A numerical approximation scheme for reachability analysis of stochastic hybrid systems with state-dependent switchings

Maria Prandini and Jianghai Hu

***Abstract*— We describe a methodology for reachability analysis of a certain class of stochastic hybrid systems, whose continuous dynamics is governed by stochastic differential equations and discrete dynamics by state-dependent probabilistic transitions. The main feature of the proposed methodology is that it rests on the weak approximation of the solution to the stochastic differential equation with random mode transitions by a Markov chain. Reachability computations then reduce to propagating the transition probabilities of the approximating Markov chain. An example of applications to system verification is presented.**

I. INTRODUCTION

Hybrid systems represent a powerful modeling framework and have been applied to a diverse range of engineering problems (air traffic management systems [15] automated highway systems [16], robotics [7], computer and communication networks [9], automotive systems [3]) and also to model biological systems [2].

The majority of the efforts in the study of hybrid systems have been devoted to deterministic hybrid systems, which, however, are not suitable for modeling practical systems with inherent uncertainty.

In stochastic hybrid systems, probabilistic laws govern the discrete/continuous dynamics of the system. As a result, each trajectory is weighted according to its likelihood as determined by these probabilistic laws. The problems one can study of stochastic hybrid systems are then more various and less 'sharp' than those of deterministic hybrid systems, and the results obtained are often more robust and less conservative. In particular, a reachability problem in the deterministic case is a yes/no problem, while in the stochastic case one faces a continuous spectrum of 'soft' problems with quantitative answers, such as the hitting probability, the expected hitting time, the hitting distribution, etc. As a counterpart for their enhanced modeling flexibility and expressiveness, the problems arising in stochastic hybrid systems are in general much more challenging: analytical solutions are difficult or impossible to obtain; and few effective general algorithms exist for the numerical simulation of stochastic hybrid systems, compared with the many software packages for deterministic hybrid systems. Many problems well studied in the deterministic case remain open for stochastic hybrid systems.

In this paper, we focus on reachability problems for stochastic hybrid systems. In [4], [5], theoretical issues regarding the measurability of the reachability events are addressed. In [5] and [6], upper bounds on the probability of reachability events are derived based on the theory of Dirichlet forms associated with a right-Markov process and on certain functions of the state of the system known as barrier certificates, respectively. In [1], stochastic reachability is addressed in the discrete time case by dynamic programming. Here, we study the reachability problem of estimating the probability that the system state will enter some region of the state space within a finite time horizon, starting from an arbitrary initial condition, for a class of continuous time stochastic hybrid systems generating switching diffusion processes, [8].

By discretizing the state space and using an interpolated Markov chain to approximate the switching diffusion weakly, we develop an algorithm to compute an estimate of the probability of interest. This numerical scheme is amenable for applications to system verification and safety analysis.

The proposed methodology for reachability analysis was introduced by the authors of the present paper in [10], and further developed in [13], with reference to systems described by stochastic differential equations with coefficients switching value at prescribed, a-priori known, time instants. The methodology is extended here to a hybrid setting, where switchings are state dependent, thus proving that the Markov chain weak approximation result in [12] is still valid in a hybrid setting. A more detailed version of this conference paper is available as a book chapter in [14].

II. SWITCHING DIFFUSIONS

We consider a continuous time stochastic hybrid system whose hybrid state $\mathbf{s}$ is characterized by a continuous component $\mathbf{x}$ and a discrete component $\mathbf{q}$, taking on value in the Euclidean space $\mathbb{R}^n$ and in the finite set $\mathcal{Q} = \{1, 2, \dots M\}$: $\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{q}(t)) \in \mathcal{S} := \mathbb{R}^n \times \mathcal{Q}$, $t \geq 0$.

The evolution of $\mathbf{q}$ is piecewise constant and right continuous, that is for each trajectory of $\mathbf{q}$, there exists a sequence of consecutive left closed, right open time interval $\{T_i,\, i = 0, 1, \dots\}$ such that $\mathbf{q}(t) = q_i$, $\forall t \in T_i$ and $q_i \neq q_{i\pm 1}$.

During each time interval T_i when $\mathbf{q}(t)$ is constant and equal to $q_i \in \mathcal{Q}$, the continuous state component $\mathbf{x}$ is governed by the stochastic differential equation (SDE)

$$d\mathbf{x}(t) = a(\mathbf{x}(t),\, q_i)dt + b(\mathbf{x}(t),\, q_i)\, \Sigma\, d\mathbf{w}(t),$$

This work was supported by MIUR (Ministero dell'Istruzione, dell'Università e della Ricerca) under the project *Identification and adaptive control of industrial systems*, and by the National Science Foundation under Grant CNS-0643805.

M. Prandini is with Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy `prandini@elet.polimi.it`

J. Hu is with School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA `jianghai@purdue.edu`

initialized with $\mathbf{x}(t_i^-) = \lim_{h\to 0^+}\mathbf{x}(t_i - h)$ at time $t_i := \sup\{\cup_{k=0}^{i-1} T_k\}$, where $a(\cdot, q_i) : \mathbb{R}^n \to \mathbb{R}^n$ and $b(\cdot, q_i) : \mathbb{R}^n \to \mathbb{R}^{n\times n}$are the drift and diffusion terms, and Σ is a diagonal matrix with positive entries, which modulates the variance of the standard n-dimensional Brownian motion $\mathbf{w}$. A jump in the discrete state may occur during the continuous state evolution with an intensity and according to a probabilistic reset map that both depend on the current value taken by $\mathbf{s}$. More specifically, $\mathbf{q}$ is a continuous time process, whose evolution at time t is conditionally independent on the past given $\mathbf{s}(t^-) = (x, q) \in \mathcal{S}$, and is governed by the transition probabilities

$$P\{\mathbf{q}(t+\Delta) = q' | \mathbf{s}(t^-) = (x, q)\} = \lambda_{qq'}(x)\Delta + o(\Delta),$$

$q' \neq q \in \mathcal{Q}$, where the transition rate $\lambda_{qq'} : \mathbb{R}^n \to \mathbb{R}$ satisfies $\lambda_{qq'}(x) \geq 0$, $\forall x \in \mathbb{R}^n$. The transition rate functions determine both the switching intensity and the reset map of the discrete state $\mathbf{q}$. During the infinitesimal time interval $[t, t+\Delta]$, $\mathbf{q}(t)$ will jump once with probability $\lambda(s)\Delta + o(\Delta)$, and two or more times with probability $o(\Delta)$, starting from $\mathbf{s}(t^-) = s$, where $\lambda : \mathcal{S} \to [0, +\infty)$

$$\lambda(s) = \sum_{q'\in\mathcal{Q}, q'\neq q} \lambda_{qq'}(x),\ s = (x, q) \in \mathcal{S}, \tag{1}$$

is the jump intensity function. If $s \in \mathcal{S}$ is such that $\lambda(s) = 0$, then no instantaneous jump can occur from s. Let $s \in \mathcal{S}$ be such that $\lambda(s) \neq 0$. Then, when a jump occurs at time t from $\mathbf{s}(t^-) = s$, the distribution of $\mathbf{q}(t)$ over $\mathcal{Q}$ is given by the reset function $R : \mathcal{S} \times \mathcal{Q} \to [0, 1]$

$$R(s, q') = \begin{cases} \frac{\lambda_{qq'}(x)}{\lambda(s)}, & q' \neq q \\ 0, & q' = q \end{cases}, \quad s = (x, q). \tag{2}$$

Assumption 1: $a(\cdot, q)$, $b(\cdot, q)$, and $\lambda_{qq'}(\cdot)$ are bounded and Lipschitz continuous for any $q, q' \in \mathcal{Q}$.

Under Assumption 1, process $\mathbf{s}(t)$, $t \geq 0$, generated by the described stochastic hybrid system initialized with $\mathbf{s}(0) = s_0 \in \mathcal{S}$ is a switching diffusion.

A formal description of a switching diffusion, with the pure jump process $\mathbf{q}$ represented by an integral with respect to a Poisson random measure, is provided next, following [8] and [11]. This is the reference representation for the proposed numerical approximation scheme for reachability analysis.

For each $x \in \mathbb{R}^n$, define the consecutive disjoint intervals $\Gamma_{ki}(x)$ of length $\lambda_{ki}(x)$, $k, i \in \mathcal{Q}$, $i \neq k$, as follows:

$$\Gamma_{12}(x) = \big[0, \lambda_{12}(x)\big)$$
$$\Gamma_{13}(x) = \big[\lambda_{12}(x), \lambda_{12}(x) + \lambda_{13}(x)\big)$$
$$\vdots$$
$$\Gamma_{1M}(x) = \big[\sum_{h=2}^{M-1} \lambda_{1h}(x), \sum_{h=2}^{M} \lambda_{1h}(x)\big)$$
$$\Gamma_{21}(x) = \big[\sum_{h=2}^{M} \lambda_{1h}(x), \sum_{h=2}^{M} \lambda_{1h}(x) + \lambda_{21}(x)\big)$$
$$\Gamma_{23}(x) = \big[\sum_{h=2}^{M} \lambda_{1h}(x) + \lambda_{21}(x), \sum_{h=2}^{M} \lambda_{1h}(x) + \lambda_{21}(x) + \lambda_{23}(x)\big)$$
$$\vdots$$

Set $\lambda_{\max} = \max_{x\in\mathbb{R}^n} \sum_{k,i=1,i\neq k}^{M} \lambda_{ki}(x)$, and let Π denote the uniform probability measure on $\Gamma_{\max} = [0, \lambda_{\max})$. Define $r_q : \mathbb{R}^n \times \mathcal{Q} \times \Gamma_{\max} \to \mathbb{Z}$ describing the entity of a jump in $\mathbf{q}$ from $(x, q) \in \mathcal{S}$ corresponding to $\gamma \in \Gamma_{\max}$ as

$$r_q(x, q, \gamma) = \begin{cases} q' - q, & \text{if } \gamma \in \Gamma_{qq'}(x) \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

With these notations, the stochastic hybrid system can be represented by the SDEs

$$\begin{aligned} d\mathbf{x}(t) &= a(\mathbf{x}(t),\, \mathbf{q}(t))dt + b(\mathbf{x}(t),\, \mathbf{q}(t))\, \Sigma\, d\mathbf{w}(t) \\ d\mathbf{q}(t) &= \int_{\Gamma_{\max}} r_q(\mathbf{x}(t^-), \mathbf{q}(t^-), \gamma)\mathbf{p}(dt, d\gamma), \end{aligned} \tag{4}$$

where $\mathbf{p}(\cdot,\cdot)$ is a Poisson random measure of intensity $\lambda_{\max} dt \times \Pi(d\gamma)$, independent of $\mathbf{w}(\cdot)$.

$\mathbf{p}(\cdot,\cdot)$ generates independently of $\mathbf{w}(\cdot)$ the sequence $\{\mathbf{t}_i,\ i \geq 1\}$ of increasing nonnegative random variables representing the jump times of a standard Poisson process with intensity $\lambda_{\max}$, and the sequence $\{\boldsymbol{\gamma}_i,\ i \geq 1\}$ of i.i.d. random variables uniformly distributed over $\Gamma_{\max}$ and independent of $\{\mathbf{t}_i,\ i \geq 1\}$. Since $\mathbf{p}(d\tau, d\gamma)$ assigns unit mass to (τ, γ) if there exists $i \geq 1$ such that $\mathbf{t}_i = \tau$ and $\boldsymbol{\gamma}_i = \gamma$, then, for any measurable $C \subseteq \Gamma_{\max}$ and any $t \in T$, $\int_0^t \int_C \mathbf{p}(d\tau d\gamma) = \sum_{i\geq 1} 1_{\{\mathbf{t}_i \leq t\}} 1_{\{\boldsymbol{\gamma}_i \in C\}}$ is a random quantity representing the number of jumps during the time interval $[0, t]$ corresponding to values for γ in C. Thus the $\mathbf{q}$ process solving (4) initialized with $s_0 = (x_0, q_0)$ is given by

$$\begin{aligned} \mathbf{q}(t) &= q_0 + \int_0^t \int_{\Gamma_{\max}} r_q(\mathbf{x}(\tau^-), \mathbf{q}(\tau^-), \gamma)\mathbf{p}(d\tau, d\gamma) \\ &= q_0 + \sum_{i\geq 1:\, \mathbf{t}_i \leq t} r_q(\mathbf{x}(\mathbf{t}_i^-), \mathbf{q}(\mathbf{t}_i^-), \boldsymbol{\gamma}_i), \end{aligned}$$

whereas, between each pair of time instants $\mathbf{t}_i$ and $\mathbf{t}_{i+1}$, the $\mathbf{x}$ component behaves as a diffusion process with local properties determined by $a(\cdot, \mathbf{q}(\mathbf{t}_i))$ and $b(\cdot, \mathbf{q}(\mathbf{t}_i))$.

Under Assumption 1, system (4) initialized with $s_0 \in \mathcal{S}$ admits a unique strong solution $\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{q}(t))$, $t \geq 0$. Moreover, $\mathbf{s}$ is a Markov process and the trajectories of the continuous component $\mathbf{x}$ are continuous.

III. Stochastic Reachability

Given a closed set $D \subset \mathbb{R}^n$, our objective is to determine the probability that the solution $\mathbf{x}(t)$ to (4) initialized with $s_0 = (x_0, q_0) \in \mathcal{S}$ reaches D during some look-ahead time horizon $T = [0, t_f]$:

$$P_{s_0}\big\{\mathbf{x}(t) \in D \text{ for some } t \in T\big\}. \tag{5}$$

This can be of interest in system verification and safety analysis, and in such cases D represents a target set rather than an unsafe set.

To evaluate the probability (5) numerically, we introduce a bounded open set $U \subset \mathbb{R}^n$ where to confine $\mathbf{x}$. If D represents an unsafe set, the domain U should be chosen large enough so that the situation can be declared safe once $\mathbf{x}$ ends up outside U. If D represents a target set, U should be chosen large enough so that the objective of reaching the target is failed in practice once $\mathbf{x}$ exits U. This makes sense, for example, in those regulation problems where the system should be driven to operate close to some reference state value $x^\star$ and deviations from the desired operating point in the transient are allowed only to some extent (see the example in Section VI).

Let U^c denote the complement of U in $\mathbb{R}^n$. Then, with reference to the domain U, the probability of entering D can be expressed as

$$\mathrm{P}_{s_0} := P_{s_0}\big\{\mathbf{x} \text{ hits } D \text{ before hitting } U^c \text{ within } T\big\}. \quad (6)$$

Hence, for the purpose of computing (6), we can assume that $\mathbf{x}$ in (4) is defined on the open domain $U \setminus D$ with initial condition $x_0 \in U \setminus D$, and that $\mathbf{x}$ is stopped as soon as it hits the boundary $\partial U^c \cup \partial D$ of $U \setminus D$.

IV. Markov Chain Approximation

The switching diffusion process generated by the stochastic hybrid system described in Section II is approximated here by the piecewise constant interpolation of a suitably defined discrete time Markov chain. The interpolation time interval Δ_δ is a positive function of a gridding scale parameter δ tending to zero faster than δ ($\Delta_\delta = o(\delta)$).

The discrete time Markov chain $\{\mathbf{v}_k, k \geq 0\}$ is characterized by a two-component state: $\mathbf{v} = (\mathbf{z}, \mathbf{m})$, where $\mathbf{m}$ takes on value in the finite set $\mathcal{Q}$, whereas $\mathbf{z}$ takes on value in a finite set $\mathcal{Z}_\delta$ obtained by gridding $U \setminus D$.

Recall that the $\mathbf{q}$ component of the switching diffusion $\mathbf{s} = (\mathbf{x}, \mathbf{q})$ is a pure jump process. The jump occurrences in $\mathbf{q}$ are governed by a standard Poisson process with intensity $\lambda_{\max}$, which is independent of the random variables determining the jump entity and of the Brownian motion affecting the $\mathbf{x}$ component. The distribution of the inter-jump times is exponential with coefficient $\lambda_{\max}$, so that the probability of a single jump within a time interval Δ is $\lambda_{\max}\Delta + o(\Delta)$. The jump entity is state dependent, and jumps of zero entity may occur. When a jump (possibly of zero entity) occurs at time t, then, the $\mathbf{x}$ component is reinitialized with the same value $\mathbf{x}(t^-)$ prior to the jump occurrence.

In order to take this into account when defining the transition probabilities of the approximating Markov chain $\{\mathbf{v}_k, k \geq 0\}$, we introduce an enlarged Markov chain process $\{(\mathbf{v}_k, \mathbf{j}_k), k \geq 0\}$, where $\{\mathbf{j}_k, k \geq 0\}$ is a Bernoulli process representing the jump occurrences: if $\mathbf{j}_k = 1$, then a jump, possibly of zero entity, occurs at time k, whereas if $\mathbf{j}_k = 0$, then no jump occurs at time k. If $\mathbf{j}_k$ is independent of $\mathbf{v}_i$, $i = 0, 1, \dots, k$, $\forall k \geq 0$, then, it is easily shown that $\{\mathbf{v}_k, k \geq 0\}$ is a Markov chain with transition probabilities given by

$$\begin{aligned} &P_\delta\big\{\mathbf{v}_{k+1} = v' \mid \mathbf{v}_k = v\big\} \\ &\quad = \sum_{j \in \{0,1\}} P_\delta\big\{\mathbf{v}_{k+1} = v' \mid \mathbf{v}_k = v,\, \mathbf{j}_k = j\big\} P_\delta\big\{\mathbf{j}_k = j\big\}. \end{aligned}$$

We now define the jump probability $P_\delta\big\{\mathbf{j}_k = 1\big\}$, the inter macro-states transition probability $P_\delta\big\{\mathbf{v}_{k+1} = v' \mid \mathbf{v}_k = v,\, \mathbf{j}_k = 1\big\}$, and the intra macro-states transition probability $P_\delta\big\{\mathbf{v}_{k+1} = v' \mid \mathbf{v}_k = v,\, \mathbf{j}_k = 0\big\}$.

Jump probability: We set

$$P_\delta\big\{\mathbf{j}_k = 1\big\} = 1 - e^{-\lambda_{\max}\Delta_\delta} = \lambda_{\max}\Delta_\delta + o(\Delta_\delta), \quad (7)$$

which tends to the jump rate of the standard Poisson process generating jumps in $\mathbf{q}$, as $\delta \to 0$.

Inter macro-states transitions: If $\mathbf{j}_k = 1$ (a jump occurs at time k), then, $\mathbf{z}_{k+1} = \mathbf{z}_k$, whereas the value of $\mathbf{m}_{k+1}$ is determined based on that of $\mathbf{v}_k$ through the (conditional) transition probabilities $p_\delta(q \to q'|z) := P_\delta\big\{\mathbf{m}_{k+1} = q' \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 1\big\}$. As a result,

$$\begin{aligned} &P_\delta\big\{(\mathbf{z}_{k+1}, \mathbf{m}_{k+1}) = (z', q') \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 1\big\} \\ &\quad = \begin{cases} 0, & z' \neq z \\ p_\delta(q \to q'|z), & z' = z. \end{cases} \end{aligned}$$

To reproduce the behavior of $\mathbf{q}$ when a jump occurs, we set

$$\begin{aligned} p_\delta(q \to q'|z) &= \int_{\{\gamma : r_q(z,q,\gamma) = q'-q\}} \Pi(d\gamma) \\ &= \begin{cases} \frac{\lambda_{q\,q'}(z)}{\lambda_{\max}}, & q' \neq q \\ 1 - \frac{1}{\lambda_{\max}} \sum_{q^* \in \mathcal{Q}, q^* \neq q} \lambda_{q\,q^*}(z), & q' = q. \end{cases} \end{aligned}$$

This way, the probability distribution of $\mathbf{m}_{k+1}$ when a jump of non-zero entity occurs at time k from (z,q) is

$$\begin{aligned} &P_\delta\big\{\mathbf{m}_{k+1} = q' \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 1, \mathbf{m}_{k+1} \neq \mathbf{m}_k\big\} \\ &\quad = \frac{\lambda_{q\,q'}(z)}{\sum_{q^* \in \mathcal{Q}, q^* \neq q} \lambda_{q\,q^*}(z)} = R((z,q), q'), \end{aligned}$$

where $R(\cdot,\cdot)$ is the reset function defined in (2). Also, the probability that a jump of non-zero entity occurs at time k from (z,q) is given by

$$\begin{aligned} &P_\delta\big\{\mathbf{j}_k = 1, \mathbf{m}_{k+1} \neq q \mid \mathbf{v}_k = (z,q)\big\} \\ &= \sum_{q' \in \mathcal{Q}, q' \neq q} P_\delta\big\{\mathbf{m}_{k+1} = q' \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 1\big\} P_\delta\big\{\mathbf{j}_k = 1\big\} \\ &= \big(1 - e^{-\lambda_{\max}\Delta_\delta}\big) \frac{\sum_{q' \in \mathcal{Q}, q' \neq q} \lambda_{q\,q'}(z)}{\lambda_{\max}} = \lambda(z)\Delta_\delta + o(\Delta_\delta), \end{aligned}$$

where $\lambda(\cdot)$ is the jump intensity function defined in (1).

Intra macro-state transitions: If $\mathbf{j}_k = 0$ (no jump occurs at time k), then, $\mathbf{m}_{k+1} = \mathbf{m}_k$, whereas the value of $\mathbf{z}_{k+1}$ is determined based on that of $\mathbf{v}_k$, through the (conditional) transition probabilities

$$p_\delta(z \to z'|q) := P_\delta\big\{\mathbf{z}_{k+1} = z' \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 0\big\} \quad (8)$$

describing the evolution of $\mathbf{z}$ within the “macro-state” $q \in \mathcal{Q}$. Then,

$$\begin{aligned} &P_\delta\big\{(\mathbf{z}_{k+1}, \mathbf{m}_{k+1}) = (z', q') \mid \mathbf{v}_k = (z,q), \mathbf{j}_k = 0\big\} \\ &\quad = \begin{cases} 0, & q' \neq q \\ p_\delta(z \to z'|q), & q' = q. \end{cases} \end{aligned}$$

For the weak convergence result to hold, the probabilities (8) should be suitably selected so as to approximate ‘locally’ the evolution of the $\mathbf{x}$ component of the switching diffusion $\mathbf{s} = (\mathbf{x}, \mathbf{q})$ with absorption on the boundary $\partial U^c \cup \partial D$ when no jump occurs in $\mathbf{q}$.

To formally define this ‘local consistency’ notion, we need first to introduce some notations. Let $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \dots, \sigma_n)$ with $\sigma_i > 0$, $i = 1, \dots, n$. Fix a grid parameter $\delta > 0$. Denote by $\mathbb{Z}^n_\delta$ the integer grids of $\mathbb{R}^n$ scaled according to the grid parameter δ and the positive diagonal entries of matrix Σ as follows $\mathbb{Z}_\delta^n = \{(m_1\eta_1\delta, m_2\eta_2\delta, \dots, m_n\eta_n\delta)| \;\; m_i \in \mathbb{Z}, i = 1,\dots,n\}$, where $\eta_i := \frac{\sigma_i}{\sigma_{\max}}$, $i = 1,\dots,n$, with $\sigma_{\max} = \max_i \sigma_i$. For each grid point $z \in \mathbb{Z}_\delta^n$, define the immediate neighbors set as a subset of $\mathbb{Z}_\delta^n$ whose distance from z along the coordinate axis x_i is at most $\eta_i\delta$, $i = 1,\dots,n$. Formally,

$$\mathcal{N}_\delta(z) = \{z + (i_1\eta_1\delta,\dots,i_n\eta_n\delta) \in \mathbb{Z}_\delta^n | \; (i_1,\dots,i_n) \in \mathcal{I}\},$$

where $\mathcal{I} \subseteq \{0,1,-1\}^n \setminus \{(0,0,\dots,0)\}$.

The finite set $\mathcal{Z}_\delta$ where $\mathbf{z}$ takes on value is defined as the set of all those grid points in $\mathbb{Z}_\delta^n$ that lie inside U but outside D: $\mathcal{Z}_\delta = (U \setminus D) \cap \mathbb{Z}_\delta^n$. $\mathcal{N}_\delta(z)$ represents the set of states to which $\mathbf{z}$ can evolve in one time step within a macro-state, starting from z. The interior $\mathcal{Z}_\delta^\circ$ of $\mathcal{Z}_\delta$ consists of all those points in $\mathcal{Z}_\delta$ which have all their neighbors in $\mathcal{Z}_\delta$. The boundary of $\mathcal{Z}_\delta$ is given by $\partial\mathcal{Z}_\delta = \mathcal{Z}_\delta \setminus \mathcal{Z}_\delta^\circ$. $\partial\mathcal{Z}_\delta$ is the union of the set $\partial\mathcal{Z}_{\delta U^c}$ of points with at least one neighbor outside U, and the set $\partial\mathcal{Z}_{\delta D}$ of points with at least one neighbor inside D. The points that satisfy both the conditions are assigned all to either $\partial\mathcal{Z}_{\delta D}$ or $\partial\mathcal{Z}_{\delta U^c}$, so as to make these two sets disjoint. The estimation error eventually introduced becomes negligible if U is chosen sufficiently large.

For each $q \in \mathcal{Q}$, we define $p_\delta(z \to z'|q)$ in (8) so that:

- each state z in $\partial\mathcal{Z}_\delta$ is an absorbing state;
- starting from a state z in $\mathcal{Z}_\delta^\circ$, $\mathbf{z}$ moves to one of its neighbors in $\mathcal{N}_\delta(z)$ or stays at the same state according to probabilities determined by its current location:

$$p_\delta(z \to z'|q) = \begin{cases} \pi_\delta(z'|(z,q)), & z' \in \mathcal{N}_\delta(z) \cup \{z\} \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

$z \in \mathcal{Z}_\delta^\circ$, where $\pi_\delta(z'|(z,q))$ are appropriate functions of the drift and diffusion terms in (4) evaluated at (z,q).

Fix some time step k and consider the conditional mean and variance of the finite difference $\mathbf{z}_{k+1} - \mathbf{z}_k$ given that $\mathbf{v}_k = (z,q)$ and $\mathbf{j}_k = 0$:

$$m_\delta(z,q) = \sum_{z' \in \mathcal{N}_\delta(z)} (z'-z)\pi_\delta(z'|(z,q)),$$
$$V_\delta(z,q) = \sum_{z' \in \mathcal{N}_\delta(z)} (z'-z)(z'-z)^T\pi_\delta(z'|(z,q)).$$

The immediate neighbors set $\mathcal{N}_\delta(z)$ and the family of distribution functions $\{\pi_\delta(\cdot|(z,q)) : \mathcal{N}_\delta(z) \cup \{z\} \to [0,1],\, z \in \mathcal{Z}_\delta^\circ\}$ should be selected so that as $\delta \to 0$,

$$\frac{m_\delta(z,q)}{\Delta_\delta} \to a(x,q), \;\; \frac{V_\delta(z,q)}{\Delta_\delta} \to b(x,q)\Sigma^2 b(x,q)^T, \tag{10}$$

for all $x \in U \setminus D$, where, for any δ, z is a point in $\mathcal{Z}_\delta^\circ$ closest to x (local consistency property).

Example: Suppose that the diffusion term in (4) is of the form $b(s) = \beta(s)I$, where $\beta : \mathcal{S} \to \mathbb{R}$ is a scalar function and I is the identity matrix of size n. The immediate neighbors set $\mathcal{N}_\delta(z)$, $z \in \mathcal{Z}_\delta$, can be confined to the set of points:

$$\begin{array}{ll} z_{1_+} = z + (+\eta_1\delta, 0, \dots, 0) & z_{1_-} = z + (-\eta_1\delta, 0, \dots, 0) \\ z_{2_+} = z + (0, +\eta_2\delta, \dots, 0) & z_{2_-} = z + (0, -\eta_2\delta, \dots, 0) \\ \vdots & \\ z_{n_+} = z + (0, 0, \dots, +\eta_n\delta) & z_{n_-} = z + (0, 0, \dots, -\eta_n\delta). \end{array}$$

The transition probability function $\pi_\delta(\cdot|v)$ over $\mathcal{N}_\delta(z) \cup \{z\}$ from $v = (z,q) \in \mathcal{Z}_\delta^\circ \times \mathcal{Q}$ can be chosen as follows:

$$\pi_\delta(z'|v) = \begin{cases} c(v)\,\xi_0(v), & z' = z \\ c(v)\,e^{+\delta\xi_i(v)}, & z' = z_{i_+},\, i = 1,\dots,n \\ c(v)\,e^{-\delta\xi_i(v)}, & z' = z_{i_-},\, i = 1,\dots,n \end{cases} \tag{11}$$

with $\xi_0(v) = \frac{2}{\rho\sigma_{\max}^2\beta(v)^2} - 2n$, $\xi_i(v) = \frac{[a(v)]_i}{\eta_i\sigma_{\max}^2\beta(v)^2}$, $i = 1,\dots,n$, $c(v) = \frac{1}{2\sum_{i=1}^n \mathrm{csh}(\delta\xi_i(v)) + \xi_0(v)}$, where for any $y \in \mathbb{R}^n$, $[y]_i$ denotes the component of y along the x_i direction, $i = 1,2,\dots,n$. ρ is a positive constant satisfying $0 < \rho \le (n\sigma_{\max}^2 \max_{s \in (U \setminus D) \times \mathcal{Q}} \beta(s)^2)^{-1}$, so that $\xi_0(v)$ is positive for all $v \in \mathcal{Z}_\delta^\circ \times \mathcal{Q}$. The interpolation time interval Δ_δ can be set equal to $\rho\delta^2$. □

Different choices are possible so as to satisfy the local consistency property, and they affect the computational complexity of the approximation. The reader is referred to [12] for more details on this.

Discrete time Markov chain interpolation: Let $\{\boldsymbol{\Delta\tau}_k, k \ge 0\}$ be an i.i.d. sequence of random variables independent of $\{\mathbf{v}_k, k \ge 0\}$, exponentially distributed with mean value Δ_δ satisfying $\Delta_\delta > 0$ and $\Delta_\delta = o(\delta)$. Denote by $\{\mathbf{v}(t), t \ge 0\}$ the continuous time stochastic process that is equal to $\mathbf{v}_k$ on the time interval $[\boldsymbol{\tau}_k, \boldsymbol{\tau}_{k+1})$ for all k, where $\boldsymbol{\tau}_0 = 0$ and $\boldsymbol{\tau}_{k+1} = \boldsymbol{\tau}_k + \boldsymbol{\Delta\tau}_k$, $k \ge 0$.

Theorem 1: Suppose that the approximating Markov chain $\{\mathbf{v}_k, k \ge 0\}$ is initialized at a point $v_0 \in \mathcal{Z}_\delta^\circ \times \mathcal{Q}$ closest to $s_0 \in (U \setminus D) \times \mathcal{Q}$ and satisfies the local consistency properties (10). Then, under Assumption 1, the process $\{\mathbf{v}(t), t \ge 0\}$ obtained by interpolation of $\{\mathbf{v}_k, k \ge 0\}$ converges weakly as $\delta \to 0$ to the solution $\{\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{q}(t)),\, t \ge 0\}$ of (4) initialized with s_0, with $\mathbf{x}(t)$ defined on $U \setminus D$ and absorption on the boundary $\partial U^c \cup \partial D$. □

Weak convergence of Markov chain approximations is proven in [12, Theorem 4.1, Chapter 10] for jump diffusion processes satisfying appropriate regularity conditions. This theorem does not directly imply Theorem 1, because it would require $r_q(x,q,\gamma)$ in (3) to be continuous as a function of x, which is not the case. However, the continuity property shown in Proposition 1 below for the integral $\int_{\mathbb{R}} r_q(\cdot,q,\gamma)\Pi(d\gamma)$ can be used in the proof of [12, Theorem 4.1, Chapter 10] in place of the continuity of $r_q(\cdot,q,\gamma)$ to assess the weak convergence result.

Proposition 1: Suppose that $\lambda_{qq'}(\cdot)$ is bounded and Lipschitz continuous for any $q,q' \in \mathcal{Q}$, $q \neq q'$. Then, there exists a constant $C > 0$ such that $\int_{\mathbb{R}} |r_q(x,q,\gamma) - r_q(x',q,\gamma)|\Pi(d\gamma) \le C\,|x - x'|$, $\forall x,x' \in \mathbb{R}^n$, $q \in \mathcal{Q}$. □

This proposition can be proven following [11]. The proof is omitted due to space limitations.

V. Reachability Computations

Given the look-ahead time horizon $T = [0,t_f]$, fix $\delta > 0$ so that $k_f := \frac{t_f}{\Delta_\delta}$ is an integer, and construct the approximating Markov chain $\{\mathbf{v}_k = (\mathbf{z}_k, \mathbf{m}_k),\, k \ge 0\}$ satisfying Theorem 1.

If the diffusion matrix $\Sigma(s) := b(x,q)\Sigma^2 b(x,q)^T$, $s = (x,q) \in \mathcal{S}$, is uniformly positive definite over $(U \setminus D) \times \mathcal{Q}$, then appropriate regularity conditions on D and U^c guarantee with probability one that those pathological situations where the trajectory of the process $\mathbf{x}$ touches the absorbing boundary $\partial D \cup \partial U$ without leaving $U \setminus D$ do not occur, [11] and [12, Chapter 10]. By [12, Chapter 9, Theorem 1.5], weak convergence of the approximating Markov chain interpolation to $\mathbf{s}$ implies convergence with probability one to the probability of interest P_{s_0} in (6) of the corresponding quantity for $\{\mathbf{v}_k = (\mathbf{z}_k, \mathbf{m}_k),\, k \geq 0\}$:

$$\hat{\mathrm{P}}_{s_0} := P_\delta\big\{\mathbf{z}_k \text{ hits } \partial\mathcal{Z}_{\delta D} \text{ before } \partial\mathcal{Z}_{\delta U^c} \text{ within } [0, k_f]\big\}.$$

We now describe an iterative algorithm to compute $\hat{\mathrm{P}}_{s_0}$.

Observe that, since the boundaries $\partial\mathcal{Z}_{\delta U^c}$ and $\partial\mathcal{Z}_{\delta D}$ are both absorbing, then $\hat{\mathrm{P}}_{s_0} = P_\delta\big\{\mathbf{z}_k \in \partial\mathcal{Z}_{\delta D} \text{ for some } k \in [0, k_f]\big\} = P_\delta\big\{\mathbf{z}_{k_f} \in \partial\mathcal{Z}_{\delta D}\big\}$.

Let introduce the set of probability maps $\hat{\mathrm{p}}^{(k)} : \mathcal{Z}_\delta \times \mathcal{Q} \to [0,1]$, $k = 0, 1, \dots, k_f$, where

$$\hat{\mathrm{p}}^{(k)}(v) := P_\delta\big\{\mathbf{z}_{k_f} \in \partial\mathcal{Z}_{\delta D} \mid \mathbf{v}_{k_f - k} = v\big\}, \qquad (12)$$

is the probability of hitting $\partial\mathcal{Z}_{\delta D}$ before $\partial\mathcal{Z}_{\delta U^c}$ within $[k_f - k, k_f]$ starting from v at time k. Then,

$$\hat{\mathrm{p}}^{(0)}(v) = \begin{cases} 1, & \text{if } v \in \partial\mathcal{Z}_{\delta D} \times \mathcal{Q} \\ 0, & \text{otherwise,} \end{cases} \qquad (13)$$

and the desired quantity $\hat{\mathrm{P}}_{s_0}$ can be expressed as $\hat{\mathrm{P}}_{s_0} = \hat{\mathrm{p}}^{(k_f)}(v_0)$. Moreover, it is easily seen that the probability map $\hat{\mathrm{p}}^{(k_f)} : \mathcal{Z}_\delta \times \mathcal{Q} \to [0,1]$, can be can be computed by iterating k_f times starting from $k = 0$ the following equation:

$$\hat{\mathrm{p}}^{(k+1)}(v) = \sum_{v' \in \mathcal{Z}_\delta \times \mathcal{Q}} p_\delta(v \to v')\hat{\mathrm{p}}^{(k)}(v'),\ v \in \mathcal{Z}_\delta \times \mathcal{Q}.$$

Recalling that any $v \in \partial\mathcal{Z}_\delta \times \mathcal{Q}$ is an absorbing state and that, for each $k \in [0, k_f]$, $\hat{\mathrm{p}}^{(k)}(v) = 1$ if $v \in \partial\mathcal{Z}_{\delta D} \times \mathcal{Q}$, and $\hat{\mathrm{p}}^{(k)}(v) = 0$ if $v \in \partial\mathcal{Z}_{\delta U^c} \times \mathcal{Q}$, we get

$$\hat{\mathrm{p}}^{(k+1)}(v) = \begin{cases} \displaystyle\sum_{v' \in \mathcal{Z}_\delta \times \mathcal{Q}} p_\delta(v \to v')\hat{\mathrm{p}}^{(k)}(v'), & v \in \mathcal{Z}_\delta^\circ \times \mathcal{Q} \\ 1, & v \in \partial\mathcal{Z}_{\delta D} \times \mathcal{Q} \\ 0, & v \in \partial\mathcal{Z}_{\delta U^c} \times \mathcal{Q}, \end{cases}$$

which can be further simplified since the transitions from $v = (z, q) \in \mathcal{Z}_\delta^\circ \times \mathcal{Q}$ are confined to $\mathcal{N}_\delta(z) \cup \{z\}$.

The proposed iterative algorithm to compute $\hat{\mathrm{p}}^{(k_f)}$ determines all the $k_f + 1$ maps $\hat{\mathrm{p}}^{(k)}$, $k = 0, 1, \dots, k_f$. Despite the computation effort, this has the advantage that, at any $t \in (0, t_f)$, an estimate of the probability of interest over the residual time horizon $[t_f - t, t_f]$ of length t is readily available, and is given by the map $\hat{\mathrm{p}}^{(\lfloor (t_f - t)/\Delta_\delta \rfloor)}$. Alternatively, if one were interested in computing only $\hat{\mathrm{P}}_{s_0}$, he/she could simply propagate forward in time the Markov chain transition probabilities starting from a v_0 closest to s_0.

In any case, the grid size δ should be chosen so as to balance the two conflicting considerations that large δ's may not allow for the simulation of fast moving processes, but for small δ's the running time may be too long. Reachability computations become more intensive as the dimension of the continuous state space grows. This is a well-known problem also for deterministic hybrid systems.

VI. Application to system verification

We consider the problem of progressively driving the temperature of a room within some desired range $\mathcal{W}^\star = (x_-^\star, x_+^\star)$, with $x_-^\star < x_+^\star$, by turning on and off a heater, [1]. Starting from an initial set $\mathcal{W}(0) \supset \mathcal{W}^\star$, the desired set $\mathcal{W}^\star$ should be reached within a certain time $t_f > 0$.

A discrete state $\mathbf{q}$ taking values in $\mathcal{Q} = \{1, 2\}$ represents the two conditions when the heater is either "on" ($\mathbf{q} = 1$) or "off" ($\mathbf{q} = 2$). The average temperature of the room $\mathbf{x}$ evolves according to the following SDEs

$$d\mathbf{x}(t) = \begin{cases} (-\frac{l}{C}(\mathbf{x}(t) - x_a) + \frac{r}{C})dt + \frac{\sigma}{C}d\mathbf{w}(t), & \mathbf{q}(t) = 1, \\ -\frac{l}{C}(\mathbf{x}(t) - x_a)dt + \frac{\sigma}{C}d\mathbf{w}(t), & \mathbf{q}(t) = 2, \end{cases}$$

where l is the average heat loss rate, C is the average thermal capacity of the room, x_a is the ambient temperature (assumed to be constant), r is the rate of heat gain supplied by the heater, and $\mathbf{w}$ is a standard Brownian motion modeling the uncertainty and disturbances affecting the temperature evolution with variance modulated by $\sigma > 0$. The temperature is measured in Fahrenheit degrees and the time in minutes. The parameters are assigned the following values: $x_a = 28$, $l/C = 0.1$, $r/C = 10$, and $\sigma/C = 1$.

The heater is turned on or off according to a threshold-based strategy. Once a switching command is issued, it takes some time for the heater to actually commute. This is modeled by introducing a state-dependent switching rate

$$\lambda_{12}(x) = \frac{\bar{\lambda}}{1 + e^{-100(\frac{x}{x_{\text{high}}} - 0.9)}} \quad \lambda_{21}(x) = \frac{\bar{\lambda}}{1 + e^{100(\frac{x}{x_{\text{low}}} - 1.1)}}$$

where value $\bar{\lambda} > 0$ is reached as soon as x gets higher than x_{high} or smaller than x_{low}, with $x_-^\star < x_{\text{low}} < x_{\text{high}} < x_+^\star$. We set $\bar{\lambda} = 10$ so that the minimum average commutation time is 0.1 minutes.

Efficacy of the threshold-based strategy can be verified by the described methodology for reachability analysis. To this purpose, we introduce a time-varying set $\mathcal{W}(t)$ shrinking from $\mathcal{W}(0)$ to $\mathcal{W}^\star$ during the time interval $T = [0, t_f]$, and estimate the probability of $\mathbf{x}(t)$ exiting $\mathcal{W}(t)$ during T

$$\mathrm{P}_{s_0}(\mathcal{W}^c(\cdot)) := P_{s_0}\big\{\mathbf{x}(t) \in \mathcal{W}^c(t) \text{ for some } t \in T\big\} \quad (14)$$

by a Markov chain approximation to the solution $\{\mathbf{s}(t) = (\mathbf{x}(t), \mathbf{q}(t)),\, t \geq 0\}$ of (4), with $\mathbf{x}(t)$ defined on $\mathcal{W}(0)$ and absorption on the boundary $\partial\mathcal{W}(0)$.

In this case $U \setminus D := \mathcal{W}(0)$ and $D := \mathcal{W}^c(0)$, $\mathcal{Z}_\delta = \mathcal{W}(0) \cap \mathbb{Z}_\delta$, and $\partial\mathcal{Z}_\delta$ is composed of those points in $\mathcal{Z}_\delta$ with at least one neighbor inside $\mathcal{W}^c(0)$.

Let $\mathcal{Z}_{\delta,t} = \mathcal{W}(t) \cap \mathbb{Z}_\delta$. Denote by $\mathcal{Z}_{\delta,t}^\circ$ the interior of $\mathcal{Z}_{\delta,t}$, i.e., the set of all those points in $\mathcal{Z}_{\delta,t}$ which have all their neighbors in $\mathcal{Z}_{\delta,t}$. Clearly, $\mathcal{Z}_{\delta,0} = \mathcal{Z}_\delta$ and $\partial\mathcal{Z}_\delta = \mathcal{Z}_\delta \setminus \mathcal{Z}_{\delta,0}^\circ$. If we set $\partial\mathcal{Z}_{\delta,k} := \mathcal{Z}_\delta \setminus \mathcal{Z}_{\delta,k}^\circ$, then, an estimate of the probability (14) is provided by

$$\hat{\mathrm{P}}_{s_0}(\mathcal{W}^c(\cdot)) := P_\delta\big\{\mathbf{z}_k \in \partial\mathcal{Z}_{\delta,k} \text{ for some } k \in [0, k_f]\big\}.$$

This expression is different from that in the time-invariant case. Yet, following the same reasoning, we can define the probabilistic map $\hat{\mathrm{p}}^{(k)} : \mathcal{Z}_\delta \times \mathcal{Q} \to [0,1]$ as $\hat{\mathrm{p}}^{(k)}(v) :=$

$P_\delta\{\mathbf{z}_h \in \partial\mathcal{Z}_{\delta,h}$ for some $h \in [k_f - k, k_f] \mid \mathbf{v}_{k_f-k} = v\}$, and derive the recursion

$$\hat{\mathrm{p}}^{(k+1)}(v) = \begin{cases} \sum_{v' \in \mathcal{Z}_\delta \times \mathcal{Q}} p_\delta(v \to v')\hat{\mathrm{p}}^{(k)}(v'), & v \in \mathcal{Z}^\circ_{\delta,k_f-k-1} \times \mathcal{Q} \\ 1, & v \in \partial\mathcal{Z}_{\delta,k_f-k-1} \times \mathcal{Q}. \end{cases}$$

$\hat{\mathrm{p}}^{(k_f)}(v_0) = \hat{\mathrm{P}}_{s_0}(\mathcal{W}^c(\cdot))$ can be computed by iterating this equation from the initialization (13) with $\partial\mathcal{Z}_{\delta D} = \partial\mathcal{Z}_\delta$.

With reference to the case when $\mathcal{W}^\star = (66, 76)$, $T = [0, 120]$, $\mathcal{W}(0) = (20, 80)$, and $\mathcal{W}(t) = (20 + (66 - 20)\frac{t}{t_f}, 80 + (76 - 80)\frac{t}{t_f})$, $t \in T$, we study the effectiveness of the threshold-base policy when (a) $x_{\text{low}} = 69$ and $x_{\text{high}} = 73$, and (b) $x_{\text{low}} = 67$ and $x_{\text{high}} = 75$. We set $\delta = (\rho \max_{x \in \mathcal{W}(0)}\{a(x, 1), a(x, 2)\})^{-1}$ with $a(x, 1) = -\frac{l}{C}(x - x_a) + \frac{r}{C}$ and $a(x, 2) = -\frac{l}{C}(x - x_a)$, and $\rho = (C/\sigma)^2$.

The results are reported in Figure 1. The three dimensional surface in each plot represents the probability that the temperature is progressively conveyed from $[20, 80]$ to the desired range $(66, 76)$ along the time horizon $[t, 120]$, $t \in [0, 120]$, as a function of the temperature value x at time t for $\mathbf{q}(t) = 2$. This surface is easily obtained from $\hat{\mathrm{p}}^{(k)}$, $k = 0, 1, \ldots, k_f$.

Not surprisingly, these plots show that the probability that the temperature is progressively conveyed to $(66, 76)$ is smaller in the case $x_{\text{low}} = 67$ and $x_{\text{high}} = 75$ (reported on the bottom of Figure 1), and this is because the heater is switched on/off when the temperature is closer to the boundaries.

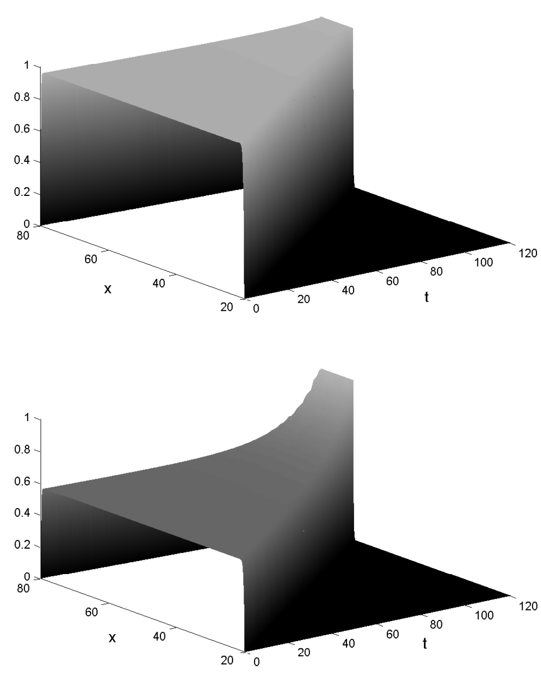

Fig. 1. Probability that the temperature is conveyed to the desired range $(66, 76)$ during $[t, 120]$ starting at time t with the heater off, as a function of the temperature value x at time t, for $t \in [0, 120]$ (on the top: $x_{\text{low}} = 69$ and $x_{\text{high}} = 73$; on the bottom: $x_{\text{low}} = 67$ and $x_{\text{high}} = 75$).

VII. Conclusions

We studied the reachability problem for stochastic hybrid system generating switching diffusions, and presented a numerical approximation scheme to estimate the probability that the switching diffusion will enter a certain subset of the state space within a finite horizon. Application to regulation problems is illustrated by a simple example.

Extensions to the infinite horizon case and to the case when the initial condition is uncertain are easy to obtain. Further work is instead needed to address more general stochastic hybrid systems, the main issue being that of coping with forced transitions due to boundary hitting.

References

[1] S. Amin, A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Reachability analysis for controlled discrete time stochastic hybrid systems. In J. Hespanha and A. Tiwari, eds., *Hybrid Systems: Computation and Control*, *LNCS series*, 3927:49–63, Springer 2006.

[2] K. Amonlirdviman, N. A. Khare, D. R. P. Tree, W.-S. Chen, J. D. Axelrod, and C. J. Tomlin. Mathematical modeling of planar cell polarity to understand domineering nonautonomy. *Science*, 307(5708):423–426, Jan. 2005.

[3] A. Balluchi, L. Benvenuti, M. D. D. Benedetto, G. M. Miconi, U. Pozzi, T. Villa, H. Wong-Toi, and A. L. Sangiovanni-Vincentelli. Maximal safe set computation for idle speed control of an automotive engine. In N. Lynch and B. H. Krogh, eds., *Hybrid Systems: Computation and Control*, *LNCS series*, 1790:32–44, Springer, 2000.

[4] M. L. Bujorianu and J. Lygeros. Reachability questions in piecewise deterministic Markov processes. In O. Maler and A. Pnueli, eds., *Hybrid Systems: Computation and Control*, *LNCS series* 2623:126–140, Springer, 2003.

[5] M.L. Bujorianu. Extended stochastic hybrid systems and their reachability problem. In R. Alur and G. Pappas, eds., *Hybrid Systems: Computation and Control*, *LNCS series*, 2993:234–249, Springer, 2004.

[6] S. Prajna, A. Jadbabaie, and G. J. Pappas. Stochastic safety verification using barrier certificates. In *Proceedings of the IEEE Conference on Decision and Control*, 2004.

[7] M. Egerstedt. Behavior based robotics using hybrid automata. In N. Lynch and B. H. Krogh, eds., *Hybrid Systems: Computation and Control*, *LNCS series*, 1790:103–116. Springer, 2000.

[8] M.K. Ghosh, A. Arapostathis, and S.I. Marcus. Optimal control of switching diffusions with application to flexible manufacturing systems. *SIAM J. Control Optim.*, 31:1183–1204, 1993.

[9] J. Hespanha. Stochastic hybrid modeling of on-off TCP flows. In C.G. Cassandras and J. Lygeros, eds., *Stochastic Hybrid Systems*, Control Engineering Series, 24:191–220, Taylor & Francis Group/CRC Press, 2006.

[10] J. Hu and M. Prandini. Aircraft conflict detection: A method for computing the probability of conflict based on Markov chain approximation. In *European Control Conference*, Cambridge, UK, Sept. 2003.

[11] J. Krystul and A. Bagchi. Approximation of first passage times of switching diffusion. In *Intern. Symposium on Mathematical Theory of Networks and Systems*, Leuven, Belgium, July 2004.

[12] H.J. Kushner and P.G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, 2001.

[13] M. Prandini and J. Hu. A stochastic approximation method for reachability computations. In H.A.P. Blom and J. Lygeros, eds., *Stochastic hybrid systems: theory and safety applications*, *LNCIS series*, 337:107–139, Springer, 2006.

[14] M. Prandini and J. Hu. Stochastic reachability: Theory and numerical approximation. In C.G. Cassandras and J. Lygeros, eds., *Stochastic Hybrid Systems*, Control Engineering Series, 24:107-138, Taylor & Francis Group/CRC Press, 2006.

[15] C. Tomlin, G.J. Pappas, and S. Sastry. Conflict resolution for air traffic management: A study in multi-agent hybrid systems. *IEEE Trans. on Automatic Control*, 43:509–521, 1998.

[16] P. Varaiya. Smart cars on smart roads: problems of control. *IEEE Trans. on Automatic Control*, 38(2):195–207, 1993.