



Delft University of Technology

Interval Markov Decision Processes with Continuous Action-Spaces

Delimpaltadakis, Giannis; Lahijanian, Morteza; Mazo, Manuel; Laurenti, Luca

DOI

[10.1145/3575870.3587117](https://doi.org/10.1145/3575870.3587117)

Publication date

2023

Document Version

Final published version

Published in

HSCC 2023 - Proceedings of the 26th ACM International Conference on Hybrid Systems

Citation (APA)

Delimpaltadakis, G., Lahijanian, M., Mazo, M., & Laurenti, L. (2023). Interval Markov Decision Processes with Continuous Action-Spaces. In *HSCC 2023 - Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control, Part of CPS-IoT Week* Article 12 Association for Computing Machinery (ACM). <https://doi.org/10.1145/3575870.3587117>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Interval Markov Decision Processes with Continuous Action-Spaces

Giannis Delimpaltadakis

i.delimpaltadakis@tue.nl

Control Systems Technology Group
Mechanical Engineering
Eindhoven University of Technology
Eindhoven, Netherlands

Manuel Mazo Jr.

m.mazo@tudelft.nl

Delft Center for Systems and Control
Mechanical Engineering
Delft University of Technology
Delft, Netherlands

Morteza Lahijanian

morteza.lahijanian@colorado.edu

Aerospace Engineering Sciences
Computer Science
University of Colorado Boulder
Boulder, Colorado, U.S.A.

Luca Laurenti

l.laurenti@tudelft.nl

Delft Center for Systems and Control
Mechanical Engineering
Delft University of Technology
Delft, Netherlands

ABSTRACT

Interval Markov Decision Processes (IMDPs) are finite-state uncertain Markov models, where the transition probabilities belong to intervals. Recently, there has been a surge of research on employing IMDPs as abstractions of stochastic systems for control synthesis. However, due to the absence of algorithms for synthesis over IMDPs with continuous action-spaces, the action-space is assumed discrete a-priori, which is a restrictive assumption for many applications. Motivated by this, we introduce continuous-action IMDPs (caIMDPs), where the bounds on transition probabilities are functions of the action variables, and study value iteration for maximizing expected cumulative rewards. Specifically, we decompose the max-min problem associated to value iteration to $|Q|$ max problems, where $|Q|$ is the number of states of the caIMDP. Then, exploiting the simple form of these max problems, we identify cases where value iteration over caIMDPs can be *solved efficiently* (e.g., with linear or convex programming). We also gain other interesting insights: e.g., in certain cases where the action set \mathcal{A} is a polytope, synthesis over a discrete-action IMDP, where the actions are the vertices of \mathcal{A} , is sufficient for optimality. We demonstrate our results on a numerical example. Finally, we include a short discussion on employing caIMDPs as abstractions for control synthesis.

CCS CONCEPTS

• **Mathematics of computing** → **Markov processes; Stochastic control and optimization**; • **Computing methodologies** → **Planning under uncertainty**; *Computational control theory*.



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

HSCC '23, May 09–12, 2023, San Antonio, TX, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0033-0/23/05.

<https://doi.org/10.1145/3575870.3587117>

KEYWORDS

uncertain Markov decision processes; bounded-parameter Markov decision processes; value iteration; planning under uncertainty; control synthesis

ACM Reference Format:

Giannis Delimpaltadakis, Morteza Lahijanian, Manuel Mazo Jr., and Luca Laurenti. 2023. Interval Markov Decision Processes with Continuous Action-Spaces. In *Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control (HSCC '23), May 09–12, 2023, San Antonio, TX, USA*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3575870.3587117>

1 INTRODUCTION

1.1 Motivation and Contributions

Interval Markov Decision Processes (IMDPs; alternatively called bounded-parameter Markov decision processes) are a class of uncertain finite-state Markov Decision Processes (MDPs), where transition probabilities between states are only known to belong to intervals [8]. Due to their modelling flexibility and the availability of efficient planning algorithms [8], IMDPs have recently gained popularity in the control and computer-science communities for verification of and control-synthesis for uncertain systems (see, e.g., [4, 5, 12, 14, 15]). However, due to the absence of computational algorithms for IMDPs with continuous action-spaces, the action set is, generally, assumed discrete [12, 14]. This is a restrictive assumption as many realistic control applications involve continuous underlying action-spaces. Furthermore, with the current available algorithms, the only way to address continuous action-spaces is either a) to discretize the action-space and deal with discrete-action IMDPs, or b) solve the optimization problems associated to control synthesis using heuristics, as recently proposed in [5]. Nonetheless, a) scales exponentially with the dimension of the action space, rendering high-dimensional problems intractable. Moreover, b) - and a) - when the discretization is carried out blindly - results into suboptimal solutions without suboptimality bounds.

Motivated by the above, we introduce continuous-action Interval Markov Decision Processes (caIMDPs), where the transition probability intervals are functions of the action variables. We study value iteration over caIMDPs for synthesizing policies that maximize finite-horizon pessimistic expected cumulative rewards (which is referred to as the *robust control problem*), and identify cases where value iteration can be *solved efficiently*. Specifically, we show that the max-min optimization problem associated to value iteration is equivalent to solving $|Q|$ maximization problems, where $|Q|$ is the number of states of the caIMDP. Then, exploiting the simple form of these maximization problems, we distinguish the following tractable cases: 1) linear (on the action variable) bounds on transition probabilities and polytopic action set \mathcal{A} , where the maximization problems are linear programs, 2) concave and convex, respectively, transition bounds and convex \mathcal{A} , where we have convex programs, and 3) convex and concave, respectively, transition bounds and polytopic \mathcal{A} , which amounts to convex maximization over polytopes. As a result, in these cases, control synthesis over caIMDPs, via value iteration, comes with guaranteed optimality and fast computation. Moreover, we gain further interesting insights. For example, in cases 1 and 3, synthesizing over discrete-action IMDPs, where the action set consists of the vertices of \mathcal{A} , is sufficient for optimality. Overall, these results provide the necessary theory to synthesize optimal policies efficiently over caIMDPs.

We showcase our results on a numerical example, which demonstrates that synthesizing over a caIMDP is not only the only way to optimality, but it is also computationally efficient, even when compared to synthesizing over discrete-action IMDPs. Finally, we briefly discuss about using caIMDP-abstractions for control synthesis for stochastic systems, touching upon constructing the abstraction and obtaining suboptimality¹ bounds on the obtained policy.

1.2 Related work

While several existing works have focused on developing control and verification algorithms for finite-state and finite-action IMDPs [8, 11, 16], efficient algorithms for control of IMDPs with continuous action-spaces are missing. A first attempt to close this gap has been proposed in [5], where the authors present an algorithm to synthesize policies that maximize a reachability probability based on suboptimal heuristics and non-convex optimization. In contrast, here we derive an exact reformulation of value iteration for caIMDPs, which leads to tractable solutions based on linear or convex programming, in many cases of interest. caIMDPs are also closely related to parametric MDPs, which are Markov decision processes with finite state-/action-spaces, but where the transition probabilities may depend on some parameters [3, 9, 13]. However, a key difference is that, in parametric MDPs, the problem is to find parameter values that are fixed and independent of the state of the system. Instead, in caIMDPs we seek feedback control policies that depend on both time and state, thus requiring substantially different approaches.

¹The policy computed over the (ca)IMDP-abstraction, even if optimal for the abstraction, will generally be suboptimal for the original system. That is because the continuous state-space has been partitioned to a finite number of subsets, which represent the (ca)IMDP's states.

1.3 Notation

\mathbb{R} (resp., $\mathbb{R}_{\geq 0}$) stands for the set of real numbers (resp., non-negative reals). Given a polytope $\mathcal{A} \subset \mathbb{R}^n$, its set of vertices is denoted by $\text{ver}(\mathcal{A})$. Given a discrete set Q , we denote its cardinality by $|Q|$. By $\mathbb{1}$ we denote vectors, of appropriate dimension, of which all entries are equal to 1. For a vector $x \in \mathbb{R}^n$, $x \geq 0$ ($x \leq 0$) denotes that all its components are non-negative (non-positive, resp.).

2 CONTINUOUS-ACTION INTERVAL MARKOV DECISION PROCESSES (caIMDPs)

2.1 The basic elements of a caIMDP

Continuous-action interval Markov decision processes (caIMDPs) generalize finite-state MDPs with interval-valued transition probabilities and continuous action-spaces.

DEFINITION 2.1 (caIMDP). *A continuous-action interval Markov decision process (caIMDP) is a tuple $I = (Q, \mathcal{A}, \check{P}, \hat{P}, R)$, where*

- Q is the finite set of states,
- $\mathcal{A} \subset \mathbb{R}^{n_{\mathcal{A}}}$ is the set of actions,
- $\check{P} : Q \times \mathcal{A} \times Q \rightarrow [0, 1]$, where $\check{P}(q, a, q')$ is the lower bound on the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in \mathcal{A}$,
- $\hat{P} : Q \times \mathcal{A} \times Q \rightarrow [0, 1]$, where $\hat{P}(q, a, q')$ is the upper bound on the transition probability from state $q \in Q$ to state $q' \in Q$ under action $a \in \mathcal{A}$,
- $R : Q \rightarrow \mathbb{R}_{\geq 0}$ is a bounded state-dependent reward function.

For all $q, q' \in Q$ and $a \in \mathcal{A}$, it holds that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and $\sum_{q' \in Q} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q, a, q')$. Given a state $q \in Q$ and an action $a \in \mathcal{A}$, a transition probability distribution $p_{q,a} : Q \rightarrow [0, 1]$ is called *feasible* if $\check{P}(q, a, q') \leq p_{q,a}(q') \leq \hat{P}(q, a, q')$ for all $q' \in Q$. The set of all feasible distributions for the state-action pair (q, a) is denoted by $\Gamma_{q,a}$. We define $\Gamma = \{\Gamma_{q,a} : (q, a) \in Q \times \mathcal{A}\}$ to be the set of all feasible distributions for all state-action pairs.

A path of a caIMDP is a sequence of states and actions $\omega = (q_0, a_0), (q_1, a_1), \dots$, where $(q_i, a_i) \in Q \times \mathcal{A}$, and we denote $\omega(k) = (q_k, a_k)$ and $\omega^q(k) = q_k$, for $k = 0, 1, \dots$. We denote the set of all finite paths by $\text{Paths}^{\text{fin}}$. For a path $\omega \in \text{Paths}^{\text{fin}}$, we denote its last element by $\omega(\text{end})$ (i.e., for an $(N+1)$ -length path $\omega(\text{end}) = \omega(N) = (q_N, a_N)$).

To describe the evolution of a caIMDP-path, we also need to introduce the notions of *policy* and *adversary*.

DEFINITION 2.2 (POLICY). *For a caIMDP, a policy $\pi : \text{Paths}^{\text{fin}} \times Q \rightarrow \mathcal{A}$ is a function that, given a finite path $\omega \in \text{Paths}^{\text{fin}}$ and a state $q \in Q$, returns an action. The set of all policies is denoted by Π . A policy is Markov, if the choice of action depends only on q and on the path's length.*

DEFINITION 2.3 (ADVERSARY). *For a caIMDP, an adversary is a function $\xi : \text{Paths}^{\text{fin}} \rightarrow \Gamma$. Given a finite path $\omega \in \text{Paths}^{\text{fin}}$, it returns a feasible distribution $p_{q,a} \in \Gamma_{q,a}$, where $(q, a) = \omega(\text{end})$. The set of all adversaries is denoted by Ξ . An adversary is Markov, if the choice of a feasible distribution depends only on $\omega(\text{end})$ and the path's length.*

Given a policy π and an adversary ξ , a caIMDP-path evolves as follows. At time i , given the finite path $(q_0, a_0), \dots, (q_{i-1}, a_{i-1})$ and

the current state q_i , the policy π chooses the action a_i . Then, the adversary ξ , given the path $(q_0, a_0), \dots, (q_i, a_i)$, chooses a feasible distribution $p_{q_i, a_i} \in \Gamma_{q_i, a_i}$. The next state of the path q_{i+1} is sampled randomly from p_{q_i, a_i} .

In words, a policy is a control strategy, that, at each time, decides the control action based on the history of the path. The adversary represents the environment: once the control action is taken, the adversary resolves the uncertainty stemming from the transition probability intervals, thus fixing the -still stochastic- environment for the current time step. Notice that, given a specific policy π and adversary ξ , the cAIMDP collapses to a time-varying Markov chain. Thus, given π, ξ , an initial state q_0 , and a horizon N , a probability measure is uniquely defined over $(N+1)$ -length paths $\omega \in \text{Paths}^{\text{fin}}$ [1].

REMARK 1. For ease of presentation, we constrain the definitions of policy and adversary to deterministic ones. In fact, they can be random, but, as Proposition 2.4 shows, the optimal rewards that we study are achieved by deterministic, Markov policies and adversaries; thus, it suffices to consider only deterministic ones.

2.2 Optimal policies, optimal rewards and value iteration

In what follows, we implicitly assume that all mentioned max min quantities (e.g., the ones in (1) and (3b)) are well-defined. Numerous sets of assumption can be employed to impose this: e.g., as a consequence of Proposition 2.4 below, it suffices that \mathcal{A} is compact and \hat{P}, \hat{P} are continuous functions of the action variables a .

Given π, ξ, N , an initial state $q_0 \in \mathcal{Q}$, and a factor $\gamma \geq 0$, we define the so-called *expected cumulative reward*:

$$\mathcal{R}_{\pi, \xi}^N(q_0) = \mathbb{E}^{q_0} \left[\sum_{i=0}^N \gamma^i R(\omega^Q(i)) \mid \pi, \xi \right]$$

where the expectation is taken w.r.t. the probability measure over $(N+1)$ -length paths in $\text{Paths}^{\text{fin}}$ starting from state q_0 . Notice that R , as defined in Definition 2.1, depends only on states, not actions.

In this work, as done commonly in the literature [8], we consider the *robust control problem*: finding a policy that maximizes the expected cumulative reward generated by the worst possible adversary (and the policy itself). In other words, we study the following max-min problem²:

PROBLEM (ROBUST CONTROL). Given a cAIMDP \mathcal{I} , a factor $\gamma \geq 0$, a horizon N , and an initial state q_0 , solve:

$$\mathcal{R}_{\star}^N(q_0) \equiv \max_{\pi \in \Pi} \min_{\xi \in \Xi} \mathcal{R}_{\pi, \xi}^N(q_0). \quad (1)$$

A policy π_{\star} that solves (1), i.e.,

$$\pi_{\star} \in \arg \max_{\pi \in \Pi} \min_{\xi \in \Xi} \mathcal{R}_{\pi, \xi}^N(q_0),$$

is called a (*pessimistically*)³ *optimal policy*. The optimal value $\mathcal{R}_{\star}^N(q_0)$ of (1) is called (*pessimistic*) *optimal (expected cumulative) reward*.

²While this work focuses on the max-min problem, max-max, min-max, and min-min problems can be addressed similarly.

³"Pessimistically", as it optimizes the reward generated by the *worst* adversary. If we considered the *best* adversary instead, which would amount to a max-max problem, the policy would be called *optimistically* optimal.

Solving (1) is computing an optimal policy π_{\star} and the optimal reward.

For MDPs and for IMDPs with discrete action-spaces, it is well-known that (1) can be solved by an iterative scheme called *value iteration*. The following proposition extends this to cAIMDPs as well:

PROPOSITION 2.4. Consider a cAIMDP \mathcal{I} . For any $j = 0, \dots, N$, the following holds:

$$\mathcal{R}_{\star}^j(q_0) = V_{N-j}(q_0) \quad (2)$$

where $V_{N-j}(q)$ is defined, for any $q \in \mathcal{Q}$, through the following iteration:

$$V_N(q) = R(q) \quad (3a)$$

$$V_{k-1}(q) = R(q) + \gamma \max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q') \quad (3b)$$

for $k = N, \dots, 1$.

PROOF. See Section 7.1. □

The process of solving (3) for all iterations is called *value iteration* and the obtained function $V_0(\cdot)$ is called *value function*. A direct corollary of Proposition 2.4, is that there exist Markov policies (and adversaries) achieving the optimal reward, defined as follows:

COROLLARY 2.5 (TO PROPOSITION (2.4)). Any Markov policy satisfying:

$$\pi_{\star}(q, N-k-1) \in \arg \max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q')$$

is optimal, where $\pi_{\star}(q, i)$ denotes the action taken if the current state is q and the current time is i .

REMARK 2. Note that, while in this work we only consider cumulative rewards, similarly to IMDPs with discrete action-spaces, our results can easily be extended to more general properties, such as bounded-time fragments of PCTL [12] or LTL [10]. Infinite-horizon (unbounded-time) extensions are also possible, but particular care is needed to guarantee convergence of the value iteration of Proposition 2.4 due to the non-finiteness of \mathcal{A} .

3 EFFICIENT FORMULATIONS OF VALUE ITERATION OVER cAIMDPs

By Proposition 2.4 and Corollary 2.5, we see that to compute the optimal reward and an optimal policy it suffices to solve (3) for all iterations $k = N, \dots, 1$. However, solving (3b) is not straightforward, as it includes solving (for each q) the max-min optimization problem $\max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} p^T V_k$, where we, abusively, denote $V_k = [V_k(q_1) \ \dots \ V_k(q_{|\mathcal{Q}|})]^T$, $p = [p(q_1) \ \dots \ p(q_{|\mathcal{Q}|})]^T$. This max-min problem can be written as:

$$\begin{aligned} & \max_{a \in \mathcal{A}} \min_{p \in \mathbb{R}^{|\mathcal{Q}|}} p^T V_k \\ & \text{s.t. : } \begin{aligned} & \hat{P}(q, a, q_i) - p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \\ & -\hat{P}(q, a, q_i) + p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \end{aligned} \quad (\text{Mm}) \\ & \sum_{i=1}^{|\mathcal{Q}|} p_i - 1 = 0 \end{aligned}$$

where p_i denotes the i -th component of vector p , and the constraints guarantee that p is a probability distribution (equality constraint) and belongs in $\Gamma_{q,a}$ (inequality constraints).

In this section, we decompose problem (Mm) to a number of simpler problems, and identify special cases of \mathcal{A} , $\check{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ in which problem (Mm) can be solved *efficiently*. First, we state and prove this paper's main result (Theorem 3.2): that problem (Mm) is equivalent to solving the $|\mathcal{Q}|$ simpler maximization problems in (MP). This gives rise to Algorithm 1, which performs value iteration over cAIMDPs, where, instead of solving (Mm), it solves (MP). Then, the simple form of the maximization problems in (MP) allows us to identify special cases in which they, and thus value iteration, are tractable. Furthermore, we show that when \mathcal{A} is polytopic and $\check{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ are either linear or convex and concave, respectively, performing synthesis over a discrete-action IMDP with its action set being $\text{ver}(\mathcal{A})$ is sufficient for optimality.

3.1 Decomposing the max-min problem to $|\mathcal{Q}|$ max problems

To prove the main result, first we transform the original max-min problem (Mm) to a maximization problem, by employing duality:

PROPOSITION 3.1. *Consider the optimization problem:*

$$\begin{aligned} \max_{a, \lambda_L, \lambda_U, v} \quad & \sum_{i=1}^{|\mathcal{Q}|} \lambda_{L_i} \check{P}(q, a, q_i) - \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} \hat{P}(q, a, q_i) - v \\ \text{s.t. :} \quad & a \in \mathcal{A}, \lambda_L \geq 0, \lambda_U \geq 0 \\ & V_k - \lambda_L + \lambda_U + v\mathbf{1} = 0 \end{aligned} \quad (4)$$

where $\lambda_L \in \mathbb{R}^{|\mathcal{Q}|}$, $\lambda_U \in \mathbb{R}^{|\mathcal{Q}|}$, $v \in \mathbb{R}$ and $\lambda_{L_i}, \lambda_{U_i}$ denote the i -th component of λ_L and λ_U , respectively. If $(a_\star, \lambda_{L_\star}, \lambda_{U_\star}, v_\star)$ solves (4), then it also solves (Mm). Moreover, the optimal values of (Mm) and (4) coincide.

PROOF. See Section 7.2. \square

The variables λ_L, λ_U and v are the so-called *Lagrange multipliers*. Generally, even though just a max problem instead of a max-min problem, optimization problem (4) is not easy to work with directly. That is because of the product terms $\lambda_{L_i} \check{P}(q, a, q_i)$ and $\lambda_{U_i} \hat{P}(q, a, q_i)$ (e.g., even when $\check{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ are linear on a , problem (4) is, generally, nonconvex). Nonetheless, by working around its special structure, we are able to dissect it into $|\mathcal{Q}|$ simpler max problems, which constitutes this work's main result:

THEOREM 3.2. *Assume, without loss of generality, that $V_k(q_i) \geq V_k(q_{i+1})$ for all $i = 1, 2, \dots, |\mathcal{Q}| - 1$, i.e., that V_k is in descending order. Consider the following optimization problem:*

$$\begin{aligned} \max_{j=1,2,\dots,|\mathcal{Q}|} \max_a \quad & \sum_{i=1}^{j-1} (V_k(q_i) - V_k(q_j)) \check{P}(q, a, q_i) \\ & + \sum_{i=j+1}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_j)) \hat{P}(q, a, q_i) + V_k(q_j) \\ \text{s.t. :} \quad & a \in \mathcal{A} \end{aligned} \quad (\text{MP})$$

If a_\star solves (MP), then it solves (Mm). Moreover, the optimal values of (MP) and (Mm) coincide.

PROOF. See Section 7.3. \square

By employing Theorem 3.2, we can devise an algorithm (Algorithm 1) to solve value iteration over cAIMDPs, where, instead of solving the stringent (Mm), we solve (MP), which can be solved efficiently in many cases of interest, as it is revealed in the next section. Observe how in each iteration of value iteration, we sort the value-function vector in decreasing order, thus imposing the assumption of Theorem 3.2 (hence why it is written that the assumption is without loss of generality).

Algorithm 1 Value Iteration on cAIMDPs

```

1: for  $q \in \mathcal{Q}$  do
2:    $V_N(q) \leftarrow R(q)$ 
3: end for
4: for  $k = N, \dots, 1$  do
5:    $[V_{\text{sorted}}, \text{indices}_{\text{sorted}}] = \text{sort}(V_k)$   $\triangleright$  sorts  $V_k$  in descending order and stores it in  $V_{\text{sorted}}$ .  $\text{indices}_{\text{sorted}}$  describes the arrangement of the elements of  $V_k$  into  $V_{\text{sorted}}$ , i.e.  $V_k = V_{\text{sorted}}(\text{indices}_{\text{sorted}})$ .
6:   for  $q$  in  $\text{indices}_{\text{sorted}}$  do
7:      $V_{k-1}(q) = R(q) + \gamma^* \text{Solve}(\text{MP})$   $\triangleright$  where in (MP) we replace  $V_k$  with  $V_{\text{sorted}}$ .
8:   end for
9: end for
```

It remains to deduce when is it easy to perform line 7 of Algorithm 1, i.e. to solve (MP). Notice how (MP) is in a much simpler form to analyze than the primal max-min problem (Mm): a) they are just max problems, b) their only decision variable is the action a and the constraint set is \mathcal{A} , and c) the functions $\check{P}(q, a, q_i)$ and $\hat{P}(q, a, q_i)$ appear linearly, which facilitates convexity arguments. It is this simple form that allows us to deduce when they can be solved efficiently, in the following section.

REMARK 3. *Notice that the derived maximization problems (MP) and value-iteration algorithm share similarities with the algorithm proposed in [8] to solve value iteration over discrete-action IMDPs. First, they both sort the value-function vector in each iteration. Furthermore, in both algorithms, given a specific action a , a combination of additions of lower and upper bounds on transition probabilities is found, which is optimal w.r.t. the ordered value-function vector. The number of possible such combinations is $|\mathcal{Q}|$ in both cases. Nevertheless, contrary to here, in [8] there is a stopping criterion, based on which the algorithm may terminate before computing all $|\mathcal{Q}|$ combinations. The existence of such a stopping criterion for problems (MP) (i.e., for the continuous-action case) is an open question, and a potential affirmative answer would significantly reduce computations.*

3.2 When can value iteration for cAIMDPs be solved efficiently?

By inspecting (MP), we are able to identify cases where value iteration for cAIMDPs can be solved efficiently.

The linear case. In the case of polytopic \mathcal{A} , and linear $\check{P}(\cdot, a, \cdot)$, $\hat{P}(\cdot, a, \cdot)$, solving problem (MP) amounts to solving $|\mathcal{Q}|$ linear programs (LPs):

COROLLARY 3.3 (TO THEOREM 3.2). *If \mathcal{A} is a polytope and both $\check{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ are linear on a , then the maximization problems in (MP) are LPs.*

Moreover, notice that an optimal a_\star lies on the vertices of \mathcal{A} , since the problems in (MP) are LPs with \mathcal{A} as the constraint set. Thus, another corollary of Theorem 3.2 is that, instead of solving (MP) in line 7 of Algorithm 1, we could solve the original max-min problem (Mm) only for $a \in \text{ver}(\mathcal{A})$, which amounts to solving $|\text{ver}(\mathcal{A})|$ LPs:

COROLLARY 3.4 (TO THEOREM 3.2). *If \mathcal{A} is a polytope and both $\check{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ are linear on a , solving the max-min problem (Mm) amounts to solving the following $|\text{ver}(\mathcal{A})|$ minimization LPs:*

$$\begin{aligned} \max_{a \in \text{ver}(\mathcal{A})} \min_{p \in \mathbb{R}^{|\mathcal{Q}|}} p^\top V_k \\ \text{s.t. : } \quad & \check{P}(q, a, q_i) - p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \\ & -\hat{P}(q, a, q_i) + p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \quad (5) \\ & \sum_{i=1}^{|\mathcal{Q}|} p_i = 1 = 0 \end{aligned}$$

REMARK 4. *Equivalently, Corollary 3.4 guarantees that it is sufficient to construct a discrete-action IMDP with its action set being $\text{ver}(\mathcal{A})$, and solve the synthesis problem over the discrete-action IMDP.*

Corollary 3.3 amounts to $|\mathcal{Q}|$ LPs with $n_{\mathcal{A}}$ decision variables and \mathcal{A} as a constraint set. On the other hand, the approach of Corollary 3.4 (or Remark 4) amounts to $|\text{ver}(\mathcal{A})|$ LPs with $|\mathcal{Q}|$ decision variables. Note, however, that to solve the LPs of Corollary 3.4, one can employ the algorithm proposed in [8], which has been shown to be very efficient (see also Remark 3). Intuitively, when $|\text{ver}(\mathcal{A})| \gg |\mathcal{Q}|$, then the approach of Corollary 3.3 might be preferred against the one of Corollary 3.4, and vice-versa. Nevertheless, a thorough experimental study has to be conducted, to decide when to use each.

REMARK 5. *Max-min problems with linear objective functions and constraints, in their general form, are studied in [6]. It is therein where it is proven that an optimal solution is attained at the vertices of the constraint set. Moreover, an algorithm is proposed that decomposes the original max-min problem to a number of LPs, which in our case would be $O(|\mathcal{Q}|^2)$ LPs. Here, we are able to derive more computationally efficient results, because we exploit our specific problem's special structure (e.g. the knowledge that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and $\sum_{q' \in \mathcal{Q}} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in \mathcal{Q}} \hat{P}(q, a, q')$).*

The concave/convex case. When \mathcal{A} is convex, $\check{P}(\cdot, a, \cdot)$ is concave and $\hat{P}(\cdot, a, \cdot)$ is convex on a , then the maximization problems in (MP) are convex programs (CPs):

COROLLARY 3.5 (TO THEOREM 3.2). *If \mathcal{A} is convex, $\check{P}(\cdot, a, \cdot)$ is concave on a and $\hat{P}(\cdot, a, \cdot)$ is convex on a , then the maximization problems in (MP) are convex programs (CPs).*

In such case, standard convex optimization algorithms (see [2]) can be employed to solve (MP) and perform value iteration over the given cAIMDP.

The convex/concave case. When \mathcal{A} is a polytope, $\check{P}(\cdot, a, \cdot)$ is convex and $\hat{P}(\cdot, a, \cdot)$ is concave on a , we have the following:

COROLLARY 3.6 (TO THEOREM 3.2). *If \mathcal{A} is a polytope, $\check{P}(\cdot, a, \cdot)$ is convex on a and $\hat{P}(\cdot, a, \cdot)$ is concave on a , then the maximization problems in (MP) amount to maximizing a convex function over a polytope. Thus, an optimal solution lies at the constraint set's vertices.*

PROOF OF COROLLARY 3.6. The fact that maximizing a convex function over a polytope attains a solution at the vertices is well-known; e.g., see [18, Theorem 32.2]. \square

In this case, we can either solve all $|\mathcal{Q}|$ problems in (MP), by evaluating the objective function on the vertices of \mathcal{A} , or we can solve the $|\text{ver}(\mathcal{A})|$ LPs (5), which, again, is equivalent to building a discrete-action IMDP, as explained in Remark 4. Either way, when $|\text{ver}(\mathcal{A})|$ is small, value iteration can be solved quickly.

4 NUMERICAL EXAMPLE

Here, we demonstrate this work's theoretical results on a numerical example. We consider the concave/convex case and we compare the results obtained by synthesizing over a cAIMDP with the ones obtained by discrete-action IMDPs (discretized versions of the cAIMDP). We constructed a randomly generated cAIMDP $\mathcal{I} = (\mathcal{Q}, \mathcal{A}, \check{P}, \hat{P}, R)$ with $|\mathcal{Q}| = 25$ states, concave/convex transition bounds and a 3-dimensional action space ($n_{\mathcal{A}} = 3$). The action set is the following cylinder:

$$\mathcal{A} = \{a \in \mathbb{R}^3 : (a_1 - 0.5)^2 + (a_2 - 0.5)^2 \leq 0.2, a_3 \in [0, 1]\}$$

The horizon $N = 10$ and $\gamma = 1$.

We computed the robust optimal policy and expected cumulative reward $\mathcal{R}_\star(q)$ (for all q) over the cAIMDP, via convex programming, as instructed by Corollary 3.5. Moreover, we obtained policies and the corresponding rewards $\mathcal{R}_s(q)$ (for all q) by performing synthesis over discrete-action IMDPs, which are discretized versions of the original cAIMDP: a discrete-action IMDP is constructed as $\mathcal{I}_{discr}^s = (\mathcal{Q}, \hat{\mathcal{A}}_s, \check{P}, \hat{P}, R)$, where $\hat{\mathcal{A}}_s \subseteq \mathcal{A}$ is a finite set of s actions randomly sampled from the continuous action set \mathcal{A} ; that is, $|\hat{\mathcal{A}}_s| = s$. Specifically, we computed $\mathcal{R}_s(q)$ for different values of s . Note that, since randomness is involved in the experiments⁴, we performed them multiple times⁵. Moreover, to perform synthesis over the discrete-action IMDPs (i.e. to solve the s associated LPs), we used the algorithm of [8]. Finally, optimality tolerance for convex programming was set to 10^{-4} .

Figure 1 depicts the optimal reward $\mathcal{R}_\star(q)$ obtained from the cAIMDP and rewards $\mathcal{R}_s(q)$ obtained from some of the discrete-action IMDPs, for different values of s . As expected, in all cases $\mathcal{R}_\star(q) \geq \mathcal{R}_s(q)$, since $\mathcal{R}_\star(q)$ is the optimal reward, whereas $\mathcal{R}_s(q)$ is suboptimal. Moreover, Table 1 shows the average percentage of suboptimality of $\mathcal{R}_s(q)$, i.e. the quantity:

$$100 \cdot \max_q \frac{\mathcal{R}_\star(q) - \mathcal{R}_s(q)}{\mathcal{R}_\star(q)} \%,$$

for all different values of s . Furthermore, it shows the average computation time for each experiment. We observe that performing

⁴In the case of value iteration over the cAIMDP with convex programming, the initial condition for the optimization is picked randomly, which affects the total computation time. In the case of discrete-action IMDPs, the sampled actions affect both the computation time and, especially, the reward $\mathcal{R}_s(q)$.

⁵Value iteration over the cAIMDP has been performed 3 times. The experiments on discrete-action IMDPs for $s = 1, 8, 27, 64, 125$ have been performed [20, 11, 9, 6, 3] times, respectively.

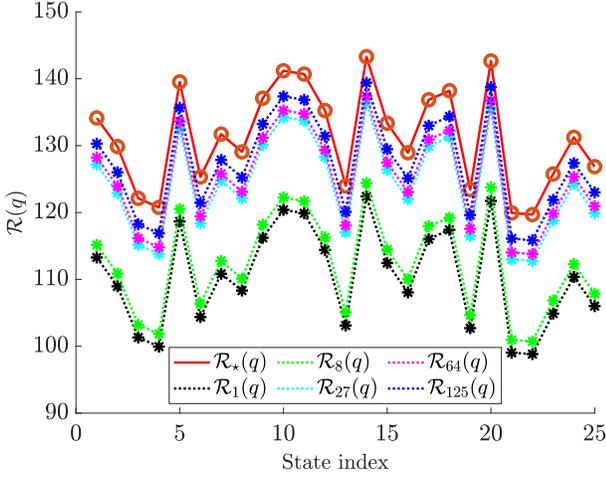


Figure 1: The optimal reward $\mathcal{R}_*(q)$ (solid line) obtained from the cAIMDP, and rewards $\mathcal{R}_s(q)$ (dashed lines) obtained from discrete-action IMDPs.

Table 1: The first five lines report average cpu times and average suboptimality percentages for value iteration over discrete-action IMDPs obtained by randomly sampling s actions from the continuous action-space. The last line reports the corresponding results for value iteration over the cAIMDP.

#samples s	CPU Time (s)	$100 \cdot \max_q \frac{\mathcal{R}_*(q) - \mathcal{R}_s(q)}{\mathcal{R}_*(q)} \%$
1	88	16.51%
8	668	7.86%
27	2170	5.17%
64	6225	4.59%
125	9736	4.32%
cAIMDP	2212	0%

synthesis on the cAIMDP is approximately as expensive as synthesizing over a discrete-action IMDP with only 27 actions (only 3 points per dimension of the action space). That is while $\mathcal{R}_{27}(q)$ is significantly suboptimal with an average suboptimality of 5.17%. In fact, even with 125 sampled actions (5 points per dimension), the suboptimality percentage is considerable. From this experiment, we deduce that synthesizing over a cAIMDP is not only the way to optimality, but also it is computationally efficient even compared to discrete-action IMDPs. In fact, for even higher-dimensional action spaces, it is expected that it is the only tractable choice.

5 DISCUSSION: cAIMDPs FOR CONTROL SYNTHESIS FOR STOCHASTIC SYSTEMS

We now, briefly, describe how this paper’s results can be employed for control synthesis for stochastic systems.

5.1 Abstracting stochastic systems via cAIMDPs

Consider the controlled stochastic system:

$$x_{k+1} = f(x_k, a_k, w_k) \quad (6)$$

where $x_k \in X$ is the state of the system at time k , $a_k \in \mathcal{A}$ is the action at time k and $w_k \in \mathbb{R}^{n_w}$ are i.i.d random variables with a known probability distribution p_w . The paths of the system are again sequences $\omega = (x_1, a_1), (x_2, a_2), \dots$, and we denote the system’s set of finite paths by $Paths_{sys}^{fin}$. Assume a state-dependent reward $R_{sys} : X \rightarrow \mathbb{R}_{\geq 0}$. Given a specific control policy $\chi : Paths_{sys}^{fin} \times X \rightarrow \mathcal{A}$, a horizon N , an initial condition $x_0 \in \mathbb{R}^{n_x}$ and a factor $\gamma \geq 0$, the expected cumulative reward over the system’s trajectories is defined as:

$$\mathbb{E}_{\mathbb{P}_w}^{x_0} \left[\sum_{i=0}^N \gamma^i R_{sys}(\omega^Q(i)) \mid \chi \right]$$

where \mathbb{P}_w is a probability measure over the system’s paths induced by p_w and the initial condition x_0 . This expected reward constitutes a metric describing the system’s behaviors under policy χ (e.g., a reachability or safety property, a quantitative measure such as convergence speed, etc.). The objective is to synthesize a policy χ_* that maximizes the expected reward.

In the literature of IMDP-abstractions for control synthesis for stochastic systems [5, 12, 14], the above problem is solved by abstracting the system via an appropriate IMDP, finding an optimal policy over the constructed IMDP and mapping this policy back to the original system. Nevertheless, as already mentioned in the introduction section, the action set \mathcal{A} is, generally, assumed discrete a-priori. To the best of our knowledge, the only exception is [5], which keeps the action-space continuous, but the max-min problem of value iteration is solved using heuristics.

With the results derived in this work, we can avoid unnecessary discretizations of the action set or heuristics with no formal guarantees. To abstract a system (6) by a cAIMDP, the state space X is partitioned into $|\mathcal{Q}|$ sets $q_i \subseteq X$, which represent the cAIMDP’s states, and the transition bounds are defined as follows for all $q_i, q_j \in \mathcal{Q}$:

$$\hat{P}(q_i, a, q_j) \leq \inf_{x \in q_i} \mathbb{P}_w(f(x, a, w) \in q_j), \quad (7)$$

$$\hat{P}(q_i, a, q_j) \geq \sup_{x \in q_i} \mathbb{P}_w(f(x, a, w) \in q_j) \quad (8)$$

The cAIMDP-reward is defined by $R(q) = \inf_{x \in q} R_{sys}(q)$. By extending [4, Theorem 4.1] to cAIMDPs, we can show that the optimal expected cumulative cAIMDP-reward lower-bounds the optimal expected cumulative reward over the system’s trajectories:

$$\max_{\pi} \min_{\xi} \mathcal{R}_{\pi, \xi}^N(q_0) \leq \max_{\chi} \mathbb{E}_{\mathbb{P}_w}^{x_0} \left[\sum_{i=0}^N \gamma^i R_{sys}(\omega^Q(i)) \mid \chi \right] \quad (9)$$

where $x_0 \in q_0$. Furthermore, the optimal reward obtained via the cAIMDP is larger than or equal to an optimal reward obtained via a discrete-action IMDP, as its action set is a subset of the cAIMDP’s action set.

In order to enable the results presented in Section 3.2 and render value iteration easy-to-solve, the bounds $\hat{P}(\cdot, a, \cdot)$ and $\hat{P}(\cdot, a, \cdot)$ have to be (piecewise) linear or concave/convex or convex/concave on the action variables. For instance, this is the case for a large class of switching diffusion models where the continuous action

controls the probability of switching between different modes of the system [19]. In more general cases, the functions $\inf_{x \in q_i}$ and $\sup_{x \in q_i} \mathbb{P}_w(f(x, a, w) \in q_j)$ may need to be lower-/upper-bounded via (piecewise) linear, convex or concave functions.

5.2 Suboptimality bounds

A policy derived through a cAIMDP-abstraction is generally sub-optimal for system (6), even though the algorithms presented here do compute a policy that is optimal w.r.t. the given cAIMDP. That is because the cAIMDP is an abstraction of the original stochastic system in (6), and not an equivalent representation. Nevertheless, cAIMDPs can be used to obtain so-called *suboptimality bounds*: bounds on the difference between the optimal reward computed via the cAIMDP and the true optimal reward:

$$\max_{\chi} \mathbb{E}_{\mathbb{P}_w}^{x_0} \left[\sum_{i=0}^N \gamma^i R_{sys}(\omega^Q(i)) \mid \chi \right] - \max_{\pi} \min_{\xi} \mathcal{R}_{\pi, \xi}^N(q_0)$$

This quantifies how suboptimal a policy obtained via a cAIMDP is; if the obtained result is not satisfactory, then the cAIMDP-abstraction has to be refined with a finer partition of the state-space X or with tighter transition bounds.

Again, by extending [4, Theorem 4.1] to cAIMDPs, it can be shown that, if we define the reward function of the cAIMDP as $R'(q) = \sup_{x \in q} R_{sys}(q)$ (previously, we used \inf), the *optimistically optimal* cAIMDP reward upper bounds the true optimal one:

$$\max_{\chi} \mathbb{E}_{\mathbb{P}_w}^{x_0} \left[\sum_{i=0}^N \gamma^i R_{sys}(\omega^Q(i)) \mid \chi \right] \leq \max_{\pi, \xi} \mathcal{R}_{\pi, \xi}^N(q_0)$$

where $\mathcal{R}_{\pi, \xi}^N(q_0) = \mathbb{E}^{q_0} \left[\sum_{i=0}^N \gamma^i R'(\omega^Q(i)) \mid \pi, \xi \right]$, and the expectation is w.r.t. the probability measure over the paths of the cAIMDP. Thus, $\max_{\pi, \xi} \mathcal{R}_{\pi, \xi}^N(q_0) - \max_{\pi} \min_{\xi} \mathcal{R}_{\pi, \xi}^N(q_0)$ provides a suboptimality bound.

The above requires solving $\max_{\pi, \xi} \mathcal{R}_{\pi, \xi}^N(q_0)$, which amounts to a value iteration, in which, instead of a max-min problem, the following max problem needs to be solved:

$$\begin{aligned} \max_{a \in \mathcal{A}, p \in \mathbb{R}^{|\mathcal{Q}|}} \quad & p^\top V_k \\ \text{s.t.} \quad & \check{P}(q, a, q_i) - p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \\ & -\hat{P}(q, a, q_i) + p_i \leq 0, \quad i = 1, 2, \dots, |\mathcal{Q}|, \\ & \sum_{i=1}^{|\mathcal{Q}|} p_i - 1 = 0 \end{aligned} \quad (10)$$

Referring to the three cases distinguished in Section 3.2, problem (10): a) in the linear case, is also an LP, b) in the convex/concave case, is a CP, but c) in the concave/convex case, its constraint set is nonconvex: the intersection of sub-level sets of convex and concave functions. Hence, in the linear and the convex/concave case, problem (10) can be solved efficiently, to generate a tight suboptimality bound. In the concave/convex case, we could derive an upper bound on the optimal value of problem (10), e.g., by using an SMT solver [7].

6 CONCLUSION

We have introduced continuous-action interval Markov decision processes (cAIMDPs), and studied value iteration for optimizing pessimistic expected cumulative rewards. Specifically, we have shown that the max-min problem associated to value iteration can be decomposed to $|\mathcal{Q}|$ maximization problems. The simple form of these max problems allowed us to distinguish cases where they, and thus value iteration, can be solved efficiently. It also provided us with further interesting insights, such as cases where synthesis over discrete-action and continuous-action IMDPs is equivalent. These results have been demonstrated on a numerical example.

Furthermore, we have, briefly, discussed how cAIMDPs can be employed for control synthesis for stochastic systems. Specifically, we have given guidelines on how to construct the cAIMDP-abstraction such that it falls in one of the aforementioned easy-to-solve cases. We have, also, shown how to obtain suboptimality bounds on the policy generated by the cAIMDP-abstraction.

Future work includes: i) studying convergence of the value iteration for infinite horizons, ii) addressing rewards that depend both on the state and the action, iii) investigating if the $|\mathcal{Q}|$ maximization problems can be simplified even further, and iv) a thorough study on cAIMDP-abstractions and on obtaining suboptimality bounds.

7 TECHNICAL PROOFS

7.1 Proof of Proposition 2.4

SKETCH OF PROOF OF PROPOSITION 2.4. By the transformation presented in [16, Appendix A] for finite-action IMDPs, the robust control problem can be equivalently transformed into a turn-based zero-sum game. Consequently, similarly to [16], we can employ the same reasoning as in the proof of [17, Theorem 1], to show that:

$$\mathcal{R}_{\star}^j(q_0) = V_{N-j}(q_0)$$

where V_{N-j} comes from the following Bellman recursion:

$$V_N(q) = R(q)$$

$$V_{k-1}(q) = R(q) + \gamma \max_{\mu \in M} \int_{\mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q') \mu(a) da(q')$$

where M is the set of probability distributions⁶ over \mathcal{A} . We now notice that for any $\mu \in M$:

$$\int_{\mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q') \mu(a) da \leq \max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q')$$

This implies that the optimal choice is a Dirac distribution centered in any $a \in \arg \max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q')$. Consequently, we can replace the maximization over M with a maximization over \mathcal{A} , i.e.:

$$\begin{aligned} \max_{\mu \in M} \int_{\mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q') \mu(a) da &= \\ \max_{a \in \mathcal{A}} \min_{p \in \Gamma_{q,a}} \sum_{q' \in \mathcal{Q}} p(q') V_k(q') & \end{aligned}$$

□

⁶Maximization takes place over the set M of probability distributions on \mathcal{A} , because, as mentioned in Remark 1, policies can be probabilistic. Nonetheless, it is shown that optimal policies are deterministic (or, in other words, Dirac distributions).

7.2 Proof of Proposition 3.1

First, let us prove the following lemma:

LEMMA 7.1. Consider the following max-min problem:

$$\max_{a \in \mathcal{A}} \min_{p \in P, f(a,p) \leq 0} c^\top p$$

Assume that it is feasible and bounded, and that strong duality holds for the inner minimization problem and its Lagrangian dual (see, e.g., [2]):

$$\forall a \in \mathcal{A} : \min_{p \in P, f(a,p) \leq 0} c^\top p = \max_{\lambda \geq 0, v} g(a, \lambda, v) \quad (11)$$

Then $\max_{a \in \mathcal{A}} \min_{p \in P, f(a,p) \leq 0} c^\top p = \max_{a \in \mathcal{A}, \lambda \geq 0, v} g(a, \lambda, v)$. Moreover:

$$(a_\star, \lambda_\star, v_\star) \in \arg \max_{a \in \mathcal{A}, \lambda \geq 0, v} g(a, \lambda, v) \implies a_\star \in \arg \max_{a \in \mathcal{A}} \min_{p \in P, f(a,p) \leq 0} c^\top p$$

PROOF OF LEMMA 7.1. We have that:

$$\forall a \in \mathcal{A} : g(a_\star, \lambda_\star, v_\star) \geq \max_{\lambda \geq 0, v} g(a, \lambda, v) = \min_{p \in P, f(a,p) \leq 0} c^\top p \quad (12)$$

where, for the inequality we used the definition of $(a_\star, \lambda_\star, v_\star)$ and for the equality we used (11). Moreover, we have:

$$g(a_\star, \lambda_\star, v_\star) = \max_{\lambda \geq 0, v} g(a_\star, \lambda, v) = \min_{p \in P, f(a_\star, p) \leq 0} c^\top p \quad (13)$$

where for the first equality we used the definition of $(a_\star, \lambda_\star, v_\star)$ and for the second one we used (11). Replacing (13) into (12) we get:

$$\forall a \in \mathcal{A} : \min_{p \in P, f(a_\star, p) \leq 0} c^\top p \geq \min_{p \in P, f(a, p) \leq 0} c^\top p$$

Since the above holds for all $a \in \mathcal{A}$ and since $a_\star \in \mathcal{A}$, then:

$$\min_{p \in P, f(a_\star, p) \leq 0} c^\top p = \max_{a \in \mathcal{A}} \min_{p \in P, f(a, p) \leq 0} c^\top p$$

which, since the left-hand side is equal to $g(a_\star, \lambda_\star, v_\star)$, proves that $\max_{a \in \mathcal{A}} \min_{p \in P, f(a,p) \leq 0} c^\top p = \max_{a \in \mathcal{A}, \lambda \geq 0, v} g(a, \lambda, v)$. Moreover, it implies that $a_\star \in \arg \max_{a \in \mathcal{A}} \min_{p \in P, f(a,p) \leq 0} c^\top p$. The proof is complete. \square

Now, we are ready to prove Proposition 3.1.

PROOF OF PROPOSITION 3.1. The Lagrange dual function that corresponds to the inner minimization problem of (Mm) is:

$$\begin{aligned} g(a, \lambda_L, \lambda_U, v) &= \inf_{p \in \mathbb{R}^{|\mathcal{Q}|}} \left[p^\top V_k + \sum_{i=1}^{|\mathcal{Q}|} \lambda_{L_i} (\check{P}(q, a, q_i) - p_i) + \right. \\ &\quad \left. + \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} (-\hat{P}(q, a, q_i) + p_i) + v \sum_{i=1}^{|\mathcal{Q}|} p_i - v \right] \\ &= \sum_{i=1}^{|\mathcal{Q}|} \lambda_{L_i} \check{P}(q, a, q_i) - \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} \hat{P}(q, a, q_i) - v \\ &\quad + \inf_{p \in \mathbb{R}^{|\mathcal{Q}|}} \left[p^\top (V_k - \lambda_L + \lambda_U + v\mathbb{1}) \right] \end{aligned}$$

For any $a \in \mathcal{A}$, the Lagrangian dual problem of the minimization problem in (Mm) is (see [2]):

$$\begin{aligned} \max_{\lambda_L, \lambda_U, v} g(a, \lambda_L, \lambda_U, v) \\ \text{s.t. : } \lambda_L \geq 0, \lambda_U \geq 0, v \in \mathbb{R} \end{aligned} \quad (14)$$

As done in [2], to maximize $g(a, \lambda_L, \lambda_U, v)$, we have to make $V_k - \lambda_L + \lambda_U + v\mathbb{1} = 0$, or otherwise the term $\inf_{p \in \mathbb{R}^{|\mathcal{Q}|}} \left[p^\top (V_k - \lambda_L + \lambda_U + v\mathbb{1}) \right]$ will always be $-\infty$. Thus, (14) becomes:

$$\begin{aligned} \max_{\lambda_L, \lambda_U, v} \sum_{i=1}^{|\mathcal{Q}|} \lambda_{L_i} \check{P}(q, a, q_i) - \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} \hat{P}(q, a, q_i) - v \\ \text{s.t. : } \lambda_L \geq 0, \lambda_U \geq 0, v \in \mathbb{R}, V_k - \lambda_L + \lambda_U + v\mathbb{1} = 0 \end{aligned} \quad (15)$$

As the inner minimization problem of (Mm) is an LP, strong duality holds between the minimization problem of (Mm) and (15). Thus, we can apply Lemma (7.1), and the proof is completed. \square

7.3 Proof of Theorem 3.2

PROOF OF THEOREM 3.2. From Proposition 3.1, we only need to solve (4), in order to calculate the optimal value of (Mm) and to find an optimal action a_\star that solves (Mm). Thus, for the following, we focus on (4).

Eliminating the λ -variables. First, we eliminate the equality constraint, and thus the variable λ_L , by replacing $\lambda_L = V_k + \lambda_U + v\mathbb{1}$ in the objective function and all other constraints:

$$\begin{aligned} \max_{a, \lambda_U, v} \sum_{i=1}^{|\mathcal{Q}|} (V_k(q_i) + \lambda_{U_i} + v) \check{P}(q, a, q_i) - \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} \hat{P}(q, a, q_i) - v = \\ \max_{a, \lambda_U, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \check{P}(q, a, q_i) + \sum_{i=1}^{|\mathcal{Q}|} \lambda_{U_i} \underbrace{(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i))}_{\leq 0} \\ + v \underbrace{\left(\sum_{i=1}^{|\mathcal{Q}|} \check{P}(q, a, q_i) - 1 \right)}_{\leq 0} \\ \text{s.t. : } a \in \mathcal{A}, \lambda_U \geq 0, v \in \mathbb{R}, V_k + \lambda_U + v\mathbb{1} \geq 0 \end{aligned} \quad (16)$$

For any given a and v , since $(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i)) \leq 0$ and $\lambda_U \geq 0$, the variables λ_{U_i} have to be made as small as possible, in order to maximize the objective function. Due to the two constraints $\lambda_U \geq 0$ and $V_k + \lambda_U + v\mathbb{1} \geq 0$, this is encoded as $\lambda_{U_i} = \max(0, -V_k(q_i) - v)$. Thus, optimization problem (16) becomes:

$$\begin{aligned} \max_{a, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \check{P}(q, a, q_i) \\ + \sum_{i=1}^{|\mathcal{Q}|} \left[\max(0, -V_k(q_i) - v) \underbrace{(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i))}_{\leq 0} \right] \\ + v \underbrace{\left(\sum_{i=1}^{|\mathcal{Q}|} \check{P}(q, a, q_i) - 1 \right)}_{\leq 0} \\ \text{s.t. : } a \in \mathcal{A}, v \in \mathbb{R} \end{aligned} \quad (17)$$

where both constraints $\lambda_U \geq 0$ and $V_k + \lambda_U + v\mathbb{1} \geq 0$ have been by-construction eliminated by fixing $\lambda_{U_i} = \max(0, -V_k(q_i) - v)$. We are now left only with the optimization variables a and v .

Deriving $|\mathcal{Q}|$ simpler max problems from (17). Problem (17) contains the product $\max\left(0, -V_k(q_i) - v\right)\left(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i)\right)$, which is hard to analyze directly. Let us further simplify it. Since V_k is in descending order, optimization problem (17) can be broken down to $|\mathcal{Q}| + 1$ simpler problems, as follows:

- (1) The 1st optimization problem is the same with (17), but v is restricted by the inequality constraint $v \leq -V_k(q_1)$.
- (2) The $(|\mathcal{Q}| + 1)$ -th optimization problem is the same with (17), but v is restricted by the inequality constraint $v \geq -V_k(q_{|\mathcal{Q}|})$.
- (3) For $j \in \{2, \dots, |\mathcal{Q}|\}$, the j -th optimization problem is the same with (17), but v is restricted by the inequality constraints $-V_k(q_{j-1}) \leq v \leq -V_k(q_j)$.

Let us examine each case separately.

Case 1: Since $v \leq -V_k(q_1)$, then $v \leq -V_k(q_i)$ for all $i = 1, 2, \dots, |\mathcal{Q}|$. Thus, $\max\left(0, -V_k(q_i) - v\right) = -V_k(q_i) - v$ for all i . Hence, optimization problem (17) becomes:

$$\begin{aligned} & \max_{a, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \check{P}(q, a, q_i) \\ & + \sum_{i=1}^{|\mathcal{Q}|} \left[\left(-V_k(q_i) - v \right) \left(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i) \right) \right] \\ & + v \left(\sum_{i=1}^{|\mathcal{Q}|} \check{P}(q, a, q_i) - 1 \right) = \\ & \max_{a, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \hat{P}(q, a, q_i) + v \underbrace{\left(\sum_{i=1}^{|\mathcal{Q}|} \hat{P}(q, a, q_i) - 1 \right)}_{\geq 0} \\ \text{s.t. : } & a \in \mathcal{A}, v \leq -V_k(q_1) \end{aligned}$$

Now, to maximize the objective function, given any a , v has to be made as big as possible, i.e. $v = -V_k(q_1)$. Hence, the optimization problem transforms to the following:

$$\begin{aligned} & \max_a \sum_{i=1}^{|\mathcal{Q}|} \left(V_k(q_i) - V_k(q_1) \right) \hat{P}(q, a, q_i) + V_k(q_1) = \\ & \max_a \sum_{i=2}^{|\mathcal{Q}|} \left(V_k(q_i) - V_k(q_1) \right) \hat{P}(q, a, q_i) + V_k(q_1) \quad (18) \\ \text{s.t. : } & a \in \mathcal{A} \end{aligned}$$

Case 2: Since $v \geq -V_k(q_{|\mathcal{Q}|})$, then $v \geq -V_k(q_i)$ for all $i = 1, 2, \dots, |\mathcal{Q}|$. Thus, $\max\left(0, -V_k(q_i) - v\right) = 0$ for all i . Hence, optimization problem (17) becomes:

$$\begin{aligned} & \max_{a, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \check{P}(q, a, q_i) + v \underbrace{\left(\sum_{i=1}^{|\mathcal{Q}|} \check{P}(q, a, q_i) - 1 \right)}_{\leq 0} \\ \text{s.t. : } & a \in \mathcal{A}, v \geq -V_k(q_{|\mathcal{Q}|}) \end{aligned}$$

To maximize the objective function, given any a , v has to be made as small as possible, i.e. $v = -V_k(q_{|\mathcal{Q}|})$. Hence, the optimization

problem transforms to the following:

$$\begin{aligned} & \max_a \sum_{i=1}^{|\mathcal{Q}|} \left(V_k(q_i) - V_k(q_{|\mathcal{Q}|}) \right) \check{P}(q, a, q_i) + V_k(q_{|\mathcal{Q}|}) = \\ & \max_a \sum_{i=1}^{|\mathcal{Q}|-1} \left(V_k(q_i) - V_k(q_{|\mathcal{Q}|}) \right) \check{P}(q, a, q_i) + V_k(q_{|\mathcal{Q}|}) = \quad (19) \\ \text{s.t. : } & a \in \mathcal{A} \end{aligned}$$

Case 3: In this case, we have that:

$$\max\left(0, -V_k(q_i) - v\right) = \begin{cases} 0, & i = 1, 2, \dots, j-1 \\ -V_k(q_i) - v, & i = j, \dots, |\mathcal{Q}| \end{cases}$$

Then, optimization problem (17) becomes:

$$\begin{aligned} & \max_{a, v} \sum_{i=1}^{|\mathcal{Q}|} V_k(q_i) \check{P}(q, a, q_i) \\ & + \sum_{i=j}^{|\mathcal{Q}|} \left[\left(-V_k(q_i) - v \right) \left(\check{P}(q, a, q_i) - \hat{P}(q, a, q_i) \right) \right] \\ & + v \left(\sum_{i=1}^{|\mathcal{Q}|} \check{P}(q, a, q_i) - 1 \right) = \\ & \max_{a, v} \sum_{i=1}^{j-1} V_k(q_i) \check{P}(q, a, q_i) + \sum_{i=j}^{|\mathcal{Q}|} V_k(q_i) \hat{P}(q, a, q_i) \\ & + v \left(\sum_{i=1}^{j-1} \check{P}(q, a, q_i) + \sum_{i=j}^{|\mathcal{Q}|} \hat{P}(q, a, q_i) - 1 \right) \\ \text{s.t. : } & a \in \mathcal{A}, -V_k(q_{j-1}) \leq v \leq -V_k(q_j) \end{aligned}$$

Now, for any given a , the term $\sum_{i=1}^{j-1} \check{P}(q, a, q_i) + \sum_{i=j}^{|\mathcal{Q}|} \hat{P}(q, a, q_i) - 1$ is either negative or positive. In the first case, we would have to make v as small as possible, i.e. $v = -V_k(q_{j-1})$. In the second case, we would have to make v as big as possible, i.e. $v = -V_k(q_j)$. Thus, we can break down the above optimization problem to the two following problems (20) and (21):

$$\begin{aligned} & \max_a \sum_{i=1}^{j-1} \left(V_k(q_i) - V_k(q_{j-1}) \right) \check{P}(q, a, q_i) \\ & + \sum_{i=j}^{|\mathcal{Q}|} \left(V_k(q_i) - V_k(q_{j-1}) \right) \hat{P}(q, a, q_i) + V_k(q_{j-1}) = \\ & \max_a \sum_{i=1}^{j-2} \left(V_k(q_i) - V_k(q_{j-1}) \right) \check{P}(q, a, q_i) \quad (20) \\ & + \sum_{i=j}^{|\mathcal{Q}|} \left(V_k(q_i) - V_k(q_{j-1}) \right) \hat{P}(q, a, q_i) + V_k(q_{j-1}) \\ \text{s.t. : } & a \in \mathcal{A} \end{aligned}$$

$$\begin{aligned}
& \max_a \sum_{i=1}^{j-1} (V_k(q_i) - V_k(q_j)) \check{P}(q, a, q_i) \\
& + \sum_{i=j}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_j)) \hat{P}(q, a, q_i) + V_k(q_j) = \\
& \max_a \sum_{i=1}^{j-1} (V_k(q_i) - V_k(q_j)) \check{P}(q, a, q_i) \\
& + \sum_{i=j+1}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_j)) \hat{P}(q, a, q_i) + V_k(q_j) \\
& \text{s.t. : } a \in \mathcal{A}
\end{aligned} \tag{21}$$

Thus far, we have that (4), and thus the primal problem (Mm), has been broken down to the combination of problem (18), problem (19) and $|\mathcal{Q}| - 1$ instances of problems (20) and (21). However, taking a closer look, we see that (21) with $j = m$ is identical to (20) with $j = m + 1$ (their objective functions are identical, as well as their constraint sets)⁷. Moreover, (19) is identical to (21) with $j = |\mathcal{Q}|$. Finally, (18) is the same problem as (20) with $j = 2$, which can also be written in the form of (21) if we put $j = 1$. In other words, it suffices to only consider (21) for $j = 1, \dots, |\mathcal{Q}|$. To conclude, solving (4) is equivalent to solving the following optimization problems for $j = 1, \dots, |\mathcal{Q}|$:

$$\begin{aligned}
& \max_a \sum_{i=1}^{j-1} (V_k(q_i) - V_k(q_j)) \check{P}(q, a, q_i) \\
& + \sum_{i=j+1}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_j)) \hat{P}(q, a, q_i) + V_k(q_j) \\
& \text{s.t. : } a \in \mathcal{A}
\end{aligned}$$

□

ACKNOWLEDGMENTS

The authors would like to thank Dr. Sergio Grammatico, Dr. Amin Sharifi Kolarijani, and Dr. Gabriel de Albuquerque Gleizer for their helpful comments.

This work was supported in part by the European Research Council through the SENTIENT project, Grant No. ERC-2017-STG #755953 (<https://cordis.europa.eu/project/id/755953>), and National Science Foundation (NSF) award 2039062.

⁷Denote by $f_j(a)$ the objective function of (21) and by $g_j(a)$ the objective function of (20). Then, we have:

$$\begin{aligned}
& f_m(a) - g_{m+1}(a) = \\
& = \sum_{i=1}^{m-1} (V_k(q_i) - V_k(q_m)) \check{P}(q, a, q_i) - \sum_{i=1}^{m-1} (V_k(q_i) - V_k(q_m)) \check{P}(q, a, q_i) \\
& + \sum_{i=m+1}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_m)) \hat{P}(q, a, q_i) - \sum_{i=m+1}^{|\mathcal{Q}|} (V_k(q_i) - V_k(q_m)) \hat{P}(q, a, q_i) \\
& + V_k(q_m) - V_k(q_m) \\
& = 0
\end{aligned}$$

REFERENCES

- [1] Dimitri P Bertsekas and Steven Shreve. 2004. *Stochastic optimal control: the discrete-time case*.
- [2] Stephen Boyd and Lieven Vandenberghe. 2004. *Convex optimization*. Cambridge university press.
- [3] Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, and Ufuk Topcu. 2021. Convex Optimization for Parameter Synthesis in MDPs. *IEEE Trans. Automat. Control* (2021).
- [4] Giannis Delimpaltadakis, Luca Laurenti, and Manuel Mazo Jr. 2022. Formal Analysis of the Sampling Behaviour of Stochastic Event-Triggered Control. *arXiv preprint arXiv:2202.10178* (2022).
- [5] Maxence Dutreix, Jeongmin Huh, and Samuel Coogan. 2022. Abstraction-based synthesis for stochastic systems with omega-regular objectives. *Nonlinear Analysis: Hybrid Systems* 45 (2022), 101204.
- [6] James E Falk. 1973. A linear max–min problem. *Mathematical Programming* 5, 1 (1973), 169–188.
- [7] Sicun Gao, Soonho Kong, and Edmund M Clarke. 2013. dReal: An SMT solver for nonlinear theories over the reals. In *International conference on automated deduction*. Springer, 208–214.
- [8] Robert Givan, Sonia Leach, and Thomas Dean. 2000. Bounded-parameter Markov decision processes. *Artificial Intelligence* 122, 1-2 (2000), 71–109.
- [9] Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. 2011. Probabilistic reachability for parametric Markov models. *International Journal on Software Tools for Technology Transfer* 13, 1 (2011), 3–19.
- [10] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. 2021. Strategy synthesis for partially-known switched stochastic systems. In *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control*. 1–11.
- [11] Xenofon Koutsoukos and Derek Riley. 2006. Computational methods for reachability analysis of stochastic hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*. Springer, 377–391.
- [12] Morteza Lahijanian, Sean B Andersson, and Calin Belta. 2015. Formal verification and synthesis for discrete-time stochastic systems. *IEEE Trans. Automat. Control* 60, 8 (2015), 2031–2045.
- [13] Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina. 2007. Parametric probabilistic transition systems for system design and analysis. *Formal Aspects of Computing* 19, 1 (2007), 93–109.
- [14] Luca Laurenti, Morteza Lahijanian, Alessandro Abate, Luca Cardelli, and Marta Kwiatkowska. 2020. Formal and efficient synthesis for continuous-time linear stochastic hybrid processes. *IEEE Trans. Automat. Control* 66, 1 (2020), 17–32.
- [15] Abolfazl Lavaei, Sadegh Soudjani, Alessandro Abate, and Majid Zamani. 2022. Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica* 146 (2022), 110617.
- [16] Arnab Nilim and Laurent El Ghaoui. 2005. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research* 53, 5 (2005), 780–798.
- [17] Andrzej S Nowak. 1984. On zero-sum stochastic games with general state space I. *Probability and Mathematical Statistics* 4, 1 (1984), 13–32.
- [18] R Tyrrell Rockafellar. 1970. *Convex analysis*. Vol. 18. Princeton university press.
- [19] G George Yin and Chao Zhu. 2009. *Hybrid switching diffusions: properties and applications*. Vol. 63. Springer Science & Business Media.