

Interpreting Primal-Dual Algorithms for Constrained Multiagent Reinforcement Learning

Daniel Tabas

DTABAS@UW.EDU

University of Washington Department of Electrical Engineering, Seattle, WA, USA

Ahmed S. Zamzam

AHMED.ZAMZAM@NREL.GOV

National Renewable Energy Laboratory, Golden, CO, USA

Baosen Zhang

ZHANGBAO@UW.EDU

University of Washington Department of Electrical Engineering, Seattle, WA, USA

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

Constrained multiagent reinforcement learning (C-MARL) is gaining importance as MARL algorithms find new applications in real-world systems ranging from energy systems to drone swarms. Most C-MARL algorithms use a primal-dual approach to enforce constraints through a penalty function added to the reward. In this paper, we study the structural effects of this penalty term on the MARL problem. First, we show that the standard practice of using the constraint function as the penalty leads to a weak notion of safety. However, by making simple modifications to the penalty term, we can enforce meaningful probabilistic (chance and conditional value at risk) constraints. Second, we quantify the effect of the penalty term on the value function, uncovering an improved value estimation procedure. We use these insights to propose a constrained multiagent advantage actor critic (C-MAA2C) algorithm. Simulations in a simple constrained multiagent environment affirm that our reinterpretation of the primal-dual method in terms of probabilistic constraints is effective, and that our proposed value estimate accelerates convergence to a safe joint policy.

Keywords: Multiagent reinforcement learning, primal-dual methods, chance constraints, conditional value at risk

1. Introduction

As reinforcement learning (RL) algorithms progress from virtual to cyber-physical applications, it will be necessary to address the challenges of safety, especially when systems are controlled by multiple agents. Examples of multiagent safety-critical systems include power grids (Cui et al., 2022), building energy management (BEM) systems (Biagioni et al., 2022), autonomous vehicle navigation (Zhou et al., 2022), and drone swarms (Chen et al., 2020). In each of these applications, agents must learn to operate in a complicated environment while satisfying various local and system-wide constraints. Such constraints, derived from domain-specific knowledge, are designed to prevent damage to equipment, humans, or infrastructure or to preclude failure to complete some task or objective.

Constrained multiagent reinforcement learning (C-MARL) poses challenges beyond the single-agent constrained reinforcement learning (C-RL) problem because the interactions between agents can influence both the satisfaction of constraints and the convergence of policies. The potential scale of C-MARL problems eliminates the possibility of directly using common model-based methods for

C-RL, such as in [Chen et al. \(2018\)](#); [Ma et al. \(2021\)](#); [Tabas and Zhang \(2022\)](#). The main strategy for tackling C-MARL problems found in the literature is the Lagrangian or primal-dual method (see, e.g. [Lu et al. \(2021\)](#); [Li et al. \(2020\)](#); [Lee et al. \(2018\)](#); [Parnika et al. \(2021\)](#) and the references therein). Our aim is to understand some potential drawbacks of this approach and some ways these drawbacks can be mitigated.

In the primal-dual approach to C-MARL, each agent receives a reward signal that is augmented with a penalty term designed to incentivize constraint satisfaction. The magnitude of the penalty term is tuned to steer policies away from constraint violations while not unnecessarily overshadowing the original reward. Although this approach has been shown to converge to a safe joint policy under certain assumptions ([Lu et al., 2021](#)), it changes the structure of the problem in a way that is not well understood, leading to two challenges.

The first challenge is that the primal-dual algorithm only enforces *discounted sum constraints* derived from the original safety constraints of the system. As we will show, discounted sum constraints guarantee safety only in expectation, which is difficult to interpret. We propose simple modifications to the penalty term that enable the enforcement of more interpretable constraints, namely, chance constraints ([Mesbah, 2016](#)) and conditional value at risk constraints ([Rockafellar and Uryasev, 2000](#)), providing bounds on the probability and the severity of future constraint violations. There have been several C-RL algorithms that work with risk sensitivities ([García and Fernández, 2015](#); [Chow et al., 2018](#)), but the multiagent context is less studied, and our contributions provide a novel understanding of the safety guarantees provided by C-MARL algorithms.

The second challenge is the fact that the reward is constantly changing as the dual variables are updated, which diminishes the accuracy of value estimates. We quantify this loss of accuracy, and we propose a new value estimation procedure to overcome it. Our proposal builds on results in [Tessler et al. \(2019\)](#) showing the affine relationship between the value function and the dual variables. We develop a novel class of temporal difference algorithms for value function estimation that directly exploits this observation, giving rise to a value estimate that maintains an accurate derivative with respect to the dual variables. Compared to existing algorithms, our estimates are much more robust to dual variable updates.

The specific C-MARL formulation we study in this paper is inspired by the BEM problem ([Molina-Solana et al., 2017](#); [Biagioni et al., 2022](#)), illustrated in Figure 1. The main objective of BEM is to control a building’s resources to minimize the cost of energy consumption while maintaining comfort and convenience for the occupants. However, when BEMs are deployed in multiple buildings, it is critical to ensure that the power network connecting them is safely operated because the uncoordinated control of buildings can cause network-level voltage or power flow violations. This mandates a level of coordination among agents in the learning stage; thus, we adopt the commonly-studied centralized training/decentralized execution (CTDE) framework ([Lowe et al., 2017](#); [Foerster et al., 2018](#)),

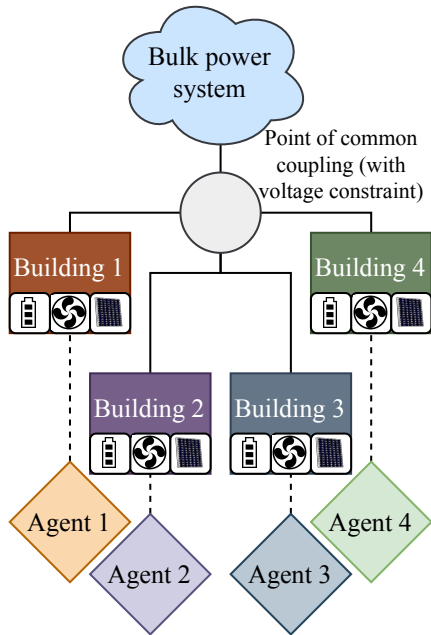


Figure 1: BEM with a voltage constraint at the point of common coupling.

in which a simulator or coordinator provides global state information, constraint evaluations, and Lagrange multipliers (dual variables) to each agent during training. During the testing (execution) phase, we assume that there is no real-time communication between the agents. This stems from the need for privacy and the lack of communication infrastructure in practical systems ¹.

The rest of the paper is organized as follows. In Section 2, we formulate the problem under consideration. In Section 3, we provide an overview of our main interpretive tool, the occupation measure (Borkar and Bhatt, 1996). In Section 4, we use the occupation measure to reformulate discounted sum constraints as probabilistic constraints. In Section 5, we study the value structure of the primal-dual problem and use the results to propose a new value estimation algorithm. In Section 6, we provide some simulation results affirming the contribution of the theoretical observations.

1.1. Notation

The natural numbers and the real numbers are denoted \mathbb{N} and \mathbb{R} , respectively. Given a measurable set \mathcal{S} , the set of all possible probability densities over \mathcal{S} is denoted as $\Delta_{\mathcal{S}}$. For any discount factor $\gamma \in (0, 1)$ and any sequence $\{y_t\}_{t=0}^T$, the discounted sum operator is $\Gamma_{t=0}^T[y_t | \gamma] = (1 - \gamma) \sum_{t=0}^T \gamma^t y_t$, and $\Gamma_{t=0}^{\infty}[y_t | \gamma] = \lim_{T \rightarrow \infty} \Gamma_{t=0}^T[y_t | \gamma]$ if the limit exists. We often drop the second argument γ for brevity. The positive component operator is $[y]_+ = \max\{y, 0\}$, and the logical indicator function $I[\cdot]$ maps $\{\text{True}, \text{False}\}$ to $\{1, 0\}$.

2. Problem formulation

2.1. Constrained MARL

We consider a noncooperative setting in which n agents pursue individual objectives while subject to global constraints (e.g., a shared resource constraint). We assume there is no real-time communication, and that each agent’s action is based only on its local observations. However, policy updates can use global information under the CTDE framework (Lowe et al., 2017; Foerster et al., 2018). In this paper, we consider the case of continuous state and action spaces.

The setting is described by the tuple $(\{\mathcal{X}_i\}_{i \in \mathcal{N}}, \{\mathcal{U}_i\}_{i \in \mathcal{N}}, \{R_i\}_{i \in \mathcal{N}}, f, C, p_0, \gamma)$, where \mathcal{N} is the index set of agents, $\mathcal{X}_i \subset \mathbb{R}^{n_x^i}$ and $\mathcal{U}_i \subset \mathbb{R}^{n_u^i}$ are the state and action spaces of agent i , and $R_i : \mathcal{X}_i \times \mathcal{U}_i \rightarrow \mathbb{R}$ is the reward function of agent i . We assume that the sets \mathcal{X}_i and \mathcal{U}_i are compact for all i . Let $\mathcal{X} = \prod_{i \in \mathcal{N}} \mathcal{X}_i$ and $\mathcal{U} = \prod_{i \in \mathcal{N}} \mathcal{U}_i$ be the joint state and action spaces of the system, respectively. Then $f : \mathcal{X} \times \mathcal{U} \rightarrow \Delta_{\mathcal{X}}$ describes the state transition probabilities, i.e., $f(\cdot | x, u)$ is a probability density function. The function $C : \mathcal{X} \rightarrow \mathbb{R}^m$ is used to describe a set of safe states, $\mathcal{S} = \{x \in \mathcal{X} | C(x) \leq 0\}$.

Let $p_0 \in \Delta_{\mathcal{X}}$ denote the initial state probability density and $\gamma \in (0, 1)$ be a discount factor. At time t , the state, action, and reward of agent i are x_t^i , u_t^i , and r_t^i , respectively, and constraint j evaluates to $c_t^j = C^j(x_t)$. Using a quantity without a superscript to represent a stacked vector ranging over all $i \in \mathcal{N}$ or all $j \in \{1, \dots, m\}$, a system trajectory is denoted $\tau = \{(x_t, u_t, r_t, c_t)\}_{t=0}^{\infty}$.

In the noncooperative C-MARL framework, each agent seeks to learn a policy $\pi_i : \mathcal{X}_i \rightarrow \Delta_{\mathcal{U}_i}$ that maximizes the expected discounted accumulation of individual rewards. We let $\pi : \mathcal{X} \rightarrow \Delta_{\mathcal{U}}$ denote the joint policy, and $f^{\pi} : \mathcal{X} \rightarrow \Delta_{\mathcal{X}}$ is the state transition probability induced by a joint policy π . The tuple (p_0, f, π) induces a state visitation probability density at each time step,

1. Even in buildings with advanced metering infrastructure or smart meters, they typically only exchange information with the utility a few times a day.

$p_t^\pi(x) = \int_{\mathcal{X}^t} p_0(x_0) \cdot \prod_{k=1}^t f^\pi(x_k | x_{k-1}) dx_0 \dots dx_{k-1}$, and we say $p_\infty^\pi(x) = \lim_{t \rightarrow \infty} p_t^\pi(x)$ for each $x \in \mathcal{X}$ if the limit exists. The collection of visitation probabilities $\{p_t^\pi\}_{t=0}^\infty$ gives rise to a probability density of trajectories τ , denoted $\mathcal{M} \in \Delta_{\prod_{t=0}^\infty (\mathcal{X} \times \mathcal{A} \times \mathbb{R}^n \times \mathbb{R}^m)}$; thus, the objective of each agent can be stated precisely as maximizing $\mathbb{E}_{\tau \sim \mathcal{M}}[\sum_{t=0}^\infty r_t^i]$.

The agents, however, must settle on a joint policy that keeps the system in the safe set \mathcal{S} . Due to the stochastic nature of the system, satisfying this constraint at all times is too difficult and in some cases too conservative. A common relaxation procedure is to formulate an augmented reward $\tilde{r}_t^i = r_t^i - \lambda^T c_t$ where $\lambda \in \mathbb{R}_+^m$, the *Lagrange multiplier* or *dual variable*, is adjusted to incentivize constraint satisfaction. This leads to the primal-dual algorithm for C-MARL, discussed in the next section. The following mild assumption facilitates the analysis.

Assumption 1 R^i, C^j , and p_t^π are bounded on \mathcal{X} for all $i \in \mathcal{N}$, all $j \in \{1, \dots, m\}$, and all $t \in \mathbb{N}$.

The boundedness of R^i and C^j is a common assumption (Lu et al., 2021; Tessler et al., 2019; Paternain et al., 2019) that we will use to exchange the order of limits, sums, and integrals using the dominated convergence theorem. The assumption of bounded p_t^π is not strictly necessary and does not change the results; however, we use it throughout the paper to simplify calculations.

2.2. Primal-dual algorithms

The augmented reward function leads to the following min-max optimization problem for agent i :

$$\min_{\lambda \geq 0} \max_{\pi_i} \mathbb{E}_{\tau \sim \mathcal{M}} \left[\sum_{t=0}^\infty [r_t^i - \lambda^T c_t] \right] \quad (1)$$

$$= \min_{\lambda \geq 0} \max_{\pi_i} \left(\mathbb{E}_{\tau \sim \mathcal{M}} \left[\sum_{t=0}^\infty [r_t^i] \right] - \lambda^T \mathbb{E}_{\tau \sim \mathcal{M}} \left[\sum_{t=0}^\infty [c_t] \right] \right) \quad (2)$$

where (2) uses absolute convergence (stemming from Assumption 1) to rearrange the terms of the infinite sum. Note that the minimization over λ is coupled across agents. Any fixed point of (2) will satisfy $\mathbb{E}_{\tau \sim \mathcal{M}}[\sum_{t=0}^\infty c_t] \leq 0$ because if $\mathbb{E}_{\tau \sim \mathcal{M}}[\sum_{t=0}^\infty c_t^j] \neq 0$, then the objective value can be reduced by increasing or decreasing λ_j , unless $\mathbb{E}_{\tau \sim \mathcal{M}}[\sum_{t=0}^\infty c_t^j] < 0$ and $\lambda_j = 0$. In other words, the primal-dual method enforces a *discounted sum constraint* derived from the safe set \mathcal{S} . Although discounted sum constraints are convenient, it is not obvious what they imply about safety guarantees with respect to the original constraints. We begin our investigation of discounted sum constraints by taking a closer look at a state visitation probability density known as the occupation measure.

3. Occupation measure

The *occupation measure* describes the average behavior of a Markov process in some sense which will be made precise shortly. As we will show, the occupation measure is instrumental in clarifying the role of discounted sum constraints. In this paper, we use a definition common for continuous-state, infinite-horizon discounted MDPs (Paternain et al., 2019; Silver et al., 2014).

Definition 2 (Occupation measure) *The occupation measure $\mu_\gamma^\pi \in \Delta_{\mathcal{X}}$ associated with discount factor γ , induced by a joint policy π , is defined for any $x \in \mathcal{X}$ as $\mu_\gamma^\pi(x) = \sum_{t=0}^\infty \gamma^t p_t^\pi(x)$.*

In this section, we provide some interpretations for the occupation measure before using it to ascribe meaning to discounted sum constraints. The first question one might ask is whether μ_γ^π is itself a pdf. It is, of course, nonnegative, and the following proposition shows it integrates to unity under mild conditions.

Proposition 3 *Under Assumption 1, $\int_{\mathcal{X}} \mu_\gamma^\pi(x) dx = 1$.*

The proof for Proposition 3 is in Appendix A². What does μ_γ^π tell us about the behavior of a system under a given policy? It describes the probability of visiting a certain state but with more weight placed on states that are likely to be visited *earlier* in time. In fact, μ_γ^π describes the near-term behavior in the following sense.

Proposition 4 *Under Assumption 1, for any $x \in \mathcal{X}$, the following statements hold:*

1. $\lim_{\gamma \rightarrow 0^+} \mu_\gamma^\pi(x) = p_0(x)$.
2. $\lim_{\gamma \rightarrow 1^-} \mu_\gamma^\pi(x) = \lim_{t \rightarrow \infty} p_t^\pi$ if the latter limit exists.

The proof for Proposition 4 is in Appendix A. Figure 2 provides an illustration of the result in Proposition 4 when p_t^π evolves as a normal distribution with mean 0.95^t and constant variance. The point at which μ_γ^π equally resembles p_0 and p_∞^π is exactly at $\gamma = 0.95$.

According to Proposition 4, the occupation measure describes a state distribution that lies between the initial and long-term behavior of the system. But where exactly does it lie in between these two extremes? The *effective horizon* of a discounted planning problem is often set to $T_1(\gamma) = \frac{1}{1-\gamma}$, which is the expected termination time if the probability of an episode terminating at any given time step is $(1 - \gamma)$ (Paternain et al., 2022); however, the concept of a random stopping time might not be sensible in all applications. Another way to define the effective horizon is to study the geometric accumulation of weights. In this case, the effective horizon can be measured as $T_2(\gamma, \varepsilon) = \min\{K \in \mathbb{N} : \prod_{t=0}^{K-1} [1] \geq 1 - \varepsilon\}$, where $\varepsilon \in (0, 1)$ is a tolerance. Using either of these two definitions, the occupation measure can be said to describe the behavior of the system from the start time up to the effective horizon. Specifically, one may truncate the sum in Definition 2 at the effective horizon to obtain a conceptual understanding of what the occupation measure describes.

Depending on the application, either T_1 or T_2 can provide a more sensible connection between discounted and finite-horizon problems. But are these two definitions related? The next proposition answers this affirmatively by showing that T_1 is actually a special case of T_2 .

Proposition 5 $T_1(\gamma) = T_2(\gamma, \varepsilon)$ when ε is set to $\gamma^{\frac{1}{1-\gamma}} \approx \frac{1}{e}$.

The proof for Proposition 5 is in Appendix A. Proposition 5 is illustrated in Figure 3, where the effective horizon is plotted as a function of γ for three different values of ε . With an understanding of the occupation measure as a visitation density describing behavior up to the effective horizon, we can begin to derive meaningful risk-related interpretations of discounted sum constraints. These interpretations lead directly to sensible recommendations for the design of C-MARL algorithms.




Figure 2: Example of the occupation measure for various levels of γ .

2. The full paper with appendix is available at <https://arxiv.org/abs/2211.16069>.

4. Discounted risk metrics

The discounted sum constraint can naturally be reinterpreted as a certain type of average constraint. In particular, Assumption 1 ensures the equivalence $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty} C(x_t)] = \mathbb{E}_{x \sim \mu_{\gamma}^{\pi}}[C(x)]$ (Paternain et al., 2019). This near-term average does not relate to any well-known risk metrics and hence does not provide a practical safety guarantee. In general, information about the mean of a distribution cannot be used to infer information about its tails; however, simple changes to the penalty function can yield information about either the *probability* of incurring a constraint violation or the expected *severity* of constraint violations.




Figure 3: Effective horizon length as a function of γ .

Proposition 6 (Near-term probability of constraint violations) *Suppose that for some $\delta_j \in [0, 1]$ and $\alpha_j \in \mathbb{R}$, we have $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty} I[C^j(x_t) \geq \alpha_j]] \leq \delta_j$. Then under Assumption 1, $\Pr\{C^j(x) \geq \alpha_j \mid x \sim \mu_{\gamma}^{\pi}\} \leq \delta_j$.*

Proof $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty} I[C^j(x_t) \geq \alpha_j]] = \mathbb{E}_{x \sim \mu_{\gamma}^{\pi}}[I[C^j(x) \geq \alpha_j]] = \Pr\{C^j(x) \geq \alpha_j \mid x \sim \mu_{\gamma}^{\pi}\}$. The first equality uses Assumption 1 to apply an equivalence established in e.g. Paternain et al. (2019). The second equality follows from the definition of expectation. ■

Proposition 6 makes it easy to enforce chance constraints using primal-dual methods. When the penalty term $C^j(x)$ is replaced by the quantity $I[C^j(x) \geq \alpha_j] - \delta_j$, the primal-dual algorithm enforces $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty} I[C^j(x_t) \geq \alpha_j]] - \delta_j \leq 0$. By Proposition 6, this guarantees that $\Pr\{C^j(x) \geq \alpha_j \mid x \sim \mu_{\gamma}^{\pi}\} \leq \delta_j$. Because the probability of constraint violations is defined with x varying over μ_{γ}^{π} , we call the resulting guarantee a *near-term* or *discounted chance constraint*. This can be repeated for each $j \in \{1, \dots, m\}$, providing a set of bounds on the probability of violating *each* constraint by more than its tolerance α_j . On the other hand, we can control the probability of violating *any* constraint as follows. Define the statement $C(x) \geq \alpha$ to be true if $C^j(x) \geq \alpha_j \forall j \in \{1, \dots, m\}$, and false otherwise. Then, applying Proposition 6 to the test condition $C(x) \geq \alpha$ will result in a bound on $\Pr\{C(x) \geq \alpha \mid x \sim \mu_{\gamma}^{\pi}\}$.

While discounted chance constraints enable one to control the *probability* of extreme events in the near future, conditional value at risk constraints (Rockafellar and Uryasev, 2000) afford control over the *severity* of such events.

Definition 7 (Rockafellar and Uryasev (2000)) *Given a risk level $\beta \in (0, 1)$, a cost $h : \mathcal{X} \rightarrow \mathbb{R}$, and a probability density μ on \mathcal{X} , the value at risk (VaR) and conditional value at risk (CVaR) are defined as:*

$$\begin{aligned} \text{VaR}(\beta, h, \mu) &= \min\{\alpha \in \mathbb{R} : \Pr\{h(x) \leq \alpha \mid x \sim \mu\} \geq \beta\}, \\ \text{CVaR}(\beta, h, \mu) &= \frac{1}{1 - \beta} \int_{h(x) \geq \text{VaR}(\beta, h, \mu)} h(x) \mu(x) dx. \end{aligned}$$

In other words, $\text{VaR}(\beta, h, \mu)$ is the least upper bound on h that can be satisfied with probability β , while $\text{CVaR}(\beta, h, \mu)$ describes the expected value in the VaR-tail of the distribution of h . CVaR

characterizes the expected severity of extreme events, which can be defined precisely as the $(1 - \beta)$ fraction of events x with the worst outcomes as ranked by the cost incurred, $h(x)$. The VaR and CVaR for $h(x) = x$, when x follows a standard normal distribution, are illustrated in Figure 4, where the shaded region has an area of $(1 - \beta)$. For the rest of the paper, we assume that the cdf of $h(x)$ is continuous when $x \sim \mu$. For further details and for cases in which this assumption does not hold, we refer the reader to [Rockafellar and Uryasev \(2002\)](#).

Proposition 8 (Near-term CVaR) *For any $\alpha_j \geq 0$, suppose that $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty}[[C^j(x_t) - \alpha_j]_+]] \leq \delta_j$. Then, $\text{CVaR}(\beta, C^j, \mu_{\gamma}^{\pi}) \leq \alpha_j + (1 - \beta)^{-1}\delta_j$.*

Proof Under Assumption 1, the identity $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty}[C^j(x_t) - \alpha_j]_+] = \mathbb{E}_{x \sim \mu_{\gamma}^{\pi}}[[C^j(x) - \alpha_j]_+]$ holds ([Paternain et al., 2019](#)). Next, we use the fact that the CVaR is the minimum value of the convex function in α_j given by $F(\alpha_j \mid \beta, C^j, \mu_{\gamma}^{\pi}) := \alpha_j + (1 - \beta)^{-1}\mathbb{E}_{x \sim \mu_{\gamma}^{\pi}}[[C^j(x) - \alpha_j]_+]$ ([Rockafellar and Uryasev, 2000](#)); thus, F provides an upper bound on CVaR. Some rearranging leads to the result. \blacksquare

Similar to the chance-constrained case, Proposition 8 makes it easy to enforce CVaR constraints in the primal-dual algorithm. Here, the penalty term used is $[C^j(x) - \alpha_j]_+ - \delta_j$. Using this penalty, the algorithm enforces $\mathbb{E}_{\tau \sim \mathcal{M}}[\Gamma_{t=0}^{\infty}[[C^j(x_t) - \alpha_j]_+]] - \delta_j \leq 0$, which by Proposition 8 implies $\text{CVaR}(\beta, C^j, \mu_{\gamma}^{\pi}) \leq \alpha_j + (1 - \beta)^{-1}\delta_j$. By repeating for each $j \in \{1, \dots, m\}$, we can bound the expected severity of the constraint violations for each of the m constraints. Because the CVaR constraint is defined with x varying over μ_{γ}^{π} , the resulting guarantee is called a *near-term* or *discounted CVaR constraint*.

To obtain a tight bound on the CVaR, α_j must be set to $\text{VaR}(\beta, C^j, \mu_{\gamma}^{\pi})$, which minimizes the function F introduced in the proof of Proposition 8 ([Rockafellar and Uryasev, 2000](#)). Unfortunately, the VaR is not known ahead of time. [Chow et al. \(2018\)](#) include α_j as an optimization variable in the learning procedure, but extending their technique to the multiagent setting is not straightforward. Our approach is to include it as a tunable hyperparameter. Simulation results in Section 6 show that it is easy to choose α_j to give a nearly tight bound.




Figure 4: Example of VaR and CVaR at risk level $\beta = 0.9$.

5. Primal-dual value functions

In this section, we investigate challenges with value estimation in the primal-dual regime. The fact that the reward to each agent is constantly changing (due to dual variable updates) makes it difficult to accurately estimate state values. To quantify this decrease in accuracy, we introduce the value functions induced by the joint policy π , $\{V_{\pi}^i : \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}\}_{i \in \mathcal{N}}$, $\{V_{R,\pi}^i : \mathcal{X} \rightarrow \mathbb{R}\}_{i \in \mathcal{N}}$, $V_{C,\pi} : \mathcal{X} \rightarrow \mathbb{R}^m$ where:

$$V_{\pi}^i(x, \lambda) = \mathbb{E}_{\tau \sim \mathcal{M}}\left[\sum_{t=0}^{\infty} r_t^i - \lambda^T c_t \mid x_0 = x\right], \quad (3)$$

$$V_{R,\pi}^i(x) = \mathbb{E}_{\tau \sim \mathcal{M}}\left[\sum_{t=0}^{\infty} r_t^i \mid x_0 = x\right], \quad V_{C,\pi}(x) = \mathbb{E}_{\tau \sim \mathcal{M}}\left[\sum_{t=0}^{\infty} c_t \mid x_0 = x\right]. \quad (4)$$

Note that c_t could be modified as indicated in Section 4, and the following results would hold for the modified penalty function.

Obviously, it is impossible to learn an accurate value function when λ is unknown and changing; however, simply making λ available to a value function approximator does not guarantee good generalization beyond previously seen values of λ . Having a good estimate of the *derivative* of the value function with respect to λ will ensure accuracy under small perturbations to the dual variables. Fortunately, this derivative is easy to obtain. Formally, under Assumption 1, we can write $V_\pi^i(x, \lambda) = V_{R,\pi}^i(x) - \lambda^T V_{C,\pi}(x)$ (Tessler et al., 2019), and therefore, $\nabla_\lambda V_\pi^i(x, \lambda) = -V_{C,\pi}(x)$. By learning $V_{R,\pi}^i$ and $V_{C,\pi}$ as separate functions and then combining them using the true value of λ , we can construct a value estimate whose derivative with respect to the dual variables is as accurate as our estimate of $V_{C,\pi}$ itself. This estimate will be more robust to small changes in λ . We will refer to this type of value estimate as a *structured value function* or a *structured critic*.

Proposition 9 *Let $\bar{c} = \mathbb{E}_{x \sim \mu_\pi^\gamma}[C(x)]$ and $\Sigma_C^2 = \mathbb{E}_{x \sim \mu_\pi^\gamma}[(\bar{c} - C(x))(\bar{c} - C(x))^T]$. Suppose λ is randomly varying with mean $\bar{\lambda}$ and variance Σ_λ^2 . Using a structured value function approximator can reduce the mean square temporal difference error by up to $\text{Tr}[\Sigma_\lambda^2 \cdot (\Sigma_C^2 + \bar{c}\bar{c}^T)]$.*

The proof of Proposition 9 is in Appendix A. Figure 5 illustrates Proposition 9 in a simple value estimation task with quadratic rewards, linear dynamics and policies, linear state constraints, and randomly varying λ . The *generic critic* (GC) is a value function modeled as a quadratic function of the state only. The *input-augmented critic* (IAC) is a value function modeled as an unknown quadratic function of the state and dual variables, while the *structured critic* (SC) is modeled using $\hat{V}_\pi^i = \hat{V}_{R,\pi}^i - \lambda^T \hat{V}_{C,\pi}$ with quadratic $\hat{V}_{R,\pi}^i$ and linear $\hat{V}_{C,\pi}$ trained on their respective signals.

The dashed line in Figure 5 is at the value $\text{Tr}[\Sigma_\lambda^2 \cdot (\Sigma_C^2 + \bar{c}\bar{c}^T)]$ predicted in Proposition 9. In this simple value estimation task, high accuracy can be achieved when conditioning on the randomly varying λ ; however, having an accurate estimate of $\nabla_\lambda V_\pi^i$ by using a structured critic is also shown to help. Although in practice $\bar{\lambda}$ and Σ_λ^2 change over time, the simulation results in Section 6 confirm that using structured critics improves performance. The loss function for value function approximation is therefore given by:

$$TDE(x, x') = [R^i(x^i) + \gamma \hat{V}_{R,\pi}^i((x^i)') - \hat{V}_{R,\pi}^i(x^i)]^2 + \|C(x) + \gamma \hat{V}_{C,\pi}(x') - \hat{V}_{C,\pi}(x)\|_2^2 \quad (5)$$

where $x \in \mathcal{X}$ and $x' \sim f^\pi(x)$. Equation (5) is simply a sum of squared temporal difference errors over the set of $m + 1$ value functions. For algorithmic details, we refer the reader to Appendix B.

6. Simulations

In our simulations, we sought to demonstrate the effectiveness of the penalty modifications and structured critic proposed in sections 4 and 5. We tested our findings in a modified multiagent particle environment³ (Lowe et al., 2017) with two agents pursuing individual objectives subject




Figure 5: Temporal difference error trajectories in a simple policy evaluation task.

3. Code for the environments is available at github.com/dtabas/multiagent-particle-envs.

to a constraint on the joint state. The state of each agent is its position and velocity in \mathbb{R}^2 , i.e. $x^i = [y^{iT} \ v^{iT}]^T$ where $y^i \in \mathbb{R}^2$ is the position and $v^i \in \mathbb{R}^2$ is the velocity of agent i . The objective of each agent is to drive its position y^i to a landmark $y^{i*} \in \mathbb{R}^2$, while making sure that the agent ensemble satisfies the safety constraint. The reward and constraint functions are given by:

$$R^i(y^i) = -\xi_i \|y^i - y^{i*}\|_2^2, \quad C(y) = \mathbf{1}^T y \quad (6)$$

where $\xi_i > 0$ is a constant and $y = [y^{1T} \ y^{2T}]^T$ is the position of the agent ensemble.

The landmark $y^* = [y^{1*T} \ y^{2*T}]^T$ is stationed outside of the safe region $\mathcal{S} = \{y \mid C(y) \leq 0\}$. Thus, the agents cannot both reach their goals while satisfying $C(y) \leq 0$. To train the agents to interact in this environment, we used a modified version of the EPyMARL codebase⁴ (Papoudakis et al., 2020). We tested several MARL algorithms, including MADDPG (Lowe et al., 2017), COMA (Foerster et al., 2018), and MAA2C (Papoudakis et al., 2020). We decided to use the MAA2C algorithm because it consistently produced the best results and because as a value function-based algorithm, it provided the most straightforward route to implementing the changes proposed in Section 5. Details of the algorithm, pseudocode, hyperparameters, and supplementary simulation results are provided in Appendix B.

For each risk metric described in Section 4, we tested the convergence of the agents to a safe policy with and without modifications to the penalty and value functions. Figure 6 shows the results when we make the substitution $C(x) \leftarrow I[C(x) \geq \alpha] - \delta$ in the penalty function to enforce a chance constraint, $\Pr\{C(x) \geq \alpha \mid x \sim \mu_\gamma^\pi\} \leq \delta$ with α and δ each set to 0.1. The modified penalty function performs the best as a chance constraint-enforcing signal (red and green lines in Figure 6). Whether or not the penalty function is modified, the structured critic finds safer policies throughout training (red vs. green and orange vs. blue lines).

Figure 7 shows the results when we make the substitution $C(x) \leftarrow [C(x) - \alpha]_+ - \delta$ in the penalty function to enforce the constraint $\text{CVaR}(\beta, C, \mu_\gamma^\pi) \leq \alpha + (1 - \beta)^{-1}\delta$. Using the modified penalty (red and green lines in Figure 7) drives the CVaR upper bound (drawn in dashed lines) to the target value, and due to the choice of α , this bound is nearly tight. On the other hand, using the original penalty results in an overly conservative policy that achieves low risk at the expense of rewards (right panel). We also point out that when using the modified penalty with the structured critic, the CVaR is lower throughout training compared to when the generic critic is used, indicating improved effectiveness in enforcing limits on risk.

We chose α using the following heuristic, to make the bound on CVaR nearly tight. The ‘‘correct’’ value of α that would achieve a tight bound is $\text{VaR}(\beta, C, \mu_\gamma^\pi)$. Moreover, the upper bound that we used is convex and continuously differentiable in α (Rockafellar and Uryasev, 2000); therefore, small errors in α will lead to small errors in the upper bound on CVaR, and any approximation of




Figure 6: $\Pr\{C(x) \geq 0.1 \mid x \sim \mu_\gamma^\pi\}$ measured throughout training. Key: SC = structured critic, MP = modified penalty (Prop. 6). Both modifications speed convergence to a safe policy. The shaded region represents ± 1 standard deviation across 5 training runs.

4. Code for the algorithms is available at github.com/dtabas/epymar1.




Figure 7: $\text{CVaR}(\beta = 0.9, C, \mu_\gamma^\pi)$ measured throughout training. Key: SC = structured critic, MP = modified penalty (Prop. 8). The dashed lines represent the CVaR upper bound used in Prop. 8. The panel on the right shows progress toward the original objective through the total original returns, $\sum_{i=1}^2 \Gamma_{t=0}^T r_t^i$, without penalty terms. The shaded region represents ± 1 standard deviation across 5 training runs. The rewards increase then decrease because the agents first learn to navigate towards the landmark, which is outside the safe region, then learn to back off to satisfy the constraint.

the VaR will suffice. We obtained an approximation simply by running the simulation once with α set to zero and computing $\text{VaR}(\beta, C, \mu_\gamma^\pi)$ over some test trajectories. If necessary, the process could be repeated additional times. Alternatively, α could be tuned adaptively by computing VaR online, but the stability of such a procedure would need further investigation.

7. Conclusion

In this paper, we studied the effect of primal-dual algorithms on the structure of C-MARL problems. First, we used the occupation measure to study the effect of the penalty term on safety. We showed that using the constraint function as the penalty enforces safety only in expectation, but by making simple modifications to the penalty term, one may enforce meaningful probabilistic safety guarantees, namely, chance and CVaR constraints. These risk metrics are defined over the occupation measure, leading to notions of safety in the near term. Next, we studied the effect of the penalty term on the value function. When the dual variable and constraint evaluation signals are available, it is easy to model the relationship between the penalty term and the value function. By exploiting this structure, the accuracy of the value function can be improved. We demonstrated the usefulness of both of these insights in a constrained multiagent particle environment, showing that convergence to a low-risk policy is accelerated. One open question is the effect of primal-dual methods on game outcomes. Some agents might pay a higher price than others for modifying their policies to satisfy system-wide constraints. Understanding and mitigating this phenomenon will be the focus of future work.

Acknowledgments

D. Tabas and A. S. Zamzam would like to thank Patrick Emami and Xiangyu Zhang for helpful discussions in the early stages of this work, and Georgios Papoudakis for advice on software implementation. This work is partially supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2140004. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This work was authored in part by the National Renewable Energy Laboratory (NREL), operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The work of A. S. Zamzam was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

- David Biagioni, Xiangyu Zhang, Dylan Wald, Deepthi Vaidhyanathan, Rohit Chintala, Jennifer King, and Ahmed S. Zamzam. PowerGridworld: A Framework for Multi-Agent Reinforcement Learning in Power Systems. *Proceedings of the 2022 13th ACM International Conference on Future Energy Systems*, pages 565–570, 2022.
- Vivek S. Borkar and Abhay G. Bhatt. Occupation Measures for Controlled Markov Processes: Characterization and Optimality. *The Annals of Probability*, 24(3):1531–1562, 1996.
- Steven Chen, Kelsey Saulnier, Nikolay Atanasov, Daniel D. Lee, Vijay Kumar, George J. Pappas, and Manfred Morari. Approximating Explicit Model Predictive Control Using Constrained Neural Networks. In *Proceedings of the American Control Conference*, pages 1520–1527, 2018.
- Yu Jia Chen, Deng Kai Chang, and Cheng Zhang. Autonomous Tracking Using a Swarm of UAVs: A Constrained Multi-Agent Reinforcement Learning Approach. *IEEE Transactions on Vehicular Technology*, 69(11):13702–13717, 2020.
- Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18: 1–51, 2018.
- Wenqi Cui, Jiayi Li, and Baosen Zhang. Decentralized safe reinforcement learning for inverter-based voltage control. *Electric Power Systems Research*, 211(108609), 2022.
- Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. *32nd AAAI Conference on Artificial Intelligence*, pages 2974–2982, 2018.
- Javier García and Fernando Fernández. A Comprehensive Survey on Safe Reinforcement Learning. *Journal of Machine Learning Research*, 16:1437–1480, 2015.

- Donghwan Lee, Hyungjin Yoon, and Naira Hovakimyan. Primal-Dual Algorithm for Distributed Reinforcement Learning: Distributed GTD. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1967–1972, 2018.
- Wenhao Li, Bo Jin, Xiangfeng Wang, Junchi Yan, and Hongyuan Zha. F2A2: Flexible Fully-decentralized Approximate Actor-critic for Cooperative Multi-agent Reinforcement Learning. *arXiv: 2004.11145*, pages 1–42, 2020.
- Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. In *31st Conference on Neural Information Processing Systems*, 2017.
- Songtao Lu, Kaiqing Zhang, Tianyi Chen, Tamer Başar, and Lior Horesh. Decentralized Policy Gradient Descent Ascent for Safe Multi-Agent Reinforcement Learning. *35th AAAI Conference on Artificial Intelligence*, pages 8767–8775, 2021.
- Haitong Ma, Jianyu Chen, Shengbo Eben, Ziyu Lin, Yang Guan, Yangang Ren, and Sifa Zheng. Model-based Constrained Reinforcement Learning using Generalized Control Barrier Function. *IEEE International Conference on Intelligent Robots and Systems*, pages 4552–4559, 2021.
- Ali Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.
- Miguel Molina-Solana, María Ros, M. Dolores Ruiz, Juan Gómez-Romero, and M. J. Martín-Bautista. Data science for building energy management: A review. *Renewable and Sustainable Energy Reviews*, 70:598–609, 2017.
- Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. In *35th Conference on Neural Information Processing Systems*, 2020.
- P. Parnika, Raghuram Bharadwaj Diddigi, Sai Koti Reddy Danda, and Shalabh Bhatnagar. Attention actor-critic algorithm for multi-agent constrained co-operative reinforcement learning. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, 3: 1604–1606, 2021.
- Santiago Paternain, Luiz F.O. Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Santiago Paternain, Miguel Calvo-Fullana, Luiz F.O. Chamon, and Alejandro Ribeiro. Safe Policies for Reinforcement Learning via Primal-Dual Methods. *IEEE Transactions on Automatic Control*, 2022.
- R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of Conditional Value-at-Risk. *Journal of Risk*, 2:21–42, 2000.
- R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471, 2002.

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. *31st International Conference on Machine Learning, ICML 2014*, 1:605–619, 2014.

Daniel Tabas and Baosen Zhang. Computationally Efficient Safe Reinforcement Learning for Power Systems. In *Proceedings of the American Control Conference*, pages 3303–3310. American Automatic Control Council, 2022.

Chen Tessler, Daniel J. Mankowitz, and Shie Mannor. Reward constrained policy optimization. In *7th International Conference on Learning Representations, ICLR 2019*, May 2019.

Wei Zhou, Dong Chen, Jun Yan, Zhaojian Li, Huilin Yin, and Wanchen Ge. Multi-agent reinforcement learning for cooperative lane changing of connected and autonomous vehicles in mixed traffic. *Autonomous Intelligent Systems*, 2(1), 2022.