2018 IEEE Conference on Decision and Control (CDC)
Miami Beach, FL, USA, Dec. 17-19, 2018

Deciding Stability of a Switched System Without Identifying It

Ayca Balkan[a] and Raphaël M. Jungers[b] and Joris Kenanian[a] and Paulo Tabuada[a]

***Abstract*— We address the problem of deciding stability of a "black-box" dynamical system (i.e., a system whose model is not known) from a set of observations. The only assumption we make on the black-box system is that it can be described by a switched linear system.**

We show that, for a given (randomly generated) set of observations, one can give a stability guarantee, for some level of confidence, with a trade-off between the quality of the guarantee and the level of confidence. We provide an explicit way of computing the best stability guarantee, as a function of both the number of observations and the required level of confidence. Our results rely on geometrical analysis and combine chance-constrained optimization theory with stability analysis techniques for switched systems[1].

I. INTRODUCTION

The problem of determining the stability of a given dynamical system has captured the control theory community's attention for decades. The most common approach to this problem is to start by assuming a model for the dynamical system to be analyzed. Yet, for industrial scale systems, even this first step is hard to establish. Cyber-physical systems are a good example to illustrate why this is the case. These systems are characterized by intricate and complex interactions of a large number of heterogeneous components. For example, while a physical component is modeled by a differential equation, a computational component might be modeled by a difference equation, a hybrid automaton or even a lookup table. Therefore, for control engineers, it is common practice to rely on simulations instead of closed-loop models to gain confidence in a given design. For the rest of the paper, we consider *black-box* systems, i.e., systems where we do not have access to the model, f. We can still indirectly learn information about f by observing pairs of points (x_k, y_k) as defined in (1). In this paper, we look at stability analysis with this reality in mind and ask the question: *Can we give formal guarantees on the stability of a system based on the information obtained via its simulation?*

Formally, we consider a dynamical system as in:

$$x_{k+1} = f(k, x_k), \tag{1}$$

where, $x_k \in X$ is the state and $k \in \mathbb{N}$ is the time index. We consider N pairs, (x_1, y_1), (x_2, y_2), …, (x_N, y_N) belonging to the behavior of the system (1), (i.e., $y_k = f(k, x_k)$ for some k), and try to assess the stability of the system (1). We assume that our observations are not correlated, that is, they are made for different runs of the system. As the system is homogeneous, the time k at which the observation is done is not relevant for our problem of stability analysis.

A straightforward attempt to solve this problem would be to first identify the dynamics, i.e., the function f, and then apply existing techniques from the model-based stability analysis literature. However, unless f is a linear function, there are two main reasons behind our quest to directly work on system behaviors and bypass the identification phase: 1) Even when the function f is known, in general, stability analysis is a very difficult problem; 2) Identification can potentially introduce approximation errors, and can be algorithmically hard as well. Note that, both these problems are well-known to exist in the particular case of switched systems [3], [15]. A fortiori, the combination of these two steps in an efficient and robust way seems suboptimal, if only feasible.

In recent years, an increasing number of researchers started addressing various verification and design problems in control of black-box systems [1], [2], [9]–[11]. In particular, the initial idea behind this paper was influenced by the recent efforts in [13], [22], and [4] on using simulation traces to find Lyapunov functions for systems with known dynamics. In these works, the main idea is that if one can construct a Lyapunov function candidate decreasing along several finite trajectories starting from different initial conditions, it should also decrease along every other trajectory. Then, once a Lyapunov function candidate is constructed, this intuition is put to test by verifying the candidate function either via off-the-shelf tools as in [22] and [13], or via sampling-based techniques as in [4]. This also relates to the notion of almost-Lyapunov functions introduced in [16], which presents a relaxed notion of stability proved via Lyapunov functions decreasing everywhere except on a small set. Note that, since we do not have access to the dynamics, these approaches cannot be directly applied to black-box systems. However, these ideas raise the following problem that we address in this paper: By observing that a candidate Lyapunov function decreases on a large number of observations, we empirically build a certain confidence that such candidate Lyapunov function is a bona-fide Lyapunov function. *Can we translate this into a confidence on the stability of the system, that is, that* all *points in the state space converge to zero asymptotically?*

As we illustrate next, this translation is nontrivial even in the case of linear systems. In particular, even though a

[a]A. Balkan, J.Kenanian, and P. Tabuada are with the ECE Dept. at University of California, Los Angeles (UCLA). Their work was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, by NSF Grants 1645824, 1521617, and by the NSF Expeditions in Computer Augmented Program Engineering.

[b]R. M. Jungers is with the ICTEAM Institute at UCLouvain, Belgium. He is a Fulbright Fellow and a F.N.R.S. Fellow. His work is supported by the Innoviris foundation and the Walloon Region.

[1]We present here part of a research project whose complete set of results is available in full version in the Technical Report [14]. This work has never been presented in a conference. All the proofs can be found in our Technical Report [14].

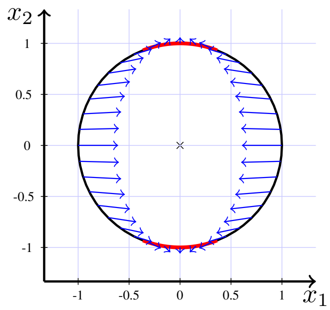

Fig. 1. A simple dynamics and the level set of an "almost Lyapunov function". Even though this function decreases at almost all points in its level set, almost all trajectories diverge to infinity.

candidate Lyapunov function is decreasing almost everywhere on its level sets, all state trajectories might be converging to infinity asymptotically. Indeed, consider the simple following 2-dimensional system:

$$x^+ = \begin{bmatrix} 0.14 & 0 \\ 0 & 1.35 \end{bmatrix} x,$$

which admits the Lyapunov function candidate $V(x) = x^T x$ on the unit circle except on the two red areas shown in Fig. 1. Moreover, the size of this "violating set" can be made arbitrarily small by changing the magnitude of the unstable eigenvalue. Nevertheless, the only trajectories that do not diverge to infinity are those starting on the stable eigenspace (the abscissa axis) that has zero measure.

In this paper, we start our quest to infer stability by constraining ourselves to the study of unknown switched linear systems, where we assume to only know the dimension of the system, n, and the number of modes, m. Note that, identifying and deciding stability of arbitrary switched linear systems is NP-hard [12]. The stability of switched systems is closely related to the *joint spectral radius* (JSR) of the matrices modeling the dynamics in each mode. Deciding stability amounts to deciding whether the JSR is less than one [12]. We present an algorithm to bound the JSR of an unknown switched linear system from a finite number N of observations. This algorithm partially relies on tools from the random convex optimization literature (also known as *chance-constrained optimization,* see [6], [7], [17]), and provides an upper bound on the JSR with a user-defined confidence level. As N increases, this bound gets tighter. Moreover, with a closed form expression, we characterize the exact trade-off between the tightness of this bound and the number of observations made. In order to assess the quality of our upper bound, we also provide a deterministic lower bound. Finally, we provide an asymptotic guarantee on the gap between the upper and the lower bound, for large N.

The organization of the paper is as follows: In Section II, we introduce notations and provide the necessary background in stability of switched systems. In Section III, we present a deterministic lower bound for the JSR. Section IV presents the main contribution of our work, where we provide a probabilistic stability guarantee for a given switched system, based on finite observations. We experimentally demonstrate the performance of the presented techniques in Section V and conclude in Section VI, while hinting at our related future work.

II. Preliminaries

A. Notation

We consider the usual finite normed vector space $(\mathbb{R}^n, \ell_2)$, $n \in \mathbb{N}_{>0}$, with ℓ_2 the classical Euclidean norm. We denote the set of linear functions from $\mathbb{R}^n$ to $\mathbb{R}^n$ by $\mathcal{L}(\mathbb{R}^n)$, and the set of real symmetric matrices of size n by $\mathcal{S}^n$. In particular, the set of positive definite matrices is denoted by $\mathcal{S}^n_{++}$. We write $P \succ 0$ to state that P is positive definite, and $P \succeq 0$ to state that P is positive semi-definite. Given a set $X \subset \mathbb{R}^n$, we denote by $X^{\mathbb{N}}$ the set of all possible sequences $(x_n)_{n\in\mathbb{N}}$, $x_n \in X$. For any $r \in \mathbb{R}_{>0}$, we write $rX := \{rx : x \in X\}$ to denote the scaling of this set. We denote by $\mathbb{B}$ (respectively $\mathbb{S}$) the ball (respectively sphere) of unit radius centered at the origin in $\mathbb{R}^n$. For any $\mathcal{A} \subset \mathbb{S}$, the *sector* of $\mathbb{B}$ defined by $\mathcal{A}$ and denoted by $\mathbb{S}^{\mathcal{A}}$ is the set $\{t\mathcal{A},\ t \in [0,1]\} \subset \mathbb{B}$. We denote the ellipsoid described by the matrix $P \in \mathcal{S}^n_{++}$ as E_P, i.e., $E_P := \{x \in \mathbb{R}^n : x^T P x = 1\}$. Finally, we denote the spherical projector on $\mathbb{S}$ by $\Pi_{\mathbb{S}} := x/\|x\|$.

We consider in this work the classical unsigned and finite uniform spherical measure on $\mathbb{S}$, denoted by σ. It is associated to $\mathcal{B}_{\mathbb{S}}$, the spherical Borelian σ-algebra, and is derived from the Lebesgue measure λ. We have $\mathcal{B}_{\mathbb{S}}$ defined by

$$\mathcal{A} \in \mathcal{B}_{\mathbb{S}} \iff \mathbb{S}^{\mathcal{A}} \in \mathcal{B}_{\mathbb{R}^n}.$$

The spherical measure σ is defined by $\forall\ \mathcal{A} \in \mathcal{B}_{\mathbb{S}}$, $\sigma(\mathcal{A}) = \frac{\lambda(\mathbb{S}^{\mathcal{A}})}{\lambda(\mathbb{B})}$. Notice that $\sigma(\mathbb{S}) = 1$.

For $m \in \mathbb{N}_{>0}$, we denote by M the set $M = \{1, 2, \ldots, m\}$. The set M is provided with the classical σ-algebra associated to finite sets: $\Sigma_M = \wp(M)$, where $\wp(M)$ is the powerset of M. We consider the uniform measure μ_M on (M, Σ_M).

We define $Z = \mathbb{S} \times M$ as the Cartesian product of the unit sphere and M. We denote the product σ-algebra $\mathcal{B}_{\mathbb{S}} \bigotimes \Sigma_M$ generated by $\mathcal{B}_{\mathbb{S}}$ and Σ_M: $\Sigma = \sigma(\pi_{\mathbb{S}}^{-1}(\mathcal{B}_{\mathbb{S}}), \pi_M^{-1}(\Sigma_M))$, where $\pi_{\mathbb{S}} : Z \to \mathbb{S}$ and $\pi_M : Z \to M$ are the standard projections. On this set, we define the product measure $\mu = \sigma^{n-1} \otimes \mu_M$. Note that, μ is a uniform measure on Z and $\mu(Z) = 1$.

B. Stability of Switched Linear Systems

A *switched linear system* with a set of modes $\mathcal{M} = \{A_i, i \in M\}$ is a time-varying discrete-time system of the form:

$$x_{k+1} = f(k, x_k), \tag{2}$$

with $f(k, x_k) = A_{\tau(k)} x_k$ and $\tau \in M^{\mathbb{N}}$ is the *switching sequence*. There are four important properties of switched linear systems that we exploit in this paper.

Property 2.1: Let $\xi(x, k, \tau)$ denote the state of the system (2) at time k starting from the initial condition x and with switching sequence τ. The dynamical system (2) is homogeneous: $\xi(\gamma x, k, \tau) = \gamma \xi(x, k, \tau)$.

Property 2.2: The dynamics given in (2) is convexity-preserving, meaning that for any set of points $X \subset \mathbb{R}^n$, at any time k, we have:

$$f(k, \text{convhull }(X)) \subset \text{convhull }(f(k, X)).$$

The joint spectral radius of the set of matrices $\mathcal{M}$ describes the worst case stability of the system (2) and is defined as follows:

Definition 2.1 (from [21], [12]): Given a finite set of matrices $\mathcal{M} \subset \mathbb{R}^{n\times n}$, its *joint spectral radius* (JSR) is given by

$$\rho(\mathcal{M}) = \lim_{k\to\infty} \max_{i_1,\ldots,i_k} \left\{ ||A_{i_1} \ldots A_{i_k}||^{1/k} : A_{i_j} \in \mathcal{M} \right\}.$$

This quantity is always well defined and finite.

In the following, by *stability* we mean *uniform asymptotic stability.*

Property 2.3 (Corollary 1.1, [12]): Given a finite set of matrices $\mathcal{M}$, the corresponding switched dynamical system is stable if and only if $\rho(\mathcal{M}) < 1$.

Property 2.4 (Proposition 1.3, [12]): Given a finite set of matrices $\mathcal{M}$, and any invertible matrix T, we have $\rho(\mathcal{M}) = \rho(T\mathcal{M}T^{-1})$, i.e., the JSR is invariant under similarity transformations (and is a fortiori a homogeneous function: $\forall\gamma > 0$, $\rho\left(\mathcal{M}/\gamma\right) = \rho(\mathcal{M})/\gamma$.

III. A Deterministic Lower Bound for the JSR

We start by computing a lower bound for ρ which is based on the following theorem from the switched linear systems literature.

Definition 3.1: Consider a finite set of matrices $\mathcal{M} \subset \mathbb{R}^{n\times n}$. A *common quadratic form (CQF)* for a system (2) with set of matrices $\mathcal{M}$, is a positive definite matrix $P \in \mathcal{S}^n_{++}$ such that for some $\gamma \geq 0$,

$$\forall A \in \mathcal{M}, A^T PA \preceq \gamma^2 P. \tag{3}$$

CQFs are useful because they can be computed, when they exist, with semidefinite programming (see [5]), and they constitute a stability guarantee (when $\gamma < 1$, they are Lyapunov functions) for switched systems as we formalize next.

Theorem 3.1: [12, Theorem 2.11] For any finite set of matrices $\mathcal{M} \subset \mathbb{R}^{n\times n}$ such that $\rho(\mathcal{M}) < \frac{1}{\sqrt{n}}$, there exists a Common Quadratic Form (CQF) with $\gamma = 1$ for $\mathcal{M}$, that is, a $P \succ 0$ such that:

$$\forall\ A \in \mathcal{M},\ A^T PA \preceq P. \tag{4}$$

Theorem 3.2: [12, Prop. 2.8] Consider a finite set of matrices $\mathcal{M}$. If there exist a $\gamma \geq 0$ and $P \succ 0$ such that $\forall\ A \in \mathcal{M}, A^T PA \preceq \gamma^2 P$, then $\rho(\mathcal{M}) \leq \gamma$.

Note that the smaller γ is, the tighter is the upper bound we get on $\rho(\mathcal{M})$. In order to properly analyze our setting, where the matrices are unknown, let us reformulate (4) in another form. We can consider the optimal solution γ^* of the following optimization problem:

$$\begin{aligned} \min\nolimits_{\gamma,P} \quad & \gamma \\ \text{s.t.} \quad & (Ax)^T P(Ax) \leq \gamma^2 x^T Px,\ \forall\ A \in \mathcal{M}, \forall x \in \mathbb{R}^n, \\ & P \succ 0. \end{aligned} \tag{5}$$

Thanks to Property 2.1, this problem is equivalent to

$$\begin{aligned} \min\nolimits_{\gamma,P} \quad & \gamma \\ \text{s.t.} \quad & (Ax)^T P(Ax) \leq \gamma^2 x^T Px,\ \forall\ A \in \mathcal{M}, \forall x \in \mathbb{S}, \\ & P \succ 0. \end{aligned} \tag{6}$$

One can now see a formalization of our problem (6): our goal amounts to finding a solution to a convex problem with an infinite number of constraints, while only sampling a finite number of them.

Even though this upper bound is more difficult to obtain in a black-box setting where only a finite number of observations are available, in this section we leverage Theorem 3.1 in order to derive a straightforward lower bound. Indeed, the following theorem shows that the existence of a finite minimum for Program (6), given N arbitrarily drawn pairs $(x_i, j_i) \in Z$, where $i \in \{1, \ldots N\}$, allows to retrieve a lower bound on the JSR of the system. Recall that in our setting, we assume that we only observe pairs of the form (x_i, y_i), but we do not observe the mode applied to the system during this time step, i.e., the values taken by the switching sequence. Modes are supposed to be randomly drawn by the system, according to a uniform law. The user does not have access to this process nor its outcomes. The assumption of the uniform law is discussed with its possible relaxations in [14]. The user's knowledge is limited to the number of modes and the dimension of the system (or an upper bound on this number).

Theorem 3.3: For a given uniform sampling:

$$\omega_N := \{(x_1, j_1), (x_2, j_2), \ldots, (x_N, j_N)\} \subset Z,$$

let $W_{\omega_N} = \{(x_1, y_1), \ldots, (x_N, y_N)\}$ be the corresponding available observations, which can be rewritten as

$$y_i = A_{j_i} x_i\ \forall (x_i, y_i) \in W_{\omega_N}.$$

Also let $\gamma^*(\omega_N)$ be the optimal solution of the following optimization problem:

$$\begin{aligned} \min\nolimits_{\gamma,P} \quad & \gamma \\ \text{s.t.} \quad & y_i^T P y_i \leq \gamma^2 x_i^T P x_i, \forall (x_i, y_i) \in W_{\omega_N} \\ & P \succ 0. \end{aligned} \tag{7}$$

Then, we have $\rho(\mathcal{M}) \geq \frac{\gamma^*(\omega_N)}{\sqrt{n}}$.

Note that, (7) can be efficiently solved by semidefinite programming and bisection on the variable γ (see [5]).

Proof: Let $\varepsilon > 0$. By definition of γ^*, there exists no matrix $P \in \mathcal{S}^n_{++}$ such that:

$$(Ax)^T P(Ax) \leq (\gamma^*(\omega_N) - \varepsilon)^2 x^T Px, \forall x \in \mathbb{R}^n, \forall A \in \mathcal{M}.$$

By Property 2.4, this means that there exists no CQF for the scaled set of matrices $\frac{\mathcal{M}}{(\gamma^*(\omega_N)-\varepsilon)}$. Then, using Theorem 3.1, we conclude:

$$\frac{\rho(\mathcal{M})}{\gamma^*(\omega_N)} \geq \frac{1}{\sqrt{n}}.$$

■

IV. A Probabilistic Stability Guarantee

In this section, we analyze a somehow converse result to Theorem 3.3. We consider a finite uniform random sampling of constraints ω_N, and $\gamma^*(\omega_N)$ the optimal solution of Problem 7 as defined in Section 3. In our proof, for technical reasons that will become clear later, we look at the slightly more involved optimization problem, denoted by Opt(ω_N):

$$\begin{aligned} \min_P \quad & \lambda_{\max}(P) \\ \text{s.t.} \quad & (A_j x)^T P (A_j x) \leq \gamma^*(\omega_N)^2 x^T P x, \\ & \qquad\qquad \forall (x,j) \in \omega_N \subset Z, \\ & P \succeq I. \end{aligned} \tag{8}$$

We denote its optimal solution by $P(\omega_N)$, and drop the explicit dependence of P on ω_N when it is clear from the context. Feasibility and other considerations on this optimization problem are examined in [14]. From the relationship between this problem and Problem 6, we show how to compute an upper bound on ρ, with a user-defined confidence $\beta \in [0,1)$. We do this by constructing a CQF which is valid with probability at least β.

We start with a classical result from random convex optimization literature, which we adapt to our problem below:

Theorem 4.1 (adapted from Theorem 3.3, [6]): Consider the optimization problem Opt(ω_N) given in (8), where ω_N is a uniform random sampling of the set Z. Let $d = \frac{n(n+1)}{2}$ be the dimension of the decision variable P of Opt(ω_N) and $N \geq d+1$. Then, for all $\varepsilon \in (0,1]$ the following holds:

$$\mu^N\{\omega_N \in Z^N : \mu(V(\omega_N)) \leq \varepsilon\} \geq 1 - \sum_{j=0}^{d} \binom{N}{j} \varepsilon^j (1-\varepsilon)^{N-j},$$

where μ^N denotes the product probability measure on Z^N, and $V(\omega_N)$ is defined by

$$\begin{aligned} V(\omega_N) = \{&(x,j) \in Z : \\ & (A_j x)^T P(\omega_N)(A_j x) > \gamma^* x^T P(\omega_N) x\}, \end{aligned}$$

i.e., it is the set of constraints of Problem 6 that are violated by the optimal solution of Opt(ω_N).

Theorem 4.1 states that the optimal solution of the sampled problem Opt(ω_N) violates no more than an ε fraction of the constraints in the original optimization problem (6) with probability β, where β goes to 1 as N goes to infinity. This means that, the ellipsoid computed by Opt(ω_N) is "almost invariant" except for a set of measure ε.

We are now able to state the main theorem of this paper. For a given level of confidence β, we prove that the upper bound $\gamma^*(\omega_N)$, which is valid solely on finitely many observations, is in fact a true upper bound, at the price of increasing it by the factor $\frac{1}{\delta(\beta,\omega_N)}$. Moreover, as expected, this factor gets smaller as we increase N and decrease β.

Theorem 4.2: Consider an n-dimensional switched linear system as in (2) and a uniform random sampling $\omega_N \subset Z$, where $N \geq \frac{n(n+1)}{2} + 1$. Let $\gamma^* = \gamma^*(\omega_N)$ be a feasible solution to (7). Then, for any given $\beta \in [0,1)$, we have with probability at least β, $\rho \leq \frac{\gamma^*}{\delta}$, where ρ is the JSR of the system and $\delta(\tilde{\varepsilon}) = \sqrt{1 - I^{-1}(2\tilde{\varepsilon}; \frac{n-1}{2}, \frac{1}{2})}$, with I the regularized incomplete beta function [18, Section 6.6.2]. Moreover, $\lim_{N\to\infty} \delta(\beta, \omega_N) = 1$ with probability 1.

We only sketch the main reasoning of the proof here, because of space constraints. A complete version of the proofs can be found in [14].

Proof: Let us consider a uniform random sampling $\omega_N \subset Z$ as in the statement of the theorem. We first suppose that the optimal solution of (8) is the unit sphere, i.e., $P = I$. By exploiting Property 2.2, we show how one can compute an upper bound on the JSR in this particular case. In the claim below, $\mathbb{S}'$ represents the set of 'bad points', that is, the small set of points that violate the constraints, because we only know of finite sampling of these (recall that Theorem 4.1 precisely gives us a bound ε on the measure of this set $\mathbb{S}'$).

Claim 1 *L*et $\varepsilon \in (0,1]$ and $\gamma \in \mathbb{R}_{>0}$. Consider the set of matrices $\mathcal{M}$ and $A \in \mathcal{M}$ satisfying:

$$(A_j x)^T (A_j x) \leq \gamma x^T x, \quad \forall x \in \mathbb{S} \setminus \mathbb{S}', \forall j \in, M, \tag{9}$$

where $\mathbb{S}' \subset \mathbb{S}$ and $\sigma(\mathbb{S}') \leq \varepsilon$, then we have:

$$\rho(\mathcal{M}) \leq \frac{\gamma}{\alpha(\varepsilon)}$$

for some computable function $\alpha(\varepsilon)$.

See [14, Equation 29] for a closed-form formula for the function $\alpha(\cdot)$. The rationale behind the above claim is that Equation (9) combined with Property 2.2 is not only valid for $(\mathbb{S} \setminus \mathbb{S}')$, but actually for the whole convex hull, that is:

$$A_j \text{ convhull } (\mathbb{S} \setminus \mathbb{S}') \subset \text{convhull } (A_j(\mathbb{S} \setminus \mathbb{S}')) \subset \gamma\mathbb{B}.$$

Thus, if one could find a number α such that convhull $(\mathbb{S} \setminus \mathbb{S}') \subset \mathbb{B}/\alpha$, one would obtain a bound on the JSR. This is actually possible, as soon as one has a bound on the measure of $\mathbb{S}'$, as claimed below:

Claim 2 *F*or any $\varepsilon > 0$, there exists a closed-form expression for the function

$$\alpha(\varepsilon) := \sup_{X \in \mathcal{X}_\varepsilon} \{r : r\mathbb{B} \subset \text{convhull } (\mathbb{S} \setminus X)\}, \tag{10}$$

where $\mathcal{X}_\varepsilon = \{X \subset \mathbb{S} : \sigma(X) \leq \varepsilon\}$. The proof is geometric, and can be found in [14, Proposition 2].

We now know how to compute an upper bound on the JSR when the "almost invariant" ellipsoid, solution to Equation (8) is the sphere $\mathbb{S}$. If, on the contrary, this set is a non-spherical ellipsoid E_P, then thanks to Property 2.4, we can simply perform a change of coordinates mapping this ellipsoid to $\mathbb{S}$ and compute the JSR in the new coordinates system instead. In order to do this, in the next claim, we bound the measure of violating constraints on $\mathbb{S}$ after the change of coordinates, in terms of the measure of the violated constraints on $\mathbb{S}$ in the original coordinates.

Claim 3 *L*et $\gamma \in \mathbb{R}_{>0}$. Consider a set of matrices $A \in \mathcal{M}$, and a matrix $P \succ 0$ satisfying:

$$(A_j x)^T P (A_j x) \leq \gamma^2 x^T P x, \ \forall (x,j) \in Z \setminus V, \tag{11}$$

for some $V \subset Z$ where $\mu(V) \leq \varepsilon$. Then, denoting $L \in \mathbb{R}^{n\times n}$ such that $L^{-T}PL = I$, and $\bar{A}_j = L^{-1}A_jL$, one also has:

$$(\bar{A}_j x)^T(\bar{A}_j x) \leq \gamma^2 x^T x, \ \forall x \in \mathbb{S} \setminus \mathbb{S}', \forall j \in M,$$

where $\mathbb{S}' \subset \mathbb{S}$ is such that $\sigma(\mathbb{S}') \leq m\varepsilon\kappa(P)$, and

$$\kappa(P) = \sqrt{\frac{\det(P)}{\lambda_{\min}(P)^n}}.$$

In the above claim, $\lambda_{\min}(P)$ denotes the minimal eigenvalue of P. The proof of Claim 3 proceeds by elementary geometry of $\mathbb{R}^n$. Again the details can be found in [14]. Now, by taking L equal to the change of coordinate mapping the non-spherical ellipsoid E_P to the sphere, one obtains a new system which satisfies the conditions in Claim 1, together with a bound on the measure of the violating set for this new system.

We now summarize our proof by putting together all the above pieces. For a given level of confidence β, we obtain a solution P to Problem (6), which is valid for almost all points, except for a measure ε. We apply a change of variables to our matrices *without changing the JSR to be computed*, and we now have a set which solves Problem (6) with a solution I (that is, the 'almost' invariant set is the sphere) and the same γ, for all points except a set of measure $m\varepsilon\kappa(P)$. We can thus apply Claim 1 and deduce a bound (with level of confidence β) on the JSR for our set equal to $\frac{\gamma}{\alpha(m\varepsilon\kappa(P))}$, and the proof is done. ∎

V. Experimental Results

We illustrate our technique on a two-dimensional switched system with 4 modes. We fix the confidence level, $\beta = 0.92$, and compute the lower and upper bounds on the JSR for $N := 15 + 15k$, $k \in \{0, \ldots, 23\}$, according to Theorem 3.3 and Theorem 4.2, respectively. We illustrate the average performance of our algorithm over 10 different runs in Fig. 2 and Fig. 3. Fig. 2 shows the evolution of $\delta(\beta, N)$ as N increases. We illustrate that δ converges to 1 as expected. In Fig. 3, we plot the upper bound and lower bound for the JSR of the system computed by Theorem 4.2 and Theorem 3.3, respectively. To demonstrate the performance of our technique, we also provide the JSR approximated by the JSR toolbox [23], which turns out to be 0.7727. Note that, the plot for the upper bound starts from $N = 45$. This is due to the fact for $N = 15$, and $N = 30$, $\delta(\beta, \omega_N) = 0$, hence it is not possible to compute a nontrivial upper bound for these small values of N. As can be seen, the upper bound approaches to a close vicinity of the real JSR with approximately 200 samples. In addition, the gap between the upper and lower bound converges to a multiplicative factor of $\frac{\rho}{\sqrt{n}}$ as expected.

Note that, if we increase the dimension of the switched system, the convergence of δ to 1 will become much slower. We confirmed this via experiments up to dimension $n = 6$. For example, for dimension $n = 4$, it took $N = 5,000$ to $N = 10,000$ points to reach $\delta = 0.9$. We nevertheless observe convergence of the upper bound to $\rho(\mathcal{M})$, and convergence of the lower bound to $\frac{\rho(\mathcal{M})}{\sqrt{n}}$. The gap between these two

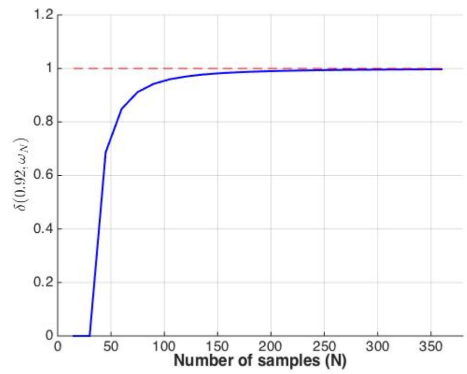

Fig. 2. Evolution of δ with increasing N.

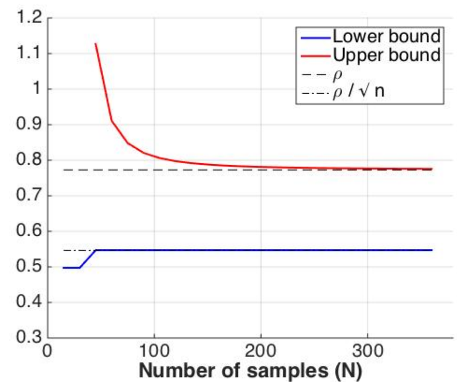

Fig. 3. Evolution of the upper and lower bounds on the JSR with increasing N, for $\beta = 0.92$.

limits is $\frac{\rho}{\sqrt{n}}$ and could be improved by considering a more general class of common Lyapunov functions, such as those that can be described by sum-of-squares polynomials [19]. We leave this for future work.

Finally, we randomly generate $10,000$ test cases with systems of dimension between 2 and 7, number of modes between 2 and 5, and size of samples N between 30 and 800. We take $\beta = 0.92$ and we check if the upper bound computed by our technique is greater than the actual JSR of the system. We get 9873 positive tests, out of $10,000$, which gives us a probability of 0.9873 of the correctness of the upper bound computed. Note that, this probability is significantly above the provided β. This is expected, since our techniques are based on worst-case analysis and thus fairly conservative.

VI. Conclusions and possible extensions

In this paper, we investigated the question of how one can conclude stability of a dynamical system when a model is not available and, instead, we only observe randomly generated state measurements. Our goal is to understand how the observation of well-behaved trajectories *intrinsically* implies stability of a system. It is not surprising that we need some standing assumptions on the system, in order to allow for any sort of nontrivial stability certificate solely from a finite number of observations. Even if we focused on a black box setting, our technique can still be used as a randomized algorithm to evaluate the JSR when a model of the system is known.

The novelty of our contribution is threefold.
First, we apply powerful techniques from chance constrained optimization. The application is not obvious, and relies on geometric properties of linear switched systems. In our view, the obtained guarantee is quite powerful, in view of the hardness of the general problem.
Second, we use as standing assumption that the unknown system can be described by a switching linear system. This assumption covers a wide range of systems of interest, and to our knowledge no such "black-box" result has been available so far on switched systems. We believe that our work opens up a new direction in general, by showing that chance-constrained optimization and the so-called scenario approach can be used for other purposes than to provide a confidence of the measure of possible bad scenarios. Indeed we show that these results can be converted into actual intrinsic *global* stability properties for the complete system. In particular, this approach is new within the field of switched systems: to the best of our knowledge, randomness has never been considered in this aspect in the field. Situations where the switching signal comes from a random process (see e.g. Markov Jump Linear Systems [8], or the Lyapunov Exponent of a switched system [20]) have been studied, but this is a completely different issue from the one addressed here.
Third, we present a hypothesis testing result for stability of complex systems, which is also new to the best of our knowledge.

Let us end this paper by mentioning a few possible generalizations of our work that we find promising.
First, we have restricted our attention to observations of the type (x_t, x_{t+1}), that is, trajectories of length one. It is well known in the switched systems literature that one can refine the computation of the JSR by considering longer trajectories, that is, observations of the type $x_t, \dots, x_{t+L}$ for some L. One could then apply this 'iterated technique' in order to refine the bounds obtained in this paper. See our Technical Report [14] for results in that direction.
Second, a key step in our main theorem consists in making a change of variables, in order to transform the computed ellipsoid into a sphere. This incurs a loss of accuracy of our bound. This loss could be alleviated if one knew in advance the solution-ellipsoid, and could then adapt the sampling accordingly. We would like to develop an adaptive learning algorithm that would allow to modify the sampling strategy online, while collecting observations.
Third, we firmly believe that the methodology applied here to switched systems could be generalized to other, general nonlinear, systems. Indeed, as pointed out above, we only make use here of general and natural geometric properties of switched systems (see in particular Properties 2.1 and 2.2).
Finally, we would like to investigate how these ideas can be translated to probabilistic analysis of computerized systems, as for instance the termination, or convergence, of computer software.

REFERENCES

[1] A. Balkan, P. Tabuada, J. V. Deshmukh, X. Jin, and J. Kapinski. Underminer: A Framework for Automatically Identifying Non-converging Behaviors in Black Box System Models. In *Proceedings of the Thirteenth ACM International Conference on Embedded Software (EMSOFT)*, EMSOFT '16, pages 1–10, 2016.

[2] F. Blanchini, G. Fenu, G. Giordano, and F. A. Pellegrino. Model-Free Plant Tuning. *IEEE Transactions on Automatic Control*, 62(6):2623–2634, 2017.

[3] V. D. Blondel and J. N. Tsitsiklis. Complexity of Stability and Controllability of Elementary Hybrid Systems. *Automatica*, 35(3):479–489, 1999.

[4] R. Bobiti and M. Lazar. A Delta-sampling Verification Theorem for Discrete-time, Possibly Discontinuous Systems. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, HSCC '15, pages 140–148. ACM, 2015.

[5] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[6] G. C. Calafiore. Random Convex Programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.

[7] M. C. Campi and S. Garatti. The Exact Feasibility of Randomized Solutions of Uncertain Convex Programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

[8] O. Costa, M. Fragoso, and R. Marques. *Discrete-time Markov Jump Linear Systems*. Springer Science & Business Media, 2006.

[9] P. S. Duggirala, S. Mitra, and M. Viswanathan. Verification of Annotated Models from Executions. In *Proceedings of the Eleventh ACM International Conference on Embedded Software*, EMSOFT '13, pages 26:1–26:10, 2013.

[10] S. Haesaert, A. Abate, and P. Van den Hof. Data-driven and Model-based Verification: A Bayesian Identification Approach. In *Proceedings of the 54th IEEE Conference on Decision and Control (CDC)*, pages 6830–6835. IEEE, 2015.

[11] Z. Huang and S. Mitra. Proofs from Simulations and Modular Annotations. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, HSCC '14, pages 183–192. ACM, 2014.

[12] R. M. Jungers. The Joint Spectral Radius, Theory and Applications. In *Lecture Notes in Control and Information Sciences*, volume 385. Springer-Verlag, 2009.

[13] J. Kapinski, J. V. Deshmukh, and N. Sankaranarayanan, S.and Arechiga. Simulation-guided Lyapunov Analysis for Hybrid Dynamical Systems. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, HSCC '14, pages 133–142. ACM, 2014.

[14] J. Kenanian, A. Balkan, R. M. Jungers, and P. Tabuada. Data Driven Stability Analysis of Black-box Switched Linear Systems. *arXiv preprint arXiv:1803.03687*, 2018.

[15] F. Lauer. On the Complexity of Switching Linear Regression. *Automatica*, 74:80–83, 2016.

[16] S. Liu, D. Liberzon, and V. Zharnitsky. On Almost Lyapunov Functions for Non-vanishing Vector Fields. In *Proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, pages 5557–5562, 2016.

[17] A. Nemirovski and A. Shapiro. Convex Approximations of Chance Constrained Programs. *SIAM Journal on Optimization*, 17(4):969–996, 2006.

[18] F. W. Olver. *NIST Handbook of Mathematical Functions Hardback and CD-ROM*. Cambridge university press, 2010.

[19] A. Papachristodoulou and S. Prajna. On the Construction of Lyapunov Functions Using the Sum of Squares Decomposition. In *Proceedings of 41st IEEE Conference on Decision and Control*, pages 3482–3487, 2002.

[20] V. Protasov and R. M. Jungers. Lower and Upper Bounds for the Largest Lyapunov Exponent of Matrices. *Linear Algebra and its Applications*, 438(11):4448 – 4468, 2013.

[21] G.-C. Rota and W. Strang. A Note on the Joint Spectral Radius. *Nederl. Akad. Wetensch. Proc.*, (vol. 22):379,Äì381, 1960.

[22] U. Topcu, A. Packard, and P. Seiler. Local Stability Analysis Using Simulations and Sum-of-Squares Programming. *Automatica*, 44(10):2669–2675, 2008.

[23] G. Vankeerberghen, J. Hendrickx, and R. M. Jungers. Jsr: A Toolbox to Compute the Joint Spectral Radius. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control*, HSCC '14, pages 151–156. ACM, 2014.