

Quantitative Automata-based Controller Synthesis for Non-Autonomous Stochastic Hybrid Systems*

Ilya Tkachev
TU Delft
i.tkachev@tudelft.nl
Joost-Pieter Katoen
RWTH Aachen
katoen@cs.rwth-aachen.de

Alexandru Mereacre
University of Oxford
mereacre@cs.ox.ac.uk
Alessandro Abate
TU Delft
a.abate@tudelft.nl

ABSTRACT

This work deals with Markov processes that are defined over an uncountable state space (possibly hybrid) and embedding non-determinism in the shape of a control structure. The contribution looks at the problem of optimization, over the set of allowed controls, of probabilistic specifications defined by automata – in particular, the focus is on deterministic finite-state automata. This problem can be reformulated as an optimization of a probabilistic reachability property over a product process obtained from the model for the specification and the model of the system. Optimizing over automata-based specifications thus leads to maximal or minimal probabilistic reachability properties. For both setups, the contribution shows that these problems can be sufficiently tackled with history-independent Markov policies. This outcome has relevant computational repercussions: in particular, the work develops a discretization procedure leading into standard optimization problems over Markov decision processes. Such procedure is associated with exact error bounds and is experimentally tested on a case study.

Categories and Subject Descriptors

G.3 [Mathematics of Computing]: PROBABILITY AND STATISTICS—*Stochastic processes*

General Terms

Theory, verification

*This work is supported by the European Commission MoVeS project FP7-ICT-2009-5 257005, by the European Commission Marie Curie grant MANTRAS 249295, by the European Commission NoE HYCON2 FP7-ICT-2009-5 257462, and by the NWO VENI grant 016.103.020.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HSCC'13, April 8–11, 2013, Philadelphia, Pennsylvania, USA.
Copyright 2013 ACM 978-1-4503-1567-8/13/04 ...\$15.00.

Keywords

stochastic hybrid systems, probabilistic reachability, formal verification, stochastic optimal control, approximate abstractions.

1. INTRODUCTION

Stochastic Hybrid Systems (SHS) are a general and widely applicable mathematical framework involving the interaction of discrete, continuous, and probabilistic dynamics. Because of their generality, SHS have been applied over many areas, including telecommunication networks, manufacturing systems, transportation, and biological systems [8, 9].

SHS can be abstractly regarded as general Markov processes defined over an uncountable (in particular, hybrid) state space and embedded with non-determinism in the sense that they are dependent on a control structure [7]. The allowed control policies are functions of the history: in other words, we allow the control inputs to depend not only on the current state, but also on past trajectory realization and on past choices of control inputs. Markov policies are important special instances of these policies and depend solely on the current state.

SHS models are properly tagged by a labeling function that maps points in the state space to elements of a finite labeling set, which we refer to as an alphabet. SHS are structurally useful in the verification of linear time specifications, for example specifications expressed as deterministic finite-state automata (DFA) with labels. The work in [4] has shown that the verification of linear time specifications (in particular, DFA properties) can be performed by solving a probabilistic reachability problem over a new SHS, which is obtained by taking the cross product between the SHS and the DFA. Probabilistic reachability can then be practically assessed by computing its dual, namely probabilistic invariance [5].

This work generalizes the results in [4] to the case of control-dependent SHS and of DFA specifications. Since the verification of DFA specifications boils down to solving a probabilistic reachability problem, we consider this setup both in its finite- and infinite-horizon formulations. As the models under study are control dependent, we look into the possible maximal and minimal probabilistic reachability formulations: these can be studied as limits of dynamic programming recursions or as solutions of integral equations [7]. Since the cost functions expressing the dynamic programming scheme for probabilistic reachability take a multiplicative form [5], the analysis of their properties is in general difficult. This has led several works in the

literature [18, 10] to try formulating the reach-avoid problem (a generalization of the reachability problem) in terms of additive costs, where the dynamic programming theory is mature [7, 13]. To the best of our knowledge, inspired by the work in [11] the present work provides such a reduction explicitly for the first time. This as a consequence leads to several important results: first of all, we show that Markov policies are sufficient for the optimal probabilistic reachability; additionally, such a technique allows obtaining Bellman recursion and Bellman fixpoint equations for the reachability probability in the most general case, whereas the results in the literature were either focused on Markov policies exclusively [5, 18] or required structural assumptions on the model [10].

Sufficiency of Markov policies has important computational implications, since optimal controls can be synthesized according to the current value of trajectories, rather than based on the entire past history. Further along this computational line, the work introduces a discretization scheme that reduces the original (uncountable) setup to a finite optimization problem over Markov decision processes. Provided some continuity assumptions on the model are valid, the work shows that the discretization scheme can be associated with exact bounds on the introduced error. The obtained error bounds, inspired by [2], are functions of tunable parameters in the model and thus can be made, at the expense of more computations, arbitrarily small.

The article is structured as follows. Section 2 introduces the model syntax and semantics, as well as the class of linear temporal specifications of interest. Section 3 discusses the problem statement (probabilistic reachability) and its alternative formulation via additive cost functions, and derives the main theoretical results in this work: it shows in particular the sufficiency of Markov policies, and elaborates the minimal and maximal optimization problems over finite and infinite horizons. Section 4 puts forward a discretization scheme for the computation of the quantities in Section 3, with an exact quantification of the introduced errors. Finally, Section 5 presents the experimental outcomes over a case study and Section 6 concludes the paper. Due to space constraints, the proofs of the statements are omitted from this manuscript.

2. PRELIMINARIES

2.1 Notations and model definition

In this section we give a brief recap on Borel spaces and related stochastic kernels. It is common in the literature on the theory of controlled discrete-time Markov processes (cdt-MP) to assume that both the state space and the control space are endowed with a certain topological structure [7, 13]. A topological space X is called a Borel space if it is homeomorphic to a Borel subset of a Polish space¹. Examples of a Borel space are the Euclidean spaces \mathbb{R}^n , its Borel subsets endowed with a subspace topology [16], as well as hybrid spaces [5]. Any Borel space X is assumed to be endowed with a Borel σ -algebra, which is denoted by $\mathcal{B}(X)$. We say that a map $f : X \rightarrow Y$ is measurable whenever it is Borel measurable.

Given two Borel spaces X, Y the stochastic kernel on X given Y is the map $P : Y \times \mathcal{B}(X) \rightarrow [0, 1]$ such that $P(\cdot|y)$ is a probability measure on X for any point $y \in Y$, and such that $P(B|\cdot)$ is a measurable function on Y for any set $B \in \mathcal{B}(X)$. Stochas-

¹A Polish space is a topological space which is separable and completely metrizable.

tic kernels provide a natural generalization of update laws for deterministic systems, as we further show below.

We adopt the notation from [13] and consider a tuple $\mathfrak{D} = (X, U, \{U(x)\}_{x \in X}, \mathbb{T})$, where X is a Borel space, referred to as the state space of the model, and U is a Borel space to be thought of as the set of controls. Furthermore, $\{U(x)\}_{x \in X}$ is a family of non-empty measurable subsets of U with the property that

$$\mathbb{K} := \{(x, u) : x \in X, u \in U(x)\}$$

is measurable in $X \times U$. Intuitively, $U(x)$ is the set of controls that are feasible at state x . Finally, \mathbb{T} is a stochastic kernel on X given \mathbb{K} : note that \mathbb{K} is a measurable subset of a Borel space $X \times U$, hence it is itself a Borel space. In order to assure that the set of control policies is not empty we require \mathbb{K} to contain the graph of a measurable function [13, Assumption 2.2.2]. In other words, we assume that there exists a measurable map $k : X \rightarrow U$ such that $k(x) \in U(x)$ for any $x \in X$.

DEFINITION 1 (cdt-MP). We call any tuple

$$\mathfrak{D} = (X, U, \{U(x)\}_{x \in X}, \mathbb{T})$$

that satisfies the assumptions above a cdt-MP.

The following notation is used throughout the paper. We denote the set of positive integers by \mathbb{N} and the set of non-negative integers by $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. Furthermore, we denote $\bar{m}, \bar{n} := \{m, m+1, \dots, n-1, n\}$ for $m, n \in \mathbb{N}_0$ such that $m < n$; \mathbb{R} stands for the set of reals.

For any set X we denote by $X^{\mathbb{N}_0}$ the Cartesian product of a countable number of copies of X , i.e. $X^{\mathbb{N}_0} = \prod_{k=0}^{\infty} X$.

2.2 Model semantics

The semantics of a cdt-MP is characterized by its *paths* or *executions*, which reflect both the history of previous states of the system and of implemented control actions. Paths (often thought as infinite paths) are used to measure the performance of the system, which is done here via model checking methods. Also, a path up to a time epoch n can be used to derive the control action on the next step: these are finite paths that we also call *histories* as in [13, Section 2.2]. Let us finally mention that for technical reasons, when dealing with uncountable state spaces, it is usual to take into consideration also non-admissible paths, i.e. those containing non-feasible controls.

DEFINITION 2 (HISTORY). Given a cdt-MP \mathfrak{D} and a number $n \in \mathbb{N}_0$, an n -history is a finite sequence

$$h_n = (x_0, u_0, \dots, x_{n-1}, u_{n-1}, x_n), \quad (1)$$

where $x_i \in X$ are state coordinates and $u_i \in U$ are control coordinates of the *history*. An n -history h_n is called admissible if $u_i \in U(x_i)$, $i \in \bar{0}, n-1$. The space of all n -histories is denoted by \tilde{H}_n , and its subspace of admissible n -histories is denoted by H_n :

$$\tilde{H}_n = (X \times U)^n \times X, \quad H_n = \mathbb{K}^n \times X.$$

Further, we denote projections by $h_n[i] := x_i$ and $h_n(i) := u_i$.

DEFINITION 3 (PATH). An infinite path of a cdt-MP \mathfrak{D} is

$$\omega = (x_0, u_0, x_1, u_1, \dots), \quad (2)$$

where $x_i \in X$ and $u_i \in U$ for all $i \in \mathbb{N}_0$. As above, let us introduce projections $\omega[i] := x_i$ and $\omega(i) := u_i$.

The space of all infinite paths $\Omega = (X \times U)^{\mathbb{N}_0}$ together with its product σ -algebra \mathcal{F} is called a canonical sample space for

a *cdt*-MP \mathfrak{D} [13, Section 2.2]. An infinite path $\omega \in \Omega$ is called admissible if $u_n \in U(x_n)$ for any $n \in \mathbb{N}_0$. The space of all infinite admissible paths is denoted by $H_\infty = \mathbb{K}^\infty$ and is a subspace of Ω .

Given a path ω or a history h_n , we assume below that x_i and u_i are their state and control coordinates respectively, unless otherwise stated. In order to emphasize the control structure, for the path ω as in (2) we also write

$$\omega = x_0 \xrightarrow{u_0} x_1 \xrightarrow{u_1} x_2 \xrightarrow{u_2} \dots$$

Clearly, for any infinite path $\omega \in \Omega$ its n -prefix (ending in a state) ω_n is an n -history. However, for notational reasons we aim on denoting finite histories through h_n and paths through ω exclusively, since they usually serve for different purposes. We are now ready to introduce the notion of control policy.

DEFINITION 4 (POLICY). A policy is a sequence $\pi = (\pi_n)_{n \in \mathbb{N}_0}$ of universally measurable stochastic kernels π_n [7, Chapter 8.1], each defined on the control set U given H_n and such that

$$\pi_n(U(x_n)|h_n) = 1 \quad (3)$$

for all $h_n \in H_n$, $n \in \mathbb{N}_0$. The set of all policies is denoted by Π .

Equation (2) shows that all policies are “admissible” by definition. More precisely, the class of non-admissible policies is of 0 probability: given a policy $\pi \in \Pi$ and an admissible n -history $h_n \in H_n$, the distribution of the next control action u_n given by $\pi(\cdot|h_n)$ is supported on $U(x_n)$. Among the class of all possible policies, we are especially interested in those with a simple structure in that they depend only on the current state, rather than on the whole history.

DEFINITION 5 (MARKOV POLICY). A policy $\pi \in \Pi$ is called a Markov policy if for any $n \in \mathbb{N}_0$ it holds that $\pi_n(\cdot|h_n) = \pi_n(\cdot|x_n)$, i.e. π_n depends on the history h_n only through the current state x_n . The class of all Markov policies is denoted by $\Pi_M \subset \Pi$.

Although a Markov policy is not history-dependent, it is not necessary stationary, i.e. it may depend on the time variable. This fact highlights the difference between Markov policies and memoryless schedulers as defined in [6, Section 10.6]. More precisely, a policy $\pi \in \Pi$ is a memoryless scheduler if and only if it is Markov and stationary, i.e. $\pi_n = \pi_0$ for all $n \in \mathbb{N}_0$.

A central measurability result is that given a *cdt*-MP \mathfrak{D} and a policy $\pi \in \Pi$, it is always possible to construct a suitable probability measure over the set of paths. Due to technical reasons, such a measure is constructed on the space (Ω, \mathcal{F}) which also contains non-admissible paths, but it is supported on the space H_∞ of admissible paths. More precisely, let a *cdt*-MP \mathfrak{D} , a policy π and a probability measure α on X be given – the latter is referred to be the initial probability distribution of the *cdt*-MP. By the theorem of Ionescu Tulcea [13], there exists a unique probability measure P_α^π on the canonical sample space (Ω, \mathcal{F}) supported on H_∞ , i.e. $P_\alpha^\pi(H_\infty) = 1$ and such that

$$\begin{aligned} P_\alpha^\pi(x_0 \in B) &= \alpha(B), \\ P_\alpha^\pi(u_n \in C|h_n) &= \pi_n(C|h_n), \\ P_\alpha^\pi(x_{n+1} \in B|h_n, u_n) &= T(B|x_n, u_n) \end{aligned}$$

for all $B \in \mathcal{B}(X)$, $C \in \mathcal{B}(U)$ and $h_n \in H_n$, $n \in \mathbb{N}_0$. In the case when the initial distribution is supported on a single point, i.e. $\alpha(\{x\}) = 1$, we write P_x^π in place of P_α^π . Note that the last equation means that T is a transition kernel for the *cdt*-MP: whenever the current state is known and the control action is chosen, T gives a distribution for the next state.

2.3 Controlled Stochastic Hybrid Systems

Controlled discrete-time Stochastic Hybrid Systems (*cdt*-SHS) are a particular subclass of *cdt*-MP with a more explicit structure that distinguishes between continuous and discrete states of the system. The class of *cdt*-SHS provides rich modeling power and is applicable to various areas [9, 5]. It has been introduced in [2] and further studied in [5, 19]. Here we introduce it directly through the discussed framework: a *cdt*-SHS is a *cdt*-MP $\mathfrak{D} = (X, U, \{U(x)\}_{x \in X}, T)$ whose state space is $X = \bigcup_{q \in Q} \{q\} \times D_q$, where Q is a finite set of discrete *modes* and the measurable sets $D_q \in \mathcal{B}(\mathbb{R}^{n(q)})$ are the continuous components of the state space.

The control space U is often taken to be some Borel space. Furthermore, the transition kernel of *cdt*-SHS is often defined through its hybrid components [5] as:

$$T(\{q'\} \times dc' | (q, c), u) = \begin{cases} T_q(q' | (q, c), u) T_x(dc' | (q, c), u), & q' = q, \\ T_q(q' | (q, c), u) T_r(dc' | (q, c, q'), u), & q' \neq q, \end{cases}$$

for any $c \in D_q$ and $q \in Q$ and where T_q is a discrete probability law, whereas T_r, T_x are continuous (reset and transition) kernels. The semantical meaning of the conditional distributions T_q, T_r and T_x is given in [5]. The hybrid structure of the kernel further allows considering the control space U to be a product of the controls affecting the continuous dynamics and of those affecting the discrete dynamics as $U = U_q \times U_c$, where U_q, U_c are some Borel spaces [5].

Markov Decision Processes (MDP) are a subclass of *cdt*-SHS characterized by finite state and control spaces. Formally, U is finite and each continuous component D_q of a MDP is a singleton, which can be identified with the discrete component q itself. The theory of MDP is mature and allows for explicit solutions of many synthesis problems [6, Section 10.6].

2.4 Deterministic Finite State Automata

We are interested in linear temporal properties of trajectories of a given *cdt*-MP. For this purpose, in this section we introduce a model known as Deterministic Finite-state Automaton (DFA).

DEFINITION 6 (DFA). A DFA is a tuple $\mathcal{A} = (Q, q_0, \Sigma, F, t)$, where Q is a finite set of locations, $q_0 \in Q$ is the initial location, Σ is a finite set, $F \subseteq Q$ is a set of accept locations, and $t : Q \times \Sigma \rightarrow Q$ is a transition function.

We call the set Σ an *alphabet* and its elements $\sigma \in \Sigma$ *letters*. We denote by $\mathfrak{S} = \Sigma^{\mathbb{N}_0}$ (by $\mathfrak{S}_{<\infty}$) the collection of all infinite (finite) words over Σ . A finite word $w = (w[0], \dots, w[n]) \in \mathfrak{S}_{<\infty}$ is accepted by a DFA \mathcal{A} if there exists a finite run $z = (z[0], \dots, z[n+1]) \in Q^{n+2}$ such that $z[0] = q_0$, $z[i+1] = t(z[i], w[i])$ for all $0 \leq i \leq n$ and $z[n+1] \in F$. Although the verification of DFA is classically based on finite words, it is here more convenient to work with infinite words and define the corresponding accepting languages over infinite words. We say that an infinite word $w \in \mathfrak{S}$ is accepted by a DFA \mathcal{A} if there exists a finite prefix of w accepted by \mathcal{A} as a finite word. This is equivalent to the following statement: an infinite word $w \in \mathfrak{S}$ is accepted by \mathcal{A} if and only if there exists an infinite run $z \in Q^{\mathbb{N}_0}$ such that $z[0] = q_0$, $z[i+1] = t(z[i], w[i])$ for all $i \in \mathbb{N}_0$ and there exists $j \in \mathbb{N}_0$ such that $z[j] \in F$. Note that $w[i]$ (resp. $z[i]$) denotes the i -th letter (resp. state of the automaton) on w (resp. z). The accepted language of \mathcal{A} , denoted $\mathfrak{L}(\mathcal{A})$, is the set of all words accepted by \mathcal{A} . We are further interested in other kinds of accepting conditions over the DFA \mathcal{A} : we say that the word w is n -accepted by the DFA \mathcal{A} if there exists a

run $z \in Q^{\mathbb{N}_0}$ such that $z[0] = q_0$, $z[i+1] = t(z[i], w[i])$ for all $i \in \mathbb{N}_0$ and $z[j] \in F$ for some $j \leq n$. We denote the set of all n -accepted words by $\mathcal{L}_n(\mathcal{A})$. It is clear that $\mathcal{L}_\infty(\mathcal{A}) = \mathcal{L}(\mathcal{A})$ and furthermore that $\mathcal{L}_n(\mathcal{A}) \subseteq \mathcal{L}_{n+1}(\mathcal{A})$, where $n \in \mathbb{N}_0$ is arbitrary.

We use the DFA \mathcal{A} to specify properties of the cdt-MP as follows. Let $L : X \rightarrow \Sigma$ be a measurable function which we call the *labeling function* for a cdt-MP \mathcal{D} . To each state $x \in X$ it assigns the letter $L(x) \in \Sigma$. In the same fashion, each path $\omega = (x_0, u_0, x_1, u_1, x_2, u_2, \dots) \in \Omega$ induces the word $w \in \mathcal{G}$ given by $w = (L(x_0), L(x_1), L(x_2), \dots)$. We define the function L_Ω which maps paths onto words, i.e., $L_\Omega(\omega) = w$, as above. Using this function we can introduce the satisfaction relation between paths of \mathcal{D} and the DFA specification as follows:

$$\omega \models \mathcal{A} \iff L_\Omega(\omega) \in \mathcal{L}(\mathcal{A}). \quad (4)$$

As a result, given a policy $\pi \in \Pi$ we can define the probability that a path of \mathcal{D} satisfies \mathcal{A} , i.e. $P_\alpha^\pi(\omega \models \mathcal{A})$. It can be shown that under the assumption made on the measurability of $L : X \rightarrow \Sigma$ it holds that $\{\omega \in \Omega : \omega \models \mathcal{A}\} \in \mathcal{F}$ and hence such probability is well-defined. For any $n \in \mathbb{N}_0$ we also define

$$\omega \models_n \mathcal{A} \iff L_\Omega(\omega) \in \mathcal{L}_n(\mathcal{A}). \quad (5)$$

By now it should be clear why we prefer dealing with infinite words: the reason for this is that runs of the cdt-MP are always infinite, i.e. the cdt-MP “never stops”. At the same time, the input to the DFA \mathcal{A} is a word $L_\Omega(\omega)$ which is infinite as well.

The work in [4] studied model-checking of automata specifications against autonomous (i.e. uncontrolled) discrete-time stochastic models over uncountable state spaces. It was shown that the computation of the probability of satisfying a DFA can be restated in terms of a *probabilistic reachability* problem over the product between the original model and the DFA. In this paper we extend this result to the case of cdt-MP. As a result, we need to introduce the probabilistic reachability problem in general terms, and discuss its solution.

3. PROBABILISTIC REACHABILITY

3.1 Problem formulation

In this section we consider a basic and important problem where the probability of reaching a goal set is to be optimized. Let us consider a cdt-MP $\mathcal{D} = (X, U, \{U(x)\}, T)$ and some goal set $G \in \mathcal{B}(X)$. For any $n \in \mathbb{N}_0$ we define

$$\diamond^{\leq n} G = \{\omega \in \Omega : x_k \in G \text{ for some } 0 \leq k \leq n\} \quad (6)$$

and further $\diamond^{\leq \infty} G = \bigcup_{n=0}^{\infty} \diamond^{\leq n} G$. We call the event in $\diamond^{\leq n} G$ an *n-bounded* reachability if $n < \infty$, and an *unbounded* reachability if $n = \infty$. Since it holds that $\diamond^{\leq n} G = \bigcup_{k=0}^n \{x_k \in G\}$ for any $n \in \tilde{\mathbb{N}}_0 = \mathbb{N}_0 \cup \{\infty\}$, we obtain that $G \in \mathcal{B}(X)$ implies that $\diamond^{\leq n} G \in \mathcal{F}$ for all $n \in \tilde{\mathbb{N}}_0$. Due to this reason, the quantities $P_\alpha^\pi(\diamond^{\leq n} G)$ are well-defined for any initial distribution α and any control policy $\pi \in \Pi$. We further refer to these quantities as *reachability probabilities*.

We assume that the target set $G \in \mathcal{B}(X)$ is given and fixed. We are interested in the problem of reachability probability optimization, being either a maximization or a minimization over all possible policies $\pi \in \Pi$. For this purpose we restrict our attention to initial distributions supported at single points and define the following *value functions*: $V_n^\pi(x) := P_x^\pi(\diamond^{\leq n} G)$ for all $n \in \tilde{\mathbb{N}}_0$. For the problem of maximal reachability, the corresponding value functions are given by

$$V_n^*(x) := \sup_{\pi \in \Pi} V_n^\pi(x) \quad (7)$$

and for the problem of minimal reachability by

$$V_{*,n}(x) := \inf_{\pi \in \Pi} V_n^\pi(x). \quad (8)$$

A policy $\pi^* \in \Pi$ is called optimal for the problem (7) if it satisfies $V_n^{\pi^*}(x) = V_n^*(x)$. Similarly, a policy $\pi_* \in \Pi$ is called optimal for the problem (8) if $V_n^{\pi_*}(x) = V_{*,n}(x)$.

The reachability problem introduced above relates to other important problems in the analysis of probabilistic systems. First of all, it is a dual to the probabilistic safety (or invariance) problem, which was studied for cdt-SHS in [5]. We can define such problem as follows: for some $S \in \mathcal{B}(X)$ and any $n \in \mathbb{N}_0$,

$$\square^{\leq n} S = \{\omega \in \Omega : x_k \in S \text{ for all } 0 \leq k \leq n\} \quad (9)$$

and $\square^{\leq \infty} S = \bigcap_{n=0}^{\infty} \square^{\leq n} S$. The duality between the safety and the reachability is given by the identity $(\square^{\leq n} S)^c = \diamond^{\leq n} S^c$, which holds for any $n \in \tilde{\mathbb{N}}_0$. As a result, we obtain that the maximal and minimal safety problems can be successfully reformulated in terms of optimal reachability value functions, i.e. if $S = G^c$

$$\begin{aligned} \sup_{\pi \in \Pi} P_x^\pi(\square^{\leq n} S) &= 1 - V_{*,n}(x), \\ \inf_{\pi \in \Pi} P_x^\pi(\square^{\leq n} S) &= 1 - V_n^*(x), \end{aligned} \quad (10)$$

for any $n \in \tilde{\mathbb{N}}_0$. To find the maximal safety one has to look for the minimal reachability over the complement of the goal set.

Another problem related to probabilistic reachability is known as probabilistic reach-avoid [18]. Let us introduce this problem for two given sets $S, G \in \mathcal{B}(X)$ as follows. For $n \in \mathbb{N}_0$ we define

$$SU^{\leq n} G = \left\{ \omega \in \Omega : \begin{array}{l} x_k \in G \text{ for some } 0 \leq k \leq n \text{ and} \\ x_j \in S \text{ for all } 0 \leq j < k \end{array} \right\}$$

and $SU^{\leq \infty} G = \bigcup_{n=0}^{\infty} SU^{\leq n} G$. The measurability of the defined events is clear. The set S is referred to as the safe set (or the set of legal states) and as above the set G is referred to as the goal set. We further introduce the corresponding reach-avoid value function for $x \in X, n \in \tilde{\mathbb{N}}_0$, as

$$W_n^*(x) := \sup_{\pi \in \Pi} P_x^\pi(SU^{\leq n} G), \quad W_{*,n}(x) := \inf_{\pi \in \Pi} P_x^\pi(SU^{\leq n} G). \quad (11)$$

It is well-known that the reach-avoid problem is more general than the reachability one, since $\diamond^{\leq n} G = XU^{\leq n} G$ for any $G \in \mathcal{B}(X)$ and any $n \in \tilde{\mathbb{N}}_0$. However, the converse statement also holds true, as we are going to show now: the idea is to state the reach-avoid problem over the cdt-MP \mathcal{D} as a reachability problem over a new cdt-MP $\tilde{\mathcal{D}}$, where the avoid set $A := (S \cup G)^c$ is identified with an auxiliary single state with a loop.

More precisely, given a cdt-MP \mathcal{D} and two sets $S, G \in \mathcal{B}(X)$ we define $\tilde{\mathcal{D}} = (\tilde{X}, U, \{\tilde{U}_x\}_{x \in \tilde{X}}, \tilde{T})$ where $\tilde{X} = S \cup G \cup \{\psi\}$ for some auxiliary state $\psi \notin X$. We further define $\tilde{U}(x) = U(x)$ for $x \in S \cup G$ and $\tilde{U}(\psi) = \tilde{u}$ where \tilde{u} is an element of U . Finally,

$$\tilde{T}(B|x, u) = \begin{cases} T(B|x, u), & \text{if } B \in \mathcal{B}(S \cup G), x \in S \cup G \\ T(A|x, u), & \text{if } B = \{\psi\}, x \in S \cup G \\ 1, & \text{if } B = \{\psi\}, x = \psi. \end{cases}$$

Let us now define the corresponding optimal reachability value functions for the goal set G over the cdt-MP $\tilde{\mathcal{D}}$, which we denote by \tilde{V}_n^* for the maximal reachability and $\tilde{V}_{*,n}$ for the minimal one. The following result relates the reach-avoid value functions over \mathcal{D} to the reachability ones over $\tilde{\mathcal{D}}$.

PROPOSITION 1. For any $x \in A$ and any $n \in \mathbb{N}_0$ it holds that $W_{*,n}(x) = W_n^*(x) = 0$. Furthermore, for any $x \in S \cup G$ and $n \in \mathbb{N}_0$

$$W_n^*(x) = \check{V}_n^*(x), \quad W_{*,n}(x) = \check{V}_{*,n}(x).$$

It follows from Proposition 1 that results obtained for the optimization over the reachability probabilities can be directly applied to the case of reach-avoid. Due to this reason, we focus on the former problem and later show explicitly how to extend the obtained results from the reachability to the reach-avoid.

3.2 Formulation with an additive cost

The work in [5] considered the optimization of reachability probabilities over the class of cdt-MP where the cost functional took a multiplicative form. Focusing exclusively on Markov policies allowed obtaining DP recursions for value functions. In this work, however, we are interested in a wider class of policies, so the results in [5] are not directly applicable. In particular, an interesting question is the following: is it sufficient to consider only Markov policies in the optimization procedure? In order to answer this question, as well as to derive DP recursions, we are going to reformulate the original optimization problem via an additive cost functional, for which the theory of DP is rather rich [7, 13]. This approach is inspired by the one in [11].

Given a goal set $G \in \mathcal{B}(X)$ we consider a new cdt-MP

$$\hat{\mathcal{D}} := (\hat{X}, U, \{\hat{U}(x, y)\}_{(x,y) \in \hat{X}}, \hat{T})$$

with an augmented state space $\hat{X} = X \times Y$, where $Y = \{0, 1\}$. The states are of the form (x, y) with coordinates being $x \in X$, $y \in Y$. The control space U is the same and we further define $\hat{U}(x, y) := U(x)$. The dynamics of $\hat{\mathcal{D}}$ are given as follows:

$$\begin{cases} x_{n+1} & \sim \mathbb{T}(\cdot | x_n, u_n) \\ y_{n+1} & = 1_{G^c}(x_n) \cdot y_n, \end{cases}$$

hence the corresponding transition kernel \hat{T} is given by

$$\hat{T}(B \times \{y'\} | x, y, u) := \begin{cases} y \cdot 1_{G^c}(x) \mathbb{T}(B | x, u), & \text{if } y' = 1, \\ (1 - y \cdot 1_{G^c}(x)) \mathbb{T}(B | x, u), & \text{if } y' = 0. \end{cases}$$

We construct a space of policies $\hat{\Pi}$ and for each $\hat{\pi} \in \hat{\Pi}$ and initial distribution $\hat{\alpha}$ on \hat{X} , a probability space $(\Omega, \mathcal{F}, \hat{\mathbb{P}}_{\hat{\alpha}}^{\hat{\pi}})$ with the expectation $\hat{\mathbb{E}}_{\hat{\alpha}}^{\hat{\pi}}$. We denote by $\hat{\Pi}_M \subset \hat{\Pi}$ the corresponding class of Markov policies for $\hat{\mathcal{D}}$.

The additive cost structure consists of a cost $c : \hat{X} \rightarrow \{0, 1\}$ given by $c(x, y) := y \cdot 1_{G^c}(x)$ and a functional

$$J_n^{\hat{\pi}}(x, y) := \hat{\mathbb{E}}_{(x,y)}^{\hat{\pi}} \left[\sum_{k=0}^n c(x_k, y_k) \right].$$

In order to relate it to the original formulation defined over the cdt-MP \mathcal{D} , we first have to establish an explicit relationship between classes of strategies Π and $\hat{\Pi}$. Clearly, we can treat Π as a subset of $\hat{\Pi}$ as any policy $\pi \in \Pi$ for the cdt-MP \mathcal{D} serves also as a policy for the cdt-MP $\hat{\mathcal{D}}$. We let $\iota : \Pi \rightarrow \hat{\Pi}$ be the inclusion map. On the other hand, we define the projection map $\theta : \hat{\Pi} \rightarrow \Pi$ by

$$\theta_i(\pi)(du_i | x_0, u_0, \dots, x_i) := \hat{\pi}_i(du_i | x_0, y_0 = 1, u_0, \dots, x_i, y_i = 1)$$

The following result relates the two optimization problems.

THEOREM 1. For any $n \in \mathbb{N}_0$, $\pi \in \Pi$ and $\hat{\pi} \in \hat{\Pi}$, it holds that

$$J_n^{\hat{\pi}}(x, 1) = V_n^{\theta(\hat{\pi})}(x), \quad V_n^{\pi}(x) = J_n^{\iota(\pi)}(x, 1). \quad (12)$$

Theorem 1 has several important corollaries. First of all, it can be used to prove that Markov policies are sufficient for the original optimal reachability problem over a bounded time horizon, i.e. in case when $n < \infty$. At the same time, the optimal policy may depend on time and thus is not necessary stationary. Note that for a MDP, a special case of cdt-MP, this fact has been already known [6, Section 10.6]. More precisely:

COROLLARY 1. For any $n \in \mathbb{N}_0$ and $\pi \in \Pi$ there exists $\pi' \in \Pi_M$ such that $V_n^{\pi} = V_n^{\pi'}$, and as a consequence $V_n^*(x) := \sup_{\pi' \in \Pi_M} V_n^{\pi'}(x)$ and $V_{*,n}(x) := \inf_{\pi' \in \Pi_M} V_n^{\pi'}(x)$.

Moreover, it follows from Theorem 1 that one can do equivalently optimization of the additive cost functionals J to solve the original optimal reachability problem. Let us further define $J_n^*(x, y) := \sup_{\hat{\pi} \in \hat{\Pi}} J_n^{\hat{\pi}}(x, y)$ and $J_{*,n}(x, y) := \inf_{\hat{\pi} \in \hat{\Pi}} J_n^{\hat{\pi}}(x, y)$.

COROLLARY 2. For any $n \in \mathbb{N}_0$, $x \in X$, the following equalities hold true: $J_n^*(x, 0) = J_{*,n}(x, 0) = 0$ and

$$V_n^*(x) = J_n^*(x, 1), \quad V_{*,n}(x) = J_{*,n}(x, 1). \quad (13)$$

Finally, we can exploit DP recursions for the additive cost functionals J_n^* and $J_{*,n}$ to study DP recursion for the optimal reachability value functions. Let us introduce the following operators

$$\mathcal{J}^* f(x) := 1_G(x) + 1_{G^c}(x) \sup_{u \in U(x)} \int_X f(y) \mathbb{T}(dy | x, u),$$

$$\mathcal{J}_* f(x) := 1_G(x) + 1_{G^c}(x) \inf_{u \in U(x)} \int_X f(y) \mathbb{T}(dy | x, u),$$

which act on the space of bounded universally measurable functions. These operators can be used to compute optimal value functions recursively, as the following result states.

COROLLARY 3. For any $n \in \mathbb{N}_0$, the functions V_n^* , $V_{*,n}$ are universally measurable. Moreover, $V_0^* = V_{*,0} = 1_G$ and for any $n \in \mathbb{N}_0$

$$V_{n+1}^* = \mathcal{J}^* V_n^*, \quad V_{*,n+1} = \mathcal{J}_* V_{*,n}. \quad (14)$$

3.3 Infinite time horizon

In the previous section we have successfully restated the original reachability problem as a classical additive-cost problem, which made it possible to derive several important results. In particular, Corollary 3 allows one to compute finite-horizon optimal reachability functions recursively, which can be done using numerical methods based on approximate abstractions: we present them later in Section 4. With focus on the infinite time horizon, it is expected that the solution can be obtained as a limit of solutions of finite-horizon optimization problems, which is a fixpoint of an appropriate operator: either \mathcal{J}^* for the maximal value function, or \mathcal{J}_* for the minimal one. However, it is known from the literature that this is not necessarily true [7]. In particular, in our case the fixpoint characterization $V_{\infty}^* = \mathcal{J}^* V_{\infty}^*$ holds in general, whereas additional assumptions are needed to show that $V_{*,\infty} = \mathcal{J}_* V_{*,\infty}$.

We start with the following result, which describes the properties of reachability probabilities with a fixed policy.

LEMMA 1. For any $n \in \mathbb{N}_0$, $\pi \in \Pi$ and an initial distribution α ,

$$\mathbb{P}_{\alpha}^{\pi}(\diamond^{\leq n} G) \leq \mathbb{P}_{\alpha}^{\pi}(\diamond^{\leq n+1} G) \leq \mathbb{P}_{\alpha}^{\pi}(\diamond^{< \infty} G). \quad (15)$$

Moreover, it holds that

$$\mathbb{P}_{\alpha}^{\pi}(\diamond^{\infty} G) = \lim_n \mathbb{P}_{\alpha}^{\pi}(\diamond^{\leq n} G) = \sup_n \mathbb{P}_{\alpha}^{\pi}(\diamond^{\leq n} G). \quad (16)$$

The monotonicity result above is crucial for the proof of the fixpoint characterization for the maximal reachability over an infinite time horizon. When considering the limit of functions V_n^* as $n \rightarrow \infty$ the key step is to swap the order of $\lim_{n \rightarrow \infty}$ and the supremum over the control actions, which comes from \mathfrak{J}^* – this can be done as the limit of an increasing sequence is a supremum itself. This leads us to the following:

LEMMA 2. For any $x \in X$, $(V_n(x))_{n \in \mathbb{N}_0}$ is a non-decreasing sequence of real numbers and there exists a point-wise limit

$$V^*(x) := \lim_n V_n(x) = \sup_n V_n(x), \quad (17)$$

which is the least non-negative fixpoint of \mathfrak{J} , i.e. $V^* = \mathfrak{J}V^*$ and if there is another fixpoint $f \in \mathcal{B}(X)$ such that $f \geq 0$ then $f \geq V^*$.

THEOREM 2. The maximal reachability value function V_∞^* is the least non-negative fixpoint of the operator \mathfrak{J}^* .

REMARK 1. To our knowledge, this is the first result on the fixpoint characterization of the maximal reachability value function for cdt-MP. Although [10, Theorem (2.10)] provides a fixpoint characterization for a (more general) reach-avoid problem, one of the assumptions required there is that under any Markov policy, for any initial condition the set of legal states is left by the path of the cdt-MP with a probability 1 in some finite time. This clearly leads to the fact that $V_\infty^* \equiv 1$. Thus, with focus on the reachability problem [10, Theorem (2.10)] can be applied only for the case when the solution is known to be constant.

Let us now consider the minimal reachability problem: in this case the $\lim_{n \rightarrow \infty}$ has to be swapped with the infimum that comes from \mathfrak{J}_* , which cannot be done in general. We then tailor a technique in [13, Section 4] to our case. In order to establish the main result we need the following assumption:

ASSUMPTION 1. The kernel T is strongly continuous: $T(A|\cdot)$ is a continuous function on \mathbb{K} for any $A \in \mathcal{B}(X)$ [13, Appendix C].

THEOREM 3. Under Assumption 1, the minimal reachability function $V_{*,\infty}$ is the least non-negative fixpoint of \mathfrak{J}_* .

Let us discuss some properties of the optimal infinite-horizon reachability value functions and of the fixpoint equations

$$f = \mathfrak{J}^*f, \quad (18)$$

$$f = \mathfrak{J}_*f. \quad (19)$$

First of all, note that $V^*(x) = V_*(x) = 1$ for all $x \in G$, so in case $V^* \equiv 1$ or $V_* \equiv 1$ we say that the corresponding optimization problem has a trivial solution. This case refers to the reachability of the goal state G in some finite time with probability 1 starting from any initial condition. However, by substitution we find that the function $f \equiv 1$ solves both equations (18) and (19), regardless of the shape of the transition kernel T . Due to this reason, we are able to formulate the following result.

PROPOSITION 2. Equations (18) and (19) have unique solutions if and only if the corresponding optimization problems have trivial solutions.

Some examples when the solutions of the optimization problems are not trivial can be constructed using appropriate notions of absorbing sets over the cdt-MP.

DEFINITION 7 (STRONGLY AND WEAKLY ABSORBING SETS). Given the cdt-MP $\mathcal{D} = (X, U, \{U(x)\}_{x \in X}, T)$, the set $A \in \mathcal{B}(X)$ is called strongly absorbing if $T(A|x, u) = 1$ for all $u \in U(x)$ and $x \in A$.

The set B is called weakly absorbing if there exists a kernel μ on U given B such that $\mu(U(x)|x) = 1$ for all $x \in B$ and such that

$$\int_{U(x)} T(B|x, u)\mu(du|x) = 1, \quad \forall x \in B.$$

Let us briefly comment on the definition above. First of all, every strongly absorbing set A is a weakly absorbing set since the required kernel μ as per Definition 7 in such case can be chosen to be a deterministic one, obtained by the restriction of the map k (defined in Section 2.1) to the set A , i.e.

$$\mu(C|x) = 1_C(k(x))$$

for any $x \in A$ and $C \in \mathcal{B}(U)$. Furthermore, in the autonomous case when the control set U is identified with a singleton, the notion of weak and strong sets coincide with that of an absorbing set [15]. Intuitively, a strongly absorbing set remains absorbing under any action whereas for a weakly absorbing set there exists a control which makes such set absorbing.

PROPOSITION 3. If $A \subseteq G^c$ is a strongly (weakly) absorbing set, then $V^*(x) = 0$ ($V_*(x) = 0$) for all $x \in A$.

The latter result in particular shows that the presence of absorbing subsets on the complement of the goal state violates the uniqueness of the fixpoint equations, and leads to a non-trivial solution of the problem. This fact is already known for the case of autonomous systems [20, 21]. There, it has been shown that under some structural assumptions on the model and on the goal set, the presence of absorbing subsets is not only a sufficient condition for the lack of uniqueness, but is also necessary. In particular, let us further mention that in the second part of [4, Theorem 3], where the infinite-horizon (autonomous) reachability value function is considered, the result holds only under the assumption that the solution of the fixpoint equation is unique – in which case as we have shown the solution is trivial and equal to a constant function 1.

Finally, let us apply the derived results to the probabilistic reach-avoid problem using Proposition 1. We keep G as the goal set, and define $S \in \mathcal{B}(X)$ to be the set of legal states. As above, let W_n^* and $W_{*,n}$ be the optimal reach-avoid value functions as in (11) and define the following operators

$$\mathfrak{R}^*f(x) := 1_G(x) + 1_A(x) \sup_{u \in U(x)} \int_X f(y)T(dy|x, u),$$

$$\mathfrak{R}_*f(x) := 1_G(x) + 1_A(x) \inf_{u \in U(x)} \int_X f(y)T(dy|x, u),$$

where as above $A = (S \cup G)^c$ is the set to be avoided. The next result follows immediately from those we have obtained for reachability and from Proposition 1.

THEOREM 4. It holds that $W_{*,0} = W_0^* = 1_G$ and for any $n \in \mathbb{N}_0$

$$W_{n+1}^* = \mathfrak{R}^*W_n^*, \quad W_{*,n+1} = \mathfrak{R}_*W_{*,n}.$$

In particular, Markov policies are sufficient for optimization:

$$W_n^*(x) := \sup_{\pi \in \Pi_M} P_x^\pi(SU^{\leq n}G), \quad W_{*,n}(x) := \inf_{\pi \in \Pi_M} P_x^\pi(SU^{\leq n}G).$$

Moreover, function W_∞^* is the least non-negative fixpoint of the operator \mathfrak{R}^* and, under Assumption 1, function $W_{*,\infty}$ is the least non-negative fixpoint of the operator \mathfrak{R}_* .

Theorem 4 generalizes several results in the literature. First, it extends the work in [18] by considering all possible policies, rather than focusing on Markov policies exclusively. Second, it extends results obtained in [10] over the infinite-time horizon by considering both optimization problems, rather than the maximization one only. Furthermore, we are able to relax the assumptions in [10, Theorem (2.10)] and show that the fixpoint characterization for W_∞^* holds in a more general case.

3.4 Automata model checking

In section 2.4 we have introduced the following problem: given a cdt-MP \mathcal{D} , a policy $\pi \in \Pi$, an initial distribution α and a DFA \mathcal{A} , find the probability $P_\alpha^\pi(\omega \models \mathcal{A})$. We are further interested in maximizing and minimizing such a probability over all possible policies $\pi \in \Pi$. For this purpose, we are going to reduce this general problem over \mathcal{D} to the reachability problem studied above over another cdt-MP $\mathcal{D} \otimes \mathcal{A}$, which we refer to as a product of the cdt-MP \mathcal{D} and the automaton \mathcal{A} . This product is defined as follows:

DEFINITION 8 (PRODUCT BETWEEN cdt-MP AND DFA). *Given a cdt-MP $\mathcal{D} = (X, U, \{U(x)\}_{x \in X}, T)$, a finite alphabet Σ , a labeling function $L : X \rightarrow \Sigma$, and a DFA $\mathcal{A} = (Q, q_0, \Sigma, F, t)$, we define the product between \mathcal{D} and \mathcal{A} to be another cdt-MP denoted as $\mathcal{D} \otimes \mathcal{A} = (\tilde{X}, U, \{U(x)\}_{x \in \tilde{X}}, \tilde{T})$. Here $\tilde{X} = X \times Q$ and*

$$\tilde{T}(A \times \{q'\} | x, q, u) = 1_{t(q, L(x))}(q') \cdot T(A | x, u).$$

We want to show that $P_x^\pi(\omega \models \mathcal{A})$ can be related to the reachability probability over the cdt-MP $\mathcal{D} \otimes \mathcal{A}$ with a goal state $G := X \times F$, as it was shown to be the case when the state space is finite [12, Proposition 1]. In order to reformulate the automaton verification as a reachability problem, we are going to follow a procedure similar to the one in Section 3.2, where the reachability problem has been reformulated by an additive cost. We again construct a space of policies $\tilde{\Pi}$ for $\mathcal{D} \otimes \mathcal{A}$, and for each $\tilde{\pi} \in \tilde{\Pi}$ and initial distribution $\tilde{\alpha}$ on \tilde{X} , a probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P}_{\tilde{\alpha}}^{\tilde{\pi}})$. As in Section 3.2 we need to relate policies in Π to those in $\tilde{\Pi}$: again, we can use the fact that $\Pi \subset \tilde{\Pi}$ hence any policy $\pi \in \Pi$ over the original cdt-MP \mathcal{D} is also a policy over $\mathcal{D} \otimes \mathcal{A}$. By $\iota : \Pi \rightarrow \tilde{\Pi}$ we can hence denote the corresponding inclusion map. For the other direction, to any policy $\tilde{\pi} \in \tilde{\Pi}$ we can assign a policy $\theta(\tilde{\pi}) \in \Pi$ as follows

$$\theta_i(\tilde{\pi})(du_i | x_0, u_0, \dots, x_i) = \tilde{\pi}_i(du_i | x_0, z_0, u_0, \dots, x_i, z_i),$$

where $z_0 = q_0$ is the initial state of \mathcal{A} and $z_{j+1} = t(z_j, L(x_j))$ is defined recursively for $j \in \overline{0, i-1}$. The following technical lemma is necessary to state the main result.

LEMMA 3. *For any $n \in \mathbb{N}_0$, $x \in X$ and $\pi \in \Pi$ it holds that*

$$P_x^\pi(\omega \models_n \mathcal{A}) = \tilde{P}_{(x, q_0)}^{\iota(\pi)}(\diamond^{\leq n+1} G),$$

where $\infty + 1 := \infty$, and for any $\tilde{\pi} \in \tilde{\Pi}$ it holds that

$$\tilde{P}_{(x, q_0)}^{\tilde{\pi}}(\diamond^{\leq n+1} G) = P_x^{\theta(\tilde{\pi})}(\omega \models_n \mathcal{A}).$$

THEOREM 5. *For any $n \in \mathbb{N}_0$ and $x \in X$ it holds that*

$$\begin{aligned} \sup_{\pi \in \Pi} P_x^\pi(\omega \models_n \mathcal{A}) &= \tilde{V}_{n+1}^*(x), \\ \inf_{\pi \in \Pi} P_x^\pi(\omega \models_n \mathcal{A}) &= \tilde{V}_{*,n+1}(x), \end{aligned} \quad (20)$$

where $\infty + 1 := \infty$ and $\tilde{V}_n^*, \tilde{V}_{*,n}$ are the optimal reachability functions that are defined over the cdt-MP $\mathcal{D} \otimes \mathcal{A}$.

Notice that in the above theorem we need to consider only Markov policies as was shown in the previous section. However, such policies are Markov with respect to $\mathcal{D} \otimes \mathcal{A}$ and are not necessarily Markov with respect to \mathcal{D} . In terms of \mathcal{D} this means that the policy is dependent on the history through the state of the automaton. The actual computation of the optimal reachability value functions can be implemented by discretizing the set of states as well as the set of controls, as elaborated in Section 4.

4. APPROXIMATE ABSTRACTIONS

Since in general there is no hope that the iterations in (14) yield value functions in an explicit form, we introduce an abstraction procedure that leads to numerical methods for the computation of such functions. Moreover, we provide an explicit upper bound on the error caused by the abstraction. We present the results with focus on the maximal reachability problem, with the understanding that a similar procedure applies to the minimal reachability one too.

Let us consider some cdt-MP $\mathcal{D} = (X, U, \{U(x)\}_{x \in X}, T)$. Since X and U are Borel spaces they are metrizable topological spaces. Let ρ_X and ρ_U be some metrics on X and U respectively, which are consistent with the given topologies of the underlying spaces. We first introduce some technical considerations that are important for the abstraction over the control space. Let $U = \bigcup_{j=1}^M U_j$ be a measurable partition of U , and let $u_j \in U_j$ for $1 \leq j \leq M$ be arbitrary representative points. Define $\Delta := \max_{1 \leq j \leq M} \text{diam}_U(U_j)$ where the diameter of a subset of U for any $C \subseteq U$ is given by

$$\text{diam}_U(C) = \sup_{u', u'' \in C} \rho_U(u', u'').$$

LEMMA 4. *Let $g, \tilde{g} : U \rightarrow \mathbb{R}$ be two functions. Define two optimization problems: $g^* := \sup_{u \in U} g(u)$ and $\tilde{g}^* := \max_{1 \leq j \leq M} \tilde{g}(u_j)$. If g is Lipschitz continuous, i.e. if there exists $K > 0$ such that*

$$|g(u') - g(u'')| \leq K \cdot \rho_U(u', u'') \quad (21)$$

for all $u', u'' \in U$, then it holds that $|g^* - \tilde{g}^*| \leq K \cdot \Delta + \|g - \tilde{g}\|_U$.

Let us further proceed with the abstraction procedure and choose $G \in \mathcal{B}(X)$ to be a target set. Clearly, the solution of any reachability problem on G is trivial, so we only need to solve these problems on G^c . For this purpose we define $G^c = \bigcup_{i=1}^N G_i$ to be a measurable partition of the set G^c , and we choose representative points $x_i \in G_i$ for $1 \leq i \leq N$ in an arbitrary way. We further denote $\delta_i := \text{diam}_X(G_i)$.

We abstract the original cdt-MP \mathcal{D} as an MDP denoted by $\tilde{\mathcal{D}} = (\tilde{X}, \tilde{U}, \{\tilde{U}(x)\}_{x \in \tilde{X}}, \tilde{T})$. The finite state and control spaces are $\tilde{X} := \{i\}_{i=1}^N \cup \{\phi\}$ and $\tilde{U} := \{j\}_{j=1}^M$, where $\phi \notin X$ is a ‘‘sink’’ state that corresponds to the target set G of the original cdt-MP. To complete the definition of $\tilde{\mathcal{D}}$ we have to specify the map \tilde{U} and the kernel \tilde{T} . We choose $\tilde{U}(i) = \tilde{U}$ for any $i \in \tilde{X}$ and

$$\begin{cases} \tilde{T}(k | i, j) := T(G_k | x_i, u_j) & \text{for all } 1 \leq j \leq M, 1 \leq i, k \leq N \\ \tilde{T}(\phi | i, j) := T(G | x_i, u_j) & \text{for all } 1 \leq j \leq M, 1 \leq i \leq N \\ \tilde{T}(\phi | \phi, j) = 1 & \text{for all } 1 \leq j \leq M. \end{cases}$$

For $\tilde{\mathcal{D}}$, let us define \tilde{V}_n^* to be the n -horizon maximal reachability value function of the target state ϕ . Clearly, $\tilde{V}_0^* = 1_{\{\phi\}}$ and it follows from Corollary 3 that

$$\tilde{V}_{n+1}^*(i) = 1_{\{\phi\}}(i) + 1_{\{\phi\}^c}(i) \max_{1 \leq j \leq M} \sum_{k \in \tilde{X}} \tilde{V}_n^*(k) \tilde{T}(k | i, j).$$

Such value functions can be computed with numerically efficient methods [14]. Clearly, we expect that such functions serve as a good approximation for functions V_n^* defined for the original cdt-MP \mathfrak{D} . However, so far they are defined over a different state space, which makes it hard to compare them. Due to this reason, we choose $\hat{V}_n^* : X \rightarrow \mathbb{R}$ to be a piece-wise constant interpolation of V_n^* , defined in the following manner:

$$\hat{V}_n^*(x) = 1_G(x) + \sum_{i=1}^N 1_{G_i}(x) \tilde{V}_n^*(i).$$

Clearly, $\hat{V}_n^*(x) = 1$ for all $x \in G$. Moreover, for all other states x the following result holds true.

LEMMA 5. For any $1 \leq i \leq N$ and any $x \in G_i$ it holds that

$$\hat{V}_{n+1}^*(x) = \max_{1 \leq j \leq M} \int_X \hat{V}_n^*(y) \mathbb{T}(dy|x_i, u_j).$$

To ensure that the difference between V_n^* and \hat{V}_n^* can be made as small as needed by tuning parameters of the partition Δ and δ_i , we need the following assumption.

ASSUMPTION 2. The action space $U(x)$ does not depend on x , i.e. it holds that $U(x) = U$ for all $x \in X$. In addition, \mathbb{T} is an integral kernel, i.e. there exists a σ -finite measure μ on X and a jointly measurable function $t : X \times X \times U \rightarrow \mathbb{R}$ such that

$$\mathbb{T}(B|x, u) = \int_B t(y|x, u) \mu(dy), \quad (22)$$

for any $B \in \mathcal{B}(X)$, $x \in X$, $u \in U$. Moreover, there exist:

1. measurable functions $\lambda_i : X \rightarrow \mathbb{R}$ for $1 \leq i \leq N$, such that for all $x', x'' \in G_i$ and for all $y \in X$, $u \in U$:

$$|t(y|x', u) - t(y|x'', u)| \leq \lambda_i(y),$$

$$\text{where } \Lambda_i := \int_X \lambda_i(y) \mu(dy) < \infty;$$

2. measurable functions $\kappa_i : X \rightarrow \mathbb{R}$ for $1 \leq i \leq N$, such that for all $u', u'' \in U$ and for all $y \in X$, $x \in G_i$:

$$|t(y|x, u') - t(y|x, u'')| \leq \kappa_i(y),$$

$$\text{where } K_i := \int_X \kappa_i(y) \mu(dy) < \infty.$$

The following result provides upper-bounds on the difference between the value functions of the original model and those computed via the MDP abstraction.

THEOREM 6. Let V_n^* and \hat{V}_n^* be the functions defined above. Introduce $r := \sup_{x \in G^c, u \in U} \mathbb{T}(G^c|x, u)$ and let Assumption 2 hold true. If $r < 1$, then for all $n \geq 1$ it holds that

$$\|V_n^* - \hat{V}_n^*\| \leq \frac{1 - r^n}{1 - r} \cdot \max_{1 \leq i \leq N} (\Lambda_i \delta_i + K_i \Delta). \quad (23)$$

If $r = 1$, then for all $n \geq 0$ it holds that

$$\|V_n^* - \hat{V}_n^*\| \leq n \cdot \max_{1 \leq i \leq N} (\Lambda_i \delta_i + K_i \Delta). \quad (24)$$

Since the case of autonomous SHS can be regarded as a cdt-SHS with a control space being a singleton, $U = \{u\}$, it is worth commenting on the relation between Theorem 6 and approximation techniques known from the literature on autonomous SHS [3, 17]. First of all, in such a case Assumption 2 is a slight

generalization of assumptions on the uniform Lipschitz continuity of kernels in [3] and on the local Lipschitz continuity in [17]. In particular, the application of Theorem 6 under Assumption 2 to the autonomous case implies [17, Theorem 6] as a special case, where the functions λ_i are assumed to be piece-wise constant. In turn, [17, Theorem 6] is further known to be a generalization of [3, Theorem 1].

Let us mention that the structure of the bounds allows applying the adaptive gridding procedure developed for the autonomous SHS in [17] over the state space discretization, by computing the local errors Λ_i and choosing the discretization size δ_i accordingly. However, the error introduced by the partition of the control space depends on the global discretization parameter $\max_{1 \leq i \leq N} K_i \Delta$, so the adaptive gridding of the control space may not improve results versus those obtained by the uniform gridding of the control space.

5. CASE STUDY: ENERGY CONTROL

Inspired by [1], we consider a resource allocation problem for an energy network comprised of two subnetworks $i = 1, 2$. The energy provider for each subnetwork is given the choice of generating energy at capacity either by a polluting device (say, a coal plant), or alternatively via local renewables: accordingly, the (normalized) decision variables are $0 \leq u_p^i \leq 1$, $i = 1, 2$, denoting the production of polluting power (u_p^i) and of renewable power (u_r^i) within the i^{th} subnetwork. There is a constraint on the maximal power generation, both globally (over the total generated polluting power) and locally (whatever is not generated via coal can be obtained from the renewables u_r^i), so that

$$u_p^1 + u_p^2 \leq 2T, \quad u_p^i + u_r^i = T, \quad i = 1, 2. \quad (25)$$

Each of the two subnetworks sets forth a time-varying demand $D_i(k)$, which is not exactly known due to the intrinsic variability of power demand.

The state variables of the model are the time-varying energy levels of each subnetwork, which we denote by E_i for $i = 1, 2$. They are driven by the following dynamics:

$$E_i(k+1) = E_i(k) + P(k)u_p^i(k) + R_i(k)u_r^i(k) - D_i(k), \quad (26)$$

where P is the actual power generated by the coal plant (polluting device) and R_i is the generated power by local renewables. Furthermore, P and R_i are independent random variables, identically distributed in time, and such that $\mathbb{E}[P] = \mu_P > \mathbb{E}[R_i] = \mu_{R_i}$, whereas $\text{Var}[P] = \sigma_P^2 < \text{Var}[R_i] = \sigma_{R_i}^2$. The above relation among parameters is suggested by assuming that the coal plant is a stronger and more reliable (though less desirable) source of power. Variables R_1 and R_2 can be correlated due to the spatial adjacency of the two subnetworks and its effect on the production based on renewables. This, along with the presence of P and of the constraints in (25), couples the two dynamics. We select the demand variables D_i to be independent and identically distributed, and independent from P and R_i , $i = 1, 2$.

We consider two scenarios expressed via specifications, and proceed optimizing over. Let us introduce two additional constraints on the energy levels for the subnetworks, namely:

$$E_1 + E_2 \leq S, \quad (27)$$

$$E_i \geq M_i, \quad i = 1, 2, \quad (28)$$

where the threshold S denotes a (constant) limit due to storage, while the second inequality refers to (constant) minimal energy

requirements. We are looking at the following problem:

$$\sup_{\pi} \mathbb{P}_x^{\pi} \left\{ \square^{\leq N} (27) \wedge (28) \right\}, \quad (29)$$

which corresponds to the following safety specification: the energy levels have to be above the thresholds and the sum of them must not exceed the storage capacity. Another problem which is interesting to us is whether the energy levels can simultaneously exceed some given value F without ever falling below the zero level beforehand. This problem can be expressed as:

$$E_i \geq F_i, i = 1, 2, \quad (30)$$

$$E_i \geq 0, i = 1, 2, \quad (31)$$

and the related synthesis problem is

$$\sup_{\pi} \mathbb{P}_x^{\pi} \left\{ (31) \cup^{\leq N} (30) \right\}, \quad (32)$$

The specification in (29) denotes maximal probabilistic invariance (equivalently, minimal reachability), whereas (32) is a maximal reach-avoid property, which can be reformulated as a maximal probabilistic reachability problem as discussed in Sec. 3.

In the implementation we have used the following model parameters: $T = 1$, $P = 2$ (constant polluting power production), whereas $R_i \sim \mathcal{N}(0.5, 2)$ and $D_i \sim \mathcal{N}(1, 2)$. None of the variables P , R and D depend on the step size k . We have selected the parameters M_i for (29) to be smaller than F_i for (32): for the first synthesis problem we have picked $M_i = 5$ and $S = 30$, whereas for the second one we have picked $F_i = 25$. We have chosen a discretization step $\delta_s = 1$ for the state variables and $\delta_c = 0.05$ for the control variables. All the experiments have been run on a 2.83GHz 4 Core(TM)2 Quad CPU with 3.7Gb of memory. The total running time for the experiments has amounted to 14.37 min.

Fig. 1(a) displays the maximal probabilistic invariance obtained (at the initial time step) for the synthesis problem in (29). It can be seen that the probability decreases close to the boundary of the region defined by the conditions in (30) and (31). This is as expected, since there is a higher probability close to the boundary to falsify the invariance property. The optimal control action at the initial time step for the synthesis problem (29) is plotted in Fig. 1(b). Note that this action suggests using some polluting power generation (rather than relying on the more uncertain renewable power) whenever the energy levels are too low and close to the thresholds M_i . Conversely, whenever the energy levels are high enough, the policy suggests relying on renewables only.

Fig. 2(a) depicts the maximal probabilistic reach-avoid (at the initial time step) for the synthesis problem in (32). As expected, the obtained probability is higher for energy levels close to F_i . The optimal control action for the synthesis problem (32) is given in Fig. 2(b) (here plotted at the final step): since the goal is essentially to maximize the energy levels, it can be seen that the obtained policy selects the reliable coal plant for maximal energy generation (except when already in the goal sets).

The computational results have been obtained by discretizing the state space and computing the discrete value functions, and can be further improved by using an adaptive gridding procedure as in [17]. On the other hand, rather than computing value functions, the obtained MDP abstractions for safety and reach-avoid can be engaged in the verification of the properties of interest using model-checking software [14].

6. CONCLUSIONS AND NEXT STEPS

The contributions of this work are twofold. On the theoretical side, the article has zoomed in on issues related to reachability and invariance optimization over both finite and infinite horizons for non-autonomous stochastic hybrid systems, showing the sufficiency of memoryless optimal policies and tackling the problem of verification of specifications expressed as deterministic, finite-state automata. On the applicative side, a new computational scheme with explicit error quantification has been introduced and applied over a controller synthesis case study from the area of power systems. Both the theoretical and the computational outcomes nicely tailor back to known special models (autonomous, Markovian) from the literature.


The authors are interested in considering models with more complicated control structures (both discrete and continuous), and in looking at verification problems over ω -regular properties expressed as Büchi automata: the latter require non-trivial measure-theoretical results dealing with infinite-horizon problems that go beyond the scope of this work.

7. ACKNOWLEDGMENTS


The authors would like to thank Sadegh E. Z. Soudjani for the discussions on the bounds on the discretization procedure.

8. REFERENCES

- [1] MoVeS website. <http://www.movesproject.eu>.
- [2] A. Abate, S. Amin, M. Prandini, J. Lygeros, and S. Sastry. Computational approaches to reachability analysis of stochastic hybrid systems. In *Hybrid Systems: Computation and Control*, pages 4–17. Springer Verlag, 2007.
- [3] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16(6):1–18, 2010.
- [4] A. Abate, J.-P. Katoen, and A. Mereacre. Quantitative automata model checking of autonomous stochastic hybrid systems. In *Proceedings of the 14th ACM international conference on Hybrid Systems: Computation and Control*, pages 83–92, Chicago, IL, April 2011.
- [5] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.
- [6] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [7] D. Bertsekas and S. Shreve. *Stochastic Optimal Control: The Discrete Time Case*, volume 139. Academic Press, 1978.
- [8] H. Blom and J. Lygeros (Eds.). *Stochastic Hybrid Systems: Theory and Safety Critical Applications*. Number 337 in Lecture Notes in Control and Information Sciences. Springer Verlag, Berlin Heidelberg, 2006.
- [9] C. Cassandras and J. E. Lygeros. *Stochastic hybrid systems*, volume 24. CRC Press, 2007.
- [10] D. Chatterjee, E. Cinquemani, and J. Lygeros. Maximizing the probability of attaining a target prior to extinction. *Nonlinear Analysis: Hybrid Systems*, 5(2):367–381, 2011.
- [11] J. Ding, A. Abate, and C. Tomlin. Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications. *Proceedings of the 32nd American Control Conference*, 2013.




(a) Maximal probabilistic invariance




(b) Optimal control action at the initial step

Figure 1: Synthesis problem formulated in (29).



(a) Maximal probabilistic reach-avoid



(b) Optimal control action at the final step

Figure 2: Synthesis problem formulated in (32).

- [12] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker. Automated verification techniques for probabilistic systems. *Formal Methods for Eternal Networked Software Systems*, pages 53–113, 2011.
- [13] O. Hernández-Lerma and J. B. Lasserre. *Discrete-time Markov control processes*, volume 30 of *Applications of Mathematics (New York)*. Springer Verlag, New York, 1996.
- [14] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In H. Hermanns and J. Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 3920 of *Lecture Notes in Computer Science*, pages 441–444. Springer Verlag, 2006.
- [15] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. Springer Verlag, 1993.
- [16] W. Rudin. *Real and complex analysis*. McGraw-Hill Book Co., New York, third edition, 1987.
- [17] S. Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Quantitative Evaluation of Systems (QEST)*, pages 59–68, Aachen, DE, 2011.
- [18] S. Summers and J. Lygeros. A Probabilistic Reach-Avoid Problem for Controlled Discrete Time Stochastic Hybrid Systems. In *IFAC Conference on Analysis and Design of Hybrid Systems, ADHS*, Zaragoza, Spain, September 2009.
- [19] S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010.
- [20] I. Tkachev and A. Abate. On infinite-horizon probabilistic properties and stochastic bisimulation functions. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 526–531, Orlando, FL, December 2011.
- [21] I. Tkachev and A. Abate. Regularization of Bellman Equations for Infinite-Horizon Probabilistic Properties. In *Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control*, pages 227–236, Beijing, PRC, April 2012.