Available online at www.sciencedirect.com

ScienceDirect

IFAC Papers Online CONFERENCE PAPER ARCHIVE

IFAC-PapersOnLine 48-27 (2015) 062–067

Safety Analysis of Stochastic Dynamical Systems

Christoffer Sloth * **Rafael Wisniewski** **

* *Section for Automation and Control, Aalborg University, Denmark (e-mail: ces@es.aau.dk).*

** *Section for Automation and Control, Aalborg University, Denmark (e-mail: raf@es.aau.dk)*

Abstract: This paper presents a method for verifying the safety of a stochastic system. In particular, we show how to compute the largest set of initial conditions such that a given stochastic system is safe with probability p.
To compute the set of initial conditions we rely on the moment method that via Haviland's theorem allows an infinite dimensional optimization problem on measures to be formulated as a polynomial optimization problem. Subsequently, the moment sequence is truncated (relaxed) to obtain a finite dimensional polynomial optimization problem. Finally, we provide an illustrative example that shows how the p-safe initial set is computed numerically.

Keywords: Safety analysis; Formal verification; Stochastic systems; Convex optimisation.

1. INTRODUCTION

A system is said to be safe if it does not violate any system constraints; thus, safety plays an important role in the evaluation of system performance in most applications. In the stochastic settings, a system is said to be p-safe if it does not violate system constraints with a probability of at least p. In this paper, we address methods for safety verification, i.e., methods demonstrating that a system will not violate system constraints with a probability of at least p.

Methods for safety verification of dynamical systems include reachability methods Mitchell et al. (2005) and the barrier certificate method Prajna et al. (2007). Due to its simple computation, the barrier certificate method has attracted much attention.The barrier certificate method has been used for verification Prajna and Rantzer (2007), for estimation of operating envelope Wisniewski et al. (2013), for safe control Wieland and Allgöwer (2007), and has been combined with control barrier functions to include constraints Ames et al. (2014); Romdlony and Jayawardhana (2014).

Work on stochastic safety verification was presented in Prajna et al. (2007) where a super martingale was used as a stochastic barrier certificate, and easy computational conditions were derived based on Doob's martingale inequality Øksendal (2000). Also Bujorianu (2012) thoroughly covers the subject of reachability of hybrid systems. It introduces a formal definition of stochastic reachability. The expected occupation- and hitting measures, also employed in this paper are introduced for the purpose of reachability.

The aim of this work is to compute the largest set of initial conditions from which a system is safe. This is similar to Wisniewski et al. (2013); however, in Wisniewski et al. (2013) the largest set is not necessarily obtained. The methodology exploited in this work is inspired by Korda et al. (2013); Henrion and Korda (2014), which compute regions of attraction for deterministic systems via solving convex programming problems. The problems are infinite dimensional; thus, finite dimensional truncations are used for the actual computations via the hierarchy of relaxations of Lasserre (2001). Note that for safety verification, it is imperative to obtain inner approximations of the set of safe initial conditions.

In contrast to Korda et al. (2013), we consider stochastic systems and study safety, and find the largest set of initial conditions from which a stochastic system is p-safe. Consequently, we do not strive to find an explicit barrier certificate as in Prajna et al. (2007). The solution to the safety verification can be found by solving a polynomial optimization problem using for example GloptiPoly Henrion et al. (2009).

The outline of the paper is as follows. Section 2 presents the safety problem that is addressed in the paper, and Section 3 provides a description of all realizations of a stochastic process initialized in some set; this is the main ingredient of the subsequent formulation of the safety problem. Section 4 derives the optimization problem that is solved to obtain the p-safe initial conditions. Finally, an illustrative example is given in Section 5 and conclusions are provided in Section 6.

2. PROBLEM FORMULATION

This section presents the considered safety problem for stochastic dynamical systems. We search for the largest set of initial conditions for which the system is safe with probability at least p.

The considered system is defined on the probability space $(\Omega, \mathcal{F}, P)$. As usual, P denotes a probability measure on

10.1016/j.ifacol.2015.11.153

the measurable space $(\Omega, \mathcal{F})$ with $\mathcal{F}$ a σ-algebra on a non-empty set Ω. Specifically, we regard a (time-homogeneous) Itô diffusion, i.e., an $\mathbb{R}^n$-valued stochastic process X_t satisfying the stochastic differential equation

$$dX_t = b(X_t)dt + \sigma(X_t)dW_t, \tag{1}$$

where W_t is an m-dimensional Brownian motion, and $b : \mathbb{R}^n \to \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ are measurable functions that satisfy conditions for existence and uniqueness of t-continuous solution X_t, see Theorem 5.2.1 in Øksendal (2000). The notation X_t^x is used to indicate the solution that starts at x, that is, $P(X_0^x = x) = 1$.

The probability $P(X_t^x \in A)$ of the process X_t^x reaching a set $A \subset \mathbb{R}^n$ at a given time t is imperative for the definition of safety. We denote by $P_t(x, A)$ the transition probabilities for the (strong Markov) process X_t, i.e., $P_t(x, A) \equiv P^x(X_t \in A)$, where P^x is the law of X_t^x, that is $P^x(X_t \in A) \equiv P(X_t^x \in A)$.

A system is given by the tuple $\Gamma = (X_t, P^x, U, U_0, U_\mathrm{u})$, where (X_t, P^x) is a strong Markov process, and $U \subseteq \mathbb{R}^n$, $U_0 \subseteq U$, $U_\mathrm{u} \subset U$ are subsets (characterised later in the text). We say that a system Γ is p-safe if the probability of the realisations initialised in the set U_0 that are contained in the set U and that reach the unsafe set U_u is less than $1 - p$. The notion of p-safety of a system Γ is formalised in Definition 1. To this end, we employ the concept of a stopping time. For a subset $V \subseteq \mathbb{R}^n$, we define $\tau_{V,T}^x$ as the first time X_t^x hits V before reaching time T

$$\tau_V^x \equiv \tau_{V,T}^x \equiv \inf(\{t \in \mathbb{R}_+ |\ X_t^x \in V\} \cup \{T\}).$$

The function $\tau_{V,T}^x$ is indeed a stopping time with respect to the filtration generated by X_t^x, for details see Example 9.17 and Lemma 9.18 in Klenke (2008). We consider an event D^x - a subset of Ω containing all the realisations of X_t^x that hit U_u before hitting the complement U^c of U (in $\mathbb{R}^n$)

$$D^x \equiv \{\omega \in \Omega |\ \tau_{U^\mathrm{c}}^x(\omega) - \tau_{U_\mathrm{u}}^x(\omega) > 0\}.$$

It follows that $Q(x) \equiv P(D^x)$ is the probability that X_t^x hits U_u before leaving U. For an initial probability measure with support in U_0, the probability that the process X_t^x with $x \in U_0$ hits U_u before it leaves U is

$$\int_U Q(x)\mu_0(dx).$$

In summary, we formalise the definition of p-safety as follows.

Definition 1. Let $\Gamma = (X_t, P^x, U, U_0, U_\mathrm{u})$ be a system.

- We say that Γ is p-safe with respect to a measure μ_0 on U concentrated on U_0 ($\mu_0(U_0) = 1$) if

$$\int_U Q(x)\mu_0(dx) < 1 - p \tag{2}$$

- We say that Γ is strong p-safe if (2) holds for all probability measures μ_0 on U concentrated on U_0.

Proposition 1. System Γ is strong p-safe if and only if for all $x \in U_0$, $Q(x) < 1 - p$.

Proof 1. If Γ is strong p-safe then for any $x \in U_0$, $\int_U Q(y)d\delta_x(y) < 1 - p$, where δ_x is Dirac measure at x; hence, $Q(x) < 1 - p$.

If $Q(x) < 1 - p\ \forall x \in U_0$ then $\int_U Q(x)d\mu_0(x) < \int_U (1 - p)d\mu_0(x) = (1 - p)$.

Proposition 1 indicates that the system Γ for which the subset $U_0 \subset \mathbb{R}^n$ contains a single point x with $Q(x) \geq 1-p$ is not strong p-safe. This does not align with our intuition of p-safety. That is, any set of 0-volume shall be discarded form safety analysis. Consequently, we provide below a modified definition of p-safety.

Definition 2. (p-safety). We say that the system $\Gamma = (X_t, P^x, U, U_0, U_\mathrm{u})$ is p-safe if for

$$\lambda(\{x \in U_0 |\ Q(x) \geq 1 - p\}) = 0,$$

where λ is the Lebesgue measure on $\mathbb{R}^n$.

In this work, the set of initial conditions U_0 is not fixed, on the contrary, we strive to find the largest set of initial conditions for which the system is p-safe.

Problem 1. (p-safety problem). Given an Itô diffusion defined by (1), a number $p \in [0, 1]$, and two subsets U and U_u. Find the set $U_0 \subseteq U$ of largest volume such that $\Gamma = (X_t, P^x, U, U_0, U_\mathrm{u})$ is p-safe.

A solution to the p-safety problem is provided in Section 4.

3. EVOLUTION OF REALIZATIONS

The purpose of this section is to describe the evolution of all realizations of an Itô diffusion process initialized in a given set. The presentation is motivated by Øksendal (2000) and Section 9.2.2 in Lasserre (2010).

We follow Definition 7.3.1 in Øksendal (2000) and define an infinitesimal generator $\mathcal{A}$ of a time-homogeneous Markov process X_t as

$$\mathcal{A}f(x) = \lim_{t\downarrow 0} \frac{E^x[f(X_t)] - f(x)}{t},$$

where E^x is the expected value computed with respect to P^x. If the process X_t is a time-homogeneous Itô diffusion given by (1) then by Theorem 7.3.3 in Øksendal (2000), the infinitesimal generator $\mathcal{A}$ of X_t is defined by

$$\mathcal{A}f(x) = \sum_i b_i \frac{\partial f}{\partial x_i} + \frac{1}{2}\sum_{i,j} (\sigma\sigma^T)_{ij}(x) \frac{\partial^2 f}{\partial x_i \partial x_j}$$

for any $f \in C_0^2(\mathbb{R}^n) \subset \mathcal{D}_\mathcal{A}$, where $C_0^2(\mathbb{R}^n)$ denotes the set of twice continuously differentiable functions with compact support, and $\mathcal{D}_\mathcal{A}$ denotes the set of functions $f : \mathbb{R}^n \to \mathbb{R}$ such that the limit exists for all $x \in \mathbb{R}^n$. It follows that for $f \in C_0^2(\mathbb{R}^n)$,

$$f(X_t) - \int_0^t \mathcal{A}f(X_s)ds$$

is an $\mathcal{F}_t$-martingale ($\mathcal{F}_t$ is the filtration generated by X_t), Theorem 8.3.1 in Øksendal (2000). To study the hitting time $\tau \equiv \tau_V^x$, we augment $\mathcal{A}$ with a time component. To this end, we follow Helmes et al. (2001) and introduce the time-space generator

$$\hat{\mathcal{A}}(\gamma f)(t, x) = \gamma(t)\mathcal{A}f(x) + \dot{\gamma}(t)f(x)$$

for $f \in C_0^2(\mathbb{R}^n)$ and $\gamma \in C_0^1(\mathbb{R}_+)$. Consequently,

$$\gamma(t)f(X_t) - \int_0^t \hat{\mathcal{A}}(\gamma f)(s, X_s)ds \tag{3}$$

is $\mathcal{F}_t$-martingale. By the optional stopping Theorem 10.15 in Klenke (2008),

$$E^x\left[\gamma(\tau)f(X_\tau) - \int_0^\tau \hat{\mathcal{A}}(\gamma f)(s, X_s)ds \middle| \mathcal{F}_0\right] = \gamma(0)f(x),$$