

Decisiveness of Stochastic Systems and its Application to Hybrid Models*

Patricia Bouyer

LSV, CNRS & ENS Paris-Saclay, Université Paris-Saclay, France

Thomas Brihaye

UMONS – Université de Mons, Belgium

Mickael Randour

F.R.S.-FNRS & UMONS – Université de Mons, Belgium

Cédric Rivière

UMONS – Université de Mons, Belgium

Pierre Vandenhove

F.R.S.-FNRS & UMONS – Université de Mons, Belgium

LSV, CNRS & ENS Paris-Saclay, Université Paris-Saclay, France

In [3], Abdulla et al. introduced the concept of decisiveness, an interesting tool for lifting good properties of finite Markov chains to denumerable ones. Later [10], this concept was extended to more general stochastic transition systems (STSs), allowing the design of various verification algorithms for large classes of (infinite) STSs. We further improve the understanding and utility of decisiveness in two ways.

First, we provide a general criterion for proving the decisiveness of general STSs. This criterion, which is very natural but whose proof is rather technical, (strictly) generalizes all known criteria from the literature. Second, we focus on stochastic hybrid systems (SHSs), a stochastic extension of hybrid systems. We establish the decisiveness of a large class of SHSs and, under a few classical hypotheses from mathematical logic, we show how to decide reachability problems in this class, even though they are undecidable for general SHSs. This provides a decidable stochastic extension of o-minimal hybrid systems [16, 22, 32].

1 Introduction

Hybrid and stochastic models. Various kinds of mathematical models have been proposed to represent real-life systems. In this article, we focus on models combining *hybrid* and *stochastic* aspects. We outline the main features of these models to motivate our approach.

The idea of *hybrid systems* originates from the urge to study systems subject to both *discrete* and *continuous* phenomena, such as digital computer systems interacting with analog data. These systems are transition systems with two kinds of transitions: continuous transitions, where some continuous variables evolve over time (e.g., according to a differential equation), and discrete transitions, where the system changes modes and variables can be reset. Much of the research about hybrid systems focuses on *non-deterministic* hybrid systems, i.e., systems modeling uncertainty by considering all possible behaviors (e.g., different possibilities for discrete transitions at a given time, arbitrary long time between discrete transitions). A typical question concerns the *safety* of such systems—if a system can reach an undesirable state, it is said to be *unsafe*; if not, it is *safe*—such a specification is called *qualitative*. However, this is limiting for two reasons. First, it does not take into account that some behaviors are more likely to

*Research supported by F.R.S.-FNRS under Grant n° F.4520.18 (ManySynth), and ENS Paris-Saclay visiting professorship (M. Randour, 2019). Pierre Vandenhove is an F.R.S.-FNRS Research Fellow. Mickael Randour is an F.R.S.-FNRS Research Associate.

occur than others. Second, risks cannot necessarily be avoided, and it is unrealistic to prevent undesirable outcomes altogether. Therefore, we want to make probabilities an integral part of our models, in order to be able to *quantify* the probability that they behave according to the specification. We thus consider the class of *stochastic hybrid systems* (SHSs, for short), hybrid systems in which a stochastic semantics replaces non-determinism.

Goals. Our interest lies in the formal analysis of continuous-time SHSs, and more specifically in *reachability* questions, i.e., concerning the likelihood that some set of states is reached. The questions we seek to answer are both of the qualitative kind (is some region of the state space *almost surely* reached, i.e., reached with probability 1?) and of the quantitative kind (what is the probability that some states are eventually reached?). Such questions are crucial, as verifying that a system works safely often reduces to verifying that some undesirable state of the system is never reached (or with a low probability), or that some desirable state is to be reached with high probability [8]. We want to give algorithms that decide, for an SHS \mathcal{H} and reachability property P , whether P is satisfied in \mathcal{H} . Such an endeavor faces multiple challenges; a first obvious one being that even for rather restricted classes of non-deterministic hybrid systems, reachability problems are undecidable [24, 25]. We want to define and consider a *class* of SHSs for which *some* reachability problems are decidable.

Methods and contributions. Our methodology consists of two main steps. In a first step, we follow the approach of Bertrand et al. [10]: we study general *stochastic transition systems* (STSs) through the *decisiveness* concept (Section 2). The class of STSs is a very versatile class of systems encompassing many well-known families of stochastic systems, such as Markov chains, generalized semi-Markov processes, stochastic timed automata, stochastic Petri nets, and stochastic hybrid systems. Decisiveness was introduced in [3] to study Markov chains, and extended to STSs in [10]. An STS is said to be *decisive* with respect to a set of states B if executions of the system almost surely reach B or a state from which B is unreachable. Decisive STSs benefit from useful properties that make possible the design of verification algorithms related to reachability properties. Our first contribution is to provide a **criterion to check the decisiveness of STSs** (Proposition 2.8), which generalizes the decisiveness criteria from [3, 10]. This generalization was mentioned as an open problem in [10].

In a second step, we focus on *stochastic hybrid systems*, which we introduce in Section 3.1. Reachability problems in hybrid systems are notoriously undecidable [24, 25]. Our contributions regarding SHSs are split into two parts.

First, we aim to use the decisiveness idea to get closer to the decidability frontier. Albeit desirable, the decisiveness of a class of SHSs is not sufficient to handle algorithmic questions about each SHS, as we need an effective way to apprehend their uncountable state space. In this regard, an often-used technique is to consider a *finite abstraction* of the system, that is, a finite partition of the state space that preserves the properties to be verified (a well-known example is the region graph for timed automata [5]). To find such an abstraction of SHSs, we borrow ideas from [16, 32]: we consider SHSs with *strong resets* (Section 3.2), a syntactic condition that decouples their continuous behavior from their discrete behavior. We show that **SHSs with adequately placed strong resets** (at least one per cycle of their discrete graph) *(i) have a finite abstraction* (Proposition 3.11), and *(ii) are decisive* (Proposition 3.9), which can be proved using our new criterion.

Second, in Section 3.3, we show, with strong resets, how to effectively compute a finite abstraction and use it to perform a reachability analysis of the original system. We proceed by assuming that the components of our systems are definable in an *o-minimal structure*. The main difficulty here is that “a

satisfactory theory of measure and integration seems to be lacking in the o-minimal context” [9]; in an o-minimal structure, the primitive function of a definable function is in general not definable, which complicates definability questions regarding probabilities. We therefore restrict the possible probability distributions, using properties of o-minimal structures to keep our class as large as possible. When the **theory** of the structure is **decidable** (as is the case for the ordered field of real numbers [43]), the **reachability problems are then decidable**. This provides a stochastic extension to the theory of o-minimal hybrid systems [32]. We study **qualitative** and **quantitative** problems. Due to space constraints, all the missing proofs and some technical details and additional results are omitted from this version and can be found in the full article [14].

Related work. Our results combine previous work on stochastic systems and hybrid systems. About stochastic systems, we build on the work of [3, 10]. Our work also takes inspiration from research about *stochastic timed automata*, a subclass of stochastic hybrid systems which already combines stochastic and timed aspects; model checking of stochastic timed automata has been studied in [7, 11, 12] and considered in the context of the decisiveness property in [18]. Fundamental results about the decidability of the reachability problem for hybrid systems can be found in [5, 24, 25]. The class of *o-minimal hybrid systems*, of which we introduce a stochastic extension, has been studied in [16, 22, 32].

The literature about SHSs often follows a more practical or numerical approach. A first introduction to the model was provided in [29]. An extensive review of the underlying theory and of many applications of SHSs is provided in books [19, 17], and a review of different possible semantics for this model is provided in [34]. Applications of SHSs are numerous: a few examples are air traffic management [40, 41], communication networks [26], biochemical processes [33, 42]. The software tool UPPAAL [20] implements a model of SHS similar to the one studied in this article and uses numerical methods to compute reachability probabilities, through numerical solving of differential equations and Monte Carlo simulation. Reachability problems have also been considered in an alternative semantics with *discrete time*; a numerical approach is for instance provided in [1, 2].

Notations. We write $\mathbb{R}^+ = \{x \in \mathbb{R} \mid x \geq 0\}$ for the set of non-negative real numbers. Let (Ω, Σ) be a measurable space. We write $\text{Dist}(\Omega, \Sigma)$ (or $\text{Dist}(\Omega)$ if there is no ambiguity) for the set of *probability distributions* over (Ω, Σ) . The complement of a set $A \in \Sigma$ is denoted by $A^c = \Omega \setminus A$. For $A \in \Sigma$, we say that two probability distributions $\mu, \nu \in \text{Dist}(\Omega)$ are (*qualitatively*) *equivalent on A* if for each $B \in \Sigma$, if $B \subseteq A$, $\mu(B) > 0$ if and only if $\nu(B) > 0$.

2 Decisiveness of Stochastic Transition Systems

In this section, we define *stochastic transition systems* (STSs, for short) as in [10]. We then describe the concept of *decisiveness*, first defined in the specific case of Markov chains [3], and then extended to STSs [10]. Decisive stochastic systems benefit from “nice” properties making their qualitative and quantitative analysis more accessible. The first contribution of our work is a new decisiveness criterion (Proposition 2.8), which generalizes existing criteria from the literature [3, 10]. It is an intuitive criterion, which was conjectured in [10] but could not be proved. We finish with a brief section on the notion of *abstraction* between STSs, which will be useful to apply our results to stochastic hybrid systems.

2.1 Stochastic Transition Systems and Decisiveness [10]

Definition 2.1 (Stochastic transition system). A *stochastic transition system* (STS) is a tuple $\mathcal{T} = (S, \Sigma, \kappa)$ consisting of a measurable space of *states* (S, Σ) , and a function $\kappa: S \times \Sigma \rightarrow [0, 1]$ (sometimes called *Markov kernel*) such that for every $s \in S$, $\kappa(s, \cdot)$ is a probability measure and for every $A \in \Sigma$, $\kappa(\cdot, A)$ is a measurable function.

The second condition on κ implies in particular that for a measurable set $B \in \Sigma$, the set $\{s \in S \mid \kappa(s, B) > 0\}$ is measurable.

We interpret STSs as systems generating runs, with a probability measure over these runs. We fix an STS $\mathcal{T} = (S, \Sigma, \kappa)$. From a state $s \in S$, a probabilistic transition is performed according to distribution $\kappa(s, \cdot)$, and the system resumes from one of the successor states; this process generates random sequences of states. A *run* of \mathcal{T} is an infinite sequence $s_0 s_1 s_2 \dots$ of states. To formally provide a probabilistic semantics to STSs, we define a probability measure over the set of runs of \mathcal{T} . From an initial distribution $\mu \in \text{Dist}(S)$, we can define in a classical way a unique probability measure $\text{Prob}_\mu^{\mathcal{T}}$ on the σ -algebra generated by all the *cylinders*, using Carathéodory's extension theorem.

Expressing properties of runs. We use standard LTL notations [39] to express properties of runs of \mathcal{T} . If $B, B' \in \Sigma$, we write in particular $\mathbf{F}B$ (resp. $\mathbf{F}_{<n}B$, $B' \mathbf{U}B$, $B' \mathbf{U}_{<n}B$, $\mathbf{G}B$, \mathbf{GFB}) for the set of runs that visit B at some point (resp. visit B in less than n steps, stay in B' until a first visit to B , stay in B' until a first visit to B in less than n steps, always stay in B , visit B infinitely often). We are interested in two kinds of reachability problems.

Definition 2.2 (Qualitative and quantitative reachability). Let $B \in \Sigma$ be a measurable set of target states, and $\mu \in \text{Dist}(S)$ be an initial distribution. The *qualitative reachability problems* consist in deciding whether $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}B) = 1$, and whether $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}B) = 0$. The *quantitative reachability problem* consists in deciding, given $\varepsilon, p \in \mathbb{Q}$ with $\varepsilon > 0$, whether $|\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}B) - p| < \varepsilon$.

Transforming probability distributions. Another useful way to reflect on STSs is as transformers of probability distributions on (S, Σ) .

Definition 2.3 (STS as a transformer). For $\mu \in \text{Dist}(S)$, its *transformation through \mathcal{T}* is the probability distribution $\Omega_{\mathcal{T}}(\mu) \in \text{Dist}(S)$ defined for $A \in \Sigma$ by

$$\Omega_{\mathcal{T}}(\mu)(A) = \int_{s \in S} \kappa(s, A) d\mu(s).$$

The value $\Omega_{\mathcal{T}}(\mu)(A)$ is the probability to reach A in one step, from the initial distribution μ .

Attractors. We will be particularly interested in the existence of *attractors* for STSs.

Definition 2.4 (Attractor). A set $A \in \Sigma$ is an *attractor for \mathcal{T}* if for every $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{F}A) = 1$.

While S is always an attractor for \mathcal{T} , we will later search for attractors with more interesting properties. The definition of *attractor* actually implies a seemingly stronger statement: an attractor is almost surely visited *infinitely often* from any initial distribution [10, Lemma 19].

Decisiveness. Before introducing decisiveness, we give the definition of an *avoid-set*: for $B \in \Sigma$, its *avoid-set* is written as $\tilde{B} = \{s \in S \mid \text{Prob}_{\delta_s}^{\mathcal{T}}(\mathbf{F}B) = 0\}$ (where δ_s is the Dirac distribution at s). The avoid-set \tilde{B} corresponds to the set of states from which runs almost surely stay out of B *ad infinitum*. One can show that the set \tilde{B} belongs to the σ -algebra Σ [10, Lemma 14]. We can now define the concept of decisiveness as in [10].

Definition 2.5 (Decisiveness). Let $B \in \Sigma$ be a measurable set. We say that \mathcal{T} is *decisive w.r.t. B* if for every $\mu \in \text{Dist}(S)$, $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B \vee \mathbf{F}\tilde{B}) = 1$.

Intuitively, the decisiveness property states that, almost surely, either B will eventually be visited, or states from which B can no longer be reached will eventually be visited.

Example 2.6 (Random walk). We consider the STS \mathcal{T} (random walk) from Figure 1. We want to find out whether \mathcal{T} is decisive w.r.t. $B = \{0\}$. We assume that the initial distribution is given by δ_1 (the Dirac distribution at 1). By the theory on random walks, if $\frac{1}{2} < p < 1$, the walk will almost surely diverge to ∞ . This entails that $\text{Prob}_{\delta_1}^{\mathcal{T}}(\mathbf{F}B) < 1$. Moreover, since $p < 1$, there is a path with positive probability from every state to 0, so $\tilde{B} = \emptyset$. Therefore, $\text{Prob}_{\delta_1}^{\mathcal{T}}(\mathbf{F}B \vee \mathbf{F}\tilde{B}) = \text{Prob}_{\delta_1}^{\mathcal{T}}(\mathbf{F}B) < 1$, which means that \mathcal{T} is not decisive w.r.t. B . If $p \leq \frac{1}{2}$, state 0 is almost surely reached from any state, that is, for any initial distribution $\mu \in \text{Dist}(S)$, we have that $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B) = 1$. Hence, in this case, STS \mathcal{T} is decisive w.r.t. B .




Figure 1: STS \mathcal{T} representing a random walk on \mathbb{N} .

A major interest of the decisiveness concept lies in the design of simple procedures for the qualitative and quantitative analysis of stochastic systems. Indeed, as presented in [3, 10], it allows, under effectiveness hypotheses, to verify qualitative properties or to compute arbitrary close approximations of the probabilities of various properties, like reachability, repeated reachability, and even ω -regular properties. The reader is referred to [10, Sections 6 & 7] for more details, but we briefly recall the approximation scheme for reachability properties in order to illustrate the usefulness of the decisiveness property. This scheme will be applied to a specific class of STSs in Section 3.3.

Let $B \in \Sigma$ be a measurable set and $\mu \in \text{Dist}(S)$ be an initial distribution. To compute an approximation of $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B)$, we define two sequences $(p_n^{\text{Yes}})_{n \in \mathbb{N}}$ and $(p_n^{\text{No}})_{n \in \mathbb{N}}$ such that for $n \in \mathbb{N}$,

$$p_n^{\text{Yes}} = \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}_{\leq n}B) \text{ and } p_n^{\text{No}} = \text{Prob}_{\mu}^{\mathcal{T}}(B^c \mathbf{U}_{\leq n}\tilde{B}).$$

These sequences are non-decreasing and converge respectively to $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B)$ and $\text{Prob}_{\mu}^{\mathcal{T}}(B^c \mathbf{U}\tilde{B})$. Observe moreover that for all $n \in \mathbb{N}$, we have that $p_n^{\text{Yes}} \leq \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B) \leq 1 - p_n^{\text{No}}$.

The main idea behind decisiveness of STSs lies in the following property [3, 10]: if \mathcal{T} is decisive w.r.t. B , then $\lim_{n \rightarrow \infty} p_n^{\text{Yes}} + p_n^{\text{No}} = 1$. Therefore, for any given $\varepsilon > 0$, for some n sufficiently large, $p_n^{\text{Yes}} \leq \text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B) \leq p_n^{\text{Yes}} + \varepsilon$. In situations where p_n^{Yes} and p_n^{No} can be effectively approximated arbitrarily closely and \mathcal{T} is decisive w.r.t. B , we can thus approximate $\text{Prob}_{\mu}^{\mathcal{T}}(\mathbf{F}B)$ up to any desired error bound.

2.2 A New Criterion for Decisiveness

Our goal is to provide new sufficient conditions for the decisiveness of STSs. To this end, we expose the following crucial technical lemma.

Lemma 2.7. *Let $B \in \Sigma$, and $A \in \Sigma$. Suppose that there is $p > 0$ such that for all $\nu \in \text{Dist}(A)$, we have $\text{Prob}_\nu^{\mathcal{T}}(\mathbf{F}B) \geq p$. Then for any $\mu \in \text{Dist}(S)$, $\text{Prob}_\mu^{\mathcal{T}}(\mathbf{G}B^c \wedge \mathbf{G}FA) = 0$.*

This result seems rather intuitive: if we visit A infinitely often, and after every passage through A we have a probability bounded from below to reach B , then the probability to stay in B^c forever is 0. An equivalent statement for Markov chains has been used without proof in [3, Lemmas 3.4 & 3.7]. A weaker version of this statement is given as part of the proof of [10, Proposition 36], where it is said that this general case was not known to be true or false. This weaker version assumes that there is a uniform upper bound k such that for all $\nu \in \text{Dist}(A)$, $\text{Prob}_\nu^{\mathcal{T}}(\mathbf{F}_{\leq k}B) \geq p$ to obtain a similar conclusion. We have removed the need for this constraint.

A possible proof of this result consists of regarding a stochastic transition system as a stochastic process, and resorting to results from martingale theory. More precisely, using *Lévy's zero-one law*, we obtain that infinite runs that never reach B are the same (up to a set of probability 0) as the infinite runs $s_0s_1\dots$ for which the probability to reach B given $s_0\dots s_n$ converges to 0 as n grows to infinity. Runs that visit A infinitely often cannot both avoid B and have a probability to visit B that converges to 0 (since for every visit to A , the probability to visit B is bounded from below by $p > 0$). Therefore, such runs will almost surely visit B .

We can now state our main contribution to decisiveness.

Proposition 2.8 (Decisiveness criterion). *Let $B \in \Sigma$ be a measurable set, and $A \in \Sigma$ be an attractor for \mathcal{T} . We denote $A' = A \cap (\tilde{B})^c$ the set of states of A from which B is reachable with a positive probability. Assume that there exists $p > 0$ such that for all $\nu \in \text{Dist}(A')$, $\text{Prob}_\nu^{\mathcal{T}}(\mathbf{F}B) \geq p$. Then \mathcal{T} is decisive w.r.t. B .*

With probability 1, every run visits attractor A infinitely often, but the hypotheses imply a dichotomy between runs. Some runs will reach a state of A from which B is almost surely non-reachable (i.e., in \tilde{B}). The other runs will go infinitely often through states of A such that the probability of reaching B is lower bounded by p (i.e., in A'), and will almost surely visit B by Lemma 2.7. This almost-sure dichotomy between runs is required to show decisiveness.

This criterion strictly generalizes those used in the literature. The criterion in [3, Lemma 3.4] assumes the existence of a finite attractor; the criteria in [10, Propositions 36 & 37] assume some finiteness property in an abstraction (see next section), which we do not. In [3, Lemma 3.7], a similar kind of property as ours is required from all the states of the STSs, not only from an attractor. We can actually prove the same result using a slightly more general notion of *attractor* (see [14, Proposition 11]).

2.3 Abstractions of Stochastic Transitions Systems

Decisiveness and abstractions are deeply intertwined concepts, so we briefly recall this notion [10] and related properties. We let $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$ be two STSs, and $\alpha: (S_1, \Sigma_1) \rightarrow (S_2, \Sigma_2)$ be a measurable function. We say that a set $B \in \Sigma_1$ is α -closed if $B = \alpha^{-1}(\alpha(B))$. To mean that B is α -closed, we also say that α is *compatible* with B . Following [13, 21], we define a natural way to transfer measures from (S_1, Σ_1) to (S_2, Σ_2) through α : the *pushforward* of α is the function $\alpha_\#: \text{Dist}(S_1) \rightarrow \text{Dist}(S_2)$ such that $\alpha_\#(\mu)(B) = \mu(\alpha^{-1}(B))$ for every $\mu \in \text{Dist}(S_1)$ and for every $B \in \Sigma_2$.

Definition 2.9 (α -abstraction). STS \mathcal{T}_2 is an α -abstraction of \mathcal{T}_1 if for all $\mu \in \text{Dist}(S_1)$, $\alpha_\#(\Omega_{\mathcal{T}_1}(\mu))$ is qualitatively equivalent to $\Omega_{\mathcal{T}_2}(\alpha_\#(\mu))$.

Informally, the two STSs then have the same “qualitative” single steps. Later, we may speak of *abstraction* instead of α -abstraction if α is clear in the context.

The objective of the notion of abstraction is that by finding an α -abstraction \mathcal{T}_2 which is somehow simpler than \mathcal{T}_1 (for example, with a smaller state space), we should be able to use \mathcal{T}_2 (with initial distribution $\alpha_{\#}(\mu)$) to analyze some properties of \mathcal{T}_1 (with initial distribution μ). To do so, we need to know which properties are preserved through an α -abstraction. As a first observation, positive probability of reachability properties is preserved. Stronger conditions are required to study almost-sure reachability properties through α -abstractions. We select a key property of abstractions about that matter which will be useful in the subsequent sections.

Definition 2.10 (Sound α -abstraction). We say that \mathcal{T}_2 is a *sound* α -abstraction of \mathcal{T}_1 if for all $B \in \Sigma_2$, $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F}B) = 1$ implies $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F}\alpha^{-1}(B)) = 1$.

Sound abstractions preserve almost-sure reachability properties from \mathcal{T}_2 to \mathcal{T}_1 . A major result from [10] is that if \mathcal{T}_1 is decisive w.r.t. B and \mathcal{T}_2 is a sound α -abstraction, then for all $\mu \in \text{Dist}(S_1)$, $\text{Prob}_{\mu}^{\mathcal{T}_1}(\mathbf{F}\alpha^{-1}(B)) = 1$ if and only if $\text{Prob}_{\alpha_{\#}(\mu)}^{\mathcal{T}_2}(\mathbf{F}B) = 1$. Additional results about abstractions are provided in [14, Section 2.4].

3 Application to Stochastic Hybrid Systems

We choose to restrict our attention to a stochastic extension of the well-studied *hybrid systems*. A *hybrid system* is a dynamical system combining discrete and continuous transitions. It can be defined as a non-deterministic automaton with a finite number of continuous variables, whose evolution is described via an infinite transition system. Hybrid systems have been widely studied since their introduction in the 1990s (e.g., [4, 23]). They are effectively used to model various time-dependent reactive systems; systems that need to take into account both continuous factors (e.g., speed, heat, time, distance) and discrete factors (e.g., events, instructions) are ubiquitous.

We define hybrid systems and give them a fully stochastic semantics, yielding the class of *stochastic hybrid systems* (SHSs). We then show that under classical definability hypotheses, we can obtain decidability results for qualitative and quantitative reachability problems in a class of SHSs (even though such problems are undecidable in full generality), using decisiveness (along with our new decisiveness criterion) as a key tool. This provides a stochastic extension to the study of *o-minimal hybrid systems* [16, 22, 32].

3.1 Hybrid Systems with Probabilistic Semantics

We proceed with the definition of a (non-deterministic) *hybrid system*.

Definition 3.1 (Hybrid system). A *hybrid system* (HS) is a tuple $\mathcal{H} = (L, X, \mathcal{A}, E, \gamma, \mathcal{I}, \mathcal{G}, \mathcal{R})$ where: L is a finite set of *locations* (discrete states); $X = \{x_1, \dots, x_n\}$ is a finite set of *continuous variables*; \mathcal{A} is a finite alphabet of *events*; $E \subseteq L \times \mathcal{A} \times L$ is a finite set of *edges*; for each $\ell \in L$, $\gamma(\ell): \mathbb{R}^n \times \mathbb{R}^+ \rightarrow \mathbb{R}^n$ is a continuous function describing the *dynamics* in location ℓ ; \mathcal{I} assigns to each location a subset of \mathbb{R}^n called *invariant*; \mathcal{G} assigns to each edge a subset of \mathbb{R}^n called *guard*; \mathcal{R} assigns to each edge e and valuation $\mathbf{v} \in \mathbb{R}^n$ a subset $\mathcal{R}(e)(\mathbf{v})$ of \mathbb{R}^n called *reset*. For $\ell \in L$, $e \in E$, we usually denote $\gamma(\ell)$ and $\mathcal{R}(e)$ by γ_{ℓ} and \mathcal{R}_e .

We denote the number of variables $|X|$ by n . Given a hybrid system \mathcal{H} , we define $S_{\mathcal{H}} = L \times \mathbb{R}^n$ as the states of \mathcal{H} . We distinguish two kinds of transitions between states:

- there is a *switch-transition* $(\ell, \mathbf{v}) \xrightarrow{a} (\ell', \mathbf{v}')$ if there exists $e = (\ell, a, \ell') \in E$ such that $\mathbf{v} \in \mathcal{I}(\ell) \cap \mathcal{G}(e)$, $\mathbf{v}' \in \mathcal{R}_e(\mathbf{v}) \cap \mathcal{I}(\ell')$;
- there is a *delay-transition* $(\ell, \mathbf{v}) \xrightarrow{\tau} (\ell, \mathbf{v}')$ if there exists $\tau \in \mathbb{R}^+$ such that for all $0 \leq \tau' \leq \tau$, $\gamma_\ell(\mathbf{v}, \tau') \in \mathcal{I}(\ell)$ and $\mathbf{v}' = \gamma_\ell(\mathbf{v}, \tau)$.

Informally, a switch-transition $(\ell, \mathbf{v}) \xrightarrow{a} (\ell', \mathbf{v}')$ means that an edge $e = (\ell, a, \ell')$ can be taken without violating any constraint: the value \mathbf{v} of the continuous variables is an element of the invariant $\mathcal{I}(\ell)$ and of the guard $\mathcal{G}(e)$, and there is a possible reset \mathbf{v}' of the variables which is an element of the invariant $\mathcal{I}(\ell')$. A delay-transition $(\ell, \mathbf{v}) \xrightarrow{\tau} (\ell, \mathbf{v}')$ means that some time τ elapses without changing the discrete location of the system—the only constraint is that all the values taken by the continuous variables during this time are in the invariant $\mathcal{I}(\ell)$.

Given $s = (\ell, \mathbf{v}) \in L \times \mathbb{R}^n$ a state of the hybrid system, and $\tau \in \mathbb{R}^+$, we denote by $s + \tau = (\ell, \gamma_\ell(\mathbf{v}, \tau))$ the new state after some *delay* τ , without changing the location.

We consider semantics given by *mixed transitions*, i.e., transitions that consist of a delay-transition (some time elapses) followed by a switch-transition (an edge is taken and the location changes). A mixed transition is denoted by $(\ell, \mathbf{v}) \xrightarrow{\tau, a} (\ell', \mathbf{v}')$ if and only if there exists $\mathbf{v}'' \in \mathbb{R}^n$ such that $(\ell, \mathbf{v}) \xrightarrow{\tau} (\ell, \mathbf{v}'') \xrightarrow{a} (\ell', \mathbf{v}')$.

We usually assume that there is a bijection between the edges E and the alphabet of events \mathcal{A} , and we omit mentioning this alphabet. If $e = (\ell, a, \ell') \in E$, we can thus denote \xrightarrow{e} (resp. $\xrightarrow{\tau, e}$) for switch-transitions (resp. mixed transitions) instead of \xrightarrow{a} (resp. $\xrightarrow{\tau, a}$).

Example 3.2. We provide in Figure 2 an example of a hybrid system (first studied in [7]). There are two continuous variables (x and y) and five locations, each of them equipped with the same simple dynamics: $\dot{x} = \dot{y} = 1$ (i.e., $\gamma_\ell((x, y), \tau) = (x + \tau, y + \tau)$ for every location $\ell \in L$). Locations ℓ_2 and ℓ_4 have the same invariant, which is $\{(x, y) \mid y < 1\}$; the other invariants are simply \mathbb{R}^2 . Guards are written next to the edge to which they are related: for instance, $\mathcal{G}(e_4) = \{(x, y) \mid y = 2\}$. The notation “ $x := 0$ ” is used to denote a deterministic reset (in this case, the value of x is reset to 0 after taking the edge). For instance, $\mathcal{R}_{e_1}(x, y) = \{(x, 0)\}$ (the value of x is preserved and y is reset to 0). If nothing else is written next to an edge e , it means that there is no reset on e , i.e., that $\mathcal{R}_e(\mathbf{v}) = \{\mathbf{v}\}$ for all $\mathbf{v} \in \mathbb{R}^n$. An example of mixed transitions of this system can be $(\ell_0, (0, 0)) \xrightarrow{0.4, e_0} (\ell_1, (0.4, 0.4)) \xrightarrow{0.6, e_1} (\ell_2, (1, 0)) \xrightarrow{0.2, e_2} (\ell_0, (0, 0.2)) \xrightarrow{1.5, e_3} (\ell_3, (1.5, 1.7))$.




Figure 2: Example of a hybrid system with two continuous variables. Each location is equipped with the dynamics $\dot{x} = \dot{y} = 1$.

We give more vocabulary to refer to hybrid systems. If there is a switch-transition $(\ell, \mathbf{v}) \xrightarrow{e} (\ell', \mathbf{v}')$, we say that edge e is *enabled* at state (ℓ, \mathbf{v}) —this means that edge e can be taken with no delay from state (ℓ, \mathbf{v}) . Given a state $s = (\ell, \mathbf{v})$ and an edge $e = (\ell, a, \ell')$ of \mathcal{H} , we define $I(s, e) = \{\tau \in \mathbb{R}^+ \mid s \xrightarrow{\tau, e} s'\}$ as the set of delays after which edge e is enabled from s , and $I(s) = \bigcup_e I(s, e)$ as the set of delays after which any edge is enabled from s . For instance, in the hybrid system from Figure 2, for $s = (\ell_0, (0, 0.2))$, we have $I(s, e_0) = [0, 0.8)$, and $I(s) = [0, 1.8) \setminus \{0.8\}$.

We say that a state $s \in L \times \mathbb{R}^n$ is *non-blocking* if $I(s) \neq \emptyset$. In the sequel, we only consider hybrid systems such that all states are non-blocking, thereby justifying why considering solely mixed transitions is doable—such a transition is available from any state.

We now replace the non-deterministic elements of hybrid systems with stochasticity.

Definition 3.3 (Stochastic hybrid system). A *stochastic hybrid system* (SHS) is defined as a tuple $\mathcal{H} = (\mathcal{H}', \mu_L, \eta_{\mathcal{R}}, \theta)$, where:

- $\mathcal{H}' = (L, X, \mathcal{A}, E, \gamma, \mathcal{I}, \mathcal{G}, \mathcal{R})$ is a hybrid system, which is referred to as the *underlying hybrid system of \mathcal{H}* . We require guards, invariants and resets to be Borel sets.
- $\mu_L: L \times \mathbb{R}^n \rightarrow \text{Dist}(\mathbb{R}^+)$ associates to each state a probability distribution on the time delay in \mathbb{R}^+ (equipped with the classical Borel σ -algebra) before leaving a location. Given $s \in L \times \mathbb{R}^n$, the distribution $\mu_L(s)$ will also be denoted by μ_s . We require that for every $s \in L \times \mathbb{R}^n$, $\mu_s(I(s)) = 1$, i.e., the probability that an edge is enabled after a delay is 1.
- $\eta_{\mathcal{R}}$ associates to each edge e and valuation \mathbf{v} a probability distribution on the set $\mathcal{R}_e(\mathbf{v}) \subseteq \mathbb{R}^n$. Given $e \in E$ and $\mathbf{v} \in \mathbb{R}^n$, the distribution $\eta_{\mathcal{R}}(e)(\mathbf{v})$ will also be denoted by $\eta_e(\mathbf{v})$.
- $\theta: L \times \mathbb{R}^n \rightarrow \text{Dist}(E)$ assigns to each state of \mathcal{H} a probability distribution on the edges. We require that $\theta(s)(e) > 0$ if and only if edge e is enabled at s . For $s \in L \times \mathbb{R}^n$, we denote $\theta(s)$ by θ_s . This distribution is only defined for states at which an edge is enabled.

Remark 3.4. The term “stochastic hybrid system” is used for a wide variety of stochastic extensions of hybrid systems throughout the literature. In this work, we consider stochastic delays, stochastic resets, a stochastic edge choice, and stochastic initial distributions. The way this probabilistic semantics is added on top of hybrid systems is very similar to how *timed automata* are converted to *stochastic timed automata* in [7, 11, 12]. Although dynamics appear to be deterministic, the model is powerful enough to emulate stochastic dynamics by assuming that extra variables are solely used to influence the continuous flow of the other variables. These variables can be chosen stochastically in each location through the reset mechanism. This is for example sufficient to consider a stochastic extension of the *rectangular automata* [24], whose variables evolve according to slopes inside an interval (such as $\dot{x} \in [1, 4]$). Our model is very close to the one of the software tool UPPAAL [20]. In Section 3.3, we will identify the need to restrict the definition of some components of SHSs to ensure their definability; this is however not required for the results of Section 3.2.

When referring to an SHS, we make in particular use of the same terminology as for hybrid systems (e.g., runs, enabled edges, allowed delays $I(\cdot)$) to describe its underlying hybrid system.

In order to apply the theory developed in Section 2, we give the semantics of an SHS \mathcal{H} as an STS $\mathcal{T}_{\mathcal{H}} = (S_{\mathcal{H}}, \Sigma_{\mathcal{H}}, \kappa_{\mathcal{H}})$. The set $S_{\mathcal{H}}$ is the set $L \times \mathbb{R}^n$ of states of \mathcal{H} , and $\Sigma_{\mathcal{H}}$ is the σ -algebra product between 2^L and the Borel σ -algebra on \mathbb{R}^n . To define $\kappa_{\mathcal{H}}$, we first explain briefly the role of each probability distribution in the definition of SHS. Starting from a state $s = (\ell, \mathbf{v})$, a *delay* τ is chosen randomly, according to the distribution μ_s . From state $s + \tau = (\ell, \mathbf{v}')$, an edge $e = (\ell, a, \ell')$ (enabled in $s + \tau$) is selected, following distribution $\theta_{s+\tau}$ (such an edge is almost surely available, as $\mu_s(I(s)) = 1$ by hypothesis). The next state will be in location ℓ' , and the values of the continuous variables are stochastically reset according to the distribution $\eta_e(\mathbf{v}')$. We can thus define $\kappa_{\mathcal{H}}$ as follows: for $s = (\ell, \mathbf{v}) \in S_{\mathcal{H}}$, $B \in \Sigma_{\mathcal{H}}$,

$$\kappa_{\mathcal{H}}(s, B) = \int_{\tau \in \mathbb{R}^+} \sum_{e=(\ell, a, \ell') \in E} \left(\theta_{s+\tau}(e) \cdot \int_{\mathbf{v}'' \in \mathbb{R}^n} \mathbf{1}_B(\ell', \mathbf{v}'') d(\eta_e(\gamma_{\ell}(\mathbf{v}, \tau)))(\mathbf{v}'') \right) d\mu_s(\tau)$$

where $\mathbf{1}_B$ is the characteristic function of B . It gives the probability to hit set $B \subseteq S_{\mathcal{H}}$ from state s in one step (representing a mixed transition). The function $\kappa_{\mathcal{H}}(s, \cdot)$ defines a probability distribution for all $s \in S_{\mathcal{H}}$.

Definition 3.5 (STS induced by an SHS). For an SHS \mathcal{H} , we define $\mathcal{T}_{\mathcal{H}} = (S_{\mathcal{H}}, \Sigma_{\mathcal{H}}, \kappa_{\mathcal{H}})$ as the STS induced by \mathcal{H} .

Thanks to the stochasticity of our models, we can reason about both *qualitative* and *quantitative* reachability problems, as defined in Definition 2.2.

Undecidability. We provide a proof that qualitative and quantitative reachability problems for SHSs with “simple” features are undecidable. The undecidability in itself is not surprising, as some of the undecidability proofs from the literature about non-deterministic hybrid systems [24, 25] can be translated directly to our stochastic setting. Our goal with this new proof is to get as close as possible to the class that we will later (in Section 3.3) show to be decidable, in order to outline as well as possible an undecidability frontier.

We prove undecidability even when constrained to purely continuous distributions on time delays from any state, and very simple guards, resets, and dynamics (that can be defined in simple mathematical structures). This requires a distinct proof from [24]. The result from [25] is close to the one we want to achieve, and we take inspiration from its proof. The proof consists of reducing the *halting problem for two-counter machines* to deciding whether a measurable set in an SHS is reached with probability 1.

Proposition 3.6 (Undecidability of SHSs). *The qualitative reachability problems and the approximate quantitative reachability problem are undecidable for stochastic hybrid systems with purely continuous distributions on time delays, guards that are linear comparisons of variables and constants, and using positive integer slopes for the flow of the continuous variables. The approximate quantitative problem is moreover undecidable for any fixed precision $\varepsilon < \frac{1}{2}$.*

Although the proof is centered on showing the undecidability of qualitative reachability problems, we get as a by-product the undecidability of the approximation. Indeed, as the systems used throughout the proof reach a target set B with a probability that is either 0 or 1, the ability to approximate $\text{Prob}_{\mu}^{\mathcal{T}_{\mathcal{H}}}(\mathbf{F}B)$ with $\varepsilon < \frac{1}{2}$ would be sufficient to solve the qualitative problem. As the proof shows that these qualitative problems are undecidable, we obtain that the approximate quantitative problem is undecidable as well. Our result also implies that deciding whether a state lies in \tilde{B} is undecidable.

3.2 The Cycle-Reset Hypothesis

The literature about non-deterministic hybrid systems suggests that to obtain subclasses for which the reachability problem becomes decidable, one must set sharp restrictions on the continuous flow of the variables and/or on the discrete transitions (via the *reset* mechanism). In this decidable spectrum lie for instance the *rectangular initialized automata*, which are quite permissive toward the continuous evolution of the variables, but need strong hypotheses about the discrete transitions [24]. Our approach lies at one end of this spectrum: we restrict the discrete behavior by considering *strong resets*, i.e., resets that forget about the previous values of the variables, decoupling the discrete behavior from the continuous behavior. We show that one strong reset per cycle of the graph is sufficient to obtain valuable results, and we name this property *cycle-reset*. This point of view has already been studied in [16, 22, 32] for non-deterministic hybrid systems, and in [11] for stochastic timed automata. Note that previous work about finite abstractions of non-deterministic hybrid systems (called *time-abstraction bisimulations*) does not

extend in general to “stochastic” abstractions (as defined in Section 2.3), as there may be transitions that have probability 0 to happen.

Definition 3.7 (Strong reset, cycle-reset SHS). Given \mathcal{H} a stochastic hybrid system and e an edge of \mathcal{H} , we say that e has a *strong reset* (or is *strongly reset*) if there exist $\mathcal{R}_e^* \subseteq \mathbb{R}^n$ and $\eta_e^* \in \text{Dist}(\mathcal{R}_e^*)$ such that for all $\mathbf{v} \in \mathcal{G}(e)$, $\mathcal{R}_e(\mathbf{v}) = \mathcal{R}_e^*$ and $\eta_e(\mathbf{v}) = \eta_e^*$. We say that an SHS \mathcal{H} is *cycle-reset* if for every simple cycle of \mathcal{H} , there exists a strongly reset edge.

If an edge e is strongly reset, it stochastically resets all the continuous variables when it is taken, and the stochastic reset does not depend on their values. We show two independent and very convenient results of cycle-reset SHSs: such SHSs are *decisive w.r.t. any measurable set* (the proof of this statement relies on the decisiveness criterion from Proposition 2.8), and *admit a finite abstraction*.

Decisiveness of cycle-reset SHSs. We motivate this section with an example of a simple non-decisive SHS, which we will use to show that our decisiveness result is tight.

Example 3.8. We add a stochastic layer to the hybrid system of Example 3.2, pictured in Figure 2. The distributions on the time delays in locations ℓ_0 , ℓ_2 and ℓ_4 are uniform distributions on the interval of allowed delays. For instance, at state $s = (\ell_0, (x, y))$, the distribution μ_s follows a uniform distribution on $[0, 2 - y)$. In locations ℓ_1 and ℓ_3 , the distributions on the delays are Dirac distributions. Reset and edge distributions are simply modeled as Dirac distributions. It is proved in [12, Section 6.2.2] that this SHS is not decisive w.r.t. $B = \{\ell_2\} \times \mathbb{R}^2$. The proof is quite technical and we do not recall it here; intuitively, at each passage through location ℓ_0 , the value of y increases but stays bounded from above by 1, which decreases the probability to take edge e_0 (and thus reach B); this decrease is too fast to ensure that B is almost surely reached.

Proposition 3.9. *Every cycle-reset SHS is decisive w.r.t. any measurable set.*

Placing (at least) one strong reset per simple cycle is an easy syntactic way to guarantee that almost surely, infinitely many strong resets are performed, which is the actual sufficient property used in the proof. As there are only finitely many edges, we can find a probability lower bound p on the probability to reach B after any strong reset, as required in the criterion of Proposition 2.8. Notice that as shown in Example 3.8, having independent flows for each variable and resetting each variable once in each cycle is not sufficient to obtain decisiveness; variables need to be reset *on the same discrete transition* in each cycle.

Existence of a finite abstraction. We show that cycle-reset SHSs admit a finite α -abstraction. We first give a simple example showing that without one strong reset per cycle, some simple systems do not admit a finite α -abstraction compatible with the locations.

Example 3.10. Consider the SHS of Figure 3. The self-loop edge of ℓ_0 is the only edge not being strongly reset. We assume that we want to have an abstraction compatible with $s^* := \{\ell_1\} \times \mathbb{R}$. In order to obtain an abstraction, we must split $\{\ell_0\} \times \mathbb{R}^n$ in $s_0 := \{\ell_0\} \times (-\infty, 0)$ and $\{\ell_0\} \times [0, +\infty)$, as all the states of s_0 can reach s^* with a positive probability in one step, but none of the states of $\{\ell_0\} \times [0, +\infty)$ can. Then, $\{\ell_0\} \times [0, +\infty)$ must also be split into $s_1 := \{\ell_0\} \times [0, 1)$ and $\{\ell_0\} \times [1, +\infty)$ because the states of s_1 can all reach s_0 with a positive probability in one step, but none of the states of $\{\ell_0\} \times [1, +\infty)$ can. By iterating this argument, we find that the smallest α -abstraction compatible with $\{\ell_1\} \times \mathbb{R}$ is countably infinite, and the partition that it induces is composed of s^* , s_0 and $s_i = \{\ell_0\} \times [i - 1, i)$ for $i \geq 1$. The underlying hybrid system actually belongs to the class of *updatable timed automata* [15], and the abstraction almost coincides with the *region graph* of the automaton. The reasoning we have used

is very close to the classical bisimulation algorithm for non-deterministic systems, and is explained in more detail in [14, Section 2.4].




Figure 3: The time delays are given by exponential distributions from any state; resets are Dirac distributions. The smallest abstraction compatible with $\{\ell_1\} \times \mathbb{R}$ is countably infinite.

The cycle-reset assumption is sufficient to guarantee the existence of a finite abstraction, as formulated in the next proposition.

Proposition 3.11. *Let \mathcal{H} be an SHS, and $B \in \Sigma_{\mathcal{H}}$. If \mathcal{H} is cycle-reset, it has a finite and sound abstraction compatible with B and with the locations.*

The proof consists of showing that a stochastic adaptation of the classical *bisimulation algorithm* terminates under the cycle-reset assumption. Combined with our decisiveness result, we can even show that this abstraction is sound.

3.3 Reachability Analysis in Cycle-Reset Stochastic Hybrid Systems

Our goal in this section is to perform a reachability analysis of cycle-reset SHSs. A first hurdle to circumvent is that arbitrary SHSs are in general difficult to apprehend algorithmically: for instance, the continuous evolution of their variables may be defined by solutions of systems of differential equations, which we do not know how to solve in general. To make the problem more accessible, we follow the approach of [16, 32] for non-deterministic hybrid systems by assuming that some key components of our systems are definable in a mathematical structure. We identify a large class of SHSs for which the finite abstraction from Section 3.2 is computable, which makes *qualitative* reachability problems decidable; namely, *cycle-reset o-minimal SHSs defined in a decidable theory*. We then identify sufficient hypotheses for the *approximate quantitative* problem to be decidable, in the form of a finite set of probabilities that have to be approximately computable.

Qualitative analysis. We assume that the reader is familiar with basic model-theoretic and logical terms—the reader is referred to [27] for an introduction to the main concepts. In what follows, by *definable*, we mean definable without parameters.

Definition 3.12 (Definable SHS). Given a structure \mathcal{M} , an SHS \mathcal{H} is said to be *defined in \mathcal{M}* if for every location $\ell \in L$, γ_ℓ is a function definable in \mathcal{M} and $\mathcal{I}(\ell)$ is a set definable in \mathcal{M} , and for every edge $e \in E$, the set $\mathcal{G}(e)$ is definable and there exists a first-order formula $\psi_e(\mathbf{x}, \mathbf{y})$ such that $\mathbf{v}' \in \mathcal{R}_e(\mathbf{v})$ if and only if $\psi_e(\mathbf{v}, \mathbf{v}')$ is true.

Note that we require that the *flow* of the dynamical system in each location is definable in \mathcal{M} , and not that it is the solution to a definable system of differential equations.

Our goal is to prove that the class of *cycle-reset o-minimal SHSs*, i.e., cycle-reset SHSs defined in an *o-minimal structure* (introduced in [44, 38]), with additional assumptions on the probability distributions, is decidable. We will use that the finite sound abstraction from Proposition 3.11 is then computable,

and that decisiveness (Proposition 3.9) gives us a strong link between reachability properties of the abstractions and the original systems.

Definition 3.13 (O-minimality). A totally ordered structure $\mathcal{M} = \langle M, <, \dots \rangle$ is *o-minimal* if every definable subset of M is a finite union of points and open intervals (possibly unbounded).

In other words, the definable subsets of M are exactly the ones that are definable with parameters in $\langle M, < \rangle$. Some well-known structures are o-minimal, such as the ordered additive group of rationals $\langle \mathbb{Q}, <, +, 0 \rangle$, the ordered additive group of reals $\mathbb{R}_{\text{lin}} = \langle \mathbb{R}, <, +, 0, 1 \rangle$, the ordered field of reals $\mathbb{R}_{\text{alg}} = \langle \mathbb{R}, <, +, \cdot, 0, 1 \rangle$, the ordered field of reals with the exponential function $\mathbb{R}_{\text{exp}} = \langle \mathbb{R}, <, +, \cdot, 0, 1, e^x \rangle$ [45]. There is no general result about the decidability of the theories of o-minimal structures. A well-known case is the Tarski-Seidenberg theorem [43], which asserts that there exists a quantifier-elimination algorithm for sentences in the first-order language of real closed fields. This result implies the decidability of the theory of \mathbb{R}_{alg} . However, it is not known whether the theory of \mathbb{R}_{exp} is decidable. Its decidability is implied by *Schanuel's conjecture*, a famous unsolved problem in transcendental number theory [36]. In this work, we define an *o-minimal SHS* as an SHS defined in an o-minimal structure.

The o-minimality of \mathcal{M} implies that definable subsets of M^n have a very “nice” structure, described notably by the *cell decomposition theorem* [31]. In particular, every subset of \mathbb{R}^n definable in an o-minimal structure belongs to the σ -algebra of Borel sets of \mathbb{R}^n [30, Proposition 1.1]. Moreover, to help with issues related to the definability of probabilities, we will use the property that in o-minimal structures, definable sets with positive Lebesgue measure coincide with definable sets with non-empty interior [30, Remark 2.1]. This is very helpful, as the property of having an empty interior is definable.

Remark 3.14. The definition of *o-minimal (non-deterministic) hybrid system* in the literature usually assumes that all edges are strongly reset. *O-minimal hybrid systems* were first introduced in [32], and further studied notably in [16]. The strong reset hypothesis was relaxed in [22] to “one strong reset per cycle”.

Let $\mathcal{M} = \langle \mathbb{R}, <, +, \dots \rangle$ be an o-minimal structure whose theory is decidable, such as \mathbb{R}_{alg} . Let $\mathcal{H} = (\mathcal{H}', \mu_L, \eta_{\mathcal{R}}, \theta)$ be a cycle-reset o-minimal SHS defined in \mathcal{M} . Let $\mu \in \text{Dist}(S_{\mathcal{H}})$ be an initial distribution. We make the following assumptions, which we denote by (\dagger) :

- for all $s = (\ell, \mathbf{v}) \in L \times \mathbb{R}^n$, if $I(s)$ is finite, μ_s is equivalent to the uniform discrete distribution on $I(s)$; if $I(s)$ is infinite, μ_s is equivalent to the Lebesgue measure on $I(s)$;
- the initial distribution μ is either equivalent to the discrete measure on some finite definable support D , or equivalent to the Lebesgue measure on a definable support D ;
- for $e \in E$, $\mathbf{v} \in \mathbb{R}^n$, we ask that $\mathcal{R}_e(\mathbf{v})$ is finite or has positive Lebesgue measure; $\eta_e(\mathbf{v})$ is resp. either equivalent to the discrete measure or the Lebesgue measure on $\mathcal{R}_e(\mathbf{v})$.

The first requirement was already a standard assumption in the case of stochastic timed automata [7, 11, 12]. Hypothesis (\dagger) is easily satisfied: for instance, exponential distributions (resp. uniform distributions on $[a, b]$) are equivalent to the Lebesgue measure on \mathbb{R}^+ (resp. $[a, b]$). We summarize what these ideas entail in the next proposition.

Proposition 3.15. *Let \mathcal{H} be a cycle-reset o-minimal SHS defined in a structure whose theory is decidable. Let $B \in \Sigma_{\mathcal{H}}$ be definable and $\mu \in \text{Dist}(S_{\mathcal{H}})$. We assume that assumption (\dagger) holds. Then one can decide whether $\text{Prob}_{\mu}^{\mathcal{F}_{\mathcal{H}}}(\mathbf{F}B) = 1$ and whether $\text{Prob}_{\mu}^{\mathcal{F}_{\mathcal{H}}}(\mathbf{F}B) = 0$.*

In particular, we can decide the qualitative reachability problems for cycle-reset SHSs defined in \mathbb{R}_{alg} and satisfying (\dagger) . Assuming Schanuel's conjecture [36], we could extend this result to \mathbb{R}_{exp} .

Approximate quantitative analysis. Under strengthened numerical hypotheses, we can solve the approximate quantitative reachability problem in cycle-reset SHSs. Let \mathcal{H} be a cycle-reset SHS, $B \in \Sigma_{\mathcal{H}}$ and $\mu \in \text{Dist}(S_{\mathcal{H}})$. Our goal is to apply the approximation scheme described in [10] in order to approximate $\text{Prob}_{\mu}^{\mathcal{F}_{\mathcal{H}}}(\mathbf{F}B)$. To do so, we require that $\mathcal{F}_{\mathcal{H}}$ is decisive w.r.t. B , which is implied by the cycle-reset hypothesis (Proposition 3.9), and the ability to compute for all $m \in \mathbb{N}$, an arbitrarily close approximation of

$$p_m^{\text{Yes}} = \text{Prob}_{\mu}^{\mathcal{F}}(\mathbf{F}_{\leq m}B) \text{ and } p_m^{\text{No}} = \text{Prob}_{\mu}^{\mathcal{F}}(B^c \mathbf{U}_{\leq m}\tilde{B}).$$

Using the cycle-reset hypothesis, we can express these probabilities as sums and products of a *finite* number of probabilities of paths with bounded length $b \in \mathbb{N}$, where b is the length of the longest path without encountering a strong reset. Details on this computation are available in [14, Section 5.2].

4 Conclusion

Summary. We presented in Section 2 how to solve reachability problems in stochastic transition systems via the *decisiveness* notion, introduced in [3, 10]. We notably solved in Lemma 2.7 a technical question that was open in [10], which helped formulate a general sufficient condition for decisiveness in Proposition 2.8, encompassing known decisiveness criteria from the literature.

In Section 3, we considered *stochastic hybrid systems* (SHSs). We showed that SHSs with one *strong reset* per cycle (*cycle-reset*) are decisive (using our new decisiveness criterion), and admit a finite abstraction. We then identified assumptions, using ideas from logic, leading to the effective computability of this abstraction, and to the decidability of qualitative and quantitative reachability problems. These assumptions pertain to the definability of the different components of the SHSs in a decidable mathematical structure. Combined with the previous decisiveness results, this abstraction can be used to decide qualitative reachability problems. In particular, these results apply to cycle-reset SHSs defined in $\mathbb{R}_{\text{alg}} = \langle \mathbb{R}, <, +, \cdot, 0, 1 \rangle$ with well-behaved distributions.

Possible extensions and future work. We identify some possible extensions of our results.

A first direction of study is to find other classes of decisive stochastic systems that can be encompassed by our decisiveness criterion (Proposition 2.8). In that respect, a good candidate is the class of *stochastic regenerative Petri nets* [28, 37]. An application of decisiveness results to *stochastic Petri nets* was briefly discussed in [10, Section 8.3], but under severe constraints; we may be able to relax part of these constraints with the generalized criterion.

In Section 3.3, we circumvent the issue of the definability of measures by using a specific property of the o-minimal structures (namely, that the Lebesgue measure of a definable set is positive if and only if the interior of that set is non-empty). However, more powerful results exist about the compatibility of o-minimal structures and measure theory [30]: some o-minimal structures are closed under integration with respect to a given measure (then called *tame* measure). This consideration may help extend our results to a larger class that is less restrictive w.r.t. probability distributions, or in which actual probabilities may be definable.

Acknowledgments. We would like to thank the anonymous reviewers for their valuable advice, which notably helped simplify the proof of Lemma 2.7.

References

- [1] Alessandro Abate, Joost-Pieter Katoen, John Lygeros & Maria Prandini (2010): *Approximate Model Checking of Stochastic Hybrid Systems*. *Eur. J. Control* 16(6), pp. 624–641, doi:10.3166/ejc.16.624-641.
- [2] Alessandro Abate, Maria Prandini, John Lygeros & Shankar Sastry (2008): *Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems*. *Automatica* 44(11), pp. 2724–2734, doi:10.1016/j.automatica.2008.03.027.
- [3] Parosh A. Abdulla, Noomene Ben Henda & Richard Mayr (2007): *Decisive Markov Chains*. *Log. Methods Comput. Sci.* 3(4), doi:10.2168/LMCS-3(4:7)2007.
- [4] Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis & Sergio Yovine (1995): *The Algorithmic Analysis of Hybrid Systems*. *Theor. Comput. Sci.* 138(1), pp. 3–34, doi:10.1016/0304-3975(94)00202-T.
- [5] Rajeev Alur & David L. Dill (1994): *A Theory of Timed Automata*. *Theor. Comput. Sci.* 126(2), pp. 183–235, doi:10.1016/0304-3975(94)90010-8.
- [6] Rajeev Alur & George J. Pappas, editors (2004): *Hybrid Systems: Computation and Control, 7th International Workshop, HSCC 2004, Philadelphia, PA, USA, March 25-27, 2004, Proceedings*. *Lecture Notes in Computer Science* 2993, Springer, doi:10.1007/b96398.
- [7] Christel Baier, Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye & Marcus Größer (2008): *Almost-Sure Model Checking of Infinite Paths in One-Clock Timed Automata*. In: *Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24-27 June 2008, Pittsburgh, PA, USA*, IEEE Computer Society, pp. 217–226, doi:10.1109/LICS.2008.25. Available at <https://ieeexplore.ieee.org/xpl/conhome/4557886/proceeding>.
- [8] Christel Baier & Joost-Pieter Katoen (2008): *Principles of model checking*. MIT Press.
- [9] Alessandro Berarducci & Margarita Otero (2004): *An additive measure in o-minimal expansions of fields*. *Q. J. Math.* 55(4), pp. 411–419, doi:10.1093/qmath/hah010.
- [10] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye & Pierre Carlier (2018): *When are stochastic transition systems tameable?* *J. Log. Algebr. Meth. Program.* 99, pp. 41–96, doi:10.1016/j.jlamp.2018.03.004.
- [11] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye & Nicolas Markey (2008): *Quantitative Model-Checking of One-Clock Timed Automata under Probabilistic Semantics*. In: *Fifth International Conference on the Quantitative Evaluation of Systems (QEST 2008), 14-17 September 2008, Saint-Malo, France*, IEEE Computer Society, pp. 55–64, doi:10.1109/QEST.2008.19. Available at <https://ieeexplore.ieee.org/xpl/conhome/4634932/proceeding>.
- [12] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, Quentin Menet, Christel Baier, Marcus Größer & Marcin Jurdzinski (2014): *Stochastic Timed Automata*. *Log. Methods Comput. Sci.* 10(4), doi:10.2168/LMCS-10(4:6)2014.
- [13] Vladimir I. Bogachev (2007): *Measure theory*. 1, Springer Science & Business Media, doi:10.1007/978-3-540-34514-5.
- [14] Patricia Bouyer, Thomas Brihaye, Mickael Randour, Cédric Rivière & Pierre Vandenhove (2020): *Decisiveness of Stochastic Systems and its Application to Hybrid Models (Full Version)*. *CoRR*.
- [15] Patricia Bouyer, Catherine Dufourd, Emmanuel Fleury & Antoine Petit (2004): *Updatable timed automata*. *Theor. Comput. Sci.* 321(2-3), pp. 291–345, doi:10.1016/j.tcs.2004.04.003.
- [16] Thomas Brihaye, Christian Michaux, Cédric Rivière & Christophe Troestler (2004): *On O-Minimal Hybrid Systems*. In Alur & Pappas [6], pp. 219–233, doi:10.1007/978-3-540-24743-2_15.
- [17] Luminita Manuela Bujorianu (2012): *Stochastic reachability analysis of hybrid systems*. Springer Science & Business Media, doi:10.1007/978-1-4471-2795-6.

- [18] Pierre Carlier (2017): *Verification of Stochastic Timed Automata. (Vérification des automates temporisés et stochastiques)*. Ph.D. thesis, University of Paris-Saclay, France. Available at <https://tel.archives-ouvertes.fr/tel-01696130>.
- [19] Christos G. Cassandras & John Lygeros (2007): *Stochastic Hybrid Systems*. CRC Press, doi:10.1201/9781420008548.
- [20] Alexandre David, Dehui Du, Kim G. Larsen, Axel Legay, Marius Mikucionis, Danny Bøgsted Poulsen & Sean Sedwards (2012): *Statistical Model Checking for Stochastic Hybrid Systems*. In Ezio Bartocci & Luca Bortolussi, editors: *Proceedings First International Workshop on Hybrid Systems and Biology, HSB 2012, Newcastle Upon Tyne, UK, 3rd September 2012, EPTCS 92*, pp. 122–136, doi:10.4204/EPTCS.92.9.
- [21] Daniel Gburek, Christel Baier & Sascha Klüppelholz (2016): *Composition of Stochastic Transition Systems Based on Spans and Couplings*. In: *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy, LIPIcs 55*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 102:1–102:15, doi:10.4230/LIPIcs.ICALP.2016.102.
- [22] Raffaella Gentilini (2005): *Reachability Problems on Extended O-Minimal Hybrid Automata*. In Paul Pettersson & Wang Yi, editors: *Formal Modeling and Analysis of Timed Systems, Third International Conference, FORMATS 2005, Uppsala, Sweden, September 26-28, 2005, Proceedings, Lecture Notes in Computer Science 3829*, Springer, pp. 162–176, doi:10.1007/11603009_14.
- [23] Thomas A. Henzinger (1996): *The Theory of Hybrid Automata*. In: *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, IEEE Computer Society, pp. 278–292, doi:10.1109/LICS.1996.561342. Available at <https://ieeexplore.ieee.org/xpl/conhome/4265/proceeding>.
- [24] Thomas A. Henzinger, Peter W. Kopke, Anuj Puri & Pravin Varaiya (1998): *What's Decidable about Hybrid Automata?* *J. Comput. Syst. Sci.* 57(1), pp. 94–124, doi:10.1006/jcss.1998.1581.
- [25] Thomas A. Henzinger & Jean-François Raskin (2000): *Robust Undecidability of Timed and Hybrid Systems*. In Lynch & Krogh [35], pp. 145–159, doi:10.1007/3-540-46430-1_15.
- [26] João P. Hespanha (2004): *Stochastic Hybrid Systems: Application to Communication Networks*. In Alur & Pappas [6], pp. 387–401, doi:10.1007/978-3-540-24743-2_26.
- [27] Wilfrid Hodges (1997): *A Shorter Model Theory*. Cambridge University Press.
- [28] András Horváth, Marco Paolieri, Lorenzo Ridi & Enrico Vicario (2012): *Transient analysis of non-Markovian models using stochastic state classes*. *Perform. Eval.* 69(7-8), pp. 315–335, doi:10.1016/j.peva.2011.11.002.
- [29] Jianghai Hu, John Lygeros & Shankar Sastry (2000): *Towards a Theory of Stochastic Hybrid Systems*. In Lynch & Krogh [35], pp. 160–173, doi:10.1007/3-540-46430-1_16.
- [30] Tobias Kaiser (2012): *First order tameness of measures*. *Ann. Pure Appl. Logic* 163(12), pp. 1903–1927, doi:10.1016/j.apal.2012.06.002.
- [31] Julia F. Knight, Anand Pillay & Charles Steinhorn (1986): *Definable sets in ordered structures. II*. *Transactions of the American Mathematical Society* 295(2), pp. 593–605, doi:10.1090/S0002-9947-1986-0833698-1.
- [32] Gerardo Lafferriere, George J. Pappas & Shankar Sastry (2000): *O-Minimal Hybrid Systems*. *Math. Control Signals Syst.* 13(1), pp. 1–21, doi:10.1007/PL00009858.
- [33] Xiangfang Li, Oluwaseyi Omotere, Lijun Qian & Edward R. Dougherty (2017): *Review of stochastic hybrid systems with applications in biological systems modeling and analysis*. *EURASIP J. Bioinformatics and Systems Biology* 2017, p. 8, doi:10.1186/s13637-017-0061-5.
- [34] John Lygeros & Maria Prandini (2010): *Stochastic Hybrid Systems: A Powerful Framework for Complex, Large Scale Applications*. *Eur. J. Control* 16(6), pp. 583–594, doi:10.3166/ejc.16.583-594.

- [35] Nancy A. Lynch & Bruce H. Krogh, editors (2000): *Hybrid Systems: Computation and Control, Third International Workshop, HSCC 2000, Pittsburgh, PA, USA, March 23-25, 2000, Proceedings*. Lecture Notes in Computer Science 1790, Springer, doi:10.1007/3-540-46430-1.
- [36] Angus Macintyre & Alex J. Wilkie (1996): *On the decidability of the real exponential field*. In: *Kreiseliana: About and around Georg Kreisel*, A. K. Peters, pp. 441–467.
- [37] Marco Paolieri, András Horváth & Enrico Vicario (2016): *Probabilistic Model Checking of Regenerative Concurrent Systems*. *IEEE Trans. Software Eng.* 42(2), pp. 153–169, doi:10.1109/TSE.2015.2468717.
- [38] Anand Pillay & Charles Steinhorn (1986): *Definable sets in ordered structures. I*. *Transactions of the American Mathematical Society* 295(2), pp. 565–592, doi:10.2307/2000052.
- [39] Amir Pnueli (1977): *The Temporal Logic of Programs*. In: *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, IEEE Computer Society, pp. 46–57, doi:10.1109/SFCS.1977.32. Available at <https://ieeexplore.ieee.org/xpl/conhome/4567914/proceeding>.
- [40] Maria Prandini & Jianghai Hu (2009): *Application of Reachability Analysis for Stochastic Hybrid Systems to Aircraft Conflict Prediction*. *IEEE Trans. Automat. Contr.* 54(4), pp. 913–917, doi:10.1109/TAC.2008.2011011.
- [41] Maria Prandini, Jianghai Hu, John Lygeros & Shankar Sastry (2000): *A probabilistic approach to aircraft conflict detection*. *IEEE Trans. Intelligent Transportation Systems* 1(4), pp. 199–220, doi:10.1109/6979.898224.
- [42] Abhyudai Singh & Joao P. Hespanha (2010): *Stochastic hybrid systems for studying biochemical processes*. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 368(1930), pp. 4995–5011, doi:10.1098/rsta.2010.0211.
- [43] Alfred Tarski (1951): *A decision method for elementary algebra and geometry*. University of California Press, Berkeley, California.
- [44] Lou van den Dries (1984): *Remarks on Tarski's problem concerning $(R, +, *, exp)$* . *Studies in Logic and the Foundations of Mathematics* 112(C), pp. 97–121, doi:10.1016/S0049-237X(08)71811-1.
- [45] Alex J. Wilkie (1996): *Model completeness results for expansions of the ordered field of real numbers by restricted Pfaffian functions and the exponential function*. *Journal of the American Mathematical Society* 9(4), pp. 1051–1094, doi:10.1090/S0894-0347-96-00216-0.