Proceedings of the
47th IEEE Conference on Decision and Control
Cancun, Mexico, Dec. 9-11, 2008

Approximate Abstractions of Discrete-Time Controlled Stochastic Hybrid Systems

Alessandro D'Innocenzo, Alessandro Abate, and Maria D. Di Benedetto

***Abstract*— This work presents a procedure to construct a finite abstraction of a controlled discrete-time stochastic hybrid system. The state space and the control space of the original system are partitioned by finite lattices, according to some refinement parameters. The errors introduced by the abstraction procedure can be explicitly computed, over time, given some continuity assumptions on the original model. We show that the errors can be arbitrarily tuned by selecting the partition accuracy. The obtained abstraction can be interpreted as a controlled Markov set-Chain, and can be used both for verification and control design purposes. We test the proposed technique to analyze a model from systems biology.**

I. Introduction

The dynamical analysis of complex, high-dimensional, hybrid and stochastic models poses many challenges, both at a theoretical and at a computational level. One relevant methodology to tackle this issue is known as abstraction: a "simpler" system with smaller state space (possibly finite), and *equivalent* to the original system, is searched for. Equivalence is usually defined by the notions of language equivalence and bisimulation [4]. Often though such (exact) notions are quite restrictive, since they require a perfect correspondence between the trajectories of the original system and those of its abstraction. To address this potential limitation, *approximate* notions of system equivalence [6], [7], [10], [13], [14] have been recently developed to relax the abstraction procedure. According to this approach, a proper metric is introduced to quantify the distance between the (trajectories of the) original system and (those of) the approximation.

The present line of research looks at abstractions of Hybrid Systems (HS), models known for their generality and complexity. The contribution in [5] proposes an algorithm to construct an approximate abstraction of a HS by means of a timed automaton, which is a model with simpler continuous dynamics. In [13] a notion of approximate bisimilarity is proposed for the class of probabilistic models known as "jump linear stochastic systems."

In this work we provide new results on approximate abstractions of discrete-time controlled stochastic hybrid systems (dt-cSHS), which encompass a number of other classes of stochastic hybrid systems (SHS). The recent work in [1] has introduced an approximate abstraction for a class of SHS, and formalized a computation of a bound on the error associated with this abstraction. By reinterpreting the new model as a Markov set-Chain (MSC) [9], the authors have investigated the asymptotic behavior of the original SHS via that of the MSC. Furthermore, the work has proposed an algorithm which, given a desired precision on the steady-state error, finds an adequate refinement parameter and synthesizes an abstraction according to that parameter. The present contribution extends that in [1] in three directions. First, the considered SHS model is more general. Second, the model is control-dependent, which requires a specific partition of the control space, and an integration of the errors on the state and on the control. Third, the work derives explicit bounds on the error between the transition probabilities of the abstracted model and those of the original model (considered over the regions of the partition) along time (and in particular in steady-state). The abstraction can be used both for verification purposes and control design tasks. This work represents another step towards a formalization of the notion of stochastic bisimulation, which is the ultimate goal of this line of research.

The paper develops as follows. Section II introduces the SHS model, namely the discrete-time, controlled SHS (dt-cSHS). Section III recalls some results on MSC, which will be utilized in the following. Section IV introduces the abstraction procedure, which turns the original dt-cSHS into a controlled MSC. Section V delves into the computation of the errors associated to the abstraction. Finally, in section VI we test the proposed abstraction technique on a model drawn from biology, which describes the biosynthesis of the antibiotic subtilin by the soil bacterium *Bacillus subtilis*. We employ the abstraction framework to investigate its steady-state.

II. Discrete-Time Controlled Stochastic Hybrid Systems

Definition 1: A discrete-time controlled stochastic hybrid system is a tuple $\mathcal{H} = (\mathcal{S}, \mathcal{A}, T_q, T_t, T_r)$, where

- $\mathcal{S} := \cup_{i \in \mathcal{Q}} \{i\} \times \mathcal{D}_i$, is the hybrid state space, that consists of a set of discrete states $\mathcal{Q} := \{q_1, q_2, \ldots, q_m\}$, for some finite $m \in \mathbb{N}$, and by a set of continuous "domains" for each mode $i \in \mathcal{Q}$, each of which is defined to be a compact subset $\mathcal{D}_i \subset \mathbb{R}^{n(i)}$. The function $n : \mathcal{Q} \to \mathbb{N}$ assigns to each $i \in \mathcal{Q}$ the dimension of the continuous state space $\mathbb{R}^{n(i)}$;
- $\mathcal{A}$ is the control space, a continuous and compact Borel subset of $\mathbb{R}^p$;
- $T_q : \mathcal{Q} \times \mathcal{S} \times \mathcal{A} \to [0, 1]$ is a discrete stochastic kernel (the "discrete transition kernel") on $\mathcal{Q}$, given $\mathcal{S} \times \mathcal{A}$, which assigns to each $(q, x) \in \mathcal{S}$ and $a \in \mathcal{A}$, a discrete probability distribution over $\mathcal{Q}$: $T_q(q|(q, x), a)$;
- $T_t : \mathcal{B}(\mathcal{D}_{(\cdot)}) \times \mathcal{S} \times \mathcal{A} \to [0, 1]$ is a Borel-measurable stochastic kernel (the "continuous transition kernel") on $\mathcal{D}_{(\cdot)}$, given $\mathcal{S} \times \mathcal{A}$, which assigns to each $(q, x) \in \mathcal{S}$ and $a \in \mathcal{A}$ a probability measure on the Borel space $(\mathcal{D}_q, \mathcal{B}(\mathcal{D}_q))$: $T_t(dx|(q, x), a)$;

This work was supported by the NIH Grant 5 R01 GM075311, by the European Commission under project IST NoE HYCON contract n. 511368 and STREP project n. TREN/07/FP6AE/S07.71574/037180 IFLY. The first two authors have equally contributed to this work.

Alessandro D'Innocenzo and Maria D. Di Benedetto are with the Department of Electrical Engineering and Computer Sciences, University of L'Aquila – {adinnoce,dibenede}@ing.univaq.it

Alessandro Abate is with the Department of Aeronautics and Astronautics, Stanford University – aabate@stanford.edu

- $T_r : \mathcal{B}(\mathcal{D}_{(\cdot)}) \times \mathcal{S} \times \mathcal{A} \times \mathcal{Q} \to [0,1]$ is a Borel-measurable stochastic kernel (the "reset kernel") on $\mathcal{D}_{(\cdot)}$, given $\mathcal{S} \times \mathcal{A} \times \mathcal{Q}$, that assigns to each $(q,x) \in \mathcal{S}$, $a \in \mathcal{A}$, and $q' \in \mathcal{Q}, q' \neq q$, a probability measure on the Borel space $(\mathcal{D}_{(q')}, \mathcal{B}(\mathcal{D}_{(q')}))$: $T_r(dx|(q,x),a,q')$. □

The system initialization at the initial time (say $k=0$) may be specified by 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 $\mathcal{D}_q$.

The model is inspired by that in [3]. However, unlike this last source, for the sole sake of simplicity we do not distinguish between a control input acting on the discrete dynamics or on the continuous ones. Moreover, this model is an extension of the similar framework in [1] by the introduction of the control set. Furthermore, in [1] the change of mode depends on the verification of spatial conditions, which are expressed as subsets of a domain of the hybrid state space: for instance, $g_{q,q'} \subset \mathcal{D}_q$ denotes a spatial condition associated with the change of discrete mode from q to q'. Instead in this work the execution is allowed to change domain (say, q) according to the probability law of the discrete kernel T_q, which is defined on the whole domain $\mathcal{D}_q$. Notice that the spatial conditions in [1] may be obtained by assuming that T_q has an indicator-like structure:

$$T_q(q'|(q,x),a) = \begin{cases} 1, & \text{if } (q,x) \in g_{q,q'}, \forall a \in \mathcal{A}, \\ 0, & \text{else,} \end{cases} \quad (1)$$

where, as in [1], $g_{q,q'} \subset \mathcal{D}_q$. Notice that according to the law T_q a "discrete event" may also be affected by the choice of a particular control $a \in \mathcal{A}$. We refer the reader to the details contained in [3] for further insights on the model. Let us report here the definition of finite and infinite execution process, after introducing the concept of control feedback.

***Definition** 2:* Given a dt-cSHS $\mathcal{H}$, a control feedback for $\mathcal{H}$ is a Borel measurable function $\nu : \mathcal{S} \to \mathcal{A}$, which associates to each hybrid state $s \in \mathcal{S}$ a control action $\nu(s) \in \mathcal{A}$. We denote by $\mathcal{C}$ the set of measurable functions $\nu : \mathcal{S} \to \mathcal{A}$.
A vector of these control functions over a time horizon $[0, \ldots, N]$ is called policy, and belongs to $\mathcal{C}^{N+1}$. □

The control functions $\nu : \mathcal{S} \to \mathcal{A}$ are assumed to be deterministic, that is they are non-randomized. Also, notice the Markov property for such control functions, i.e. their dependence solely on the current value of the state.

***Definition** 3:* Given a dt-cSHS $\mathcal{H}$, an initial distribution π_0 and a control policy $\nu \in \mathcal{C}^{N+1}$, an execution of $\mathcal{H}$ is a stochastic process[1] $\{\mathbf{s}(k) = (\mathbf{q}(k), \mathbf{x}(k)) : \forall k = 0, \ldots, N+1, \mathbf{s}(k) \in \mathcal{S}\}$ generated by 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,*

extract a value $q_{k+1} \in \mathcal{Q}$ for $\mathbf{q}(k+1)$, according to $T_q(\cdot\,|s_k, \nu(s_k))$;

***if** $q_{k+1} \neq q_k \in \mathcal{Q}$,*

***then** extract a value $x_{k+1} \in \mathcal{D}_{q_{k+1}}$ for $\mathbf{x}(k+1)$ according to $T_r(\cdot\,|s_k, \nu(s_k), q_{k+1})$;*

***else** extract a value $x_{k+1} \in \mathcal{D}_{q_{k+1}}$ for $\mathbf{x}(k+1)$, according to $T_t(\cdot\,|s_k, \nu(s_k))$;*

***end**.* □

[1]In this work bold symbols denote (stochastic) processes, while a regular typeset is used for points on the state space.

It is understood that, when $N = \infty$, the algorithm does not terminate. We make use of the following shortened notation for the probability kernels, where $q, q' \in \mathcal{Q}, q \neq q', s = (q,x) \in \mathcal{S}, a \in \mathcal{A}$:

$$T(ds|(q,x),a,q') = \begin{cases} T_q(q'|(q,x),a) T_t(dx|(q,x),a), \\ \\ T_q(q'|(q,x),a) T_r(dx|(q,x),a,q'). \end{cases}$$

III. Markov set-Chains

We recall here the concept of Markov set-Chain, which in this paper is later leveraged to prove properties of the abstraction. The results are from [9] and references therein.

***Definition** 4:* [9, Definition 2.5] Let $P, Q \in \mathbb{R}^{n \times n}$ be nonnegative matrices (not necessarily stochastic) with $P \leq Q$. We define a transition set as:

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

□

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 at each time step within the compact transition set $[\Pi]$. More formally,

***Definition** 5:* [9, Definition 2.5] 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, as in Definition 4. 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.

***Definition** 6:* [9, Definition 1.2] 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, a_j are the i–th, j–th rows of A. □
It can be shown that the condition $\mathcal{T}(A) < 1$, along with the irreducibility of the chain, implies the existence of a unique and invariant limiting distribution for the associated Markov Chain [9]. The above definition can be directly extended to Markov set-Chains:

***Definition** 7:* [9, Definition 3.1] For any transition set $[\Pi]$, its coefficient of ergodicity is defined as follows:

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

Since $\mathcal{T}(\cdot)$ is a continuous function and $[\Pi]$ a compact set, the corresponding maximum argument exists. Similar to the simpler case of Markov chains, the quantity $\mathcal{T}([\Pi]) \in [0,1]$ provides a measure of the "contractive" nature of the Markov set-Chain: the smaller $\mathcal{T}([\Pi])$, the "more contractive" the MSC. This will become clear when studying the asymptotic properties of the MSC, and is related to the regularity properties of the matrices that build up the MSC [9]. The following notion connects to Definition 7:

***Definition** 8:* [9, Definition 3.3] Let a MSC with transition set $[\Pi]$ be given. Suppose r is a positive integer such that $\mathcal{T}(A_1 \dots A_r) < 1, \forall A_1, \cdots, A_r \in [\Pi]$. Then $[\Pi]$ is said to be product scrambling and r its scrambling integer. $\square$

We are interested in the minimum possible scrambling integer, since for any positive $k, r, s \in \mathbb{N} : k = r+s, \mathcal{T}([\Pi]^k) \leq \mathcal{T}([\Pi]^r)\mathcal{T}([\Pi]^s)$. For this reason, given a MSC $[\Pi]$, it is possible that $\mathcal{T}([\Pi]) = 1$, yet that there exists $k > 1$ such that $\mathcal{T}([\Pi]^k) < 1$. We now illustrate some results on the convergence of MSC.

***Theorem** 1:* [9, Theorem 3.4] Given a product scrambling MSC with transition set $[\Pi]$ and initial distribution set $[\pi_0]$, there exists a unique limit set $[\pi_\infty]$ that is an invariant, i.e. such that $[\pi_\infty][\Pi] = [\pi_\infty]$. Thus

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

Moreover, let r be the scrambling integer. Then for any positive integer k,

$$d_h([\pi_k],[\pi_\infty]) \leq \alpha\beta^k, \qquad (2)$$

where $\alpha = [\mathcal{T}([\Pi]^r)]^{-1} d_h([\pi_0],[\pi_\infty]), \beta = \mathcal{T}([\Pi]^r)^{\frac{1}{r}} < 1$ and d_h is the Hausdorff distance. $\square$

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 upper bound for the diameter of the limit set $[\pi_\infty]$ when it exists.

***Theorem** 2:* [9, Theorems 3.9, 3.11] Given a product scrambling Markov set-Chain with transition set $[\Pi] = [P, Q]$ and such that $\mathcal{T}([\Pi]) < 1$, then

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

Notice that the results in Theorem 2 can be intuitively extended to the case of ergodicity with finite scrambling integer $r > 1$, by replacing $\mathcal{T}([\Pi])$ with $\mathcal{T}([\Pi]^r)$, and by considering P, Q as the bounding matrices of $[\Pi]^r$.

IV. Abstraction Procedure

The abstraction for state and control spaces proposed in this work can be used with two possible goals (see section V-A).

From a *verification* perspective, chosen a particular continuous control feedback over a continuous plant, we can exploit a finite abstraction to verify properties of the original system. It is desirable to claim that, in the limit as the approximation error associated with the abstraction goes to zero, the property to be verified is true on the original system if and only if it is true on the abstraction [7]. In this verification instance, we propose a Markov set-Chain [9] as the abstraction framework for the original *closed-loop* system.

On the other hand, from a *design* point of view, given a continuous controllable plant we can obtain a finite abstraction over which we can synthesize a quantized control feedback (e.g. using classical MDP algorithms [15]), while guaranteeing that the closed-loop behavior of the original system is close to that of the abstract system with a required precision. In this case, we propose a controlled Markov set-Chain (cMSC) [11] as the abstraction framework for the original *control* system. This structure is equivalent to that of Bounded-parameter Markov Decision Process (BMDP) [8].

In order to attain the convergence properties mentioned above, it is necessary to enforce some continuity on the dynamics of the dt-cSHS. Suppose that the continuous stochastic kernels T_t, T_r admit densities t, r. Let us raise the following conditions:

***Assumption** 1:*

1) $|T_q(\bar{q}|s,a) - T_q(\bar{q}|s',a')| \leq L_q\|x - x'\| + M_q\|a - a'\|$, for all $s = (q,x), s' = (q,x') \in \mathcal{D}_q, a, a' \in \mathcal{A}$ and $\bar{q} \in \mathcal{Q}$;
2) $|t(\bar{x}|s,a) - t(\bar{x}|s',a')| \leq L_t\|x - x'\| + M_t\|a - a'\|$, for all $s = (q,x), s' = (q,x') \in \mathcal{D}_q, a, a' \in \mathcal{A}$, and $(q,\bar{x}) \in \mathcal{D}_q$;
3) $|r(\bar{x}|s,a,\bar{q}) - r(\bar{x}|s',a',\bar{q})| \leq L_r\|x - x'\| + M_r\|a - a'\|$, for all $s = (q,x), s' = (q,x') \in \mathcal{D}_q, a, a' \in \mathcal{A}$, $(\bar{q},\bar{x}) \in \mathcal{D}_{\bar{q}}$, and $\bar{q} \in \mathcal{Q}, \bar{q} \neq q$;

$L_q, L_t, L_r, M_q, M_t, M_r$ are finite positive constants. $\square$

Consider now any control feedback ν defined on the system $\mathcal{H}$. Let us raise the following additional continuity assumption:

***Assumption** 2:* For any control feedback $\nu \in \mathcal{C}$,

$$|\nu(s) - \nu(s')| \leq L_a\|x - x'\|, \forall s = (q,x), s' = (q,x') \in \mathcal{D}_q,$$

where L_a is a finite and positive constant. $\square$

A. State Space Partition

Let us recall that $\mathcal{D}_q \subset \mathbb{R}^{n(q)}$ is required to be a compact set, for each $q \in \mathcal{Q}$. We introduce a finite partition of the hybrid state space $\mathcal{S} = \cup_{q\in\mathcal{Q}}\{q\} \times \mathcal{D}_q$. For each $q \in \mathcal{Q}$, such a partition $\{\mathcal{D}_q\}_\delta$ is made up of non-overlapping subsets of $\mathcal{D}_q$ and can be general in its shape. However it is usually introduced as a uniform grid of width δ in $\mathbb{R}^{n(q)} \cap \mathcal{D}_q$, where δ is the diameter of the partition, that is the maximum distance between any two points in the same equivalence class (i.e., in the same element of the partition). In the following, unless otherwise stated, we will be working with a uniform partition, parameterized by its diameter δ. The set $\{\mathcal{S}\}_\delta = \bigcup_{q\in\mathcal{Q}} \{q\} \times \{\mathcal{D}_q\}_\delta$ is then a partition of the whole $\mathcal{S}$.

Given any $s = (q,x) \in \mathcal{S}$ there exists an element $\langle s \rangle \in \{\mathcal{S}\}_\delta$ such that the point $s \in \langle s \rangle$. It is clear that any $\langle s \rangle \in \{\mathcal{S}\}_\delta$ is a subset of the hybrid state space, i.e. $\langle s \rangle \subseteq \mathcal{S}$, and that $\langle s \rangle$ belongs to only a single domain $\mathcal{D}_q$. Let us select any point $\bar{s} = (\bar{q},\bar{x}) \in \langle s \rangle \subseteq \mathcal{S}$ to be the *representative* point of the set $\langle s \rangle$. For instance, we may select its centroid. The following relates any point s with its representative one $\bar{s}$, within their equivalence class $\langle s \rangle$:

$$\forall s \in \mathcal{S},\ \exists \langle s \rangle \in \{\mathcal{S}\}_\delta : (s, \bar{s} \in \langle s \rangle) \wedge (q = \bar{q}) \wedge (\|x - \bar{x}\| \leq \delta).$$

Given any $q \in \mathcal{Q}$ and any subset $W \subseteq \mathbb{R}^{n(q)}$, we denote the measure of the volume of W as $\lambda_W = \mathcal{L}(W)$, where $\mathcal{L}$ is the Lebesgue measure. λ_W is finite if $W \subseteq \mathcal{D}_q$. The volume of the hybrid state space is defined to be $\lambda_{\mathcal{S}} = \sum_{q \in \mathcal{Q}} \lambda_{\mathcal{D}_q}$. Since in this work a partition of $\mathcal{D}_q$ is defined as a grid of width δ, then $\forall s = (q, x) \in \mathcal{S}, \lambda_{\langle s \rangle} = \delta^{n(q)}$. It follows that the cardinality of the complete partition $\{\mathcal{S}\}_\delta$ is given by:

$$|\{\mathcal{S}\}_\delta| = \sum_{q \in \mathcal{Q}} \frac{\lambda_{D_q}}{\delta^{n(q)}}.$$

If we assume, for the sake of simplicity and without much loss of generality, that $\forall q \in \mathcal{Q}, n(q) = n$, then $|\{\mathcal{S}\}_\delta| = \frac{\lambda_{\mathcal{S}}}{\delta^n}$.

B. Control Space Partition

Let $\mathcal{A} \subset \mathbb{R}^p$ be a compact set: we introduce a finite partition $\{\mathcal{A}\}_\eta$ of the control space $\mathcal{A}$ by defining a grid of width η of $\mathbb{R}^p \cap \mathcal{A}$ and its associated non-overlapping sets, as it was illustrated above for $\mathcal{S}$.

Given any point $a \in \mathcal{A}$ there exists an element of $\{\mathcal{A}\}_\eta$, which we denote as $\langle a \rangle$, such that $a \in \langle a \rangle$. Any element $\langle a \rangle \in \{\mathcal{A}\}_\eta$ is a subset of the control space, i.e. $\langle a \rangle \subseteq \mathcal{A}$. Let $\bar{a} \in \mathcal{A}$ be a *representative* point of $\langle a \rangle$, e.g. its centroid. The following holds on $a, \bar{a}, \langle a \rangle$:

$$\forall a \in \mathcal{A},\ \exists \langle a \rangle \in \{\mathcal{A}\}_\eta : (a, \bar{a} \in \langle a \rangle) \wedge (\|a - \bar{a}\| \leq \eta).$$

As before, we define the volume of the control space as $\lambda_{\mathcal{A}}$, and the volume of each element of the control space partition $\lambda_{\langle a \rangle} = \eta^p$. The cardinality of $\{\mathcal{A}\}_\eta$ is given by $|\{\mathcal{A}\}_\eta| = \frac{\lambda_{\mathcal{A}}}{\eta^p}$.

V. Error Analysis of the Abstraction

In this section, we quantify the precision of the abstraction by providing a bound for the approximation error. This error bound is associated to an interval which, along with the transition probabilities computed on the partition, defines the MSC. We also investigate the actual dynamics in time of these errors, and show that, if the original system is ergodic, it is possible to obtain at any point in time an abstraction with arbitrary precision by tuning the parameters associated with the partition.

Let us recall that for any hybrid state $s = (q, x) \in \mathcal{S}$, we denote with $\langle s \rangle$ the corresponding element in the state space partition $\{\mathcal{S}\}_\delta$, and with $\bar{s} = (q, \bar{x})$ the representative point of $\langle s \rangle$. For any control feedback $\nu : \mathcal{S} \to \mathcal{A}$ and a partition $\{\mathcal{A}\}_\eta$ of $\mathcal{A}$, we define the corresponding quantized control feedback $\bar{\nu} : \mathcal{S} \to \{\mathcal{A}\}_\eta$ as follows: $\forall s \in \mathcal{S}, \nu(s) = a \in \mathcal{A} \Rightarrow \bar{\nu}(s) = \bar{a} \in \{\mathcal{A}\}_\eta$. In particular, this quantization is also applied on points $s = \bar{s} \in \{\mathcal{S}\}_\delta$.

Given a hybrid state $s = (q, x) \in \mathcal{S}$, a control feedback $\nu : \mathcal{S} \to \mathcal{A}$, and any target set $\langle s' \rangle \in \{\mathcal{S}\}_\delta$, the main idea is to approximate, for any $k \geq 0$, the one-step transition probability

$$p(\mathbf{s}(k+1) \in \langle s' \rangle \mid \mathbf{s}(k) = s, \nu(s)), \tag{3}$$

with the transition probability

$$p(\mathbf{s}(k+1) \in \langle s' \rangle \mid \mathbf{s}(k) = \bar{s}, \bar{\nu}(\bar{s})). \tag{4}$$

The computation of the above quantities involves the use either of the continuous transition kernel T_t, or of the continuous reset kernel T_r, depending on the mode selected by the discrete distribution T_q. More generally, by selecting a feedback policy ν over a k-step time horizon, we shall use the notation $p^k_{s,\nu}(\langle s' \rangle)$. We will omit the apex in p^k when $k = 1$. We will study the probability of events for the dt-cSHS by considering the control-dependent system defined on the quotient spaces and induced by the partition.

A. Single step error

Let us start by selecting a hybrid state $s = (q, x) \in \mathcal{S}$ and a control feedback ν. For any set $\langle s' \rangle \in \mathcal{S}$ (and, in particular, $\langle s' \rangle \in \{\mathcal{S}\}_\delta$), where $s' = (q', x')$, and assuming that $q' = q \in \mathcal{Q}$, we can derive the following bound:

$$\begin{aligned}
&|p_{s,\nu(s)}(\langle s' \rangle) - p_{\bar{s},\bar{\nu}(\bar{s})}(\langle s' \rangle)| \leq \\
&\left| \int_{\langle s' \rangle} T(dz|(q,x),\nu(q,x),q') - \int_{\langle s' \rangle} T(dz|(q,\bar{x}),\bar{\nu}(q,\bar{x}),q') \right| \\
&\qquad \leq \lambda_{\langle s' \rangle} \big((M_t + M_q)\|\nu(q,x) - \bar{\nu}(q,\bar{x})\| \\
&\qquad + ((L_t + L_q) + L_a(M_t + M_q))\|x - \bar{x}\| \big).
\end{aligned}$$

A similar bound can be derived if instead it is the case that $s' = (q', x')$, $\langle s' \rangle \subseteq \mathcal{D}_{q'}, q' \neq q$, by using the reset kernel T_r. Let us introduce the new constants $L = \max\{L_t + L_q, L_r + L_q\}$, $M = \max\{M_t + M_q, M_r + M_q\}$. Thus, in general, we can state that, $\forall s, s' \in \mathcal{S}$, for any control feedback ν:

$$\begin{aligned}
&|p_{s,\nu(s)}(\langle s' \rangle) - p_{\bar{s},\bar{\nu}(\bar{s})}(\langle s' \rangle)| \\
&\qquad \leq \delta^n (M\eta + (L + L_a M)\delta) \doteq \delta^n \varepsilon.
\end{aligned} \tag{5}$$

By virtue of the state and control spaces partition procedure (which is parameterized by the pair (δ, η)), and of the computed error bound, it is possible to associate to the dt-cSHS $\mathcal{H}$ a controlled MSC, call it $[\mathcal{M}]$, with transition intensities $p_{\bar{s},\bar{\nu}(\bar{s})}(\langle s' \rangle)$ and intervals given in (5).

B. Error dynamics

Recall the notations introduced above, $\forall s_0 \in \mathcal{S}, \langle s \rangle \in \{\mathcal{S}\}_\delta$ and for any continuous control feedback $\nu \in \mathcal{C}$, $p^k_{s_0,\nu}(\langle s \rangle)$ denotes the k-step transition probability into $\langle s \rangle$, starting from s_0 and following policy ν. Instead, $p^k_{\bar{s}_0,\bar{\nu}}$ is the corresponding probability distribution over the spaces $\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta$, and is generated by the definition of the Markov set-Chain $[\mathcal{M}]$. As we previously discussed, the distribution $p^k_{s_0,\nu}$ over the sets of the partition $\{\mathcal{S}\}_\delta$ is derived, for any k, from that of the SHS $\mathcal{H}$. This distribution is not known a-priori and depends on the particular initial state s_0. It can be thought of as a (unknown) non-homogeneous Markov chain evolving on the partition space $\{\mathcal{S}\}_\delta$: it will be denoted by $\Pi(k)$. Let us define a monotonically increasing function f: $f(\varepsilon, n, k) = (\varepsilon n + 1)^k f(\varepsilon, n, 0) + \varepsilon \sum_{l=1}^{k} (\varepsilon n + 1)^{l-1}$, $f(\varepsilon, n, 0) = d_h(p^0_{s_0}, p^0_{\bar{s}_0})$, where $p^0_{s_0}, p^0_{\bar{s}_0}$ are initial probability distributions respectively of the original system (considered over the regions of the partition) and of the Markov set-Chain abstraction.

Theorem 3: Given a dt-cSHS $\mathcal{H}$ and upholding Assumptions 1 and 2, suppose that there exists a partition defined by the parameters (δ, η) such that the corresponding Markov set-Chain abstraction $[\mathcal{M}]$ is ergodic with coefficient of ergodicity $\mathcal{T}([\mathcal{M}]) < 1$. Then, for any $s_0 \in \mathcal{S}, \langle s \rangle \in \{\mathcal{S}\}_\delta$ and any control feedback $\nu, \forall k \geq 0$,

$$\begin{aligned}
&d_h(p^k_{s_0,\nu}(\langle s \rangle), p^k_{\bar{s}_0,\bar{\nu}}(\langle s \rangle)) \leq \\
&\min\left\{ f(\varepsilon, n, k), 2\alpha\beta^k + \tfrac{2\lambda_{\mathcal{S}}\varepsilon}{1 - \mathcal{T}([\mathcal{M}])} \right\}
\end{aligned} \tag{6}$$

where α, β and d_h are defined as in Theorem 1. □

Remark 1: The study can be extended to events defined over time on Borel sets $C \subseteq \mathcal{B}(\mathcal{S})$, which do not necessarily coincide with the partition sets $\langle s \rangle \in \{\mathcal{S}\}_\delta$. The computation will resort to an under- or an over-approximation of C by sets of the partition. We leave the details to the interested reader. As already discussed, our attention will instead be focused on the probability distribution of the original SHS $\mathcal{H}$ over the sets of the partition, and on its asymptotics. □

Equation (6) provides a time-dependent bound for the approximation error, which is finite for each time step $k \geq 0$ if the abstraction induced by the partition parameters (δ, η) is ergodic. However, it is not always the case that there exist (δ, η) such that the abstraction is ergodic. In the following theorem we prove that, if the original system $\mathcal{H}$ is endowed with some ergodicity, there always exist (δ, η) such that the abstraction is ergodic. This fact implies that it is possible to achieve any desired precision on the approximation error of the abstraction over time by tuning the partition parameters (δ, η).

Theorem 4: Given $\mathcal{H}$ and any control feedback $\nu : \mathcal{S} \to \mathcal{A}$, if a stationary probability distribution $\mathfrak{p}_\nu^\infty$ of $\mathcal{H}$ exists,[2] then there exist partition parameters $(\delta > 0, \eta > 0)$, such that the induced abstraction $[\mathcal{M}]$ satisfies $\mathcal{T}([\mathcal{M}]) < 1$. □

Remark 2: The results in Theorems 3 and 4 can be extended to the case where the scrambling integer for the obtained MSC $[\mathcal{M}]$ is a finite integer $r > 1$, according to the corresponding bounds discussed in section III. □

VI. Abstraction and Analysis of a Stochastic Model for Bacterial Antibiotic Biosynthesis

The following model describes the production of the antibiotic subtilin by the bacterium *Bacillus subtilis*. The original model from [12] is slightly simplified in its structure by exploiting some symmetry in its organization, as observed in [2]. The model presents four variables: y = [SigH] (concentration of a sigma factor in the environment), z = [SpaS] (concentration of subtilin), X (nutrient level), D (population level). The first two entities are at the cellular level and have probabilistic dynamics, while the last two are deterministic averaged dynamics. Time is discrete, with time-step Δ. The positive variables are bounded above by the quantities y_M, z_M, X_M, D_M.

The level of the sigma factor SigH follows a probabilistic switching behavior according to:

$$y(k+1) = \begin{cases} y(k) - \lambda_1 y(k)\Delta + w_1(k) & \text{if } X \geq \eta D_M \\ y(k) + (k_3 - \lambda_1 y(k))\Delta + w_2(k) & \text{if } X < \eta D_M \end{cases}$$

which hinges on the nutrient-dependent spatial condition $\{X = \eta D_M, 0 \leq \eta \leq 1\}$. The terms w_1, w_2 are independent normal variables with zero mean and variance Δ.

Next, the concentration of the protein SpaS depends on one of two possible states of a switch S_1 as follows:

$$z(k+1) = \begin{cases} z(k) - \lambda_3 z(k)\Delta + v_1(k) & \text{if } S_1 \text{ is } OFF_z \\ z(k) + (k_5 - \lambda_3 z(k))\Delta + v_2(k) & \text{if } S_1 \text{ is } ON_z. \end{cases}$$

Again v_1, v_2 are independent normal variables with zero mean and variance Δ. The structure of $S_1 = \{OFF_z, ON_z\}$ is assumed to be that of a Markov Chain, whose transition probability matrix is $P_1 = \begin{bmatrix} 1-b_0 & b_0 \\ b_1 & 1-b_1 \end{bmatrix}$. The parameters b_0, b_1 depend on [SigH] according to $b_0(y) = \frac{\alpha y}{1+\alpha y}, b_1(y) = 1 - b_0(y)$. The quantity $\alpha = e^{-\Delta G_{rk}/RT}$ depends on the Gibbs free energy, a gas constant R, and the environment temperature T [12].

The variation in the population level is modeled by a logistic equation as follows:

$$D(k+1) = D(k) + rD(k)\left(1 - \frac{D(k)}{D_\infty}\right)\Delta, r > 0. \quad (7)$$

This is a quadratic equation, with two equilibria. The first ($D = 0$) is unstable, while the second ($D = D_\infty$) is stable. The non-trivial (and stable) equilibrium relation depends on the quantity D_∞ (the *carrying capacity*). Let us intuitively define it to be $D_\infty = \frac{X}{X_M} D_M$. The nutrient dynamics follow the difference equation

$$X(k+1) = X(k) + (k_2 \nu y - k_1 D(k) X(k))\Delta, \nu < 1. \quad (8)$$

The second equilibrium point for the population level corresponds a stable equilibrium for the nutrient level.

The above set of dynamical relations can be reframed as a SHS. The model has four modes, $\mathcal{Q} = \{q_1 = (ON_y, ON_z), q_2 = (ON_y, OFF_z), q_3 = (OFF_y, ON_z), q_4 = (OFF_y, OFF_z)\}$, where the pairs refer to the activities of the variables y, z: $ON_y = \{X : 0 \leq X < \eta D_M\}, OFF_y = \{\eta D_M \leq X \leq X_M\}; ON_z = \{S_1 \text{ is } ON_z\}, OFF_z = \{S_1 \text{ is } OFF_z\}$. The continuous part of the state space is also four dimensional, in each of the discrete domains, and it reflects the bounds on the four variables: $\mathcal{D} = \{[0, y_M] \times [0, z_M] \times [0, X_M] \times [0, D_M]\}$.

Let us introduce the stochastic kernels relative to the probabilistic dynamics at the cellular level:

$$\begin{aligned} T_t(dy|(OFF_y, y)) &= \mathcal{N}(dy, y - \lambda_1 y\Delta, \Delta); \\ T_t(dy|(ON_y, y)) &= \mathcal{N}(dy, y + (k_3 - \lambda_1 y)\Delta, \Delta); \\ T_t(dz|(OFF_z, z)) &= \mathcal{N}(dz, z - \lambda_3 z\Delta, \Delta); \\ T_t(dz|(ON_z, z)) &= \mathcal{N}(dz, z + (k_5 - \lambda_3 z)\Delta, \Delta). \end{aligned}$$

Here $\mathcal{N}(\cdot, x, \sigma)$ is the probability measure associated to a normal distribution of mean x and variance σ. The reset kernels T_r are trivial. Furthermore, the discrete kernels have the following form:

$$\begin{aligned} T_q(q_2|((ON_y, y), (ON_z, z)) &= P_1(2,1)\mathbf{1}_{ON_y}; \\ T_q(q_3|((ON_y, y), (OFF_z, z)) &= P_1(1,2)\mathbf{1}_{OFF_y}; \\ T_q(q_4|((OFF_y, y), (ON_z, z)) &= P_1(2,1)\mathbf{1}_{OFF_y}; \\ T_q(q_1|((OFF_y, y), (OFF_z, z)) &= P_1(1,2)\mathbf{1}_{ON_y}, \end{aligned}$$

where the other possible transition probabilities are obtained analogously, or by complementation.

The constants for the locally Lipschitz T_t kernels, with regards to their density t, are those in [1]:

$$|t(\cdot|q, y) - t(\cdot|q, y')| \leq \frac{1}{\sqrt{2}} \frac{e^{y_M} - 1}{y_M} |y - y'|, \forall q \in \mathcal{Q}.$$

Those for the T_q kernels can be directly found by inspecting P_1:

$$|P_1(i,j)(y') - P_1(i,j)(y'')| \leq \alpha |y' - y''|, \forall y', y'' \in [0, y_M].$$

Let us introduce a uniform partition of the state space according to a grid of width δ, which we assume to be a divisor of the quantity ηD_M. The order of the cardinality of the partition is easily $|\mathcal{Q}|(\max\{y_M, z_M, X_M, D_M\}/\delta)^4$.

[2]Notice the difference between the measure $\mathfrak{p}_\nu^\infty$, which refers to the original dt-cSHS $\mathcal{H}$, and p_ν^∞, which refers to $\Pi(k)$.

	Population Level	Nutrient Level	Time
Simulation with full Monte Carlo	0.88	3.6	127
Simulation with Reduced Abstraction	0.78	3.1	37
Simulation with Full Abstraction	0.5	2.7	80

Clearly $\lambda_S = y_M z_M X_M D_M$. The introduced error is $\varepsilon_1 = \Delta^5 \left(\frac{1}{\sqrt{2}} \frac{e^{y_M}-1}{y_M} + \alpha \right)$. Let us consider for simplicity a probability for the initial points that is uniformly distributed over the space. As in [12], we have chosen $y_M = 4, z_M = 4, X_M = 10, D_M = 1$. The reference simulations have been implemented with a Monte Carlo approach: ten simulations have been run from the starting states corresponding to the representative points of the abstraction. The outputs for population level and nutrient are in the Table.

We have also implemented a full abstraction of the above dynamics with discretization level $\delta = 0.5$ (see Table). The MC has 1260 rows. While the steady states of the population and of the nutrient appear to be close to the desired ones, those of the cellular dynamics are not satisfactory. This is possibly due to the sparsity of the obtained MC, which is ascribed to the presence of deterministic dynamics. We have then decided to reduce the abstraction at the level of the cellular dynamics. More precisely, we have come up with two different abstractions for the two regions ON_y, OFF_y. This has also expedited the computation time for the abstraction, so that we can push the discretization level to be quite small. The outputs for population and nutrient are close to those of the Monte Carlo Simulations (See Figures 1, 2). Notice the computational improvement that the abstraction procedure gains, despite the required partitioning procedure and calculation of the steady state.

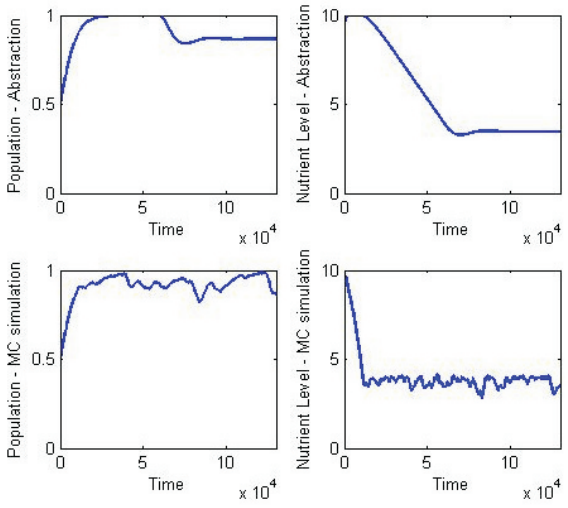

Fig. 1. Comparison between Monte Carlo Simulations and Abstraction Steady-State, for a first initial condition.

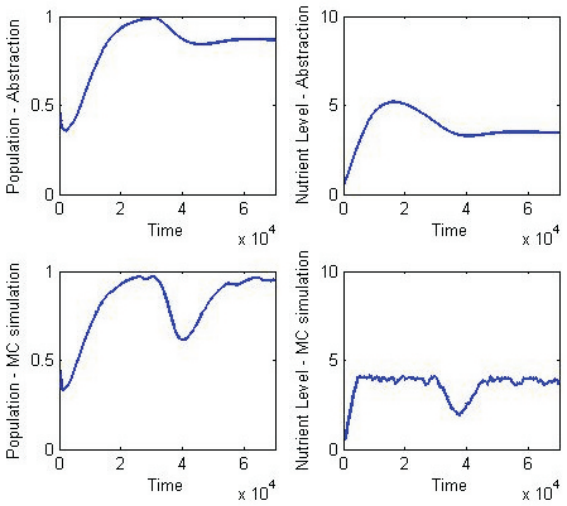

Fig. 2. Comparison between Monte Carlo Simulations and Abstraction Steady-State, for a second initial condition.

REFERENCES

[1] A. Abate, A. D'Innocenzo, M.D. Di Benedetto, and S. Sastry. Markov set-Chains as abstractions of stochastic hybrid systems. In M. Egerstedt and B. Misra, editors, *Hybrid Systems: Computation and Control*, LNCS 4981. Springer Verlag, 2008.

[2] A. Abate, J. Lygeros, and S. Sastry. Probabilistic safety and optimal control for survival analysis of *Bacillus subtilis*. In *Proceedings of the 2nd Conference on Foundations of Systems Biology in Engineering*, pages 527–532, Stuttgart, DE, 2007.

[3] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, Nov 2008. In press.

[4] R. Alur, T. Henzinger, G. Lafferriere, and G. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(2):971–984, July 2000.

[5] A. D'Innocenzo, A.A. Julius, M.D. Di Benedetto, and G.J. Pappas. Approximate timed abstractions of hybrid automata. In *Proceedings of the* 46^{th} *IEEE Conference on Decision and Control. New Orleans, Louisiana, USA*, 12–14 December 2007.

[6] Antoine Girard. Approximately bisimilar finite abstractions of stable linear systems. In *Hybrid Systems: Computation and Control*, volume 4416 of *Lecture Notes in Computer Science*, pages 231–244. Springer Verlag, 2007.

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

[8] Robert Givan, Sonia Leach, and Thomas Dean. Bounded parameter Markov decision processes. *Artificial Intelligence*, 122(1-2):71–109, 2000.

[9] Harald J. Hartfiel. *Markov Set-Chains*, volume 1695 of *Lecture Notes in Mathematics*. Springer-Verlag Berlin Heidelberg, 1998.

[10] T. A. Henzinger, R. Majumdar, and V. Prabhu. 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*, pages 226–241. Springer, 2005.

[11] Masanori Hosaka, Masayuki Horiguchi, and Masami Kurano. Controlled Markov set-Chains under average criteria. *Applied Mathematics and Computation*, 120:195–209, 2001.

[12] J. Hu, W.-C. Wu, and S. Sastry. Modeling subtilin production in *bacillus subtilis* using stochastic hybrid systems. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 417–431. Springer Verlag, 2004.

[13] A.A. Julius and G.J. Pappas. Approximate abstraction of stochastic hybrid systems. *IEEE Trans. Automatic Control*, provisionally accepted.

[14] Giordano Pola, Antoine Girard, and Paulo Tabuada. Symbolic models for nonlinear control systems using approximate bisimulations. In *Proceedings of the* 46^{th} *Conference on Decision and Control, New Orleans, LA, USA*, December 2007.

[15] M.L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 1994.