

Robust Control for Dynamical Systems With Non-Gaussian Noise via Formal Abstractions

Thom Badings

Radboud University, Nijmegen, the Netherlands

THOM.BADINGS@RU.NL

Licio Romao

Alessandro Abate

David Parker

University of Oxford, Oxford, United Kingdom

LICIO.ROMAO@CS.OX.AC.UK

ALESSANDRO.ABATE@CS.OX.AC.UK

DAVID.PARKER@CS.OX.AC.UK

Hasan A. Poonawala

University of Kentucky, Kentucky, USA

HASAN.POONAWALA@UKY.EDU

Marielle Stoelinga

Radboud University, Nijmegen, the Netherlands

University of Twente, Enschede, the Netherlands

M.STOELINGA@CS.RU.NL

Nils Jansen

Radboud University, Nijmegen, the Netherlands

N.JANSEN@SCIENCE.RU.NL

Abstract

Controllers for dynamical systems that operate in safety-critical settings must account for stochastic disturbances. Such disturbances are often modeled as process noise in a dynamical system, and common assumptions are that the underlying distributions are known and/or Gaussian. In practice, however, these assumptions may be unrealistic and can lead to poor approximations of the true noise distribution. We present a novel controller synthesis method that does not rely on any explicit representation of the noise distributions. In particular, we address the problem of computing a controller that provides probabilistic guarantees on safely reaching a target, while also avoiding unsafe regions of the state space. First, we abstract the continuous control system into a finite-state model that captures noise by probabilistic transitions between discrete states. As a key contribution, we adapt tools from the scenario approach to compute probably approximately correct (PAC) bounds on these transition probabilities, based on a finite number of samples of the noise. We capture these bounds in the transition probability intervals of a so-called interval Markov decision process (iMDP). This iMDP is, with a user-specified confidence probability, robust against uncertainty in the transition probabilities, and the tightness of the probability intervals can be controlled through the number of samples. We use state-of-the-art verification techniques to provide guarantees on the iMDP and compute a controller for which these guarantees carry over to the original control system. In addition, we develop a tailored computational scheme that reduces the complexity of the synthesis of these guarantees on the iMDP. Benchmarks on realistic control systems show the practical applicability of our method, even when the iMDP has hundreds of millions of transitions.

1. Introduction

Dynamical systems. Dynamical systems are widely used to model the state dynamics of complex systems (Kulakowski, Gardner, & Shearer, 2007). For example, an unmanned aerial vehicle (UAV) can be modeled as a dynamical system, in which the (possibly continuous)

state reflects its current position and velocity, and the (possibly continuous) *control inputs* reflect choices that may change the state over time (Kulakowski et al., 2007). The dynamical system is *linear* if the state transition is linear in the current state and control input.

Controller synthesis. Verifying that a controlled dynamical system satisfies a desired property is of paramount importance, especially in safety-critical applications. Traditional methods from control theory are powerful tools for addressing properties about stability and (asymptotic) convergence, but these methods by and large do not provide formal guarantees about richer, temporal properties (Baier & Katoen, 2008; Fan, Qin, Mathur, Ning, Mitra, & Viswanathan, 2022), namely requirements on state trajectories of the system in time. A common example is the *reach-avoid property*, where the task is to reach a desirable region within a given time horizon, while always remaining in a (possibly non-convex) safe region. The general *controller synthesis problem* is to compute (in an automated fashion) a *feedback controller* for the dynamical system, such that any state trajectory that the system generates satisfies the given reach-avoid property (Kwiatkowska & Parker, 2013).

Modeling process noise. However, for a system such as a UAV, factors including turbulence and wind gusts cause *uncertainty* affecting the state dynamics of the system (Blackmore, Ono, Bektassov, & Williams, 2010). We model such uncertainty as *process noise*, which is an additive random variable (with possibly infinitely many outcomes) entering the dynamical system and thus affecting state transitions. Due to this additional noise term, it is generally impossible to guarantee that *every* state trajectory satisfies a reach-avoid property. Instead, we reason over the *probability* that a reach-avoid property is satisfied, and the synthesis problem is then to compute a controller that maximizes this probability.

Distribution of the noise. A common assumption to achieve computational tractability for the controller synthesis problem is that the process noise follows a Gaussian distribution (Park, Serpedin, & Qaraqe, 2013), e.g., as is classically assumed in linear-quadratic-Gaussian control (Anderson & Moore, 2007). However, in realistic problems, such as a UAV operating under turbulence, this assumption yields a poor approximation of the uncertainty (Blackmore et al., 2010). Distributions may even be *unknown*, meaning that one cannot derive a set-bounded or a precise probabilistic representation of the noise. In this case, it is generally hard or even impossible to derive *hard guarantees* on the probability that a given controller ensures the satisfaction of the considered reach-avoid property.

Problem statement. In this paper, we consider the controller synthesis problem for dynamical systems with additive process noise of an unknown distribution. For the synthesized controller, we provide a *probably approximately correct (PAC)* guarantee on the probability of satisfying a given reach-avoid problem. As such, we solve the following problem:

Given a linear dynamical system perturbed by additive noise of unknown distribution, compute a controller under which, with a user-specified confidence level, the probability to satisfy a reach-avoid problem is above a given threshold value.

Finite-state abstraction. We solve this problem by computing a finite-state abstraction of the original dynamical system (Soudjani & Abate, 2013), which we obtain from a *partition* of its continuous state space into a set of disjoint convex *regions*. Actions in this abstraction correspond to continuous control inputs that yield transitions between these regions. Due

to the process noise, the outcome of an action is stochastic, rendering transitions probabilistic. We capture these probabilities in a Markov decision process (MDP) (Puterman, 1994). A defining characteristic of our approach is that we leverage *backward reachability computations* on the dynamical system to determine which actions are enabled at each discrete region. By contrast, most other abstraction methods (see the related work in Sect. 7 and the survey article Lavaei, Soudjani, Abate, & Zamani, 2022) rely on *forward* reachability computations, which are associated with errors that grow with the time horizon of the considered property. Our backward scheme avoids such abstraction errors, at the cost of requiring slightly more restrictive assumptions on the system dynamics.

Probability intervals. Since the distribution of the noise is unknown, it is not possible to compute the transition probabilities of the abstract MDP exactly, e.g., as in Soudjani and Abate (2013). Instead, we estimate the probabilities based on a finite number of *samples* of the noise, which may be obtained from a high fidelity (black box) simulator, from historical data, or from (physical or numerical) experiments. To be *robust* against estimation errors in these probabilities, we formulate the error estimation process as a so-called scenario optimization problem and leverage state-of-the-art tools (Romao, Papachristodoulou, & Margellos, 2022) from the *scenario approach*, which is a methodology to deal with stochastic optimization in a data-driven fashion (Campi & Garatti, 2008; Campi, Carè, & Garatti, 2021). We compute *upper and lower bounds* on the transition probabilities with a desired *confidence level*, which we choose up front. These bounds are *PAC*, as they contain the true probabilities with at least this confidence level.

Interval MDPs. We formalize our abstractions with the PAC probability intervals using so-called interval Markov decision processes (iMDPs), which are an extension of MDPs with intervals of probabilities (Givan, Leach, & Dean, 2000). More generally, iMDPs are a particular case of uncertain or non-deterministic MDPs; see the references in Sect. 7 for more details. iMDPs have recently been proposed as an alternative to standard MDPs for abstracting stochastic dynamical systems (Cauchi, Laurenti, Lahijanian, Abate, Kwiatkowska, & Cardelli, 2019; Lavaei et al., 2022). We show explicitly how to lift the PAC guarantees on individual transition probabilities to a correctness guarantee on the whole iMDP. Policies for iMDPs have to robustly account for all possible probabilities within the intervals, and one usually provides upper and lower bounds on maximal or minimal reachability probabilities or expected rewards (Hahn, Hashemi, Hermanns, Lahijanian, & Turrini, 2017; Puggelli, Li, Sangiovanni-Vincentelli, & Seshia, 2013; Wolff, Topcu, & Murray, 2012). Note that given the PAC-correct iMDP, these bounds on the reachability probabilities are exact and thus do not introduce another error bound. For MDPs, mature tool support exists, e.g., via PRISM (Kwiatkowska, Norman, & Parker, 2011), and this support was previously extended to iMDPs in Badings, Abate, Jansen, Parker, Poonawala, and Stoelinga (2022a).

Infinite-horizon properties. One notable feature of our abstraction scheme is that we can handle reach-avoid properties over *infinite-time horizons*, also known as unbounded properties (Tkachev & Abate, 2014). Most existing abstraction-based controller synthesis methods cause abstraction errors that grow with the time horizon and are thus limited to finite-horizon properties (Abate, Katoen, Lygeros, & Prandini, 2010; Soudjani & Abate, 2013; Lavaei et al., 2022; Zikelic, Lechner, Henzinger, & Chatterjee, 2022). By contrast, we

provide a PAC guarantee on each transition probability interval of the iMDP being correct, without errors that grow with the horizon of the considered properties.

Iterative abstraction improvement. The tightness of the probability intervals in the iMDP depends on the number of noise samples. Hence, we propose an *iterative abstraction scheme* to improve these intervals by sequentially increasing the number of noise samples used to compute probability intervals. Our scheme can be summarized as follows. We first compute a robust policy that maximizes the probability of safely reaching the goal states. Then, we check whether this optimal reach-avoid probability is satisfactory or not with respect to the pre-defined threshold value on the given property. In case it is not, we collect additional samples to reduce the uncertainty in the probability intervals; otherwise, we extract and use the optimal policy to compute “on the fly” a feedback controller for the original dynamical system. The specified confidence level reflects the likelihood that the optimal reach-avoid probability on the iMDP is a *lower bound* for the probability that the dynamical system satisfies the reach-avoid problem under this derived controller.

Improved policy synthesis scheme. The generated iMDP abstractions can potentially have hundreds of millions of transitions, especially if the state dimension is high. To reduce the computational complexity related to the policy synthesis algorithm on the iMDP, we propose an algorithmic optimization. Instead of employing the large iMDP abstraction, we soundly merge states that show similar reach-avoid probabilities. In particular, we associate a merged state with the minimum reach-avoid probability of the set of states it is comprised of, resulting in a conservative but sound approximation of the original iMDP. This strategy reduces the overall size of the iMDP significantly and additionally enables us to solve more complex reach-avoid problems, as shown in our experiments.

Contributions. Our contributions are threefold: (1) We propose a novel method to compute controllers with PAC guarantees for dynamical systems with unknown noise distributions. Specifically, the probability of satisfying an (in)finite-horizon reach-avoid problem is guaranteed, with a user-specified confidence probability, to exceed a pre-defined threshold. (2) We propose a sound and scalable policy synthesis algorithm that allows us to verify abstract models with hundreds of millions of transitions. (3) We apply our method to multiple realistic control problems and benchmark against two other tools: StocHy and SReachTools. We demonstrate that the guarantees obtained for the iMDP abstraction carry over to the dynamical system of interest. Moreover, we show that using probability intervals instead of point estimates of probabilities yields significantly more robust results.

This paper extends Badings et al. (2022a) in several ways. First, we expand on the intuition and presentation of our abstraction method, comparing the scenario optimization with other methods for computing probability intervals. Second, we provide deeper theoretical results on the correctness of our method by formalizing the relationship between the dynamical system and the iMDP. Moreover, we lift the PAC guarantees on individual transitions, as done in Badings et al. (2022a), to a correctness guarantee on *the whole* iMDP abstraction. Finally, we present a new algorithmic extension and an additional experiment on an autonomous satellite rendezvous problem from Jewison and Erwin (2016).

2. Foundations and Outline

A *discrete probability distribution* over a finite set X with cardinality $|X|$ is a function $\text{prob}: X \rightarrow [0, 1]$ with $\sum_{x \in X} \text{prob}(x) = 1$. The set of all distributions over X is $\text{Dist}(X)$. A *probability density function* over a random variable x conditioned on y is written as $p(x|y)$. All vectors $\mathbf{x} \in \mathbb{R}^n$, $n \in \mathbb{N}$, are column vectors and are denoted by bold letters. We use the term *controller* when referring to a map selecting inputs for dynamical systems, while we use the term *policy* for (i)MDPs.

2.1 Linear Dynamical Systems

We consider discrete-time, continuous-state systems, where the progression of the state $\mathbf{x} \in \mathbb{R}^n$ depends *linearly* on the current state, on a control input, and on a process noise term. Given a state \mathbf{x}_k at discrete time $k \in \mathbb{N}$, the successor state at time $k + 1$ is given as

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k + \mathbf{q}_k + \mathbf{w}_k, \quad (1)$$

where $\mathbf{u}_k \in \mathcal{U} \subset \mathbb{R}^p$ is the (continuous) control input at time k , $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times p}$ are appropriate matrices, and $\mathbf{q}_k \in \mathbb{R}^n$ is a deterministic disturbance. The term $\mathbf{w}_k \in \Delta \subset \mathbb{R}^n$ is an arbitrary additive process noise term, which is a random variable defined on a probability space $(\Delta, \mathcal{D}, \mathbb{P})$, with σ -algebra \mathcal{D} and probability measure \mathbb{P} defined over \mathcal{D} . In our formulation, the sample space Δ is an abstract set of possible noise values at any time k , and the probability measure \mathbb{P} is unknown but time-invariant. To deal with this lack of knowledge, we employ a sampling-based approach, for which it suffices to obtain a finite number of samples of the process noise, whose distribution is independent of time. This underlying source of uncertainty induces a probabilistic model over the successor state \mathbf{x}_{k+1} . We denote the probability density function over successor states for a given state \mathbf{x}_k and control input \mathbf{u}_k as $p_{\mathbf{w}_k}(\mathbf{x}_{k+1} | \mathbf{x}_k, \mathbf{u}_k)$.

The linear dynamical system defined by Eq. (1) is controlled by a *time-varying linear feedback controller* of the following form:

Definition 1. A *piece-wise linear feedback controller* is a function $\phi: \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$, which maps a state $\mathbf{x}_k \in \mathbb{R}^n$ and a time step $k \in \mathbb{N}$ to a control input $\mathbf{u}_k \in \mathcal{U}$.

We use time-varying controllers because, for finite-horizon control objectives, the optimal control generally depends on the time index. For infinite-horizon properties instead, the optimal control is independent of the time index (Baier & Katoen, 2008).

Example 1. The longitudinal dynamics of a UAV are modeled as

$$\mathbf{x}_{k+1} = \begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} \mathbf{u}_k + \mathbf{w}_k, \quad (2)$$

where p_k and v_k are the position and velocity at time k , and the control is bounded by $\mathbf{u}_k \in \mathcal{U} = [-4, 4]$. The distribution of the noise \mathbf{w}_k is unknown, thus possibly not Gaussian.

Remark 1 (Restriction to linear systems). *Our methods are theoretically amenable to work with nonlinear drift dynamics rather than the linear drift terms we see in Eqs. (1) and (2). However, this requires more advanced 1-step reachability computations, which are not core to our main contributions. Hence, we restrict ourselves to the linear drift of the model in Eq. (1) and discuss extensions to nonlinear models in Sect. 8.*

2.2 Problem Statement

We consider control objectives expressed as *reach-avoid properties* over an infinite or a finite horizon. A reach-avoid property $\varphi_{\mathbf{x}_0}^K$ is satisfied if, starting from an initial state $\mathbf{x}_0 \in \mathbb{R}^n$ at time $k = 0$, the system reaches a desired *goal region* $\mathcal{X}_G \subset \mathbb{R}^n$ within a horizon of $K \in \mathbb{N} \cup \infty$ steps, while avoiding a *critical region* $\mathcal{X}_C \subset \mathbb{R}^n$. We write the probability of satisfying a reach-avoid property $\varphi_{\mathbf{x}_0}^K$ under a controller ϕ as $Pr^\phi(\varphi_{\mathbf{x}_0}^K)$. We formally state the problem that we solve in this paper as follows.

Compute a controller ϕ for the dynamical system in Eq. (1) that, with a confidence probability of at least $\alpha \in [0, 1]$, guarantees that $Pr^\phi(\varphi_{\mathbf{x}_0}^K) \geq \eta$, where $\eta \in [0, 1]$ is a pre-defined probability threshold.

2.3 Markov Decision Processes

We solve the problem above via a finite-state abstraction of the dynamical system, which we formalize as a variant of MDPs (Puterman, 1994):

Definition 2 (MDP). *A Markov decision process (MDP) is a tuple $\mathcal{M} = (S, Act, s_I, P)$ where S is a (finite) set of states, Act is a (finite) set of actions, s_I is the initial state, and $P: S \times Act \rightarrow Dist(S)$ is the (partial)¹ probabilistic transition function.*

The set of actions enabled in state $s \in S$ is $Act(s) \subseteq Act$. We call (s, a, s') with probability $P(s, a)(s') > 0$ a *transition*. A time-varying deterministic policy for an MDP \mathcal{M} is a function $\pi: S \times \mathbb{N} \rightarrow Act$ mapping states and time indices to actions (Puterman, 1994). The set of all possible policies for \mathcal{M} is denoted by $\Pi_{\mathcal{M}}$.

A *probabilistic reach-avoid property* $Pr^\pi(\varphi_{s_I}^K)$ for an MDP describes the probability of reaching a set of goal states $S_G \subset S$ within $K \in \mathbb{N} \cup \infty$ steps under policy $\pi \in \Pi_{\mathcal{M}}$, while avoiding a set of critical states $S_C \subset S$, where $S_G \cap S_C = \emptyset$. An optimal policy $\pi^* \in \Pi_{\mathcal{M}}$ for MDP \mathcal{M} maximizes the *reach-avoid probability*:

$$\pi^* = \arg \max_{\pi \in \Pi_{\mathcal{M}}} Pr^\pi(\varphi_{s_I}^K). \quad (3)$$

Remark 2 (Dynamical systems as continuous MDPs). *The dynamical system in Eq. (1) can equivalently be seen as an MDP with an uncountably infinite number of states S (representing all $\mathbf{x}_k \in \mathbb{R}^n$), and an infinite number of actions Act (representing all $\mathbf{u}_k \in \mathcal{U}$). Note that every action is enabled at any state, i.e. $Act(s) = Act, \forall s \in S$. Moreover, each state-control pair $(\mathbf{x}_k, \mathbf{u}_k)$, i.e., applying control input \mathbf{u}_k in state \mathbf{x}_k induces a probability distribution (or in fact, a density function) over successor states $\mathbf{x}_{k+1} \in \mathbb{R}^n$.*

We now relax the assumption that the transition probabilities of MDPs are precisely given.

Definition 3 (iMDP). *An interval Markov decision process (iMDP) is a tuple $\mathcal{M}_{\mathbb{I}} = (S, Act, s_I, \mathcal{P})$ where S , Act , and s_I are defined by Def. 2, and where the uncertain (partial) probabilistic transition function $\mathcal{P}: S \times Act \times S \rightarrow \mathbb{I} \cup \{[0, 0]\}$ is defined over intervals $\mathbb{I} = \{[a, b] \mid a, b \in (0, 1] \text{ and } a \leq b\}$.*

1. The transition function is partial in general, meaning that not all state-action pairs $(s, a) \in S \times Act$ are in the domain of P . We need to define this partial transition function because not all actions may be enabled in each state.

Note that an interval cannot have a lower bound of 0 except for the $[0, 0]$ interval. Since the transition function \mathcal{P} is partial, we do not require each action to be enabled at any state. An iMDP defines a (possibly empty) set of MDPs that vary only in their transition function. In particular, for an MDP with transition function P , we write $P \in \mathcal{P}$ if for all $s, s' \in S$ and $a \in Act$ we have $P(s, a)(s') \in \mathcal{P}(s, a, s')$ and $P(s, a) \in Dist(S)$.

Remark 3 (Geometry of uncertainty sets). *For each state-action pair (s, a) , the set $\{P(s, a) \mid P \in \mathcal{P}\}$ of feasible probability distributions is a probability simplex with interval (box) constraints defined by the intervals in \mathcal{P} . The resulting uncertainty set $\{P(s, a) \mid P \in \mathcal{P}\}$ is a convex polytope by construction. Moreover, the transition probabilities $P(s, a)$ for different state-action pairs (s, a) are independent and thus respect the so-called rectangularity assumption, which is common in the literature on robust MDPs (Wiesemann, Kuhn, & Sim, 2014) and robust optimization (Bertsimas, Brown, & Caramanis, 2011), and is necessary to guarantee computational tractability.*

For iMDPs, a policy needs to be *robust* against all possible transition functions $P \in \mathcal{P}$. We employ a robust variant of value iteration from Wolff et al. (2012) to compute a policy $\underline{\pi}^* \in \Pi_{\mathcal{M}_{\mathbb{I}}}$ for iMDP $\mathcal{M}_{\mathbb{I}}$ that maximizes the lower bound $\underline{Pr}^{\pi}(\varphi_{s_I}^K)$ of on the reach-avoid probability within horizon K :

$$\underline{\pi}^* = \arg \max_{\pi \in \Pi_{\mathcal{M}_{\mathbb{I}}}} \underline{Pr}^{\pi}(\varphi_{s_I}^K) = \arg \max_{\pi \in \Pi_{\mathcal{M}_{\mathbb{I}}}} \min_{P \in \mathcal{P}} Pr^{\pi}(\varphi_{s_I}^K). \quad (4)$$

Similarly, we can also maximize an upper bound $\overline{Pr}^{\pi}(\varphi_{s_I}^K)$ on the reach-avoid probability:

$$\overline{\pi}^* = \arg \max_{\pi \in \Pi_{\mathcal{M}_{\mathbb{I}}}} \overline{Pr}^{\pi}(\varphi_{s_I}^K) = \arg \max_{\pi \in \Pi_{\mathcal{M}_{\mathbb{I}}}} \max_{P \in \mathcal{P}} Pr^{\pi}(\varphi_{s_I}^K). \quad (5)$$

We will use the lower bound in Eq. (4) to compute robust optimal policies, while we use Eq. (5) to determine if the threshold η specified in the formal problem statement is satisfiable at all. Note that deterministic policies suffice to obtain optimal values for (i)MDPs (Puterman, 1994; Puggelli et al., 2013), and particularly also for reach-avoid properties (Abate, Prandini, Lygeros, & Sastry, 2008).

2.4 An Iterative Abstraction Scheme

Our approach is summarized in Fig. 1 and consists of an *offline planning phase* in which we generate the abstraction and obtain guarantees on the abstract model, and an *online control phase* in which we derive a controller for the dynamical system on the fly. We describe in Sect. 3 how, for a given linear dynamical system, we generate a finite-state abstraction as an MDP by partitioning the continuous state space. We then describe in Sect. 4 how we obtain PAC bounds on the MDP's transition probabilities based on a finite set of N samples of the noise. The resulting abstraction is an iMDP, for which we use Eqs. (4) and (5) to compute optimal policies $\underline{\pi}^*$ and $\overline{\pi}^*$ that maximize the lower/upper bounds on the probability of satisfying the given property. We then proceed in one of the following three ways:

1. If $\underline{Pr}^{\underline{\pi}^*}(\varphi_{s_I}^K) \geq \eta$, the formal problem is satisfied. Thus, we extract the optimal policy $\underline{\pi}^*$, which we use to determine a controller ϕ for the dynamical system on the fly.




Figure 1: Our iterative approach between abstraction and verification, where N is the number of samples used for the abstraction, and η is the threshold reach-avoid probability.

2. If $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K) < \eta$, the reach-avoid problem is (with high confidence) unsatisfiable. In this case, increasing the number of samples does not help, so we terminate the scheme.
3. If $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) < \eta$ but $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K) \geq \eta$, we obtain additional samples by increasing N by a fixed factor $\gamma > 1$. The updated iMDP has tighter probability intervals but may also have more transitions. However, since the states and actions of the abstraction are independent of N and defined at the first iteration, the MDP abstraction is only performed once.

In Sect. 5 we describe each step of our approach in more detail, as well as a complete algorithm for solving the formal problem. We perform numerical experiments on benchmarks from multiple application domains in Sect. 6. We discuss the related work in Sect. 7.

3. Finite-State MDP Abstraction

In this section, we describe how we generate an MDP abstraction of the linear dynamical system in Eq. (1). First, we partition the state space into a set of discrete convex regions. We then use this partition to build a finite-state MDP abstraction of the dynamical system.

3.1 State Space Discretization

We choose a *partition* $\mathcal{R} = \{R_1, \dots, R_v\}$ of a bounded portion $\mathcal{X} \subset \mathbb{R}^n$ of the continuous state space \mathbb{R}^n into a set of v disjoint *regions*, such that $\bigcup_{i=1}^v R_i = \mathcal{X}$. In addition, we define a single *absorbing region* R_* , representing $\mathbb{R}^n \setminus \mathcal{X}$. The absorbing region R_* captures the event that the continuous state leaves the bounded portion of the state space over which we plan. We consider the regions in \mathcal{R} to be n -dimensional bounded, convex polytopes. Thus, each region $R_i \in \mathcal{R}$ is the solution set of m linear inequalities parameterized by $M_i \in \mathbb{R}^{m \times n}$ and




Figure 2: A fragment of a rectangular partitioning of $\mathcal{X} \subset \mathbb{R}^2$, showing the backward reachable set $\mathcal{G}(d_a)$ of one particular action $a \in \text{Act}$ with target point d_a . Action a is enabled in states s and s' , since $R_s \subseteq \mathcal{G}(d_a)$ and $R_{s'} \subseteq \mathcal{G}(d_a)$, respectively.

$\mathbf{b}_i \in \mathbb{R}^m$, yielding $R_i = \{\mathbf{x} \in \mathbb{R}^n \mid M_i \mathbf{x} \leq \mathbf{b}_i\}$. In addition, the following assumption allows us to translate properties for the dynamical system to properties on the iMDP abstraction:

Assumption 1. *The continuous goal region \mathcal{X}_G and critical region \mathcal{X}_C are aligned with the union of a subset of regions in \mathcal{R} , i.e., $\mathcal{X}_G = \cup_{i \in I} R_i$ and $\mathcal{X}_C = \cup_{j \in J} R_j$ for index sets $I, J \subset \{1, 2, \dots, |\mathcal{R}|\}$.*

3.2 MDP Abstractions

We formalize the dynamical system discretized under \mathcal{R} as an MDP $\mathcal{M} = (S, \text{Act}, s_I, P)$, by defining its states, actions, and transition probabilities. We assume that the initial state $s_I \in S$ is known, and we capture time constraints by the reach-avoid property.

States. We define an MDP state for every region of partition \mathcal{R} , as well as for the absorbing region R_* . Thus, we obtain the set of MDP states $S = \{s_i \mid i = 1, \dots, |\mathcal{R}|\} \cup \{s_*\}$, which consists of $|S| = |\mathcal{R}| + 1$ states: one for every region in \mathcal{R} , plus one state s_* corresponding to the absorbing region R_* where the only outgoing transition leads back to s_* . We denote by $R_s \in \mathcal{R}$ the region of the partition associated with MDP state $s \in S$. Formally, we define a function $T: \mathbb{R}^n \rightarrow S$ that maps continuous states $\mathbf{x} \in \mathbb{R}^n$ to MDP states $s \in S$:²

Definition 4. *A continuous state $\mathbf{x} \in \mathbb{R}^n$ and MDP state $s \in S$ are related, i.e., $T(\mathbf{x}) = s$, if $\mathbf{x} \in R_s$. The inverse mapping $T^{-1}(s) = \{\mathbf{x} \in \mathbb{R}^n \mid (\mathbf{x}, s) \in T\}$ is equivalent to region R_s .*

The function T expresses a relation from any state $\mathbf{x} \in \mathbb{R}^n$ to a (unique) discrete state $s \in S$, and is, therefore, also called an *abstraction function*. Moreover, its inverse $T^{-1}(s) = R_s$ maps any state $s \in S$ to a set of continuous states and is equal to the region R_s of s itself.

2. Alternatively, we may define $T \subseteq \mathbb{R}^n \times S$ as a binary relation between each continuous state $\mathbf{x} \in \mathbb{R}^n$ and MDP state $s \in S$. In this paper, we use the function notation since we apply T as a function only.

Actions. Discrete actions correspond to the execution of a control input $\mathbf{u}_k \in \mathcal{U}$ in the dynamical system in Eq. (1). We define $q \in \mathbb{N}$ MDP actions, so $Act = \{a_1, \dots, a_q\}$. Let the noiseless successor state of \mathbf{x}_k be defined as $\hat{\mathbf{x}}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k + \mathbf{q}_k$ (i.e., \mathbf{x}_{k+1} under the assumption that $\mathbf{w}_k = 0$). Every action a is associated with a fixed continuous *target point* $\mathbf{d}_a \in \mathbb{R}^n$, and is defined such that its noiseless successor state $\hat{\mathbf{x}}_{k+1} = \mathbf{d}_a$. While not a restriction of our approach, we define one action for every MDP state $s \in S$ (except for the absorbing state s_*), and choose the target point to be the center of its region R_s .³

The MDP must form a *sound abstraction* of the dynamical system. Thus, action $a \in Act$ only exists in an MDP state $s \in S$ if, for every continuous state $\mathbf{x}_k \in R_s$, there exists a control $\mathbf{u}_k \in \mathcal{U}$, such that $\hat{\mathbf{x}}_{k+1} = \mathbf{d}_a$. To impose this constraint, we define the *one-step backward reachable set* $\mathcal{G}(\mathbf{d}_a)$ of action $a \in Act$:

$$\mathcal{G}(\mathbf{d}_a) = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{d}_a = A\mathbf{x} + B\mathbf{u}_k + \mathbf{q}_k, \mathbf{u}_k \in \mathcal{U}\}. \quad (6)$$

Then, action $a \in Act$ exists in state $s \in S$ if and only if $R_s \subseteq \mathcal{G}(\mathbf{d}_a)$. As an example, Fig. 2 shows the backward reachable set of a particular action $a \in Act$, which contains regions R_s and $R_{s'}$, and hence this action is enabled in states s and s' . We write the set of actions that exist in state $s \in S$ as $Act(s) = \{a \in Act \mid R_s \subseteq \mathcal{G}(\mathbf{d}_a)\}$.

Note that the existence of an action in an MDP state merely implies that *for every continuous state* in the associated region, *there exists* a feasible control input that induces this transition. Intuitively, this means that any possible transition in the MDP must also be possible (with the same probability) in the dynamical system. We make the following assumption on the controllability of the dynamical system (Ogata et al., 2010).

Assumption 2. *The dynamical system in Eq. (1) is controllable, meaning that the matrix $\mathcal{C} = [B \ AB \ A^2B \ \dots \ A^{n-1}B]$ has full row rank.*

Intuitively, the dynamics of a controllable system can be ‘excited’ (that is, we can make state transitions) to a subset of the n -dimensional state space that has a nonempty interior over a finite number of time steps. As a result, under Assumption 2 we can, by construction, derive a dynamical system for which the backward reachable set $\mathcal{G}(\mathbf{d}_a)$ has a non-empty interior. For instance, in Eq. (2), the dimension of the control space ($p = 1$) is lower than the dimension of the state space ($n = 2$). In this case, the backward reachable set $\mathcal{G}(\mathbf{d}_a)$ is a line segment in \mathbb{R}^2 and thus has an empty interior. Hence, no region $R_s \in \mathcal{R}$ of the partition can be contained in $\mathcal{G}(\mathbf{d}_a)$, so no action can ever exist in the MDP. However, since the system is assumed to be controllable (note that Assumption 2 is indeed satisfied for Eq. (2)), we can group together two discrete time steps, such that the dimension of the control space becomes equal to that of the state space, i.e., $p = n = 2$. Concretely, we redefine the dynamical system in Eq. (2) as

$$\begin{aligned} \mathbf{x}_{k+2} &= A^2\mathbf{x}_k + [AB \ B] \begin{bmatrix} \mathbf{u}_k \\ \mathbf{u}_{k+1} \end{bmatrix} + A\mathbf{w}_k + \mathbf{w}_{k+1} \\ &= \bar{A}\mathbf{x}_k + \bar{B}\mathbf{u}_{k,k+1} + \mathbf{w}_{k,k+1}, \end{aligned} \quad (7)$$

where \mathbf{x}_k and $\mathbf{u}_{k,k+1}$ now have equal dimension, making $\mathcal{G}(\mathbf{d}_a)$ have a non-empty interior.

3. If, on the one hand, this choice for defining actions results in an MDP that is too large, we may reduce the number of actions; if on the other the results are unsatisfactory, we may define additional actions.

Transition probabilities. To compute the control input \mathbf{u}_k associated with action a in a continuous state \mathbf{x}_k at time k , we replace the successor state \mathbf{x}_{k+1} by target point \mathbf{d}_a in Eq. (1) and solve for \mathbf{u}_k , yielding a control parameterized by state \mathbf{x}_k and action a :

$$\mathbf{u}_k(\mathbf{x}_k, a) = B^+(\mathbf{d}_a - \mathbf{q}_k - A\mathbf{x}_k), \quad (8)$$

where B^+ is the pseudoinverse of B .⁴ It is easily verified that for every state $\mathbf{x}_k \in R_s$ where action $a \in Act(s)$ exists, there exists a \mathbf{u}_k such that Eq. (8) holds (depending on B , it may not be unique). Due to the process noise \mathbf{w}_k , the continuous successor state \mathbf{x}_{k+1} upon choosing an action $a \in Act(s)$ in state $s \in S$ is a random variable, which is written as

$$\begin{aligned} \mathbf{x}_{k+1} &= A\mathbf{x}_k + B\mathbf{u}_k(\mathbf{x}_k, a) + \mathbf{q}_k + \mathbf{w}_k \\ &= \mathbf{d}_a + \mathbf{w}_k. \end{aligned} \quad (9)$$

Remark 4 (Independence of density functions). *Recall that the probability density function of \mathbf{x}_{k+1} under action a is denoted by $p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a))$. Importantly, we observe that by construction of Eq. (9), the continuous successor state \mathbf{x}_{k+1} (and thus also the probability density function) upon choosing any action $a \in Act(s)$, with $s = T(\mathbf{x}_k)$, is independent of the current state \mathbf{x}_k (as long as a is enabled in s). We shall see how this simplifies the complexity of the abstract MDP.*

To define the transition probability function P of the MDP, we want to determine, for every pair of states $s, s' \in S$ and action $a \in Act(s)$, the probability that the state-action pair (s, a) leads to a continuous successor state $\mathbf{x}_{k+1} \in R_{s'}$. This transition probability is equal to the cumulative density function $P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a))$ that the successor state \mathbf{x}_{k+1} takes on a value in region $R_{s'}$ upon choosing action a .

Definition 5. *The transition probability $P(s, a)(s')$ for $s, s' \in S$, $a \in Act(s)$ is defined as*

$$\begin{aligned} P(s, a)(s') &= P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) \\ &= \int_{R_{s'}} p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) d\mathbf{x}_{k+1}. \end{aligned} \quad (10)$$

where $\mathbf{u}_k(\mathbf{x}_k, a)$ is defined as per Eq. (8).

Due to its importance in the arguments of this paper, note that the cumulative density function, denoted by capital $P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'}) \in [0, 1]$ and being a probability, is the integral of the probability density function over region $R_{s'}$, denoted by small $p_{\mathbf{w}_k}(\mathbf{x}_{k+1})$ and being a function that assigns a probability to every possible value of \mathbf{x}_{k+1} .

Remark 5 (Number of transition probabilities of the MDP). *For a given action $a \in Act$ and successor state $s' \in S$, the transition probability $P(s, a)(s')$ is equal for any MDP state s for which $a \in Act(s)$. In other words, for any two states $s, \tilde{s} \in S$ and an action a for which $a \in Act(s)$ and $a \in Act(\tilde{s})$, we have that $P(s, a)(s') = P(\tilde{s}, a)(s')$ for any $s' \in S$. Thus, the MDP has at most $|Act| \cdot |S|$ unique transition probabilities (rather than at most $|S| \cdot |Act| \cdot |S|$ probabilities), which reduces the complexity of generating abstractions.*

4. Even though we assume B to be full row rank, it may have more columns than rows, so we use the pseudoinverse in Eq. (8).

3.3 Soundness of the MDP Abstractions

The generated abstract MDP \mathcal{M} is an *underapproximation* of the dynamical system, in the sense that every transition (s, a, s') of \mathcal{M} is contained in the dynamical system (note that the opposite is not true: the dynamical system contains transitions that are not in the MDP). We formalize this observation using the concept of probabilistic (bi)simulation, originally proposed for probabilistic transition systems by Larsen and Skou (1991).

Bisimilar states. We claim that any two continuous states $\mathbf{x}_k, \mathbf{x}'_k \in T^{-1}(s) = R_s$ are *probabilistically bisimilar* under the set of discrete actions $Act(s)$. Intuitively, a sufficient condition for states $\mathbf{x}_k, \mathbf{x}'_k$ to be probabilistically bisimilar, is that the probability of transitioning to any $\mathbf{x}_{k+1} \in R_{s'}, s' \in S$ is the same for all actions $a \in Act(s)$ (Desharnais, Edalat, & Panangaden, 2002). Mathematically, this is written as follows.

Corollary 1. *Any two states $\mathbf{x}_k, \mathbf{x}'_k \in T^{-1}(s) = R_s$ are probabilistically bisimilar, because for all actions $a \in Act(s)$ and for all successor states $s' \in S$, it holds that*

$$P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) = P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'} \mid \mathbf{x}'_k, \mathbf{u}_k(\mathbf{x}'_k, a)). \quad (11)$$

The proof of Corollary 1 follows directly from Remark 4, which states that the density function $p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) d\mathbf{x}_{k+1}$, and thus also its integral over any discrete region of partition \mathcal{R} , is independent of the current state. Corollary 1 carries an important message: we can abstract any continuous state $\mathbf{x}_k \in R_s$ into a single MDP state $s \in S$ *without losing any information about the probabilistic behavior of the model*.

Probabilistic simulations. Second, we claim that our abstraction procedure creates a *probabilistic feedback refinement relation* (and thus a *probabilistic simulation relation*; see Hermanns, Parma, Segala, Wachter, & Zhang, 2011) from the generated MDP to the dynamical system (Reissig, Weber, & Rungger, 2017; Haesaert, Soudjani, & Abate, 2017).⁵ Intuitively, a sufficient condition for such a relation to exist is that for any pair (\mathbf{x}, s) of related states, i.e., for which $T(\mathbf{x}) = s$, and for all actions $a \in Act(s)$, there exists a control input $\mathbf{u} \in \mathcal{U}$ such that the probability of transitioning to any state s' in the MDP equals the probability of transitioning to any $\mathbf{x}' \in T^{-1}(s') = R_{s'}$ in the dynamical system. Based on the definition from Reissig et al. (2017),⁶ we mathematically write this condition as follows.

Corollary 2. *The abstraction function $T: \mathbb{R}^n \rightarrow S$ induces a probabilistic feedback refinement relation from the MDP to the dynamical system, because $T(\mathbf{x}_0) = s_I$ (i.e., the initial states match), and for any pair (\mathbf{x}_k, s) of states related as $T(\mathbf{x}_k) = s$, it holds that*

$$\forall a \in Act(s), \forall s' \in S : P(s, a)(s') = \int_{R_{s'}} p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) d\mathbf{x}_{k+1}. \quad (12)$$

5. The latter article employs alternating approximate relations, rather than simulation relations. The resulting refinement step is, however, very similar.

6. The definition in Reissig et al. (2017) also contains a condition for matching observations between the two models, which we drop because we consider fully observable systems. Similarly, Haesaert et al. (2017) has an error metric on observations that does not apply to our setting. We also expand the original definition from non-probabilistic to probabilistic systems. Note that Corollary 2 defines sufficient conditions for a feedback refinement relation, which are stricter than those in Reissig et al. (2017).

Corollary 2 follows directly from Def. 5. Crucially, this probabilistic feedback refinement relation implies that *any reach-avoid problem satisfiable for the MDP is also satisfiable for the dynamical system with the same probability bound* (Hermanns et al., 2011).

4. Sampling-Based Probability Intervals

Recall that the probability density function $p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a))$ is unknown, making a direct evaluation of Eq. (10) impossible. In this paper, we use a sampling-based method to estimate the transition probabilities described in Eq. (10), which requires a finite set of N observations of the process noise, $\mathbf{w}_k^{(i)} \in \Delta$, $i = 1, \dots, N$. Each sample has a unique index $i = 1, \dots, N$ and is associated with a possible successor state $\mathbf{x}_{k+1}^{(i)} = \mathbf{d}_a + \mathbf{w}_k^{(i)}$ under a particular action a . We assume that these samples are available from experimental data or simulations. As such, we can generate these samples by inferring the process noise from obtained state trajectories of the dynamical system in Eq. (1), or we may sample from the noise distribution directly (e.g., if a simulator is available). Thus, in our setting, samples of the noise can be obtained at a relatively low cost.

Recall that the process noise affecting Eq. (1) is considered to be independent and identically distributed (i.i.d.). Due to the importance of this practically reasonable assumption to the arguments we develop in this section, we formally enshrine it in the following.

Assumption 3. *The process noise \mathbf{w}_k is i.i.d., and we assume to have a collection of N samples from it, which we denote by the discrete set $\mathbf{w}_k^{(i)} \in \Delta$, $i = 1, \dots, N$.*

Due to Assumption 3, the set $\mathbf{w}_k^{(1)}, \dots, \mathbf{w}_k^{(N)}$ of N samples is a random element from the probability space Δ^N equipped with the product probability \mathbb{P}^N and the product σ -algebra.

A frequentist approach to estimation. As an example, we want to estimate the probability $P(s, a)(s')$ that some enabled state-action pair (s, a) induces a transition to state $s' \in S$. A common approach to approximate the transition probability is to count the number of samples $N_{s'}^{\text{in}} \leq N$ leading to a transition to state s' and divide it by the total of N samples. This approach is known as a *frequentist approach*, and is statistically justified by the strong law of large numbers. Concretely, the value of $N_{s'}^{\text{in}}$ is obtained as follows.

Definition 6. *The cardinality $N_{s'}^{\text{in}} \in \{0, \dots, N\}$ of the index set of the samples leading to a successor state $\mathbf{x}_{k+1} \in R_{s'}$ associated with MDP state $s' \in S$ is defined as*

$$N_{s'}^{\text{in}} = \left| \{i \in \{1, \dots, N\} \mid (\mathbf{d}_a + \mathbf{w}_k^{(i)}) \in R_{s'}\} \right|. \quad (13)$$

Similarly, we define $N_{s'}^{\text{out}} = N - N_{s'}^{\text{in}}$ as the number of samples for which $\mathbf{d}_a + \mathbf{w}_k^{(i)} \notin R_{s'}$.

Note that $N_{s'}^{\text{in}}$ and $N_{s'}^{\text{out}}$ depend on both the set of noise samples $\mathbf{w}_k^{(1)}, \dots, \mathbf{w}_k^{(N)}$ and on the action taken. The frequentist approach is simple, but may lead to estimates that deviate critically from their true values if the number of samples is limited (we illustrate this issue in the UAV experiment in Sect. 6.2). In what follows, we discuss how to render our method robust against such estimation errors.

4.1 Sampling Techniques From the Scenario Approach

As a key contribution, we present a method based on the scenario approach (Campi & Garatti, 2018) to compute *intervals of probabilities* instead of precise estimates. Specifically, for every transition (s, a, s') , we compute an upper and lower bound, i.e., an interval, that contains the quantity $P(s, a)(s')$ with a user-specified (high) confidence probability. We formalize the resulting abstraction as an iMDP, where these intervals enter the uncertain transition function $\mathcal{P}: S \times Act \times S \rightarrow \mathbb{I}$. As the intervals are PAC, this iMDP is a robust abstraction of the dynamical system.

Outline: PAC intervals for transition probabilities. Intuitively, we recast the estimation of transition probabilities into a convex optimization problem that includes a constraint for each element of a subset of noise samples. To obtain PAC intervals on transition probabilities, we leverage recent results from Romao et al. (2022) on the probability of *constraint violation* for such optimization problems. Crucially, we show that the problem can be solved based on its geometry by an analytical *counting argument*. This counting argument makes our approach applicable to abstractions with hundreds of millions of transitions, as we do not have to solve any optimization problem explicitly. Toward our main result for computing probability intervals (which is presented in Theorem 1), we introduce a number of core concepts of the scenario approach. Interestingly, due to the counting argument, Theorem 1 does not directly involve these concepts (only its derivation does, which we provide in Appendix A), so the reader may decide to skip directly to Sect. 4.2.

Risk of violation. First, we introduce the concept of *risk* (or *violation probability*), which is a measure of the probability that a successor state \mathbf{x}_{k+1} is not in a certain subset $\tilde{R} \subset \mathbb{R}^n$ upon choosing an action $a \in Act(s)$ in state $s \in S$ (Campi & Garatti, 2008).

Definition 7. The risk $P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \notin \tilde{R} | \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a))$ that \mathbf{x}_{k+1} is not in region \tilde{R} is

$$P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \notin \tilde{R} | \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) = \mathbb{P}\{\mathbf{w}_k \in \Delta : \mathbf{d}_a + \mathbf{w}_k \notin \tilde{R}\}. \quad (14)$$

Crucially, observe from Eq. (10) that the transition probability $P(s, a)(s')$ that we aim to estimate is the complement of the violation probability over the fixed region $R_{s'}$, i.e.,

$$P(s, a)(s') = P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \in R_{s'} | \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) = 1 - P_{\mathbf{w}_k}(\mathbf{x}_{k+1} \notin R_{s'} | \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)). \quad (15)$$

Scenario optimization problem. The scenario approach enables us to bound the risk over the feasible set of the optimal point of a so-called *scenario optimization problem*. Specifically, we consider scenario problems with *discarded constraints*; see Campi and Garatti (2011) for details. Roughly speaking, solving this problem amounts to finding, over a scalar decision variable $\lambda \in \mathbb{R}_+$, a convex set $R_s(\lambda)$ of minimal size that contains a particular number of successor state samples (shortly, we shall define the geometry of $R_s(\lambda)$ in relation to state $s \in S$ and λ). Concretely, the optimization problem is given as

$$\begin{aligned} \mathfrak{L}_Q : \text{minimize } \lambda & \\ \lambda \in \mathbb{R}_+ & \\ \text{subject to } \mathbf{d}_a + \mathbf{w}_k^{(i)} \in R_s(\lambda) \quad \forall i \in \{1, \dots, N\} \setminus Q, & \end{aligned} \quad (16)$$




Figure 3: Polytope $R_{s'}$ of state $s' \in S$ has a Chebyshev center $\tilde{\mathbf{h}}_{s'}$ (note that it is not unique, as the circle can be shifted while remaining within $R_{s'}$). Polytope $R_{s'}(1.2)$ is scaled by a factor $\lambda = 1.2$ and is computed using Eq. (17).




Figure 4: Bounding region $R_{s'} = [-1, 1]$ using $N = 10$ successor state samples ($N_{s'}^{\text{out}} = 5$). Discarding $N_{s'}^{\text{out}} = 5$ samples defines the red region $R_{s'}(\lambda_{N_{s'}^{\text{out}}}^*) \subseteq R_{s'}$; discarding one less sample defines the blue region $R_{s'}(\lambda_{N_{s'}^{\text{out}}-1}^*) \supset R_{s'}$.

where we explicitly write the dependency on the set $Q \subset \{1, \dots, N\}$, which is a subset of samples whose constraints have been discarded. The optimal solution $\lambda_{|Q|}^*$ to problem \mathfrak{L}_Q parameterizes a feasible set $R_s(\lambda_{|Q|}^*)$ over which we can bound the risk using the scenario approach theory. However, to compute a transition probability $P(s, a)(s')$, we must shape the feasible set such that $R_s(\lambda_{|Q|}^*)$ is closely related to the fixed region $R_{s'}$. As explained below, we achieve this by 1) defining $R_s(\lambda)$ as a scaled version of R_s , and 2) appropriately choosing the subset of discarded samples Q .

Scaled polytopes. We define $R_s(\lambda)$ as a version of polytope $R_s \in \mathcal{R}$ (recall from Sect. 3 that R_s is defined by matrix M_s and vector \mathbf{b}_s) which is *scaled* by a factor $\lambda \geq 0$ relative to a so-called Chebyshev center $\tilde{\mathbf{h}}_s \in \mathbb{R}^n$ (Boyd & Vandenberghe, 2014). As such, we obtain

$$R_s(\lambda) = \{\mathbf{x} \in \mathbb{R}^n \mid M_s \mathbf{x} \leq \lambda(\mathbf{b}_s + \tilde{\mathbf{h}}_s) - \tilde{\mathbf{h}}_s\}. \quad (17)$$

Note that $R_s(1) = R_s$, and that shifting by $\tilde{\mathbf{h}}_s$ ensures that we are scaling around a point that is by construction contained in R_s , such that $R_s(\lambda_1) \subset R_s(\lambda_2)$ for every $0 \leq \lambda_1 < \lambda_2$. A visualization of the scaling for an arbitrary region $R_{s'} \in \mathcal{R}$ associated with a discrete successor state $s' \in S$ is shown in Fig. 3. Note that, in this example, the Chebyshev center is not unique since the circle can be shifted while remaining within $R_{s'}$.

Set of discarded samples. Let us now determine the subset Q of samples whose constraints have been discarded from problem \mathfrak{L}_Q . We obtain this set by iteratively removing individual samples based on the following rule:

Proposition 1. *The sample removal set $Q \subset \{1, \dots, N\}$ is obtained by iteratively removing the active constraints from Eq. (16). Thus, given N samples and any two removal sets with cardinalities $|Q_1| < |Q_2|$, it holds that $Q_1 \subset Q_2$. Moreover, any discarded sample $i \in Q$ violates the solution λ_Q^* to Eq. (16), i.e., $\mathbf{d}_a + \mathbf{w}_k^{(i)} \notin R_{s'}(\lambda_{|Q|}^*)$, with probability one.*

Intuitively, the successor state of a sample associated with an active constraint is on the boundary of the optimal solution $R_{s'}(\lambda_{|Q|}^*)$ of Eq. (16). Under the following *non-accumulation* assumption, an active constraint exists and is unique, as long as $|Q| < N$ (i.e., not all samples have been discarded).

Assumption 4. *Given a noise sample $w_k \in \Delta$, the probability that the successor state x_{k+1} is on the boundary of any polytope $R_{s'}(\lambda)$ is zero, for any $s' \in S$ and $\lambda \in \mathbb{R}_+$.*

Note that Assumption 4 holds for most of the smooth probability distributions commonly encountered in practice, e.g., Gaussian distributions, and is thus easily satisfied. Assumption 4 implies that the solution to Eq. (16) is unique with probability one and that the number of active constraints is equal to one (given $N > 0$ and $|Q| < N$), as samples accumulate on the boundary of the polytope with probability zero.

Under/over-approximating regions. Recall from Def. 6 that $N_{s'}^{\text{out}}$ is the number of samples leading to a successor state outside of region $R_{s'}$ for discrete state $s' \in S$. The following lemma uses $N_{s'}^{\text{out}}$ to define an under- or over-approximation of region $R_{s'}$.

Lemma 1. *When discarding $|Q| = N_{s'}^{\text{out}}$ noise samples as per proposition 1, it holds that $R_{s'}(\lambda_{N_{s'}^{\text{out}}}^*) \subseteq R_{s'}$. When discarding $|Q| = N_{s'}^{\text{out}} - 1$ samples, it holds that $R_{s'}(\lambda_{N_{s'}^{\text{out}}-1}^*) \supset R_{s'}$*

Proof. proposition 1 states that at every step, we may only discard the sample of the active constraint. By construction, after discarding $|Q| = N_{s'}^{\text{out}}$ samples, all remaining successor state samples lie within $R_{s'}$, so $\lambda_{N_{s'}^{\text{out}}}^* \leq 1$, so $R_{s'}(\lambda_{N_{s'}^{\text{out}}}^*) \subseteq R_{s'}$. When we discard one less sample, i.e., $|Q| = N_{s'}^{\text{out}} - 1$, we must have one sample outside of $R_{s'}$, so $\lambda_{N_{s'}^{\text{out}}-1}^* > 1$, meaning that $R_{s'}(\lambda_{N_{s'}^{\text{out}}-1}^*) \supset R_{s'}$. This concludes the proof. \square

Intuitively, $R_{s'}(\lambda_{N_{s'}^{\text{out}}}^*)$ contains *exactly* those samples in $R_{s'}$, while $R_{s'}(\lambda_{N_{s'}^{\text{out}}-1}^*)$ additionally contains *the sample closest outside of $R_{s'}$* , as visualized in Fig. 4 for a 1-dimensional example. By using the scenario approach theory to bound the risk over these regions, we can thus also obtain bounds on the transition probability $P(s, a)(s')$, as per Eq. (15).

4.2 Bounds for the Transition Probabilities

Based on the techniques from the scenario approach introduced above, we state the main contribution of this section, as a non-trivial variant of Romao et al. (2022, Theorem 5), adapted to our new context. Specifically, for a given transition (s, a, s') and the resulting number of samples $N_{s'}^{\text{out}}$ outside of region $R_{s'}$ (as per Def. 6), Theorem 1 returns an interval $[\underline{p}, \bar{p}]$ that contains $P(s, a)(s')$ with at least a pre-defined confidence probability $(1 - \beta)$.

Theorem 1 (PAC probability intervals). *For $N \in \mathbb{N}$ samples of the noise, fix a confidence parameter $\beta \in (0, 1)$. Given $N_{s'}^{\text{out}}$, the transition probability $P(s, a)(s')$ is bounded by*

$$\mathbb{P}^N \left\{ \underline{p} \leq P(s, a)(s') \leq \bar{p} \right\} \geq 1 - \beta, \quad (18)$$

where $\underline{p} = 0$ if $N_{s'}^{\text{out}} = N$, and otherwise \underline{p} is the solution of

$$\frac{\beta}{2N} = \sum_{i=0}^{N_{s'}^{\text{out}}} \binom{N}{i} (1 - \underline{p})^i \underline{p}^{N-i}, \quad (19)$$

and $\bar{p} = 1$ if $N_{s'}^{out} = 0$, and otherwise \bar{p} is the solution of

$$\frac{\beta}{2N} = 1 - \sum_{i=0}^{N_{s'}^{out}-1} \binom{N}{i} (1-\bar{p})^i \bar{p}^{N-i}. \quad (20)$$

We present the proof in Appendix A. Theorem 1 states that, with a probability of at least $1 - \beta$, the probability $P(s, a)(s')$ is bounded by the obtained interval $[\underline{p}, \bar{p}]$. Importantly, this claim holds for *any* Δ and \mathbb{P} (given the previous assumptions), so we can bound the probability in Eq. (10), even when the probability distribution of the noise is unknown.

Remark 6 (Beta distribution). *Note that Eqs. (19) and (20) are cumulative distribution functions of a beta distribution with parameters $N_{s'}^{out} + 1$ (or $N_{s'}^{out}$) and $N - N_{s'}^{out}$ (or $N - N_{s'}^{out} - 1$), respectively (Campi & Garatti, 2018), which can directly be solved numerically for \underline{p} or \bar{p} up to arbitrary precision. To speed up the computations at run-time, we apply a tabular approach to compute the intervals for all relevant values of N , β , and $N_{s'}^{out}$ upfront.*

Counting argument. As explained in Appendix A, the proof of Theorem 1 is based on the solutions to a set of N optimization problems. Interestingly, as also shown in the proof, we can solve these optimization problems analytically based on their geometry. As a result, Theorem 1 only requires the sample count $N_{s'}^{out}$, the total number of samples N , and the confidence parameter β , which are all quantities that are *independent of the solutions to these optimization problems*. Remarkably, this implies that we can compute PAC probability intervals without solving any optimization program explicitly using a solver. Since we only need the values of $N_{s'}^{out}$, N , and β , our method is in practice almost as simple as the frequentist approach but has the notable advantage that we obtain robust intervals of probabilities. As shown in our experiments in Sect. 6, this means we can use our abstraction procedure to generate iMDPs with hundreds of millions of transitions.

4.3 Tightness of Probability Intervals

Let us now explore the tightness of the obtained probability intervals. We can interpret a transition probability $P(s, a)(s')$ as the probability of a Bernoulli random variable. In this context, we may use Hoeffding’s inequality, a well-known concentration inequality, to infer PAC bounds on this probability (Boucheron, Lugosi, & Massart, 2013). In particular, given N successor state samples of which $N_{s'}^{in}$ are contained in region $R_{s'}$, Hoeffding’s inequality states that, for a pre-defined $\beta \in (0, 1)$, it holds that

$$\mathbb{P}^N \left\{ \frac{N_{s'}^{in}}{N} - \varepsilon \leq P(s, a)(s') \leq \frac{N_{s'}^{in}}{N} + \varepsilon \right\} \geq 1 - \beta, \quad (21)$$

where $\varepsilon = \sqrt{\frac{1}{2N} \log(\frac{2}{\beta})}$. In what follows, we evaluate the tightness of our intervals obtained from Theorem 1, versus those obtained from Hoeffding’s inequality.

Number of noise samples N . To illustrate how the choice for the number of samples N affects the tightness of the intervals, consider a system with a 1-dimensional state $\mathbf{x}_k \in \mathbb{R}$, where the probability density function $p_{w_k}(\mathbf{x}_{k+1} | \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a))$ for a specific action $a \in Act$ with target point \mathbf{d}_a is given by a uniform distribution over the domain $[-4, 4]$. For a

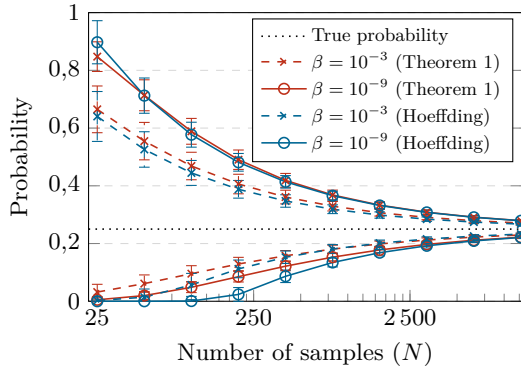


Figure 5: Probability intervals (with standard deviation) obtained from Theorem 1 versus Hoeffding’s inequality, with $\beta = 10^{-3}$ or 10^{-9} on a transition with a true probability of 0.25 (note the log scale).




Figure 6: Probability intervals obtained from Theorem 1, with $\beta = 10^{-9}$ and different values for $N_{s'}^{\text{out}} \in [0, N]$, versus probability intervals obtained from Hoeffding’s inequality for the same value of β .

given region $R_{s'} = [-1, 1]$ (also shown in Fig. 4), we want to evaluate the probability that $\mathbf{x}_{k+1} \in R_{s'}$, which is 0.25. To this end, we apply Theorem 1 for different numbers of samples $N \in [25, 12\,800]$ and a confidence level of $\beta = 10^{-3}$ or 10^{-9} . The obtained probability bounds are random variables through their dependence on the samples, so we repeat each experiment 100 000 times, resulting in the probability intervals shown in Fig. 5. We observe that the uncertainty in the transition probability is reduced by increasing the number of samples. Moreover, the lower bounds obtained from Theorem 1 are better than those obtained from Hoeffding’s inequality, while the converse holds for the upper bounds.

Sample count $N_{s'}^{\text{out}}$. To explain why Theorem 1 yields tighter lower bounds, while Hoeffding’s inequality yields tighter upper bounds, we plot in Fig. 6 the resulting probability intervals for $N = 800$ samples and different values of $N_{s'}^{\text{out}} \in [0, N]$ (recall that $N = N_{s'}^{\text{out}} + N_{s'}^{\text{in}}$). We observe our scenario-based approach results in *significantly tighter* intervals for values of $N_{s'}^{\text{out}}$ close to 0 and N . On the other hand, Hoeffding’s inequality leads to *slightly better* intervals for moderate values of $N_{s'}^{\text{out}}$ around $\frac{N}{2}$. As observed from Eq. (21), Hoeffding’s inequality yields bounds given by the *sample mean* $N_{s'}^{\text{in}}$, plus or minus a *fixed* value of ε which is *independent of the sample mean*. This property explains why Hoeffding’s inequality leads to poor probability intervals if the probability $P(s, a)(s')$ is close to zero or one.

4.4 iMDP Abstractions With PAC Guarantees

We describe how we apply Theorem 1 to solve the overall problem statement. Since the process noise is independent of the state and time, we require *only N samples in total*.⁷ We

7. If the noise distribution is time-varying, the transition probabilities are time-dependent. In this case, we have at most $|\text{Act}| \cdot |\mathcal{S}| \cdot K$ unique probabilities and must use N noise samples for each step $k = 0, \dots, K-1$. For finite-horizon properties, adapting our abstraction method for this case is straightforward.




Figure 7: Bounds $[p, \bar{p}]$ on the probabilities $P(s, a)(s')$ for 3 regions using $N = 100$ samples (black ticks) and $\beta = 0.01$. The probability density function over successor states is $p_{w_k}(x_{k+1} | x_k, u_k(x_k, a))$. Point estimate probabilities are computed as $N_{s'}^{\text{in}}/N$.

can use the same samples to compute multiple intervals by shifting these samples by the appropriate target point d_a of each action $a \in \text{Act}$. Recall from Remark 5 that the abstract MDP has at most $|\text{Act}| \cdot |S|$ unique transition probabilities. For every unique probability, we determine the successor state samples $x_{k+1}^{(1)}, \dots, x_{k+1}^{(N)}$ under the N samples of the noise. For every possible successor state $s' \in S$, we then determine $N_{s'}^{\text{out}}$ and invoke Theorem 1 to compute the PAC bounds on $P(s, a)(s')$ that hold for every $s \in S$. Fig. 7 shows this process, where every tick is a successor state x_{k+1} under a noise sample. This figure also shows point estimates of the probabilities derived using the frequentist approach. If no samples are observed in a region, we assume that $P(s, a)(s') = 0$.

Theorem 2. *Given a dynamical system in Eq. (1), generate an iMDP abstraction $\mathcal{M}_{\mathbb{I}}$, and compute the robust reach-avoid probability $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K)$ under the optimal policy π^* . Under the controller ϕ that applies control inputs according to Eq. (8) for each $k < K$, it holds that*

$$\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) \geq \eta \implies \mathbb{P}^N \left\{ Pr^{\phi}(\varphi_{x_0}^K) \geq \eta \right\} \geq 1 - \alpha, \quad (22)$$

where $\alpha = \beta \cdot |\text{Act}| \cdot |S|$ upper bounds the number of unique probability intervals.

Proof. Denote by $\mathcal{P}: S \times \text{Act} \times S \rightarrow \mathbb{I} \cup \{[0, 0]\}$ the uncertain transition function of the generated iMDP, and denote by $P: S \times \text{Act} \rightarrow \text{Dist}(S)$ the unknown transition function, defined by Eq. (10), of the underlying MDP. For each $a \in \text{Act}$ and $s' \in S$, it holds that

$$\mathbb{P}^N \left\{ P(s, a)(s') \in \mathcal{P}(s, a, s'), \forall s \in S \right\} \geq 1 - \beta. \quad (23)$$

Recall from Remark 5 that the iMDP has at most $|\text{Act}| \cdot |S|$ unique probability intervals. Using Boole's inequality,⁸ we thus obtain the following correctness guarantee for the iMDP:

$$\mathbb{P}^N \left\{ P(s, a)(s') \in \mathcal{P}(s, a, s), \forall s, s' \in S, a \in \text{Act} \right\} = \mathbb{P}^N \{ P \in \mathcal{P} \} \geq 1 - \alpha, \quad (24)$$

8. Boole's inequality (or the union bound) says that the probability that at least one of a finite set of events happens, is upper bounded by the sum of the probabilities of these events (Casella & Berger, 2021).

where $\alpha = \beta \cdot |Act| \cdot |S|$. Next, observe that for any MDP instantiation $P \in \mathcal{P}$, it holds by construction of Eq. (4) that $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) \leq Pr^{\pi^*}(\varphi_{s_I}^K)$. Moreover, Corollary 2 asserts that $Pr^{\pi^*}(\varphi_{s_I}^K) = Pr^\phi(\varphi_{x_0}^K)$ (the reach-avoid probability on the abstract MDP equals the reach-avoid probability on the dynamical system), so we obtain

$$\mathbb{P}^N \left\{ \underline{Pr}^{\pi^*}(\varphi_{s_I}^K) \leq Pr^\phi(\varphi_{x_0}^K) \right\} \geq 1 - \alpha. \quad (25)$$

Letting $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) \geq \eta$, we arrive at the implication in Eq. (22), so we conclude the proof. \square

We remark that the tightness of Eq. (22) in Theorem 2 depends on the value of α , which in turn depends on the size of the generated iMDP (through $|S|$ and $|Act|$), and the confidence level β . For large iMDPs, one must choose a stronger confidence level (i.e., β closer to zero), to obtain an informative bound $1 - \alpha$, which is close to one. The following corollary provides a tighter bound if the iMDP is generated in a very specific manner.

Corollary 3. *If the iMDP in Theorem 2 is generated under a uniform rectangular partition \mathcal{R} (i.e., a grid) into r_i regions in each dimension $i = 1, \dots, n$ of the state space \mathbb{R}^n , and with one action $a \in Act$ per region $R \in \mathcal{R}$ whose target point \mathbf{d}_a is the center of R , then the confidence parameter becomes $\alpha = \beta(\prod_{i=1}^n (2r_i - 1) + |Act|)$.*

The proof of Corollary 3 is provided in Appendix B and is analogous to the proof of Theorem 2, with the only difference that the number of unique probability intervals is reduced. The key observation is that transition probabilities $P(s, a)(s')$ of the abstraction are defined by the relative position $R_{s'} - \mathbf{d}_a$ between region $R_{s'}$ and target point \mathbf{d}_a . In particular, the corollary exploits the symmetry between the partition and the target points, which reduces the number of unique transition probabilities significantly and leads to stronger confidence levels compared to Theorem 2.

5. Overall Robust Control Algorithm

In Sects. 3 and 4, we have introduced tools for generating an iMDP abstraction of the dynamical system in Eq. (1), which is correct with a user-specified confidence probability. We now discuss in more detail how we use these tools to solve the overall problem statement posed in Sect. 2.2. Recall that our approach, shown in Fig. 1, consists of an offline and an online phase. In what follows, we present an algorithm for both phases.

5.1 Offline Planning Phase

The offline planning phase is presented in Algorithm 1. We first generate an MDP abstraction by defining its states S , actions Act , and initial state $s_I \in S$ (line 1). For each state $s \in S$, we compute the set $Act(s)$ of enabled actions (line 2) and define the initial values for the sample size N and for the reachability probabilities (line 3). Recall that computing the transition probabilities of this MDP is not possible, so we instead compute intervals of probabilities and formalize the abstract model as an iMDP instead.

We then enter the iterative part of our approach. We first obtain N samples of the noise (line 5). Recall from Sect. 4 that we can obtain these samples by inferring the process noise from previously generated state trajectories of the dynamical system or by sampling the

Algorithm 1 Offline planning (generate iMDP and compute optimal policy)

Input: Linear dynamical model; property $\varphi_{s_I}^K$ with probability threshold η
Params: Partition \mathcal{R} ; confidence lvl. β ; increment factor γ
Output: Optimal policy π^*

```

1: Define states  $S$ , actions  $Act$ , and initial state  $s_I = T(\mathbf{x}_k)$ 
2: Define enabled actions  $Act(s) \subseteq Act$  for all  $s \in S$ 
3: Set initial values:  $N = N_0$ ,  $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) = 0$ ,  $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K) = 1$ 
4: while  $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) < \eta \leq \overline{Pr}^{\pi^*}(\varphi_{s_I}^K)$  do
5:   Obtain  $N$  samples  $\mathbf{w}_k^{(1)}, \mathbf{w}_k^{(N)} \in \Delta$ 
6:   for all actions  $a$  in  $Act$  do
7:     for all successor states  $s'$  in  $S$  do
8:       Compute  $[p, \bar{p}]$  by applying Theorem 1 for  $N$ ,  $N_{s'}^{out}$ , and  $\beta$ 
9:       Let  $\mathcal{P}(s, a, s') = [p, \bar{p}] \forall s \in S$  such that  $a \in Act(s)$ 
10:    end for
11:  end for
12:  Store iMDP  $\mathcal{M}_{\mathbb{I}} = (S, Act, s_I, \mathcal{P})$ 
13:  Compute  $\pi^*$ ,  $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K)$ , and  $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K)$  on  $\mathcal{M}_{\mathbb{I}}$  using PRISM
14:  Let  $N = \gamma N$ 
15: end while
16: if  $\eta > \overline{Pr}^{\pi^*}(\varphi_{s_I}^K)$  then
17:   return Unsatisfiable
18: else
19:   return optimal policy  $\pi^*$ 
20: end if
    
```

noise distribution directly using a simulator. For every action $a \in Act$ and successor state $s' \in S$, we compute a PAC transition probability interval $[p, \bar{p}]$ using Theorem 1 (line 8). These intervals are used to define the transition function \mathcal{P} of the abstract iMDP (line 9).

Recall from Sect. 2.4 that for the resulting iMDP (stored in line 12), we leverage PRISM to obtain the optimal policies that maximize lower and upper bounds on the reach-avoid probability using Eqs. (4) and (5) (line 13). If the maximum lower bound reach-avoid probability $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K)$ under this policy is above the required threshold η , then the algorithm returns the optimal policy π^* and proceeds to the online control phase. Otherwise, if $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) < \eta$ but the problem is still satisfiable (i.e., $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K) \geq \eta$) we obtain additional samples by increasing N by a fixed factor $\gamma > 1$ (line 14) and repeat the while loop. If instead, it holds that $\overline{Pr}^{\pi^*}(\varphi_{s_I}^K) < \eta$, the reach-avoid problem is not satisfiable for any MDP instantiated by $P \in \mathcal{P}$. In this case, the algorithm terminates without returning a policy.

The main computational complexity of Algorithm 1 lies within the double for loop. In particular, generating one iMDP abstraction⁹ (lines 5-12 of Algorithm 1) has the following complexity with respect to the numbers of states $|S|$ and actions $|Act|$ (which are directly controlled through the partitioning of the state space), and the number of noise samples N .

9. Verifying iMDPs can be done in polynomial time, and we refer to Puggelli et al. (2013) for details.

Algorithm 2 Online control (on-the-fly controller synthesis)

Input: Optimal policy π^* ; initial state \mathbf{x}_0 ; property $\varphi_{\mathbf{x}_0}^K$ **Output:** SAT (Boolean)

```

1: Let  $k = 0$ 
2: while  $k \leq K$  do
3:   if  $\mathbf{x}_k \in \mathcal{X}_G$  then
4:     return SAT = True
5:   else if  $\mathbf{x}_k \in \mathcal{X}_C$  then
6:     return SAT = False
7:   else
8:     Let optimal action  $a^* = \pi^*(s, k)$  for state  $s = T(\mathbf{x}_k)$ 
9:     Compute control input  $u_k(a^*)$  using Eq. (8)
10:    Sample state  $\mathbf{x}_{k+1} \sim p_{w_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(a))$ 
11:   end if
12:    $k = k + 1$ 
13: end while
14: return SAT = False

```

Theorem 3 (Complexity of generating abstractions). *The worst-case complexity of generating one iMDP abstraction using Algorithm 1 is $\mathcal{O}(N \cdot |\text{Act}| + |S|^2 \cdot |\text{Act}|)$.*

Proof. For every action $a \in \text{Act}$, we must check for each state $s \in S$ if a is enabled in s , leading to $\mathcal{O}(|S| \cdot |\text{Act}|)$ operations. Next, for each action $a \in \text{Act}$, we must determine for each of the N samples to which state it belongs, leading to $\mathcal{O}(N \cdot |\text{Act}|)$, and subsequently, we compute for each action $a \in \text{Act}$ and successor state $s' \in S$ the probability interval $[p, \bar{p}]$, leading to $\mathcal{O}(|\text{Act}| \cdot |S|)$. Finally, we store each transition of the iMDP, which has at most $|S|^2 \cdot |\text{Act}|$ transitions. By summing these contributions and only keeping the highest order terms, we obtain the order of complexity in Theorem 3 and conclude the proof. \square

We remark that the number of states $|S|$ and $|\text{Act}|$ depend on the partition and target points used to generate the abstraction; see Sect. 3 for details. For example, under the rectangular partitioning described by Corollary 3, the number of states is $|S| = \prod_{i=1}^n + 1$ and the number of actions is $|\text{Act}| = \prod_{i=1}^n$. This result shows that our abstraction procedure (like any other discretization-based technique) suffers from the well-known *curse of dimensionality* (Lavaei et al., 2022), i.e., the complexity is exponential in the dimension n of the continuous state space \mathbb{R}^n .

5.2 Online Control Phase

In the online control phase, we synthesize a controller for the dynamical system on the fly, based on the policy π^* returned by Algorithm 1. Recall that this policy is a time-varying map from iMDP states to actions. Intuitively, we translate the policy to a time-varying feedback controller that is piece-wise linear (namely, linear in the state \mathbf{x}_k within each region of the partition). Concretely, at each time step, we use Def. 4 to determine the current iMDP state $s = T(\mathbf{x}_k)$ to which the continuous state \mathbf{x}_k belongs, and then retrieve

the optimal action $a^* = \pi^*(s, k)$ from the policy (line 8). We compute the associated control input using Eq. (8), which is valid by construction, followed by sampling the successor state \mathbf{x}_{k+1} (lines 9-10). The algorithm returns $\text{SAT} = \text{True}$ if the goal region is reached within K steps (i.e., the reach-avoid problem is satisfied), while it returns $\text{SAT} = \text{False}$ if either the critical region is reached or the system fails to reach the goal region within K steps.

Remark 7 (Backup controller). *The controller obtained via Algorithm 1 only yields control inputs for states in the bounded portion of the state space $\mathcal{X} = \mathbb{R}^n \setminus R_*$, which is covered by the partition \mathcal{R} . To alleviate this conservatism, one may additionally design a backup controller, which is activated upon leaving \mathcal{X} and aims to return the state \mathbf{x}_k to the bounded portion \mathcal{X} for which the controller generated control inputs. In the abstract iMDP, visiting the absorbing state s_* (which corresponds to the absorbing region $R_* = \mathbb{R}^n$) means that the reach-avoid property is violated by definition. As such, deploying the backup controller can only increase the probability $Pr^\phi(\varphi_{\mathbf{x}_0}^K)$ of satisfying the reach-avoid property, and thus, Theorem 2 still holds regardless of the designed backup controller.*

5.3 Reducing the Complexity of the iMDP Policy Synthesis

The iMDP abstractions generated using our approach can potentially have hundreds of millions of transitions, especially if the state dimension is high (see Tables 1 and 2). The main bottleneck that leads to such large models, is that the number of transitions in the iMDP is (worst-case) quadratic in the number of states, and linear in the number of actions. Indeed, note that Algorithm 1 consists of a double for-loop, enumerating over both the set of actions and the set of states, which can thus be expensive.

To alleviate this limitation, we develop an improved policy synthesis scheme that reduces the complexity of the Bellman iterations required to compute the optimal policy over the iMDP.¹⁰ This scheme is *sound*, in the sense that the maximum lower bound reach-avoid probability $\underline{Pr}^{\pi^*}(\varphi_{s_f}^K)$ that we obtain under the proposed scheme can be more conservative, but the correctness guarantee in Theorem 2 still holds. In practice, instead of working with one large abstraction across the whole horizon of the reach-avoid property, we work with a smaller iMDP at each time step, by merging states that are associated to similar reach-avoid probabilities computed via the Bellman iterations. In what follows, we first explain how we generate and verify this reduced model for a single time step, by merging (aggregating) states with similar reach-avoid probabilities. Thereafter, we describe how we iterate backward over the whole time horizon, as needed to synthesize the policy.

Remark 8 (Restriction to finite horizons). *As the improved policy synthesis scheme unfolds the iMDP over each time step, the scheme is only applicable to finite-horizon properties.*

State aggregation for a single time step. As an illustrative example, consider the iMDP in Fig. 8, which consists of four states: s_1 is a goal state, s_4 is a critical state, and s_2 and s_3 are neither. Recall that K denotes the time horizon of the reach-avoid property. If we are at time step $k = K - 1$ (i.e., only one action to go), the only way to satisfy the reach-avoid problem is to reach state s_1 . In other words, the reward (being the probability

¹⁰. While the proposed scheme reduces the complexity of computing optimal policies on the iMDP, it does not alleviate the complexity stated in Theorem 3 for generating abstractions.




Figure 8: The improved policy synthesis scheme (for clarity, we have omitted actions), in which we merge states based on their lower bound reach-avoid probabilities, computed by Eq. (4) (shown for $\rho = 10$ here). These probabilities are shown within the nodes of each state, while merged states are shown as dotted regions with their lower bound reach-avoid probability written above. If a state has multiple outgoing transitions to the same merged state, e.g., s_4 at time $k = K - 1$, we replace these transitions with a single transition whose probability interval is computed as the sum over the original lower/upper bounds.

of satisfying the reach-avoid property) of reaching state s_1 is one, while the reward of the other states is zero. Based on this intuition, we can merge states $\{s_2, s_3, s_4\}$ as a single successor state in the iMDP (at the final time step, it does not matter if we end up in a critical state or in any other non-goal state: the property is not satisfied in either case). More generally, we *partition* the iMDP states into $\rho \in \mathbb{N}$ discrete *bins*, based on their lower bound reach-avoid probability. For example, if we choose $\rho = 100$ (bins of size 1%), then we have at most 100 successor states, which is often significantly less than the term $|S|$ in the original scheme. The reach-avoid probability of a merged state is the minimum of the lower bound reach-avoid probabilities of the original states it is comprised of. The probability interval $[\underline{p}, \bar{p}]$ of transitioning under state-action pair (s, a) to merged state comprised of a subset of states $M \subset S$ is the union of the intervals over all states in M :

$$\underline{p} = \sum_{s' \in M} \mathcal{P}(s, a, s')_{\text{low}}, \quad \bar{p} = \sum_{s' \in M} \mathcal{P}(s, a, s')_{\text{up}}, \quad (26)$$

where $\mathcal{P}(s, a, s')_{\text{low}}$ and $\mathcal{P}(s, a, s')_{\text{up}}$ denote the lower and upper bounds of these intervals.

Iterating backward over multiple time steps. Given the state aggregation described above for a certain time step, e.g., $k = K$, we can compute the maximum lower bound

reach-avoid probabilities $\underline{Pr}^{\pi^*}(\varphi_{s_t}^K)$ and the corresponding optimal policy π^* at time $K - 1$. For each state $s \in S$, we store the optimal action of the policy as $\pi^*(s, K - 1)$. We then go one step backward in time (according to Bellman iterations) and create another state aggregation over the time steps $k = K - 2$ to $K - 1$. We again partition the range of reach-avoid probabilities (this time those at time $k = K - 1$) into ρ discrete bins and merge the successor states accordingly. As shown in Fig. 8, we thus *iterate backward in time* to compute the optimal policy and lower bound reachability probabilities, until we have reached the initial time step of $k = 0$, and thus have computed the whole policy π^* .

Soundness of the improved scheme. Due to the optimal policy being a Markovian mapping from the state and time step, we can decompose the policy synthesis into individual time steps by iterating backward in time. Theorem 2 is preserved under the improved policy synthesis scheme due to two reasons: 1) the partitioning of states based on their lower bound reach-avoid probabilities leads to an *under-approximation* of the actual reach-avoid probabilities, and 2) summing up the probability intervals over merged states *preserves the original PAC guarantees* on the original abstract model. However, under the proposed scheme, the number of unique intervals increases linearly with the time horizon K . In particular, the iMDPs for all K steps combined have at most $|Act|\rho K$ unique probability intervals (rather than $|Act| \cdot |S|$, cf. Remark 5). Hence, under the proposed scheme, we change the confidence parameter in Theorem 2 to $\alpha = \beta\rho K \cdot |Act|$. Under this modification of the confidence parameter, the proposed scheme is sound.

Using ρ as a tuning parameter. Under the proposed scheme, ρ is a tuning parameter that provides a trade-off between the size of the state space obtained from the iMDPs, versus the level of conservatism in the obtained reach-avoid guarantees introduced by aggregating states. A typical choice for the tuning parameter ρ is $\rho = 100$ (i.e., partitioning states based on their lower bound reach-avoid probabilities with a precision of 1% in reach-avoid probability). If $|S| > \rho K$, then the worst-case number of transitions under the proposed improved policy synthesis scheme is lower than the worst-case number of transitions under the original scheme. However, even if this condition is not satisfied, the improved scheme may be beneficial, because the memory requirements for solving the original iMDP can be excessively large, as demonstrated in the satellite rendezvous benchmark in Sect. 6.1.

6. Numerical Examples

We implement our iterative abstraction method in Python, and tailor the model checker PRISM (Kwiatkowska et al., 2011) for iMDPs to compute robust optimal policies. At every iteration, the obtained iMDP is fed to PRISM, which computes the optimal policy associated with the maximum reach-avoid probability, as per Eq. (4). Our code is available via <https://github.com/LAVA-LAB/DynAbs>, and all experiments are run on a computer with 32 3.7GHz cores and 64 GB of RAM. We report the performance of our method on: (1) a UAV motion control, (2) a building temperature regulation, and (3) a new satellite rendezvous benchmark, which was not in the earlier paper by Badings et al. (2022a). In addition, we benchmark our method against SReachTools and StocHy, which are two other tools for controller synthesis based on formal abstractions. Finally, we demonstrate the applicability of our techniques to infinite-horizon properties.




Figure 9: State trajectory for the satellite benchmark ($N = 3\,200$ with improved policy synthesis scheme). The chaser satellite (white) must navigate to the target (green) while not colliding with the one in red.




Figure 10: UAV reach-avoid problem (goal in green; obstacles in red), plus trajectories under the optimal iMDP-based controller from initial state $\mathbf{x}_0 = [-14, 0, 6, 0, -6, 0]^T$, under high and low turbulence.

6.1 Satellite Rendezvous Problem

We consider the satellite benchmark problem from Jewison and Erwin (2016). More specifically, we consider phase 1 of their satellite rendezvous problem, in which a *chaser* satellite needs to dock with a *target* satellite while being in orbit. The relative motion of the chaser satellite with respect to the target is modeled by the so-called Clohessy-Wiltshire-Hill (CWH) equations (Clohessy & Wiltshire, 1960). We present the full 6-dimensional linear dynamical system in Appendix C.1. Notably, we partition the 6-dimensional state space into $11 \times 23 \times 5 \times 5 \times 5 \times 5 = 158\,125$ discrete regions and define the same number of actions. We define a time horizon of $K = 16$ steps, which becomes 8 steps after grouping every two discrete time steps as described in Sect. 3; see Appendix C.1 for details. The original problem from Jewison and Erwin (2016) is a reachability problem, which we extend to a reach-avoid problem by adding a third satellite that must be avoided (as shown in Fig. 9). We assume that this third satellite is located between the chaser and target satellite and has a fixed position in the CWH frame, yielding a stationary critical region.¹¹

Correct-by-construction control with PAC guarantees. First, let us show how to use Theorem 2 in practice to determine the confidence parameter β needed on individual transition probabilities to solve the overall problem statement with a desired confidence probability. We choose an overall confidence probability of $1 - \alpha = 0.95$. Since we use a uniform rectangular partition of the state space, we can use Corollary 3 to compute the corresponding confidence parameter β on individual transitions. For this particular experiment, we find that a confidence parameter of $\beta = \frac{\alpha}{(22-1)(46-1)(10-1)^4 + 158,125} = 7.86 \times 10^{-9}$ is sufficient to obtain an overall confidence of $1 - \alpha = 0.95$.

11. Note that, in principle, we can also model moving obstacles (or goal regions) by modeling time explicitly in the iMDP abstraction and changing the set of critical (goal) regions at each time step.

Table 1: Model sizes and run times for the spacecraft benchmark, with $N = 3\,200$ and 20 000 noise samples, and either the default abstraction scheme (which generates a single iMDP) or the improved policy synthesis scheme proposed in Sect. 5.3 (which generates an iMDP for every time step $k = 0, \dots, K$ of the property horizon, of which we show every second step in the table). The default scheme with $N = 20\,000$ leads to a memory overflow.

N	Policy synthesis scheme	iMDP size		Run times [s]		
		Transitions	Comp. intervals	Write iMDP	Compute π^*	
3.2k	Default (Single iMDP)	338 847 144	301.23	1103.87	1040.42	
3.2k	Improved synthesis scheme	$k = 8$	813 724	249.46	82.79	136.41
		$k = 6$	26 892 509	29.39	160.64	163.73
		$k = 4$	43 868 194	31.80	212.21	180.77
		$k = 2$	51 816 689	34.85	251.02	195.52
20k	Default (Single iMDP)	560 426 004	–	–	–	
20k	Improved synthesis scheme	$k = 8$	1 156 654	1654.64	86.61	142.70
		$k = 6$	31 914 917	142.52	188.10	177.27
		$k = 4$	52 517 316	144.88	253.31	193.56
		$k = 2$	62 934 064	148.10	286.33	203.09

Improved policy synthesis scheme solves larger problems. We apply our method with $N = 3\,200$ and 20 000 samples, either with or without the improved policy synthesis scheme proposed in Sect. 5.3. One simulated state trajectory of the dynamical system ($N = 3\,200$ and with the improved synthesis scheme enabled) is shown in Fig. 9. The figure shows that the chaser satellite (in white) is indeed able to navigate to the target (in green) without colliding with the third satellite shown in red. For this particular case, the reach-avoid guarantee on the iMDP is $\underline{Pr}^{\pi^*}(\varphi_{s_I}^K) = 0.56$, while the empirical satisfaction of the reach-avoid property (obtained via Monte Carlo simulations) is 0.74.¹²

The number of transitions (s, a, s') of the iMDPs and the run times for all cases are presented in Table 1. The time to compute the states and (enabled) actions is around 30 min and is equal for all cases. Under the default synthesis scheme, we generate a single iMDP over the whole horizon of the reach-avoid property, resulting in very large iMDPs of about 338 and 560 million transitions for $N = 3\,200$ and 20 000, respectively. In the latter instance, the default policy synthesis scheme failed due to a memory overflow (note that we used 64 GB of RAM). By contrast, the improved policy synthesis scheme, which generates a significantly smaller iMDP for each time step $k = 0, \dots, K$ of the horizon, is able to solve both cases, even though the total run time (over all iterations $k = 0, \dots, K$) for $N = 3\,200$ is around 20% higher than with the default scheme. As shown in Table 1, the size of the iMDPs under the improved synthesis scheme increase per iteration, since there is more variety among the reach-avoid probabilities, and thus we can aggregate fewer states. Nevertheless, the results show that with the improved policy synthesis scheme, we can solve reach-avoid problems leading to iMDPs that would otherwise be infeasibly large.

12. Note that we have deliberately chosen a high process noise strength in this benchmark, leading to relatively lower reach-avoid probabilities.




Figure 11: Reach-avoid guarantees on the iMDPs (blue) and MDPs (orange) for their respective policies, versus the resulting empirical (simulated) performance (dashed lines) on the dynamical system. Shaded areas show the standard deviation across 10 iterations. The empirical performance obtained from the MDPs violates the guarantees, whereas that from the iMDPs does not.

6.2 UAV Motion Planning

We consider the reach-avoid problem for a UAV operating under turbulence, which was introduced in Sect. 1. The goal is to compute a controller that guarantees (with high confidence) that the probability to reach a goal area while also avoiding unsafe regions, is above a performance threshold of $\eta = 0.75$. We consider a horizon of 64 time steps, and the problem layout is displayed in Fig. 10, with the goal and unsafe regions shown in green and red, respectively. We model the UAV as a system of 3 double integrators (see Appendix D for details). The state $\mathbf{x}_k \in \mathbb{R}^6$ encodes the position and velocity components, and control inputs $\mathbf{u}_k \in \mathbb{R}^3$ model actuators that change the velocity. The effect of turbulence on the state causes (non-Gaussian) process noise, which we model using a Dryden gust model (Bøhn, Coates, Moe, & Johansen, 2019; Dryden, 1943). We compare two cases: 1) a low turbulence case, and 2) a high turbulence case. We partition the state space into 25 515 regions, using Theorem 1 with $\beta = 0.01$, and apply the iterative scheme with $\gamma = 2$, starting at $N = 25$, with an upper bound of 12 800 samples.

Scalability. We report the model sizes and run times in Appendix D.2. The number of iMDP states equals the size of the partition. Depending on the number of samples N , the iMDP has 9 – 24 million transitions. The mean time to compute the set of iMDP actions (which is only done in the first iteration) is around one minute.¹³ Computing the probabilities plus the verification in PRISM takes 1 – 8 min, depending on the number of samples N .

13. We exploit the block-diagonal structure of matrices A and B of the dynamical system to speed up the computation of the enabled state-action pairs. In particular, we can do all necessary reachability computations separately in the spatial (x, y, z) dimensions and compose the full iMDP afterward.




Figure 12: Cross section (for radiator temp. of 38 °C) of the maximum lower bound probabilities to reach the goal of 20 °C from any initial state, for either 50, 200 or 800 samples.

Accounting for noise matters for probabilistic safety. In Fig. 10, we show state trajectories under the optimal controller derived from the optimal iMDP policy, under high and low turbulence (noise). Under low noise, the controller prefers the short but narrow path; however, with the high noise level, the longer but safer path is preferred since the risk of colliding with an obstacle is too high. Thus, accounting for process noise is important to obtain controllers that are safe.

iMDPs yield safer guarantees than MDPs. To show the importance of using robust abstractions, we compare, under high turbulence, our robust iMDP approach against a naïve MDP abstraction. This MDP has the same states and actions as the iMDP, but uses precise probabilities, which are computed using the frequentist approach introduced in Sect. 4. The maximum reach-avoid probabilities (guarantees) for both methods are shown in Fig. 11. For every value of N , we apply the resulting controllers to the dynamical system in Monte Carlo simulations with 10 000 iterations, to determine the empirical reach-avoid probability as the fraction of trajectories that satisfies the reach-avoid property. Fig. 11 shows that the non-robust MDPs yield *poor and unsafe performance guarantees*: the actual reach-avoid probability of the controller on the dynamical system is much lower than the reach-avoid guarantees obtained from PRISM. By contrast, our robust iMDP-based approach consistently yields safe lower bound guarantees on the actual performance of controllers. The performance threshold of $\underline{Pr}^{\pi^*}(\varphi_{st}^K) \geq 0.75$ is guaranteed for $N = 3\,200$ and higher.

6.3 Building Temperature Regulation

Inspired by Cauchi and Abate (2018), we consider a temperature control problem for a building with two rooms, both having their own radiator and air supply. The reach-avoid goal is to maximize the probability to reach a temperature within 19.8–20.2°C in both zones within 32 steps of 15 minutes each, while avoiding temperatures below 17.8 or above 22.2 °C. The state $\mathbf{x}_k \in \mathbb{R}^4$ of the system (see Appendix E for details) reflects the temperatures of both zones and radiators, and control inputs $\mathbf{u}_k \in \mathbb{R}^4$ change the air supply and boiler temperatures in both zones. The deterministic heat gain through zone walls is modeled by the disturbance $\mathbf{q}_k \in \mathbb{R}^4$. The noise $\mathbf{w}_k \in \mathbb{R}^4$ has a Gaussian distribution (but this assumption is not required for our approach). We partition the state space into 35 721 regions (21 values for zone temperatures and 9 for radiator temperatures), and we use the same values for β and N as in the UAV benchmark in Sect. 6.2.




Figure 13: Maximum lower bound probabilities to reach the goal zone temperature of 21 °C from any initial state within 64 steps, for our approach ($N = 12\,800$) and StocHy.

More samples means less uncertainty. In Fig. 12, we show (for fixed radiator temperatures) the maximum lower bound probabilities obtained from PRISM, to reach the goal from any initial state within the safe set. The results clearly show that better reach-avoid guarantees are obtained when more samples are used to compute the iMDP probability intervals. The higher the value of N , the lower the uncertainty in the intervals, leading to better reach-avoid guarantees. Notably, as reported in Appendix E.2, the largest iMDP has around 200 million transitions, showing that our approach can effectively generate and verify large abstract models.

6.4 Benchmarks Against Other Control Synthesis Tools

StocHy. We benchmark our method on a building temperature regulation problem against StocHy (Cauchi & Abate, 2019), a verification and synthesis tool based on formal abstractions (see Appendix F.1 for details on the setup and results). Similar to our approach, StocHy also derives robust iMDP abstractions. However, StocHy requires precise knowledge of the noise distribution, and it discretizes the control input space of the dynamical system, to obtain a finite action space. The maximum probabilities to reach the goal zone temperature from any initial state obtained for both methods are presented in Fig. 13. The obtained results are qualitatively similar and, close to the goal zone temperature, our lower bound reach-avoid guarantees are *slightly higher* than those obtained from StocHy. However, when starting at temperatures close to the boundary (e.g., at both low radiator and zone temperature), the guarantees obtained from our approach are *slightly more conservative*. This is due to the fact that our approach relies on PAC guarantees on the transition probabilities, while StocHy gives straight probabilistic outcomes, thanks to the assumed precise knowledge of the noise distribution. While both methods yield results that are qualitatively similar, our approach is an order of magnitude faster (45 min for StocHy, vs. 3 – 9 s for our approach; see Appendix F.1 and Table 2 for details).




Figure 14: Simulated state trajectories for the spacecraft docking problem, under low and high noise covariance. Our feedback controllers are more robust, as shown by the smaller error in the state trajectories over time (the Voronoi method under high covariance failed to generate a solution).




Figure 15: Obtained lower bound reach-avoid guarantees for the 1-room temperature control problem with time horizons between $K = 4$ steps (left) and infinity (right).

SReachTools. We apply our method to the spacecraft docking benchmark¹⁴ (see Fig. 14) of SReachTools (Vinod, Gleason, & Oishi, 2019), an optimization-based toolbox for probabilistic reach-avoid problems (see Appendix F.2 for details). While we use samples to generate a model abstraction, SReachTools employs sample-based methods over the properties directly. Distinctively, SReachTools does not create abstractions (as in our case) and is thus generally faster than our method. However, its complexity is exponential in the number of samples (versus the linear complexity for our method). Importantly, we derive *feedback* controllers, while the sampling-based methods of SReachTools compute *open-loop* controllers. Feedback controllers respond to state observations over time and are, therefore, more robust against strong disturbances from noise, as also shown in Fig. 14.

14. Note that the dynamical system used in this spacecraft docking benchmark differs significantly from the system used in the satellite rendezvous problem in Sect. 6.1; see Appendices C and F.2 for details.

6.5 Infinite-Horizon Properties

To demonstrate that our techniques are also applicable to infinite-horizon properties, we consider again the 1-room version of the temperature control problem (on which we benchmarked against StocHy). We consider reach-avoid properties with increasing time bounds of $K \in \{4, 8, 16, \infty\}$ and show the corresponding maximum lower bounds on the reach-avoid probability in Fig. 15. As expected, these figures show that the reach-avoid probability increases with the time horizon (even though the difference between $K = 16$ and $K = \infty$ is marginal). The time to verify the iMDP (which has 383 states and 64 029 transitions) in PRISM is about 1.7s for the property with a horizon of $K = 4$ steps, versus 3.4s for the infinite-horizon property.

7. Related Work

We summarize the literature most related to this paper. We start with a discussion of robust MDPs, (probabilistic) bisimulation, and controller synthesis with formal guarantees. Then, we introduce the scenario approach, as well as other sampling techniques and distributionally robust optimization. Finally, we discuss the concept of PAC and briefly touch upon safe and robust learning methods.

Robust MDPs. In robust (also called uncertain) MDPs, the transition function (in particular, each distribution $P(s, a)$ over successor states, cf. Def. 2) is only known to be an element of an uncertainty set (Wiesemann et al., 2014; Xu & Mannor, 2010). This set is often assumed to be convex for computational reasons, but generalizations to nonconvex sets exist as well (Nilim & Ghaoui, 2005). An iMDP is a special case of robust MDP, where the uncertainty set is described by interval constraints on each transition probability. Specifically, iMDPs have uncertainty sets as *probability simplexes with interval constraints*, which are convex polytopes by construction (Ben-Tal, Ghaoui, & Nemirovski, 2009).

Probabilistic bisimulation. Bisimulation is a key concept used to establish equivalence in the behavior between different systems with respect to a particular logic property (Baier & Katoen, 2008). Probabilistic (or stochastic) notions of bisimulation have been studied for MDPs (Ferns, Precup, & Knight, 2014; Givan, Dean, & Greig, 2003) and other Markov models (Larsen & Skou, 1991; Desharnais et al., 2002). Since exact probabilistic bisimulation requires transition probabilities to agree *exactly*, various metrics and notions of approximate bisimulation have been developed as well (Ferns, Panangaden, & Precup, 2004; Abate, 2011). More recently, probabilistic bisimulation has been studied for interval MDPs (Hashemi, Hermanns, Song, Subramani, Turrini, & Wojciechowski, 2016; Hahn, Hashemi, Hermanns, & Turrini, 2016), and notions of approximate probabilistic bisimulation have been leveraged for robust model checking of probabilistic computation tree logic (PCTL) for interval Markov models (D’Innocenzo, Abate, & Katoen, 2012; Hashemi, Hatefi, & Krcál, 2014; Lun, Wheatley, D’Innocenzo, & Abate, 2018).

In Corollary 2, we establish a probabilistic feedback refinement relation from the abstract MDP to the dynamical system. This relation is exact and yields a *simulation* rather than a *bisimulation* relation: due to the abstraction, not every move by the dynamical system can be matched by the abstract MDP. Estimating the abstract MDP by an iMDP is, on the other hand, approximate since (as we have shown in Theorem 2) there is a probability of at

most β that any of the probability intervals is incorrect. As such, our approach creates a probabilistic closeness guarantee on the satisfaction of a reach-avoid property between the dynamical system and the iMDP, similar to the relations presented in (Lavaei et al., 2022)

Formal controller synthesis. Formal verification and controller synthesis for reachability and reach-avoid problems in stochastic continuous-state systems is an active field of research in safety-critical engineering (Abate et al., 2008; Lavaei et al., 2022). Most approaches are either based on formal abstractions of these continuous systems (Alur, Henzinger, Lafferriere, & Pappas, 2000; Lahijanian, Andersson, & Belta, 2015; Soudjani & Abate, 2013) or work in the continuous domain directly, e.g., using Hamilton-Jacobi reachability analysis (Bansal, Chen, Herbert, & Tomlin, 2017; Herbert, Chen, Han, Bansal, Fisac, & Tomlin, 2017) or optimization (Rosolia, Singletary, & Ames, 2022). Several controller synthesis tools have been developed in this research area, such as StocHy (Cauchi & Abate, 2019), ProbReach (Shmarov & Zuliani, 2015) and SReachTools (Vinod et al., 2019). Despite intense activity, however, the majority of these papers and tools rely on *full knowledge* of the probabilistic models.

The scenario approach. Building upon this motivation, we break away from this literature by developing a method that can be used to generate formal abstractions *without requiring any knowledge of the noise distribution*. The main theoretical tool that we leverage in this paper is the scenario approach (Calafiore & Campi, 2005; Campi & Garatti, 2018). The scenario approach has been used for the verification of MDPs (Badings, Cubuktepe, Jansen, Junges, Katoen, & Topcu, 2022b) and continuous-time Markov chains (Badings, Jansen, Junges, Stoelinga, & Volk, 2022c) with uncertain parameters, albeit only for finite-state systems. However, the non-trivial connection that we make in this paper between the scenario approach theory and techniques for formal abstractions had not been established.

Other sampling techniques. The aforementioned tool SReachTools also exhibits a sampling-based method but relies on Hoeffding’s inequality to obtain confidence guarantees (Sartipizadeh, Vinod, Açikmese, & Oishi, 2019), so the noise is assumed to be sub-Gaussian (Boucheron et al., 2013). By contrast, the scenario approach is *completely distribution-free* (Campi & Garatti, 2018). Moreover, as we have shown in Figs. 5 and 6, the scenario approach may lead to abstract models with significantly better probability intervals compared to Hoeffding’s inequality. In addition, SReachTools is limited to problems with convex safe sets (a restrictive assumption in many problems) and its sampling-based methods can only synthesize open-loop controllers, which may undermine the robustness of the overall approach.

Another body of relevant literature entails sampling-based feedback motion planning algorithms, e.g., LQR-Trees (Tadrake, Manchester, Tobenkin, & Roberts, 2010). However, sampling in LQR-Trees relates to random exploration of the state space and not to stochastic noise affecting the dynamics as in our setting (Reist, Preiswerk, & Tadrake, 2016).

Monte Carlo methods (e.g., particle methods) have also been used to solve stochastic reach-avoid problems (Blackmore et al., 2010; Lesser, Oishi, & Erwin, 2013). These methods simulate the system via many samples of the uncertain variable (Smith, 2013). Monte Carlo methods *approximate* stochastic problems but do not provide rigorous *bounds* with a desired *confidence level* on the obtained results as our approach does.

Distributionally robust optimization. In distributionally robust optimization (DRO), decisions are robust with respect to *ambiguity sets* of distributions (Esfahani & Kuhn, 2018; Goh & Sim, 2010; Wiesemann et al., 2014). While the scenario approach uses samples of the uncertain variable, DRO works on the domain of uncertainty directly, thus involving potentially complex ambiguity sets (Garatti & Campi, 2022). Designing robust policies for iMDPs with known uncertainty sets was studied by Puggelli et al. (2013), and Wolff et al. (2012). Finally, hybrid methods between the scenario approach and robust optimization also exist (Margellos, Goulart, & Lygeros, 2014).

PAC literature. The term PAC refers to obtaining, with high probability, a hypothesis that is a good approximation of some unknown phenomenon (Haussler, 1990). PAC learning methods for discrete-state MDPs are developed in Brafman and Tennenholtz (2002), Fu and Topcu (2014), and Kearns and Singh (2002), and PAC statistical model checking for MDPs in Ashok, Kretínský, and Weininger (2019).

Safe and robust learning methods. We briefly discuss the emerging field of safe learning (Brunke, Greeff, Hall, Yuan, Zhou, Panerati, & Schoellig, 2022; García & Fernández, 2015). Recent works use Gaussian processes for learning-based model predictive control (Hewing, Kabzan, & Zeilinger, 2020; Koller, Berkenkamp, Turchetta, & Krause, 2018) or reinforcement learning with safe exploration (Berkenkamp, Turchetta, Schoellig, & Krause, 2017), and control barrier functions to reduce model uncertainty (Taylor, Singletary, Yue, & Ames, 2020). Safe learning control concerns learning unknown, *deterministic* system dynamics, while imposing strong assumptions on stochasticity (Fisac, Akametalu, Zeilinger, Kaynama, Gillula, & Tomlin, 2019). By contrast, our problem setting is fundamentally different: we reason about *stochastic* noise of a completely unknown distribution.

Finally, various methods have been proposed to render reinforcement learning more robust against differences in the dynamics between the training and testing environments (Morimoto & Doya, 2005; Moos, Hansel, Abdulsamad, Stark, Clever, & Peters, 2022). For example, Peng, Andrychowicz, Zaremba, and Abbeel (2018) train neural network controllers for robotic control tasks on simulation models with randomized dynamics. The authors show empirically that this randomization leads to controllers that are significantly more robust against discrepancies in the dynamics of the real robot on which they are deployed. Similarly, Pinto, Davidson, Sukthankar, and Gupta (2017) and Vinitzky, Du, Parvate, Jang, Abbeel, and Bayen (2020) propose adversarial methods to improve the robustness of reinforcement learning by introducing an adversary which applies maximally destabilizing disturbances to the system. While these methods have been shown empirically to improve the robustness of reinforcement learning, we remark that providing formal guarantees about safety or robustness (as we do in this paper) is rarely possible.

8. Concluding Remarks and Future Work

We have presented a novel approach for robust control of continuous-state dynamical systems with (stochastic) process noise of unknown distribution. Our approach is based on a finite abstraction of the continuous-state system in the form of an iMDP. As a key contribution, we have developed a rigorous method to use the scenario approach theory for computing PAC probability intervals for such an abstract iMDP. The PAC-correctness of

these intervals, and thus of the whole abstract iMDP, is valid irrespective of the noise distribution. We have shown how to use the iMDP to compute feedback controllers with PAC guarantees on the performance on the continuous system. Our experiments have confirmed this claim, showing that our method effectively solves realistic problems and provides safe lower-bound guarantees on the performance of controllers.

State space discretization. The discretization of the state space influences the quality of the reach-avoid guarantees. Partitions into regions that are smaller are more easily contained in the backward reachable sets of actions, thus enabling more actions in the abstraction. Moreover, defining more target points leads to an abstract model that captures more actions in general. Hence, a more fine-grained partition with more target points yields an abstraction that is a more accurate representation of the dynamical system but also increases the computational complexity. Whilst the time-dependent improved policy synthesis scheme based on state aggregation, as discussed in Sect. 5.3, balances this trade-off to some extent, we plan to employ more general and sophisticated adaptive discretization schemes in the future, such as those proposed in Soudjani and Abate (2013).

Extensions to general PCTL properties. As described in Sect. 2.2, we have considered control objectives in the form of (in)finite-horizon reach-avoid properties. Formally, these reach-avoid properties contain formulae expressed in probabilistic computation tree logic, or PCTL (Ciesinski & Größer, 2004; Hansson & Jonsson, 1994). Our abstraction approach is valid for properties over both finite and infinite horizons, except for the improved policy synthesis scheme proposed in Sect. 5.3, which only applies to finite horizons (as this scheme relies on modeling each time step separately). By contrast, techniques that result in abstraction errors accumulating with time are generally not applicable to infinite-horizon properties (Lavaei et al., 2022).

Moreover, while we have left out rewards for brevity, our approach is also amenable to expected reward properties (Baier & Katoen, 2008) with state-based¹⁵ reward functions (mapping states to rewards) as long as the reward function is aligned with the partition, similar to Assumption 1. More precisely, all states $\mathbf{x}_k \in R_s$ belonging to the region of abstract state $s \in S$ must be assigned the same reward. Under this assumption, our approaches are equally applicable to general PCTL properties, although we leave a formal derivation of these results as an avenue for future work.

Non-stationary noise distributions. For our controller synthesis approach to be correct, the samples used for computing the probability intervals of the iMDP must be i.i.d. and drawn from the same distribution as the noise \mathbf{w}_k affecting the linear dynamical system in Eq. (1). Thus, our approach requires that the probability distribution of the noise is fixed, which may be a questionable assumption in practice (Brunke et al., 2022). As such, one open research question is to what extent the correctness of our approach can be guaranteed if the noise distribution is not stationary but instead perturbed by some (bounded) distance.

Extensions to systems with uncertain dynamics. In this paper, we have considered linear dynamical systems whose dynamics (apart from the distribution of the noise) are

15. We cannot consider action-based reward functions because there is no direct correspondence between control inputs $\mathbf{u}_k \in \mathcal{U}$ for the dynamical system and actions $a \in Act$ of the abstract iMDP.

precisely known. Therefore, for physics-based systems, such as a UAV, system parameters, including its mass or friction coefficient, must be known exactly. Our ongoing research, recently presented in Badings, Romao, Abate, and Jansen (2022), lifts this restrictive assumption by considering systems with both stochastic noise and uncertain dynamics, e.g., due to imprecisely known parameters. In particular, such uncertainty in the dynamics causes *epistemic uncertainty*, which is fundamentally different from the probability distributions over outcomes due to process noise, leading to *aleatoric uncertainty* (Sullivan, 2015).

In a setting with both epistemic and aleatoric uncertainty, we must simultaneously learn about the unknown deterministic dynamics and the stochastic noise. Learning deterministic dynamics is common in safe learning control (Brunke et al., 2022), but enabling safe exploration requires strong assumptions on stochastic uncertainty. As such, we see this direction as a challenging avenue for future work.

Extensions to dynamical systems with both aleatoric and epistemic uncertainty also open the door to capturing other types of uncertainty beyond process noise. Besides the uncertainty in system parameters described above, we may, for example, deal with state/control-dependent process noise or systems with nonlinear dynamics (see below). Moreover, such an extension may also enable us to lift Assumption 2, which is needed because our current method requires the system to be fully actuated.

Nonlinear systems. While we have focused on linear systems, we wish to develop extensions to nonlinear systems, as discussed in Remark 1. Such extensions are non-trivial and may require more involved reachability computations (Bansal et al., 2017; Chen, Ábrahám, & Sankaranarayanan, 2013). Specifically, the challenge is to compute the enabled iMDP actions via the backward reachable set defined in Eq. (6), which may become non-convex under nonlinear dynamics. Note that computing the PAC probability intervals remains unchanged, as the scenario approach relies on the convexity of the target set only, and not on that of the backward reachable set. Alternatively, we may apply our method on a linearized version of the nonlinear system, in which case we must account for any linearization error to preserve guarantees. Similar to the extension with uncertain parameters, this linearization error may potentially be captured by epistemic uncertainty in the system dynamics.

Acknowledgments

This work was partially funded by the NWO grant NWA.1160.18.238 (*Prima Vera*), the 2022 JPMorgan Chase Faculty Research Award “Learning and Reasoning in Repeated Games with Partial Information”, the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101008233, the ERC Consolidator Grant 864075 (*CAESAR*), the EPSRC IAA Award EP/X525777/1, and the ERC Advanced Grant 834115 (*FUN2MODEL*).

Appendix A. Proof of Theorem 1

The proof of Theorem 1 is adapted from Romao et al. (2022, Theorem 5), which is based on three key assumptions: (1) the considered scenario problem belongs to the class of so-called *fully-supported problems* (see Campi and Garatti (2008) for a definition), (2) its solution is unique, and (3) discarded samples violate all interim optimal solutions with probability one.

In our case, requirement (2) is implied by Assumption 4, (3) is satisfied by proposition 1, and the problem is fully-supported because the number of decision variables is one. Under these requirements, Romao et al. (2022, Theorem 5) states that the risk associated with an optimal solution $\lambda_{|Q|}^*$ to Eq. (16) for $|Q|$ discarded samples satisfies the following expression:

$$\mathbb{P}^N \left\{ P(\mathbf{x} \notin R_j(\lambda_{|Q|}^*)) \leq \epsilon \right\} = 1 - \sum_{i=0}^{|Q|} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i}, \quad (27)$$

where we omit the subscripts in $P_{\mathbf{w}_k}(\cdot)$ and \mathbf{x}_{k+1} for brevity. Eq. (27) is the cumulative distribution function (CDF) of a beta distribution with parameters $|Q| + 1$ and $N - |Q|$. We denote this CDF by $F_{|Q|}(\epsilon)$, where we explicitly write the dependency on $|Q|$. Hence, we obtain

$$F_{|Q|}(\epsilon) = \mathbb{P}^N \left\{ P(\mathbf{x} \notin R_j(\lambda_{|Q|}^*)) \leq \epsilon \right\} = \tilde{\beta}. \quad (28)$$

Thus, if we discard $|Q|$ samples, Eq. (28) returns the confidence probability $\tilde{\beta} \in (0, 1)$ by which the probability of $\mathbf{x}_{k+1} \notin R_j(\lambda_{|Q|}^*)$ is upper bounded by ϵ . Conversely, we can compute the value of ϵ needed to obtain a confidence probability of $\tilde{\beta}$, using the percent point function (PPF) $G_{|Q|}(\tilde{\beta})$:

$$\mathbb{P}^N \left\{ P(\mathbf{x} \notin R_j(\lambda_{|Q|}^*)) \leq G_{|Q|}(\tilde{\beta}) \right\} = \tilde{\beta}. \quad (29)$$

The PPF is the inverse of the CDF, so by definition, we have

$$\epsilon = G_{|Q|}(\tilde{\beta}) = G_{|Q|}(F_{|Q|}(\epsilon)). \quad (30)$$

Note that $P(\mathbf{x} \in R) + P(\mathbf{x} \notin R) = 1$, so Eq. (28) equals

$$F_{|Q|}(\epsilon) = \mathbb{P}^N \left\{ P(\mathbf{x} \in R_j(\lambda_{|Q|}^*)) \geq 1 - \epsilon \right\} = \tilde{\beta}. \quad (31)$$

By defining $p = 1 - \epsilon$, Eqs. (30) and (31) are combined as

$$\mathbb{P}^N \left\{ 1 - G_{|Q|}(\tilde{\beta}) \leq P(\mathbf{x} \in R_j(\lambda_{|Q|}^*)) \right\} = 1 - \sum_{i=0}^{|Q|} \binom{N}{i} (1-p)^i p^{N-i} = \tilde{\beta}. \quad (32)$$

In what follows, we separately use Eq. (31) to prove the lower and upper bound of the probability interval in Theorem 1.

Lower bound. There are N possible values for $|Q|$, ranging from 0 to $N - 1$. The case $|Q| = N$ (i.e. all samples are discarded) is treated as a special case in Theorem 1. We fix $\tilde{\beta} = 1 - \frac{\beta}{2N}$ in Eq. (32), yielding the series of equations

$$\begin{aligned} \mathbb{P}^N \left\{ 1 - G_0 \left(1 - \frac{\beta}{2N} \right) \leq P(\mathbf{x} \in R_j(\lambda_0^*)) \right\} &= 1 - \frac{\beta}{2N} \\ \vdots & \vdots \\ \mathbb{P}^N \left\{ 1 - G_{N-1} \left(1 - \frac{\beta}{2N} \right) \leq P(\mathbf{x} \in R_j(\lambda_{N-1}^*)) \right\} &= 1 - \frac{\beta}{2N}. \end{aligned} \quad (33)$$

Denote the event that $1 - G_n(1 - \frac{\beta}{2N}) \leq P(\mathbf{x} \in R_j(\lambda_n^*))$ for $n = 0, \dots, N - 1$ by \mathcal{A}_n . Regardless of n , this event has a probability of $\mathbb{P}^N\{\mathcal{A}_n\} = 1 - \frac{\beta}{2N}$, and its complement \mathcal{A}'_n of $\mathbb{P}^N\{\mathcal{A}'_n\} = \frac{\beta}{2N}$. Via Boole's inequality, we know that

$$\mathbb{P}^N\left\{\bigcup_{i=0}^{N-1} \mathcal{A}'_i\right\} \leq \sum_{i=0}^{N-1} \mathbb{P}^N\{\mathcal{A}'_i\} = \frac{\beta}{2N}N = \frac{\beta}{2}. \quad (34)$$

Thus, for the intersection of all events in Eq. (33) we have

$$\mathbb{P}^N\left\{\bigcap_{i=0}^{N-1} \mathcal{A}_i\right\} = 1 - \mathbb{P}^N\left\{\bigcup_{i=0}^{N-1} \mathcal{A}'_i\right\} \geq 1 - \frac{\beta}{2}. \quad (35)$$

After observing the samples at hand, we replace $|Q|$ by the actual value of N_j^{out} (as per Def. 6), giving one of the expressions in Eq. (33). The probability that this expression holds cannot be smaller than that of the intersection of all events in Eq. (35). Thus, we obtain

$$\mathbb{P}^N\left\{\underline{p} \leq P(\mathbf{x} \in R_j(\lambda_{N_j^{\text{out}}}^*))\right\} \geq 1 - \frac{\beta}{2}, \quad (36)$$

where $\underline{p} = 0$ if $N_j^{\text{out}} = N$ (which trivially holds with probability one), and otherwise $\underline{p} = 1 - G_{N_j^{\text{out}}}(1 - \frac{\beta}{2N})$ is the solution for p to Eq. (32), with $|Q| = N_j^{\text{out}}$ and $\tilde{\beta} = 1 - \frac{\beta}{2N}$:

$$1 - \frac{\beta}{2N} = 1 - \sum_{i=0}^{N_j^{\text{out}}} \binom{N}{i} (1-p)^i p^{N-i}, \quad (37)$$

which is equivalent to Eq. (19).

Upper bound. Eq. (32) is rewritten as an upper bound as

$$\mathbb{P}^N\left\{P(\mathbf{x} \in R_j(\lambda_{|Q|}^*)) < 1 - G_{|Q|}(\tilde{\beta})\right\} = 1 - \tilde{\beta}, \quad (38)$$

where again, $|Q|$ can range from 0 to $N - 1$. However, to obtain high-confidence guarantees on the upper bound, we now fix $\tilde{\beta} = \frac{\beta}{2N}$, which yields the series of equations

$$\begin{aligned} \mathbb{P}^N\left\{P(\mathbf{x} \in R_j(\lambda_0^*)) < 1 - G_0\left(\frac{\beta}{2N}\right)\right\} &= 1 - \frac{\beta}{2N} \\ &\vdots \\ \mathbb{P}^N\left\{P(\mathbf{x} \in R_j(\lambda_{N-1}^*)) < 1 - G_{N-1}\left(\frac{\beta}{2N}\right)\right\} &= 1 - \frac{\beta}{2N}. \end{aligned} \quad (39)$$

Analogous to the lower bound case, Boole's inequality implies that the intersection of all expressions in Eq. (39) has a probability of at least $1 - \frac{\beta}{2}$. After observing the samples at hand, we replace $|Q|$ by $N_j^{\text{out}} - 1$, yielding one of the expressions in Eq. (39). For this expression, it holds that

$$\mathbb{P}^N\left\{P(\mathbf{x} \in R_j(\lambda_{N_j^{\text{out}}-1}^*)) \leq \bar{p}\right\} \geq 1 - \frac{\beta}{2}, \quad (40)$$

where $\bar{p} = 1$ if $N_j^{\text{out}} = 0$ (which trivially holds with probability one), and otherwise $\bar{p} = 1 - G_{N_j^{\text{out}}-1}(\frac{\beta}{2N})$ is the solution for p to Eq. (32), with $|Q| = N_j^{\text{out}} - 1$ and $\tilde{\beta} = \frac{\beta}{2N}$:

$$\frac{\beta}{2N} = 1 - \sum_{i=0}^{N_j^{\text{out}}-1} \binom{N}{i} (1-p)^i p^{N-i}, \quad (41)$$

which is equivalent to Eq. (20).

Probability interval. We invoke Lemma 1, which states that $R_j(\lambda_{N_j^{\text{out}}}^*) \subseteq R_j \subseteq R_j(\lambda_{N_j^{\text{out}}-1}^*)$, so we have

$$\begin{aligned} P(\mathbf{x} \in R_j(\lambda_{N_j^{\text{out}}}^*)) &\leq P(\mathbf{x} \in R_j) = P(s_i, a_l)(s_j) \\ &< P(\mathbf{x} \in R_j(\lambda_{N_j^{\text{out}}-1}^*)). \end{aligned} \quad (42)$$

We use Eq. (42) to write Eqs. (36) and (40) in terms of the transition probability $P(s_i, a_l)(s_j)$. Finally, by applying Boole's inequality, we combine Eqs. (36) and (40) as follows:

$$\begin{aligned} \mathbb{P}^N \left\{ \underline{p} \leq P(s_i, a_l)(s_j) \cap P(s_i, a_l)(s_j) \leq \bar{p} \right\} &= \mathbb{P}^N \left\{ \underline{p} \leq P(s_i, a_l)(s_j) \leq \bar{p} \right\} \\ &\geq 1 - \beta, \end{aligned} \quad (43)$$

which is equal to Eq. (18), so we conclude the proof.

Appendix B. Proof of Corollary 3

Proof. The derivation of Corollary 3 is analogous to the proof of Theorem 2, with the only difference that the number of unique probability intervals is reduced. The key observation is that the successor state \mathbf{x}_{k+1} in Eq. (9) is linear in the target point \mathbf{d}_a and the noise \mathbf{w}_k . Thus, by denoting $p(\mathbf{w}_k)$ as the (unknown) probability density function of the noise, we rewrite Eq. (10) as

$$\begin{aligned} P(s, a)(s') &= \int_{R_{s'}} p_{\mathbf{w}_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \mathbf{u}_k(\mathbf{x}_k, a)) d\mathbf{x}_{k+1} \\ &= \int_{R_{s'} - \mathbf{d}_a} p(\mathbf{w}_k) d\mathbf{w}_k, \end{aligned} \quad (44)$$

that is, for a fixed density function of \mathbf{w}_k the transition probability $P(s, a)(s')$ is defined by the relative position $R_{s'} - \mathbf{d}_a$ between the region and the target point. Thus, for any two regions $R_{s'}$, $R_{s''}$ and two target points \mathbf{d}_a , $\mathbf{d}_{a'}$ whose relative positions are the same, i.e. $R_{s'} - \mathbf{d}_a = R_{s''} - \mathbf{d}_{a'}$, it holds that

$$P(s, a)(s') = \int_{R_{s'} - \mathbf{d}_a} p(\mathbf{w}_k) d\mathbf{w}_k = \int_{R_{s''} - \mathbf{d}_{a'}} p(\mathbf{w}_k) d\mathbf{w}_k = P(s, a')(s''). \quad (45)$$

Now, observe that under the proposed partitioning and target points, the number of pairs $(R_{s'}, \mathbf{d}_a)$ with equal relative position is maximized. Formally, let $\mu_{s'}$ be the center of $R_{s'}$ for

each $s' \in S \setminus \{s_\star\}$, and let w_i be the width of the rectangular partitioning in each dimension $i = 1, \dots, n$. Then, the difference $\mu_{s'} - \mathbf{d}_a$ for any $s' \in S$, $a \in Act$ can be written as

$$\mu_{s'} - \mathbf{d}_a = \begin{bmatrix} \xi_1 w_1 \\ \vdots \\ \xi_n w_n \end{bmatrix}, \quad \xi_i \in \{-r_i + 1, \dots, r_i - 1\} \forall i = 1, \dots, n, \quad (46)$$

except when s' is the absorbing state. Since each integer variable ξ_i can take one of $2r_i - 1$ different values, we observe that $\mu_{s'} - \mathbf{d}_a$ can take one of $\prod_{i=1}^n (2r_i - 1)$ different values. Thus, the iMDP has exactly $\prod_{i=1}^n (2r_i - 1)$ unique transition probabilities (and hence the same number of unique probability intervals), plus one unique probability $P(s, a)(s_\star)$ of transitioning to the absorbing state s_\star for each $a \in Act$. Plugging in this total number of $\prod_{i=1}^n (2r_i - 1) + |Act|$ unique probability intervals, we obtain the desired expression in Eq. (44) and thus conclude the proof. \square

Appendix C. Details on Satellite Rendezvous Benchmark

C.1 Explicit Model Formulation

The dynamical system of this benchmark problem is defined on the so-called Hill's frame, which represents the relative coordinate frame describing the difference in position and velocity between a chaser and target satellite (Jewison & Erwin, 2016). The dynamics of this 6-dimensional system are defined as follows:

$$\begin{aligned} \mathbf{x}_{k+1} = & \begin{bmatrix} 4 - 3 \cos(n\tau) & 0 & 0 & \frac{1}{n} \sin(n\tau) & \frac{2}{n}(1 - \cos(n\tau)) & 0 \\ 6(\sin(n\tau) - n\tau) & 1 & 0 & \frac{-2}{n}(1 - \cos(n\tau)) & \frac{4}{n} \sin(n\tau) - 3n\tau & 0 \\ 0 & 0 & \cos(n\tau) & 0 & 0 & \frac{1}{n} \sin(n\tau) \\ 3n \sin(n\tau) & 0 & 0 & \cos(n\tau) & 2 \sin(n\tau) & 0 \\ -6n(1 - \cos(n\tau)) & 0 & 0 & -2 \sin(n\tau) & 4 \cos(n\tau) - 3 & 0 \\ 0 & 0 & -n \sin(n\tau) & 0 & 0 & \cos(n\tau) \end{bmatrix} \mathbf{x}_k \\ & + \begin{bmatrix} \frac{1}{n} \sin(n\tau) & \frac{2}{n}(1 - \cos(n\tau)) & 0 \\ \frac{-2}{n}(1 - \cos(n\tau)) & \frac{1}{n}(4 \sin(n\tau) - 3n\tau) & 0 \\ 0 & 0 & \frac{1}{n} \sin(n\tau) \\ \cos(n\tau) & 2 \sin(n\tau) & 0 \\ -2 \sin(n\tau) & 4 \cos(n\tau) - 3 & 0 \\ 0 & 0 & \cos(n\tau) \end{bmatrix} \mathbf{u}_k + \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \text{diag}\left(\begin{bmatrix} 0.1 \\ 0.1 \\ 0.01 \\ 0.01 \\ 0.01 \\ 0.01 \end{bmatrix}\right)\right), \end{aligned}$$

where τ is the discretization time, and n is a constant; see Jewison and Erwin (2016) for details. Every element of the 3-dimensional control input vector \mathbf{u}_k is constrained to the interval $[-2, 2]$. Since the model in Appendix C.1 is not fully actuated (it has only 3 controls, versus a state of dimension 6), we group every two discrete time steps together (as described in Sect. 3) and rewrite the model as follows:

$$\mathbf{x}_{k+2} = \bar{A} \mathbf{x}_k + \bar{B} \mathbf{u}_{k,k+1} + \mathbf{w}_{k,k+1}, \quad (47)$$

where $\bar{A} \in \mathbb{R}^{6 \times 6}$ and $\bar{B} \in \mathbb{R}^{6 \times 6}$ are properly redefined matrices, and $\mathbf{u}_{k,k+1} \in \mathbb{R}^6$ and $\mathbf{w}_{k,k+1} \in \mathbb{R}^6$ reflect the control inputs and process noise at both time steps k and $k + 1$, combined in one vector.

Appendix D. Details on UAV Benchmark

D.1 Explicit Model Formulation

The 6-dimensional state vector of the UAV model is $\mathbf{x} = [p_x, v_x, p_y, v_y, p_z, v_z]^\top \in \mathbb{R}^6$, where p_i and v_i are the position and velocity in the direction i . Every element of the vector of control inputs $\mathbf{u} = [f_x, f_y, f_z]^\top \in \mathbb{R}^3$ is constrained to the interval $[-4, 4]$. The discrete-time dynamics are an extension of the lower-dimensional case in Badings, Jansen, Poonawala, and Stoelinga (2021), and are written in the form of Eq. (1) as follows:

$$\mathbf{x}_{k+1} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0.5 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0.5 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0.5 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{u}_k + \mathbf{w}_k, \quad (48)$$

where \mathbf{w}_k is the effect of turbulence on the continuous state, which we model using a Dryden gust model; we use a Python implementation by Bøhn et al. (2019). Note that the drift term $\mathbf{q}_k = 0$, and is thus omitted from Eq. (48). Similar to the dynamical system of the satellite benchmark in Appendix C.1, we group every two discrete time steps together to render the system in Eq. (48) fully actuated.

The objective is to reach the goal region, which is also depicted in Fig. 10, within 64 steps (i.e. 32 steps with the model in Eq. (47)). This goal region is written as the following set of continuous states (for brevity, we omit an explicit definition of the critical regions):

$$\mathcal{X}_G = \{\mathbf{x} \in \mathbb{R}^6 \mid 11 \leq p_x \leq 15, \quad 1 \leq p_y \leq 5, \quad -7 \leq p_z \leq -3\}. \quad (49)$$

D.2 Run Times and Size of Model Abstraction

The average run times and iMDP model sizes (over 10 runs) for all iterations of the UAV benchmark are presented in Table 2. Computing the states and actions of the underlying MDP took one minute but since this step is only performed in the first iteration, we omit the value from the table. The run times per iteration are the times for those steps within the while-loop in Algorithm 1. The number of states (which equals the number of regions in the partition) and choices (the total number of state-action pairs) of the iMDP are independent of the value of N . By contrast, the number of transitions increases with N , because the additional noise samples may reveal more possible outcomes of a state-action pair.

Appendix E. Details on Building Temperature Regulation Benchmark

E.1 Explicit Model Formulation

This model is a variant of the two-zone building automation system benchmark with stochastic dynamics in Cauchi and Abate (2018). The state vector of the model is $\mathbf{x} = [T_1^z, T_2^z, T_1^r, T_2^r]^\top \in \mathbb{R}^4$, reflecting both room (zone) temperatures (T^z) and both radiator temperatures (T^r). The control inputs $\mathbf{u} = [T_1^{\text{ac}}, T_2^{\text{ac}}, T_1^{\text{boil}}, T_2^{\text{boil}}]^\top \in \mathbb{R}^4$ are the air conditioning (ac) and boiler temperatures, respectively, which are constrained within

Table 2: Run times per iteration (which excludes the computation of states and actions) and model sizes for the UAV (under high turbulence), and 2-zone and 1-zone building temperature regulation (BAS) benchmarks.

Instance		Run times (per iteration, in seconds)			iMDP model size		
Benchmark	Samples N	Update intervals	Write iMDP	Compute π^*	States	Choices	Transitions
UAV	25	11.4	29.0	39.2	25 516	1 228 749	9 477 795
UAV	50	15.4	34.4	46.0	25 516	1 228 749	11 350 692
UAV	100	18.6	39.3	51.8	25 516	1 228 749	13 108 282
UAV	200	24.3	44.4	59.1	25 516	1 228 749	14 921 925
UAV	400	35.4	50.6	68.5	25 516	1 228 749	17 203 829
UAV	800	55.4	58.3	80.2	25 516	1 228 749	19 968 663
UAV	1600	92.2	64.9	90.6	25 516	1 228 749	22 419 161
UAV	3200	164.1	68.6	94.4	25 516	1 228 749	23 691 107
UAV	6400	312.1	69.8	95.5	25 516	1 228 749	23 958 183
UAV	12800	611.2	69.8	95.1	25 516	1 228 749	23 969 028
<hr/>							
BAS (2-zone)	25	13.9	73.6	105.9	35 722	1 611 162	24 929 617
BAS (2-zone)	50	43.2	107.7	160.5	35 722	1 611 162	37 692 421
BAS (2-zone)	100	50.8	153.2	230.2	35 722	1 611 162	53 815 543
BAS (2-zone)	200	59.1	205.8	315.2	35 722	1 611 162	72 588 298
BAS (2-zone)	400	72.3	260.6	423.2	35 722	1 611 162	91 960 970
BAS (2-zone)	800	96.2	316.8	504.3	35 722	1 611 162	112 028 449
BAS (2-zone)	1600	132.6	370.0	614.8	35 722	1 611 162	131 844 146
BAS (2-zone)	3200	200.1	432.9	789.4	35 722	1 611 162	154 155 445
BAS (2-zone)	6400	331.8	513.2	982.2	35 722	1 611 162	182 175 229
BAS (2-zone)	12800	545.5	601.8	1177.8	35 722	1 611 162	218 031 384
<hr/>							
BAS (1-zone)	25	1.62	0.06	1.39	381	1 511	20 494
BAS (1-zone)	50	1.81	0.08	1.39	381	1 511	27 327
BAS (1-zone)	100	1.87	0.09	1.40	381	1 511	33 871
BAS (1-zone)	200	1.96	0.11	1.40	381	1 511	40 368
BAS (1-zone)	400	2.06	0.13	1.42	381	1 511	47 333
BAS (1-zone)	800	2.26	0.14	1.40	381	1 511	54 155
BAS (1-zone)	1600	2.59	0.16	1.39	381	1 511	59 686
BAS (1-zone)	3200	3.24	0.17	1.40	381	1 511	65 122
BAS (1-zone)	6400	4.50	0.18	1.40	381	1 511	70 245
BAS (1-zone)	12800	7.01	0.20	1.38	381	1 511	76,076

$14 \leq T^{\text{ac}} \leq 26$ and $65 \leq T^{\text{boil}} \leq 85$. The changes in the temperature of both zones are governed by the following thermodynamics:

$$\begin{aligned} \dot{T}_1^z &= \frac{1}{C_1} \left[\frac{T_2^z - T_1^z}{R_{1,2}} + \frac{T_{\text{wall}} - T_1^z}{R_{1,\text{wall}}} + mC_{pa}(T_1^{\text{ac}} - T_1^z) + P_{\text{out}}(T_1^r - T_1^z) \right] \\ \dot{T}_2^z &= \frac{1}{C_2} \left[\frac{T_1^z - T_2^z}{R_{1,2}} + \frac{T_{\text{wall}} - T_2^z}{R_{2,\text{wall}}} + mC_{pa}(T_2^{\text{ac}} - T_2^z) + P_{\text{out}}(T_2^r - T_2^z) \right], \end{aligned}$$

where C_i is the thermal capacitance of zone i , and $R_{i,j}$ is the resistance between zones i and j , m is the air mass flow, C_{pa} is the specific heat capacity of air, and P_{out} is the rated output of the radiator. Similarly, the dynamics of the radiator in room i are governed by the following equation:

$$\dot{T}_i^r = k_1(T_i^z - T_i^r) + k_0w(T_i^{\text{boil}} - T_i^r), \quad (51)$$

where k_0 and k_1 are constant parameters, and w is the water mass flow from the boiler. For the precise parameter values used, we refer to our codes, which are provided in the

supplementary material. By discretizing Eqs. (50) and (51) by a forward Euler method at a time resolution of 15 min, we obtain the following model in explicit form:

$$\begin{aligned} \mathbf{x}_{k+1} = & \begin{bmatrix} 0.8425 & 0.0537 & -0.0084 & 0.0000 \\ 0.0515 & 0.8435 & 0.0000 & -0.0064 \\ 0.0668 & 0.0000 & 0.8971 & 0.0000 \\ 0.0000 & 0.0668 & 0.0000 & 0.8971 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0.0584 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0599 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0362 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0362 \end{bmatrix} \mathbf{u}_k \\ & + \begin{bmatrix} 1.2291 \\ 1.0749 \\ 0.0000 \\ 0.0000 \end{bmatrix} + \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0.01 & 0 & 0 & 0 \\ 0 & 0.01 & 0 & 0 \\ 0 & 0 & 0.01 & 0 \\ 0 & 0 & 0 & 0.01 \end{bmatrix}\right), \end{aligned}$$

where $\mathcal{N}(\boldsymbol{\mu}, \Sigma) \in \mathbb{R}^n$ is a multivariate Gaussian random variable with mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance matrix Σ .

The goal in this problem is to reach a temperature of 19.9 – 20.1 °C in both zones within 32 discrete time steps of 15 min. As such, the goal region \mathcal{X}_G and critical region \mathcal{X}_C are:

$$\mathcal{X}_G = \{\mathbf{x} \in \mathbb{R}^2 \mid 19.9 \leq T_1^z \leq 20.1, 19.9 \leq T_2^z \leq 20.1\}, \quad \mathcal{X}_C = \emptyset.$$

E.2 Run Times and Size of Model Abstraction

The run times and iMDP model sizes for all iterations of the 2-zone building temperature regulation benchmark are presented in Table 2. Computing the states and actions took 5.5 min, but we omit this value from the table as this step is only performed in the first iteration. The discussion of model sizes is analogous to Appendix D.2 on the UAV benchmark.

Appendix F. Benchmarks Against Other Tools

F.1 Benchmark Against StocHy

We provide a comparison between our approach and StocHy (Cauchi & Abate, 2019) on a reduced version of the temperature regulation problem in Appendix E.1 with only one zone. Thus, the state vector becomes $\mathbf{x} = [T^z, T^r]^\top$, and the control inputs are $\mathbf{u} = [T^{\text{ac}}, T^{\text{boil}}]^\top$, which are constrained to $14 \leq T^{\text{ac}} \leq 28$ and $-10 \leq T^{\text{boil}} \leq 10$ (note that T^{boil} is now relative to the nominal boiler temperature of 75 °C). Using the same dynamics as in Eqs. (50) and (51), we obtain the following model in explicit form:

$$\mathbf{x}_{k+1} = \begin{bmatrix} 0.8820 & 0.0058 \\ 0.0134 & 0.9625 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0.0584 & 0.0000 \\ 0.0000 & 0.0241 \end{bmatrix} \mathbf{u}_k + \begin{bmatrix} 0.9604 \\ 1.3269 \end{bmatrix} + \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0.02 & 0 \\ 0 & 0.1 \end{bmatrix}\right).$$

The objective in this problem is to reach a zone temperature of 20.9 – 21.1 °C within 64 discrete time steps of 15 min. As such, the goal region \mathcal{X}_G and critical region \mathcal{X}_C are:

$$\mathcal{X}_G = \{\mathbf{x} \in \mathbb{R}^2 \mid 20.9 \leq T^z \leq 21.1\}, \quad \mathcal{X}_C = \emptyset. \quad (52)$$

We partition the continuous state space in a rectangular grid of 19×20 regions of width 0.2 °C, which is centered at $T^z = 21$ °C and $T^r = 38$ °C. As such, the partition covers zone temperatures between 19.1 – 22.9 °C, and radiator temperatures between 36 – 40 °C.

Table 3: Maximum lower bound reach-avoid probability obtained using our method, the average simulated reach-avoid probability, and run times for our method and SReachTools, under the default (x1) and increased (x10) noise covariance.

	Our method					SReachTools	
	Init. abstr.	$N = 25$	$N = 100$	$N = 400$	$N = 1600$	Particle	Voronoi
Noise covariance x1							
Reach-avoid probability	—	0.44	0.81	0.95	0.98	0.88	0.85
Simulated probability	—	1.00	1.00	1.00	1.00	0.83	0.86
Run time (s)	3.420	6.966	9.214	11.016	17.294	0.703	3.462
Noise covariance x10							
Reach-avoid probability	—	0.19	0.36	0.52	0.62	0.40	NaN
Simulated probability	—	0.59	0.63	0.59	0.65	0.19	NaN
Run time (s)	3.372	11.433	15.909	18.671	25.452	2.495	NaN

Abstraction method. Similar to our approach, StocHy generates robust iMDP abstractions of the model above. However, while our approach also works when the distribution of the noise is unknown, StocHy requires precise knowledge of the noise distribution. Another difference is that StocHy discretizes the control inputs of the dynamical system, to obtain a finite action space. An iMDP action is then associated with the execution of a specific value of the control input \mathbf{u}_k . In our approach, an iMDP action instead reflects an attempted transition to a given target point, and the actual control input \mathbf{u}_k (calculated using the control law Eq. (8)) depends on the precise continuous state \mathbf{x}_k .

Our approach is faster. The run times and model sizes of our approach are reported in Table 2 under BAS (1-zone). For our approach, the time to compute the states and actions (needed only in the first iteration) is 0.1 s, and run times per iteration (excluding verification times) vary from 1.7 – 7.2 s, depending on the value of N . Notably, the StocHy run times are *orders of magnitude higher*: the abstraction procedure of StocHy (excluding verification time) took 45 min, compared to 7.2 s for our approach with the maximum of $N = 12800$ samples (as reported in Table 2). We omit a comparison of the verification times, since StocHy does not leverage PRISM like our method does.

These results provide a clear message: our approach is more general and significantly faster, and (as shown earlier in Fig. 13) generates results that are qualitatively similar to those obtained from StocHy.

F.2 Benchmark Against SReachTools

We benchmark our approach against the spacecraft docking problem supplied with the MATLAB toolbox SReachTools (Vinod et al., 2019). In this problem, one spacecraft must dock with another within 8 time steps, while remaining in a target tube. The goal is to compute a controller that maximizes the probability of achieving this objective. The 4-dimensional state $\mathbf{x} = [p_x, p_x, v_x, v_y]^T \in \mathbb{R}^4$ describes the position and velocity in both directions. We adopt the same discrete-time dynamics as used in Vinod et al. (2019) and assume the same Gaussian noise (see the reference for details).

This problem is a finite-horizon reach-avoid problem with a rectangular goal region (target set), while the safe set is a convex tube, as also shown in Fig. 14. For our approach, we partition the state space into 3 200 rectangular regions. It is fair to note that the safe set in SReachTools is smooth, while ours is not (due to our partition-based approach). While our current implementation is limited to regions as hyper-rectangles and parallelepipeds, our method can in theory be used with all convex polytopic regions.

Reach-avoid guarantees. We compare our method (with $N = 25, \dots, 1600$ samples of the noise) to the results obtained via the different methods in SReachTools. We only consider the particle and Voronoi methods of SReachTools, because only these methods are sampling-based like our approach (although only the Voronoi method can give confidence guarantees on the results). The reach-avoid guarantees and run times are presented in Table 3, and simulated trajectories in Fig. 14. As expected, our reach-avoid guarantees, as well as for the Voronoi method, are a lower bound on the average simulated performance (and are thus safe), while this is not the case for the particle method of SReachTools.

Complexity and run time. As shown in Table 3, the grid-free methods of SReachTools are generally faster than our abstraction-based approach. However, we note that our method was designed for non-convex problems, which cannot be solved by SReachTools. Interestingly, the complexity of the particle (Lesser et al., 2013) and Voronoi partition method (Sartipizadeh et al., 2019) increases *exponentially* with the number of samples, because their optimization problems have a binary variable for every particle. By contrast, the complexity of our method increases only *linearly* with the number of samples, because it suffices to count the number of samples in every region, as discussed in Sect. 4.2.

Controller type. The particle and Voronoi methods of SReachTools synthesize *open-loop* controllers, which means that they cannot act in response to observed disturbances of the state. Open-loop controllers do not consider any feedback, making such controllers unsafe in settings with significant noise levels (Åström & Murray, 2010). By contrast, we derive *piecewise linear feedback* controllers. This structure is obtained because our controllers are based on a state-based control policy that maps every continuous state within the same region to the same abstract action.

Robustness against disturbances. To demonstrate the importance of feedback control, we test both methods under stronger disturbances, by increasing the covariance of the noise by a factor of 10. As shown in Table 3, the particle method yields a low reach-avoid probability (with the simulated performance being even lower), and the Voronoi method was not able to generate a solution at all. By contrast, our method still provides safe lower bound guarantees, which become tighter for increasing sample sizes. Our state-based controller is robust against the stronger noise, because it is able to act in response to observed disturbances, unlike the open-loop methods of SReachTools that results in a larger error in the simulated state trajectories, as also shown in Fig. 14.

References

Abate, A., Katoen, J., Lygeros, J., & Prandini, M. (2010). Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16(6), 624–641.

- Abate, A. (2011). Approximation metrics based on probabilistic bisimulations for general state-space markov processes: A survey. In *Hybrid Autonomous Systems@ETAPS*, Vol. 297 of *Electronic Notes in Theoretical Computer Science*, pp. 3–25. Elsevier.
- Abate, A., Prandini, M., Lygeros, J., & Sastry, S. (2008). Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11), 2724 – 2734.
- Alur, R., Henzinger, T. A., Lafferriere, G., & Pappas, G. J. (2000). Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(7), 971–984.
- Anderson, B. D., & Moore, J. B. (2007). *Optimal control: linear quadratic methods*. Courier Corporation.
- Ashok, P., Kretínský, J., & Weininger, M. (2019). PAC statistical model checking for markov decision processes and stochastic games. In *CAV (1)*, Vol. 11561 of *Lecture Notes in Computer Science*, pp. 497–519. Springer.
- Åström, K. J., & Murray, R. M. (2010). *Feedback systems: an introduction for scientists and engineers*. Princeton university press.
- Badings, T. S., Abate, A., Jansen, N., Parker, D., Poonawala, H. A., & Stoelinga, M. (2022a). Sampling-based robust control of autonomous systems with non-gaussian noise. In *AAAI*, pp. 9669–9678. AAAI Press.
- Badings, T. S., Cubuktepe, M., Jansen, N., Junges, S., Katoen, J., & Topcu, U. (2022b). Scenario-based verification of uncertain parametric mdps. *Int. J. Softw. Tools Technol. Transf.*, 24(5), 803–819.
- Badings, T. S., Jansen, N., Junges, S., Stoelinga, M., & Volk, M. (2022c). Sampling-based verification of ctmc with uncertain rates. In *CAV (2)*, Vol. 13372 of *Lecture Notes in Computer Science*, pp. 26–47. Springer.
- Badings, T. S., Jansen, N., Poonawala, H. A., & Stoelinga, M. (2021). Filter-based abstractions with correctness guarantees for planning under uncertainty. *CoRR*, abs/2103.02398.
- Badings, T. S., Romao, L., Abate, A., & Jansen, N. (2022). Probabilities are not enough: Formal controller synthesis for stochastic dynamical models with epistemic uncertainty. *CoRR*, abs/2210.05989.
- Baier, C., & Katoen, J. (2008). *Principles of model checking*. MIT Press.
- Bansal, S., Chen, M., Herbert, S. L., & Tomlin, C. J. (2017). Hamilton-jacobi reachability: A brief overview and recent advances. In *CDC*, pp. 2242–2253. IEEE.
- Ben-Tal, A., Ghaoui, L. E., & Nemirovski, A. (2009). *Robust Optimization*, Vol. 28 of *Princeton Series in Applied Mathematics*. Princeton University Press.
- Berkenkamp, F., Turchetta, M., Schoellig, A. P., & Krause, A. (2017). Safe model-based reinforcement learning with stability guarantees. In *NIPS*, pp. 908–918.
- Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Rev.*, 53(3), 464–501.

- Blackmore, L., Ono, M., Bektassov, A., & Williams, B. C. (2010). A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Trans. Robotics*, 26(3), 502–517.
- Bøhn, E., Coates, E. M., Moe, S., & Johansen, T. A. (2019). Deep reinforcement learning attitude control of fixed-wing uavs using proximal policy optimization. In *ICUAS*, pp. 523–533. IEEE.
- Boucheron, S., Lugosi, G., & Massart, P. (2013). *Concentration Inequalities - A Nonasymptotic Theory of Independence*. Oxford University Press.
- Boyd, S. P., & Vandenberghe, L. (2014). *Convex Optimization*. Cambridge University Press.
- Brafman, R. I., & Tennenholtz, M. (2002). R-MAX - A general polynomial time algorithm for near-optimal reinforcement learning. *J. Mach. Learn. Res.*, 3, 213–231.
- Brunke, L., Greeff, M., Hall, A. W., Yuan, Z., Zhou, S., Panerati, J., & Schoellig, A. P. (2022). Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems*, 5(1), 411–444.
- Calafiore, G. C., & Campi, M. C. (2005). Uncertain convex programs: randomized solutions and confidence levels. *Math. Program.*, 102(1), 25–46.
- Campi, M. C., Carè, A., & Garatti, S. (2021). The scenario approach: A tool at the service of data-driven decision making. *Annu. Rev. Control.*, 52, 1–17.
- Campi, M. C., & Garatti, S. (2008). The exact feasibility of randomized solutions of uncertain convex programs. *SIAM J. Optim.*, 19(3), 1211–1230.
- Campi, M. C., & Garatti, S. (2011). A sampling-and-discarding approach to chance-constrained optimization: Feasibility and optimality. *J. Optim. Theory Appl.*, 148(2), 257–280.
- Campi, M. C., & Garatti, S. (2018). *Introduction to the scenario approach*. SIAM.
- Casella, G., & Berger, R. L. (2021). *Statistical inference*. Cengage Learning.
- Cauchi, N., & Abate, A. (2018). Benchmarks for cyber-physical systems: A modular model library for building automation systems (extended version). *CoRR*, abs/1803.06315.
- Cauchi, N., & Abate, A. (2019). Stochy: Automated verification and synthesis of stochastic processes. In *TACAS (2)*, Vol. 11428 of *Lecture Notes in Computer Science*, pp. 247–264. Springer.
- Cauchi, N., Laurenti, L., Lahijanian, M., Abate, A., Kwiatkowska, M., & Cardelli, L. (2019). Efficiency through uncertainty: scalable formal synthesis for stochastic hybrid systems. In *HSCC*, pp. 240–251. ACM.
- Chen, X., Abraham, E., & Sankaranarayanan, S. (2013). Flow*: An analyzer for non-linear hybrid systems. In *CAV*, Vol. 8044 of *Lecture Notes in Computer Science*, pp. 258–263. Springer.
- Ciesinski, F., & Größer, M. (2004). On probabilistic computation tree logic. In *Validation of Stochastic Systems*, Vol. 2925 of *Lecture Notes in Computer Science*, pp. 147–188. Springer.

- Clohessy, W., & Wiltshire, R. (1960). Terminal guidance system for satellite rendezvous. *Journal of the Aerospace Sciences*, 27(9), 653–658.
- Desharnais, J., Edalat, A., & Panangaden, P. (2002). Bisimulation for labelled markov processes. *Inf. Comput.*, 179(2), 163–193.
- D’Innocenzo, A., Abate, A., & Katoen, J. (2012). Robust PCTL model checking. In *HSCC*, pp. 275–286. ACM.
- Dryden, H. L. (1943). A review of the statistical theory of turbulence. *Quarterly of Applied Mathematics*, 1(1), 7–42.
- Esfahani, P. M., & Kuhn, D. (2018). Data-driven distributionally robust optimization using the wasserstein metric: performance guarantees and tractable reformulations. *Math. Program.*, 171(1-2), 115–166.
- Fan, C., Qin, Z., Mathur, U., Ning, Q., Mitra, S., & Viswanathan, M. (2022). Controller synthesis for linear system with reach-avoid specifications. *IEEE Trans. Autom. Control.*, 67(4), 1713–1727.
- Ferns, N., Panangaden, P., & Precup, D. (2004). Metrics for finite markov decision processes. In *AAAI*, pp. 950–951. AAAI Press / The MIT Press.
- Ferns, N., Precup, D., & Knight, S. (2014). Bisimulation for markov decision processes through families of functional expressions. In *Horizons of the Mind*, Vol. 8464 of *Lecture Notes in Computer Science*, pp. 319–342. Springer.
- Fisac, J. F., Akametalu, A. K., Zeilinger, M. N., Kaynama, S., Gillula, J. H., & Tomlin, C. J. (2019). A general safety framework for learning-based control in uncertain robotic systems. *IEEE Trans. Autom. Control.*, 64(7), 2737–2752.
- Fu, J., & Topcu, U. (2014). Probably approximately correct MDP learning and control with temporal logic constraints. In *Robotics: Science and Systems*.
- Garatti, S., & Campi, M. C. (2022). Risk and complexity in scenario optimization. *Math. Program.*, 191(1), 243–279.
- García, J., & Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *J. Mach. Learn. Res.*, 16, 1437–1480.
- Givan, R., Dean, T. L., & Greig, M. (2003). Equivalence notions and model minimization in markov decision processes. *Artif. Intell.*, 147(1-2), 163–223.
- Givan, R., Leach, S. M., & Dean, T. L. (2000). Bounded-parameter markov decision processes. *Artif. Intell.*, 122(1-2), 71–109.
- Goh, J., & Sim, M. (2010). Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1), 902–917.
- Haesaert, S., Soudjani, S., & Abate, A. (2017). Verification of general markov decision processes by approximate similarity relations and policy refinement. *SIAM Journal on Control and Optimisation*, 55(4), 2333–2367.
- Hahn, E. M., Hashemi, V., Hermanns, H., Lahijanian, M., & Turrini, A. (2017). Multi-objective robust strategy synthesis for interval markov decision processes. In *QEST*, Vol. 10503 of *Lecture Notes in Computer Science*, pp. 207–223. Springer.

- Hahn, E. M., Hashemi, V., Hermanns, H., & Turrini, A. (2016). Exploiting robust optimization for interval probabilistic bisimulation. In *QEST*, Vol. 9826 of *Lecture Notes in Computer Science*, pp. 55–71. Springer.
- Hansson, H., & Jonsson, B. (1994). A logic for reasoning about time and reliability. *Formal Aspects Comput.*, 6(5), 512–535.
- Hashemi, V., Hatefi, H., & Krcál, J. (2014). Probabilistic bisimulations for PCTL model checking of interval mdps (extended version). In *SynCoP*, Vol. 145 of *EPTCS*, pp. 19–33.
- Hashemi, V., Hermanns, H., Song, L., Subramani, K., Turrini, A., & Wojciechowski, P. (2016). Compositional bisimulation minimization for interval markov decision processes. In *LATA*, Vol. 9618 of *Lecture Notes in Computer Science*, pp. 114–126. Springer.
- Hausser, D. (1990). Probably approximately correct learning. In *AAAI*, pp. 1101–1108. AAAI Press / The MIT Press.
- Herbert, S. L., Chen, M., Han, S., Bansal, S., Fisac, J. F., & Tomlin, C. J. (2017). Fastrack: A modular framework for fast and guaranteed safe motion planning. In *CDC*, pp. 1517–1522. IEEE.
- Hermanns, H., Parma, A., Segala, R., Wachter, B., & Zhang, L. (2011). Probabilistic logical characterization. *Information and Computation*, 209(2), 154–172.
- Hewing, L., Kabzan, J., & Zeilinger, M. N. (2020). Cautious model predictive control using gaussian process regression. *IEEE Trans. Control. Syst. Technol.*, 28(6), 2736–2743.
- Jewison, C., & Erwin, R. S. (2016). A spacecraft benchmark problem for hybrid control and estimation. In *CDC*, pp. 3300–3305. IEEE.
- Kearns, M. J., & Singh, S. P. (2002). Near-optimal reinforcement learning in polynomial time. *Mach. Learn.*, 49(2-3), 209–232.
- Koller, T., Berkenkamp, F., Turchetta, M., & Krause, A. (2018). Learning-based model predictive control for safe exploration. In *CDC*, pp. 6059–6066. IEEE.
- Kulakowski, B. T., Gardner, J. F., & Shearer, J. L. (2007). *Dynamic modeling and control of engineering systems*. Cambridge University Press.
- Kwiatkowska, M. Z., Norman, G., & Parker, D. (2011). PRISM 4.0: Verification of probabilistic real-time systems. In *CAV*, Vol. 6806 of *Lecture Notes in Computer Science*, pp. 585–591. Springer.
- Kwiatkowska, M. Z., & Parker, D. (2013). Automated verification and strategy synthesis for probabilistic systems. In *ATVA*, Vol. 8172 of *Lecture Notes in Computer Science*, pp. 5–22. Springer.
- Lahijanian, M., Andersson, S. B., & Belta, C. (2015). Formal verification and synthesis for discrete-time stochastic systems. *IEEE Trans. Autom. Control.*, 60(8), 2031–2045.
- Larsen, K. G., & Skou, A. (1991). Bisimulation through probabilistic testing. *Inf. Comput.*, 94(1), 1–28.

- Lavaei, A., Soudjani, S., Abate, A., & Zamani, M. (2022). Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica*, 146, 110617.
- Lesser, K., Oishi, M. M. K., & Erwin, R. S. (2013). Stochastic reachability for control of spacecraft relative motion. In *CDC*, pp. 4705–4712. IEEE.
- Lun, Y. Z., Wheatley, J., D’Innocenzo, A., & Abate, A. (2018). Approximate abstractions of markov chains with interval decision processes. In *ADHS*, Vol. 51 of *IFAC-PapersOnLine*, pp. 91–96. Elsevier.
- Margellos, K., Goulart, P., & Lygeros, J. (2014). On the road between robust optimization and the scenario approach for chance constrained optimization problems. *IEEE Transactions on Automatic Control*, 59(8), 2258–2263.
- Moos, J., Hansel, K., Abdulsamad, H., Stark, S., Clever, D., & Peters, J. (2022). Robust reinforcement learning: A review of foundations and recent advances. *Mach. Learn. Knowl. Extr.*, 4(1), 276–315.
- Morimoto, J., & Doya, K. (2005). Robust reinforcement learning. *Neural Comput.*, 17(2), 335–359.
- Nilim, A., & Ghaoui, L. E. (2005). Robust control of markov decision processes with uncertain transition matrices. *Oper. Res.*, 53(5), 780–798.
- Ogata, K., et al. (2010). *Modern control engineering*, Vol. 5. Prentice hall Upper Saddle River, NJ.
- Park, S., Serpedin, E., & Qaraqe, K. A. (2013). Gaussian assumption: The least favorable but the most useful [lecture notes]. *IEEE Signal Process. Mag.*, 30(3), 183–186.
- Peng, X. B., Andrychowicz, M., Zaremba, W., & Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In *ICRA*, pp. 1–8. IEEE.
- Pinto, L., Davidson, J., Sukthankar, R., & Gupta, A. (2017). Robust adversarial reinforcement learning. In *ICML*, Vol. 70 of *Proceedings of Machine Learning Research*, pp. 2817–2826. PMLR.
- Puggelli, A., Li, W., Sangiovanni-Vincentelli, A. L., & Seshia, S. A. (2013). Polynomial-time verification of PCTL properties of mdps with convex uncertainties. In *CAV*, Vol. 8044 of *Lecture Notes in Computer Science*, pp. 527–542. Springer.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley.
- Reissig, G., Weber, A., & Rungger, M. (2017). Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Trans. Autom. Control.*, 62(4), 1781–1796.
- Reist, P., Preiswerk, P., & Tedrake, R. (2016). Feedback-motion-planning with simulation-based lqr-trees. *Int. J. Robotics Res.*, 35(11), 1393–1416.
- Romao, L., Papachristodoulou, A., & Margellos, K. (2022). On the exact feasibility of convex scenario programs with discarded constraints. *IEEE Transactions on Automatic Control*, to appear.
- Rosolia, U., Singletary, A., & Ames, A. D. (2022). Unified multirate control: From low-level actuation to high-level planning. *IEEE Transactions on Automatic Control*, 67(12), 6627–6640.

- Sartipizadeh, H., Vinod, A. P., Açikmese, B., & Oishi, M. (2019). Voronoi partition-based scenario reduction for fast sampling-based stochastic reachability computation of linear systems. In *ACC*, pp. 37–44. IEEE.
- Shmarov, F., & Zuliani, P. (2015). Probreach: verified probabilistic delta-reachability for stochastic hybrid systems. In *HSCC*, pp. 134–139. ACM.
- Smith, A. (2013). *Sequential Monte Carlo methods in practice*. Springer Science & Business Media.
- Soudjani, S. E. Z., & Abate, A. (2013). Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM J. Appl. Dyn. Syst.*, *12*(2), 921–956.
- Sullivan, T. J. (2015). *Introduction to uncertainty quantification*, Vol. 63. Springer.
- Taylor, A. J., Singletary, A., Yue, Y., & Ames, A. D. (2020). Learning for safety-critical control with control barrier functions. In *L4DC*, Vol. 120 of *Proceedings of Machine Learning Research*, pp. 708–717. PMLR.
- Tedrake, R., Manchester, I. R., Tobenkin, M. M., & Roberts, J. W. (2010). Lqr-trees: Feedback motion planning via sums-of-squares verification. *Int. J. Robotics Res.*, *29*(8), 1038–1052.
- Tkachev, I., & Abate, A. (2014). Characterization and computation of infinite horizon specifications over markov processes. *Theoretical Computer Science*, *515*, 1–18.
- Vinitsky, E., Du, Y., Parvate, K., Jang, K., Abbeel, P., & Bayen, A. M. (2020). Robust reinforcement learning using adversarial populations. *CoRR*, *abs/2008.01825*.
- Vinod, A. P., Gleason, J. D., & Oishi, M. M. K. (2019). Sreachtools: a MATLAB stochastic reachability toolbox. In *HSCC*, pp. 33–38. ACM.
- Wiesemann, W., Kuhn, D., & Sim, M. (2014). Distributionally robust convex optimization. *Oper. Res.*, *62*(6), 1358–1376.
- Wolff, E. M., Topcu, U., & Murray, R. M. (2012). Robust control of uncertain markov decision processes with temporal logic specifications. In *CDC*, pp. 3372–3379. IEEE.
- Xu, H., & Mannor, S. (2010). Distributionally robust markov decision processes. In *NIPS*, pp. 2505–2513. Curran Associates, Inc.
- Zikelic, D., Lechner, M., Henzinger, T. A., & Chatterjee, K. (2022). Learning control policies for stochastic systems with reach-avoid guarantees. *CoRR*, *abs/2210.05308*.