Formula-free Finite Abstractions for Linear Temporal Verification of Stochastic Hybrid Systems[*]

Ilya Tkachev
TU Delft - Delft University of Technology
i.tkachev@tudelft.nl

Alessandro Abate
TU Delft - Delft University of Technology
a.abate@tudelft.nl

ABSTRACT

Results on approximate model-checking of Stochastic Hybrid Systems (SHS) against general temporal specifications lead to abstractions that structurally depend on the given specification or with a state cardinality that crucially depends on the size of the specification. In order to cope with the associated issues of generality and scalability, we propose a specification-free abstraction approach that is general, namely it allows constructing a single abstraction to be then used for a whole cohort of problems. It furthermore computationally outperforms specification-dependent abstractions over linear temporal properties, such as bounded LTL (BLTL). The proposed approach unifies techniques for the approximate abstraction of SHS over different classes of properties by explicitly relating the error introduced by the approximation to the distance between transition kernels of abstract and concrete models, and by propagating the error in time over the horizon of the specification. The new technique is compared over a case study to related results in the literature.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Stochastic processes

Keywords

Markov processes, stochastic hybrid systems, formal verification, probabilistic model-checking, linear temporal specifications, finite abstractions, approximate bisimulations.

1. INTRODUCTION

Stochastic Hybrid Systems (SHS) provide a powerful modeling framework for diverse application areas such as systems biology, air traffic control, power networks and telecommunication systems [8, 18]. The reliable employment of SHS models demands solid foundations for their analysis and verification.

One of the most prominent tools for the verification of finite-state systems is model checking. In particular, with focus on the discrete-time case, the verification theory of finite-state probabilistic models known as discrete-time Markov Chains (dt-MC) is mature [5]. The formal verification of dt-MC is enabled by probabilistic model checking: in this instance verification problems allow for explicit solutions or answers that can be obtained in a numerically efficient way leveraging dedicated probabilistic model checking software [13, **?**]. On the other hand, the price to pay for the descriptive generality of SHS, models that are characterized by an uncountable state space, is the lack of explicit solutions and the undecidability for most verification problems [1]. A possible approach to overcome this issue is based on the concept of *abstraction*, namely "a quotient system that preserves some properties of interest, while ignoring details" [17]. Abstractions are ideally *finite*, as this often leads to problem decidability and to explicit solutions. For SHS, finite abstractions are naturally dt-MC. Notice however that whenever the original state space is infinite (as for SHS models), it is often only possible to synthesize a finite model that is an *approximate* abstraction of the concrete one [11].

Many properties of interest for SHS can be expressed as PCTL formulae or as linear temporal (LT) specifications [5]. With focus on the former, the work in [18] has formally related the verification of PCTL formulae to the computation of corresponding *value functions* defined over the state space of a SHS. Given an initial state, such a value function represents the probability that the execution of SHS satisfies a given PCTL path formula. Thus one can relate the quality of an approximate abstraction with respect to a given property to the difference between value functions computed respectively over the abstraction and over the concrete model [1].

So far only *formula-dependent* techniques have been developed to find approximate finite abstractions of SHS. The first step of these techniques is to leverage dynamic programming (DP) [6] principles to derive DP-like recursions for the value function related to a given formula. The second step is to build an abstraction in order to numerically compute integrals involved in the DP recursions, with explicit bounds on the approximation error. The work in [1] has developed this approach and applied it to the problem of probabilistic safety (or invariance) within the class of PCTL formulae. Later, [20] has further improved these results by relaxing some of the model assumptions and by finding tighter error bounds, which in turn led to a lower cardinality of the abstraction required to match a given precision. Both works have used DP procedures for bounded-horizon

[*] 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.

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

safety value function developed in [3]. Such recursions have also been developed for the probabilistic reach-avoid problem in [21]. Although PCTL path formulae for safety and reach-avoid are part of a more general class of LT specifications, it has not been clear yet whether DP recursions could be also developed for other LT specifications. Due to this reason, [2] has suggested a new approach for the verification of LT specifications, by reducing the original problem to the safety one defined over a new SHS, the latter being the product between the original SHS and the automaton corresponding to the specification of interest. Let us mention that the approximate abstraction methods discussed above are limited to the verification of bounded-horizon specifications – the work in [23, 24] argued that direct abstraction may not work for infinite-time problems, and developed alternative techniques to tackle them.

Notice that all the described methods require building a brand new approximate abstraction for each given different formula. This contribution is thus challenged to develop *formula-free* finite abstractions over SHS. More precisely: given a SHS $\mathfrak{D}$, a bounded time horizon n and a precision level ε, we provide an explicit way to build a dt-MC $\hat{\mathfrak{D}}$ which allows computing value functions of any n-bounded LT specification with an error that does not exceed ε. This result has several important features. Firstly, no matter how many properties are to be model-checked against $\mathfrak{D}$, one has to construct only a single formula-free abstraction $\hat{\mathfrak{D}}$; one can then use any desired model-checking software to do verification on $\hat{\mathfrak{D}}$ [13, **?**]. Secondly, the approach we propose is especially useful when one needs to look into LT specifications that are richer than PCTL path formulae, for example BLTL specifications (their applicative importance was recently emphasized in [15]). For such problems, the only technique available in the literature requires solving the safety problem over the product between a SHS and an automaton expressing the formula [2]. However, the error for the computation of the safety value function depends on the size of the state space, thus the overall error is crucially dependent on the size of the automaton: this is not the case for the proposed new formula-free abstraction method. Lastly, the approach we use to quantify the error of the formula-free abstraction is directly extendable from LT specifications to other verification problems, such as those based on reward properties, and it allows developing a unified technique for the approximate abstraction of SHS over diverse classes of specifications.

For notational convenience, results in this paper are stated for discrete-time Markov processes (dt-MP), a class of models that is more general than discrete-time SHS. The structure of the paper is the following: Section 2 introduces classes of models and specifications of interest, and formalizes the model-checking of LT specifications against dt-MP. Section 3 describes the abstraction technique for BLTL and compares its performance with results from the literature. An extension of the technique from BLTL to other specifications is presented in Section 4. Computational examples are given in Section 5, whereas Section 6 contains the conclusions. Due to space constraints, the proofs of the statements are omitted from this manuscript.

2. MODELS AND SPECIFICATIONS

2.1 Notations

Let us recall some concepts from measure theory – for a detailed exposition the interested reader is referred to the books [7, Chapters 1-3] and [10, Chapters 1-3].

We use $\mathbb{N} = \{1, 2, \dots\}$ to denote the set of natural numbers and write $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$ and $\overline{m,n} = \{m, m+1, \dots, n\}$ whenever $m, n \in \mathbb{N}_0$ and $m < n$. We also use the notation $\mathbb{R}$ for the set of real numbers and $\bar{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ for the set of extended reals.

For any set X and collection of its subsets $\mathscr{C} \subseteq 2^X$, the σ-algebra generated by $\mathscr{C}$ is denoted by $\sigma(\mathscr{C})$. For example, $\mathscr{B}(\mathbb{R})$ is the *Borel* σ-algebra on $\mathbb{R}$, and is generated by the class of all open subsets of $\mathbb{R}$. We always assume $\mathbb{R}$ to be endowed with its Borel σ-algebra. Given two measurable spaces $(X, \mathscr{X})$ and $(Y, \mathscr{Y})$ the map $f : X \to Y$ is $\mathscr{X}/\mathscr{Y}$*-measurable* if $f^{-1}(A) \in \mathscr{X}$ for any $A \in \mathscr{Y}$. In the case of a $f : X \to \mathbb{R}$ we say that f is $\mathscr{X}/\mathscr{B}(\mathbb{R})$-measurable. For any function $f : X \to \mathbb{R}$ we denote its sup-norm by $\|f\| := \sup_{x\in X} |f(x)|$. We denote by $b\mathscr{X}$ the space of all bounded $\mathscr{X}$-measurable functions.

If (X, ρ) is a metric space, then $\mathrm{diam}(A) = \sup_{x,y\in A} \rho(x, y)$ denotes the diameter of a set $A \subseteq X$. Let $I \subseteq \mathbb{N}_0$ be some index set and $(X_i, \mathscr{X}_i)$ be a measurable space for any $i \in I$. We denote the corresponding *product measurable space* by $\prod_{i\in I}(X_i, \mathscr{X}_i)$.

We call a function $\mu : \mathscr{X} \to \bar{\mathbb{R}}$ a *measure* on $(X, \mathscr{X})$ if $\mu(\emptyset) = 0$, if μ takes at most one of the values $\pm\infty$, and if $\mu(\bigcup_{n\in\mathbb{N}} A_n) = \sum_{n=1}^{\infty} \mu(A_n)$ for any sequence of disjoint sets $(A_n)_{n\in\mathbb{N}} \subseteq \mathscr{X}$, where the series converges absolutely if $\mu(\bigcup_{n\in\mathbb{N}} A_n)$ is finite. Such measures are also called *signed* measures, in contrast to *positive* measures (namely measures taking values over a subset of $\mathbb{R}_+$). A positive measure $\mu : \mathscr{X} \to \mathbb{R}_+$ is called a *probability* measure (or *distribution*) whenever it holds that $\mu(X) = 1$.

The last notion to be considered is that of a *kernel*: given two measurable spaces $(X, \mathscr{X})$ and $(Y, \mathscr{Y})$, a kernel Q on $(Y, \mathscr{Y})$ given $(X, \mathscr{X})$ is a function $Q : X \times \mathscr{Y} \to \bar{\mathbb{R}}$ such that $Q_x(\cdot)$ is a measure on $(Y, \mathscr{Y})$ for all $x \in X$, and such that the function $x \mapsto Q_x(A)$ is $\mathscr{X}$-measurable for any $A \in \mathscr{Y}$. We say that Q is a stochastic kernel if for any $x \in X$ the measure Q_x is a probability measure. If $(Y, \mathscr{Y}) = (X, \mathscr{X})$ we simply say that Q is a kernel on $(X, \mathscr{X})$, in that case we sometimes write $Q(x, A)$ for $Q_x(A)$.

2.2 Discrete-time Markov processes

This work considers a class of models known as discrete-time Markov processes (dt-MP). Any dt-MP $\mathfrak{D}$ can be uniquely characterized by a triple $(E, \mathscr{E}, P)$, where $(E, \mathscr{E})$ is a measurable space and $P : E \times \mathscr{E} \to [0, 1]$ is a stochastic kernel [19]. The *state space* of $\mathfrak{D}$ is $(E, \mathscr{E})$ and the elements $x \in E$ of the state space are the *states* of $\mathfrak{D}$. P is said to be a *transition kernel* of $\mathfrak{D}$ and the quantity $P(x, A)$ represents the probability of going from the state x to the set $A \in \mathscr{E}$. The work in [3] provides details of the embedding of discrete-time SHS into the dt-MP framework.

The space of trajectories of $\mathfrak{D}$ is given by the product space $(\Omega, \mathscr{F}) := \prod_{k=0}^{\infty}(E, \mathscr{E})$, with a generic trajectory denoted by

$$\omega = (\omega_0, \omega_1, \dots) \in \Omega,$$

where for any $n \in N_0$, $\omega_n \in E$ represents the state of the system modeled by the dt-MP $\mathfrak{D}$ at time epoch n. It further follows from [19, Theorem 2.8] that there exists a unique kernel P defined on $(\Omega, \mathscr{F})$ given $(E, \mathscr{E})$ that satisfies, for any $n \in \mathbb{N}_0$,

$$\mathsf{P}_x\left(\prod_{i=0}^{n} A_i \times \prod_{i=n+1}^{\infty} E\right) = 1_{A_0}(x) \int_{A_1} \dots \int_{A_n} P(x_{n-1}, \mathrm{d}x_n) \dots P(x, \mathrm{d}x_1),$$

where $A_i \in \mathscr{E}$ are arbitrary sets and $i \in \overline{0,n}$. The measure P_x tells which *events* (measurable sets of trajectories) are more probable to happen for $\mathfrak{D}$ than others, given that the initial state is x. By slight abuse of notation, we say that $(\Omega, \mathscr{F}, \mathsf{P})$ as above is a *canonical probability space* for the dt-MP $\mathfrak{D} = (E, \mathscr{E}, P)$.

Any set $F \in \mathscr{F}$ is called an event. We are particularly interested in the following classes of events: given $n \in \mathbb{N}_0$, $F \in \mathscr{F}$ is an *n-horizon* event if $\omega \in F$ together with $\omega_i = \omega'_i$, $i \in \overline{0,n}$ implies $\omega' \in F$. The σ-algebra of n-horizon events is given by

$$\mathscr{F}_n := \sigma\left\{ \prod_{i=0}^{n} A_i \times \prod_{i=n+1}^{\infty} E \,\middle|\, A_i \in \mathscr{E},\ i \in \overline{0,n} \right\}, \tag{2.1}$$

and it represents the history of the observations of the dt-MP $\mathfrak{D}$ up to the time epoch n. We say that $(\mathscr{F}_n)_{n\in\mathbb{N}_0}$ is the *natural filtration* of the dt-MP $\mathfrak{D}$. The *Markov property* $\mathsf{P}_x(\omega_{n+1} \in A|\mathscr{F}_n) = P(\omega_n, A)$ suggests that the distribution of the next state of $\mathfrak{D}$ depends on its history only through the current state.

The following concept is important for the abstraction procedure given in Section 3: a dt-MP $\mathfrak{D} = (E, \mathscr{E}, P)$ is called *finitely generated* (f.g.) whenever $\mathscr{E}$ is finite. As an example, we call $\mathfrak{D}$ a discrete-time Markov Chain (dt-MC) if E is finite. Clearly, the finiteness of E implies the finiteness of $\mathscr{E} \subseteq 2^E$, hence we have that a dt-MC is a f.g. dt-MP. However, the inverse statement is not necessarily true: even if $\mathfrak{D}$ is finitely generated, the set E can be uncountable. Such f.g. dt-MP is an artificial object used below as an intermediate step in the abstraction of a general dt-MP into a dt-MC (cf. Figure 1).

2.3 Linear temporal specifications

The class of dt-MP models has been introduced to be model-checked against linear temporal (LT) specifications, which satisfiability can be explicitly decided over any given trajectory $\omega \in \Omega$. Thus, the satisfaction relation is defined over the set $\models \subseteq \Omega \times \mathtt{LT}$ – here the symbol $\mathtt{LT}$ designates an abstract class of LT specifications, which will be further detailed below. Each LT specification φ can be characterized by its sat-set as

$$\mathtt{Sat}_\Omega(\varphi) := \{\omega \in \Omega : \omega \models \varphi\},$$

which is the subset of the space Ω containing exactly those trajectories that satisfy the specification φ.

In accordance to [5, Section 10.3], we define the probabilistic model-checking problem over LT specifications as follows. Given a dt-MP $\mathfrak{D} = (E, \mathscr{E}, P)$, an initial state $x \in E$, and an LT specification φ, find the probability that the trajectory of $\mathfrak{D}$ starting from state x satisfies φ. More precisely, one has to evaluate

$$\mathsf{P}_x\big(\mathtt{Sat}_\Omega(\varphi)\big). \tag{2.2}$$

Recall that for any initial state $x \in E$ the probability measure P_x is only defined over the σ-algebra $\mathscr{F}$, and not over arbitrary collections of trajectories. Due to this reason, we say that the probabilistic model-checking problem is well posed for $\mathfrak{D}$ if and only if the quantity in (2.2) is defined, that is if $\mathtt{Sat}_\Omega(\varphi) \in \mathscr{F}$.

Whenever the probability in (2.2) is well defined for any φ in a given class of linear temporal specifications $\mathtt{LT}$, one can follow the procedure described in [4, Section 9.1.3] and do equivalently a probabilistic model-checking of *state* specifications, which would yield a `true/false` answer instead of an arbitrary number in the interval $[0,1]$, as in the case of (2.2). More precisely, let us define the class of state specifications as $\mathtt{LT}_{\mathrm{state}} = \mathtt{LT} \times 2^{[0,1]}$ with elements (φ, I) where $\varphi \in \mathtt{LT}$ is an LT specification and I is a subset of $[0,1]$. The satisfaction relation can then be defined on the product set $\models \subseteq E \times \mathtt{LT}_{\mathrm{state}}$ by

$$x \models (\varphi, I) \iff \mathsf{P}_x\big(\mathtt{Sat}_\Omega(\varphi)\big) \in I.$$

Since the quantity in (2.2) needs to be well defined to be evaluated, we first discuss measurability issues.

Let us focus on a particular class of LT specifications comprising automata [5, Chapter 4] and LTL [5, Chapter 5]. In both cases, specifications are expressed via languages over certain alphabets.[1] Thus, it is sufficient to consider measurability properties of such languages, without focusing on a particular modal logic, thereafter tailoring the developed results to the special cases of LTL or automata, if needed.

We call an *alphabet* some finite set Σ, we call *letters* its elements $\sigma \in \Sigma$ and we call *words* finite or infinite sequences of letters. Let $\mathfrak{S} = \Sigma^{\mathbb{N}_0}$ be the set of all infinite words over the alphabet Σ. The generic element of $\mathfrak{S}$ is denoted by

$$\pi = (\pi_0, \pi_1, \dots) \in \mathfrak{S}, \quad \pi_i \in \Sigma, \quad i \in \mathbb{N}_0.$$

The infinite *language* φ over Σ is an arbitrary collection of infinite words, i.e. $\varphi \subseteq \mathfrak{S}$. We regard words as *traces* of trajectories of a dt-MP, as already done for the case of dt-MC [5, Section 10.3] and non-probabilistic systems [22]. Note that the canonical trajectory space Ω contains only infinite trajectories. It is thus convenient to focus on infinite words and languages, since their finite counterparts can be easily embedded in this framework: to each finite word $\pi' = (\pi'_0, \dots, \pi'_n)$ there corresponds an infinite language $\{\pi'\} \times \prod_{i=n+1}^{\infty} \Sigma$ (we call such a language a *basic* language). The embedding of a finite language into an infinite one can be done in a similar way, word by word. As a result, we shall deal only with infinite words and languages and omit the word "infinite" in both cases.

We regard each language as a specification over a dt-MP as follows. In order to characterize the satisfaction relation $\models$ between trajectories $\omega \in \Omega$ and specifications (or languages) $\varphi \subseteq \mathfrak{S}$, let us introduce the labeling map $\mathsf{L} : E \to \Sigma$. As a result, to each state $x \in E$ of the dt-MP we assign a letter $\mathsf{L}(x) \in \Sigma$. While the system described by a dt-MP evolves in time, it produces a trajectory $\omega_0, \omega_1, \dots$ which in turn produces the word $\mathsf{L}(\omega_0)\mathsf{L}(\omega_1)\dots$ called the *trace* of ω [5, Section 3.2.2]. We say that a trajectory satisfies the specification expressed as an infinite language if its trace belongs to such a language.

More formally, we denote by $\mathsf{L}_* : \Omega \to \mathfrak{S}$ the element-wise extension of L given by $\mathsf{L}_*(\omega_0, \omega_1, \dots) := (\mathsf{L}(\omega_0), \mathsf{L}(\omega_1), \dots)$. We define the satisfaction relation as follows:

$$\omega \models \varphi \iff \mathsf{L}_*(\omega) \in \varphi. \tag{2.3}$$

It follows from (2.3) that $\mathtt{Sat}_\Omega(\varphi) = \mathsf{L}_*^{-1}(\varphi)$ for all $\varphi \in \mathfrak{S}$.

Having characterized sat-sets $\mathtt{Sat}_\Omega$ through the labeling map, we can state the main result about measurability of the sat-sets used in our framework. For this purpose, we introduce the important concept of *measurable language*. Let us endow the alphabet Σ with a discrete σ-algebra 2^Σ, which makes $(\Sigma, 2^\Sigma)$ a measurable space. Hence, $\mathfrak{S}$ can be endowed with its product σ-algebra, which is further denoted by $\mathscr{S}$.

Definition 1. *[16] We say that the language φ over the alphabet Σ (so that $\varphi \subseteq \mathfrak{S}$) is* measurable*, whenever $\varphi \in \mathscr{S}$.*

Obviously, the collection of all measurable languages is just the σ-algebra $\mathscr{S}$, which is closed under intersections and complementations by definition. The following theorem is crucial for our further considerations.

Theorem 1. *If L is a $\mathscr{E}/2^\Sigma$-measurable map, then the sat-set of any measurable language $\varphi \in \mathscr{S}$ is a measurable subset of Ω.*

[1] In our case there is no substantial difference whether to start from a finite set of *atomic propositions* $\mathtt{AP}$ and define an alphabet as $\Sigma = 2^{\mathtt{AP}}$, or to start directly from some finite set Σ as an alphabet. For ease of notation we have chosen the latter.

Whenever the map L is $\mathscr{E}/2^{\Sigma}$-measurable, we call a quintuple $(E, \mathscr{E}, P, \Sigma, \mathsf{L})$ a *labeled discrete-time Markov process* and write `ldt`-MP for short. This notion is different from that of Labeled Markov process (LMP) defined in [**?**], where the primary goal of using labels is to model the non-determinism in transitions. However, an `ldt`-MP is similar to a general Labeled Markov Chain [14, Definition 1] with the only difference that in the latter case $\mathsf{L}^{-1} : \Sigma \to \mathscr{E}$ is said to be the labeling map. We say that the `ldt`-MP is finitely generated (f.g.) if the σ-algebra $\mathscr{E}$ is finite; in particular, if the state space E is finite we use the name *labeled discrete-time Markov Chain* (`ldt`-MC) in place of `ldt`-MP.

Theorem 1 states that the model-checking of measurable languages $\varphi \in \mathscr{S}$ over an `ldt`-MP is a well-posed problem in the sense that (2.2) is well-defined. Although not all infinite languages are measurable [16, Example 8], the important class of ω-regular languages satisfies the measurability property.

Proposition 1. *[16] If $\varphi \subseteq \mathfrak{S}$ is ω-regular, then $\varphi \in \mathscr{S}$.*

2.4 BLTL specifications

Although Proposition 1, together with Theorem 1, implies that the probabilistic model-checking of `ldt`-MP against ω-regular properties, such as LTL formulae and Büchi automata [5], is a well posed problem, its solution is in general difficult to find: as it was shown in previous work [23, 24], the solution of each particular infinite time-horizon problem depends on structural features of the `dt`-MP, such as the presence of absorbing sets. Due to this reason, we focus on a general class of bounded time-horizon specifications, which are still very important for applications, for instance in in systems biology [15] and in financial mathematics [25, Part III].

Let us first formalize what the horizon of a specification is. A specification $\varphi \subseteq \mathfrak{S}$ has a horizon equal to $n \in \mathbb{N}_0$ if, for any word $\pi \in \mathfrak{S}$, the value of the letters in π beyond position n does not affect whether $\pi \in \varphi$. More precisely, we call a language $\varphi \subseteq \mathfrak{S}$ *bounded* if there exists $n \in \mathbb{N}_0$ such that

$$(\pi \in \varphi) \wedge \left(\pi_i = \pi'_i,\ i \in \overline{0,n}\right) \quad \Rightarrow \quad (\pi' \in \varphi) \tag{2.4}$$

holds true for all words $\pi, \pi' \in \mathfrak{S}$. Clearly, if φ satisfies (2.4) for some $n \in \mathbb{N}_0$, then it also satisfies it for $n+1$. Thus, it is natural to define the horizon of $\varphi \subseteq \mathfrak{S}$ as follows:

$$\mathsf{H}(\varphi) := \inf \left\{ n \in \mathbb{N}_0 : (2.4) \text{ holds true for } n \right\}.$$

In other words $\mathsf{H}(\varphi)$ is the smallest $n \in \mathbb{N}_0$ which makes (2.4) hold true for φ, if such n exists, whereas $\mathsf{H}(\varphi) = \infty$ otherwise, where as usual $\inf(\emptyset) := \infty$. As an example, each basic language $\varphi' = \{\pi'\} \times \prod_{i=n+1}^{\infty} \Sigma$, where $\pi' = (\pi'_0, \dots, \pi'_n) \in \Sigma^n$ is bounded, and $\mathsf{H}(\varphi') = n$ whenever Σ has more than one letter. Conversely, it follows from the finite cardinality of the alphabet Σ that each bounded language is a finite union of basic languages. As a result, since any basic language is measurable, so is each bounded language. The equivalent formula for H follows:

$$\mathsf{H}(\varphi) = \inf \left\{ n \in \mathbb{N}_0 : \mathsf{L}_*^{-1}(\varphi) \in \mathscr{F}_n \right\}, \tag{2.5}$$

where $\mathscr{F}_n$ is given by (2.1). Thus, $\mathscr{S}_n := \{\varphi \subseteq \mathfrak{S} : \mathsf{H}(\varphi) \leq n\}$ – the collection of all languages with an horizon not exceeding n is a sub-σ-algebra of $\mathscr{S}$, and hence it is closed under intersections, unions and complementations.

Clearly, each bounded language can be written via the finite number of basic languages that it contains, which in turn can be written via the corresponding finite words. It is possible to consider some alternative, compact representations of bounded languages. For instance, they appear as accepting languages of Deterministic Finite Automata (DFA) [5] taking only runs that are bounded by some a-priori integer $n \in N_0$ [2]: we give the precise definition later, in Section 3.4.

Another way to compactly encode a bounded language is via BLTL formulae: we now tailor to our study the definition of this logic given for a different class of models in [15]. The syntax of BLTL over alphabet Σ is given by the following grammar:

$$\Phi \quad ::= \quad \sigma \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \mathsf{X}\Phi. \tag{2.6}$$

We define the semantics of BLTL by introducing the satisfaction relation between BLTL formulae and infinite words over Σ:

$$\begin{aligned} \pi \models \sigma \quad &\Leftrightarrow \quad \pi_0 = \sigma \\ \pi \models \neg\Phi \quad &\Leftrightarrow \quad \pi \not\models \Phi \\ \pi \models \Phi \wedge \Psi \quad &\Leftrightarrow \quad \pi \models \Phi \text{ and } \pi \models \Psi \\ \pi \models \mathsf{X}\Phi \quad &\Leftrightarrow \quad \theta(\pi) \models \Phi, \end{aligned}$$

where the *shift operator* $\theta : \mathfrak{S} \to \mathfrak{S}$ is given as follows:

$$\theta(\pi_0, \pi_1, \pi_2, \dots) = (\pi_1, \pi_2, \dots).$$

For any BLTL formula Φ, we define its *accepting language* $\mathfrak{L}(\Phi)$ to be the collection of all infinite words that satisfy this formula, namely $\mathfrak{L}(\Phi) := \{\pi \in \mathfrak{S} : \pi \models \Phi\}$.

From the basic BLTL grammar in (2.6) we define the disjunction of two formulae as $\Phi_1 \vee \Phi_2 := \neg(\neg\Phi_1 \wedge \neg\Phi_2)$ and the truth formula as $\texttt{true} := \bigvee_{\Sigma} \sigma$. The "neXt" temporal operator X allows defining the "bounded Until" one. We first introduce powers of X inductively by $\mathsf{X}^0\Phi := \Phi$ and $\mathsf{X}^n\Phi := \mathsf{X}(\mathsf{X}^{n-1}\Phi)$.

We further define $\Phi_1 \mathsf{U}^{\leq n} \Phi_2 := \Phi_2 \vee \bigvee_{i=1}^{n} \left(\bigwedge_{j=0}^{i-1} \mathsf{X}^j \Phi_1 \wedge \mathsf{X}^i \Phi_2 \right)$ for $n \in N_0$. This formula has the following familiar semantics:

$$\begin{aligned} \pi \models \Phi_1 \mathsf{U}^{\leq n} \Phi_2 \quad \Leftrightarrow \quad &\pi \models \Phi_2 \text{ or } \theta^i \pi \models \Phi_2 \text{ for some } i \in \overline{1,n} \text{ and} \\ &\theta^j \pi \models \Phi_1 \text{ for all } 0 \leq j < i. \end{aligned}$$

Other temporal modalities can be defined using the bounded until operator, e.g. "bounded eventually" as $\Diamond^{\leq n}\Phi := \texttt{true}\mathsf{U}^{\leq n}\Phi$ and "bounded always" as $\Box^{\leq n}\Phi := \neg(\Diamond^{\leq n}\neg\Phi)$. Using BLTL we can then pose well-known verification problems, such as probabilistic reach-avoid (using $\mathsf{U}^{\leq n}$), reachability (using $\Diamond^{\leq n}$), and safety (using $\Box^{\leq n}$). As an example, the specification induced by the language $\mathfrak{L}(\Box^{\leq n}\sigma)$ is equivalent to a finite-horizon safety one [1]. Moreover, BLTL allows to consider more complex properties: let $\Sigma = \{\alpha, \beta, \gamma\}$ and consider the following formula:

$$\Phi = \Box^{\leq 100}(\alpha \vee \beta) \wedge \Diamond^{\leq 50}\Box^{\leq 50}\alpha.$$

Supposing that $\{\alpha, \beta\}$ corresponds to the safe set and $\{\alpha\}$ to the target set, formula Φ reads as "the system will be safe for at least 100 steps and within the following 50 iterations it will end up spending at least 50 consecutive steps in the target set." For more instances of BLTL formulae see e.g. [15].

Finally, the horizon of accepting languages of BLTL formulae can be found as follows. Clearly, we have that $\mathsf{H}(\mathfrak{L}(\sigma)) = 0$ and the following relations hold

$$\mathsf{H}(\mathfrak{L}(\neg\Phi)) = \mathsf{H}(\mathfrak{L}(\Phi)), \quad \mathsf{H}(\mathfrak{L}(\mathsf{X}\Phi)) = \mathsf{H}(\mathfrak{L}(\Phi)) + 1$$

$$\mathsf{H}(\mathfrak{L}(\Phi_1 \wedge \Phi_2)) = \max(\mathsf{H}(\mathfrak{L}(\Phi_1)), \mathsf{H}(\mathfrak{L}(\Phi_2))).$$

By induction, for any BLTL formula the horizon of its accepting language is finite and hence by (2.5) such a language is measurable, which leads to the well-posedness of probabilistic model checking of BLTL. Note also that for each basic language φ there exists a formula Φ such that $\varphi = \mathfrak{L}(\Phi)$. As a result, BLTL allows describing all possible bounded languages.

3. FINITE ABSTRACTIONS OF LDT-MP

In order to progressively introduce the main results presented in this work, let us first discuss how one can perform verification of BLTL formulae against stochastic models with finite state spaces, specifically `ldt`-MC. Notice that, from (2.6), the grammar of BLTL is a fragment of LTL. Thus, any BLTL formula can be expressed via an automaton, and its verification over an `ldt`-MC is known [5, Chapter 10.3]. On the other hand, any BLTL formula Φ can be directly expressed via the basic components (finite words) of its accepting language $\mathfrak{L}(\Phi)$: one can further compute probabilities of sat-sets for each word and find the sum thereof to obtain the probability of the sat-set for $\mathfrak{L}(\Phi)$.

With focus on the general case of `ldt`-MP, automata model-checking was studied in [2]. However, as it has been mentioned in the introduction, the error for the approximate solution depends on the size of the automaton (cfr. Section 3.4). This is especially important in case of BLTL, which often leads to automata with large state spaces (cfr. Section 5).

To cope with the issues described above, this contribution provides a formula-free abstraction technique made up of two steps. We show that any BLTL model-checking problem over f.g. `ldt`-MP can be explicitly reduced to the same problem over a certain `ldt`-MC. Perhaps not a striking result per se, it motivates looking for finitely generated approximate abstractions of general `ldt`-MP. The overall abstraction scheme is depicted in Figure 1: the general `ldt`-MP is *approximately* abstracted as a f.g. `ldt`-MP, which in turn is *exactly* abstracted as a `ldt`-MC.

3.1 Quotient ldt-MC of a f.g. ldt-MP

A f.g. `ldt`-MP with an infinite state space is an artificial object that is used as an intermediate step between a general `ldt`-MP and a `ldt`-MC in the abstraction procedure. Intuitively, a finitely generated abstraction is useful since it has the same uncountable state space as the original model but only a discrete measurability structure given by its finite σ-algebra. To be more precise, let us first comment on the structure of some arbitrary f.g. `ldt`-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$. Since the σ-algebra $\mathscr{E}$ is finite, it follows that there exists a finite measurable partition of E which generates $\mathscr{E}$, i.e. there exists a finite collection of disjoint non-empty sets $E_1, \dots, E_N$ satisfying $\bigcup_{i=1}^N E_i = E$, and such that $\mathscr{E} = \sigma(E_1, \dots, E_N)$. Although the state space E can still be an uncountable set, the finite structure of $\mathscr{E}$ in particular implies that all measurable maps are constant when restricted to the partition sets. This follows directly from the definition of measurability and the fact that $\mathscr{E}$ is generated by a finite partition. For example, for the stochastic kernel of $\mathfrak{D}$ it holds that

$$P(x', A) = P(x'', A) \quad \forall\ x', x'' \in E_i,\ i \in \overline{1,N},\ A \in \mathscr{E}. \tag{3.1}$$

Moreover, any set $A \in \mathscr{E}$ admits a unique representation of the form $A = \bigcup_{i \in I} E_i$ where $I \subseteq \overline{1,N}$ is some index set, e.g. it is empty for the case $A = \emptyset$. As a result, the stochastic kernel P is uniquely determined by the matrix with entries given by

$$p_{ij} := P(x_i, E_j), \tag{3.2}$$

where x_i can be *any* point in E_i, as it follows from (3.1).

Notice that the construction above means that only the sets E_i, rather than single states $x \in E$ or general subsets of E, are "observable" locations. For example, if $(\Omega, \mathscr{F}, \mathsf{P})$ is a probability space of $\mathfrak{D}$ then the probability that ω_1 belongs to E_i is well-defined and is given by $\mathsf{P}_x(\omega_1 \in E_i) = P(x, E_i)$. However, for any non-empty $E' \subsetneq E_i$ the probability $\mathsf{P}_x(\omega_1 \in E')$ is not defined since $E' \notin \mathscr{E}$. This can be interpreted as follows: we know the one-step transition probability for entering the set E_i, but nothing can be said about the transition probability into a generic subset of E_i.

The above discussion leads to regard the partition sets E_i as equivalence classes of states in E, and to construct a *quotient* `ldt`-MC over the finite state space made up by the collection of such equivalence classes. Such an `ldt`-MC is characterized by transition probabilities derived from the discrete structure of the kernel P given in (3.1). In order to formally present this object, let us introduce the *indexing* map $\mathtt{I} : E \to \overline{1,N}$, defined uniquely by the formula $x \in E_{\mathtt{I}(x)}$, that assigns to each $x \in E$ the index of the partition set that state x belongs to.

Definition 2. *Given a f.g.* `ldt`*-MP* $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ *we define the* quotient `ldt`*-MC by* $\hat{\mathfrak{D}} = (\hat{E}, \hat{\mathscr{E}}, \hat{P}, \Sigma, \hat{\mathsf{L}})$*, where the state space is* $\hat{E} = \overline{1,N}$ *and* $\hat{\mathscr{E}} = 2^{\hat{E}}$*;* $\hat{P}$ *is defined by the stochastic matrix* $\hat{P}(i, \{j\}) = p_{ij}$*, with* p_{ij} *given by* (3.2)*; the labeling map is* $\hat{\mathsf{L}}(i) = \mathsf{L}(x)$ *where* x *is any element of* E_i*.*

Note that in Definition 2 the new labeling map $\hat{\mathsf{L}}$ is well-defined since $\mathsf{L} : E \to \Sigma$ is $\mathscr{E}/2^\Sigma$-measurable and hence its restriction to any partition set E_i is constant. Let us emphasize that we have used the name *quotient* because $\hat{E}$ can be thought of as a finite collection of equivalence classes of states in the original state space E with an equivalence relation $\sim$ generated by the partition $E_1, \dots, E_N$, i.e. $x' \sim x''$ if and only if $\mathtt{I}(x') = \mathtt{I}(x'')$. Let $(\hat{\Omega}, \hat{\mathscr{F}}, \hat{\mathsf{P}})$ denote the canonical probability space of $\hat{\mathfrak{D}}$. The main result on the quotient `ldt`-MC is stated as follows.

Theorem 2. *For any specification* $\varphi \in \mathscr{S}$ *it holds that*

$$\mathsf{P}_x\left(\mathtt{Sat}_\Omega(\varphi)\right) = \hat{\mathsf{P}}_{\mathtt{I}(x)}\left(\mathtt{Sat}_{\hat{\Omega}}(\varphi)\right).$$

3.2 F.g. abstraction of a general ldt-MP

We have shown that the probabilistic model-checking of a f.g. `ldt`-MP can be reduced to that of its quotient `ldt`-MC. This motivates us to look for finitely generated abstractions of general `ldt`-MP. Obviously, such an abstraction is in general not exact, hence there is no hope for equivalence results as in Theorem 2. The best one can do is constructing an abstract f.g. `ldt`-MP that is designed to approximate the value in (2.2) for the original `ldt`-MP. This leads to the introduction of an appropriate notion of distance between probability measures.

Definition 3. *Let* $\mu : \mathscr{X} \to \bar{\mathbb{R}}$ *be a signed measure defined on a measurable space* $(X, \mathscr{X})$*; its* total variation norm *is given by*

$$\|\mu\|_{\mathscr{X}} := \sup_{A \in \mathscr{X}} \left(\left|\mu(A)\right| + \left|\mu(A^c)\right| \right).$$

If Q *is a kernel on* $(X, \mathscr{X})$ *given* $(Y, \mathscr{Y})$*, we use the same notation for the induced norm:* $\|Q\|_{\mathscr{X}} := \sup_{y \in Y} \|Q_y\|_{\mathscr{X}}$*.*

Let us now consider an `ldt`-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$. In order to construct a finitely generated abstraction, we are going to retain its state space E and its logical structure (given by Σ and L): this is done in order to avoid the necessity of abstracting specifications in addition to the model. As a result, the abstraction is obtained modifying $\mathscr{E}$ and P into some finite $\tilde{\mathscr{E}} \subseteq \mathscr{E}$ and $\tilde{P}$, thus resulting in model $\tilde{\mathfrak{D}} = (E, \tilde{\mathscr{E}}, \tilde{P}, \Sigma, \mathsf{L})$. Note that it is always possible to choose some finite $\tilde{\mathscr{E}}$ and $\tilde{P}$, e.g. one can start with $\tilde{\mathscr{E}}$ generated by the labels, then define $\tilde{P}$ such that every label is absorbing. Although this is rarely an optimal choice, it can be further refined as discussed in Section 3.3.

Let $(\Omega, \tilde{\mathscr{F}}, \tilde{\mathsf{P}})$ be the probability space of $\tilde{\mathfrak{D}}$ and let $(\tilde{\mathscr{F}}_n)_{n \in N_0}$ be its natural filtration. The following result shows that the distance between measures P and $\tilde{\mathsf{P}}$ propagates at most linearly in time via the distance between transition kernels P and $\tilde{P}$.

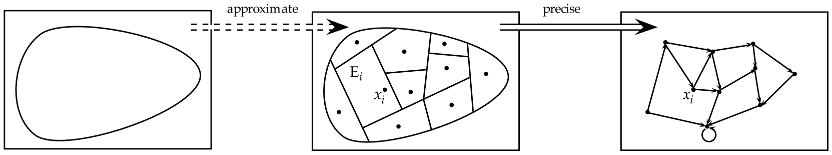

Figure 1: Two-step abstraction procedure: from a general `ldt`-MP to f.g. `ldt`-MP, to abstract `ldt`-MC

Lemma 1. *For any $n \in N_0$ the following inequality holds true:*

$$\|\mathsf{P} - \tilde{\mathsf{P}}\|_{\mathscr{F}_n} \le n \cdot \|P - \tilde{P}\|_{\tilde{\mathscr{E}}}.$$

We are now ready to state the main result, which deals with the approximate BLTL model-checking over a general `ldt`-MP using a finite `ldt`-MC abstraction obtained via a f.g. `ldt`-MP.

Theorem 3. *Let $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ be a given `ldt`-MP and let $\tilde{\mathfrak{D}} = (E, \tilde{\mathscr{E}}, \tilde{P}, \Sigma, \mathsf{L})$ be its finitely generated abstraction, and let $\hat{\mathfrak{D}}$ be the quotient `ldt`-MC of $\tilde{\mathfrak{D}}$. Then, for any $x \in E$ and $\varphi \in \mathscr{S}$,*

$$\left|\mathsf{P}_x\left(\mathtt{Sat}_\Omega(\varphi)\right) - \hat{\mathsf{P}}_{\mathtt{I}(x)}\left(\mathtt{Sat}_{\hat{\Omega}}(\varphi)\right)\right| \le \mathsf{H}(\varphi) \cdot \|P - \tilde{P}\|_{\tilde{\mathscr{E}}}\,.$$

Theorem 3 states that any BLTL probabilistic model-checking problem over an `ldt`-MP can be approximately solved using an appropriate `ldt`-MC abstraction. The derived error bounds are clearly useful only for bounded-horizon specifications φ, whereas for infinite-horizon specifications the distance between kernels P and $\tilde{P}$ in general cannot be employed to control the error. Still, we show in Section 4.1 that for some infinite-horizon specifications bounds on the error can be derived.

The results in Theorem 2 and Theorem 3 can be related to notions of precise and approximate bisimulation, respectively, which have been introduced for `ldt`-MC e.g. in [9].

Let us now focus on the bounded-horizon case $\mathsf{H}(\varphi) < \infty$ and provide an explicit construction of finitely-generated approximations for a given `ldt`-MP $(E, \mathscr{E}, P, \Sigma, \mathsf{L})$. We also show how to upper-bound the distance between kernels. This procedure has two ingredients: the choice of the finite σ-algebra $\tilde{\mathscr{E}}$ and the choice of the corresponding kernel $\tilde{P}$. Consider a finite collection of non-empty $\mathscr{E}$-measurable sets $(E_1, \dots, E_N)$, such that $E_i \cap E_j = \emptyset$ for all $i \neq j$ and $E = \bigcup_{i=1}^N E_i$, and such that for any index $i \in \overline{1, N}$ it holds that $\mathsf{L}|_{E_i} \equiv const$. We define

$$\tilde{\mathscr{E}} = \sigma(E_1, \dots, E_N) \tag{3.3}$$

to be the σ-algebra generated by this partition. To introduce the kernel $\tilde{P}$, we choose representative points $x_i \in E_i$ and define

$$\tilde{P}(x, A) := \sum_{i=1}^{n} 1_{E_i}(x) P(x_i, A) \tag{3.4}$$

for any set $A \in \tilde{\mathscr{E}}$. Note that $\tilde{P}$ given by (3.4) is uniquely determined by the matrix with entries $\tilde{p}_{ij} := P(x_i, E_j)$ (cf. (3.2)). We call a collection $(E_i, x_i)_{i=1}^N$ defined as above a *tagged partition* of the `ldt`-MP $(E, \mathscr{E}, P, \Sigma, \mathsf{L})$. Note that any tagged partition $(E_i, x_i)_{i=1}^N$ generates the pair $(\tilde{\mathscr{E}}, \tilde{P})$ by formulae (3.3), (3.4), hence for a given `ldt`-MP it defines uniquely its finitely generated abstraction $(E, \tilde{\mathscr{E}}, \tilde{P}, \Sigma, \mathsf{L})$.

3.3 Bounds on the distance between kernels

1. General BLTL specifications. Let us now discuss how to find upper bounds on the distance $\|P - \tilde{P}\|_{\tilde{\mathscr{E}}}$, and when is it possible to control it by choice of the tagged partition $(E_i, x_i)_{i=1}^N$. We define $\kappa_i(\tilde{P}) := \sum_{j=1}^N \sup_{x \in E_i} |P(x, E_j) - P(x_i, E_j)|$, $i \in \overline{1, N}$. The next proposition gives bounds on $\|P - \tilde{P}\|_{\tilde{\mathscr{E}}}$ in terms of κ_i.

Proposition 2. *For any tagged partition $(E_i, x_i)_{i=1}^N$:*

$$\|P - \tilde{P}\|_{\tilde{\mathscr{E}}} \le \max_{i \in \overline{1,N}} \kappa_i(\tilde{P}). \tag{3.5}$$

Although bounds in (3.5) are explicit and do not require any assumptions on the model, it may be impractical to find them and to control them by tuning the partition. Due to this reason, let us restrict our attention to the important class of integral kernels.

Assumption 1. *Let (E, ρ) be a metric space and $\mathscr{E}$ be its Borel σ-algebra. Assume that P is an integral kernel, i.e. that there exists a σ-finite* basis measure *μ on $(E, \mathscr{E})$ and a jointly measurable function $p : E \times E \to \mathbb{R}_+$ such that $P(x, \mathrm{d}y) = p(x, y)\mu(\mathrm{d}y)$, i.e. $P(x, A) = \int_A p(x, y)\mu(\mathrm{d}y)$ for any $x \in E$, $A \in \mathscr{E}$.*

To make our results sharper, we need to assume that the density p satisfies certain Lipschitz-like conditions. The work in [1] raised uniform Lipschitz continuity assumptions in order to control the bound for the computation of value functions characterizing probabilistic safety. This assumption was further relaxed in [20] into local Lipschitz continuity. In this work we further generalize the latter assumption as follows.

Assumption 2. *Under Assumption 1, for some tagged partition $(E_i, x_i)_{i=1}^N$ there exist measurable functions $\lambda_i : E \to \mathbb{R}_+$ such that $\Lambda_i := \int_E \lambda_i(y)\mu(\mathrm{d}y) < \infty$ for all $i \in \overline{1, N}$, and such that*

$$|p(x', y) - p(x'', y)| \le \lambda_i(y)\rho(x', x'') \quad \forall x', x'' \in E_i, \quad \forall y \in E.$$

Proposition 3. *If Assumption 2 is satisfied, then for any index $i \in \overline{1, N}$ it holds that $\kappa_i \le \Lambda_i \delta_i$ where $\delta_i = \mathrm{diam}(E_i)$.*

The latter result together with Theorem 3 leads to:

Corollary 1. *If Assumption 2 is satisfied, then for any $x \in E$ and $\varphi \in \mathscr{S}$ the following bound holds true:*

$$\left|\mathsf{P}_x\left(\mathtt{Sat}_\Omega(\varphi)\right) - \hat{\mathsf{P}}_{\mathtt{I}(x)}\left(\mathtt{Sat}_{\hat{\Omega}}(\varphi)\right)\right| \le \mathsf{H}(\varphi) \cdot \max_{i \in \overline{1,N}} \Lambda_i \delta_i.$$

One can notice that the bounds provided in Proposition 3 are similar in shape to those in [20, Theorems 4,6]. This further allows tailoring the sequential and adaptive gridding algorithm in [20, Section V] to our case. Indeed, whenever a precision level ε is fixed, one can start with some partition $(E_i, x_i)_{i=1}^N$ satisfying Assumption 2 and further refine it until the abstraction error becomes smaller than ε. If each refinement reduces the diameter of the partition sets at least by the factor of 2, then clearly the maximum number of refinements necessary to reach the precision level can be upper-bounded (for details see [20]).

It is important to note that results in Proposition 3 hold only if the state space E is bounded with respect to its metric ρ. Indeed, the condition on the finite cardinality of the partition N

that is used for the construction of the *finite* quotient ldt-MC implies that $\mathrm{diam}(E) \le \sum_{i=1}^N \delta_i < \infty$. There are two ways to cope with this restriction: first of all, one can always transform the original metric into an equivalent bounded one, for example $\rho'(x,y) := \frac{\rho(x,y)}{1+\rho(x,y)}$. The transformation of the metric leads to a change in the functions λ_i in Assumption 2: one shall further look for conditions on the original kernel in order to assure that the corresponding integrals Λ_i are bounded. Alternatively, one can introduce an additional assumption on the kernel P in order to deal with unbounded state spaces, as follows. The idea is that the original state space can be approximated with a bounded set, say B_ε, with any precision level $\varepsilon > 0$. More precisely:

Assumption 3. *Under Assumption 1, assume that there exists $\lambda \in \mathbb{R}$ such that for any points $x', x'', y \in E$ it holds that*

$$|p(x',y) - p(x'',y)| \le \lambda \cdot \rho(x',x''),$$

and for any $\varepsilon > 0$ there exists a bounded set B_ε such that:

$$\begin{aligned} P(x, B_\varepsilon^c) &\le \varepsilon, \quad \forall\, x \in E, \\ |p(x',y) - p(x'',y)| &\le \varepsilon, \quad \forall\, x', x'' \in B_\varepsilon^c,\ y \in B_\varepsilon. \end{aligned}$$

Proposition 4. *If Assumption 3 is satisfied with an $\varepsilon > 0$, let us consider $(E_i, x_i)_{i=1}^N$ to be a tagged partition with $E_N = B_\varepsilon^c$. Denote $\delta = \max_{i\in\overline{1,N-1}} \delta_i$, then*

$$\|P - \tilde{P}\|_{\tilde{\mathscr{E}}} \le \varepsilon + \frac{1}{2}\max\left\{\varepsilon \cdot \mu(B_\varepsilon), \lambda\delta \cdot \mu(B_\varepsilon)\right\}. \tag{3.6}$$

Proposition 4 allows reaching any desired precision only if there is a choice of B_ε such that $\lim_{\varepsilon\to 0} \varepsilon \cdot \mu(B_\varepsilon) = 0$. Indeed, in such case it is possible to make ε and $\varepsilon \cdot \mu(B_\varepsilon)$ in (3.6) as small as needed, and then to further construct a partition of B_ε in such a way that $\lambda\delta\mu(B_\varepsilon) < \varepsilon \cdot \mu(B_\varepsilon)$ by tuning δ appropriately.

2. Special case: bounded-horizon probabilistic safety. Let us now tailor the results above for the important case of bounded-horizon safety. Consider an ldt-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ with $\Sigma = \{\alpha, \beta\}$. As discussed earlier, we can formulate the probabilistic safety problem via the BLTL formula $\Phi_n = \Box^{\le n}\alpha$, which allows the application of the results above. However, notice that in such a case one has to partition the whole state space, whereas it is known from [1, 20] that it is sufficient to partition only the safe set $A = \mathsf{L}^{-1}(\alpha)$. Let us show how this problem can be studied in the new framework – in other words, below we consider a *formula-dependent* abstraction technique for the safety as a special case of the formula-free one presented above.

Let $\Delta \notin E$ be some auxiliary state introduced to represent set A^c and define a new ldt-MP $\mathfrak{D}' = (E', \mathscr{E}', P', \Sigma, \mathsf{L}')$, where $E' = A \cup \{\Delta\}$ and $\mathscr{E}' = \sigma(\mathscr{E}_A, \{\Delta\})$, where $\mathscr{E}_A = \{B \subseteq A : B \in \mathscr{E}\}$ is the subspace σ-algebra of A. The kernel P' is given by $P'(x,B) = P(x,B)$ for $x \in A$, $B \in \mathscr{E}_A$, and $P'(x,\Delta) = P(x,A^c)$ for $x \in A$, and $P'(\Delta,\{\Delta\}) = 1$. Finally, $\mathsf{L}'(x) = \alpha$ for $x \in A$ and $\mathsf{L}'(\Delta) = \beta$.

Let us denote by $(\Omega', \mathscr{F}', \mathsf{P}')$ the probability space of $\mathfrak{D}'$, then

$$\mathsf{P}_x\left(\mathtt{Sat}_\Omega(\mathfrak{L}(\Phi_n))\right) = \mathsf{P}'_x\left(\mathtt{Sat}_{\Omega'}(\mathfrak{L}(\Phi_n))\right)$$

for all $x \in A$ and $n \in \mathbb{N}_0$, as it follows from the integral representation of the safety probability [1]. As a result, rather than doing verification of safety over $\mathfrak{D}$ we can do this over a simpler ldt-MP $\mathfrak{D}'$, and still obtain the same result. Let us show how Assumption 2 changes in such a case for $\mathfrak{D}$.

First of all, in order to solve the safety problem over $\mathfrak{D}'$ we need to partition E', which reduces to partitioning only A. Suppose now that Assumption 1 holds for $\mathfrak{D}$ and that $\mathrm{diam}(A) < \infty$. We define a metric ρ' over E' by: $\rho'(x,y) = \rho(x,y)$ if $x,y \in A$, and $\rho'(x,\Delta) = \mathrm{diam}(A) + 1$ if $x \in A$. We further derive a new basis measure μ' from μ as $\mu'(B) = \mu(B)$ if $B \in \mathscr{E}_A$ and $\mu'(\{\Delta\}) = 1$. As a result, from $P(x,\mathrm{d}y) = p(x,y)\mu(\mathrm{d}y)$ and $P'(x,\mathrm{d}y) = p'(x,y)\mu'(\mathrm{d}y)$ we obtain the following density p': $p'(x,y) = p(x,y)$ if $x,y \in A$, and $p'(x,\Delta) = P'(x,\{\Delta\})$. Thus, if $\mathfrak{D}$ satisfies Assumption 1 with parameters (ρ,μ,p) then $\mathfrak{D}'$ also satisfies it with parameters (ρ',μ',p') defined above. Thus, for Assumption 2 to hold for $\mathfrak{D}'$, the *original* process $\mathfrak{D}$ only has to satisfy the following relaxed version of this assumption.

Assumption 4. *Under Assumption 1 for $\mathfrak{D}$: for some partition $(E_i,x_i)_{i=1}^N$ of A there exist measurable functions $\lambda_i' : A \to \mathbb{R}_+$ such that $\Lambda_i' := \int_A \lambda_i'(y)\mu(\mathrm{d}y) < \infty$ for all $i \in \overline{1,N}$, and such that*

$$|p(x',y) - p(x'',y)| \le \lambda_i'(y)\rho(x',x'') \quad \forall x',x'' \in E_i, \quad \forall y \in A.$$

To summarize, the discussion above suggests that in order to solve a probabilistic safety problem over the original ldt-MP $\mathfrak{D}$, one can partition only the safe set A and lump A^c into a single state Δ. Furthermore, Assumption 4 is sufficient to yield bounds as in Corollary 1. Note that although these bounds would hold for any bounded language over $\mathfrak{D}'$, in general only the safety specification over $\mathfrak{D}'$ can be related to that over $\mathfrak{D}$.

3.4 Connections with the literature

1. Discrete-time Stochastic Hybrid Systems. Above we have shown how to bound the quantity $\|P - \tilde{P}\|_{\tilde{\mathscr{E}}}$ needed for the abstraction technique, and which assumptions are sufficient to control this bound. The results have been stated in measure-theoretical terms, for instance dealing with abstract basis measures and densities. In order to further elucidate the meaning of these results over concrete models, as well as to highlight their connection with recent literature, this section focuses on models expressed as discrete-time SHS. We say that an ldt-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ is an ldt-SHS if $E = \bigcup_{q\in Q}\{q\} \times D_q$, where D_q are Borel subsets of $\mathbb{R}^{m_q}$ and $m_q \in \mathbb{N}$ for all $q \in Q$. Q is a finite set of *modes* (or locations) and E is a hybrid state space.

We can endow E with a disjoint union topology and choose $\mathscr{E}$ to be its Borel σ-algebra. In other words, any $B \in \mathscr{E}$ can be decomposed uniquely as $B = \bigcup_{q\in Q}\{q\} \times B_q$, where B_q is a Borel subset of D_q. It is thus natural to define a basis measure μ on $(E,\mathscr{E})$ by $\mu(B) = \sum_{q\in Q}\ell^{m_q}(B_q)$, where ℓ^m stands for the Lebesgue measure on $\mathbb{R}^m$. It is common to define the transition kernel of ldt-SHS through its hybrid components [1] as

$$P((q,c),\{q'\}\times \mathrm{d}c') = \begin{cases} T_q(q'|(q,c))T_x(\mathrm{d}c'|(q,c)), & q' = q, \\ T_q(q'|(q,c))T_r(\mathrm{d}c'|(q,c),q'), & q' \ne q, \end{cases}$$

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, T_x is given in [1]. Let us now show how the density p can be constructed given the co-product basis measure μ as above, and densities t_x, t_r of T_x, T_r respectively:

$$p((q,c),(q',c')) = \begin{cases} T_q(q'|(q,c))t_x(c'|(q,c)), & q' = q, \\ T_q(q'|(q,c))t_r(c'|(q,c),q'), & q' \ne q. \end{cases}$$

We have just explicitly embedded the densities of a ldt-SHS into the general measure-theoretical framework we use in this contribution. In particular, we obtain that the Lipschitz assumption on T_q, t_x, t_r as per [20, Assumption 2] is indeed a special case of Assumption 4 of this contribution, as it follows from

$$|f(x_1)g(x_1) - f(x_2)g(x_2)| \leq \|f\| \cdot |g(x_1) - g(x_2)| \\ + \|g\| \cdot |f(x_1) - f(x_2)|,$$

for any functions $f, g : E \to \mathbb{R}$ and any $x_1, x_2 \in E$. Thus, with focus on the safety problem, Corollary 1 under Assumption 4 in this contribution implies [20, Theorem 6] as a special case, where functions λ'_i have a piecewise-constant shape.

2. Specifications expressed as DFA. Let us discuss the verification of DFA specifications over a general `ldt`-MP. A DFA is a tuple $\mathscr{A} = (Q, q_0, \Sigma, \mathtt{t}, F)$ where Q is a finite set of states, $q_0 \in Q$ is the initial state, Σ is a finite alphabet, $\mathtt{t} : Q \times \Sigma \to Q$ is a transition function and $F \subseteq Q$ is a set of accepting states. Given an infinite word $\pi \in \mathfrak{S}$, the corresponding trajectory $\eta \in Q^{\mathbb{N}_0}$ is defined by $\eta_0 = q_0$ and $\eta_{i+1} = \mathtt{t}(\eta_i, \pi_i)$ for all $i \in \mathbb{N}_0$. For any $n \in \mathbb{N}_0 \cup \{\infty\}$ we define the accepting language $\mathfrak{L}_n(\mathscr{A}) \subseteq \mathfrak{S}$ as follows: the word $\pi \in \mathfrak{S}$ is n-accepted by $\mathscr{A}$ if its corresponding trajectory $\eta_i \in F$ for some $i \leq n$. Clearly, in case $n = \infty$ we obtain the usual accepting condition for DFA, else $\mathfrak{L}_n(\mathscr{A}) \in \mathscr{S}$ and $\mathsf{H}(\mathfrak{L}_n(\mathscr{A})) \leq n$. As a result, the model-checking problem is well-posed and we can apply a formula-free abstraction in the case $n < \infty$. Alternatively, we can follow the formula-dependent approach given in [2], which we now recall and compare. Given a `ldt`-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ we define a new `ldt`-MP $\mathfrak{D}^{\mathscr{A}} = (E^{\mathscr{A}}, \mathscr{E}^{\mathscr{A}}, P^{\mathscr{A}}, \{\alpha, \beta\}, \mathsf{L}^{\mathscr{A}})$ as follows: $E^{\mathscr{A}} = Q \times E$ and $\mathscr{E}^{\mathscr{A}}$ is the corresponding product σ-algebra. Also $\mathsf{L}^{\mathscr{A}}(q, x) = \beta$ if $q \in F$ and $\mathsf{L}^{\mathscr{A}}(q, x) = \alpha$ otherwise. Finally:

$$P^{\mathscr{A}}((q, x), \{q'\} \times \mathrm{d}x') = 1_{\mathtt{t}(q, \mathsf{L}(x))}(q') \cdot P(x, \mathrm{d}x'). \tag{3.7}$$

In this case, for any $n \in \mathbb{N}_0 \cup \{\infty\}$ [2]

$$\mathsf{P}_x(\mathtt{Sat}_\Omega(\mathfrak{L}_n(\mathscr{A}))) = 1 - \mathsf{P}^{\mathscr{A}}_x(\mathtt{Sat}_{\Omega^{\mathscr{A}}}(\mathfrak{L}(\Box^{\leq n+1}\alpha))), \tag{3.8}$$

where $(\Omega^{\mathscr{A}}, \mathscr{F}^{\mathscr{A}}, \mathsf{P}^{\mathscr{A}})$ denotes the probability space of $\mathfrak{D}^{\mathscr{A}}$. Such an object is called the *product* between the `ldt`-MP and the DFA, and is alternatively denoted by $\mathfrak{D} \otimes \mathscr{A}$ [2].

To do verification of $\mathscr{A}$ we can either leverage a formula-free abstraction technique to find the left-hand side of (3.8), or a formula-dependent one to evaluate safety in its right-hand side. Note, however, that in the latter case we still have to partition the whole original state space E as α corresponds to $(Q \setminus F) \times E$.

In order to compare the two techniques in more detail, let us suppose that Assumption 2 holds true for $\mathfrak{D}$. By Corollary 1, we have that the error introduced by the formula-free abstraction is equal to $\epsilon_1 = n \cdot \max_i \Lambda_{i,\mathfrak{D}} \delta_i$. Now, with focus on the formula-dependent approach, notice that Assumption 2 for $\mathfrak{D}$ implies Assumption 4 for $\mathfrak{D}^{\mathscr{A}}$. In particular, without having any additional information but Assumption 2 over $\mathfrak{D}$, we can only say that $\lambda'_{i,\mathfrak{D}^{\mathscr{A}}}(q, y) = \lambda_{i,\mathfrak{D}}(y)$ for all pairs (q, y) that belong to the safe set $A = (Q \setminus F) \times E$. As a result, we obtain that

$$\Lambda'_{i,\mathfrak{D}^{\mathscr{A}}} = \int_A \lambda'_{i,\mathfrak{D}^{\mathscr{A}}}(q, y) \mu^{\mathscr{A}}(\{q\} \times \mathrm{d}y) \\ = \sum_{Q \setminus F} \int_E \lambda_i(y) \mu(\mathrm{d}y) = \#(Q \setminus F) \cdot \Lambda_i,$$

where $\#$ denotes the cardinality of a set and $\mu^{\mathscr{A}}$ is a product measure of the measure μ on E and the counting measure on Q. Hence, the error of the formula-dependent approach is

$$\epsilon_2 = (n+1) \cdot \max_i \Lambda'_{i,\mathfrak{D}^{\mathscr{A}}} \delta_i \geq \#(Q \setminus F) \cdot \epsilon_1.$$

Such an error is in most cases larger than the error introduced by the formula-free abstraction – this highlights yet another advantage of the proposed approach. If we fix the precision level and further refine the partition for the formula-dependent abstraction to reach this precision, whenever $\mathfrak{D}$ is an `ldt`-SHS the cardinality of the corresponding finite abstraction will be $\mathscr{O}(\#(Q \setminus F)^{m+1})$ bigger than that of the formula-free abstraction, where m is the largest dimension of the continuous components of the hybrid state space. To further elucidate this scalability assessment, we provide a concrete example in Section 5.

Clearly, we have supposed that no additional information about the structure of the original system is used. Although such assumption is relevant in many applications where e.g. it is not possible to compute Lipschitz-like functions λ adaptively for any new partition, it motivates exploiting the structure of the DFA $\mathscr{A}$ which possibly may help reducing ϵ_2 – for example one can try using methods from [12]. However, it is by no means clear whether such attempts would in general enable overcoming the factor $\#(Q \setminus F)$, which can be a large integer.

4. FURTHER APPLICATIONS

The previous section has shown how to build a formula-free finite abstraction tailored to the goal of probabilistic BLTL model-checking against `ldt`-MP models. Let us now discuss how the technique we used to prove the main result in Theorem 3 can lead to other important applications. This technique is given in Lemma 1, which relates the *one-step error* $\|P - \tilde{P}\|_{\tilde{\mathscr{E}}}$ to the *final error* $\|\mathsf{P} - \tilde{\mathsf{P}}\|_{\tilde{\mathscr{F}}_n}$ by showing that the one-step error propagates in time at most linearly. This idea essentially extends a similar method developed for approximate model-checking of particular PCTL specifications, such as safety [1]. One of the possible advantages of the current approach is that the one-step error can be related not only to the final error for safety and BLTL, but also to the final error for other verification problems: such a relation allows working on improvements of the one-step error, rather than on computing the final error for each verification problem. The new results on the one-step error would be then applicable to all verification problems where the relation between the one-step and final errors is known.

To further emphasize the usefulness of the proposed approach, let us recapitulate that the one-step error has been successfully related to the final error in the BLTL model-checking of an `ldt`-MP. Below we show other examples of verification problems where it is as well worth applying the newly introduced approach based on the one-step error. Note that in each such example this error can be bounded using any of Propositions 2, 3 and 4.

4.1 Infinite-horizon reach-avoid

As already mentioned, the formula-free abstraction technique presented in this work is not directly applicable to general unbounded time instances, since the one-step error cannot be linearly accrued over an infinite horizon. However, the abstraction described above can still be applied over particular instances of infinite-horizon problems.

Consider an `ldt`-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$ where $\Sigma = \{\alpha, \beta, \gamma\}$. We are interested in the infinite-horizon reach-avoid problem, which can be stated using the “unbounded Until” operator from LTL or PCTL. Such an operator can be easily expressed using a disjunction over BLTL formulae $\alpha \mathsf{U} \beta := \bigvee_{n=0}^{\infty} \alpha \mathsf{U}^{\leq n} \beta$. Note that the infinite disjunction is needed for the definition of U, which is not part of BLTL where only finite disjunctions are allowed. Moreover, the corresponding accepting language $\mathfrak{L}(\alpha \mathsf{U} \beta)$ is not bounded, being an ω-regular language. For this specification we define the following value function $w(x) := \mathsf{P}_x(\mathtt{Sat}_\Omega(\mathfrak{L}(\alpha \mathsf{U} \beta)))$, which is known to be a solution of the Bellman equation [18]

$$w(x) = 1_B(x) + 1_A(x)Pw(x), \tag{4.1}$$

where $A := \mathsf{L}^{-1}(\alpha)$ and $B := \mathsf{L}^{-1}(\beta)$. Let $r := \sup_{x\in A} P(x,A)$. Whenever $r < 1$, equation (4.1) admits a unique solution [18].

Let now $\tilde{\mathfrak{D}} = (E, \tilde{\mathscr{E}}, \tilde{P}, \Sigma, \mathsf{L})$ be some finitely-generated abstraction of $\mathfrak{D}$ and let us denote the corresponding value function by $\tilde{w}(x) := \tilde{\mathsf{P}}_x(\mathtt{Sat}_\Omega(\mathfrak{L}(\alpha \mathsf{U} \beta)))$.

THEOREM 4. *If $r < 1$ the following bound holds true:*

$$\|w - \tilde{w}\| \le \frac{\|P - \tilde{P}\|_{\tilde{\mathscr{E}}}}{1 - r}. \tag{4.2}$$

As in the case of the safety problem in Section 3.3, to solve the reach-avoid problem one can consider a version $\mathfrak{D}'$ of $\mathfrak{D}$ where states corresponding to α and γ are lumped into two single states. Owing to this simpler structure, it is easier for $\mathfrak{D}'$ to satisfy assumptions needed to control the one-step error.

4.2 Rewards

One useful extension of the results in Section 3 stems out of the following observation. Let $\mathfrak{D} = (E, \mathscr{E}, P)$ be some `dt`-MP and let $\tilde{\mathfrak{D}} = (E, \tilde{\mathscr{E}}, \tilde{P})$ be its finitely-generated abstraction. Let $(\Omega, \mathscr{F}, \mathsf{P})$ and $(\Omega, \tilde{\mathscr{F}}, \tilde{\mathsf{P}})$ represent the canonical probability spaces of $\mathfrak{D}$ and $\tilde{\mathfrak{D}}$ respectively, and let E_x and $\tilde{\mathsf{E}}_x$ denote the corresponding expectations.

THEOREM 5. *If $\xi : \Omega \to \mathbb{R}$ is $\tilde{\mathscr{F}}_n$-measurable then for any $x \in E$*

$$\left|\mathsf{E}_x[\xi] - \tilde{\mathsf{E}}_x[\xi]\right| \le n \cdot \|\xi\| \cdot \|P - \tilde{P}\|_{\tilde{\mathscr{E}}}$$

This fact can be applied to the approximate computation of reward functionals (cf. [5, Chapter 10.5] for `dt`-MC), since they are real-valued maps ξ defined over trajectories ω. The $\tilde{\mathscr{F}}_n$-measurability assumption requires the reward ξ to depend only on the first $n+1$ coordinates of ω, i.e. $\omega_0, \dots, \omega_n$. Such an assumption holds for a wide range of problems, for instance in finance [25].

Let us show how to approximately compute the expected value of rewards in the following example. The first hitting time of a set $A \in \mathscr{E}$ is a defined by $\tau_A(\omega) := \inf\{n \ge 0 : \omega_n \in A\}$. For a function $g \in \mathrm{b}\mathscr{E}$, set $A \in \mathscr{E}$ and $n \in \mathbb{N}_0$, let us define a reward of the following form: $\xi^n_{g,A}(\omega) := \sum_{k=0}^{\tau_A \wedge n} g(\omega_k)$. For example, if $g = 1_B$ where $B \in \mathscr{E}$ is some set, then $\xi^n_{g,A}$ is the time that a trajectory ω spends in the set B prior to hitting set A and within the time epoch n. Let us show now how to approximately compute $\mathsf{E}_x[\xi^n_{g,A}]$ using finite abstractions. Let $(E_i, x_i)_{i=1}^n$ be a tagged partition that generates $\tilde{\mathfrak{D}}$ – since there is no labeling structure, we only require that $A \in \tilde{\mathscr{E}}$, which is equivalent to $A = \bigcup_{i\in I} E_i$ for some index set i. We cannot apply immediately Theorem 5 to $\xi^n_{g,A}$ as in general it is only $\mathscr{F}_n$-measurable rather than $\tilde{\mathscr{F}}_n$-measurable. Thus, we first approximate the original reward with an $\tilde{\mathscr{F}}_n$-measurable one: we define $\tilde{g}(x) := \sum_{i=1}^N 1_{E_i}(x) g(x_i)$, which is clearly $\tilde{\mathscr{E}}$-measurable, hence $\xi^n_{\tilde{g},A}$ is $\tilde{\mathscr{F}}_n$-measurable.

ASSUMPTION 5. *Let (E, ρ) be a metric space and $\mathscr{E}$ be its Borel σ-algebra. Denote $\delta = \max_{1\le i\le N} \mathrm{diam}(E_i)$ and assume that*

$$|g(x') - g(x'')| \le \eta \cdot \rho(x', x''), \quad \forall x', x'' \in E_i, \quad \forall i \in \overline{1, N},$$

for some constant $\eta > 0$.

THEOREM 6. *If Assumption 5 is satisfied then for any $x \in E$*

$$|\mathsf{E}_x[\xi^n_{g,A}] - \tilde{\mathsf{E}}_x[\xi^n_{\tilde{g},A}]| \le (n+1)\left(n \cdot \|g\| \cdot \|P - \tilde{P}\|_{\tilde{\mathscr{E}}} + \eta\delta\right).$$

5. CASE STUDY

To elucidate the techniques developed throughout this work, let us consider the following model. Let the state space E be the interval $[0, 10]$, endowed with a Euclidian metric, and let $\mathscr{E}$ be the corresponding Borel σ-algebra. We construct the transition kernel P as an integral kernel with a basis Lebesgue measure $\mu = \ell^1$ and a density p being a weighted sum of two components: $p(x, y) = w(x)p_1(y) + (1 - w(x))p_2(y)$. The weighting function w is chosen to be the relative distance to the center of the interval: $w(x) := \frac{1}{5}|x - 5|$. The function $p_1(y)$ corresponds to a truncated Gaussian distribution given by $p_1(y) := K \cdot \mathrm{e}^{-\frac{1}{2}(y-5)^2}$, and $p_2(y) = \frac{1}{10}$ corresponds to the uniform distribution. Here K is a normalization constant defined by $\int_E p_1(y)\mathrm{d}y = 1$, so that $K \approx 0.3989$. The shape of the density p suggests that the closer the current state to the center of the interval, the more impact the truncated Gaussian term has whereas if the current state is far from the center of the interval, the dynamics are affected by the uniform term in p. Note that the `dt`-MP $(E, \mathscr{E}, P)$ satisfies Assumption 1 by construction.

We introduce the alphabet $\Sigma = \{\alpha, \beta\}$ and the labeling map $\mathsf{L}(x) = \alpha$ if $x \in [4, 6]$ and $\mathsf{L}(x) = \beta$ otherwise. In order to build the formula-free abstraction of the `ldt`-MP $\mathfrak{D} = (E, \mathscr{E}, P, \Sigma, \mathsf{L})$, we fix a time horizon $T = 100$ and select the precision level to be equal to $\varepsilon = 0.1$. We are going to apply Proposition 3 to find the required size of the partition sets, so we need to check if Assumption 2 holds in this case. It holds that

$$\begin{aligned}|p(x', y) - p(x'', y)| &= |p_2(y) - p_1(y)| \cdot |w(x') - w(x'')| \\ &\le |p_2(y) - p_1(y)| \cdot |x' - x''|,\end{aligned}$$

so we can select $\lambda(y) = |p_2(y) - p_1(y)|$ as per Assumption 2, regardless of the choice of the partition. In this case it holds that $\Lambda := \int_0^{10} \lambda(y)\mathrm{d}y \approx 1.1422$.

Let us now consider some arbitrary partition $(E_i, x_i)_{i=1}^N$ of $\mathfrak{D}$ and let $\delta = \max_i \mathrm{diam}(E_i)$. It follows from Proposition 3 that the one-step error is given by $\Lambda\delta$ and the final error is equal to $T\Lambda\delta$. As a result, to reach the desired precision level we need to select a partition size $\delta \le \frac{\varepsilon}{T\Lambda}$ and the cardinality of the partition results in $N_1 = 10/\delta = 11422$. Let us emphasize that this partition leads to a `ldt`-MC that can be used for the model-checking of any linear temporal property with horizon T.

As a second example, in order to further clarify the statements made for formula-dependent abstraction techniques in Section 3.4, let us consider a particular specification given by a DFA $\mathscr{A} = (Q, q_0, \Sigma, \mathtt{t}, F)$ where $Q = \overline{0, M}$, $q_0 = 0$ and $F = \{M\}$. Further, let the transition function be given by $\mathtt{t}(q, \alpha) = q + 1$ if $q \notin F$, $\mathtt{t}(M, \alpha) = M$ and $\mathtt{t}(q, \beta) = 0$. Such an automaton expresses the BLTL formula $\Phi = \Diamond^{\le 100-M}\Box^{\le M}\alpha$ in the sense that $\mathfrak{L}_{M+1}(\mathscr{A}) = \mathfrak{L}(\Phi)$. For simplicity, let us select $M = 50$ from now on. As we have discussed, the formula-free abstraction yields an `ldt`-MC $\hat{\mathfrak{D}}$ with $N = 11422$ states and introduces the error $\epsilon_1 = 0.1$ over this specification. We can then verify the formula by solving a safety problem on $\hat{\mathfrak{D}}^{\mathscr{A}}$, which requires dealing with $C_1 = \#(Q \setminus F)N + 1 = 50 \cdot 11422 + 1 \approx 6 \cdot 10^5$ states.

As an alternative, we can do formula-dependent abstraction and try to solve the safety problem over $\mathfrak{D}^{\mathscr{A}}$. Based on the discussion in Section 3.4, in order to have the same precision level of 0.1, we need to take partition sets that are of size 50 times smaller, compared to the formula-free abstraction. Since the dimension of the problem is $m = 1$, this results in a cardinality of the obtained Markov Chain equal to $C_2 \approx 50^2 \cdot N_1 = 50 \cdot C_1 \approx 3 \cdot 10^6$ states, which is a number substantially larger than C_1 and as such possibly critical for computations. As a remark, if the dimension of the state space would be bigger, say $m = 2$, then we would have that $C_2 \approx 2500 C_1$. If we increased the parameter M (e.g., $M = 70$), as a result we would obtain $C_2 \approx 5 \cdot 10^3 C_1$.

6. CONCLUSIONS

This contribution has presented a formula-free approach for approximate finite abstractions of SHS tailored to BLTL model checking. The work has shown that in a number of problems, for example when dealing with specifications expressed as automata, this approach can mitigate scalability issues related to formula-dependent abstractions – this motivates further need for more tailored and precise formula-dependent abstraction methods. The approach is based on the propagation of the difference in transition kernels introduced by the abstraction (the "one-step error") as a global error over the complete verification problem. Besides the application on formula-free abstractions over BLTL, this technique can also be used over probabilistic safety (PCTL) of SHS, and allows for extensions beyond BLTL.

Results on formula-free abstraction do not hold over general infinite-horizon problems, which calls for specific techniques to tackle these problems [24]. The authors are also looking into extensions of the developed techniques to the controlled case.

7. ACKNOWLEDGMENTS

The authors would like to thank Theo Buehler for the helpful discussions on measure theory, Joost-Pieter Katoen and Manuel Mazo Jr. for comments on model-checking, and Sadegh E.Z. Soudjani for help with the automaton expression for $\Diamond^{\leq n}\Box^{\leq m}$.

8. REFERENCES

[1] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16:624–641, 2010.

[2] A. Abate, J.-P. Katoen, and A. Mereacre. Quantitative automata model checking of autonomous stochastic hybrid systems. In *Proceedings of the 14th international conference on Hybrid Systems: Computation and Control*, HSCC '11, pages 83–92, New York, NY, USA, 2011. ACM.

[3] 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.

[4] C. Baier. *On algorithmic verification methods for probabilistic systems*. PhD thesis, Universität Mannheim, 1998.

[5] C. Baier and J.-P. Katoen. *Principles of model checking*. The MIT Press, 2008.

[6] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control: The discrete time case*, volume 139. Academic Press, 1978.

[7] V. I. Bogachev. *Measure theory. Vol. I, II*. Springer-Verlag, Berlin, 2007.

[8] C.G. Cassandras and J. (Eds.) Lygeros. *Stochastic hybrid systems*, volume 24. CRC Press, 2007.

[9] A. D'Innocenzo, A. Abate, and J.P. Katoen. Robust PCTL model checking. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control*, pages 275–286. ACM, 2012.

[10] G.B. Folland. *Real analysis*. Pure and Applied Mathematics. John Wiley & Sons Inc., New York, second edition, 1999.

[11] A. Girard and G. J. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782 –798, 2007.

[12] H. Hermanns, B. Wachter, and L. Zhang. Probabilistic CEGAR. In *Computer Aided Verification*, pages 162–175. Springer, 2008.

[13] 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.

[14] M. Huth. On finite-state approximants for probabilistic computation tree logic. *Theoretical Computer Science*, 346(1):113–134, 2005.

[15] S. Jha, E. Clarke, C. Langmead, A. Legay, A. Platzer, and P. Zuliani. A Bayesian approach to model checking biological systems. In *Computational Methods in Systems Biology*, pages 218–234. Springer, 2009.

[16] N.H. Lâm and L. Vân. Measure of infinitary codes. *Acta Cybernetica*, 11(3):127–137, 1994.

[17] G.J. Pappas. Bisimilar linear systems. *Automatica*, 39(12):2035–2047, 2003.

[18] F. Ramponi, D. Chatterjee, S. Summers, and J. Lygeros. On the connections between PCTL and dynamic programming. In *Proceedings of the 13th ACM international conference on Hybrid Systems: Computation and Control*, pages 253–262, 2010.

[19] D. Revuz. *Markov chains*. North-Holland Publishing, Amsterdam, second edition, 1984.

[20] S. Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Proceedings of the 2011 Eighth International Conference on Quantitative Evaluation of SysTems*, QEST '11, pages 59–68, Washington, DC, USA, 2011. IEEE Computer Society.

[21] S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961, 2010.

[22] P. Tabuada. *Verification and control of hybrid systems: A symbolic approach*. Springer Verlag, New York, 2009.

[23] I. Tkachev and A. Abate. On infinite-horizon probabilistic properties and stochastic bisimulation functions. In *2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, pages 526–531, 2011.

[24] 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*, HSCC '12, pages 227–236, New York, NY, USA, 2012. ACM.

[25] P. Wilmott, S. Howison, and J. Dewynne. *The mathematics of financial derivatives: a student introduction*. Cambridge University Press, 1995.