

# Symbolic Controller Synthesis for Büchi Specifications on Stochastic Systems

Rupak Majumdar  
MPI-SWS, Germany.  
rupak@mpi-sws.org

Kaushik Mallik  
MPI-SWS, Germany.  
kmallik@mpi-sws.org

Sadegh Soudjani  
Newcastle University, UK.  
sadegh.soudjani@ncl.ac.uk

## ABSTRACT

We consider the policy synthesis problem for continuous-state controlled Markov processes evolving in discrete time, when the specification is given as a Büchi condition (visit a set of states infinitely often). We decompose computation of the maximal probability of satisfying the Büchi condition into two steps. The first step is to compute the maximal *qualitative winning set*, from where the Büchi condition can be enforced with probability one. The second step is to find the maximal probability of reaching the already computed qualitative winning set. In contrast with finite-state models, we show that such a computation only gives a lower bound on the maximal probability where the gap can be non-zero.

In this paper we focus on approximating the qualitative winning set, while pointing out that the existing approaches for unbounded reachability computation can solve the second step. We provide an abstraction-based technique to approximate the qualitative winning set by simultaneously using an over- and under-approximation of the probabilistic transition relation. Since we are interested in qualitative properties, the abstraction is non-probabilistic; instead, the probabilistic transitions are assumed to be under the control of a (fair) adversary. Thus, we reduce the original policy synthesis problem to a Büchi game under a fairness assumption and characterize upper and lower bounds on winning sets as nested fixed point expressions in the  $\mu$ -calculus. This characterization immediately provides a symbolic algorithm scheme. Further, a winning strategy computed on the abstract game can be refined to a policy on the controlled Markov process.

We describe a concrete abstraction procedure and demonstrate our algorithm on two case studies. We show that our techniques are able to provide tight approximations to the qualitative winning set for the Van der Pol oscillator and a 3-d Dubin's vehicle.

## 1 INTRODUCTION

Decision making under stochastic uncertainty has many applications in science, engineering, and economics. Typically, one models a system with uncertainty as a *controlled Markov*

*process* evolving in time. Such a process consists in a (possibly uncountable) set of states and actions. At a given state, an agent picks an action and the state and the action together determine the distribution over the next states. The choice of the control action depends on the history of states seen so far and may be randomized; the decision rule that assigns to each history a distribution of control actions is called a *policy*. Given a temporal specification over trajectories, the goal of the agent is to find an *optimal* policy: one that maximizes the probability that the resulting trajectory of the system satisfies the specification. The control problem asks, given a controlled Markov process and a temporal specification (given, e.g., in linear temporal logic), to design an optimal policy. In the finite-state setting, the control problem can be solved algorithmically based on graph traversal and linear programming [1, 3, 6]. A lot of recent research has focused on extending algorithmic policy synthesis techniques to *continuous-state* systems. The goal for continuous-state systems is to provide approximations to the probability of satisfaction, while providing formal guarantees on convergence of the error.

While synthesis for reachability and safety properties have been studied in this setting both for infinite horizon [9, 30] and for finite horizon [12, 13, 15–17, 23, 31], there are few techniques for synthesis against *Büchi* specifications, which requires the trajectory to visit a given set of states infinitely often. In this paper, we consider the problem of synthesizing controllers for controlled Markov processes for properties specified as Büchi conditions.

The key aspect of the solution in the finite-state case is to separate a synthesis problem into a *qualitative* part (find the set of states from which the agent has a policy to satisfy the property almost surely) and a *quantitative* part (find the policy that maximizes the probability of reaching the qualitatively winning states). Given the qualitative solution, one can iteratively compute the quantitative solution by solving a reachability problem, where the target is the absorbing set given by the qualitative solution.

Our first contribution is to show that a similar decomposition for Büchi properties does not hold for continuous state systems in general: we provide an example of a Markov

process over continuous state space for which the qualitative winning set (from which there is a policy that ensures the Büchi property holds almost surely) is empty but the maximal probability of satisfying the property has a non-zero solution. Moreover, we show that such a decomposition is able to provide a lower bound on the quantitative part of the problem. Thus, if one can compute (an approximation of) the winning set, a lower bound on the quantitative solution can be obtained by Bellman iteration or by other techniques for (unbounded) reachability in the continuous-state setting [9, 30].

As our second contribution, we provide symbolic algorithms for computing under- and over-approximations of the qualitative winning set. We compute finite-state abstractions of the continuous-state system. Our abstraction uses two transition relations: an over-approximation and an under-approximation of the continuous transitions. For qualitative probabilistic analysis on finite-state systems, one can replace the probabilistic transitions by an adversarial scheduler with a fairness requirement [10]. Accordingly, our abstractions are non-probabilistic and only require the knowledge of the support of the stochastic kernel associated to the process. We characterize the qualitative winning states as a nested fixed point expression in the  $\mu$ -calculus [14]; such an expression naturally gives a symbolic implementation. Since the abstraction is non-probabilistic, the symbolic implementation avoids numerical issues and can use standard encodings based on satisfiability checkers or binary decision diagrams.

Our fixed point characterization is similar to the characterization of qualitative almost-sure winning in concurrent games [7], but the use of two kinds of transitions acting as upper and lower bounds is a key distinguishing ingredient in our characterization. We show through examples why both are required. The qualitative winning states in the original continuous-state process are not characterized by a similar fixed point: it is well known that unlike the finite-state case, the actual probability values matter in deciding qualitative winning in finite-state systems [1, pp. 779-780].

We demonstrate our approach on the Van der Pol oscillator and the 3-d Dubin's vehicle, both in the presence of stochastic perturbation. Our computation shows that when the disturbance is treated as a worst case adversary, there exists a deterministic value of the disturbance for which the specification is violated for all initial states. On the contrary, when the disturbance is treated as random, we are able to satisfy the specification almost surely. Moreover, we empirically show that the difference between the over- and under-approximation reduces as we pick finer abstractions.

There are relatively few results on the algorithmic analysis of liveness properties for controlled Markov processes. A theoretical study of the Büchi objective  $B$  is conducted by Tkachev et al. [30] via persistence properties  $B$ . It is

shown that  $P_S^1 \cap B^0$  can be characterized by two fixed-point equations but no computational method is provided. In particular, their techniques do not provide a way to solve the qualitative Büchi problem that we solve. Our computational approach is similar in nature to results that employ Interval MDP or Interval MC as abstractions. For example, recent work by Dutreix and Coogan [8] utilizes Interval MC for verification of a particular class of systems. The method of Dutreix and Coogan requires numerical computations of lower and upper bounds of the probabilities and provides an enumerative algorithm. Our approach generalizes their construction of the over- and under-approximations to the setting of controlled Markov processes and the synthesis problem. We focus on the winning region of the specification, which allows us to write symbolic algorithms purely on non-probabilistic structures, thus avoiding numerical optimization procedures. Once the winning region is approximated, a quantitative reachability can be solved using standard techniques.

## 2 CONTROLLED MARKOV PROCESSES

### 2.1 Preliminaries

We consider a probability space  $(\Omega; \mathcal{F}; P)$ , where  $\Omega$  is the sample space,  $\mathcal{F}$  is a sigma-algebra on  $\Omega$  comprising subsets of  $\Omega$  as events, and  $P$  is a probability measure that assigns probabilities to events. We assume that random variables introduced in this article are measurable functions of the form  $X : \Omega \rightarrow \mathbb{R}^n$ . Any random variable  $X$  induces a probability measure on its space  $\mathbb{R}^n$  as  $P_X(A) = P(\omega \in \Omega \mid X(\omega) \in A)$  for any  $A \subseteq \mathbb{R}^n$ . We often directly discuss the probability measure  $P_X$  without explicitly mentioning the underlying probability space and the function  $X$  itself.

A topological space  $S$  is called a Borel space if it is homeomorphic to a Borel subset of a Polish space (i.e., a separable and completely metrizable space). Examples of a Borel space are the Euclidean space  $\mathbb{R}^n$ , its Borel subsets endowed with a subspace topology, as well as hybrid spaces. Any Borel space  $S$  is assumed to be endowed with a Borel sigma-algebra, which is denoted by  $\mathcal{B}^1 S$ . We say that a map  $f : S \rightarrow Y$  is measurable whenever it is Borel measurable.

We denote the set of nonnegative integers  $\mathbb{N}_0 := \{0, 1, 2, \dots\}$ .

### 2.2 Controlled Markov Processes

We consider controlled Markov processes (CMP) in discrete time defined over a general state space, characterized by a tuple  $S = (S; U; T_S^0)$ , where  $S$  is a Borel space as the state space of the process. We denote  $\mathcal{B}^1 S^0$  as the measurable space with  $\mathcal{B}^1 S^0$  being the Borel sigma-algebra on the

state space.  $\mathcal{U}$  is a finite space as the input space of the process.  $T_s : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \times \mathcal{U} \rightarrow [0, 1]$ , is a conditional stochastic kernel that assigns to any  $s \in \mathcal{S}$  and  $u \in \mathcal{U}$  a probability measure  $T_s(\cdot|s, u)$  on the measurable space  $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$  so that for any set  $A \in \mathcal{B}(\mathcal{S})$ ,  $P_{s,u}(A) = \int_A T_s(ds|s, u)$ , where  $P_{s,u}$  denotes the conditional probability  $P(\cdot|s, u)$ .

**REMARK 1.** *The input space  $\mathcal{U}$  in general can be any Borel space and the set of valid inputs can be state dependent. We have considered that  $\mathcal{U}$  is a finite set and all elements of this set can be taken at any state. This choice is motivated by the digital implementation of control policies and also facilitates concise presentation of the results.*

### 2.3 Semantics of controlled Markov processes

The semantics of a CMP is characterized by its *paths* or executions, which reflect both the history of previous states of the system and of implemented control inputs. Paths are used to measure the performance of the system.

*Definition 2.1.* Given a CMP  $\mathfrak{S}$ , a *finite path* is a sequence

$$w_n = (s_0, u_0, \dots, s_{n-1}, u_{n-1}, s_n), \quad n \in \mathbb{N},$$

where  $s_i \in \mathcal{S}$  are state coordinates and  $u_i \in \mathcal{U}$  are control input coordinates of the path. The space of all paths of length  $n$  is denoted by  $\text{PATH}_n := \mathcal{K}^n \times \mathcal{S}$  with  $\mathcal{K} := \mathcal{S} \times \mathcal{U}$ . Further, we denote projections by  $w_n[i] := s_i$  and  $w_n(i) := u_i$ . An *infinite path* of the CMP  $\mathfrak{S}$  is the sequence  $w = (s_0, u_0, s_1, u_1, \dots)$ , where  $s_i \in \mathcal{S}$  and  $u_i \in \mathcal{U}$  for all  $i \in \mathbb{N}$ . As above, let us introduce  $w[i] := s_i$  and  $w(i) := u_i$ . The space of all infinite paths is denoted by  $\text{PATH}_\infty := \mathcal{K}^\infty$ .

Given an infinite path  $w$  or a finite path  $w_n$ , we assume below that  $s_i$  and  $u_i$  are their state and control coordinates respectively, unless otherwise stated. For any infinite path  $w \in \text{PATH}_\infty$ , its  $n$ -prefix (ending in a state)  $w_n$  is a finite path of length  $n$ , which we also call  $n$ -*history*. We are now ready to introduce the notion of control policy.

*Definition 2.2.* A *policy* is a sequence  $\rho = (\rho_0, \rho_1, \rho_2, \dots)$  of universally measurable stochastic kernels  $\rho_n$  [2], each defined on the input space  $\mathcal{U}$  given  $\text{PATH}_n$ . The set of all policies is denoted by  $\Pi$ .

Given a policy  $\rho \in \Pi$  and a finite path  $w_n \in \text{PATH}_n$ , the distribution of the next control input  $u_n$  is given by  $\rho_n(\cdot|w_n)$ . In this work, we restrict our attention to the class of *stationary* policies.

*Definition 2.3.* A policy  $\rho$  is *stationary* if there is a universally measurable function  $C : \mathcal{S} \rightarrow \mathcal{U}$  such that at any time epoch  $n \in \mathbb{N}$ , the input  $u_n$  is taken to be  $C(s_n) \in \mathcal{U}$ . Namely, the stochastic kernel  $\rho_n(\cdot|w_n)$ ,  $n \in \mathbb{N}$ , in Definition 2.2 is a Dirac delta measure centered at  $C(s_n)$  with  $s_n = w_n[n]$

being the last element of  $w_n$ . We denote the class of stationary policies by  $\Pi_S \subset \Pi$  and a stationary policy just by the function  $C \in \Pi_S$ . The function  $C$  is also called *state feedback controller* in control theory.

For a CMP  $\mathfrak{S}$ , any policy  $\rho \in \Pi$  together with an initial probability measure  $\alpha : \mathcal{B}(\mathcal{S}) \rightarrow [0, 1]$  of the CMP induces a unique probability measure on the canonical sample space of paths [11] denoted by  $P_\alpha^\rho$  with the expectation  $\mathbb{E}_\alpha^\rho$ . In the case when the initial probability measure is supported on a single point, i.e.,  $\alpha(s) = 1$ , we write  $P_s^\rho$  and  $\mathbb{E}_s^\rho$  in place of  $P_\alpha^\rho$  and  $\mathbb{E}_\alpha^\rho$ , respectively. We denote the set of probability measures on  $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$  by  $\mathfrak{D}$ .

## 3 PROBLEM DEFINITION

**Liveness specification.** We consider liveness or repeated reachability specification as the synthesis objective, which is formalized using a predicate over the state space. Given a predicate  $B : \mathcal{S} \rightarrow \{0, 1\}$  over the state space, the liveness specification is denoted by  $\Box\Diamond B$  in linear temporal logic (LTL) notation [1]. An infinite path  $w \in \text{PATH}_\infty$  of a CMP  $\mathfrak{S}$  satisfies the liveness specification  $\Box\Diamond B$  if for all  $k_0 \in \mathbb{N}$ , there exists  $k_1 \in \mathbb{N}$  such that  $k_1 > k_0$  and  $B(w[k_1]) = 1$ . This requires that the path visits the set  $B^{-1}(1) \subseteq \mathcal{S}$  infinitely often. We indicate the set of all infinite paths  $w \in \text{PATH}_\infty$  of  $\mathfrak{S}$  that satisfy the property  $\Box\Diamond B$  by  $\mathfrak{S} \models \Box\Diamond B$ . With abuse of notation, we use  $B$  for both the predicate and the set  $B^{-1}(1)$ .

We are interested in the probability that the liveness specification can be satisfied by paths of a CMP  $\mathfrak{S}$  under different policies. Given a policy  $\rho \in \Pi$  and initial state  $s \in \mathcal{S}$ , we define the satisfaction probability as

$$f(s, \rho) := P_s^\rho(\mathfrak{S} \models \Box\Diamond B), \quad (1)$$

and the supremum satisfaction probability

$$f^*(s) := \sup_{\rho \in \Pi} P_s^\rho(\mathfrak{S} \models \Box\Diamond B). \quad (2)$$

**PROBLEM 1 (POLICY SYNTHESIS).** *Given  $\mathfrak{S}$  and the predicate  $B$ , find the optimal policy  $\rho^*$  along with  $f^*(s)$  s.t.  $P_s^{\rho^*}(\mathfrak{S} \models \Box\Diamond B) = f^*(s)$ .*

Measurability of the event  $\{\mathfrak{S} \models \Box\Diamond B\}$  in the canonical sample space of paths under the probability measure  $P_s^\rho$  is proved in [30]. An initial attempt is also made to study the properties of the function  $f^*(\cdot)$ . For instance, it is shown that  $f \equiv 1$  if and only if the probability that the path  $w$  reaches  $B$  is one for all initial states. We anticipate that the sets where  $f^*(s) = 1$  plays a crucial role in the computation of  $f^*$ .

*Definition 3.1 (Almost sure winning region).* Given the CMP  $\mathfrak{S}$ , the policy  $\rho$ , and the predicate  $B$ , the state  $s \in \mathcal{S}$  wins the specification  $\Box\Diamond B$  almost surely (a.s. in short) if  $f(s, \rho) = 1$ . The a.s. *winning region* of the policy  $\rho$  is defined as

$$\text{WinDom}(\mathfrak{S}, \rho) := \{s \in \mathcal{S} \mid P_s^\rho(\mathfrak{S} \models \Box\Diamond B) = 1\}. \quad (3)$$

We also define the *maximal a.s. winning region* as

$$\text{WinDom}^*(\mathfrak{S}) := \{s \in \mathcal{S} \mid \sup_{\rho \in \Pi} P_s^\rho(\mathfrak{S} \models \square \diamond B) = 1\}. \quad (4)$$

**PROPOSITION 3.2.** *The set  $\text{WinDom}^*(\mathfrak{S})$  is universally measurable. The set  $\text{WinDom}(\mathfrak{S}, C)$  is also universally measurable for any stationary policy  $C \in \Pi_S$ . The set  $W := \text{WinDom}(\mathfrak{S}, C)$  is an absorbing set, i.e., the paths starting from this set will stay in the set with probability one.*

The proof can be found in the appendix.

In the sequel, we restrict our attention to stationary policies  $C \in \Pi_S$  and decompose the computation of  $P_s^C(\mathfrak{S} \models \square \diamond B)$  into the computation of the winning set  $\text{WinDom}(\mathfrak{S}, C)$  and then computation of reachability probability  $P_s^C(\mathfrak{S} \models \diamond \text{WinDom}(\mathfrak{S}, C))$ . This is formalized next.

**ASSUMPTION 1.** *Stationary policies are sufficient for the computation of (2). Namely,  $f^*(s) = \sup_{C \in \Pi_S} f(s, C)$ .*

**THEOREM 3.3.** *We have for any policy  $C \in \Pi_S$  on CMP  $\mathfrak{S}$*

$$\begin{cases} P_s^C(\mathfrak{S} \models \square \diamond B) = 1 & \text{if } s \in W \\ P_s^C(\mathfrak{S} \models \square \diamond B) \geq P_s^C(\mathfrak{S} \models \diamond W) & \text{if } s \notin W, \end{cases} \quad (5)$$

with  $W := \text{WinDom}(\mathfrak{S}, C)$ .

The proof can be found in the appendix.

Computation of the reachability probability has been studied extensively in the literature for both infinite horizon [9, 28–30] and finite horizon [12, 13, 15–17, 24–26, 31] using different abstract models and computational methods. These approaches can be used to provide a lower bound on the probability of satisfaction of the Büchi condition. So from this point onward, we mostly consider the first half of (5) which is formalized as in the following:

**PROBLEM 2 (MAXIMAL WINNING REGION).** *Given  $\mathfrak{S}$  and predicate  $B$ , find a stationary policy  $C^*$  such that*

$$\text{WinDom}(\mathfrak{S}, C) \subseteq \text{WinDom}(\mathfrak{S}, C^*) \text{ almost surely}$$

for all  $C \in \Pi_S$ .

The maximal policy defined in Problem 2 is not necessarily unique, but the winning region associated to such maximal policies is unique. A formal treatment of this claim can be found in the appendix in Sec. 9.3.

Let us denote the winning region associated to the maximal policy by  $W^*$ . In the following, we mainly focus on the approximate computation  $W^*$  with a suitable policy.

For the reachability part, it has been shown that linking the infinite-horizon reachability to the finite-horizon one requires knowledge of the *absorbing sets* from which the trajectory cannot escape. So we briefly discuss how the absorbing sets can be over-approximated to enable linking the reachability in (5) to its finite-horizon version.

**A solution outline for Problem 2.** Computing the exact maximal policy for  $\mathfrak{S}$  is difficult in general. We propose an approximation procedure using a three-step abstraction-based method, outlined in the following:

**Abstraction.** First, the given CMP  $\mathfrak{S}$  is approximated using a finite state transition system  $\mathcal{A}$ , called the *abstraction*, usually by means of state space discretization [19, 21, 27]. The specification—which is the Büchi condition in our case—is also approximated using the discretized state space of the abstract transition system, and is called the *abstract specification*.

**Synthesis.** Second, the policy synthesis problem is posed as a zero-sum *game* on  $\mathcal{A}$  between the controller and an adversary, where in each step the controller chooses a control input, and the adversary chooses a successor allowed by the transitions in  $\mathcal{A}$ . The goal of the control player is a.s. satisfaction of the Büchi objective from as many states as possible, whereas the goal of the environment player is the complement of the same. The outcome of the game, when played from the perspective of the control player, is an *abstract controller*  $\hat{C}$ .

**Controller refinement.** Third, the abstract controller  $\hat{C}$  is mapped back to the continuous state space using a process called *controller refinement*. This results in a continuous controller  $C$  that can be paired with  $\mathfrak{S}$  and a continuous winning domain  $\text{WinDom}(\mathfrak{S}, C)$ .

## 4 A GENERIC FINITE STATE ABSTRACTION

Our proposed solution relies on constructing an abstraction  $\mathcal{A}$  which uses two transition functions to approximate the transition kernel of  $\mathfrak{S}$ .

**Definition 4.1.** A transition system is a tuple  $(Q, \Sigma, \delta_1, \delta_2)$  where  $Q$  is a finite set of states,  $\Sigma$  is a finite input alphabet, and  $\delta_i : Q \times \Sigma \rightarrow 2^Q$ ,  $i \in \{1, 2\}$  are two transition functions with the property that  $\delta_1(q, u) \supseteq \delta_2(q, u)$  for all  $(q, u) \in Q \times \Sigma$ .

The abstraction  $\mathcal{A}$  is constructed based on a finite partition of the state space. Therefore, we require a bounded state space  $\mathcal{S}$ . If  $\mathcal{S}$  is unbounded, we truncate it to a measurable set  $\mathcal{S}'$ , which serves as the working region; the rest is represented by a symbolic *sink* state  $\phi$ . The state  $\phi$  models all the out-of-domain behaviors of  $\mathfrak{S}$ , and this will be useful while doing synthesis in Sec. 5. The new CMP will be  $\mathfrak{S}' := (\mathcal{S}' \cup \{\phi\}, \mathcal{U}, T'_s)$  with

$$T'_s(A | s, u) = \begin{cases} T_s(A | s, u) & \text{if } s \in \mathcal{S}' \\ 0 & \text{if } s = \phi, \end{cases} \quad (6)$$

for any  $A \in \mathcal{B}(\mathcal{S}')$ , and  $T'_s(\phi | s, u) = 1 - T'_s(\mathcal{S}' | s, u)$ . In order to avoid a change of notation, we work in the sequel with

$\mathfrak{S} := (\mathcal{S}, \mathcal{U}, T_s)$  where the state space  $\mathcal{S}$  is bounded but may also include a symbolic sink state  $\phi$ . We also assume that  $B$  is fully contained in the working region.

#### 4.1 The abstraction

We propose a new type of abstraction, which will later be used in Sec. 5 to compute the two approximations of the maximal winning region and the associated policies.

First, we introduce some notation. Given the state space  $\mathcal{S}$  of the CMP  $\mathfrak{S}$ , we define a finite partition of  $\mathcal{S}$  denoted by  $\widehat{\mathcal{S}} := \{\widehat{x}_i\}_{i \in I}$  s.t.  $\mathcal{S} = \bigcup_{i \in I} \widehat{x}_i$  and  $\widehat{x}_i \cap \widehat{x}_j = \emptyset$  for all  $\widehat{x}_i, \widehat{x}_j \in \widehat{\mathcal{S}}$ ,  $i \neq j$ . The set  $\widehat{\mathcal{S}}$  will be called the abstract state space (state space of the abstraction), and each element  $\widehat{x}_i$  is an abstract state.

**REMARK 2.** *For the theory that is going to be presented in this paper, the abstract states need not be of the same size. However, for practical implementation, partition sets are chosen to be hyper-rectangular of the form  $\widehat{x} = \llbracket a, b \rrbracket$  where  $a, b \in \mathcal{S}$  are vectors. The partition sets are uniformly sized and their boundaries are assigned to only one partition element. We have also used a hyper-rectangular state space with an additional symbolic state  $\phi$  that is also an element of the abstract state space.*

**Definition 4.2.** Let  $\mathfrak{S} = (\mathcal{S}, \mathcal{U}, T_s)$  be a CMP, and  $\widehat{\mathcal{S}}$  be a finite partition of  $\mathcal{S}$ . Suppose  $\mathcal{A} = (\widehat{\mathcal{S}}, \mathcal{U}, \overline{F}, \underline{F})$  be a transition system with two modes of transitions  $\overline{F} : \widehat{\mathcal{S}} \times \mathcal{U} \rightarrow 2^{\widehat{\mathcal{S}}}$  and  $\underline{F} : \widehat{\mathcal{S}} \times \mathcal{U} \rightarrow 2^{\widehat{\mathcal{S}}}$ . The system  $\mathcal{A}$  will be called an *abstraction* of  $\mathfrak{S}$  if for all  $\widehat{x} \in \widehat{\mathcal{S}}$  and all  $u \in \mathcal{U}$ ,

$$\overline{F}(\widehat{x}, u) \supseteq \{\widehat{x}' \in \widehat{\mathcal{S}} \mid \exists s \in \widehat{x} . T_s(\widehat{x}' \mid s, u) > 0\}, \quad (7)$$

$$\underline{F}(\widehat{x}, u) \subseteq \{\widehat{x}' \in \widehat{\mathcal{S}} \mid \exists \varepsilon > 0 . \forall s \in \widehat{x} . T_s(\widehat{x}' \mid s, u) \geq \varepsilon\}. \quad (8)$$

In words, in the presence of the stochastic disturbance and given an abstract state  $\widehat{x}$  and a control input  $u \in \mathcal{U}$ ,  $\overline{F}(\widehat{x}, u)$  represents an over-approximation of the set of all abstract states which can be reached with positive probability from *some* continuous state in  $\widehat{x}$ , and  $\underline{F}(\widehat{x}, u)$  represents a subset of  $\overline{F}(\widehat{x}, u)$  under-approximating the set of states that can be reached with probability bounded away from 0 from *all* the states in  $\widehat{x}$ . We defer the actual computation of the abstract transition system until Sec. 6.

**REMARK 3.** *The abstract transition systems in the usual abstraction based control methods [19, 21, 27] play the role of a game graph for a two-player zero-sum game between the controller and an imaginary adversary, where the adversary is an accumulation of the external perturbation and the discretization-induced non-determinism. The rule of the game is that at each discrete step, the controller plays a control input, to which the adversary responds by choosing one of the many*

*non-deterministic successors. A control policy is synthesized for the controller by treating the adversary actions in worst-case fashion.*

*In our case, the controller effectively plays simultaneously against two imaginary adversaries who use two different types of actions: The first adversary—called the random adversary—uses the external random noise, while the second adversary—called the non-deterministic adversary—uses the discretization-induced non-determinism. This separation enables the controller to somewhat relax the worst-case treatment of the problem, by assuming that the random adversary is fair in choosing its actions, meaning all the noise values in the support of the distribution will appear always eventually. The non-deterministic adversary is still treated in the worst-case fashion.*

*Keeping this two-adversary interpretation in perspective, given some control action, one can interpret a  $\overline{F}$ -transition as a joint colluding move of the two adversaries, while one can interpret a  $\underline{F}$ -transition as a move of the lone random adversary. The rule of the game in our case is that at each discrete step, the controller plays a control input, to which either the adversaries jointly respond by choosing one of the  $\overline{F}$ -successors (which is possibly not an  $\underline{F}$ -successor), or the random adversary independently chooses an  $\underline{F}$ -successor. Since the random adversary is fair in its moves, hence it will not collude with the non-deterministic adversary all the time, and all the  $\underline{F}$ -transitions will be chosen at some point in the long run. This additional fairness assumption in the underlying game create favorable condition for the controller in many cases, as will be shown in the next section.*

#### 4.2 Almost sure progress

The fairness of the random adversary is materialized using the  $\varepsilon$  in the definition of  $\underline{F}$ , which guarantees that a trajectory eventually exits from an abstract state  $\widehat{x}$  in the long run, even when there is a non-zero probability for a single-step successor of a continuous state  $s \in \widehat{x}$  to stay within  $\widehat{x}$ . This feature is a central element of our synthesis method that will be presented in Sec. 5.

Following is a simple example on a general continuous state Markov chain that demonstrates the fact that in the absence of the bound  $\varepsilon$  in the definition of  $\underline{F}$ , trajectories could get trapped inside  $\widehat{x}$  forever.

**Example 4.3.** For simplicity and a focused exposition of the actual issue, we use a system admitting no control over state trajectories (can be alternatively thought of as a system with a control input space that has a single element). Consider a one dimensional CMP with state space  $\mathcal{S} = [0, 2]$ , and with



property, we have  $P_s^C(\mathcal{E} \models \diamond \square \widehat{x}) = \lim_{n \rightarrow \infty} W_n(s)$  with

$$W_{n+1}(s) = P_s^C(\mathcal{E} \models \square \widehat{x}) + \int_{S \setminus \widehat{x}} W_n(s') T_s(ds' | s, u),$$

which results in  $W_n(s) = 0$  for all  $n$ .  $\square$

We presented Ex. 4.3 to show the importance of having a uniform lower bound  $\varepsilon$  for the definition of  $\underline{F}$  in (8) and demonstrated that having  $T_s(\widehat{x}' | s, u) > 0$  for all  $s \in \widehat{x}$  is not sufficient. It can be shown that this condition is enough under proper continuity assumptions on the stochastic kernel.

**PROPOSITION 4.5.** *Suppose the kernel  $T_s$  is continuous, i.e.,  $g(s, u) = \int_S f(s') T_s(ds' | s, u)$  is upper semi-continuous for any upper semicontinuous function  $f$ . Moreover, the partition sets  $\widehat{S} = \{\widehat{x}_i\}_{i \in I}$  are such that  $T_s(\text{int}(\widehat{x}') | s, u) = T_s(\text{cl}(\widehat{x}') | s, u)$  for all  $s \in \widehat{x}$ ,  $u \in \mathcal{U}$ , and any  $\widehat{x}, \widehat{x}' \in \widehat{S}$ , where  $\text{int}(\cdot)$  and  $\text{cl}(\cdot)$  indicate respectively interior and closure of a set. Then,  $\underline{F}(\widehat{x}, u)$  can be defined alternatively as*

$$\underline{F}(\widehat{x}, u) \subseteq \{\widehat{x}' \in \widehat{S} \mid \forall s \in \text{cl}(\widehat{x}) . T_s(\widehat{x}' | s, u) > 0\}.$$

## 5 CONTROLLER SYNTHESIS AND REFINEMENT

With the abstraction  $\mathcal{A}$  of the given CMP  $\mathcal{E}$  computed in Sec. 4, we now propose algorithms to approximate—from above and below—the maximal a.s. winning region  $W^*$ . As a by-product, we will also obtain a suitable control policy. We first lift the specification  $\square \diamond B$  to an abstract specification that can be specified using the states of  $\mathcal{A}$ . For that we define an under-approximation  $\underline{B}$  and an over-approximation  $\overline{B}$  of the support of the predicate  $B$  using the state space of  $\mathcal{A}$ :

$$\begin{aligned} \underline{B} &:= \{\widehat{x} \in \widehat{S} \mid \forall s \in \widehat{x} . B(s) = 1\} \\ \overline{B} &:= \{\widehat{x} \in \widehat{S} \mid \exists s \in \widehat{x} . B(s) = 1\}. \end{aligned}$$

Note that the sink state  $\phi \notin \underline{B}$  since we assume the support of the predicate  $B$  is fully contained in the working region of the CMP. Hence, the satisfaction of  $\square \diamond \underline{B}$  would ensure that  $\phi$  is always avoided, as otherwise the specification will be violated.

In order to formalize the synthesis process, we first introduce four operators:

**Controllable predecessor:** Define  $Cpre_F : 2^{\widehat{S}} \rightarrow 2^{\widehat{S}}$  for  $F \in \{\overline{F}, \underline{F}\}$ ,

$$Cpre_F : T \mapsto \{\widehat{x} \in \widehat{S} \mid \exists u \in \mathcal{U} . F(\widehat{x}, u) \subseteq T\}. \quad (10)$$

**Cooperative predecessor:** Define  $Pre_F : 2^{\widehat{S}} \rightarrow 2^{\widehat{S}}$  for  $F \in \{\overline{F}, \underline{F}\}$ ,

$$Pre_F : T \mapsto \{\widehat{x} \in \widehat{S} \mid \exists u \in \mathcal{U} . F(\widehat{x}, u) \cap T \neq \emptyset\}. \quad (11)$$

**Almost sure predecessor:** Define  $Apr_e : 2^{\widehat{S}} \times 2^{\widehat{S}} \rightarrow 2^{\widehat{S}}$ ,

$$Apr_e : (T, S) \mapsto \{\widehat{x} \in \widehat{S} \mid \exists u \in \mathcal{U} . \overline{F}(\widehat{x}, u) \subseteq T \wedge \underline{F}(\widehat{x}, u) \cap S \neq \emptyset\}. \quad (12)$$

**Uncertain predecessor:** Define  $Upre : 2^{\widehat{S}} \times 2^{\widehat{S}} \rightarrow 2^{\widehat{S}}$ ,

$$Upre : (T, S) \mapsto \{\widehat{x} \in \widehat{S} \mid \exists u \in \mathcal{U} . \underline{F}(\widehat{x}, u) \subseteq T \wedge \overline{F}(\widehat{x}, u) \cap S \neq \emptyset\}. \quad (13)$$

### 5.1 Warm-up: reachability specification

As a warm-up, we first consider under-approximation of the largest winning domain for a.s. satisfaction of the reachability specification  $\diamond B$ . Motivated by the fixed point of a.s. reachability in concurrent two-player game [7] and for stochastic disturbances, a.s. reachability is implemented on the abstract system  $\mathcal{A}$  using the following nested fixed-point:

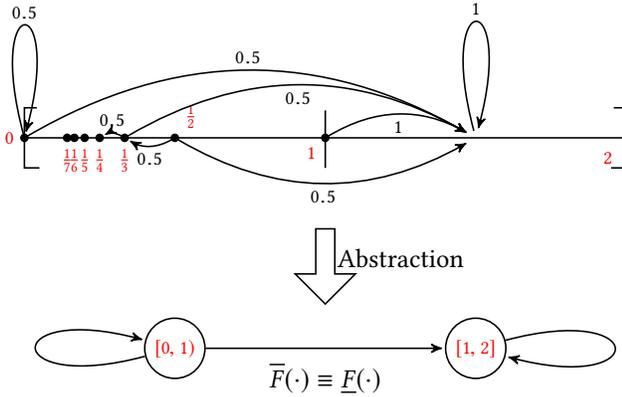
$$\nu Y . \mu Z . (\underline{B}^c \cap (Apr_e(Y, Z) \cup Cpre_{\overline{F}}(Z))) \cup \underline{B}, \quad (14)$$

where  $\underline{B}^c$  represents the complement of the set  $\underline{B}$ . Intuitively, the above fixed point computes the largest set  $Y$  s.t. from every state  $y \in Y \setminus \underline{B}$  there exists control input sequence s.t. either (a) there is a finite sequence of  $\underline{F}$ -transitions to  $\underline{B}$  and no  $\overline{F}$ -transition outside  $Y$ , or (b) there is a bounded finite sequence of  $\overline{F}$ -transitions all of whose non-deterministic branches reach  $\underline{B}$ . On the CMP level this means: for all  $y \in Y \setminus \underline{B}$  and for all  $s \in y$  there exist control input sequence s.t. either (a) there exist paths that enter  $\underline{B}$  with positive probability bounded away from zero and all paths stay inside  $Y$  with probability 1, or (b) there exist paths that enter  $\underline{B}$  with probability 1. It can be shown, by repeated use of Prop. 4.4 on the  $\underline{F}$ -transitions, that  $\underline{B}$  (and hence  $B$ ) will be reached a.s. by such a control sequence.

The fixed point (14) is in stark contrast with the usual reachability fixed point for worst case disturbances, which is given as  $\mu Z . Cpre_{\overline{F}}(Z) \cup \underline{B}$ , where it is required that from every state  $z \in Z$ , all the non-deterministic branches of  $\overline{F}$  reach  $\underline{B}$  in at most some finite number of steps. Note that the solution of (14) subsumes the solution of the usual fixed point, and in practice the usual reachability fixed point is much stronger than its stochastic counterpart. Following is an illustrative example that captures this intuition.

*Example 5.1.* Consider the CMP  $\mathcal{E}_2$  defined in Ex. 4.3, with stochastic kernel (9) and  $a(s) = 0.5$ . The transition probabilities from states in  $[0, 1)$  to the interval  $[1, 2]$  have a lower bound 0.5, and so unlike  $\mathcal{E}_1$  of Ex. 4.3,  $\underline{F}([0, 1)) = \{[0, 1), [1, 2]\}$ . Moreover, since there exist transitions with positive probability from all the state in  $[0, 1)$  to  $[1, 2]$ , hence  $\overline{F}([0, 1)) = \{[0, 1), [1, 2]\}$  as well.

Assume that the abstract specification for reachability is given as  $\diamond[1, 2]$ . If we start from some state in  $[0, 1)$  and treat the adversary (resolving the non-determinism) as worst-case,



**Figure 2: The top figure represents  $\mathfrak{S}_2$  and the bottom figure represents both transition relations  $\underline{F}([0, 1])$  and  $\overline{F}([0, 1])$  (same in this case) of  $\mathcal{A}$ . It is not possible to reach  $[1, 2]$  from  $[0, 1]$  if the non-determinism in  $\underline{F}$  or  $\overline{F}$  is treated as a worst-case adversary: the adversary can choose the loop at  $[0, 1)$  all the time. In the top figure, the labels on the edges are the probabilities of the associated transitions, and in both figures the labels in red are continuous states.**

then by using either of  $\underline{F}$  or  $\overline{F}$ , we will always loop in  $[0, 1)$  and would not be able to reach  $[1, 2]$  in any finite number of steps. Formally, the fixed point for  $\diamond[1, 2]$  will converge to the singleton set  $[1, 2]$ , since  $Cpre_{\overline{F}}(\{[1, 2]\}) = \{[1, 2]\}$ .

On the other hand, if we treat the disturbance as stochastic noise, then from Prop. 4.4 we know that if we loop in  $[0, 1)$  indefinitely long, then in the long run a.s. the system is going to move to  $[1, 2]$  (recall the interpretation of fair random adversary). So the winning region in this case should be the whole state space  $\{[0, 1), [1, 2]\}$ . Indeed, (14) will include the state  $[0, 1)$  since  $[0, 1) \in Apre(\{[0, 1), [1, 2]\}, \{[1, 2]\})$ .

**REMARK 4.** Nilsson et al. [19] introduced a particular type of augmented transition system as abstraction of non-stochastic system, which embeds liveness information in progress groups. If a set of abstract states form a progress group under some control action, then it is known that the system eventually leaves the progress group under repeated use of this particular control action. Even though our work deals with a completely unrelated problem, it is worthwhile to note that our fairness assumption on the random adversary allowing the CMP to make progress outside a given abstract state a.s. (Prop. 4.4) has very similar flavor.

## 5.2 Under-approximation of the maximal a.s. winning region

We build up on the intuition of the solution of the a.s. reachability specification, and present the computation of a sound under-approximation of the maximal a.s. winning region  $W^*$  with a suitable abstract controller  $\widehat{C}$ . In  $\mu$ -calculus notation, this under-approximation can be computed as:

$$\underline{W} := \nu Y . \mu Z . \left[ (\underline{B}^c \cap (Apre(Y, Z) \cup Cpre_{\overline{F}}(Z))) \cup (\underline{B} \cap Cpre_{\overline{F}}(Y)) \right]. \quad (15)$$

Note that the only new term in (15) as compared to (14) is the intersection of  $\underline{B}$  with  $Cpre_{\overline{F}}(Y)$ . This additional term makes sure that each time  $\underline{B}$  is reached, the winning region  $Y$  is not left in the next step to make sure that  $\underline{B}$  can be reached once again.

The fixed point (15) and the associated abstract controller  $\widehat{C}$  can be computed as the nested iteration given in Alg. 1. The controller  $\widehat{C}$  is a partial function from  $\widehat{\mathcal{S}}$  to  $\mathcal{U}$ , and we use the notation  $\text{dom } \widehat{C}$  to denote the domain of the controller  $\widehat{C}$ .

---

### Algorithm 1 Computation of $\underline{W}$

---

**Input:**  $\underline{B} \subseteq \widehat{\mathcal{S}}$

**Output:**  $\underline{W}, \widehat{C}$

```

1:  $Y \leftarrow \widehat{\mathcal{S}}, Y' \leftarrow \emptyset$ 
2: while  $Y \neq Y'$  do
3:    $Y' \leftarrow Y$ 
4:    $Z \leftarrow \emptyset, Z' \leftarrow \widehat{\mathcal{S}}$ 
5:    $\forall \widehat{x} \in \widehat{\mathcal{S}} . \widehat{C} : \widehat{x} \mapsto \emptyset$ 
6:   while  $Z \neq Z'$  do
7:      $Z' \leftarrow Z$ 
8:      $Z'' \leftarrow (\underline{B}^c \cap (Apre(Y, Z) \cup Cpre_{\overline{F}}(Z))) \cup (\underline{B} \cap Cpre_{\overline{F}}(Y))$ 
9:      $\forall \widehat{x} \in Z'' \setminus (\underline{B} \cup \text{dom } \widehat{C}) . \widehat{C} : \widehat{x} \mapsto u \text{ s.t. } \overline{F}(\widehat{x}, u) \subseteq Y \wedge \underline{F}(\widehat{x}, u) \cap Z \neq \emptyset \text{ or } \overline{F}(\widehat{x}, u) \subseteq Z$ 
10:     $Z \leftarrow Z''$ 
11:   end while
12:    $Y \leftarrow Z$ 
13: end while
14:  $\underline{W} \leftarrow Y$ 
15:  $\forall \widehat{x} \in \underline{W} \cap \underline{B} . \widehat{C} : \widehat{x} \mapsto u \text{ s.t. } \overline{F}(\widehat{x}, u) \subseteq \underline{W}$ 
16: return  $\underline{W}, \widehat{C}$ 

```

---

Note that, the existence of the control input  $u$  in Line (9) and (15) is guaranteed because of the definition of  $Cpre$  and  $Apre$ .

**PROPOSITION 5.2.** *The set  $\underline{W}$  is an under-approximation of the maximal a.s. winning region  $W^*$ .*

PROOF. The goal is to prove  $\overline{W} \subseteq W^*$ . Let  $q \in \widehat{S}$  be a state s.t.  $q \in \overline{W}$ . We show that  $q \in W^*$ . In the last iteration of the outer while loop in Alg. 1, we obtain a growing sequence of states  $Z_0 \subset Z_1 \subset \dots \subset Z_k = Y$ , where  $Z_0 = \emptyset$  and  $Y = \overline{W}$ . Since  $\text{Apre}(Y, \emptyset) = \text{Cpre}_{\overline{F}}(\emptyset) = \emptyset$ , hence  $Z_1 \subseteq \underline{B}$ . For all the other states  $\widehat{x}$  in  $Z_i$  for  $i \in (2; k]$ , one of the two cases happen: Either (a) by  $\text{Cpre}_{\overline{F}}(Z_{i-1})$ , it is ensured that  $Z_{i-1}$  is surely reached from  $\widehat{x}$  in one step, or (b) by  $\text{Apre}(Y, Z_{i-1})$ , it is ensured that  $Y$  (same as  $\overline{W}$  in the last iteration) is not left from  $\widehat{x}$ , and additionally (follows from Prop. 4.4) from all the (continuous) states inside  $\widehat{x}$ , transition to  $Z_{i-1}$  happens almost surely in the long run. Thus, for every  $q \in \overline{W} \setminus \underline{B}$  and for all  $x \in q$ ,  $\underline{B}$  is reached almost surely in the long run.

Moreover, the operator  $\text{Cpre}_{\overline{F}}(Y)$  ensures that  $Y$ —same as  $\overline{W}$  in the last iteration—is not left in the one step from  $Z_1$ . Hence almost surely  $\underline{B}$  is visited infinitely often.  $\square$

REMARK 5. The definition of the operators *Apre* and the fix-point (15) are highly motivated by how a.s. winning strategies are synthesized for Büchi specification in two-player concurrent games [7]. In the qualitative case of a two-player concurrent game with Büchi objective, the optimal strategy for each player is to play an action that surely keeps the game within the winning region, while making a progress towards the target with positive probability.

At a very high level, we use the same insight to express  $\overline{W}$  in (15), though for us the underlying game structure is totally different (see Rem. 3). It turns out that winning the game almost surely in our case means to either stay in the winning region using all  $\overline{F}$ -successors, while at the same time making progress using some  $\underline{F}$ -successor.

### 5.3 Controller refinement

The state space discretization during the abstraction process induces the following quantizer map  $Q : \mathcal{S} \rightarrow \widehat{S}$  s.t.  $Q : s \mapsto \widehat{x}$  when  $s \in \widehat{x}$ . Given the abstract controller  $\widehat{C} : \widehat{S} \rightarrow \mathcal{U}$ , we can obtain a continuous controller  $C : \mathcal{S} \rightarrow \mathcal{U}$  as  $C \equiv \widehat{C} \circ Q$ , where “ $\circ$ ” is the function composition. The following theorem states that the controller  $C$  is sound.

THEOREM 5.3. Consider the control policy  $C \equiv \widehat{C} \circ Q$  and the set  $\overline{W}$ , where  $\overline{W}$  and  $\widehat{C}$  are the outputs of Alg. 1 and  $Q$  is the quantizer map. Then,  $P_s^C(\text{C} \models \square \diamond B) = 1$  for all  $s \in \overline{W}$ .

The proof of the above theorem directly follows from the proof of Prop. 5.2, and hence is omitted.

### 5.4 Over-approximation of the maximal a.s. winning region

The over-approximation of  $W^*$  is given by the solution of the fixed-point

$$\overline{W} := \nu Y . \mu Z . \left[ (\overline{B}^c \cap \text{Upre}(Y, Z)) \cup \left( \overline{B} \cap (\text{Cpre}_{\overline{F}}(Y) \cup \text{Pre}_{\overline{F} \setminus \underline{F}}(Y)) \right) \right]. \quad (16)$$

The expression (16) can be solved in the same way as Alg. 1 by replacing the update in Line 8 with the update in the r.h.s. of (16). Also, Line 9 and 15 can be avoided, as the control policy in this case does not serve any useful purpose.

PROPOSITION 5.4. The set  $\overline{W}$  is a superset of the maximal a.s. winning region  $W^*$ .

PROOF. Let  $x^* \in W^*$ . Since  $\widehat{S}$  creates a cover of the state space  $\mathcal{S}$ , hence there exists a non-empty set  $\widehat{S}^* := \{\widehat{x} \in \widehat{S} \mid x^* \in \widehat{x}\}$ . We need to show that  $\widehat{S}^* \cap \overline{W} \neq \emptyset$ . For the sake of contradiction, assume that  $\widehat{S}^* \subseteq \overline{W}^c$ . We will show that this cannot happen.

The fixed point computation (16) produces a shrinking sequence of states  $\widehat{S} = Y^0 \supseteq Y^1 \supseteq \dots \supseteq Y^k = \overline{W}$ . Let  $\widehat{x}^* \in \widehat{S}^*$  be any abstract state and  $i \in \mathbb{N}$  be the round index when  $\widehat{x}^*$  was excluded from  $Y$  for the first time i.e.,  $\widehat{x}^* \in Y^{i-1}$  but  $\widehat{x}^* \notin Y^i$ . Consider the following two possible cases: (a) When  $\widehat{x}^* \in \overline{B}$ , then this means that for all  $u \in \mathcal{U}$ ,  $\underline{F}(\widehat{x}^*, u) \not\subseteq Y^{i-1}$  (all states in  $\widehat{x}^*$  leave  $Y^{i-1}$  with positive probability) and  $(\overline{F}(\widehat{x}^*, u) \setminus \underline{F}(\widehat{x}^*, u)) \cap Y^{i-1} = \emptyset$  (no state in  $\widehat{x}^*$  can stay in  $Y^{i-1}$  with positive probability). (b) When  $\widehat{x}^* \notin \overline{B}$ , then this means that for all  $u \in \mathcal{U}$ , either  $\underline{F}(\widehat{x}^*, u) \not\subseteq Y^{i-1}$  (all states in  $\widehat{x}^*$  leave  $Y^{i-1}$  with positive probability), or from all states  $x \in \widehat{x}^*$  there does not exist any path to  $\overline{B}$ . Both (a) and (b) mean that from all the continuous states  $x \in \widehat{x}^*$ , the specification will be violated with positive probability after  $i$  time steps. This is a contradiction to our assumption that  $x^* \in W^*$ , since we know that from  $x^* \in \widehat{x}^*$  the specification can be satisfied for infinite duration with probability 1. Hence, it must hold that  $\widehat{x}^* \in \overline{W}$ .  $\square$

### 5.5 Over-approximation of the minimal a.s. losing region for reachability

Once we have a tight approximation of the a.s. winning region, we can compute a lower-bound of the satisfaction probability for the quantitative version of the  $\square \diamond B$  through Eqn. (5):

$$P_s^C(\text{C} \models \square \diamond B) \geq P_s^C(\text{C} \models \diamond W^*) \geq P_s^C(\text{C} \models \diamond \overline{W}), \quad \forall s \in \overline{W}^c, \quad (17)$$

where  $\overline{W}^c$  is the complement of  $\overline{W}$ . Efficient computation of maximal reachability requires computation of minimal a.s.

losing region, i.e., the set

$$L := \{s \in \mathcal{S} \mid \sup_{C \in \Pi_{\mathcal{S}}} P_s^C(\mathfrak{E} \models \diamond W) = 0\}.$$

We can compute the over-approximation of  $L$  using the following fixed point

$$\bar{L} := [\mu X. \text{Pre}_{\mathfrak{F}}(X) \cup W]^c.$$

**THEOREM 5.5.**  $\bar{L}$  is an over-approximation of  $L$ .

**PROOF.** We show the contra-positive, i.e.  $\bar{L}^c \subseteq L^c$ . Consider any abstract state  $\hat{x} \in \bar{L}^c$ . By construction, from all the continuous states  $x \in \hat{x}$ ,  $\underline{W}$  is reached with non-zero probability. Hence  $\hat{x} \notin L$ .  $\square$

Once we have an over-approximation of  $L$ , the stochastic kernel of the CMP becomes contractive over  $\mathcal{S} \setminus (\bar{L} \cup \underline{W})$  under mild continuity assumptions. Then approximate computational techniques in the literature on finite-horizon reachability can be utilised to find  $P_s^C(\mathfrak{E} \models \diamond \underline{W})$  with tunable error bounds [24, 25, 30].

## 6 COMPUTATION OF THE ABSTRACTION

**The dynamical system.** We consider sampled-time continuous state dynamical system with additive stochastic disturbance. The system is formalized using the tuple  $\Sigma = (\mathcal{S}, \mathcal{U}, f, t_w)$ , where  $\mathcal{S} \subset \mathbb{R}^n$  is the state space,  $\mathcal{U} \subset \mathbb{R}^m$  is the finite input space,  $f : \mathcal{S} \times \mathcal{U} \rightarrow \mathcal{S}$  is the nominal state transition function and  $t_w : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$  is the density function of the stochastic disturbance. The state update of  $\Sigma$  is given as:

$$s(k+1) = f(s(k), u(k)) + w(k), \quad k \in \mathbb{N}, \quad (18)$$

where  $s(k) \in \mathcal{S}$  and  $u(k) \in \mathcal{U}$  are the state and input at the  $k^{\text{th}}$  time instant,  $w(k)$  is a random variable with the density function  $t_w(\cdot)$ , and  $s(k+1)$  is the state at the  $(k+1)^{\text{st}}$  time instant.

The random variables  $\{w(k)\}_{k \in \mathbb{N}}$  are pairwise independent with the same density function  $t_w(\cdot)$ . We can write the system as a CMP  $\mathfrak{S} = (\mathcal{S}, \mathcal{U}, T_{\mathfrak{S}})$  with the stochastic kernel  $T_{\mathfrak{S}}(A \mid s, u) = \int_A t_w(s' - f(s, u)) ds'$  for all  $A \in \mathcal{B}(\mathcal{S})$ . For the construction of the abstraction we assume that  $t_w(\cdot)$  is piecewise continuous and  $f(\cdot, u)$  is continuous for all  $u \in \mathcal{U}$ .

**The abstraction.** We assume that  $\mathcal{S} = \mathcal{S}' \cup \{\phi\}$ , where  $\mathcal{S}'$  is a compact hyper-rectangle as the working region of the system and  $\phi$  is a sink state representing the complement of  $\mathcal{S}'$ . The disturbance also has a compact support  $D \subset \mathbb{R}^n$ . Suppose  $\hat{\mathcal{S}}'$  be a hyper-rectangular partition of  $\mathcal{S}'$ . The overall abstract state space is  $\hat{\mathcal{S}} = \hat{\mathcal{S}}' \cup \{\phi\}$ . Given an abstract

state  $\hat{x} = \llbracket a, b \rrbracket \in \hat{\mathcal{S}}'$  and a control input  $u \in \mathcal{U}$ , we denote the approximate *nominal* reachable set of  $\mathfrak{S}$  by  $\Phi(\hat{x}, u)$  s.t.

$$\Phi(\hat{x}, u) \supseteq \bigcup_{s \in \text{cl}(\hat{x})} f(s, u), \quad (19)$$

where  $\text{cl}(\hat{x})$  is the closure of the set  $\hat{x}$ . Note that  $\Phi(\hat{x}, u)$  can be computed using any reachability analysis method for deterministic dynamical systems [4, 5].

Define two functions  $S_1, S_2 : \hat{\mathcal{S}} \times \mathcal{U} \rightarrow 2^{\mathbb{R}^n}$  s.t.

$$S_1 : (\hat{x}, u) \mapsto \bar{D} \oplus \Phi(\hat{x}, u) \quad \text{and} \quad (20)$$

$$S_2 : (\hat{x}, u) \mapsto \underline{D} \ominus (-\Phi(\hat{x}, u)), \quad (21)$$

where  $\bar{D} \supseteq D$  is any over-approximation of the support of disturbance  $D$  and  $\underline{D} \subseteq D$  is any compact under-approximation of  $D$  over which  $t_w(\cdot)$  is strictly positive. The operators  $\oplus$  and  $\ominus$  are Minkowski sum and Minkowski difference of two sets and the minus sign in  $(-\Phi(\hat{x}, u))$  is applied to all elements.

**THEOREM 6.1.** Let  $\Sigma = (\mathcal{S}, \mathcal{U}, f, t_w)$  be a dynamical system and  $\mathfrak{S} = (\mathcal{S}, \mathcal{U}, T_{\mathfrak{S}})$  be the CMP induced by  $\Sigma$ . Define  $\mathcal{A} = (\hat{\mathcal{S}}, \mathcal{U}, \bar{T}, \underline{T})$  s.t.:

$$\begin{aligned} \bar{T}(\hat{x}, u) := \{ \hat{x}' \in \hat{\mathcal{S}} \mid (\hat{x}' \neq \phi) \Rightarrow (\hat{x}' \cap S_1(\hat{x}, u) \neq \emptyset) \wedge \\ (\hat{x}' = \phi) \Rightarrow (S_1(\hat{x}, u) \not\subseteq \mathcal{S}') \}, \end{aligned} \quad (22)$$

$$\begin{aligned} \underline{T}(\hat{x}, u) := \{ \hat{x}' \in \hat{\mathcal{S}} \mid (\hat{x}' \neq \phi) \Rightarrow (\lambda(\hat{x}' \cap S_2(\hat{x}, u)) > 0) \wedge \\ (\hat{x}' = \phi) \Rightarrow (\lambda(S_2(\hat{x}, u) \setminus \mathcal{S}') > 0), \end{aligned} \quad (23)$$

where  $\lambda(\cdot)$  gives the Lebesgue measure (volume) of a set. Then  $\mathcal{A}$  is an abstraction of  $\mathfrak{S}$ .

**PROOF.** We show that  $\bar{T}$  and  $\underline{T}$  satisfy the properties of  $\bar{F}$  and  $\underline{F}$  as formalized in Def. 4.2. For (7), consider any pair of abstract states  $\hat{x}, \hat{x}' \in \hat{\mathcal{S}}'$  and input  $u \in \mathcal{U}$  and there exists  $s \in \hat{x}$ ,  $T_{\mathfrak{S}}(\hat{x}' \mid s, u) > 0$ . We show that  $\hat{x}' \in \bar{T}(\hat{x}, u)$ :

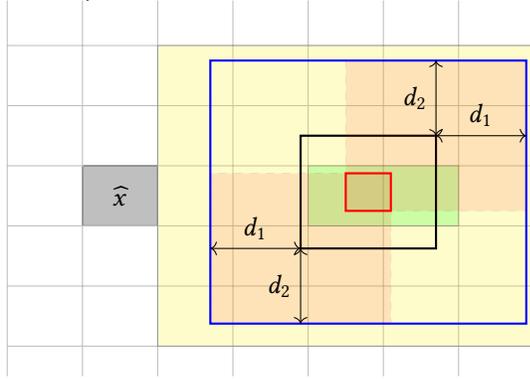
$$\begin{aligned} \int_{\hat{x}'} t_w(s' - f(s, u)) ds' > 0 &\Rightarrow \int_{\hat{x}' \ominus \{f(s, u)\}} t_w(w) dw > 0 \\ &\Rightarrow (\hat{x}' \ominus \{f(s, u)\}) \cap D \neq \emptyset \\ &\Rightarrow \exists w \in D, \exists s' \in \hat{x}' \text{ s.t. } w = s' - f(s, u) \\ &\Rightarrow \exists w \in D, \exists s' \in \hat{x}' \text{ s.t. } s' = f(s, u) + w. \end{aligned}$$

At the same time we know that  $f(s, u) \in \Phi(\hat{x}, u)$  since  $s \in \hat{x}$ . Then,

$$s' \in S_1(\hat{x}, u) \Rightarrow \hat{x}' \cap S_1(\hat{x}, u) \neq \emptyset \Rightarrow \hat{x}' \in \bar{T}(\hat{x}, u).$$

A similar reasoning holds for the case of  $\hat{x}' = \phi$ .

Now we show that  $\underline{T}$  satisfies the condition given in (8). Take  $\hat{x} \in \hat{\mathcal{S}}'$ , input  $u \in \mathcal{U}$ , and  $\hat{x}' \in \underline{T}(\hat{x}, u)$  s.t.  $\hat{x}' \neq \emptyset$ . Then  $\lambda(\hat{x}' \cap S_2(\hat{x}, u)) > 0$  according to (23). For any  $s' \in \hat{x}' \cap S_2(\hat{x}, u)$ ,



**Figure 3: Illustration of the steps of abstraction computation:** given the abstract state  $\hat{x}$  (filled with grey) and some control input  $u$ , first the nominal reachable set is over-approximated (black rectangle). Next, the two sets  $S_1$  (blue rectangle) and  $S_2$  (red rectangle) are computed. Finally, the images of the transition functions  $\bar{F}$  (filled with yellow) and  $\underline{F}$  (filled with green) are identified as the set of abstract states which intersect with  $S_1$  and  $S_2$  respectively.

we have

$$\begin{aligned} s' \in S_2(\hat{x}, u) &\Rightarrow \{s'\} \oplus (-\Phi(\hat{x}, u)) \subseteq \underline{D} \\ &\Rightarrow s' - f(s, u) \subseteq \underline{D} \quad \forall s \in \text{cl}(\hat{x}) \\ &\Rightarrow T_s(\hat{x}' | s, u) \geq \int_{\hat{x}' \cap S_2(\hat{x}, u)} t_w(s' - f(s, u)) ds' > 0 \end{aligned}$$

The right-hand side is strictly positive since the integrand is strictly positive and the domain of integration has a positive measure. It is also assumed that  $f$  is continuous and  $t_w$  piecewise continuous. Therefore, we have a positive function over the compact domain  $\text{cl}(\hat{x})$ , which will have a positive minimum:

$$\exists \varepsilon > 0 . \forall s \in \text{cl}(\hat{x}) . T_s(\hat{x}' | s, u) \geq \varepsilon \Rightarrow \hat{x}' \in \underline{F}(\hat{x}, u). \quad \square$$

The abstraction procedure can be summarized as follows: first compute the approximate nominal reachable set  $\Phi(\hat{x}, u)$  in (19), then take the Minkowski sum and difference for  $S_1, S_2$  in (20)-(21), and finally compute the transition relations (23)-(22). Fig. 3 illustrates the abstraction procedure for a 2-d system and when  $D$  is of the form  $[-d_1, d_1] \times [-d_2, d_2]$ .

### 6.1 Computation for mixed-monotone systems

We discuss in this subsection that the nominal reachable set  $\Phi(\hat{x}, u)$  can be computed without any numerical optimization if the function  $f(\cdot, u)$  is mixed-monotone for every  $u \in U$  and

the partition sets are hyper-rectangles. Recall the definition of mixed-monotonicity from the existing literature [4].

*Definition 6.2.* Let  $g : \mathcal{S} \rightarrow \mathcal{S}$  be a function, and  $\leq_{\mathcal{S}}$  be an order relation on  $\mathcal{S}$  induced by positive cones. The function  $g$  is called mixed-monotone w.r.t.  $\leq_{\mathcal{S}}$  (or simply mixed-monotone if  $\leq_{\mathcal{S}}$  is obvious from the context) if there exists a function  $h : \mathcal{S} \times \mathcal{S} \rightarrow \mathcal{S}$ —called the decomposition function—with the following properties:

- (1)  $\forall x \in \mathcal{S} . h(x, x) = g(x)$ ,
- (2)  $\forall x_1, x_2, y \in \mathcal{S} . (x_1 \leq x_2) \Rightarrow (h(x_1, y) \leq h(x_2, y))$ , and
- (3)  $\forall x, y_1, y_2 \in \mathcal{S} . (y_1 \leq y_2) \Rightarrow (h(x, y_2) \leq h(x, y_1))$ .

Intuitively, a mixed-monotone function can be decomposed into an increasing and a decreasing component. This phenomenon can be seen from the definition of the decomposition function. The following proposition [4, Thm. 1] shows a method for computing a fast over-approximation of the image of a rectangular set under a mixed-monotone function.

*PROPOSITION 6.3.* Let  $g$  be a mixed-monotone function with the decomposition function  $h$ , and  $\llbracket a, b \rrbracket \subseteq \mathcal{S}$  be any hyper-rectangle. The image of  $\llbracket a, b \rrbracket$  under  $g$  can be over-approximated as  $\llbracket h(a, b), h(b, a) \rrbracket$ .

For mixed-monotone  $f(\cdot, u)$  with decomposition function  $h_u$ ,  $\Phi(\hat{x}, u)$  can be computed using Prop. 6.3 as  $\Phi(\hat{x}, u) = \llbracket h_u(a, b), h_u(b, a) \rrbracket$  for  $\hat{x} = \llbracket a, b \rrbracket$ .

## 7 EXAMPLES

We implemented our algorithm symbolically on top of SCOTS [22]. All the experiments were performed on a computer equipped with 3GHz Intel Xeon E7-8857 v2 processor and 1.5 TB memory.

### 7.1 Perturbed Van der Pol Oscillator

We first demonstrate the applicability of our synthesis method on a verification problem which is to find the maximal set of initial states of a perturbed Van der Pol oscillator [18] such that a Büchi specification is satisfied almost surely including an additional safety objective  $\mathcal{S}' = [-0.5, 0.5] \times [-0.5, 0.5]$ . The state evolution of the oscillator is given by:

$$x_1(k+1) = x_1(k) + x_2(k)\tau + w_1(k)$$

$$x_2(k+1) = x_2(k) + (-x_1(k) + (1 - x_1(k)^2)x_2(k))\tau + w_2(k),$$

where the sampling time  $\tau$  is set to 0.1s and  $(w_1(k), w_2(k))$  is a pair of stochastic noise signals at time  $k$  drawn from a piecewise continuous density function with a compact support  $D = [-0.02, 0.02] \times [-0.02, 0.02]$ . Note that for the computation of the winning set, we do not need the actual density function as discussed in the previous section. The target set  $B$  for the Büchi specification  $\square \diamond B$  is given by a

polytope (green box) as shown in Fig. 4a. Our algorithm is able to compute the under and over-approximation of the set of a.s. winning region. In Fig. 4a, the under-approximation  $\underline{W}$  is shown in grey and  $\overline{W} \setminus \underline{W}$  is shown in blue.

It turns out that when the noise is treated as a worst case adversary, then there exists a deterministic value of the disturbance for which the trajectory of the oscillator *never* reaches the target for all the initial states inside the domain, thus violating the specification. So the winning region in this case is empty. A trajectory with a fixed deterministic perturbation that misses the target all the time is shown in black in Fig. 4a.

On the other hand, when the noise is treated as stochastic, then there are initial states from where the perturbed trajectory always eventually reaches the target polytope. Hence, the specification is satisfied. A trajectory with stochastic perturbation and the initial state  $I$  is shown in red in Fig. 4a. Table 1 summarizes the abstraction parameters used in our experiment, ratio of the computed volume of  $\underline{W}$  to the computed volume of  $\overline{W}$ , and the computation time for  $\underline{W}$  and  $\overline{W}$ .

## 7.2 Controlled perturbed vehicle

Our second example is a controller synthesis problem for a perturbed sampled-time version of the 3-d Dubin's vehicle [20]. We consider almost sure satisfaction of the Büchi specification while avoiding obstacles in the state space. Although we did not discuss avoidance of obstacle in the theory part, this can be easily handled by redefining the working region  $\mathcal{S}'$  of the system by excluding the obstacles. Thus given a hyper-rectangular working region  $[0, 2] \times [0, 3] \times [-\pi, \pi]$ , and an obstacle  $[0.8, 1.2] \times [1, 1.4] \times [-\pi, \pi]$  within that working region, we define  $\mathcal{S}' = [0, 2] \times [0, 3] \times [-\pi, \pi] \setminus [0.8, 1.2] \times [1, 1.4] \times [-\pi, \pi]$ . The system dynamics is given as: when  $u \neq 0$ ,

$$\begin{aligned} x_1(k+1) &= x_1(k) + V \sin(x_3(k) + u\tau)/u - V \sin(x_3(k))/u \\ &\quad + w_1(k) \\ x_2(k+1) &= x_2(k) - V \cos(x_3(k) + u\tau)/u + V \cos(x_3(k))/u \\ &\quad + w_2(k) \\ x_3(k+1) &= x_3(k) + u\tau + w_3(k), \end{aligned}$$

and when  $u = 0$ ,

$$\begin{aligned} x_1(k+1) &= x_1(k) + V \cos(x_3(k))\tau + w_1(k) \\ x_2(k+1) &= x_2(k) - V \sin(x_3(k))\tau + w_2(k) \\ x_3(k+1) &= x_3(k) + w_3(k), \end{aligned}$$

where the sampling time  $\tau = 1s$ , the constant forward velocity  $V = 0.1$  (maintained by e.g. a low level cruise control system), and  $(w_1(k), w_2(k), w_3(k))$  is a collection of stochastic noise samples drawn from a piecewise continuous density

function with the support  $D = [-0.06, 0.06] \times [-0.06, 0.06] \times [-0.06, 0.06]$ . It is due to this fixed velocity that the vehicle cannot stay stationary (or near stationary) after reaching the target, which makes the synthesis problem with Büchi specification much more challenging than the same with normal reachability specification.

When the noise is treated as a worst case adversary, the winning region is empty. However, when the noise is treated as stochastic, the approximate winning regions  $\underline{W}$  and  $\overline{W}$  are non-empty as shown in Fig. 4b. Fig. 4c shows the simulated trajectory of the vehicle using the synthesized controller. It was observed that even though the trajectory moves away from the target from time to time, either due to the external noise or due to the constant velocity, it always returns to the target eventually.

We performed the computation for 4 different levels of discretization granularity, and the results are summarized in Table 1. It can be observed (from the ratio  $\lambda(\underline{W})/\lambda(\overline{W})$ ) that the gap between  $\overline{W}$  and  $\underline{W}$  monotonically shrinks as we keep reducing the size of the abstract states, which means that the approximations  $\overline{W}$  and  $\underline{W}$  get progressively better with refinement of the state space partition. However, we pay the price of higher computation time as we make the discretization finer.

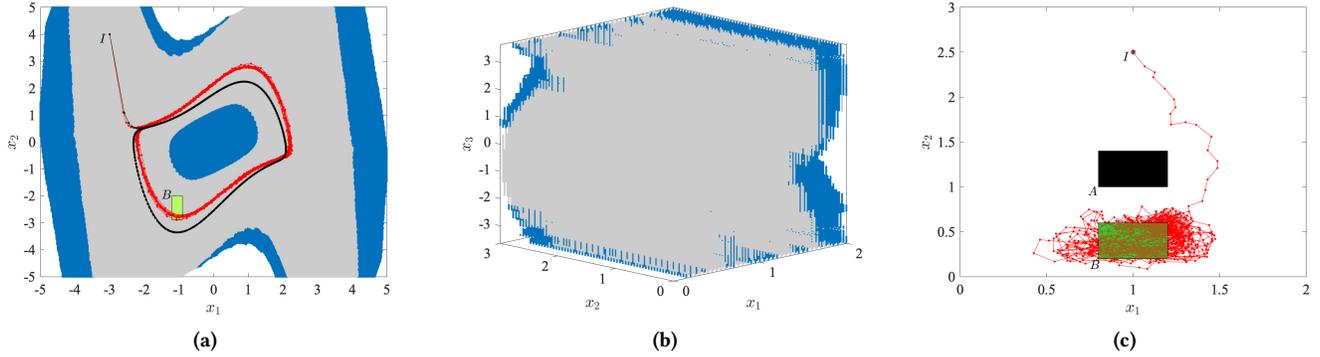
## 7.3 A note on computation time

The computation time for  $\underline{W}$  reported in Table 1 is based on a warm-start of Alg. 1 by replacing  $Y \leftarrow \widehat{\mathcal{S}}$  in Line 1 with  $Y \leftarrow \overline{W}$ . The intuition is that since it is known upfront that  $\underline{W} \subseteq \overline{W}$ , hence we do not need to consider the set  $\widehat{\mathcal{S}} \setminus \overline{W}$  in Alg. 1. In practice, the computation time for  $\underline{W}$  would be higher than the numbers reported in Table 1 had we started with  $\widehat{\mathcal{S}}$ .

In general, we observed that the computation of  $\underline{W}$  takes much longer than the computation of  $\overline{W}$ . Our hypothesis is that this is due to the properties of the operators defined in (10)-(13), and how they are used in the computation of  $\underline{W}$  and  $\overline{W}$ . For example, because  $\overline{F}(\widehat{x}, u)$  is a superset of  $\underline{F}(\widehat{x}, u)$  for all  $\widehat{x}, u$ , it can be shown that for a given  $Y, Z \subseteq \widehat{\mathcal{S}}$ ,  $\text{Apre}(\underline{W}, Z) \subseteq \text{Upre}(\underline{W}, Z)$ . Similarly  $\text{Cpre}_{\overline{F}}(Z) \subseteq \text{Upre}(\underline{W}, Z)$  when  $Y \supseteq Z$ . Moreover,  $\underline{B} \cap \text{Cpre}_{\overline{F}}(Y) \subseteq \underline{B} \cap (\text{Cpre}_{\underline{F}}(Y) \cup \text{Pre}_{\overline{F} \setminus \underline{F}}(Y))$ . Thus each iteration in the inner “ $\mu$ ” fixed point would add possibly fewer states in case of  $\underline{W}$  than in case of  $\overline{W}$ . Since ultimately the size of  $\overline{W}$  and  $\underline{W}$  are not very far apart, as shown in Col. 4 of Table 1, hence the iterations for  $\underline{W}$  would take many more number of steps than  $\overline{W}$ .

## 8 FUTURE WORK

We are working on three different extensions of this work. First, we plan to develop computation techniques for the



**Figure 4:** (a) The approximate winning region and trajectories of a Van der Pol oscillator. In figure:  $B$  (green box) is the target,  $\underline{W}$  is in grey,  $\overline{W} \setminus \underline{W}$  is in blue, and  $I$  is the initial state for simulation. The trajectory with stochastic perturbation is shown in red, and the trajectory with a fixed deterministic perturbation that misses the target all the time is shown in black. (b) The approximate a.s. winning domain for the Dubin’s vehicle example (for abstract state size set to  $0.05 \times 0.05 \times 0.05$ ):  $\underline{W}$  is in grey and  $\overline{W} \setminus \underline{W}$  is in blue. (c) Simulation of perturbed trajectory for the Dubin’s vehicle projected onto  $x_1, x_2$  plane. The black box  $A$  is the obstacle and the green box  $B$  is the target.

System	Size of abstract states	Volume of the worst-case winning region	$\lambda(\overline{W})/\lambda(\underline{W})$	Computation time		
				Abstraction	$\overline{W}$	$\underline{W}$
Van der pol oscillator	$0.02 \times 0.02$	empty	73.0%	< 1m	12m	416m
Dubin’s vehicle	$0.1 \times 0.1 \times 0.1$	empty	80.0%	1m	2m	11m
	$0.07 \times 0.07 \times 0.07$	empty	83.8%	4m	23m	35m
	$0.05 \times 0.05 \times 0.05$	empty	87.2%	18m	124m	140m
	$0.04 \times 0.04 \times 0.04$	empty	88.9%	37m	132m	128m

**Table 1:** Performance evaluation of our method on the Van der Pol oscillator and the Dubin’s vehicle. The 2nd column shows the size of each hyper-rectangular abstract state in the underlying uniform grid, the 3rd column shows the volume of the approximate winning domain when the noise is treated in the usual worst-case sense, the 4th column shows the ratio of the Lebesgue measure (volume) of  $\underline{W}$  to  $\overline{W}$ , and the 5th, 6th, and 7th columns show the computation times of different phases of our algorithm in minutes. Note that in our implementation, the computation of  $\underline{W}$  was warm started with already computed  $\overline{W}$ . Had  $\underline{W}$  been computed from scratch, the computation time for  $\underline{W}$  would be higher than what is shown in the last column.

qualitative winning regions for more general Rabin or parity conditions. Second, we are working on formulating conditions to guarantee convergence of the computations to the actual winning region when the discretization gets finer. Finally, we plan to improve the scalability of the approach using multi-resolution abstractions.

## REFERENCES

- [1] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [2] D. Bertsekas and S. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, 1996.
- [3] A. Bianco and L. de Alfaro. Model checking of probabilistic and non-deterministic systems. In P. S. Thiagarajan, editor, *Foundations of Software Technology and Theoretical Computer Science*, pages 499–513, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.
- [4] S. Coogan and M. Arcak. Efficient finite abstraction of mixed monotone systems. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, pages 58–67. ACM, 2015.
- [5] T. Dang and R. Testylier. Reachability analysis for polynomial dynamical systems using the Bernstein expansion. *Reliable Computing*, 17(2):128–152, 2012.
- [6] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Department of Computer Science, Stanford University, 1998.
- [7] L. de Alfaro and T. A. Henzinger. Concurrent omega-regular games. In *Proceedings Fifteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No. 99CB36332)*, pages 141–154. IEEE, 2000.
- [8] M. Dutreix and S. Coogan. Specification-guided verification and abstraction refinement of mixed-monotone stochastic systems, 2019.

- [9] S. Haesaert and S. Soudjani. Robust dynamic programming for temporal logic control of stochastic systems. *CoRR*, abs/1811.11445, 2018.
- [10] S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent program. *ACM Trans. Program. Lang. Syst.*, 5(3):356–380, July 1983.
- [11] O. Hernández-Lerma and J. B. Lasserre. *Discrete-time Markov control processes*, volume 30 of *Applications of Mathematics*. Springer, 1996.
- [12] P. Jagtap, S. Soudjani, and M. Zamani. Formal synthesis of stochastic systems via control barrier certificates, 2019.
- [13] N. Kariotoglou, M. Kamgarpour, T. H. Summers, and J. Lygeros. The linear programming approach to reach-avoid problems for Markov decision processes. *J. Artif. Int. Res.*, 60(1):263–285, Sept. 2017.
- [14] D. Kozen. Results on the propositional  $\mu$ -calculus. *Theoretical Computer Science*, 27(3):333 – 354, 1983. Special Issue Ninth International Colloquium on Automata, Languages and Programming (ICALP) Aarhus, Summer 1982.
- [15] M. Lahijanian, S. B. Andersson, and C. Belta. Formal verification and synthesis for discrete-time stochastic systems. *IEEE Transactions on Automatic Control*, 60(8):2031–2045, Aug 2015.
- [16] A. Lavaei, S. Soudjani, and M. Zamani. Compositional construction of infinite abstractions for networks of stochastic control systems. *Automatica*, 107:125 – 137, 2019.
- [17] K. Lesser and M. Oishi. Approximate safety verification and control of partially observable stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 62(1):81–96, Jan 2017.
- [18] H. Leung. Stochastic transient of a noisy van der pol oscillator. *Physica A: Statistical Mechanics and its Applications*, 221(1):340 – 347, 1995.
- [19] P. Nilsson, N. Ozay, and J. Liu. Augmented finite transition systems as abstractions for control synthesis. *Discrete Event Dynamic Systems*, 27(2):301–340, 2017.
- [20] R. Pepy, A. Lambert, and H. Mounier. Path planning using a dynamic vehicle model. In *2006 2nd International Conference on Information Communication Technologies*, volume 1, pages 781–786, April 2006.
- [21] G. Reissig, A. Weber, and M. Runger. Feedback refinement relations for the synthesis of symbolic controllers. *TAC*, 62(4):1781–1796, 2017.
- [22] M. Runger and M. Zamani. SCOTS: A tool for the synthesis of symbolic controllers. In *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*, HSCC '16, pages 99–104, New York, NY, USA, 2016. ACM.
- [23] S. Soudjani. *Formal Abstractions for Automated Verification and Synthesis of Stochastic Systems*. PhD thesis, Technische Universiteit Delft, The Netherlands, 2014.
- [24] S. Soudjani and A. Abate. Higher-Order Approximations for Verification of Stochastic Hybrid Systems. In S. Chakraborty and M. Mukund, editors, *Automated Technology for Verification and Analysis*, volume 7561 of *Lecture Notes in Computer Science*, pages 416–434. Springer Verlag, Berlin Heidelberg, 2012.
- [25] S. Soudjani and A. Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.
- [26] S. Soudjani, A. Abate, and R. Majumdar. Dynamic Bayesian networks for formal verification of structured stochastic processes. *Acta Informatica*, 54(2):217–242, Mar 2017.
- [27] P. Tabuada. *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, 2009.
- [28] I. Tkachev and A. Abate. Regularization of Bellman equations for infinite-horizon probabilistic properties. In *Proceedings of the 15th ACM international conference on Hybrid Systems: computation and control*, pages 227–236, Beijing, PRC, April 2012.
- [29] I. Tkachev and A. Abate. Characterization and computation of infinite-horizon specifications over markov processes. *Theoretical Computer Science*, 515:1–18, 2014.

- [30] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate. Quantitative model-checking of controlled discrete-time Markov processes. *Information and Computation*, 253:1 – 35, 2017.
- [31] A. P. Vinod and M. M. K. Oishi. Scalable underapproximative verification of stochastic LTI systems using convexity and compactness. In *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (Part of CPS Week)*, HSCC '18, pages 1–10, New York, NY, USA, 2018. ACM.

## 9 APPENDIX

### 9.1 Proof of Prop. 3.2

Following the steps utilized in [30, Theorem 7], we have that  $1 - f^*(s)$  is lower semi-analytic. Then  $\{s \in \mathcal{S} \mid 1 - f^*(s) < c\}$  is an analytic subset of  $\mathcal{S}$  for all  $c \in \mathbb{R}$ . Take a positive sequence  $\{c_n \rightarrow 0\}$ . The set  $\bigcap_n \{s \in \mathcal{S} \mid f^*(s) > 1 - c_n\} = \{s \in \mathcal{S} \mid f^*(s) = 1\}$  is also analytic. Every analytic set is universally measurable.

### 9.2 Proof of Thm. 3.3

PROOF OF THM. 3.3. We already know that  $P_s^C(\mathfrak{S} \models \square \diamond B) = 1$  for all  $s \in \text{WinDom}(\mathfrak{S}, C)$  by definition of the winning set. Take any  $s \notin W := \text{WinDom}(\mathfrak{S}, C)$ . We make the event conditional on  $\tau$  which is the first time the path hits a state in  $W$ . Then we have

$$\begin{aligned} P_s^C(\mathfrak{S} \models \square \diamond B) &= \mathbb{E}_s^C [P_s^C(\mathfrak{S} \models \square \diamond B \mid s_1, s_2, \dots, s_n, \tau = n)] \\ &= \sum_{n=0}^{\infty} P_s^C(s_1, s_2, \dots, s_{n-1} \in \mathcal{S} \setminus W, s_n \in W) \\ &\quad + P_s^C(\mathfrak{S} \models \square \diamond B \text{ and } \mathfrak{S} \models \square \mathcal{S} \setminus W). \end{aligned}$$

The sum is the reachability probability and the last term is always non-negative.  $\square$

### 9.3 Properties of the winning region

PROPOSITION 9.1. *For any control policy  $C$ , The set  $W := \text{WinDom}(\mathfrak{S}, C)$  is an absorbing set, i.e., the paths starting from this set will stay in the set a.s.:*

$$P_s^C(w[k+1] \in W \mid w[k]) = 1,$$

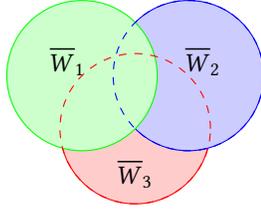
for all  $w[k] \in W$ .

PROOF. For any  $s \in W$ , we have

$$\begin{aligned} P_s^C(\mathfrak{S} \models \square \diamond B) &= \int_{\mathcal{S}} P_{s_1}^C(\mathfrak{S} \models \square \diamond B) T_s(ds_1 \mid s, C(s)) \\ &= \int_W T_s(ds_1 \mid s, C(s)) + \int_{\mathcal{S} \setminus W} P_{s_1}^C(\mathfrak{S} \models \square \diamond B) T_s(ds_1 \mid s, C(s)). \end{aligned}$$

This means

$$\begin{aligned} \int_{\mathcal{S} \setminus W} (1 - P_{s_1}^C(\mathfrak{S} \models \square \diamond B)) T_s(ds_1 \mid s, C(s)) &= 0, \\ P_s^C \left[ (1 - P_{s_1}^C(\mathfrak{S} \models \square \diamond B)) \mathbf{1}_{\mathcal{S} \setminus W}(s_1) \geq \epsilon \right] &\leq \frac{0}{\epsilon} = 0, \end{aligned}$$



**Figure 5: Illustration of construction of  $\bar{W}_1$  (green filled part),  $\bar{W}_2$  (blue filled part), and  $\bar{W}_3$  (red filled part) from  $\text{WinDom}(\mathfrak{S}, C_1)$  (green circle),  $\text{WinDom}(\mathfrak{S}, C_2)$  (blue circle), and  $\text{WinDom}(\mathfrak{S}, C_3)$  (red circle).**

where the last inequality is a consequence of Markov's inequality for non-negative random variables. By taking the union over a monotone positive sequence  $\{\epsilon_n \rightarrow 0\}$ , we get

$$\begin{aligned} P_s^C [(1 - P_{s_1}^C(\mathfrak{S} \models \square \diamond B)) \mathbf{1}_{S \setminus W}(s_1) > 0] &= 0, \\ P_s^C [s_1 \in S \setminus W \text{ and } P_{s_1}^C(\mathfrak{S} \models \square \diamond B) < 1] &= 0, \\ P_s^C [s_1 \in S \setminus W] &= 0. \end{aligned}$$

□

**PROPOSITION 9.2.** *Given a countable sequence of stationary policies  $\{C_1, C_2, \dots\}$  for the system  $\mathfrak{S}$  with winning regions  $\{\text{WinDom}(\mathfrak{S}, C_n), n = 1, 2, \dots\}$ , there is a controller  $C$  with winning region  $\text{WinDom}(\mathfrak{S}, C) = \bigcup_{n=1}^{\infty} \text{WinDom}(\mathfrak{S}, C_n)$ .*

**PROOF.** Define the sets  $\{\bar{W}_n, n = 1, 2, \dots\}$  inductively as  $\bar{W}_1 := \text{WinDom}(\mathfrak{S}, C_1)$  and  $\bar{W}_n := \text{WinDom}(\mathfrak{S}, C_n) \setminus \bigcup_{i=1}^{n-1} \bar{W}_i$  for all  $n \in \{2, 3, \dots\}$ . This construction is illustrated in Fig. 5. Also, define the new stationary policy:

$$C(s) := \begin{cases} C_1(s) & \text{if } s \in \bar{W}_1 \\ C_2(s) & \text{if } s \in \bar{W}_2 \\ \vdots & \end{cases} \quad (24)$$

for all non-empty sets  $\bar{W}_n$ . It is easy to show that the sets  $\{\bar{W}_i\}$  are non-intersecting and  $\bigcup_{i=1}^n \bar{W}_i = \bigcup_{i=1}^n \text{WinDom}(\mathfrak{S}, C_i)$ . Then for any initial state  $s \in \bigcup_{i=1}^{\infty} \text{WinDom}(\mathfrak{S}, C_i)$ , there is some  $n$  such that  $s \in \bar{W}_n$ . Note that all sets  $\text{WinDom}(\mathfrak{S}, C_i)$  are absorbing under their respective policy. The path starting from  $s$  with  $C_n$  either stay in  $\text{WinDom}(\mathfrak{S}, C_n)$  or will reach some  $\bar{W}_i$  with  $i < n$ . In the first case, the path satisfies the specification with probability one. The same argument can be applied a finite number of times until reaching the lowest index  $i = 1$ .

**measurability of  $C$ .** Note that the sets  $\{\bar{W}_1, \bar{W}_2, \dots\}$  are universally measurable and the policies  $\{C_1, C_2, \dots\}$  are universally measurable functions. We also have  $C^{-1}(A) =$

$\bigcup_{i=1}^{\infty} [C_i^{-1}(A) \cap \bar{W}_i]$ , which means  $C^{-1}(A)$  is universally measurable for any universally measurable  $A$ . Therefore,  $C$  is a universally measurable function.

□