

Approximate Abstractions of Discrete-Time Controlled Stochastic Hybrid Systems

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

Abstract—This work proposes 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 approximation errors can be explicitly computed, over time, given proper continuity assumptions on the model. We show that the errors can be arbitrarily chosen by increasing the partition accuracy. Similar bounds can be provided if a particular feedback control policy is selected and quantized. The obtained abstraction can be interpreted as a Bounded-parameters Markov Decision Process, or a controlled Markov set-Chain, and can be used both for verification and control design purposes. We finally test the approximate abstraction technique on a model from systems biology.

I. INTRODUCTION

The dynamical analysis of complex, high-dimensional, possibly stochastic models often poses many challenges, both at a theoretical and at a computational level. Abstraction techniques are often the method of choice, when not the only viable option, for the analysis of such systems.

According to this approach, a system with a smaller (possibly finite) state space is obtained, which is equivalent to the original system under study. Systems equivalence is usually defined via the notions of language equivalence and bisimulation [4], [14]. Recently, *approximate* notions of equivalence [8] have been developed, where a metric is introduced to quantify the distance between the original system and its abstraction.

The present line of research looks at abstractions of Hybrid Systems (HS), models which often require to be abstracted for analysis. The contribution in [6] 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 [7] it has been proved that stable linear systems admit finite approximately bisimilar abstractions with arbitrary precision. In [15] the same has been done for a class of stable non-linear systems. In [13] a notion of approximate bisimilarity is proposed for a class of Stochastic Hybrid Systems (SHS), that is HS which are endowed with probabilistic terms.

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) [10], 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 a refinement parameter and synthesizes an abstraction, according to that parameter, which abides by the desired precision bound.

The present contribution extends that in [1] in three directions. First, the SHS model is more general, in that an execution is allowed to change mode not just according to a spatial condition, but to a state-dependent probability distribution (which, as a special case, could reproduce a spatial guard). Second, the model is control-dependent, which requires a proper partitioning of the control space, and an integration of the errors on the state and on the control. Both the open-loop and the feedback control structures will be considered. 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), for each time instant (and in particular in steady-state). This bounds represent another step towards a definition of stochastic bisimulation, which is the ultimate goal of this line of research.

Our abstraction can be used both for verification purposes (e.g. given a continuous control policy on a continuous plant, verify properties of the quantized implementation of the control on a discretized state space) and for design purposes (e.g. given a continuous system, design a quantized control policy on the abstraction, using classical algorithms for Markov Decision Processes (MDP) [16], which may guarantee the correct behavior of the original plant).

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 bounded-parameters MDP (BMDP). Once a policy is fixed, the BMDP reduces to an 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 asymptotic properties.

II. CONTROLLED DISCRETE TIME STOCHASTIC HYBRID SYSTEMS

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

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

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

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 $\mathcal{Q} := \{q_1, q_2, \dots, q_m\}$, 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 $n : \mathcal{Q} \rightarrow \mathbb{N}$ 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 ;
- $T_q : \mathcal{Q} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a discrete stochastic kernel (the “discrete transition kernel”) on \mathcal{Q} given $\mathcal{S} \times \mathcal{A}$, which assigns to each $s \in \mathcal{S}$ and $a \in \mathcal{A}$, a discrete probability distribution over \mathcal{Q} : $T_q(q|s, a)$;
- $T_t : \mathcal{B}(\mathcal{D}_{(\cdot)}) \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a Borel-measurable stochastic kernel (the “continuous transition kernel”) on $\mathcal{D}_{(\cdot)}$, given $\mathcal{S} \times \mathcal{A}$, which assigns to each $s = (q, x) \in \mathcal{S}$ and $a \in \mathcal{A}$ a probability measure on the Borel space $(\mathcal{D}_q, \mathcal{B}(\mathcal{D}_q))$: $T_t(dx|(q, x), a)$;
- $T_r : \mathcal{B}(\mathcal{D}_{(\cdot)}) \times \mathcal{S} \times \mathcal{A} \times \mathcal{Q} \rightarrow [0, 1]$ 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 $s = (q, x) \in \mathcal{S}$, $a \in \mathcal{A}$, and $q' \in \mathcal{Q}$, $q' \neq q$, a probability measure on the Borel space $(\mathcal{D}_{(q')}, \mathcal{B}(\mathcal{D}_{(q')}))$: $T_r(dx|(q, x), a, q')$. \square

The system initialization at the initial time (say $k = 0$) may be specified by some probability measure $\pi_0 : \mathcal{B}(\mathcal{S}) \rightarrow [0, 1]$ 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 a control input acting on the discrete or on the continuous dynamics. This model is an extension of the similar framework in [1] by the introduction of the control set. Furthermore, unlike in [1], where the change in mode depends on the verification of spatial conditions, here the execution is allowed to change its operating mode according to the probability law of the discrete kernel T_q , which is defined on the whole state space. 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], we denoted with $g_{q, q'}$ a subset of the domain \mathcal{D}_q , which is associated to a jumping condition to mode $q' \in \mathcal{Q}$, $q' \neq q$. 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 concepts of control string and policy.

Definition 2 (Control String): Given a dt-cSHS \mathcal{H} , a control string is a sequence $a = (a_0, a_1, \dots, a_N)$, of cardinality $N + 1$, $N \in \mathbb{N}$, of values $a_k \in \mathcal{A}$. \square

Definition 3 (Execution): Given a dt-cSHS \mathcal{H} , an initial distribution π_0 and a control string $a = (a_0, a_1, \dots, a_N)$, an execution of \mathcal{H} is a stochastic

process¹ $\{\mathbf{s}(k) = (\mathbf{q}(k), \mathbf{x}(k)) : \forall k = 0, \dots, N + 1, s_k \in \mathcal{S}\}$ 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 $T_q(\cdot | s_k, a_k)$;

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

then extract a value $s_{k+1} \in \mathcal{S}$ for $\mathbf{s}(k + 1)$, according to $T_r(\cdot | s_k, a_k, q_{k+1})$;

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

end. \square

It is clear that, when $N = \infty$, the algorithm does not terminate. We make use of the following shortened notation, where $q, q' \in \mathcal{Q}$:

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

Definition 4 (Markov Control Policy): Given a dt-cSHS \mathcal{H} , a (state) feedback control policy ν on a time horizon $[0, \dots, N]$ for \mathcal{H} is a sequence of functions $\nu_i : \mathcal{S} \rightarrow \mathcal{A}$, $i \in [0, \dots, N]$, which associates to each hybrid state $s \in \mathcal{S}$ a control action $\nu_i(s) \in \mathcal{A}$. \square

We will often employ stationary (or homogeneous) policies, that is policies that are time invariant and thus simply characterized by a function $\nu : \mathcal{S} \rightarrow \mathcal{A}$. The controls are all assumed to be deterministic.

Definition 5 (Feedback Execution): Given a dt-cSHS \mathcal{H} , an initial distribution π_0 and a homogeneous Markov control policy ν , a feedback execution of \mathcal{H} is a stochastic process $\{\mathbf{s}(k) : \forall k = 0, \dots, N + 1, s_k \in \mathcal{S}\}$ generated by the algorithm in Definition 3, where $a_k = \nu(s_k)$. \square

III. MARKOV SET-CHAINS

We recall here the concept of Markov set-chains, which in this paper is later leveraged to prove properties of the abstraction. The results are from [10].

Definition 6: 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:”

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

\square

In the proceeding, we assume that the set $[P, Q] \neq \emptyset$. When the “bounding matrices” will be clear by the context, we will use the more compact notation [II]. We can define a Markov set-chain as a non-homogeneous discrete-time Markov chain,

¹Bold symbols will denote processes, while regular typesets points on the state space.

where the transition probabilities vary non-deterministically within a compact transition set $[\Pi]$ at each time step. More formally,

Definition 7: 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

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

a Markov set-chain, where $[\Pi]^k$ is defined by induction as the set of all possible products A_1, \dots, A_k , such that $\forall i = 1, \dots, k, A_i \in [\Pi]$.

Let $[\pi_0]$ be a compact set of $1 \times n$ stochastic vectors, introduced as in Definition 6. We call $[\pi_0]$ the initial distribution set. \square

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

$$[\pi_0], [\pi_0][\Pi], \dots$$

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

Definition 8: 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 . \square

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

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

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

Since $\mathcal{T}(\cdot)$ is a continuous function and $[\Pi]$ a compact set, the corresponding maximum argument exists. Notice that $\mathcal{T}([\Pi]) \in [0, 1]$ provides a measure of the ‘‘contractive’’ nature of the Markov set-chain: the smaller $\mathcal{T}([\Pi])$, the more contractive the MSC. The following notion is important for characterizing the convergence of a MSC:

Definition 10: Suppose r is an integer such that $\mathcal{T}(A_1, \dots, A_r) < 1, \forall A_1, \dots, A_r \in [\Pi]$. Then $[\Pi]$ is said to be product scrambling and r its scrambling integer. \square

We now illustrate some results on the convergence of MSC.

Theorem 1: 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 k ,

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

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. Thus

$$\lim_{k \rightarrow \infty} [\pi_k] = \lim_{k \rightarrow \infty} [\pi_0][\Pi]^k = [\pi_\infty], \quad \square$$

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 $[\pi_\infty]$.

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

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

IV. ABSTRACTION PROCEDURE

In digital and embedded control systems, often the control action comes in quantized form. The discrete nature of this action has required the understanding of how it relates to continuous quantities (which is a topic of classical industrial control and automation), and how concepts such as stability, reachability, and robustness are affected by this [5]. Similar comments apply to the state space of a system.

With the above frame of mind, the abstraction we propose for the state and the control spaces can be used as a double-edged weapon (see section V-A). From a *verification* perspective, given a continuous control policy on a continuous plant, we can use finite abstractions to verify properties of the quantized implementation of the control. On the other hand, from a *design* point of view, given a continuous system, we can resort to its finite abstraction as a framework over which to synthesize an implementable (quantized) control policy (e.g. using classical MDP algorithms [16]), while guaranteeing that the closed-loop behavior of the original system is similar to that of the abstracted system.

In this section we introduce an abstraction of a dt-cSHS \mathcal{H} into a Bounded-parameter Markov Decision Process (BMDP) [9]. This model is closely related to the so-named controlled Markov set-Chain (cMSC) [11]. Given a BMDP (or a cMSC) and a feedback control policy, the resulting controlled system is effectively a Markov set-Chain [10]. The MSC framework has been leveraged in [1], where a less general and uncontrolled SHS has been abstracted to within that structure.

The abstraction of the probabilistic dynamics of \mathcal{H} is obtained by approximating some of its elements [1, sections 4 and 5]. Let us suppose that the stochastic kernels T_t, T_r admit densities t, r . Let us raise the following conditions:

Assumption 1 (Continuity of the Stochastic Kernels):

- 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

A. State Space Partition

Let us recall that \mathcal{D}_q is 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 $q \in \mathcal{Q}$, one can define a partition $\{\mathcal{D}_q\}_\delta$ made up of subsets of \mathcal{D}_q . We allow the partition to be general in its shape. However it is usually introduced as a grid of width δ in $\mathbb{R}^{n(a)} \cap \mathcal{D}_q$, where δ is the diameter

of the partition, that is the maximum distance between any two points in the same equivalence class. 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 of $\{\mathcal{S}\}_\delta$, that is a set denoted by $\langle s \rangle$, such that $s \in \langle s \rangle$. 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 a single domain \mathcal{D}_q . Let us select any point $\bar{s} = (\bar{q}, \bar{x}) \in \mathcal{S}$ to be the *representative* point of the set $\langle s \rangle$. For instance, we may select its centroid. The following holds:

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

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. We also define the volume of the hybrid state space as $\lambda_{\mathcal{S}} = \sum_{q \in \mathcal{Q}} \lambda_{\mathcal{D}_q}$. Since we defined a partition of \mathcal{D}_q as a grid of width δ , then $\forall s = (q, x) \in \mathcal{S}, \lambda_{\langle s \rangle} = \delta^{n(q)}$. For these reasons, the cardinality of the complete partition $\{\mathcal{S}\}_\delta$ is given by:

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

If we assume, for the sake of simplicity and without 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 $\mathcal{A} \subseteq \mathbb{R}^p$ be a compact set: we define a finite partition $\{\mathcal{A}\}_\eta$ of the control space \mathcal{A} by defining a grid of width η of $\mathbb{R}^p \cap \mathcal{A}$, as it was illustrated above for \mathcal{S} .

Given any $a \in \mathcal{A}$ 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:

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

As we did above, 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

We recall that for any hybrid state $s = (q, x) \in \mathcal{S}$, we denote $\langle s \rangle$ as the corresponding element in the state space partition $\{\mathcal{S}\}_\delta$ and \bar{s} as the representative element of $\langle s \rangle$. Moreover, for any control $a \in \mathcal{A}$ we denote $\langle a \rangle$ as the corresponding element in the control space partition $\{\mathcal{A}\}_\eta$ and \bar{a} as the representative element of $\langle a \rangle$. We will denote with $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$ the combined partition of state and control spaces, which depends on the pair (δ, η) .

Given a hybrid state $s = (q, x) \in \mathcal{S}$, a control value $a \in \mathcal{A}$, and any set $\langle s' \rangle \in \{\mathcal{S}\}_\delta$, we will approximate the one step transition probability

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

with the transition probability

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

for any $k \geq 0$. The computation of the above quantities involves the use of either the transition kernel T_t , or of the reset kernel T_r , depending on the mode selected by T_q . We use the following notation, where $C \subseteq \mathcal{B}(\mathcal{S})$:

$$p_{s,a}(C) = p(\mathbf{s}(k+1) \in C \mid \mathbf{s}(k) = s, a(k) = a),$$

A. One-step error

Select a hybrid state $s = (q, x) \in \mathcal{S}$ and a control value $a \in \mathcal{A}$. For any set $\langle s' \rangle, s' = (q', x'), q' \neq q$ we can derive the following bound:

$$\begin{aligned} & |p_{s,a}(\langle s' \rangle) - p_{\bar{s},\bar{a}}(\langle s' \rangle)| \\ &= \left| \int_{\langle s' \rangle} T(dz|(q, x), a, q') - \int_{\langle s' \rangle} T(dz|(q, \bar{x}), \bar{a}, q') \right| \\ &\leq \int_{\langle s' \rangle} \left| T_q(q'|(q, x), a) T_r(dz|(q, x), a) \right. \\ &\quad \left. - T_q(q'|(q, x), a) T_r(dz|(q, \bar{x}), \bar{a}) \right| \\ &\leq \int_{\langle s' \rangle} \left\{ \left| T_q(q'|(q, x), a) T_r(dz|(q, x), a) \right. \right. \\ &\quad \left. \left. - T_q(q'|(q, \bar{x}), a) T_r(dz|(q, x), a) \right| \right. \\ &\quad \left. + \left| T_q(q'|(q, \bar{x}), a) T_r(dz|(q, x), a) \right. \right. \\ &\quad \left. \left. - T_q(q'|(q, \bar{x}), \bar{a}) T_r(dz|(q, x), a) \right| \right. \\ &\quad \left. + \left| T_q(q'|(q, \bar{x}), \bar{a}) T_r(dz|(q, x), a) \right. \right. \\ &\quad \left. \left. - T_q(q'|(q, \bar{x}), \bar{a}) T_r(dz|(q, \bar{x}), \bar{a}) \right| \right. \\ &\quad \left. + \left| T_q(q'|(q, \bar{x}), \bar{a}) T_r(dz|(q, \bar{x}), a) \right. \right. \\ &\quad \left. \left. - T_q(q'|(q, x), a) T_r(dz|(q, \bar{x}), \bar{a}) \right| \right\} \\ &\leq \lambda_{\langle s' \rangle} ((M_r + M_q) \|a - \bar{a}\| + (L_r + L_q) \|x - \bar{x}\|). \end{aligned}$$

If it is instead the case that $s' = (q, x'), \langle s' \rangle \subseteq \mathcal{D}_q$, by using the kernel T_t the following holds:

$$\begin{aligned} & |p_{s,a}(\langle s' \rangle) - p_{\bar{s},\bar{a}}(\langle s' \rangle)| \\ &\leq \lambda_{\langle s' \rangle} ((M_t + M_q) \|a - \bar{a}\| + (L_t + L_q) \|x - \bar{x}\|). \end{aligned}$$

In the following, we will use the new constants:

$$L = \max\{L_t + L_q, L_r + L_q\}, M = \max\{M_t + M_q, M_r + M_q\}.$$

In general, we can state that, $\forall s, s' \in \mathcal{S}, \forall a \in \mathcal{A}$:

$$\begin{aligned} |p_{s,a}(\langle s' \rangle) - p_{\bar{s},\bar{a}}(\langle s' \rangle)| &\leq \lambda_{\langle s' \rangle} (M \|a - \bar{a}\| + L \|x - \bar{x}\|) \\ &\leq \delta^n (M\eta + L\delta) \triangleq \varepsilon_1. \end{aligned} \quad (3)$$

Consider now a static Markov control policy $\nu : \mathcal{S} \rightarrow \mathcal{A}$ defined on the system \mathcal{H} . Let us raise the following additional continuity assumption:

Assumption 2 (Continuity of the Control): For any static Markov policy $\nu : \mathcal{S} \rightarrow \mathcal{A}$,

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

where L_a is a finite and positive constant. \square

Based on Assumption 2, calculations similar to those that yielded inequality (3) lead now to the bound:

$$\begin{aligned} & |p_{s,\nu(s)}(\langle s' \rangle) - p_{\bar{s},\nu(\bar{s})}(\langle s' \rangle)| \\ &= \left| \int_{\langle s' \rangle} T(dz|s, \nu(s), q') - \int_{\langle s' \rangle} T(dz|\bar{s}, \nu(\bar{s}), q') \right| \\ &\leq \lambda_{\langle s' \rangle} (ML_a + L) \|x - \bar{x}\| \leq \delta^n (ML_a + L) \delta. \end{aligned}$$

In this instance, we have referred the abstraction exclusively to the state space. We will use the notation $(\{\mathcal{S}, \mathcal{A}\})_\delta$ to stress this. Instead, we may be given a Markov control policy $\nu : \mathcal{S} \rightarrow \mathcal{A}$ and want to construct a discrete abstraction of the closed loop system \mathcal{H} by defining a quantized policy $\bar{\nu} : \{\mathcal{S}\}_\delta \rightarrow \{\mathcal{A}\}_\eta$ on the abstraction. Given $\nu(s) = a$, for any $s \in \mathcal{S}$, let us introduce \bar{a} so that $\bar{a}, a \in \langle a \rangle$. This introduces a further abstraction approximation error, because of the approximation of the Markov control policy on the original system.

$$\begin{aligned} & |p_{s,\nu(s)}(\langle s' \rangle) - p_{\bar{s},\bar{a}}(\langle s' \rangle)| \\ &= \left| \int_{\langle s' \rangle} T(dz|s, \nu(s), q') - \int_{\langle s' \rangle} T(dz|\bar{s}, \bar{a}, q') \right| \\ &\leq \int_{\langle s' \rangle} \left\{ |T(dz|s, \nu(s), q') - T(dz|\bar{s}, \nu(s), q')| \right. \\ &\quad \left. + |T(dz|\bar{s}, \nu(s), q') - T(dz|\bar{s}, \bar{a}, q')| \right\} \\ &\leq \lambda_{\langle s' \rangle} (L\|x - \bar{x}\| + M(\|\nu(s) - \nu(\bar{s})\| + \|\nu(\bar{s}) - \bar{a}\|)) \\ &\leq \delta^n (L\delta + ML_a\delta + M\eta) \triangleq \varepsilon_2. \end{aligned} \quad (4)$$

B. Multi-step error

In the following we aim at generalizing the above calculations by computing the approximate probability of an event, over a finite time horizon, with the associated error. For the sake of clarity, we derive explicit formulas for the two-steps case, then generalize them.

Pick any hybrid state $s = (q, x) \in \mathcal{S}$ and a two-step static control policy $a \in \mathcal{A}$. Consider the probability $p_{s,a}^2(\langle s'' \rangle)$ to transition to any set $\langle s'' \rangle \subseteq \mathcal{S}$, $s'' = (q'', x'')$, $q'' \neq q$ in two time steps. Let us also consider the same quantity, computed over the abstraction $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$, and denoted with $p_{\bar{s},\bar{a}}^2(\langle s'' \rangle)$. The error can be quantified as follows:

$$\begin{aligned} & |p_{s,a}^2(\langle s'' \rangle) - p_{\bar{s},\bar{a}}^2(\langle s'' \rangle)| \\ &= \left| \int_{\langle s'' \rangle} \int_{\mathcal{S}} T(dz'|q', z, a, q'') T(dz|(q, x), a, q') \right. \\ &\quad \left. - \int_{\langle s'' \rangle} \sum_{j \in \{\mathcal{S}\}_\delta} \int_{\langle s''_j \rangle} T(dz'|q', \bar{x}_j, \bar{a}, q'') T(dz|(q, \bar{x}), \bar{a}, q') \right| \\ &\leq \sum_{j \in \{\mathcal{S}\}_\delta} \int_{\langle s'' \rangle} \int_{\langle s''_j \rangle} \left| T(dz'|q', z, a, q'') T(dz|(q, x), a, q') \right. \\ &\quad \left. - T(dz'|q', z, a, q'') T(dz|(q, \bar{x}), \bar{a}, q') \right| \\ &\quad + \left| T(dz'|q', z, a, q'') T(dz|(q, \bar{x}), \bar{a}, q') \right. \\ &\quad \left. - T(dz'|q', \bar{x}_j, \bar{a}, q'') T(dz|(q, \bar{x}), \bar{a}, q') \right| \\ &\leq 2\mathcal{L}(\langle s'' \rangle) \lambda_{\mathcal{S}} (L\delta + M\eta) = 2\lambda_{\mathcal{S}} \varepsilon_1. \end{aligned}$$

By proceeding similarly, we derive that the error associated with the k^{th} -step is

$$|p_{s,a}^k(\langle s' \rangle) - p_{\bar{s},\bar{a}}^k(\langle s' \rangle)| \leq (2\lambda_{\mathcal{S}})^{k-1} \varepsilon_1.$$

Remark 1: The calculations for the single- and multi-step errors can be extended to events $C \subseteq \mathcal{B}(\mathcal{S})$ that do not necessarily coincide with the partition sets $\langle s \rangle \in \{\mathcal{S}\}_\delta$. This 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 our attention will be focused on the probability distribution of the original SHS \mathcal{H} over the sets of the partition, as the following paragraph further develops. \square

C. Error dynamics

We will analyze in this section the approximation error dynamics of the abstraction defined in section IV. We will distinguish two cases.

In the first case (“synthesis”), we assume that a Markov control policy ν is **not** selected on the system \mathcal{H} . The abstraction $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$ is made up of a MDP–and associated errors– with transition probability matrix $P = \{p_{ij}(\bar{a})\}$, where $p_{ij}(\bar{a}) = p_{\bar{s}_i, \bar{a}}(\langle s_j \rangle)$ for each $\langle s_i \rangle, \langle s_j \rangle \in \{\mathcal{S}\}_\delta$: $\bar{s}_i \in \langle s_i \rangle$, and for each $\langle a \rangle \in \{\mathcal{A}\}_\eta$: $\bar{a} \in \langle a \rangle$. Given any feedback control policy introduced on $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$, call it $\bar{\nu} : \{\mathcal{S}\}_\delta \rightarrow \{\mathcal{A}\}_\eta$, the controlled MDP is actually a MC \mathcal{M} , with transition probability matrix $P = \{p_{ij}\}$, $p_{ij} = p_{\bar{s}_i, \bar{\nu}(\langle s_i \rangle)}(\langle s_j \rangle)$. We define our abstraction as a Markov set-Chain $[\mathcal{M}]$, with transition probability interval $[P] = \{[p_{ij} - \varepsilon_1, p_{ij} + \varepsilon_1]\}$, according to the error in (3).

In the second case (“verification”), we assume that a static Markov control policy $\mu : \mathcal{S} \rightarrow \mathcal{A}$ is already selected on the SHS \mathcal{H} . Our abstraction consists of a MDP and a quantized Markov control policy $\bar{\mu} : \{\mathcal{S}\}_\delta \rightarrow \{\mathcal{A}\}_\eta$. Using the same reasoning as above, we consider an abstraction $[\mathcal{M}]$, which is a Markov set-Chain with transition probability interval $[P] = \{[p_{ij} - \varepsilon_2, p_{ij} + \varepsilon_2]\}$, according to (4).

We will prove the following theorem for the error bounds of the first case, but the same result directly applies to the second case by using the bound ε_2 , instead of ε_1 . Recall the following notations introduced above, $\forall \langle s \rangle \in \{\mathcal{S}\}_\delta, \forall h \geq 0$:

$$\begin{aligned} p_{s_0, \nu}^k(\langle s \rangle) &= p(\mathbf{s}(k+h) \in \langle s \rangle \mid s(h) \in \langle s_0 \rangle, \nu), \\ p_{s_0, \bar{\nu}}^k(\langle s \rangle) &= p(\langle s \rangle(k+h) = \langle s \rangle \mid \langle s \rangle(h) = \langle s_0 \rangle, \bar{\nu}) \end{aligned}$$

where $p_{s_0, \bar{\nu}}^k$ is the probability distribution over the state space $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$ and is generated by the Markov set-chain $[\mathcal{M}]$. The distribution $p_{s_0, \nu}^k$ over the sets of the partition is derived from that of the SHS \mathcal{H} , and can be thought to be generated by a non-homogeneous Markov chain $\Pi(k)$, with the same state space of $[\mathcal{M}]$.

Theorem 3: Assume that there exist a partition $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$ such that the corresponding Markov set-chain abstraction $[\mathcal{M}]$ is ergodic, with coefficient of ergodicity $\mathcal{T}([\mathcal{M}]) < 1$ and scrambling integer r . Then, for any $s_0 \in \mathcal{S}$, $\langle s \rangle \in \{\mathcal{S}\}_\delta$ and any policy $\bar{\nu}$ and $\nu(s) = \bar{\nu}(\langle s \rangle)$:

$$\begin{aligned} & d_h(p_{s_0, \nu}^k(\langle s \rangle), p_{s_0, \bar{\nu}}^k(\langle s \rangle)) \leq \\ & \leq \min \left\{ \left(\frac{\lambda_{\mathcal{S}}}{\delta^n} \right)^k \sum_{i=1}^k \varepsilon_1^i, 2\alpha\beta^k + \frac{2\lambda_{\mathcal{S}}(M\eta + L\delta)}{1 - \mathcal{T}([\mathcal{M}])} \right\}, \end{aligned} \quad (5)$$

where α, β , and the metric d_h are defined as in Theorem 1. \square

Proof: It can be shown by direct calculation that the approximation error increments with k according to the following upper bound:

$$|p_{s_0, \nu}^k(\langle s \rangle) - p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)| \leq \left(\frac{\lambda_S}{\delta^n} \right)^k \sum_{i=1}^k \varepsilon_1^i. \quad (6)$$

This bound corresponds to an uncertainty interval explosion, when elevating the interval matrix $[\mathcal{M}]$ to the power ok . The ergodicity assumption on $[\mathcal{M}]$ implies, by Theorem 1, that there exist a steady-state distribution $p_{\bar{\nu}}^\infty$ and constants α, β such that:

$$d_h(p_{\bar{\nu}}^\infty, p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)) \leq \alpha \beta^k, \quad (7)$$

We recall that $p_{\bar{\nu}}^\infty$ and $p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)$ are intervals of probability distributions.

By the computation of the transition probabilities and bounds of our abstraction, the stochastic behavior $p_{s_0, \bar{\nu}}^k$ generated by $[\mathcal{M}]$ is conservative with respect to the stochastic behavior $p_{s_0, \nu}^k$ generated by $\Pi(k)$. In fact $\forall k \geq 0, \Pi(k) \in [\mathcal{M}]$ by construction of $[\mathcal{M}]$. Definition 9 implies that:

$$\mathcal{T}(\Pi(k)) \leq \mathcal{T}([\mathcal{M}]). \quad (8)$$

Equation (8) implies that a limit $p_{\bar{\nu}}^\infty$ exists and belongs to the steady state interval of the abstraction

$$d_h(p_{\bar{\nu}}^\infty, p_{\nu}^\infty) \leq \Delta(p_{\bar{\nu}}^\infty), \quad (9)$$

and that the convergence speed of \mathcal{H} is bounded by the convergence speed of \mathcal{M}

$$d_h(p_{\nu}^\infty, p_{s_0, \nu}^k(\langle s \rangle)) \leq \alpha \beta^k. \quad (10)$$

By the triangular inequality, and by equations (7), (10) and (9), the following holds:

$$\begin{aligned} & d_h(p_{s_0, \nu}^k(\langle s \rangle), p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)) \\ & \leq d_h(p_{s_0, \nu}^k(\langle s \rangle), p_{\nu}^\infty) + d_h(p_{\nu}^\infty, p_{\bar{\nu}}^\infty) + d_h(p_{\bar{\nu}}^\infty, p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)) \\ & \leq 2\alpha \beta^k + \Delta(p_{\bar{\nu}}^\infty) \end{aligned} \quad (11)$$

Equations (6) and (11) imply the following inequality:

$$\begin{aligned} & d_h(p_{s_0, \nu}^k(\langle s \rangle), p_{\bar{s}_0, \bar{\nu}}^k(\langle s \rangle)) \leq \\ & \leq \min \left\{ \left(\frac{\lambda_S}{\delta^n} \right)^k \sum_{i=1}^k \varepsilon_1^i, 2\alpha \beta^k + \Delta(p_{\bar{\nu}}^\infty) \right\} \end{aligned} \quad (12)$$

Theorem 2 implies that:

$$\Delta(p_{\bar{\nu}}^\infty) \leq |\{\mathcal{S}\}_\delta| \varepsilon_1 \leq \frac{\lambda_S}{\delta^n} \frac{2\delta^n(M\eta + L\delta)}{1 - \mathcal{T}([\mathcal{M}])} = \frac{2\lambda_S(M\eta + L\delta)}{1 - \mathcal{T}([\mathcal{M}])},$$

and the result follows. \blacksquare

Equation (5) provides a bound for the approximation error, for each time step $k \geq 0$. If the abstraction is endowed with some ergodicity, by tuning the partition parameters (δ, η) of $(\{\mathcal{S}\}_\delta, \{\mathcal{A}\}_\eta)$ it is thus possible to achieve any desired precision on the error.

Theorem 4: Given \mathcal{H} and an homogeneous policy ν , if a stationary probability distribution \mathbf{p}_ν of \mathcal{H} exists, then there exist $\delta > 0, \eta > 0$, such that $\mathcal{T}([\mathcal{M}]) < 1$. \square

Proof: Let us call $P_{s_0, \nu}^k$ the probability measure associated with the SHS \mathcal{H} at time k , associated with a deterministic initial condition s_0 and with the control ν [3]. Recall that, for a given a ν , a stationary distribution for \mathcal{H} $\mathbf{p}_\nu : \mathcal{S} \rightarrow [0, 1]$ is such that, $\forall (q, x), (q', x') \in \mathcal{S}, \mathbf{p}_\nu(q, x) = \int_{\mathcal{S}} T((dx', q')|(q, x), \nu, q') \mathbf{p}_\nu(q, x)$. Furthermore, the following holds, for any $s_0 \in \mathcal{S}$:

$$\lim_{k \rightarrow \infty} d(P_{s_0, \nu}^k, \mathbf{p}_\nu) = 0, \quad (13)$$

where d is a proper distance between probability measures. As before, let $\Pi(k)$ be the transition probability matrix of the non homogeneous Markov chain that generates $p_{s_0, \nu}^k(\langle s \rangle)$, obtained as the restriction of the distribution $P_{s_0, \nu}^k$ on \mathcal{S} on the sets $\langle s \rangle \in \{\mathcal{S}\}_\delta$. Since for each $k \geq 0$

$$p_{s_0, \nu}^k(\langle s \rangle) = \int_{\langle s \rangle} P_{s_0, \nu}^k(s) ds,$$

then equation (13) implies that a steady state probability distribution for $\Pi(k)$ exists, and thus $\mathcal{T}(\Pi(k)) < 1$. Since $\lim_{\delta, \eta \rightarrow 0} [\mathcal{M}]$ is a stochastic matrix (and not an interval) and $\forall k \geq 0, \Pi(k) \in [\mathcal{M}]$, the following holds:

$$\lim_{\delta, \eta \rightarrow 0} \mathcal{T}([\mathcal{M}]) = \lim_{\delta, \eta \rightarrow 0} \max_{A \in [\mathcal{M}]} \mathcal{T}(A) = \mathcal{T}(\Pi(k)) < 1.$$

This implies that for any $\gamma > 0$, there exist finite $\delta > 0, \eta > 0$ such that $|\mathcal{T}([\mathcal{M}]) - \mathcal{T}(\Pi(k))| < \gamma$. Thus, for $\bar{\gamma} < 1 - \mathcal{T}(\Pi(k))$, there exist $\bar{\delta} > 0, \bar{\eta} > 0$ such that $\mathcal{T}([\mathcal{M}]) < 1$. \blacksquare

The theorem above guarantees that, given a stochastic system \mathcal{H} with known asymptotics, it is possible to tune the parameters pair (δ, η) in order to construct an approximate abstraction that is arbitrarily close to \mathcal{H} .

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: $y = [\text{SigH}]$ (concentration of a sigma factor in the environment), $z = [\text{SpaS}]$ (concentration of subtilin), X (food 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, and the dynamics are obtained according to a first order Euler scheme, with time-step Δ . The variables are bounded below by zero and 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 food-dependent spatial condition $\{X = \eta D_M\}$. Here $0 \leq \eta \leq 1$. 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:

$$P_1 = \begin{bmatrix} 1 - b_0 & b_0 \\ b_1 & 1 - b_1 \end{bmatrix}. \quad (14)$$

The parameters b_0, b_1 depend on [SigH] according to [12]

$$\begin{aligned} b_0(y) &= \frac{\alpha y}{1 + \alpha y}, \\ b_1(y) &= 1 - b_0(y). \end{aligned}$$

The quantity $\alpha = e^{-\Delta G_{rk}/RT}$ depends on the Gibbs free energy, a gas constant and the environment temperature.

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

$$D(k+1) = D(k) + rD(k) \left(1 - \frac{D(k)}{D_\infty}\right) \Delta, r > 0. \quad (15)$$

This is a quadratic equation, with two equilibria. The non-trivial (and stable) equilibrium relation depends on the quantity D_∞ , which is known as the *carrying capacity*. Let us a priori define it to be equal to $D_\infty = \frac{X}{X_M} D_M$. In other words, as intuitive, the steady state dynamically depends on the relative quantity of food in the environment. The food dynamics follow the difference equation

$$X(k+1) = X(k) + (k_2 \nu z(k) - k_1 D(k) X(k)) \Delta, \nu < 1. \quad (16)$$

Notice that the population level has two equilibria. The first ($D = 0$) is unstable, while the second ($D = D_\infty$) is stable. To this second equilibrium point corresponds a stable equilibrium for the food level X .

The above set of dynamical relations can be reframed as a SHS. The new 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 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 (14):

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

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 .² The order of the cardinality of the partition is easily $|\mathcal{Q}|(\max\{y_M, z_M, X_M, D_M\}/\Delta)^4$. Clearly $\lambda_S = y_M z_M X_M D_M$. The introduced error is $\varepsilon_1 = \Delta^5 \left(\frac{1}{\sqrt{2}} \frac{e^{y_M} - 1}{y_M} + \alpha\right)$.

As in [12], we have chosen $y_M = 4, z_M = 4, X_M = 10, D_M = 1, \lambda_1 = 0.2, \lambda_3 = 0.2, k_1 = 0, k_3 = 0.5, k_5 = 1, r = 0.02, \nu = 0.1, k_2 = 0.4, \alpha = 0.4$. Δ has been chosen to be equal to 0.01. The reference simulations have been implemented with a Monte Carlo approach, by running ten simulations with starting states corresponding to the representative points of the abstraction, thus assuming uniformly distributed probability for the initial points.

We have then implemented a complete abstraction of the above SHS dynamics with discretization level $\delta = 0.5$. The MC has 1260 rows. The outputs for population level and food are in Table I and are compared to those of the Monte Carlo Simulations. While the steady states of the population and of the food appear to be close to the desired ones, those of the cellular dynamics, here not reported, are not satisfactory. This is possibly due to the sparsity of the obtained MC: the presence of deterministic dynamics in fact introduce a number of null terms.

We have then decided to reduce the abstraction down to the cellular dynamics. This has been motivated by their relatively fast dynamics, as well as their noise level. More precisely, we have come up with two different abstractions, one for each of the two regions ON_y, OFF_y , for the dynamics of y and z ([SigH] and [SpaS]). The outputs for population and food are close to those of the Monte Carlo Simulations (see Table I). This has also speeded up the abstraction computation time so that we can push the discretization level to be quite small. Notice the computational improvement that the abstraction procedure gains, despite the required procedure and calculation of the steady state.

In Figures 1 and 2 we plot single realizations of the reduced-abstraction dynamics (top rows) and a single realization out

²We thus avoid that spatial boundaries split partition cells.

TABLE I
STEADY STATE COMPUTATIONS

	Population Level	Food 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

of the Monte Carlo simulations (bottom rows), representing the long-term behavior of population and food levels.

The take-away point of the case study is that the abstraction procedure appears to be significant especially in the presence of fully stochastic dynamics.




Fig. 1. Comparison between reduced-abstraction simulations (top row) and single realization of the actual dynamics (bottom row), for initial conditions $D(0) = 0.5$ (population, left column) and $X(0) = 9.5$ (food, right column).

VII. CONCLUSIONS

This work has extended an abstraction procedure for control Stochastic Hybrid Systems, using Markov set-Chains as abstraction class. By raising some continuity assumptions on the stochastic kernels, we have derived an explicit relation between the state and control space partition accuracy and the approximation error of the abstraction. We have furthermore derived a bound for the *distance* between the transition probabilities of the abstract system and those of the original system, evaluated at each time instant on the regions of the partition. Finally, we have tested the approximate abstraction technique on a dynamical model from systems biology. The main avenue for future research is to provide a definition of approximate bisimulation, in order to employ the present framework to stochastic model checking.

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. Egerst-




Fig. 2. Comparison between reduced-abstraction simulations (top row) and single realization of the actual dynamics (bottom row), for initial conditions $D(0) = 0.5$ (population, left column) and $X(0) = 0.5$ (food, right column).

edt 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*, 2007. Accepted for publication.

[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. Bicchi, A. Marigo, and B. Piccoli. On the reachability of quantized control systems. *IEEE Transactions on Automatic Control*, 47-4:546–563, 2002.

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

[7] 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.

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

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

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

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

[15] Giordano Pola, Antoine Girard, and Paulo Tabuada. Symbolic models for nonlinear control systems using approximate bisimulations. In

Proceedings of the 46th Conference on Decision and Control, New Orleans, LA, USA, December 2007.

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