Markov Set-Chains as abstractions of Stochastic Hybrid Systems *

Alessandro Abate[1], Alessandro D'Innocenzo[2], Maria D. Di Benedetto[2], and Shankar S. Sastry[1]

[1] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley - USA
{aabate,sastry}@eecs.berkeley.edu
[2] Department of Electrical Engineering and Computer Science, Center of Excellence DEWS, University of L'Aquila - Italy
{adinnoce,dibenede}@ing.univaq.it

Abstract. The objective of this study is to introduce and formalize an abstraction procedure that applies to a rather general class of dynamical systems, that is to models known as discrete-time stochastic hybrid systems (dt-SHS). The procedure abstracts the original dt-SHS into a Markov set-chain (MSC) in two steps. First, a Markov chain (MC) is obtained by partitioning the hybrid state space, according to a controllable parameter, into non-overlapping domains and computing transition probabilities for these nodes according to the dynamics of the dt-SHS. Second, explicit error bounds for the abstraction that depend on the parameter are derived, and are associated to the values of the MC, thus obtaining a MSC. We show that one can arbitrarily increase the accuracy of the abstraction by tuning the refinement parameter, albeit at an increase of the cardinality of the MSC. The application of a number of results from the MSC literature enables the analysis of the dynamics of the original dt-SHS. In this work, the asymptotics of the dt-SHS dynamics are assessed within the abstracted framework.

1 Introduction and Objectives

Hybrid Systems (HS) are dynamical system with interleaved continuous and discrete behaviors. Their great expressive power and complexity of allowed dynamical behaviors are offset, as expected, by two main issues. The first is the subtlety of their theoretical investigation. Much research has been directed to further the understanding of their properties from the perspective of Systems Theory. The second is the issue of scalability, in particular with respect to computational complexity. For instance, the formal verification of properties of the system (e.g., through the use of *model checking* procedures [1]), whereby the state space is semi–exhaustively searched, is complicated by the continuity of the space, on the one hand, and by the possible increase in its dimension, on the other.

* This work was partially supported by European Commission under Project IST NoE HYCON contract n. 511368, STREP project n. TREN/07/FP6AE/S07.71574/037180 IFLY, and by the NSF grant CCR-0225610.

That of model checking of HS specifications [2] is an intensely investigated area. A noteworthy technique which is often employed is that of *abstraction*. According to this approach, a system with a smaller state space (possibly finite) is obtained, which is equivalent, in some sense, to the original system. Systems equivalence is usually defined via the notions of language equivalence and bisimulation [3, 4]. Recently, *approximate* notions of system equivalence [5, 6] have been developed to generalize the abstraction problem, where a metric is introduced to quantify the distance between the original system and the abstracted one. [7] proposes an algorithm and implement a tool to construct an approximate abstraction of a HS by means of a timed automaton. In [8] a notion of approximate bisimilarity is proposed for a class of stochastic hybrid systems: these are HS which are endowed with probabilistic terms.

The present contribution introduces a formal abstraction procedure for a general class of *Stochastic* Hybrid Systems (SHS). We investigate the dynamics from a perspective, which is different than the classical one of the "solution process." Instead, we look at the transition probability function to describe the likelihood of the presence of an execution in a particular region of the state space, throughout time. The evolution of this likelihood follows the Kolmogorov forward equation. In Physics, this approach is adopted by the study of the Fokker-Planck equation. Recently [9] tried to extend to the SHS framework a computational approach to solve this equation. This tentative has shown that this computation is often mathematically intractable. Another instance of this analytical and computational difficulty is given by the solution of the Master equation [10], which arises in the description of chemical and biological networks and is a particular instance of the Fokker-Planck relation.

This work leverages a discrete time framework and the presence of spatial guards in a class of SHS (dt-SHS) to be able to express the transition probability function in a compact way. After introducing a partitioning procedure for the hybrid state space, the transition probabilities between these partitions are approximately computed by employing the concept of probabilistic reachability. Hence, a Markov chain (MC) is built up. Furthermore, by raising some continuity assumptions on the entities that characterize the dynamics of the dt-SHS, explicit error bounds are associated to the transition probabilities. These error bounds depend on the diameter of the introduced partitions. This allows to formally set up a Markov set-chain (MSC) associated to the partitioning procedure. The asymptotics of the MSC can be then related to that of the dt-SHS.

The present technique can be related to the work in [11], where the weak convergence to the original process of a discretization of the continuous dynamics into that of a MC defined on a set of grid points is proven. No error bounds are explicitly derived though. Both this work and [11] approximate the original process with a probabilistic discrete graphical structure. This connects to efforts on automatic verification (*model checking*) of stochastic timed automata (that is, a subclass of SHS), which has been investigated in [12–15].

A general understanding of the area of probabilistic model checking for SHS is however still far. As a first result towards this goal, we have shown the ability to construct a finite state abstraction that allows us to efficiently compute the steady state of the original system with arbitrary precision.

The paper is organized as follows. In Section 2 the dt-SHS model is introduced. In Section 3 an overview on MSC is given. In Section 4 we look into the dynamics of the dt-SHS, by discussing the notion of probabilistic reachability. In

Sections 5 and 6 we introduce the abstraction procedure, and employ it to investigate the asymptotics of the original system. We apply an implementation of the algorithm to a case study in Section 7, and conclude in Section 8.

2 dt-SHS Model

In this section we introduce the dt-SHS model. This mathematical framework is inspired by that introduced in [16], albeit with some modifications on the events generating mechanisms. Here we model the presence of a physical, forcing guards set, rather than introducing state-dependent transition probabilities, as in [16]—extensions to this more general case are discussed in Section 8. The use of a discrete time framework is motivated by the simplicity in handling measurability issues for events on the underlying probability space, as well as by the possibility to directly compute transition probabilities for the whole domain.

Definition 1 (dt-SHS). *A discrete time stochastic hybrid system is a tuple* $\mathcal{H} = (\mathcal{Q}, n, \mathcal{G}, T, R)$*, where*

- $\mathcal{Q} := \{q_1, q_2, \ldots, q_m\}$*, for some finite* $m \in \mathbb{N}$*, is the discrete state space;*
- $n : \mathcal{Q} \to \mathbb{N}$ *assigns to each* $q \in \mathcal{Q}$ *the dimension of the continuous state space* $\mathbb{R}^{n(q)}$*. The continuous part of the dynamics evolves within a "domain" for each mode* q*, which is defined to be a compact subset* $\mathcal{D}_q^* \subset \mathbb{R}^{n(q)}$*. The whole hybrid state space is given by* $\mathcal{S}^* := \cup_{q\in\mathcal{Q}}\{q\} \times \mathcal{D}_q^*$*;*
- $\mathcal{G} := \cup_{q\in\mathcal{Q}}\{q\} \times \mathcal{G}_q, \mathcal{G}_q = \{g_{ij}; i, j \in \mathcal{Q}, i \neq j, g_{ij} \subseteq \mathcal{D}_i^*\}$ *is the set of spatial guards. We assume that* $\forall i, j, k \in \mathcal{Q}, i \neq j \neq k, g_{ij} \cap g_{ik} = \varnothing$*, and that the guards have non-trivial volume:* $\mathcal{L}(g_{ij}) \neq 0, \forall i, j \in \mathcal{Q}, i \neq j$*, where* $\mathcal{L}(\cdot)$ *denotes the Lebesgue measure associated to any Borel subset of a particular domain. Let us further introduce the set* $\mathcal{D}_i := \mathcal{D}_i^* \backslash \{\cup_{j\neq i, j\in Q} g_{ij}\}$*, the "invariant" of mode* i*, and* $\mathcal{S} := \cup_{q\in\mathcal{Q}}\{q\} \times \mathcal{D}_q$*;*
- $T : \mathcal{B}(\mathcal{D}_{(\cdot)}^*) \times \mathcal{S} \to [0, 1]$ *is a Borel-measurable stochastic kernel (the "transition kernel") on* $\mathcal{D}_{(\cdot)}^*$ *given* $\mathcal{S}$*, which assigns to each* $s = (q, x) \in \mathcal{S}$ *a probability measure on the Borel space* $(\mathcal{D}_q^*, \mathcal{B}(\mathcal{D}_q^*))$*:* $T(dx|(q, x))$*;*
- $R : \mathcal{B}(\mathcal{D}_{(\cdot)}^*) \times \mathcal{G} \times \mathcal{Q} \to [0, 1]$ *is a Borel-measurable stochastic kernel (the "reset kernel") on* $\mathcal{D}_{(\cdot)}^*$*, given* $\mathcal{G} \times \mathcal{Q}$*, that assigns to each* $s = (q, x) \in \mathcal{G}$*, and* $q' \in \mathcal{Q}, q' \neq q$*, a probability measure on the Borel space* $(\mathcal{D}_{(q')}^*, \mathcal{B}(\mathcal{D}_{(q')}^*))$*:* $R(dx|(q, x), q')$*.* □

The system initialization at the initial time $k = 0$ is specified through some probability measure $\pi_0 : \mathcal{B}(\mathcal{S}^*) \to [0, 1]$ on the Borel space $(\mathcal{S}^*, \mathcal{B}(\mathcal{S}^*))$. Here $\mathcal{B}(\mathcal{S}^*)$ is the σ-field generated by the subsets of $\mathcal{S}^*$ of the form $\cup_q\{q\} \times B_q$, with B_q denoting a Borel set in $\mathcal{D}_q^*$. For details on measurability and metric properties of $\mathcal{H}$, the reader is invited to refer to [16, 17]. Notice that the transition and reset kernels have a different domain of definition, but the same support. Next, we define the notion of execution for the above model (throughout the paper, random processes will be denoted in bold fonts, while random variables in normal typesets).

Definition 2 (Execution). *Consider a dt-SHS* $\mathcal{H} = (\mathcal{Q}, n, \mathcal{G}, T, R)$*. An execution for* $\mathcal{H}$*, associated with an initial distribution* π_0*, is a stochastic process* $\{\mathbf{s}(k), k \in [0, N], N \in \mathbb{N}\}$ *with values in* $\mathcal{S}^*$*, whose sample paths are obtained according to the following algorithm:*

extract from $\mathcal{S}^*$ a value $s_0 = (q_0, x_0)$ for $\mathbf{s}(0)$, according to the distribution π_0;

for $k = 0$ to $N - 1$,

if there is a $j \neq q_k, j \in \mathcal{Q}$, such that $x_k \in g_{q_k,j}$,

then extract a value $s_{k+1} \in \mathcal{S}^*$ for $\mathbf{s}(k+1)$, according to $R(\cdot\,|s_k, j)$;

else extract a value $s_{k+1} \in \mathcal{S}^*$ for $\mathbf{s}(k+1)$, according to $T(\cdot\,|s_k)$;

end. □

Remark 1 (a note on Event Detection). The assumptions in Def. 1 on the shape of the guards set allows to disregard potential issues (such as non-determinism) with the existence of "events," that is conditions in time where an execution belongs to the guards set. We shall not further pursue the interpretation of an event "pointwise," that is by associating an event time and an event point to a particular (deterministic) initial condition and a realization of the hybrid execution, as suggested by the procedure in Definition 2. This approach, which is usual for the deterministic HS case, is quite more complicated to implement in the stochastic case. Instead, as already argued in Section 1, we compute certain transition probabilities, which will be referred to any point in $\mathcal{S}^*$. That is, we shall compute transition probabilities between invariants and guards of $\mathcal{H}$, from guards back to invariants, as well as combinations or products (forward in time) thereof, thus characterizing probabilistically the evolution of a hybrid execution on $\mathcal{S}^*$. For general stochastic models (for instance, the SHS framework in [18]), determining or computing transition probabilities may not be easy. However, in the proceeding of this work we show that for the current dt-SHS setup, such a computation is feasible. □

As mentioned, the introduced (autonomous) dt-SHS is related to the (controlled) SHS in [16], where the additional presence of a transition kernel allows for jumps within the invariant set. The theory developed in this work can be extended to account for similar terms, which we do not do for the sake of simplicity. The reader is invited to refer to [16] for connections between this model, and other SHS models in the literature, as well as for the proof of further properties (e.g., the Markov property).

3 Markov set-chains

We define here the concept of Markov set-chains, which in this paper is used as an abstraction framework for dt-SHS. We recall some results from [19], which are of interest for the present work.

Definition 3. *Let $P, Q \in \mathbb{R}^{n\times n}$ be nonnegative matrices (not necessarily stochastic) with $P \leq Q$. We define the following " transition set:"*

$$[P, Q] = \{A \in \mathbb{R}^{n\times n} : A \text{ is a stochastic matrix and } P \leq A \leq Q\} \qquad \square$$

In the proceeding, we assume that the set $[P, Q] \neq \varnothing$. When the "bounding matrices" will be clear by the context, we will use the more compact notation $[\Pi]$. We can define a Markov set-chain as a non-homogeneous discrete-time Markov chain, where the transition probabilities vary non-deterministically within a compact transition set $[\Pi]$ at each time step. More formally,

Definition 4. *Let $[\Pi]$ be a transition set, i.e. a compact set of $n \times n$ stochastic matrices. Consider the set of all non-homogeneous Markov chains having all their transition matrices in $[\Pi]$. We call the sequence*

$$[\Pi], [\Pi]^2, \cdots$$

a Markov set-chain, where $[\Pi]^k$ is defined by induction as the set of all possible products $A_1, \cdots, A_k$, such that $\forall i = 1, \cdots, k, A_i \in [\Pi]$.
Let $[\pi_0]$ be a compact set of $1 \times n$ stochastic vectors, introduced similarly as in Def. 3. We call $[\pi_0]$ the initial distribution set. □

The compact set $[\pi_k] = [\pi_0][\Pi]^k$ is the k-th distribution set and

$$[\pi_0], [\pi_0][\Pi], \cdots$$

is the Markov set-chain with initial distribution set $[\pi_0]$.

It can be shown that each element $[\pi_k]$ is a convex polytope if $[\pi_0]$ is a convex polytope and $[\Pi]$ is a transition set. It should be noticed that the number of vertices of $[\pi_k]$ increases with k, thus the computational burden to obtain $[\pi_k]$ for large values of k should be accounted for. However, it is possible to compute *tight* (see [19]) upper and lower bounding matrices L_k, H_k for $[\pi_k]$ in a very efficient way, in particular the computation of L_k, H_k can be recursively obtained from L_{k-1}, H_{k-1}.

Definition 5. *For any stochastic matrix A, its coefficient of ergodicity is defined as follows:*

$$\mathcal{T}(A) = \frac{1}{2} \max_{i,j} ||a_i - a_j||,$$

where a_i is the i–th row of A. □

The above definition can be directly extended to Markov set-chains:

Definition 6. *For any transition set $[\Pi]$, its coefficient of ergodicity is defined as follows:*

$$\mathcal{T}([\Pi]) = \max_{A \in [\Pi]} \mathcal{T}(A).$$

□

Notice that since $\mathcal{T}(\cdot)$ is a continuous function and $[\Pi]$ a compact set, the maximum argument of $\mathcal{T}([\Pi])$ exists. Notice that $\mathcal{T}([\Pi]) \in [0, 1]$: this value provides a measure of the "contractive" nature of the Markov set-chain: the smaller $\mathcal{T}([\Pi])$, the more contractive the MSC. The exact value of $\mathcal{T}([\Pi])$ can be hard to be compute, but it can be upper bounded as follows, [19].

Theorem 1. *Let $[\Pi]$ be the interval $[P, Q]$, then:*

$$\mathcal{T}([\Pi]) \leq \frac{1}{2} \max_{i,j} \sum_{k=1}^{n} \max\{|p_{ik} - q_{jk}|, |p_{jk} - q_{ik}|\} \triangleq \mathcal{T}^*([\Pi]).$$

The following notion is important for characterizing the convergence of a MSC:

Definition 7. *Suppose r is an integer such that $\mathcal{T}(A_1, \cdots, A_r) < 1$, $\forall A_1, \cdots, A_r \in [\Pi]$. Then $[\Pi]$ is said to be product scrambling and r its scrambling integer.* □

We now illustrate some results on the convergence of MSC.

Theorem 2. *Given a product scrambling MSC with transition set* $[\Pi]$ *and initial distribution set* $[\pi_0]$*, then there exists a unique limit set* $[\pi_\infty]$ *such that* $[\pi_\infty][\Pi] = [\pi_\infty]$*. Moreover, let* r *be the scrambling integer. Then for any positive integer* h*,*

$$d([\pi_k], [\pi_\infty]) \leq K\beta^h \tag{1}$$

where $K = [\mathcal{T}([\Pi]^r)]^{-1} d([\pi_0], [\pi_\infty])$ *and* $\beta = \mathcal{T}([\Pi]^r)^{\frac{1}{r}} < 1$*. Thus*

$$\lim_{k\to\infty} [\pi_k] = \lim_{k\to\infty} [\pi_0][\Pi]^k = [\pi_\infty].$$

□

As we argued before, the exact computation of $[\pi_\infty]$ can be expensive. However, it is possible to use the upper and lower bounding matrices L_k, H_k defined above to obtain an accurate estimate of $[\pi_\infty]$ with a reasonable computational complexity. In fact, L_k, H_k converge to a value L_∞, H_∞ such that $[\pi_\infty] \subseteq L_\infty, H_\infty$.
Define the diameter of a compact set (referred to either matrices or vectors) as

$$\Delta([\Pi]) = \max_{A,A' \in [\Pi]} ||A - A'||.$$

The following result provides an efficient procedure to compute an upper bound for the diameter of the limit set $[\pi_\infty]$.

Theorem 3. *Given a product scrambling Markov set-chain with transition set* $[\Pi] = [P, Q]$ *and such that* $\mathcal{T}([\Pi]) < 1$*, then*

$$\Delta([\pi_\infty]) \leq \frac{\Delta([\Pi])}{1 - \mathcal{T}([\Pi])} \leq \frac{||Q - P||}{1 - \mathcal{T}^*([\Pi])}.$$

□

4 Probabilistic Dynamics

The model described in Definition 1 is quite general and allows for a wealth of possible behaviors. Even some further knowledge of the structure of the dynamics, beyond the general stochastic kernels T, R that characterize it, are in general not translatable into a closed-form expression for the solution process of $\mathcal{H}$. Then, as anticipated in Section 2 (see Remark 1), in order to study the dynamical properties of $\mathcal{H}$, two avenues can be pursued. The first looks at the ensemble of allowed realizations that spring out of the support of the initial distribution, according to the steps in Definition 2, possibly averaged over the initial distribution. Monte Carlo simulations are an example of this approach. The second instead tries to characterize probabilistically the presence of the solution process in certain regions of $\mathcal{S}^*$, as time progresses. This characterization can leverage the ability of defining and computing quantities related to the concept of *probabilistic reachability* [16]. More precisely, it is of interest to characterize the following likelihood: given a point $s_0 \in \mathcal{S}^*$, what is the probability that the solution process $\mathbf{s}(\cdot)$ of $\mathcal{H}$, starting from s_0, is located in the set $A \in \mathcal{B}(\mathcal{S}^*)$ at time $k > 0$? Similarly, given a point $s_0 \in \mathcal{S}^*$, what is the probability that the solution process $\mathbf{s}(\cdot)$ of $\mathcal{H}$ stays within the set $A \in \mathcal{B}(\mathcal{S}^*), s_0 \in A$, for all the

time $k \in [0, N], N < \infty$?
It is not a case that these stochastic reachability problems are the related to the two deterministic reachability approaches taken in [20, 7] for constructing finite abstractions of deterministic HS.
Given the dependence of the definition of the two sets of probabilistic kernels on, respectively, the invariants and the guards sets, we are in particular interested in computing the probability of reaching these subsets of the hybrid state space. For instance, considering two modes $q, q' \in \mathcal{Q}$, we call $p_{q,q'}(x)$ the probability that a trajectory, starting from a point $(q, x) \in \mathcal{S}$, has to transition (according to $T(\cdot|(q, x))$) in a time step into any other domain $q' \neq q$, or possibly to continue evolving in $q' = q$:

$$p_{q,q'}(x) \triangleq \int_{g_{q,q'}} T(dy|(q, x)), \text{ if } q' \neq q, \tag{2}$$
$$p_{q,q}(x) \triangleq \int_{\mathcal{D}_q} T(dy|(q, x)) = 1 - \sum_{\substack{q' \in \mathcal{Q} \\ q' \neq q}} \int_{g_{q,q'}} T(dy|(q, x)).$$

Similar is the case where $(q, x) \in \mathcal{S}^* \backslash \mathcal{S}$, where we may be interested in the probability that the trajectory is reset into an invariant $q' \neq q$ by the application of $R(\cdot|(q, x), q')$: let us denote this probability $p_{(q,q'),q'}(x)$:

$$p_{(q,q'),q'}(x) \triangleq \int_{\mathcal{D}_{q'}} R(dy|(q, x), q'). \tag{3}$$

Notice that, as the support of T and of R coincides, the contribution of both terms is similar, except for the fact that T is associated with a one-step continuous motion, while R to an instantaneous reset.
Investigating similar quantities for more complex dynamics, or over a longer time interval, involves conditioning the probability backwards in time and ending up with the quantities discussed above. For instance, we may be interested in the following transition, for $q, r, s \in \mathcal{Q}, q \neq r, r \neq s$: $x \in g_{q,r} \xrightarrow{R} \mathcal{D}_r \xrightarrow{T} g_{r,s}$: $p_{(q,r),r}(x)p_{r,s}(\cdot)$. This event is related to the following probability:

$$\begin{aligned} \mathcal{P}\left(\mathbf{s}(1) \in g_{r,s} | \mathbf{s}(0) = (q, x) \in g_{q,r}\right) &= \int_{\mathcal{D}_r} \int_{g_{r,s}} R(dy|(q, x), r) T(dz|(r, y)) \\ &= \int_{\mathcal{D}_r} R(dy|(q, x), r) p_{r,s}(y) \end{aligned} \tag{4}$$

The above quantity shows that the contributions of the one-step probabilities have to be necessarily "averaged" over the influence of the stochastic kernels that precede them. As already mentioned, the interplay between transition and reset probabilities is a characteristic feature of SHS. This will also hold with reference to a particular initial distribution π_0 which may characterize, with its support $A \in \mathcal{B}(\mathcal{S}^*)$, the allowed starting points of the hybrid execution.
The terms in (2)-(3), and multiplications thereof, are then characteristic of the computations we want to perform to study the dynamics of the SHS $\mathcal{H}$. In principle, we may be able to associate a transition probability to each couple of

elements taken from the set of invariants and guards. This would allow to abstract the dynamics of $\mathcal{H}$ into those of a discrete m^2-dimensional MC (where $m = \texttt{card}(\mathcal{Q})$). However, by closely looking at the quantity in (2) [resp. (3)], it can be realized that it is necessary to compute it over the whole invariant $\mathcal{D}_q$ [over the whole guard $g_{q,q'}$], averaged over the contribution of the incoming reset maps $R(\cdot|(\cdot,\cdot),q)$ [the transition kernel $T(\cdot|(q,\cdot))$]. To make sense, these last quantities would have to depend on other probabilities, and so on backwards, until integrating over an initial distribution. This computation is rather unfeasible, and its bottleneck hinges on the dependence of T and R on the continuous component of the hybrid state space.
Rather that aiming at abstracting the dynamics of the SHS $\mathcal{H}$ into an m^2-dimensional MC, we may allow an abstraction into a higher dimensional structure, while improving the computability of the procedure. The technique to achieve this is based on a continuity assumption over the dynamics, and a state-partitioning procedure. This approach is described in the following.

5 Abstraction Procedure

In the following, an abstraction procedure for the SHS model $\mathcal{H}$, as introduced in Section 2, is given. The SHS $\mathcal{H}$ will be abstracted into a Markov set-chain, described by a *one-step* transition set $[\Pi] = [P, Q]$. The computations involved in obtaining the abstraction are reduced to integrations over the continuous part of the hybrid state space. It is argued that this procedure requires some necessary approximations, however explicit bounds on the errors will be obtained. Associating error bounds to the computed transition probabilities directly connects with the theory of MSC, providing a direct definition of the transition set $[\Pi]$. The precision of the abstraction will depend on a parameter δ . It is desirable for the abstraction to show, in the limit as $\delta \to 0$, some convergent properties to the original SHS $\mathcal{H}$.

Approximation of State-dependent Transitions and Resets.
As we discussed in Section 4, the dependence of transition and reset kernels on respectively the invariants and the guards set, and their extended supports, makes the computation of transition probabilities via nested integrals of product terms as in (2)-(3) computationally expensive. Introducing some "regularity assumptions" on the probabilistic kernels, it is be possible to achieve a "memoryless" approximation for these transition probabilities, where their calculation does not depend on the continuous part of the hybrid space $\mathcal{S}$. Indeed, if we aim at obtaining a finite abstraction of $\mathcal{H}$, it makes sense to get a procedure which does not depend an uncountable component of $\mathcal{S}$.
Let us suppose that the stochastic kernels T and R, which depend on the continuous component of the hybrid state in the Definition 1 of the SHS, admit densities respectively t and r. Similarly, let us assume the initial probability distribution π_0 has a density p_0. It is assumed that p_0, t and r satisfy the following Lipschitz condition.

Assumption 1 (Continuity of the Stochastic Kernels)

1. $|p_0(s) - p_0(s')| \leq k_0\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in \mathcal{D}_q^*$,
2. $|t(\bar{x}|s) - t(\bar{x}|s')| \leq k_T\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in \mathcal{D}_q$, *and* $(q,\bar{x}) \in \mathcal{D}_q^*$,

3. $|r(\bar{x}|s,\bar{q}) - r(\bar{x}|s',\bar{q})| \leq k_R\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in \mathcal{D}_q^* \backslash \mathcal{D}_q$, $(\bar{q},\bar{x}) \in \mathcal{D}_q^*$, *and* $\bar{q} \neq q$,

where k_0, k_T *and* k_R *are finite Lipschitz constants.* □

For the sake of computations, we also stress the implicit assumption, raised in Definition 1, that for each $q \in \mathcal{Q}$, the continuous component $\mathcal{D}_q^*$ associated to such mode is a bounded subset of $\mathbb{R}^{n(q)}$.
Let us introduce the following quantities (see Table 1 for a compendium of them), describing the (finite) volume measures of particular subsets of the domains: $\lambda_q^* = \mathcal{L}(\mathcal{D}_q^*), \lambda_q = \mathcal{L}(\mathcal{D}_q), \lambda_{q,r} = \mathcal{L}(g_{q,r}), \lambda = \sum_{q\in\mathcal{Q}} \mathcal{L}(D_q), \lambda^* = \sum_{q\in\mathcal{Q}} \mathcal{L}(D_q^*)$, where $\mathcal{L}$ is the Lebesgue measure of a bounded subset of a Euclidean space. Since $\mathcal{D}_q^* = \mathcal{D}_q \cup \mathcal{G}_q$, it follows that $\forall q \in \mathcal{Q}, \lambda_q = \lambda_q^* - \sum_{\substack{r\in\mathcal{Q}\\ r\neq q}} \lambda_{q,r}$.

Consider a mode $q \in \mathcal{Q}$. Pick any two points $(q,x), (q,x') \in \mathcal{D}_q$. Then, with reference to the quantity in (2) and according to Assumption 1, let us compute, $\forall r \in \mathcal{Q}, r \neq q$,

$$\begin{aligned} |p_{q,r}(x) - p_{q,r}(x')| &= \left| \int_{g_{q,r}} T(dz|(q,x)) - \int_{g_{q,r}} T(dz|(q,x')) \right| \\ &\leq \int_{g_{q,r}} |T(dz|(q,x)) - T(dz|(q,x'))| \leq \lambda_{q,r} k_T \|x - x'\|. \end{aligned}$$

A similar bound is obtained for the case $r = q$, which depends on the quantity λ_q. Furthermore, a similar bound can be found for the quantity in (3): selecting any two points $(q,x), (q,x') \in g_{q,r} \subset \mathcal{D}_q^*, r \neq q$, we have:

$$\begin{aligned} |p_{(q,r),r}(x) - p_{(q,r),r}(x')| &= \left| \int_{\mathcal{D}_r} R(dz|(q,x),r) - \int_{\mathcal{D}_r} R(dz|(q,x'),r) \right| \\ &\leq \int_{\mathcal{D}_r} |R(dz|(q,x),r) - R(dz|(q,x'),r)| \leq \lambda_q k_R \|x - x'\|. \end{aligned}$$

Hybrid State Space Partition.
Let us now introduce a partition of the hybrid state space $\mathcal{S}$ (see Table 1). Recall that $\mathcal{S}$ can be written as $\mathcal{S} = \cup_{q\in\mathcal{Q}}\{q\} \times \mathcal{D}_q^* = \cup_{q\in\mathcal{Q}}\{q\} \times \{\cup_{\substack{r\in\mathcal{Q}\\ r\neq q}} g_{q,r} \cup \mathcal{D}_q\}$.
With regards to a particular mode $q \in \mathcal{Q}$, let us introduce a partition of $\mathcal{D}_q^*$ of cardinality $c_q^\delta = d_q^\delta + \sum_{\substack{r\in\mathcal{Q}\\ r\neq q}} e_{q,r}^\delta$, where the first term d_q^δ refers to the invariant $\mathcal{D}_q$, while the other addends $e_{q,r}^\delta$ refer to the corresponding guard sets $g_{q,r}$. These terms are all greater than or equal to one. Let us introduce their respective measures λ_q^j and $\lambda_{q,r}^k$, so that $\lambda_q = \sum_{j=1}^{d_q^\delta} \lambda_q^j$, and $\lambda_{q,r} = \sum_{k=1}^{e_{q,r}^\delta} \lambda_{q,r}^k$. The above dependence on the parameter δ will be made clear shortly.
We do not impose any structure on the partition, but only require it to respect (that is, to not intersect) the boundaries between domain and guards, and those between each couple of adjacent guards. It is then possible to express the domain $\mathcal{D}_q^*$, associated to mode $q \in \mathcal{Q}$, as the union of the following disjoint sets:

$$\mathcal{D}_q^* = \{\cup_{\substack{r\in\mathcal{Q}\\ r\neq q}} \{\cup_{j=1}^{e_{q,r}^\delta} g_{q,r}^j\}\} \cup \{\cup_{j=1}^{d_q^\delta} \mathcal{D}_q^j\}.$$

We associate a discrete mode to each of these partitions: let us introduce mode q^j for $\mathcal{D}_q^j$, and mode q_r^j for $g_{q,r}^j$. The parameter δ is defined to be

$$\delta = \max_{q\in\mathcal{Q}}\{\max(\epsilon_q, \gamma_q)\}, \text{ where}$$
$$\epsilon_q = \max_{\substack{r\in\mathcal{Q}\\ r\neq q}}{}_{j=1,\ldots,e_{q,r}^\delta} \sup\{\|x-x'\| : x, x' \in g_{q,r}^j\} = \max_{\substack{r\in\mathcal{Q}\\ r\neq q}}{}_{j=1,\ldots,e_{q,r}^\delta} \epsilon_{q,r}^j,$$
$$\gamma_q = \max_{j=1,\ldots,d_q^\delta} \sup\{\|x-x'\| : x, x' \in \mathcal{D}_q^j\} = \max_{j=1,\ldots,d_q^\delta} \gamma_q^j.$$

In other words, δ represents the largest diameter of the partitions defined on $\mathcal{S}$. Let us choose a *representative* point within each single cell introduced through the partition: $\forall q \in \mathcal{Q}, \forall j = 1,\ldots,d_q^\delta$ select a point $\bar{x}_q^j \in \mathcal{D}_q^j$; $\forall q \in \mathcal{Q}, r \neq q, \forall j = 1,\ldots,e_{q,r}^\delta$ select a point $\bar{x}_{q,r}^j \in g_{q,r}^j$.
We will now revisit the computation of the probabilistic quantities (2-3) introduced above, in order to define a MSC that embeds the original SHS $\mathcal{H}$. To this aim, we associate to each element of the above partition a distinct state of the MSC. The upper and lower bounds of the transition probabilities directly define the transition set. As mentioned at the end of Section 4, we show that these bounds for the transition probabilities can be approximately computed by raising memoryless assumptions on the continuous part of the quantities into play. It is clear that the values of these bounds depend on the partition grid given by δ, and on the structure of the dynamics of $\mathcal{H}$. To be more precise, we shall approximate the quantities in (2-3) with ones that will be based on computations performed on the representative points. The quantification of the introduced errors is based on Assumption 1. The new transition probabilities will be intuitively denoted in a similar fashion as the relations in (2-3).
Let us start from the relation in (2):

$$\forall x \in \mathcal{D}_q^j,\ p_{q^j,r^k}(x) = \int_{g_{q,r}^k} T(dy|(q,x)) \approx p_{q^j,r^k}(\bar{x}_q^j), \tag{5}$$

More precisely, $|p_{q^j,q_r^k}(x) - p_{q^j,q_r^k}(\bar{x}_q^j)| \leq \lambda_{q,r}^k k_T \gamma_q^j \leq \lambda^* k_T \delta$. Notice that, if $x \in \mathcal{D}_q^j \subseteq \mathcal{D}_q, T(dy|(q,x)) = T(dy|(q^j,x))$.
Now, focusing on equation (3), we have:

$$\forall x \in g_{q,r}^j,\ p_{(q,r)^j,r^k}(x) = \int_{\mathcal{D}_r^k} R(dy|(q,x),r) \approx p_{(q,r)^j,r^k}(\bar{x}_{q,r}^j), \tag{6}$$

where $|p_{(q,r)^j,r^k}(x) - p_{(q,r)^j,r^k}(\bar{x}_{q,r}^j)| \leq \lambda_r^k k_R \epsilon_{q,r}^j \leq \lambda^* k_R \delta$. Notice that, if $x \in g_{q,r}^j \subseteq \mathcal{D}_q^*, R(dy|(q,x),r) = T(dy|(g_{q,r}^j,x), D_r^k)$.
Similar transition probabilities and bounds can be referred to the initial distribution π_0.

6 Steady state computation using the MSC abstraction

In this section we show that it is possible to infer the asymptotic behavior of the SHS $\mathcal{H}$ using the introduced Markov set-chain abstraction $\mathcal{M}$.
We start by providing an intuitive justification of why the abstraction procedure is able to yield some conclusions about the asymptotic dynamics of the original

component	form	parts	partitions	cardinality	size	diameter
hybrid space of $\mathcal{H}$	$\mathcal{S}^* = \bigcup_{q\in\mathcal{Q}} \{q\} \times \mathcal{D}_q^*$	$\mathcal{D}_q^*$		$\mathfrak{c}^\delta$	λ^*	δ
domain	$\mathcal{D}_q^* = \mathcal{D}_q \cup \mathcal{G}_q$	$\mathcal{D}_q, \mathcal{G}_q$		c_q^δ	λ_q^*	$\gamma_q \vee \epsilon_q$
invariant	$\mathcal{D}_q$		$\mathcal{D}_q^j$	d_q^δ	λ_q	γ_q
invariant sections	$\mathcal{D}_q^j$			1	λ_q^j	γ_q^j
guards	$\mathcal{G}_q = \bigcup_{\substack{r\in\mathcal{Q}\\ r\neq q}} \{g_{q,r}\}$	$g_{q,r}$	$g_{q,r}^j$	$\sum_{\substack{r\in\mathcal{Q}\\ r\neq q}} e_{q,r}^\delta$	$\lambda_q^* - \lambda_q$	ϵ_q
guard sections	$g_{q,r}^j$			1	$\lambda_{q,r}^j$	$\epsilon_{q,r}^j$

Table 1. Components of the SHS and elements of the partition that yields the MSC, with corresponding quantities of interest.

system.

As argued before, the possibility of finding explicit bounds for the errors associated to the approximate computations of the transition probabilities of the MSC allows to introduce a "conservative estimate" of the actual transition probabilities between regions of the state space of the original dt-SHS. Qualitatively, the key point is that if we select a small enough bound for the diameters of the partition, the contractive nature of $\mathcal{M}$ dominates over the approximation errors. The contractivity of $\mathcal{M}$ depends on the dynamics and the structure of $\mathcal{H}$. Furthermore, the procedure suggests that, by tuning the parameter for the partition, we can estimate the asymptotic behavior of $\mathcal{H}$ with arbitrary precision. In passing, this concept is most likely related to the fact that stable linear systems admit approximately bisimilar finite abstractions [21].

We now make the above discussion quantitative. Given a desired precision $\varepsilon > 0$, we integrate the procedure for the partition of $\mathcal{H}$ into an algorithm, which computes the steady of the MSC abstraction $\mathcal{M}$ until a maximum error of $\varepsilon > 0$ is obtained for a partition parameter $\delta(\varepsilon)$. As discussed above, the steady state vector $[\pi_\infty]$ for $\mathcal{M}$ is an estimate of the invariant measure π_∞ for $\mathcal{H}$, with a confidence bound given by the diameter $\Delta([\pi_\infty])$. The first idea would be to initialize a partition according to a value $\delta(\varepsilon)$, which guarantees a precision ε for the steady state computation. To this aim, we need to relate δ and ε. The transition set $[\Pi] = [P, Q]$ as constructed in the previous section has the following property:

$$\Delta([\Pi]) \leq ||Q - P|| = (\lambda^* k_T \delta)^{\mathfrak{c}^\delta} .$$

with λ^*, k_T constants as defined in the above section. In order to achieve a precision ε in the steady state computation, a sufficient condition to achieve $\Delta([\pi_\infty]) \leq \varepsilon$ is the following:

$$(\lambda^* k_T \delta)^{\mathfrak{c}^\delta} \leq \varepsilon \left(1 - \mathcal{T}(|[\Pi]|)\right). \tag{7}$$

It is clear that if $\mathcal{T}(|[\Pi]|) < 1$ there always exists a value of $\delta(\varepsilon)$ that satisfies this inequality, since the LHS expression goes to zero as δ goes to zero. The main issue is however that, without an idea of the transition probabilities that

define $[\Pi]$, one cannot estimate $\mathcal{T}([\Pi])$. Since $0 \leq \mathcal{T}([\Pi]) \leq 1$, the set of feasible values for δ that satisfy equation (7) ranges from a finite upper bound δ_0 (when $\mathcal{T}([\Pi]) = 0$) to 0 (when $\mathcal{T}([\Pi]) = 1$). This makes sense: until we have no information about the contractive nature of $[\Pi]$, there is no possibility to estimate the limit set behavior. For this reason, it is impossible to establish a priori a value for δ that guarantees a desired precision in the steady state computation using the abstraction $\mathcal{M}$. However, we can choose an "optimistic" initial value δ_0 as defined above, and apply the following iterative algorithm:

Algorithm 1 (Compute steady state of $\mathcal{H}$ with precision ε)

input: $(\mathcal{H}, \varepsilon)$;

initialize integer $k = 0$ and real $\mathcal{T}^*(0) = 0$;

for $k \geq 0$

if $\mathcal{T}^*(k) = 1$ or $\mathcal{T}^*(k) = \mathcal{T}^*(k-1)$

then set $\delta(k) = a\delta(k-1), a < 1$;

else set $\delta(k)$ such that equation (7) holds with $\mathcal{T}([\Pi(k)]) = \mathcal{T}^*(k)$;

compute $\mathcal{M}(k)$ defined by $[\Pi(k)], [\pi_0(k)]$ according to $\delta(k)$;

if $\frac{\Delta([\Pi(k)])}{1-\mathcal{T}^*([\Pi(k)])} \leq \varepsilon$ **then** exit;

else set $\mathcal{T}^*(k+1) \triangleq \mathcal{T}^*([\Pi(k)])$ and $k \triangleq k+1$;

end ;

compute L_∞, H_∞ bounding $[\pi_\infty]$ for $\mathcal{M}(k)$;

output: (L_∞, H_∞). □

Notice that, $\forall k \geq 0, \delta(k+1) < \delta(k)$. Iterating the algorithm, we obtain a sequence $\{[\Pi(k)]\}_{k\geq 0}$, with $[\Pi(k+1)] \subset [\Pi(k)]$. This implies that $\mathcal{T}([\Pi(k)])$ is non-increasing with k. Moreover, if $\mathcal{T}_\infty = \lim_{k\to+\infty} \mathcal{T}([\Pi(k)]) < 1$, then $\mathcal{T}([\Pi(k)])$ is decreasing. Namely, if $\mathcal{T}_\infty < 1$ it follows that we can arbitrarily increase the accuracy of our abstraction until

$$(\lambda^* k_T \delta)^{c^\delta} \leq \varepsilon (1 - \mathcal{T}_\infty). \tag{8}$$

When this happens the algorithm terminates, and we compute $[\pi_\infty]$ using the upper and lower bounding matrices L_∞, H_∞, as described in Section 3.
We now discuss the computational burden of our procedure. It is clear that the main bottlenecks are (1) the abstraction procedure for the partitioning of the hybrid state space; and (2) the limit set computation on the abstraction MSC. The first computation directly depends on the parameter δ, which is related to $\mathcal{T}_\infty$ by (8). The second computation depends on two parameters: the cardinality of the MSC and the convergence speed. The state cardinality is $\mathfrak{c}^\delta$, and depends on δ, while the convergence speed can be related to $\mathcal{T}_\infty$ by (1). The main weight in the computational complexity of our abstraction procedure is $\mathcal{T}_\infty$.
For the above arguments, it can be interesting to interpret $\mathcal{T}_\infty$ as the coefficient of ergodicity $\mathcal{T}(\mathcal{H})$ of the SHS $\mathcal{H}$, and possibly compare this value with other convergence bounds directly derived on the structure and the dynamics of $\mathcal{H}$.

7 Numerical Study

We implement the proposed abstraction procedure on a simple one-dimensional dynamical system, whose dynamics is described by the following SDE, defined for $t \geq 0$:

$$d\mathbf{X}_t = f(\mathbf{X}_t)dt + d\mathbf{B}_t, \text{ with } \mathbf{X}_0 \sim \mathcal{U}(A). \tag{9}$$

The drift depends on a function $f : \mathbb{R} \to \mathbb{R}$, assumed to be continuous and bounded. The term $\mathbf{B}_t$ denotes a standard Wiener process. $\mathcal{U}(A)$ is the uniform distribution, over some compact set $A \subset \mathbb{R}$.

The study in [22] proposes an abstraction of the system in relation (9), which is obtained by discretizing the state space $\mathbb{R}$ into overlapping intervals, and by introducing a dt-SHS model. The discrete structure of such a dt-SHS is a birth-death MC. The transition probabilities between a node and its neighbors of the MC are computed through the first hitting times for the solution process within the domain associated to that node. Interestingly, these probabilities have a closed form in the one-dimensional case. This procedure would also naturally induce a discretization of time, as the probabilistic dynamics are only looked at the transition instants. Some convergence properties of the discretized model to (9) are shown.

In the following, we again introduce a dt-SHS as an abstraction of (9). While a discretization of the state space similar to [22] is possible, we decide to introduce an easier one. Again, a MC structure is obtained, whose transition probabilities are computed by the technique proposed in this work. This will also yield some explicit error bounds. Unlike the abstraction in [22], the one proposed in this work is prone to be extended to continuous spaces of dimension higher than one. Time will be directly discretized at the onset: rather non-stringent sufficient conditions can be given for the convergence of this procedure to the original continuous-time process [23]. More precisely, the SDE in (9) is discretized in time according to a first-order Euler-Maruyama scheme, with discretization step $\Delta > 0$, to obtain the following, for any $n \geq 0$: $\mathbf{X}_{(n+1)\Delta} \sim \mathcal{N}\left(\mathbf{X}_{n\Delta} + \Delta f(\mathbf{X}_{n\Delta}), \Delta\right)$, where $\mathcal{N}(m, \sigma)$ is a normal random variable with mean m and variance σ^2.

For computational necessity, we shall introduce some approximation outside the compact interval $\mathcal{K} = [-K, K]$. Let us partition this interval $\mathcal{K}$ into $2l$ sections of length 2δ, where $\delta = K/l$, and centered at the points $\bar{x}_k = -K + (2k-1)\delta, k = l, 2, \ldots, l$. Call these partitions $\mathcal{D}_k = [-K + 2(k-1)\delta, -K + 2k\delta]$.

Additionally, consider two regions for the open intervals $\mathcal{D}_{lb} = (-\infty, -K], \mathcal{D}_{ub} = [K, +\infty)$, "centered" at the points $\bar{x}_{lb,ub} = \{\pm(K+\delta)\}$. Consider for convenience the extended index set $\mathcal{Q} = \{l, 2, \ldots, l, lb, ub\}$. Conditional on $\mathbf{X}_{n\Delta} = \bar{x}_k$, for

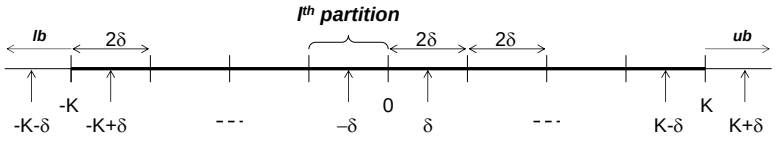

Fig. 1. Abstraction procedure for the one-dimensional system in (9).

any $k \in \mathcal{Q}$, the process at time $(n+1)\Delta$ is distributed according to $T(\cdot|\bar{x}_k) \triangleq \mathcal{N}(\cdot; m_k, \Delta)$, where $m_k = \bar{x}_k + \Delta f(\bar{x}_k)$.
This discretization procedure induces a dt-SHS, where the $l+2$ domains make up the state space as $\mathcal{S} = \bigcup_{k\in\mathcal{Q}}\{k\} \times \mathcal{D}_k$. The continuous dynamics are characterized, for any $(k,x) \in \mathcal{S}$, by the kernels $T(\cdot|(k,x)) = \mathcal{N}(\cdot; x + \Delta f(x), \Delta)$. The guards set is intuitively made up of the points in $\mathcal{G} = \{-K + 2(k-1)\delta, k = l, 2, \ldots, l\}$, and the resets coincide with the transition kernels, that is, for any $k, k' \in \mathcal{Q}, k' \neq k, R(\cdot|(k,x),k') = T(\cdot|(k,x))$.
Given a standard random variable $x \sim \mathcal{N}(0,1), x \in \mathbb{R}$, with its associated density $\phi(x)$ and distribution $\Phi(x)$, we can express the distribution of $\tilde{x} \sim \mathcal{N}(m,\sigma)$ as $\Phi_{\tilde{x}}(x) = \Phi(\frac{x-m}{\sigma})$. The Gaussian distributions associated to the transition kernels are locally Lipschitz. Given two points $x, x' \in \mathbb{R}$, it is easy to set up the following bound:

$$|\phi(x) - \phi(x')| \leq \frac{1}{\sqrt{2\pi}} \frac{|e^K - 1|}{K} \|x - x'\| \leq \frac{2\delta}{\sqrt{2\pi}} \frac{|e^K - 1|}{K}. \tag{10}$$

Consider a mode $k \in \mathcal{Q}\backslash\{lb, ub\}$. Let us compute the approximate transition probabilities between the different modes of the introduced dt-SHS as follows, for any $h \in \mathcal{Q}\backslash\{lb, ub\}$, based on the representative point $\bar{x}_k \in \mathcal{D}_k$:

$$p_{k,h}(\bar{x}_k) = \Phi\left(\frac{(\bar{x}_h + \delta) - \bar{x}_k}{\Delta}\right) - \Phi\left(\frac{(\bar{x}_h - \delta) - \bar{x}_k}{\Delta}\right),$$
$$p_{k,lb}(\bar{x}_k) = \Phi\left(\frac{-K - \bar{x}_k}{\Delta}\right), \quad p_{k,ub}(\bar{x}_k) = 1 - \Phi\left(\frac{K - \bar{x}_k}{\Delta}\right).$$

Based on (10) and on the size of the domains, we associate to each transition probability the RHS error $\frac{(2\delta)^2}{\sqrt{2\pi}} \frac{|e^K-1|}{K}$. [3]

Ornstein-Uhlenbeck process. In the following, we implement some computations for the very special linear drift case, i.e. where $f(x) = -\mu x, \mu > 0$. The knowledge of a closed form distribution for this process [24] enables a comparison of it with the outcome of the simulations. We have chosen the following parameters: $K = 15, \Delta = 1, m = 0, \sigma = 1$. Choosing a $\mu = 0.5$, the solution process of (9) is trivially distributed as $\mathcal{N}(0,1)$. We have implemented our abstraction procedure and the MSC basic algorithms on Matlab. Figure 2 illustrates the result we obtained using an abstracting MSC with a state cardinality of 50. The Table shows that by augmenting the state space of the MSC in our simulations, that is by lowering the refinement constant, the error bounds for the steady state converge to zero.

8 Conclusions

The main contribution of this paper is to provide an abstraction algorithm for dt-SHS. The abstraction is into the class of MSC. We derive a relation between

[3] Again, we inescapably introduce an approximation on the otherwise possibly infinite error bounds for the probabilities associated to transition to and from the domains lb, ub.

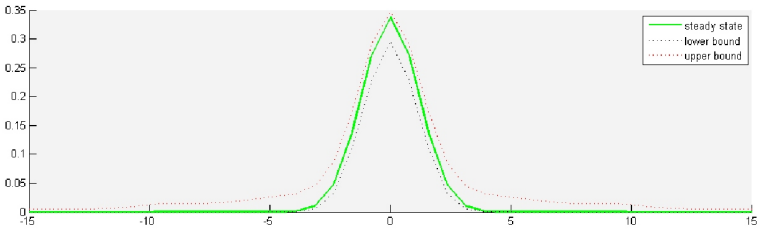

MSC state space cardinality	30	50	100	200
Steady state estimation error	0.06	0.05	0.04	0.03

Fig. 2. Simulation outputs.

the number of discrete states of the abstraction, the structure of the continuous dynamics of the original system, and the desired precision in estimating the steady state. As a first result towards more general probabilistic model checking of SHS, the abstraction is employed to assess the steady state of the original system with arbitrary precision. Other future directions aim at extending the results to more general SHS models.

References

1. Clarke, E., Grumberg, O., Peled, D.: Model Checking. The MIT Press, Cambridge, Massachusetts (2002)
2. Alur, R., Henzinger, T., Ho, P.H.: Automatic symbolic verification of embedded systems. IEEE Trans. on Software Engineering **22** (1996) 181–201
3. Alur, R., Henzinger, T., Lafferriere, G., Pappas, G.: Discrete abstractions of hybrid systems. Proceedings of the IEEE **88(2)** (July 2000) 971–984
4. Pappas, G.: Bisimilar linear systems. Automatica **39(12)** (2003) 2035–2047
5. Girard, A., Pappas, G.: Approximation metrics for discrete and continuous systems. IEEE Transactions on Automatic Control (accepted)
6. Henzinger, T.A., Majumdar, R., Prabhu, V.: Quantifying similarities between timed systems. In: Proceedings of the Third International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS). Volume 3829 of Lecture Notes in Computer Science., Springer (2005) 226–241
7. D'Innocenzo, A., Julius, A., Di Benedetto, M., Pappas, G.: Approximate timed abstractions of hybrid automata. In: Proceedings of the 46^{th} IEEE Conference on Decision and Control. New Orleans, Louisiana, USA. (2007)
8. Julius, A., Pappas, G.: Approximate abstraction of stochastic hybrid systems. IEEE Trans. Automatic Control (accepted)
9. Lichtenberg, G., Rostalski, P.: Using path integral short time propagators for numerical analysis of stochastic hybrid systems. In: Proceedings of the 2nd IFAC Conference on Analysis and Design of Hybrid Systems, Alghero, Italy. (2006) 179–184
10. Gillespie, D.: Exact stochastic simulation of coupled chemical reactions. Physical Chemistry **81,25** (1977) 2340–2361

11. Kushner, H.J.: Approximation and Weak Convergence Methods for Random Processes with Applications to Stochastic Systems Theory. MIT Press, Cambridge, Massachussets (1984)
12. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.: Model-checking continuous time markov chains. ACM Trans. on Comp. Logic **1(1)** (2000) 162–170
13. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time markov chains. IEEE Transactions on Software Engineering **29(6)** (2003) 524–541
14. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In Thiagarajan, P., ed.: Proc. 15th Conference on Foundations of Software Technology and Theoretical Computer Science. Volume 1026 of Lecture Notes in Computer Science. Springer Verlag (1995) 499–513
15. Kwiatkowska, M., Norman, G., Parker, D.: Stochastic model checking. In Bernardo, M., Hillston, J., eds.: Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM07). Volume 4486 (Tutorial Volume) of Lecture Notes in Computer Science. Springer (2007) 220–270 To appear.
16. Amin, S., Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Reachability analysis for controlled discrete time stochastic hybrid systems. In Hespanha, J., Tiwari, A., eds.: Hybrid Systems: Computation and Control. Lecture Notes in Computer Science 3927. Springer Verlag (2006) 49–63
17. Davis, M.H.A.: Markov Models and Optimization. Chapman & Hall/CRC Press, London (1993)
18. Bujorianu, M.L., Lygeros, J.: Toward a general theory of stochastic hybrid systems. In Blom, H., Lygeros, J., eds.: Stochastic Hybrid Systems. LNCIS 337. Springer Verlag (2006) 3–30
19. Hartfiel, H.J.: Markov Set-Chains. Volume 1695 of Lecture Notes in Mathematics. Springer-Verlag Berlin Heidelberg (1998)
20. D'Innocenzo, A., Di Benedetto, M.D., Di Gennaro, S.: Observability of hybrid automata by abstraction. In Hespanha, J., Tiwari, A., eds.: Hybrid Systems: Computation and Control. Volume 3927 of Lecture Notes in Computer Science. Springer Verlag (2006) 169–183
21. Girard, A.: Approximately bisimilar finite abstractions of stable linear systems. In: Hybrid Systems: Computation and Control. Volume 4416 of Lecture Notes in Computer Science. Springer Verlag (2007) 231–244
22. Hu, J., Lygeros, J., Sastry, S.: Towards a theory of stochastic hybrid systems. In Lynch, N., Krogh, B., eds.: Hybrid Systems: Computation and Control. Lecture Notes in Computer Science 1790. Springer Verlag (2000) 160–173
23. Milstein, G.: Numerical Integration of Stochastic Differential Equations. Kluwer Academic Publishing (1994)
24. Øksendal, B.: Stochastic Differential Equations: An Introduction with Applications. 6th edn. Springer-Verlag (2003)