yet without online communication. To this end, our approach integrates shield synthesis and assume-guarantee reasoning.
- We propose (shielded) *cascading learning*, a scalable MARL approach for systems with acyclic dependency structure, which further benefits from assume-guarantee reasoning.

- We evaluate our approaches in two case studies. First, we demonstrate that distributed shielding is scalable and, thanks to the integration of assume-guarantee reasoning, applicable. Second, we demonstrate that shielded cascading learning is efficient and achieves state-of-the-art performance.

1.1 Related Work

Shielding. As mentioned, shielding is a technique that computes a *shield*, which prevents an agent from taking unsafe actions. Thus, any policy under a shield is safe, which makes it attractive for safety both during learning and after deployment. Shields are typically based on game-theoretic results, where they are called *winning strategies* [8]. Early applications of shields in learning were proposed for timed systems [14] and discrete systems [3]. The idea has since been extended to probabilistic systems [18, 36], partial observability [12], and continuous-time dynamics [9, 10]. For more background we refer to surveys [19, 20]. In this work, we focus on discrete but multi-agent systems, which we now review in detail.

Multi-agent shielding. An early work on multi-agent enforcement considered a very restricted setting with deterministic environments where the specification is already given in terms of valid actions and not in terms of states [7]. Thus, the shield does not reason about the dynamics and simply overrides forbidden actions.

Model-predictive shielding assumes a backup policy together with a set of recoverable states from which this policy can guarantee safety. Such a backup policy may for instance be implemented by a shield, and is combined with another (typically learned) policy. First, a step with the second policy is simulated and, when the target state is recoverable, this step is executed; otherwise, the fallback policy is executed. Crucially, this assumes that the environment is deterministic. Zhang et al. proposed a multi-agent version [39], where the key insight is that only some agents need to use the backup policy. For scalability, the authors propose a greedy algorithm to identify a sufficiently small subset of agents. However, the “shield” is centralized, which makes this approach not scalable.

Another work computes a safe policy online [31], which may be slow. Agents in close proximity create a communication group, and they communicate their planned trajectories for the next k steps. Each agent has an agreed-on priority in which they have to resolve safety violations, but if that is not possible, agents may disturb higher-priority agents. The approach requires strong assumptions like deterministic system dynamics and immediate communication.

One work suggests to directly reinforcement-learn policies by simply encouraging safety [30]. Here, the loss function encodes a safety proof called *barrier certificate*. But, as with any reward engineering, this approach does not guarantee safety in any way.

Another way to scale up shielding for multi-agent systems is a so-called *factored shield*, which safeguards only a subset of the state space, independent of the number of agents [15]. When an agent moves, it joins or leaves a shield at border states. However, this approach relies on very few agents ever interacting with each other, as otherwise, there is no significant scalability gain.

Factored shields were extended to *dynamic shields* [35]. The idea is that, in order to reduce the communication overhead, an agent’s shield should “merge” dynamically with the shields of other agents

in the proximity. Since the shields are computed with a k -step lookahead only, safety is not guaranteed invariantly.

Multi-agent verification. *Rational verification* proposes to study specifications only from initial states in Nash equilibria, i.e., assuming that all agents act completely rationally [1]. While that assumption may be useful for rational/optimal agents, we typically have learned agents in mind, which do not always act optimally.

The tool *Verse* lets users specify multi-agent scenarios in a Python dialect and provides black-box (simulations) and white-box (formal proofs; our setting) analysis for time-bounded specifications [23].

Assume-guarantee reasoning has been applied to multi-agent systems in [27] and in [24], but not yet to (multi-agent) shielding.

Outline. In the next section, we define basic notation. In Section 3, we introduce distributed shielding based on projections and extend it with assume-guarantee reasoning. In Section 4, we develop cascading learning, tailored to systems with acyclic dependencies. In Section 5, we evaluate our approaches in two case studies. In Section 6, we conclude and discuss future work.

2 PRELIMINARIES

2.1 Transition Systems (MDPs & LTSs)

We start with some basic definitions of transition systems.

Definition 1 (Labeled transition system). A *labeled transition system* (LTS) is a triple $\mathcal{T} = (S, Act, T)$ where S is the finite state space, Act is the action space, and $T \subseteq S \times Act \times S$ is the transition relation with no dead ends, i.e., for all $s \in S$ there exists some $a \in Act$ and $s' \in S$ such that $(s, a, s') \in T$.

Definition 2 (Markov decision process). A *Markov decision process* (MDP) is a triple $\mathcal{M} = (S, Act, P)$ where S is the finite state space, Act is the action space, and $P: S \times Act \times S \rightarrow [0, 1]$ is the probabilistic transition relation satisfying $\sum_{s' \in S} P(s, a, s') \in \{0, 1\}$ for all $s \in S$ and $a \in Act$, and for at least one action, the sum is 1.

We will view an LTS as an abstraction of an MDP where probabilities are replaced by possibilities.

Definition 3 (Induced LTS). Given an MDP $\mathcal{M} = (S, Act, P)$, the *induced LTS* is $\mathcal{T}_{\mathcal{M}} = (S, Act, T)$ with $(s, a, s') \in T$ iff $P(s, a, s') > 0$.

Definition 4 (Run). Assume an LTS $\mathcal{T} = (S, Act, T)$ and a finite alternating sequence of states and actions $\rho = s_0 a_0 s_1 a_1 \dots$; then, ρ is a *run* of \mathcal{T} if $(s_i, a_i, s_{i+1}) \in T$ for all $i \geq 0$. Similarly, for an MDP $\mathcal{M} = (S, Act, P)$, ρ is a *run* of \mathcal{M} if $P(s_i, a_i, s_{i+1}) > 0$ for all $i \geq 0$.

We distinguish between strategies and policies in this work. A strategy prescribes a nondeterministic choice of actions in each LTS state. Similarly, a policy prescribes a probabilistic choice of actions in each MDP state. Before defining them formally, we need a notion of restricting the actions to sensible choices.

Definition 5 (Enabled actions). Given an LTS, $\mathcal{E}(s) = \{a \in Act \mid \exists s' : (s, a, s') \in T\}$ denotes the *enabled actions* in state s . Similarly, given an MDP, $\mathcal{E}(s) = \{a \in Act \mid \exists s' : P(s, a, s') > 0\}$.

Definition 6 (Strategy; policy). Given an LTS, a (nondeterministic) *strategy* is a function $\sigma: S \rightarrow 2^{Act}$ such that $\emptyset \neq \sigma(s) \subseteq \mathcal{E}(s)$ for all $s \in S$. Given an MDP, a (probabilistic) *policy* is a

function $\pi: S \times Act \rightarrow [0, 1]$ such that $\sum_{a \in \mathcal{E}(s)} \pi(s, a) = 1$ and $\bigwedge_{a' \in Act \setminus \mathcal{E}(s)} \pi(s, a') = 0$ for all $s \in S$.

Note that our strategies and policies are memoryless. This is justified as we will only consider safety properties in this work, for which memory is not required [8]. Strategies and policies restrict the possible runs, and we call these runs the outcomes.

Definition 7 (Outcome). A run $\rho = s_0 a_0 s_1 a_1 \dots$ of an LTS is an *outcome* of a strategy σ if $a_i \in \sigma(s_i)$ for all $i \geq 0$. Similarly, a run $\rho = s_0 a_0 s_1 a_1 \dots$ of an MDP is an outcome of a policy π if $\pi(s_i, a_i) > 0$ for all $i \geq 0$.

2.2 Safety and Shielding

In this work, we are interested in safety properties, which are characterized by a set of safe (resp. unsafe) states. The goal is to stay in the safe (resp. avoid the unsafe) states. In this section, we introduce corresponding notions, in particular (classical) shields and how they can be applied.

Definition 8 (Safety property). A *safety property* is a set of states $\phi \subseteq S$.

Definition 9 (Safe run). Given a safety property $\phi \subseteq S$, a run $s_0 a_0 s_1 a_1 \dots$ is *safe* if $s_i \in \phi$ for all $i \geq 0$.

Given an LTS, a safety property $\phi \subseteq S$ partitions the states into two sets: the *winning states*, from which a strategy exists whose outcomes are all safe, and the complement. The latter can be computed as the attractor set of the complement $S \setminus \phi$ [8]. Since it is hopeless to ensure safe behavior from the complement states, in the following we will only be interested in outcomes starting in winning states, which we abstain from mentioning explicitly.

A shield is a (typically nondeterministic) strategy that ensures safety. In game-theory terms, a shield is called a *winning strategy*.

Definition 10 (Shield). Given an LTS (S, Act, T) and a safety property $\phi \subseteq S$, a *shield* $\sqsupset[\phi]$ is a strategy whose outcomes starting in any winning state are all safe wrt. ϕ .

We often omit ϕ and just write \sqsupset . Among all shields, it is known that there is a “best” one that allows the most actions.

Definition 11 (Most permissive shield). Given an LTS and a safety property ϕ , the *most permissive shield* $\sqsupset^*[\phi]$ is the shield that allows the largest set of actions for each state $s \in S$.

Lemma 1 ([8]). \sqsupset^* is unique and obtained as the union of all shields \sqsupset for $\phi: \sqsupset^*(s) = \{a \in Act \mid \exists \sqsupset: a \in \sqsupset(s)\}$.

The standard usage of a shield is to restrict the actions of a policy for guaranteeing safety. In this work, we also compose it with another strategy. For that, we introduce the notion of composition of strategies (recall that a shield is also a strategy). We can, however, only compose strategies that are compatible in the sense that they allow at least one common action in each state (otherwise the result is not a strategy according to our definition).

Definition 12 (Composition). Two strategies σ_1 and σ_2 over an LTS (S, Act, T) are *compatible* if $\sigma_1(s) \cap \sigma_2(s) \neq \emptyset$ for all $s \in S$.

Given compatible strategies σ and σ' , their composition $\sigma \sqcap \sigma'$ is the strategy $(\sigma \sqcap \sigma')(s) = \sigma(s) \cap \sigma'(s)$.

We write $\sqcap_{i < j} \sigma_i$ to denote $\sigma_1 \sqcap \dots \sqcap \sigma_{j-1}$, and $\sqcap_i \sigma_i$ to denote $\sigma_1 \sqcap \dots \sqcap \sigma_n$ when n is clear from the context.

Given a strategy σ and a compatible shield \sqsupset , we also use the alternative notation of the *shielded strategy* $\sqsupset(\sigma) = \sigma \sqcap \sqsupset$.

Given a set of states ϕ , we are interested whether an LTS ensures that we will stay in that set ϕ , independent of the strategy.

Definition 13. Assume an LTS \mathcal{T} and a set of states ϕ . We write $\mathcal{T} \models \phi$ if for all strategies σ , all corresponding outcomes $s_0 a_0 s_1 a_1 \dots$ satisfy $s_i \in \phi$ for all $i \geq 0$.

We now use a different view on a shield and apply it to an LTS in order to “filter out” those actions that are forbidden by the shield.

Definition 14 (Shielded LTS). Given an LTS $\mathcal{T} = (S, Act, T)$, a safety property ϕ , and a shield $\sqsupset[\phi]$, the *shielded LTS* $\mathcal{T}_{\sqsupset} = (S, Act, T_{\sqsupset})$ with $T_{\sqsupset} = \{(s, a, s') \in T \mid a \in \sqsupset(s)\}$ is restricted to transitions whose actions are allowed by the shield.

The next proposition asserts that a shielded LTS is safe.

Proposition 1. Given an LTS \mathcal{T} , a safety property ϕ , and a corresponding shield $\sqsupset[\phi]$, all outcomes of any strategy for \mathcal{T}_{\sqsupset} are safe.

In other words, $\mathcal{T}_{\sqsupset} \models \phi$. We analogously define shielded MDPs.

Definition 15 (Shielded MDP). Given an MDP $\mathcal{M} = (S, Act, P)$, a safety property ϕ , and a shield \sqsupset for $\mathcal{T}_{\mathcal{M}}$, the *shielded MDP* $\mathcal{M}_{\sqsupset} = (S, Act, P_{\sqsupset})$ is restricted to transitions with actions allowed by \sqsupset : $P_{\sqsupset}(s, a, s') = P(s, a, s')$ if $a \in \sqsupset(s)$, and $P_{\sqsupset}(s, a, s') = 0$ otherwise.

Proposition 2. Assume an MDP \mathcal{M} , a safety property ϕ , and a corresponding shield $\sqsupset[\phi]$ for $\mathcal{T}_{\mathcal{M}}$. Then all outcomes of any policy for \mathcal{M}_{\sqsupset} are safe.

The last proposition explains how standard shielding is applied to learn safe policies. Given an MDP \mathcal{M} , we first compute a shield \sqsupset over the induced LTS $\mathcal{T}_{\mathcal{M}}$. Then we apply the shield to the MDP \mathcal{M} to obtain \mathcal{M}_{\sqsupset} and filter unsafe actions. The shield guarantees that the agent is safe both during and after learning.

From now on we mainly focus on computing shields from an LTS, as the generalization to MDPs is straightforward.

2.3 Compositional Systems

Now we turn to compositional systems (LTSs and MDPs) with multiple agents. We restrict ourselves to k -dimensional state spaces S , i.e., products of variables $S = \times_i S_i$. We allow for sharing some of these variables among the agents by projecting to observation subspaces. The following is the standard definition of projecting out certain variables while retaining others. We use the notation that, given an n -vector $v = (v_1, \dots, v_n)$, $v[i]$ denotes the i -th element v_i .

Definition 16 (Projection). A *projection* is a mapping $prj: S \rightarrow O$ that maps k -dimensional vectors $s \in S$ to j -dimensional vectors $o \in O$, where $j \leq k$. Formally, prj is associated with a sequence of j indices $1 \leq i_1 < \dots < i_j \leq k$ such that $prj(s) = (s[i_1], \dots, s[i_j])$. Additionally, we define $prj(\phi) = \bigcup_{s \in \phi} \{prj(s)\}$.

Definition 17 (Extension). Given projection $prj: S \rightarrow O$, the set of states projected to o is the *extension* $\uparrow(o) = \{s \in S \mid prj(s) = o\}$.

Later we will also use an alternative projection, which we call *restricted*. The motivation is that the standard projection above sometimes retains too many states. The restricted projection instead only keeps those states such that the extension of the projection

($\uparrow(\cdot)$) is contained in the original set. For instance, for the state space $S = \{0, 1\}^2$, the set of states $\phi = \{(0, 0), (0, 1), (1, 0)\}$, and the one-dimensional projection $prj(s) = s[1]$, we have that $prj(\phi) = \{0, 1\}$. The restricted projection removes 1 as $(1, 1) \notin \phi$.

Definition 18 (Restricted projection). A *restricted projection* is a mapping $\overline{prj}: 2^S \rightarrow 2^O$ that maps sets of k -dimensional vectors $s \in S$ to sets of j -dimensional vectors $o \in O$, where $j \leq k$. Formally, \overline{prj} is associated with a sequence of j indices $1 \leq i_1 < \dots < i_j \leq k$. Let prj be the corresponding (standard) projection and $\phi \subseteq S$. Then $\overline{prj}(\phi) = \{o \in O \mid \{s \in S \mid prj(s) = o\} \subseteq \phi\}$. Again, we define $\overline{prj}(\phi) = \bigcup_{s \in \phi} \{\overline{prj}(s)\}$.

We will apply \overline{prj} only to safety properties ϕ . The following alternative characterization may help with the intuition: $\overline{prj}(\phi) = prj(\overline{\phi}) = O \setminus prj(S \setminus \phi)$, where $\overline{\phi}$ denotes the complement $S \setminus \phi$ (resp. $O \setminus \phi$) of a set of states $\phi \subseteq S$ (resp. observations $\phi \subseteq O$).

Crucially, prj and \overline{prj} coincide if $\uparrow(prj(\phi)) = \phi$, i.e., if the projection of ϕ preserves correlations. We will later turn our attention to agent safety properties, where this is commonly the case.

Now we can define a multi-agent LTS and MDP.

Definition 19 (n -agent LTS/MDP). An n -agent LTS (S, Act, T) or an n -agent MDP (S, Act, P) have an n -dimensional action space $Act = Act_1 \times \dots \times Act_n$ and a family of n projections $prj_i, i = 1, \dots, n$. Each agent i is associated with the projection $prj_i: S \rightarrow O_i$ from S to its observation space O_i .

We note that the observation space introduces partial observability. Obtaining optimal strategies/policies for partial observability is difficult and generally requires infinite memory [13]. Since this is impractical, we restrict ourselves to memoryless strategies/policies.

We can apply the projection function prj to obtain a “local” LTS, modeling partial observability.

Definition 20 (Projected LTS). For an n -agent LTS $\mathcal{T} = (S, Act, T)$ and an agent i with projection function $prj_i: S \rightarrow O_i$, the *projected LTS to agent i* is $\mathcal{T}^i = (O_i, Act_i, T_i)$ where $Act_i = \{a[i] \mid a \in Act\}$ and $T_i = \{(prj_i(s), a[i], prj_i(s')) \mid (s, a, s') \in T\}$.

3 DISTRIBUTED SHIELD SYNTHESIS

We now turn to shielding in a multi-agent setting. The straightforward approach is to consider the full-dimensional system and compute a global shield. This has, however, two issues. First, a global shield assumes communication among the agents, which we generally do not want to assume. Second, and more importantly, shield computation scales exponentially in the number of variables.

To address these issues, we instead compute *local* shields, one for each agent. A local shield still keeps its agent safe. But since we only consider the agent’s observation space, the shield does not require communication, and the computation is much cheaper.

3.1 Projection-Based Shield Synthesis

Rather than enforcing the global safety property, local shields will enforce agent-specific properties, which we characterize next.

Definition 21 (n -agent safety property). Given an n -agent LTS or MDP with state space S , a safety property $\phi \subseteq S$ is an n -agent *safety property* if $\phi = \bigcap_{i=1}^n \phi_i$ consists of *agent safety properties* ϕ_i for each agent i .

Note that we can let $\phi_i = \phi$ for all i , so this is not a restriction. But typically we are interested in properties that can be accurately assessed in the agents' observation space (i.e., $\text{prj}_i(\phi_i) = \overline{\text{prj}}_i(\phi_i)$).

Next, we define a local shield of an agent, which, like the agent, operates in the observation space.

Definition 22 (Local shield). Given an n -agent LTS $\mathcal{T} = (S, \text{Act}, T)$ with observation spaces O_i and an n -agent safety property $\phi = \bigcap_{i=1}^n \phi_i \subseteq S$, let $\sqsupset_i: O_i \rightarrow 2^{\text{Act}_i}$ be a shield for \mathcal{T}^i wrt. $\text{prj}_i(\phi_i)$, for some agent $i \in \{1, \dots, n\}$, i.e., $\mathcal{T}^i \models_{\sqsupset_i} \overline{\text{prj}}_i(\phi_i)$. We call \sqsupset_i a *local shield* of agent i .

We define an operation to turn a j -dimensional (local) shield into a k -dimensional (global) shield. This global shield allows all global actions whose projections are allowed by the local shield.

Definition 23 (Extended shield). Assume an n -agent LTS $\mathcal{T} = (S, \text{Act}, T)$ with projections prj_i , an n -agent safety property $\phi = \bigcap_{i=1}^n \phi_i \subseteq S$, and a corresponding local shield \sqsupset_i . The *extended shield* $\uparrow(\sqsupset_i)$ is defined as $\uparrow(\sqsupset_i)(s) = \{a \in \text{Act} \mid a[i] \in \sqsupset_i(\text{prj}_i(s))\}$.

The following definition is just syntactic sugar to ease reading.

Definition 24. Assume an LTS \mathcal{T} , a set of states ϕ , and a shield \sqsupset for ϕ . We write $\mathcal{T} \models_{\sqsupset} \phi$ as an alternative to $\mathcal{T}_{\sqsupset} \models \phi$.

The following lemma says that it is sufficient to have a local shield ensuring the *restricted* projection $\overline{\text{prj}}_i(\phi_i)$ of an agent safety property ϕ_i in order to guarantee safety of the extended shield.

Lemma 2. Assume an n -agent LTS \mathcal{T} , a safety property ϕ_i , and a local shield \sqsupset_i such that $\mathcal{T}^i \models_{\sqsupset_i} \overline{\text{prj}}_i(\phi_i)$. Then $\mathcal{T} \models_{\uparrow(\sqsupset_i)} \phi_i$.

PROOF. The proof is by contraposition. Assume that there is an unsafe outcome ρ in \mathcal{T} (starting in a winning state) under the extended shield $\uparrow(\sqsupset_i)$, i.e., ρ contains a state $s \notin \phi$. Then the projected run $\text{prj}_i(s_0) a[i] \text{prj}_i(s_1) \dots$ is an outcome of \mathcal{T}^i under local shield \sqsupset_i , and $\text{prj}_i(s) \notin \overline{\text{prj}}_i(\phi_i)$ by the definition of $\overline{\text{prj}}$. This contradicts that \sqsupset_i is a local shield. \square

The following example shows that the *restricted* projection is necessary. Consider the LTS \mathcal{T} where $S = \{0, 1\}^2$, $\text{Act} = \{z, p\}^2$, and $T = \{((0, 0), (z, z), (0, 0)), ((0, 0), (z, p), (0, 1)), ((0, 0), (p, z), (1, 0)), ((0, 0), (p, p), (1, 1))\}$. For $i = 1, 2$ let $\phi_i = \{(0, 0), (0, 1), (1, 0)\}$ and prj_i project to the i -th component O_i . Then $\text{prj}_i(\phi_i) = \{0, 1\} = \text{prj}_i(S)$, i.e., all states in the projection are safe, and hence a local shield may allow $\sqsupset_i(0) = \{z, p\}$. But then the unsafe state $(1, 1)$ would be reachable in \mathcal{T} .

If $\sqsupset = \sqcap_i \uparrow(\sqsupset_i)$ exists, we call it a *distributed shield*. This terminology is justified in the next theorem, which says that we can synthesize n local shields in the projections and then combine these local shields to obtain a safe shield for the global system.

Theorem 1 (Projection-based shield synthesis). Assume an n -agent LTS $\mathcal{T} = (S, \text{Act}, T)$ and an n -agent safety property $\phi = \bigcap_{i=1}^n \phi_i \subseteq S$. Moreover, assume local shields \sqsupset_i for all $i = 1, \dots, n$. If $\sqsupset = \sqcap_i \uparrow(\sqsupset_i)$ exists, then \sqsupset is a shield for \mathcal{T} wrt. ϕ (i.e., $\mathcal{T}_{\sqsupset} \models \phi$).

PROOF. By definition, each local shield \sqsupset_i ensures that the (*restricted* projected) agent safety property ϕ_i holds in \mathcal{T}^i . Since \mathcal{T}^i is a projection of \mathcal{T} , any distributed shield with i -th component \sqsupset_i also preserves ϕ_i in \mathcal{T} (by Lemma 2). Hence, $\sqsupset = \sqcap_i \uparrow(\sqsupset_i)$ ensures all agent safety properties ϕ_i and thus $\phi = \bigcap_{i=1}^n \phi_i$. \square

Unfortunately, the theorem is often not useful in practice because the local shields may not exist. The projection generally removes the possibility to coordinate with other agents. By *coordination* we do not mean (online) communication but simply (offline) agreement on “who does what.” Often, this coordination is necessary to achieve agent safety. We address this lack of coordination in the next section.

3.2 Assume-Guarantee Shield Synthesis

Shielding an LTS removes some transitions. Thus, by repeatedly applying multiple shields to the same LTS, we obtain a sequence of more and more restricted LTSs.

Definition 25 (Restricted LTS). Assume two LTSs $\mathcal{T} = (S, \text{Act}, T)$, $\mathcal{T}' = (S, \text{Act}, T')$. We write $\mathcal{T} \leq \mathcal{T}'$ if $T \subseteq T'$.

Lemma 3. Let $\mathcal{T} \leq \mathcal{T}'$ be two LTSs. Then $\mathcal{T}' \models \phi \implies \mathcal{T} \models \phi$.

PROOF. As T' contains all transitions of T , it has at least the same outcomes. If no outcome of \mathcal{T}' leaves ϕ , the same holds for \mathcal{T} . \square

We now turn to the main contribution of this section. For a safety property ϕ' , we assume an n -agent safety property $\phi = \bigcap_{i=1}^n \phi_i$ is given such that $\phi \subseteq \phi'$ (i.e., ϕ is more restrictive). We use these agent safety properties ϕ_i to filter out behavior during shield synthesis. They may contain additional guarantees, which are used to coordinate responsibilities between agents.

Crucially, in our work, the guarantees are given in a certain order. We assume wlog that the agent indices are ordered from 1 to n such that agent i can only rely on the safety properties of all agents $j < i$. Thus, agent i guarantees ϕ_i by assuming $\bigcap_{j < i} \phi_j$. This is important to avoid problems with (generally unsound) circular reasoning. In particular, agent 1 cannot rely on anything, and ϕ_n is not relied on.

The theorem then states that if each agent guarantees its safety property ϕ_i , and only relies on guarantees ϕ_j such that $j < i$. The result is a (safe) distributed shield. The described condition is formally expressed as $\left(\mathcal{T}_{\sqsupset^*[\bigcap_{j < i} \phi_j]}\right)^i \models_{\sqsupset_i} \overline{\text{prj}}_i(\phi_i)$, where we use the most permissive shield \sqsupset^* for unicity.

Theorem 2 (Assume-guarantee shield synthesis). Assume an n -agent LTS $\mathcal{T} = (S, \text{Act}, T)$ with projections prj_i and an n -agent safety property $\phi = \bigcap_i \phi_i$. Moreover, assume (local) shields \sqsupset_i for all i such that $\left(\mathcal{T}_{\sqsupset^*[\bigcap_{j < i} \phi_j]}\right)^i \models_{\sqsupset_i} \overline{\text{prj}}_i(\phi_i)$. Then, if $\sqsupset = \sqcap_i \uparrow(\sqsupset_i)$ exists, it is a shield for \mathcal{T} wrt. ϕ (i.e., $\mathcal{T}_{\sqsupset} \models \phi$).

PROOF. Assume \mathcal{T} , ϕ , and local shields \sqsupset_i as in the assumptions. Observe that for $i = 1$, $\bigcap_{j < i} \phi_j = S$, and that $\mathcal{T}_{\sqsupset^*[S]} = \mathcal{T}$. Then:

$$\begin{aligned} & \bigwedge_i \left(\mathcal{T}_{\sqsupset^*[\bigcap_{j < i} \phi_j]} \right)^i \models_{\sqsupset_i} \overline{\text{prj}}_i(\phi_i) \\ & \xRightarrow{\text{Lem. 2}} \bigwedge_i \mathcal{T}_{\sqsupset^*[\bigcap_{j < i} \phi_j]} \models_{\uparrow(\sqsupset_i)} \phi_i \xRightarrow{(*)} \bigwedge_i \mathcal{T}_{\sqsupset_{j < i} \uparrow(\sqsupset_j)} \models_{\uparrow(\sqsupset_i)} \phi_i \\ & \xRightarrow{\text{Def. 24}} \bigwedge_i \mathcal{T} \models_{\sqsupset_{j \leq i} \uparrow(\sqsupset_j)} \phi_i \implies \mathcal{T} \models_{\sqcap_i \uparrow(\sqsupset_i)} \phi \xRightarrow{\text{Def. 24}} \mathcal{T}_{\sqsupset} \models \phi \end{aligned}$$

Step $(*)$ holds because the composition $\sqcap_{j \leq i} \uparrow(\sqsupset_j)$ of the local shields up to index i satisfy ϕ_i under the previous guarantees ϕ_j , $j < i$. Thus, $\mathcal{T}_{\sqcap_{j < i} \uparrow(\sqsupset_j)} \leq \mathcal{T}_{\sqsupset^*[\bigcap_{j < i} \phi_j]}$, and the conclusion follows by applying Lemma 3. \square

Finding the local safety properties ϕ_i is an art, and we leave algorithmic synthesis of these properties to future work. But we will show in our case studies that natural choices often exist, sometimes directly obtained from the (global) safety property.

4 CASCADING LEARNING

In the previous section, we have seen how to efficiently compute a distributed shield based on assume-guarantee reasoning. In this section, we turn to the question how and under which condition we can efficiently learn multi-agent policies in a similar manner.

We start by defining the multi-agent learning objective.

Definition 26 (*n-agent cost function*). Given an *n-agent* MDP $\mathcal{M} = (S, Act, P)$ with projections $prj_i: S \rightarrow O_i$, an *n-agent cost function* $c = (c_1, \dots, c_n)$ consists of (local) cost functions $c_i: O_i \times Act_i \rightarrow \mathbb{R}$. The total immediate cost $c: S \times Act \rightarrow \mathbb{R}$ is $c(s, a) = \sum_{i=1}^n c_i(prj_i(s), a[i])$ for $s \in S$ and $a \in Act$.

An agent policy is obtained by projection, analogous to a local shield. Next, we define the notion of instantiating an *n-agent* MDP with a policy, yielding an $(n-1)$ -agent MDP.

Definition 27 (*Instantiating an agent*). Given an *n-agent* MDP $\mathcal{M} = (S, Act, P)$ and agent policy $\pi: O_i \times Act_i \rightarrow [0, 1]$, the *instantiated MDP* is $\mathcal{M}_\pi = (S, Act', P')$, where $Act' = Act_1 \times \dots \times Act_{i-1} \times Act_{i+1} \times \dots \times Act_n$ and, for all $s, s' \in S$ and $a' \in Act'$, $P'(s, a', s') = \sum_{a_i} \pi(prj_i(s), a_i) \cdot P(s, (a'[1], \dots, a'[i-1], a_i, a'[i], \dots, a'[n-1]), s')$.

We will need the concept of a projected, local run of an agent.

Definition 28 (*Local run*). Given a run $\rho = s_0 a_0 s_1 a_1 \dots$ over an *n-agent* MDP (S, Act, P) , the projection to agent *i* is the *local run* $prj_i(\rho) = prj_i(s_0) a_0[i] prj_i(s_1) a_1[i] \dots$.

Given a policy $\pi: S \times Act \rightarrow [0, 1]$, the probability of a finite local run $prj_i(\rho)$ being an outcome of π is the sum of the probabilities of outcomes of π whose projection to *i* is $prj_i(\rho)$.

The probability of a run ρ of length ℓ being an outcome of policy π is $Pr(\rho \mid \pi) = \prod_{i=0}^{\ell-1} \pi(s_i, a_i) \cdot P(s_i, a_i, s_{i+1})$. We say that agent *i* depends on agent *j* if agent *j*'s action choice influences the probability for agent *i* to observe a (local) run.

Definition 29 (*Dependency*). Given an *n-agent* MDP (S, Act, P) , agent *i* depends on agent *j* if there exists a local run $prj_i(\rho)$ of length ℓ and *n-agent* policies π, π' that differ only in the *j*-th agent policy, i.e., $\pi = (\pi_1, \dots, \pi_n)$ and $\pi' = (\pi_1, \dots, \pi_{j-1}, \pi'_j, \pi_{j+1}, \dots, \pi_n)$, such that the probability of observing $prj_i(\rho)$ under π and π' differ:

$$\sum_{\rho': prj_i(\rho')=prj_i(\rho)} Pr(\rho' \mid \pi) \neq \sum_{\rho': prj_i(\rho')=prj_i(\rho)} Pr(\rho' \mid \pi')$$

where we sum over all runs ρ' of length ℓ with the same projection.

In practice, we can typically perform an equivalent syntactic check. Next, we show how to arrange dependencies in a graph.

Definition 30 (*Dependency graph*). The *dependency graph* of an *n-agent* MDP is a directed graph (V, E) where $V = \{1, \dots, n\}$ and $E = \{(i, j) \mid i \text{ depends on } j\}$.

As the main contribution of this section, Algorithm 1 shows an efficient multi-agent learning framework, which we call *cascading learning*. In order to apply the algorithm, we require an acyclic dependency graph (otherwise, an error is thrown in line 5). Then,

Algorithm 1: Cascading shielded learning of *n-agent* policies

Input : Shielded *n-agent* MDP \mathcal{M}_\square ,
n-agent cost function $c = (c_1, \dots, c_n)$
Output: *n-agent* policy (π_1, \dots, π_n)

- 1 Build dependency graph G of \mathcal{M}_\square ;
- 2 Let $\mathcal{M}' := \mathcal{M}_\square$;
- 3 **while** true **do**
- 4 **if** there is no node in G with no outgoing edges **then**
- 5 error("Cyclic dependencies are incompatible.");
- 6 Let i be a node in G with no outgoing edges;
- 7 Train agent policy π_i on the MDP $sandbox(\mathcal{M}', i)$ wrt. cost function c_i ;
- 8 Update G by removing node i and all incoming edges;
- 9 **if** G is empty **then return** (π_1, \dots, π_n) ;
- 10 Update $\mathcal{M}' := \mathcal{M}'_{\pi_i}$; // (i.e., instantiated shielded MDP)

we train the agents in the order suggested by the dependencies, which, as we will see, leads to an attractive property.

To draw the connection to the distributed shield, the crucial insight is that we can again use it for assume-guarantee reasoning to prevent behaviors that may otherwise create a dependency.

The procedure $sandbox(\mathcal{M}, i)$ in line 7 takes an *n-agent* MDP \mathcal{M} and an agent index $i \in \{1, \dots, n\}$. The purpose is to instantiate every agent except agent *i*. Since agent *i* does not depend on these agents, we arbitrarily choose a uniform policy for the instantiation.

Next, we show an important property of Algorithm 1: it trains policies in-distribution.

Definition 31 (*In-distribution*). Given two 1-agent MDPs $\mathcal{M} = (S, Act, P)$ and $\mathcal{M}' = (S, Act, P')$, an agent policy π is *in-distribution* if the probability of any local run in \mathcal{M} is the same as in \mathcal{M}' .

Now we show that the distribution of observations an agent policy π_i makes during training in Algorithm 1 is identical with the distribution of observations made in \mathcal{M}^* , the instantiation with *all other* agent policies computed by Algorithm 1.

Theorem 3. Let \mathcal{M} be an *n-agent* MDP with acyclic dependency graph. For every agent *i*, the following holds. Let \mathcal{M}^* be the 1-agent MDP obtained by iteratively instantiating the original MDP \mathcal{M} with policies π_j for all $j \neq i$. The agent policy π_i trained with Algorithm 1 is in-distribution wrt. $sandbox(\mathcal{M}', i)$ (from line 7) and \mathcal{M}^* .

PROOF. Fix a policy π_i . If π_i is the last trained policy, the statement clearly holds. Otherwise, let $\pi_j \neq \pi_i$ be a policy that has not been trained at the time when π_i is trained. The algorithm asserts that π_i has no dependency on π_j . Thus, training π_i yields the same policy no matter how π_j behaves. \square

Note that, despite trained in-distribution, the policies are not globally optimal. This is because each policy acts egoistically and optimizes its local cost, which may yield suboptimal global cost.

What we can show is that the agent policies (π_1, \dots, π_n) are *Pareto optimal* [2], i.e., they cannot all be strictly improved without raising the cost of at least one agent. That is, there is no policy π_i

that can be replaced by another policy π'_i without strictly increasing the expected local cost of at least one agent. Indeed:

Theorem 4. *If the learning method in line 7 of Algorithm 1 converged to the (local) optima, and these optima are unique, then the resulting policies are Pareto optimal.*

PROOF. The proof is by induction. Assume *wlog* that the policies are trained in the order 1 to n . By assumption, π_1 is locally optimal and unique. Hence, replacing π_1 by another policy would strictly increase its total cost. Now assume we have shown the claim for the first $i - 1$ agents. Algorithm 1 trained policy π_i wrt. the instantiation with the policies π_1, \dots, π_{i-1} , and by assumption, π_i is also locally optimal and unique. Thus, again, we cannot replace π_i . \square

5 EVALUATION

We consider two environments with discretized state spaces.¹ All experiments were repeated 10 times; solid lines in plots represent the mean cost of these 10 repetitions, while ribbons mark the minimum and maximum costs. Costs are evaluated as the mean of 1,000 episodes. We use the learning method implemented in UPPAAL STRATEGO [17] because the implementation has a native interface for shields. This method learns a policy by partition refinement of the state space. With this learning method, only few episodes are needed for convergence. We also compare to the (deep) MARL approach MAPPO [37] later.

5.1 Car Platoon with Adaptive Cruise Controls

Recall the car platoon model from Section 1. The front car follows a random distribution depending on v_n (described in [11, appendix]).

The individual cost of an agent is the sum of the observed distances to the car immediately in front of it, during a 100-second episode (i.e., keeping a smaller distance to the car in front is better).

The decision period causes delayed reaction time, and so the minimum safe distance to the car in front depends on the velocity of both cars. An agent must learn to drive up to this distance, and then maintain it by predicting the acceleration of the car in front.

For this model, all agents share analogous observations O_i and safety properties ϕ_i . Hence, instead of computing $n - 1$ local shields individually, it is sufficient to compute only one local shield and reuse it across all agents (by simply adapting the variables).

5.1.1 Relative scalability of centralized and distributed shielding. We compare the synthesis of distributed and (non-distributed) classical shields. We call the latter *centralized* shields, as they reason about the global state. Hence, they may permit more behavior and potentially lead to better policies, as the agents can coordinate to take jointly safe actions. Beside this (often unrealistic) coordination assumption, a centralized shield suffers from scalability issues. While the size of a single agent’s observation space is modest, the global state space is often too large for computing a shield.

We interrupted the synthesis of a centralized shield with $n = 3$ cars (i.e., 2 agents) and a full state space after 12 hours, at which point the computation showed less than 3% progress. In order to obtain a centralized shield, we reduced the maximum safe distance




Figure 2: Graph of learning outcomes. Comparison of different learning methods on the 10-car platoon. The centralized and the MAPPO policy were trained for the total episodes indicated, while these episodes were split evenly between each agent in the cascading case.

from 200 to just 50, shrinking the state space significantly. Synthesizing a centralized shield took 78 minutes for this property, compared to just 3 seconds for a corresponding distributed shield.

Because of the exponential complexity to synthesize a centralized shield, we will only consider distributed shields in the following. Synthesizing a shield for a single agent covering the full safety property ($0 < d_i < 200$) took 6.5 seconds, which we will apply to a platoon of 10 cars, well out of reach of a centralized shield.

5.1.2 Comparing centralized, cascading and MAPPO learning. Given a distributed shield, we consider the learning outcomes for a platoon of 10 cars (9 agents), using the learning method of UPPAAL STRATEGO. We train both a shielded *centralized* policy, which picks a joint action for all cars, and individual shielded policies using cascading learning (Algorithm 1). As expected from shielded policies, no safety violations were observed while evaluating them.

In the results shown in Figure 2, the centralized policy does not improve with more training. While it could theoretically outperform distributed policies through communication, the high dimensionality of the state and action space likely prevents that. It only marginally improves over the random baseline, which has an average cost of 71,871. On the other hand, cascading learning quickly converges to a much better cost as low as 26,435.

To examine how cascading learning under a distributed shield compares to traditional MARL techniques, we implemented the platoon environment in the benchmark suite BenchMARL [6] and trained an unshielded policy with MAPPO [37], a state-of-the-art MARL algorithm based on PPO [32], using default hyperparameters. To encourage safe behavior, we added a penalty of 1,600 to the cost function for every step upon reaching an unsafe state. (This value was obtained by starting from 100 and doubling it until safety started degrading again.) Recall that shielded agents are safe.

We include the training outcomes for MAPPO in Figure 2. Due primarily to the penalty of safety violations, the agents often have a cost greater than 100,000, even at the end of training. However, the best MAPPO policy achieved a cost of just 16,854, better than the cascading learning method. We inspected that policy and found that the cars drive very closely, accepting the risk of a crash. Overall, there is a large variance of the MAPPO policies in different runs, whereas cascading learning converges to very similar policies, and does so much faster. This is likely because of the smaller space in

¹Available online at <https://github.com/AsgerHB/N-player-shield>.




Figure 3: Percentage of safe runs with the MAPPO policy in the 10-car platoon. Blue bars show the mean of 10 repetitions, while black intervals give min and max values.

which the policies are learned, due to the distributed shield. Thus, cascading learning is more effective.

Since the MAPPO policy is not safe by construction, Figure 3 shows the percentage of safe episodes, out of 1,000 episodes. The agents tend to be safer with more training, but there is no inherent guarantee of safety, and a significant amount of violations remain.

5.2 Chemical Production Plant

In the second case study, we demonstrate that distributed shielding applies to complex dependencies where agents influence multiple other agents asymmetrically. We consider a network of interconnected chemical production units, each with an internal storage.

Figure 4 shows the graph structure of the network. Numbered nodes (1 to 10) denote controlled production units, while letter-labeled nodes (A, B) denote uncontrolled consumers with periodically varying demand. Arrows from source to target nodes denote potential flow at no incurred cost. Arrows without a source node denote potential flow from external providers, at a cost that individually and periodically varies. Consumption patterns and examples of the cost patterns are shown in [11, appendix]. The flow rate in all arrows follows a uniform random distribution in the range $[2.15, 3.15]$ ℓ/s .




Figure 4: Layout of plant network.

Each agent i is associated with a production unit (1 to 10), with internal storage volume v_i . Beside a global periodic timer, each agent can only observe its own volume. At each decision period of 0.5 seconds, an agent can open or close each of the three input flows (i.e., there are $|Act_i| = 9$ actions per agent and hence $|Act| = 9^{10}$ global actions), but cannot prevent flow from outgoing connections.

The individual cost of an agent is incurred by buying from external providers. Agents must learn to take free material from other units, except for agents 1 to 3, which instead must learn to buy from their external providers periodically when the cost is low.

Units must not exceed their storage capacity, and units 9 to 10 must also not run empty to ensure the consumers' demand is met. That is, the safety property is $\phi = \{s \mid \bigwedge_i v_i < 50 \wedge 0 < v_9 \wedge 0 < v_{10}\}$.

5.2.1 Shielding. The property $0 < v_9$ cannot be enforced by a local shield for agent 9 without additional assumptions that the other agents do not run out. This is because the (single) external provider is not enough to meet the potential (dual) demand of consumer A.




Figure 5: Comparison of different learning methods on the chemical production plant. The centralized policy was trained for the total episodes indicated, while these episodes were split evenly between each agent in the cascading case.

This yields the local safety properties $\phi_i = \{s \mid 0 < v_i < 50\}$. Here, agents 1 to 3 do not make assumptions, while agents 4 and 5 depend on agents 1 to 3 not running out, etc. For this model, we do not use the same shield for all agents, since they differ in the number of outgoing flows (either 1 or 2). Still, it is sufficient to compute two types of shields, one for each variant, and adapt them to analogous agents. Computing a centralized shield would again be infeasible, while computing the distributed shield took less than 1 second.

5.2.2 Comparing centralized and cascading learning. Thanks to the guarantees given by the distributed shields, agents 9 to 10 are only affected by the behavior of the consumers, agents 6 to 8 only depend on agents 9 to 10, etc. Thus, the agent training order is 10, 9, 8...

We compare the results of shielded cascading learning, shielded centralized learning, MAPPO, and shielded random agents in Figure 5. Centralized learning achieved a cost of 292. The lowest cost overall, 172, was achieved by cascading learning. We compare this to the (unshielded) MAPPO agents, whose lowest cost was 291. More background information is given in [11, appendix].

6 CONCLUSION

In this paper, we presented distributed shielding as a scalable MA approach, which we made practically applicable by integrating assume-guarantee reasoning. We also presented cascading shielded learning, which, when applicable, is a scalable MARL approach. We demonstrated that distributed shield synthesis is highly scalable and that coming up with useful guarantees is reasonably simple.

While we focused on demonstrating the feasibility in this work by providing the guarantees manually, a natural future direction is to learn them. As discussed, this is much simpler in the classical setting [16] because the agents/components are fixed. We believe that in our setting where both the guarantees and the agents are not given, a trial-and-error approach (e.g., a genetic algorithm) is a fruitful direction to explore. Another relevant future direction is to generalize our approach to continuous systems [10].

ACKNOWLEDGMENTS

This research was partly supported by the Independent Research Fund Denmark under reference number 10.46540/3120-00041B, DI-REC - Digital Research Centre Denmark under reference number 9142-0001B, and the Villum Investigator Grant S4OS under reference number 37819.

REFERENCES

- [1] Alessandro Abate, Julian Gutierrez, Lewis Hammond, Paul Harrenstein, Marta Kwiatkowska, Muhammad Najib, Giuseppe Perelli, Thomas Steeples, and Michael J. Wooldridge. 2021. Rational verification: game-theoretic verification of multi-agent systems. *Appl. Intell.* 51, 9 (2021), 6569–6584. <https://doi.org/10.1007/S10489-021-02658-Y>
- [2] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marll-book.com>
- [3] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. 2018. Safe Reinforcement Learning via Shielding. In *AAAI*. AAAI Press, 2669–2678. <https://doi.org/10.1609/AAAI.V32I1.11797>
- [4] Marc G. Bellemare, Salvatore Candido, Pablo Samuel Castro, Jun Gong, Marcos C. Machado, Subhodeep Moitra, Sameera S. Ponda, and Ziyu Wang. 2020. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nat.* 588, 7836 (2020), 77–82. <https://doi.org/10.1038/S41586-020-2939-8>
- [5] Albert Benveniste, Benoît Caillaud, Dejan Nickovic, Roberto Passerone, Jean-Baptiste Raclet, Philipp Reinkemeier, Alberto L. Sangiovanni-Vincentelli, Werner Damm, Thomas A. Henzinger, and Kim G. Larsen. 2018. Contracts for System Design. *Found. Trends Electron. Des. Autom.* 12, 2-3 (2018), 124–400. <https://doi.org/10.1561/10000000053>
- [6] Matteo Bettini, Amanda Prorok, and Vincent Moens. 2023. BenchMARL: Benchmarking Multi-Agent Reinforcement Learning. *CoRR* (2023). arXiv:2312.01472
- [7] Suda Bharadwaj, Roderick Bloem, Rayna Dimitrova, Bettina Könighofer, and Ufuk Topcu. 2019. Synthesis of Minimum-Cost Shields for Multi-agent Systems. In *ACC*. IEEE, 1048–1055. <https://doi.org/10.23919/ACC.2019.8815233>
- [8] Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. 2018. Graph Games and Reactive Synthesis. In *Handbook of Model Checking*. Springer, 921–962. https://doi.org/10.1007/978-3-319-10575-8_27
- [9] Asger Horn Brorholt, Andreas Holck Hoeg-Petersen, Kim Guldstrand Larsen, and Christian Schilling. 2024. Efficient Shield Synthesis via State-Space Transformation. In *AISoLA (LNCS, Vol. 15217)*. Springer, 206–224. https://doi.org/10.1007/978-3-031-75434-0_14
- [10] Asger Horn Brorholt, Peter Gjol Jensen, Kim Guldstrand Larsen, Florian Lorber, and Christian Schilling. 2023. Shielded Reinforcement Learning for Hybrid Systems. In *AISoLA (LNCS, Vol. 14380)*. Springer, 33–54. https://doi.org/10.1007/978-3-031-46002-9_3
- [11] Asger Horn Brorholt, Kim Guldstrand Larsen, and Christian Schilling. 2024. Compositional Shielding and Reinforcement Learning for Multi-Agent Systems. *CoRR*. arXiv:2410.10460
- [12] Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. 2023. Safe Reinforcement Learning via Shielding under Partial Observability. In *AAAI*. AAAI Press, 14748–14756. <https://doi.org/10.1609/AAAI.V37I12.26723>
- [13] Krishnendu Chatterjee, Martin Chmelik, and Mathieu Tracol. 2016. What is decidable about partially observable Markov decision processes with ω -regular objectives. *J. Comput. Syst. Sci.* 82, 5 (2016), 878–911. <https://doi.org/10.1016/J.JCSS.2016.02.009>
- [14] Alexandre David, Peter Gjol Jensen, Kim Guldstrand Larsen, Axel Legay, Didier Lime, Mathias Grund Sørensen, and Jakob Haahr Taankvist. 2014. On Time with Minimal Expected Cost!. In *ATVA (LNCS, Vol. 8837)*. Springer, 129–145. https://doi.org/10.1007/978-3-319-11936-6_10
- [15] Ingy Elsayed-Aly, Suda Bharadwaj, Christopher Amato, Rüdiger Ehlers, Ufuk Topcu, and Lu Feng. 2021. Safe Multi-Agent Reinforcement Learning via Shielding. In *AAMAS*. ACM, 483–491. <https://doi.org/10.5555/3463952.3464013>
- [16] Dimitra Giannakopoulou, Kedar S. Namjoshi, and Corina S. Pasareanu. 2018. Compositional Reasoning. In *Handbook of Model Checking*. Springer, 345–383. https://doi.org/10.1007/978-3-319-10575-8_12
- [17] Manfred Jaeger, Peter Gjol Jensen, Kim Guldstrand Larsen, Axel Legay, Sean Sedwards, and Jakob Haahr Taankvist. 2019. Teaching Stratego to Play Ball: Optimal Synthesis for Continuous Space MDPs. In *ATVA (LNCS, Vol. 11781)*. Springer, 81–97. https://doi.org/10.1007/978-3-030-31784-3_5
- [18] Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem. 2020. Safe Reinforcement Learning Using Probabilistic Shields. In *CONCUR*, Vol. 171. 3:1–3:16. <https://doi.org/10.4230/LIPICS.CONCUR.2020.3>
- [19] Bettina Könighofer, Roderick Bloem, Rüdiger Ehlers, and Christian Pek. 2022. Correct-by-Construction Runtime Enforcement in AI - A Survey. In *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday (LNCS, Vol. 13660)*. Springer, 650–663. https://doi.org/10.1007/978-3-031-22337-2_31
- [20] Hanna Krasowski, Jakob Thumm, Marlon Müller, Lukas Schäfer, Xiao Wang, and Matthias Althoff. 2023. Provably Safe Reinforcement Learning: Conceptual Analysis, Survey, and Benchmarking. *Trans. Mach. Learn. Res.* 2023 (2023). <https://openreview.net/forum?id=mcN0ezbnzO>
- [21] Leslie Lamport. 1977. Proving the Correctness of Multiprocess Programs. *IEEE Trans. Software Eng.* 3, 2 (1977), 125–143. <https://doi.org/10.1109/TSE.1977.229904>
- [22] Kim Guldstrand Larsen, Marius Mikucionis, and Jakob Haahr Taankvist. 2015. Safe and Optimal Adaptive Cruise Control. In *Correct System Design - Symposium in Honor of Ernst-Rüdiger Olderog on the Occasion of His 60th Birthday (LNCS, Vol. 9360)*. Springer, 260–277. https://doi.org/10.1007/978-3-319-23506-6_17
- [23] Yangge Li, Haoqing Zhu, Katherine Braught, Keyi Shen, and Sayan Mitra. 2023. Verse: A Python Library for Reasoning About Multi-agent Hybrid System Scenarios. In *CAV (LNCS, Vol. 13964)*. Springer, 351–364. https://doi.org/10.1007/978-3-031-37706-8_18
- [24] Lukasz Mikulski, Wojciech Jamroga, and Damian Kurpiewski. 2022. Assume-Guarantee Verification of Strategic Ability. In *PRIMA (LNCS, Vol. 13753)*. Springer, 173–191. https://doi.org/10.1007/978-3-031-21203-1_11
- [25] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmash Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nat.* 518, 7540 (2015), 529–533. <https://doi.org/10.1038/NATURE14236>
- [26] Susan S. Owicki and David Gries. 1976. Verifying Properties of Parallel Programs: An Axiomatic Approach. *Commun. ACM* 19, 5 (1976), 279–285. <https://doi.org/10.1145/360051.360224>
- [27] Alireza Partovi and Hai Lin. 2014. Assume-guarantee cooperative satisfaction of multi-agent systems. In *ACC*. IEEE, 2053–2058. <https://doi.org/10.1109/ACC.2014.6859441>
- [28] Amir Pnueli. 1984. In Transition From Global to Modular Temporal Reasoning about Programs. In *Logics and Models of Concurrent Systems (NATO ASI Series, Vol. 13)*. Springer, 123–144. https://doi.org/10.1007/978-3-642-82453-1_5
- [29] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley. <https://doi.org/10.1002/9780470316887>
- [30] Zengyi Qin, Kaiqing Zhang, Yuxiao Chen, Jingkai Chen, and Chuchu Fan. 2021. Learning Safe Multi-agent Control with Decentralized Neural Barrier Certificates. In *ICLR*. OpenReview.net. https://openreview.net/forum?id=P6_q1BRxY8Q
- [31] Dhananjay Raju, Sudarshan Bharadwaj, Franck Djeumou, and Ufuk Topcu. 2021. Online Synthesis for Runtime Enforcement of Safety in Multiagent Systems. *IEEE Trans. Control. Netw. Syst.* 8, 2 (2021), 621–632. <https://doi.org/10.1109/TCNS.2021.3061900>
- [32] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR* (2017). arXiv:1707.06347
- [33] Eugene W. Stark. 1985. A Proof Technique for Rely/Guarantee Properties. In *FSTTCS (LNCS, Vol. 206)*. Springer, 369–391. https://doi.org/10.1007/3-540-16042-6_21
- [34] Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press. <http://incompleteideas.net/book/the-book-1st.html>
- [35] Wenli Xiao, Yiwei Lyu, and John M. Dolan. 2023. Model-based Dynamic Shielding for Safe and Efficient Multi-agent Reinforcement Learning. In *AAMAS*. ACM, 1587–1596. <https://doi.org/10.5555/3545946.3598814>
- [36] Wen-Chi Yang, Giuseppe Marra, Gavin Rens, and Luc De Raedt. 2023. Safe Reinforcement Learning via Probabilistic Logic Shields. In *IJCAI*. ijcai.org, 5739–5749. <https://doi.org/10.24963/IJCAI.2023/637>
- [37] Chao Yu, Akash Velu, Eugene Vinytsky, Jiaxuan Gao, Yu Wang, Alexandre M. Bayen, and Yi Wu. 2022. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In *NeurIPS*. http://papers.nips.cc/paper_files/paper/2022/hash/9c1535a02f0ce079433344e14d910597-Abstract-Datasets_and_Benchmarks.html
- [38] Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. 2019. Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms. *CoRR* (2019). arXiv:1911.10635
- [39] Wenbo Zhang and Osbert Bastani. 2019. MAMPS: Safe Multi-Agent Reinforcement Learning via Model Predictive Shielding. *CoRR* (2019). arXiv:1910.12639