Formal Methods for Computing Hyperbolic Invariant Sets for Nonlinear Systems

Guillaume O. Berger and Raphaël M. Jungers

***Abstract*— Hyperbolicity is a cornerstone of nonlinear dynamical systems theory. Hyperbolic dynamics are characterized by the presence of expanding and contracting directions for the derivative along the trajectories of the system. Hyperbolic dynamical systems enjoy many interesting properties like structural stability, ergodicity, transitivity, etc. In this paper, we describe a Hybrid Systems framework to compute invariant sets with a hyperbolic structure for a given dynamical system. The method relies on an abstraction (aka. symbolic image or bisimulation) of the state space of the system, and on path-complete 'Lyapunov-like' techniques to compute the expanding and contracting directions for the derivative along the trajectories of the system. The method is illustrated on a numerical example: the Ikeda map for which an invariant set with hyperbolic structure is computed using the framework.**

I. Introduction

Dynamical systems encountered in real-world applications are generally subject to modeling uncertainties and parameter variation. The *robustness* or *structural stability* of a dynamical system is the property that the qualitative behavior of the system will not be affected by a small perturbation of the model or a small change of parameters. A classical example of robust property of dynamical systems are their hyperbolic fixed points (number and location): it is well known from bifurcation theory that a fixed point x can appear/disappear, or become stable/unstable only if the Jacobian matrix $Df_x = f'(x)$ at x has an eigenvalue on the unit circle (discrete-time case) or on the imaginary axis (continuous-time case); in other words, a bifurcation can only occur at non-hyperbolic fixed points.

The concept of hyperbolicity was introduced in the 1960's, by Dmitri Anosov and Stephen Smale, as part of a general effort to study dynamical systems that are structurally stable not only at single fixed points but on more general subsets, e.g., on their whole domain or on invariant sets. Invariant sets with hyperbolic structure are characterized by the presence of expanding and contracting directions for the derivative along the trajectories of the system. Therefore, they generalize the notion of hyperbolic fixed point—whose Jacobian matrix is a linear operator with a stable (contracting) and an unstable (expanding) eigenspace.

Hyperbolicity, which was first developed for flows and diffeomorphisms (i.e., smooth invertible discrete-time systems), has rapidly become a cornerstone of dynamical systems theory and finds applications in many different areas (e.g., chaos, ergodic theory, entropy, structurally stability, etc.). For instance, it can be shown that, under some mild assumptions (e.g., Axiom A, or no-cycle condition, etc.), the structurally stable dynamical systems are precisely the ones that are hyperbolic on some distinguished sets (e.g., limit set, chain-recurrent set, etc.). We refer the reader to [21], [9] for a comprehensive survey of results related to hyperbolic flows and diffeomorphisms. Hyperbolicity has been generalized in several directions (e.g., with partial hyperbolicity, nonuniform hyperbolicity, hyperbolic endomorphisms) allowing one to analyze a broader class of systems while retaining the main features of hyperbolic dynamics [11], [2], [4], [17].

In recent years, hyperbolicity has also been successfully applied in different areas of control theory (e.g., symbolic control, quantized control, etc.). Indeed, the robustness of hyperbolic dynamics to system perturbations makes them particularly suitable for numerical simulation and verification; see, e.g., [5]. Moreover, the existence of expanding and contracting directions for the derivative can be used to define a partition of the state space that is adapted to the system (Markov Partition), or to estimate the entropy of the system [6]. We refer the reader to [15] for a comprehensive introduction to "hyperbolic control theory".

Although the *behavior* of hyperbolic dynamical systems is now well understood, the question of deciding *whether* a dynamical system is hyperbolic or not remains a challenging task, and to the best of the authors' knowledge, only a few results on the formal verification of hyperbolicity with numerical methods are available in the literature. (See also Subsection IV-D for related works.)

In this paper, we draw upon modern optimization and control techniques to propose a novel approach for the systematic verification of hyperbolicity of dynamical systems. Our framework combines ideas from *symbolic control* (aka. *bisimulation* or *abstraction* approach) with algorithmic techniques from path-complete Lyapunov theory [14], and dominance [8], [3], to derive a new set of Linear Matrix Inequalities for the characterization of hyperbolic dynamics. This results in a sound algorithm for the automatic computation of invariant sets with hyperbolic structure for nonlinear dynamical systems. In section V, we show on a numerical example the efficiency of our approach.

The paper is organized as follows: in Section II, we introduce the fundamental concepts related to hyperbolic dynamics. In Section III, we introduce the *quadratic cone field criterion* as a sufficient condition for a dynamical system to have an invariant set with hyperbolic structure. In Section IV, we provide an algorithmic framework for the

G. Berger is a FNRS/FRIA Fellow. R. Jungers is a FNRS Research Associate. He is supported by the Walloon Region and the Innoviris Foundation.
Both are with ICTEAM institute, UCLouvain, Louvain-la-Neuve, Belgium. {guillaume.berger, raphael.jungers}@uclouvain.be.

computation of the quadratic cone field criterion that relies on an abstraction of the system and on LMIs. Finally, in Section V, we illustrate the use of the algorithmic framework on a numerical example.

II. Hyperbolic dynamical systems

In this work, we consider a discrete-time dynamical system

$$x(t+1) = f(x(t)), \quad x \in M,$$

where $M \subseteq \mathbb{R}^d$ and $f : M \to M$ is a continuous map. If D is a subset of M, $f(D)$ denotes its image $\{f(x) : x \in D\}$. A subset $D \subseteq M$ is said to be *invariant* for f if $f(D) = D$. If $f : M \to M$ is bijective and both f and f^{-1} are C^1 functions, then we say that f is a *diffeomorphism*.

Let us now introduce the notion of hyperbolicity. Therefore, we let $\|\cdot\|$ be any vector norm on $\mathbb{R}^d$. The derivative (aka. Jacobian matrix) of f at x is denoted by $Df_x \in \mathbb{R}^{d\times d}$. If f is a diffeomorphism, note that, since $(f \circ f^{-1})(x) = x$, we have that Df_x^{-1} exists and is equal to $(Df_{f^{-1}(x)})^{-1}$. If $E \subseteq \mathbb{R}^d$, then $Df_x(E)$ denotes its image by Df_x.

Definition 1. [21], [10]. Let f be a diffeomorphism, and $\Lambda \subseteq M$ be an invariant set for f. Then, Λ is said to have a *hyperbolic structure* for f (or f is *hyperbolic* on Λ) if (i) for every $x \in M$, there exists a splitting $\mathbb{R}^d = E^u_x \oplus E^s_x$ where E^u_x and E^s_x are linear subspaces; (ii) the splitting is invariant under the action of the derivative: $Df_x(E^u_x) = E^u_{f(x)}$ and $Df_x(E^s_x) = E^s_{f(x)}$; and (iii) there exist $0 < \lambda < 1$ and $C \geq 1$ independent of x such that, for every $n \geq 0$,

- $\|Df^n_x v\| \leq C\lambda^n \|v\|$ for every $v \in E^s_x$;
- $\|Df^{-n}_x v\| \leq C\lambda^n \|v\|$ for every $v \in E^u_x$.

Remark 1. Properties (ii) and (iii) implies that the subspaces E^u_x and E^s_x in Definition 1 are unique and depend continuously on x; see, e.g., Proposition 1.3.7 in [10]. This implies, among other things, that the dimensions of E^u_x and E^s_x are constant on every connected components of Λ. ◁

As mentioned in the introduction, hyperbolic dynamics enjoy many interesting properties in terms of *structural stability* (aka. *robustness* to system perturbations). For instance, it was shown by M. Hirsch and C. Pugh [13] that invariant sets with a hyperbolic structure (for a given diffeomorphism) have the same structural stability properties as hyperbolic fixed points. We refer the reader to [21] for a comprehensive survey of structural stability results related to hyperbolic dynamics.

Remark 2. The notion of hyperbolicity is generally defined for the more general class of dynamical systems on smooth Riemannian manifolds [21], [10]. For the sake of simplicity, we have restricted ourselves to the case of $M \subseteq \mathbb{R}^d$ in this paper. The reader will verify that, by means of atlases and local coordinate systems (see, e.g., [16]), all the results presented in this paper can be generalized to dynamical systems defined on smooth Riemannian manifolds. ◁

Example 1 (The hyperbolic toral automorphism). A classical example of hyperbolic diffeomorphism is the *hyperbolic toral automorphism* (aka. *Arnold's cat map*):

$$f(x) = Ax \mod 1, \quad A = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}, \quad M = \mathbb{R}^2/\mathbb{Z}^2.$$

The eigenvalues of A are equal to $\lambda_\pm \coloneqq (1 \pm \sqrt{5})/2$. At every $x \in M$, the derivative of f^n is given by $Df^n_x = A^n$. The stable subspace E^s_x in Definition 1 is then given by the eigenspace associated to $\lambda_- \approx -0.618$ while the unstable subspace E^u_x is given by the eigenspace associated to $\lambda_+ \approx 1.618$. The diffeomorphism f is thus hyperbolic on its whole domain. ◁

III. Quadratic cone field criterion

In this section, we introduce a sufficient condition for a dynamical system to be hyperbolic on a given invariant set. Connections of this criterion with other concepts from dynamical systems theory, like the *Alekseev cone field criterion* or the notion of *dominance* for continuous-time systems and linear systems, are discussed at the end of this section.

A. Description of the criterion

The criterion relies on the contraction property of a field of quadratic cones defined at every point of the invariant set. Quadratic cones are defined by means of symmetric matrices with fixed inertia. (The *inertia* of a symmetric matrix S, denoted by $\mathrm{In}(S)$, is the triplet (i_-, i_0, i_+) where i_-, i_0 and i_+ are respectively the number of negative, zero, and positive eigenvalues of S.) In the sequel, we let p be a fixed integer in $\{1, \ldots, d-1\}$. We will say that $S \in \mathbb{R}^{d\times d}$ is a *p-matrix* if S is symmetric and has inertia $(p, 0, d-p)$.

Definition 2. A *field of p-matrices* on $\Lambda \subseteq M$ is a function Φ that associates a p-matrix Φ_x to each $x \in \Lambda$. Moreover, we will assume that the field Φ is *bounded*, i.e., there is $K > 0$ such that $|v^\top \Phi_x v| \leq K\|v\|^2$ for every $x \in \Lambda$ and every $v \in \mathbb{R}^d$.

Definition 3 (Quadratic cone field criterion). Let $f : M \to M$ be a diffeomorphism, and $\Lambda \subseteq M$ be invariant for f. Let Φ be a field of p-matrices on Λ. We say that f *satisfies the cone field criterion* with respect to Φ (or that Φ is *contracting* for f) if there is $\varepsilon > 0$ such that, for every $x \in \Lambda$,

$$Df_x^\top \Phi_{f(x)} Df_x - \Phi_x \preceq -\varepsilon I \qquad (1)$$

where I is the $d \times d$ identity matrix.

The geometric interpretation of Definition 3 is the following. If we define $\mathcal{K}_x$ as the negative level set of Φ_x:

$$\mathcal{K}_x = \{v \in \mathbb{R}^d : v^\top \Phi_x v \leq 0\},$$

then it is not hard to see that $\mathcal{K}_x$ is a cone: that is, $v \in \mathcal{K}_x$ implies that $\alpha v \in \mathcal{K}_x$ for every $\alpha \geq 0$. Because it is defined from a p-matrix, we call $\mathcal{K}_x$ a *quadratic p-cone*. In fact, p is also equal to the maximal dimension of a linear subspace contained in $\mathcal{K}_x$ (e.g., the eigenspace associated to the p negative eigenvalues of Φ_x). Finally, (1) implies that $\{\mathcal{K}_x\}_x$ is forward invariant by Df, i.e., for every $x \in \Lambda$, $\mathcal{K}_x$ is mapped by Df_x into $\mathcal{K}_{f(x)}$:

$$Df_x(\mathcal{K}_x) \subseteq \mathcal{K}_{f(x)}.$$

Similarly, if we let $\mathcal{K}^c_x$ be the "dual cone" of $\mathcal{K}_x$:

$$\mathcal{K}^c_x = \{v \in \mathbb{R}^d : v^\top \Phi_x v \geq 0\} = \mathrm{cl}(\mathbb{R}^d \setminus \mathcal{K}_x),$$

then $\mathcal{K}_x^c$ is a quadratic $(d-p)$-cone. Moreover, (1) implies that $\{\mathcal{K}_x^c\}_x$ is backward invariant by Df, i.e., $\mathcal{K}_x^c$ is mapped by Df_x^{-1} into the cone $\mathcal{K}_{f^{-1}(x)}^c$:

$$Df_x^{-1}(\mathcal{K}_x^c) \subseteq \mathcal{K}_{f^{-1}(x)}^c. \tag{2}$$

The following lemma on the minimal growth rate of the derivative along trajectories in forward and backward time is instrumental:

Lemma 1. Let f, $\Lambda \subseteq M$ and Φ be as in Definition 3, and $\{\mathcal{K}_x\}_x$ and $\{\mathcal{K}_x^c\}_x$ be as above. Then, there exist $C \geq 1$ and $\mu > 1$ such that, for every $x \in \Lambda$ and every $n \geq 0$,

- $\|Df_x^n v\| \geq C\mu^n \|v\|$ for every $v \in \mathcal{K}_x$;
- $\|Df_x^{-n} v\| \geq C\mu^n \|v\|$ for every $v \in \mathcal{K}_x^c$.

Proof: First, let $v \in \mathcal{K}_x$. Since Φ is bounded, (1) implies

$$\begin{aligned} v^\top Df_x^\top \Phi_{f(x)} Df_x v &\leq v^\top \Phi_x v - \varepsilon \|v\|^2 \\ &\leq v^\top \Phi_x v + \varepsilon K^{-1} v^\top \Phi_x v \leq \gamma v^\top \Phi_x v \end{aligned}$$

with $1 < \gamma \leq 1 + \varepsilon K^{-1}$. Thus, for $n \geq 0$,

$$\begin{aligned} -K\|Df_x^n v\|^2 &\leq v^\top (Df_x^n)^\top \Phi_{f^n(x)} Df_x^n v \\ &\leq \gamma^{n-1} v^\top Df_x^\top \Phi_{f(x)} Df_x v \\ &\leq \gamma^{n-1}(-\varepsilon\|v\|_x^2 + v^\top \Phi_x v) \leq -\varepsilon\gamma^{n-1}\|v\|^2. \end{aligned}$$

Now, let $v \in \mathcal{K}_x^c$. With a similar reasoning, we find

$$v^\top (Df_x^{-1})^\top \Phi_{f^{-1}(x)} Df_x^{-1} v \geq \gamma v^\top \Phi_x v$$

with $1 < \gamma \leq (1 - \varepsilon K^{-1})^{-1}$. Thus, if $n \geq 0$,

$$\begin{aligned} K\|Df_x^{-n} v\|^2 &\geq v^\top (Df_x^{-n})^\top \Phi_{f^{-n}(x)} Df_x^{-n} v \\ &\geq \gamma^{n-1} v^\top (Df_x^{-1})^\top \Phi_{f^{-1}(x)} Df_x^{-1} v \\ &\geq \gamma^{n-1}(\varepsilon\|v\|_x^2 + v^\top \Phi_x v) \geq \varepsilon\gamma^{n-1}\|v\|^2. \end{aligned}$$

It is now straightforward to conclude the proof. ∎

The developments above lead to the following theorem stating that the quadratic cone field criterion is a sufficient condition for hyperbolicity:

Theorem 2. Let $f : M \to M$ be a diffeomorphism, and let $\Lambda \subseteq M$ be an invariant set for f. If there exists a field of p-matrices defined on Λ that is contracting for f, then Λ has a hyperbolic structure for f.

Proof: Let $x \in \Lambda$. We show the existence of the subspace E_x^s in Definition 1 (the proof of the existence of the subspace E_x^u is similar by considering f^{-1} instead of f). Define E_x^s as the set of vectors $v \in \mathbb{R}^d$ such that $Df_x^n v \in \mathcal{K}_{f^n(x)}^c$ for every $n \geq 0$. From (2) and the definition of E_x^s, it is clear that E_x^s satisfies (ii) in Definition 1: $Df_x(E_x^s) = E_{f(x)}^s$. The main trick of the proof is to show that E_x^s is a q-dimensional subspace (where $q = d - p$ for simplicity of notation).

To show this, first observe that E_x^s is the intersection of the sets $\mathcal{S}_n := Df_{f^n(x)}^{-n}(\mathcal{K}_{f^n(x)}^c)$ for $n \geq 0$. Now, (2) implies that $\mathcal{S}_1 \supseteq \ldots \supseteq \mathcal{S}_n \supseteq \ldots$. Moreover, each $\mathcal{S}_n$ includes a q-dimensional subspace (because it is the linear image of a quadratic q-cone). This implies that $E_x^s = \bigcap_n \mathcal{S}_n$ includes a q-dimensional subspace (by compactness of the set of all q-dimensional linear subspaces of $\mathbb{R}^d$ with respect to the Grassmann metric). We will show in the last part of the proof that E_x^s is actually a q-dimensional linear subspace.

Before this, we show that E_x^s satisfies the property (iii) of Definition 1, i.e., that $\|Df_x^n v\| \leq C'\lambda^n\|v\|$ for all $v \in E_x^s$ and $n \geq 0$. This is direct from the fact that, if $v \in E_x^s$ and $w = Df_x^n v$, then $w \in \mathcal{K}_{f^n(x)}^c$ by definition of E_x^s. Hence, by Lemma 1, $\|v\| = \|Df_x^{-n} w\| \geq C\mu^n\|w\|$. It suffices to take $C' = C^{-1}$ and $\lambda = \mu^{-1}$.

Finally, we show that E_x^s is a q-dimensional subspace. Therefore, let V^s be a q-dimensional subspace included in E_x^s, and V^u be a p-dimensional subspace included in $\mathcal{K}_x$ (which is a quadratic p-cone). Assume that $E_x^s \neq V^s$. Then, there exists $v \in E_x^s$ such that $v = v^s + v^u$ with $v^s \in V^s$ and $v^u \in V^u \setminus \{0\}$. Then, Lemma 1 implies that $\|Df_x^n(v - v^s)\| = \|Df_x^n v^u\| \geq C\mu^n\|v^u\|$ for all $n \geq 0$. A contradiction with the previous paragraph, and the fact that $v - v^s \in E_x^s$. This concludes the proof of the theorem. ∎

B. Connections with the literature

The quadratic cone field criterion has a strong connection with the *Alekseev cone field criterion* introduced by V. Alekseev in 1968 [1]. Indeed, the proof of Theorem 2 is grounded in the result that the Alekseev cone field criterion provides a sufficient condition for a dynamical system to be hyperbolic on a given invariant set; see, e.g., Theorem 2 in [19] or Theorem 3.10 in [10]. However, whereas Alekseev only provides definitions of properties, with no algorithms for verifying these properties in a systematic way, our characterization of hyberbolicity, on the other hand, is meant to be translated into efficient algorithms via modern optimization techniques.

The use of symmetric matrices and Linear Matrix Inequalities to express the contraction and expansion of the derivative along the trajectories of the dynamical system is inspired from the work on p-dominant continuous-time systems by F. Forni and R. Sepulchre [8]. The novelty of our approach is to increase the expressiveness by moving from a uniform quadratic cone to a field of quadratic cones while providing a computational framework for the computation of the cone field. This requires the introduction of an abstraction of the system and tools from path-complete Lyapunov theory, as explained in Section IV.

Finally, the field of p-matrices Φ_x can be regarded as a *Finsler–Lyapunov function*, that is, a "Lyapunov" function acting on the *augmented* system $(x, \delta x) \mapsto (f(x), Df_x \delta x)$, by defining the function $V(x, \delta x) = \delta x^\top \Phi_x \delta x$ on $\Lambda \times \mathbb{R}^d$. Finsler–Lyapunov functions have been successfully applied for the contraction (aka. incremental stability, or δ-ISS) analysis of nonlinear dynamical systems; see, e.g., [18], [7]. The difference of our approach is that the functions $V(x, \delta x) = \delta x^\top \Phi_x \delta x$ are not necessarily positive-definite (whereas this is a requirement for contraction analysis), thereby allowing for directions in which the system is expanding.

IV. Computational Framework

In this section, we describe an algorithmic framework for computing a field of p-matrices Φ_x that is contracting for a given dynamical system. By assuming that the field of

p-matrices is piecewise constant, the computation can be reduced to the feasibility of a finite set of Linear Matrix Inequalities. The restriction to a piecewise constant field is performed by discretizing the state space into a finite set of regions, as explained in the following subsection.

A. Abstraction of a dynamical system

In this subsection, $f : M \to M$ is a continuous map (not necessarily diffeomorphic). A *finite covering* of $\Omega \subseteq M$ is a finite collection $\mathcal{M} = \{M_1, \dots, M_N\}$ of compact regions $M_i \subseteq M$ such that $\Omega \subseteq \bigcup_i M_i$. (In particular, this implies that Ω is compact.)

Definition 4 (Abstraction, aka. Symbolic Image). An *abstraction* of the dynamical system $f : M \to M$ on $\Omega \subseteq M$ is an ordered pair $(\mathcal{M}, E)$ where $\mathcal{M} = \{M_1, \dots, M_N\}$ is a finite covering of Ω, and $E \subseteq \{1, \dots, N\}^2$ is a set of "edges" satisfying: for every $i, j \in \{1, \dots, N\}$, $f(M_i) \cap M_j \neq \varnothing$ implies that $(i, j) \in E$.

If $(\mathcal{M}, E)$ is an abstraction, we denote by $G = G(\mathcal{M}, E)$ the directed graph whose set of vertices is equal to $\{1, \dots, N\}$, and whose edges are defined by E: that is, there is an edge $i \to j$ in G if and only if $(i, j) \in E$. See Fig. 1 for an illustration.

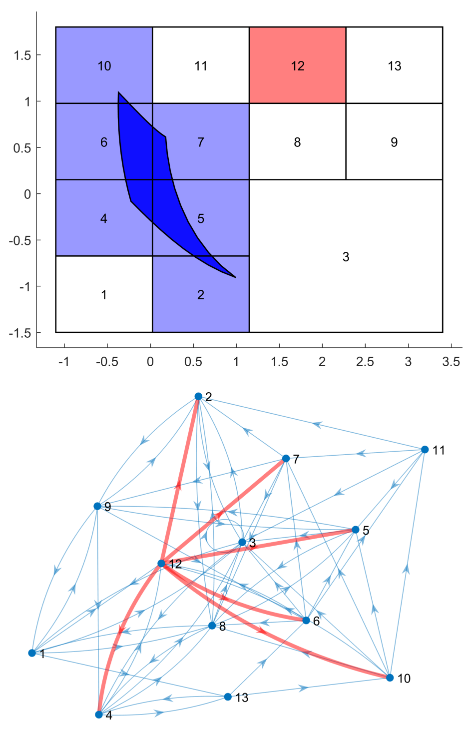

Fig. 1. *Top:* Abstraction of the Ikeda mapping (presented in Section V) on $\Omega = [-1.1, 3.4] \times [-1.5, 1.8]$. The image of the region M_{12} (in red) is represented in dark blue. The different regions that intersect the image of M_{12} are represented in light blue. *Bottom:* Graph representing the transitions (edges) between the different regions of the abstraction. The outgoing edges from vertex 12 are highlighted in red.

Definition 5 (Recurrent vertex). A vertex v of a directed graph G is called *recurrent* if there is a nontrivial (i.e., containing at least one edge) path from v to v in G. (See Fig. 2 for an illustration.)

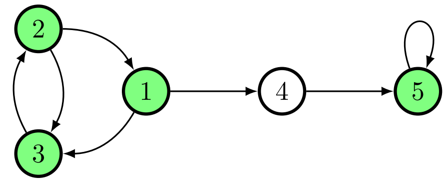

Fig. 2. Directed graph. The vertices $1, 2, 3, 5$ are recurrent.

The next proposition allows one to compute an over- (or outer-) approximation of the maximal invariant set contained in Ω if one has an abstraction of f on $\Omega \subseteq M$; see, e.g., Theorem 44 in [20] for more details. The property in the proposition will also be crucial in the proof of the correctness of the algorithm (Theorem 6 below).

Proposition 3. Let $(\mathcal{M}, E)$ be an abstraction of f on $\Omega \subseteq M$ and let $\Lambda \subseteq \Omega$ be invariant for f. Then, for every vertex i of $G = G(\mathcal{M}, E)$ such that $M_i \cap \Lambda \neq \varnothing$, there exist two recurrent vertices j_1 and j_2 such that there is a path from j_1 to i in G and there is a path from i to j_2 in G.

Proof: Let $x \in M_i \cap \Lambda \neq \varnothing$. Then, $f^n(x) \in \Lambda$ for every $n \geq 0$. This implies that there exists a forward infinite path in G starting from i. Since the number of vertices in G is finite, there is at least one vertex that is visited twice along the path. This vertex is recurrent. Similarly, because $f(\Lambda) = \Lambda$, for every $n \geq 0$ there is an $x_n \in \Lambda$ such that $f^n(x_n) = x$. Hence, there exists a backward infinite path in G ending at i, and for the same reasons as above, this backward path must contain a recurrent vertex. ∎

B. Computation of the quadratic cones

In this subsection, $f : M \to M$ is a diffeomorphism and $\Lambda \subseteq M$ is a compact invariant set for f. We let $(\mathcal{M}, E)$ be an abstraction of f on Λ. We assume that $M_i \cap \Lambda \neq \varnothing$ for each $M_i \in \mathcal{M}$ (otherwise it suffices to remove the regions with $M_i \cap \Lambda = \varnothing$). We will explain how to compute a contracting field of p-matrices that is "adapted" to this abstraction.

Definition 6 (Path-complete contracting set of p-matrices). Let f and $(\mathcal{M}, E)$ be as above. Let $\{S_1, \dots, S_N\} \subseteq \mathbb{R}^{d\times d}$, with $N = |\mathcal{M}|$, be a set of p-matrices. We say that $\{S_i\}_i$ is *path-complete contracting* with respect to f and $(\mathcal{M}, E)$ if, for every $(i, j) \in E$ and every $x \in M_i \cap f^{-1}(M_j)$,

$$Df_x^\top S_j Df_x - S_i \prec 0. \tag{3}$$

Theorem 4. Let f, $\Lambda \subseteq M$ and $(\mathcal{M}, E)$ be as above, and suppose there exists a set of p-matrices $\{S_i\}_i \subseteq \mathbb{R}^{d\times d}$ that is path-complete contracting with respect to f and $(\mathcal{M}, E)$. Then, f is hyperbolic on Λ.

Proof: We define a field of p-matrices Φ on Λ as follows: for each $x \in \Lambda$, define $\Phi_x = S_{i(x)}$ where $i(x)$ is the smallest integer $i \in \{1, \dots, |\mathcal{M}|\}$ such that $x \in M_i$. Because Df_x is continuous in x, $M_i \cap f^{-1}(M_j)$ is compact and the set $\{S_i\}_i$ is finite, we have that (i) Φ is bounded, and (ii) the right-hand term of (3) can be replaced by $-\varepsilon I$ for $\varepsilon > 0$ small enough. This shows that Φ satisfies the hypothesis of Theorem 2, concluding the proof of the theorem. ∎

Condition (3) cannot be directly handled by a computer because it involves an infinite number of LMIs. To overcome this limitation, we assume that for every edge $(i,j) \in E$, we have an approximation $\bar{A}_{i,j}$ of Df_x on $M_i \cap f^{-1}(M_j)$:

Definition 7 (δ-approximation of Df). Let f and $(\mathcal{M}, E)$ be as above. For every edge $(i,j) \in E$, let $\bar{A}_{i,j}$ be a $d \times d$ matrix. For $\delta > 0$, we say that the family of matrices $\{\bar{A}_{i,j}\}$, indexed by $(i,j) \in E$, is a *δ-approximation of Df* if, for every $(i,j) \in E$ and every $x \in M_i \cap f^{-1}(M_j)$,

$$\|Df_x - \bar{A}_{i,j}\|_2 \leq \delta \,\min\{\|\bar{A}_{i,j}\|_2^{-1}, 1\}.$$

where $\|\cdot\|_2$ denotes the matrix spectral norm.

Now, let $(\mathcal{M}, E)$ and $\{\bar{A}_{i,j}\}_{(i,j)\in E}$ be as in Definition 7, and consider the following feasibility problem:

$$\begin{array}{ll} \text{find} & S_i \in \mathbb{R}^{d\times d} \text{ symmetric, } \varepsilon \in \mathbb{R} \\ \text{subject to} & \bar{A}_{i,j}^\top S_j \bar{A}_{i,j} - S_i \preceq -\varepsilon I, \quad (i,j) \in E, \\ & -I \preceq S_i \preceq I, \quad 1 \leq i \leq N, \\ & \varepsilon > 2\delta + \delta^2. \end{array} \tag{4}$$

The following theorem makes the link between Theorem 4 and the feasibility of (4). (Remember that $\mathrm{In}(S)$ denotes the inertia of S.)

Theorem 5. Let $\delta > 0$, and assume that $\{\bar{A}_{i,j}\}_{(i,j)\in E}$ is a δ-approximation of Df. If (4) admits a feasible solution $(\{S_i\}_i, \varepsilon)$ with $\mathrm{In}(S_i) = (p, 0, d-p)$ for every $1 \leq i \leq N$, then $\{S_i\}_i$ is path-complete contracting with respect to f and $(\mathcal{M}, E)$; and thus f is hyperbolic on Λ.

Proof: Let $x \in M_i \cap f^{-1}(M_j)$, and denote $A = \bar{A}_{i,j}$ for simplicity of notation. By Definition 7, we have that $Df_x = A + \Delta$ where $\|\Delta\|_2 \leq \delta \,\min\{\|A\|_2^{-1}, 1\}$. Hence,

$$\begin{aligned} Df_x^\top S_j Df_x - S_i &= A^\top S_j A + \Delta^\top S_j A \\ &\quad + A^\top S_j \Delta + \Delta^\top S_j \Delta - S_i \\ &\preceq -\varepsilon I + \Delta^\top S_j A + A^\top S_j \Delta + \Delta^\top S_j \Delta \\ &\preceq -\varepsilon I + 2\|\Delta\|_2 \,\|A\|_2 I + \|\Delta\|_2^2 I \\ &\preceq -\varepsilon I + 2\delta I + \delta^2 I. \end{aligned}$$

Thus, $\{S_i\}_i$ satisfies (3). ∎

Theorem 6 below states that the output of (4) can be used to decide the existence of a path-complete contracting set of matrices with respect to $(\mathcal{M}, E)$, although no constraints on the inertia of the matrices $\{S_i\}_i$ are formulated in (4). This is in fact the main asset of the computational framework as it allows one to use standard SDP solvers to compute a path-complete contracting set of matrices.

Theorem 6. If (4) admits a feasible solution $(\{S_i\}_i, \varepsilon)$ with $\mathrm{In}(S_i) = (p, 0, d-p)$ for every $1 \leq i \leq N$, then every feasible solution $(\{S'_i\}_i, \varepsilon')$ satisfies $\mathrm{In}(S'_i) = (p, 0, d-p)$ for all $1 \leq i \leq N$.

Proof: The proof relies on the following result, sometimes referred to as the *Main Inertia Theorem*, due to O. Taussky [22] and R. Hill [12] (we do not provide a proof here):

Lemma 7 (Main Inertia Theorem). Let $A \in \mathbb{R}^{d\times d}$. There exists a symmetric matrix $S \in \mathbb{R}^{d\times d}$ satisfying $A^\top SA - S \prec 0$ if and only if A has no eigenvalues with $|\lambda| = 1$. Moreover, in this case, S has inertia $(r, 0, d-r)$, where r is the number of eigenvalues of A with $|\lambda| > 1$.

Using Lemma 7, we will show this *key property*: "$\mathrm{In}(S_i)$ at the recurrent vertices i is uniquely determined by $G = G(\mathcal{M}, E)$ and $\{\bar{A}_{i,j}\}$." This will imply that if (4) admits a solution with $\mathrm{In}(S_i) = (p, 0, d-p)$ for every $1 \leq i \leq N$, then any other feasible solution $(\{S'_i\}_i, \varepsilon')$ will satisfy $\mathrm{In}(S'_i) = (p, 0, d-p)$ at the recurrent vertices i.

To show the above *key property*, let $(\{S_i\}_i, \varepsilon)$ be a feasible solution of (4). Let $i \in \{1, \ldots, N\}$ be a recurrent vertex, and fix a path $P : i = i_0 \to i_1 \to \ldots \to i_k = i$ from i to i in G. Define $A_P = \bar{A}_{i_{k-1},i_k} \cdots \bar{A}_{i_1,i_2} \bar{A}_{i_0,i_1}$ and observe that the first set of constraints in (4) implies that $A_P^\top S_i A_P - S_i \prec 0$. Hence, by Theorem 7, we get that the inertia of S_i is uniquely determined by the eigenvalues of A_P.

To complete the proof, it remains to show that $\mathrm{In}(S'_i) = (p, 0, d-p)$ also holds at the non-recurrent vertices. Using Proposition 3, we would be done if we can show that: (a) "if there is a path from i to j in $G(\mathcal{M}, E)$ and $\mathrm{In}(S'_j) = (p, 0, d-p)$, then S'_i has at least p negative eigenvalues"; and (b) in the other direction: "if there is a path from j to i in $G(\mathcal{M}, E)$ and $\mathrm{In}(S'_j) = (p, 0, d-p)$, then S'_i has at least $d-p$ positive eigenvalues." For a proof of (a) and (b), we refer the reader to Proposition 4 in [3]. ∎

C. Discussion of the algorithm

Putting together the results of Subsections IV-A and IV-B, we discuss the completeness and computational complexity of the algorithm.

1) Termination of the algorithm: The two parameters that appear in the algorithm are the way the abstraction of $\Omega \subseteq M$ is built, and the choice of the Df-approximations $\bar{A}_{i,j}$. The first parameter will have an impact on how accurate the outer-approximation of the maximal invariant set Λ in Ω will be; and both parameters will influence the feasibility of (4). Moreover, the existence of a path-complete contracting set of p-matrices has been presented only as a sufficient criterion for hyperbolicity (Theorem 4), so that nothing guarantees that the algorithm will terminate in finite time.

However, it can be shown that the "path-complete contracting set of p-matrices" criterion is *asymptotically non-conservative*, meaning that, provided the accuracy of the abstraction of Λ is good enough (this can be achieved, e.g., by reducing the size of the regions), there will always exist a path-complete contracting set of p-matrices if f is hyperbolic on Λ. (The proof is left for a further paper; we refer the interested reader to [10] for related results on the sufficiency and *necessity* of the Alekseev cone criterion.)

This implies that the algorithm is *semi-complete*. This means that, if f is hyperbolic on its maximal invariant set Λ contained in $\Omega \subseteq M$, then by computing fine enough abstractions of Ω, the algorithm will always be able to prove that f is hyperbolic on an outer-approximation of Λ.

2) Computational complexity: The complexity of the algorithm is mainly driven by the complexity of computing abstractions of the invariant set Λ. For a given size of the regions, this grows in the worst case as a power of the dimension of the system; this is the curse of dimensionality of the abstraction approach. On the other hand, once the abstraction is computed, it suffices to run a SDP solver to find whether there is or not a path-complete contracting set of matrices adapted to this abstraction. The SDP problem will involve $N = |\mathcal{M}|$ matrix variables of dimension $d \times d$ and $m = |E| + 2N$ constraints; typically, $m \in \mathcal{O}(N)$.

D. *Related works*

The hyperbolicity verification problem has been addressed by George Osipenko in [20]; this is the only other work on the algorithmic hyperbolicity verification we are aware of. Osipenko's approach relies on constructing abstractions of the *augmented* system $(x, \delta x) \mapsto (f(x), Df_x \delta x)$. This requires to discretize the state space M and the "tangent space" $\mathbb{R}^d$ (more precisely, the projective space $\mathbb{PR}^{d-1}$) of the system. The Morse spectrum of the system can then be over-approximated by bounding the minimal and maximal growth rate of the derivative along cycles in the graph of the abstraction. A certificate of hyperbolicity of the system is then obtained if the over-approximation of the Morse spectrum keeps away from zero.

This approach also suffers from the curse of dimensionality since it requires to construct abstractions of a space with dimension $2d - 1$. It is difficult to have a further comparison between the two methods because this will highly depend on the size of the abstraction of Λ, which can be smaller than $\mathcal{O}(\eta^d)$, where η is the size of the regions M_i and d the dimension of the system, if Λ is low-dimensional.

V. Numerical example

In this section, we illustrate the use of the computational framework described above on a numerical example. Therefore, we consider the *modified Ikeda mapping*:

$$f(x, y) = \big(r + a(x \cos \tau - y \sin \tau), b(x \sin \tau + y \cos \tau)\big)$$

with $\tau = C_1 - C_3/(1 + x^2 + y^2)$, $r = 2$, $C_1 = 0.4$, $C_3 = 6$, $a = 0.9$, $b = -0.9$. The modified Ikeda mapping is known to have an invariant set in $\Omega = [-1.1, 3.4] \times [-1.5, 1.8]$ with a hyperbolic structure; see [20].

We have considered abstractions of the *maximal invariant set* Λ contained in Ω as represented in Fig. 3. In order to obtain an abstraction that δ-approximates Df with $\delta = 0.08$, we have used $h_x = 0.0023$ and $h_y = 0.0017$. This leads to an abstraction with 2454 vertices and 9390 edges. For this abstraction, (4) is feasible and all feasible solutions $(\{S_i\}_i, \varepsilon)$ satisfy $\mathrm{In}(S_i) = (1, 0, 1)$ for every i. Hence, f is hyperbolic on its maximal invariant set Λ contained in Ω.

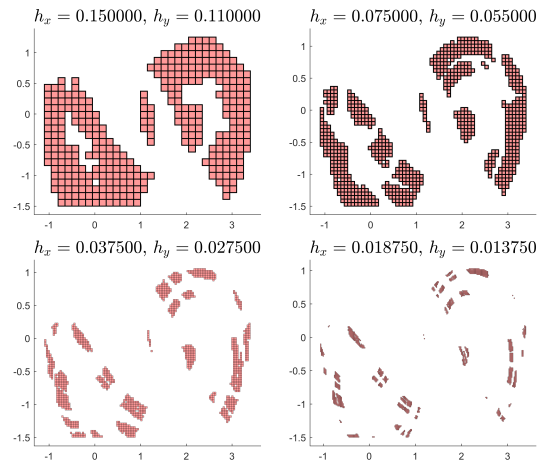

Fig. 3. Abstractions based on "square" discretizations of Ω with horizontal stepsize h_x and vertical stepsize h_y. The regions in red over-approximate the maximal invariant set Λ contained in Ω. The finer the discretization, the closer is the over-approximation to Λ. (Using square discretizations is not the best choice in general; however the purpose of this section is merely to show that the field of p-matrices can be computed using Theorem 6.)

VI. Acknowledgments

The authors would like to thank Fulvio Forni for insightful discussions on this paper.

References

[1] Vladimir Mihkaılovich Alekseev. Quasirandom dynamical systems I. Quasirandom diffeomorphisms. *Mathematics of the USSR-Sbornik*, 5(1):73, 1968. In Russian.

[2] Luis Barreira and Yakov Pesin. Smooth ergodic theory and nonuniformly hyperbolic dynamics. In *Handbook of dynamical systems*, volume 1B, pages 57–263. Elsevier, 2006.

[3] Guillaume O Berger, Fulvio Forni, and Raphaël M Jungers. Path-complete p-dominant switching linear systems. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 6446–6451. IEEE, 2018.

[4] Pierre Berger and Alvaro Rovella. On the inverse limit stability of endomorphisms. In *Annales de l'Institut Henri Poincare Non Linear Analysis*, volume 30, pages 463–475. Elsevier, 2013.

[5] Fritz Colonius and Weihua Du. Hyperbolic control sets and chain control sets. *Journal of dynamical and control systems*, 7(1):49–59, 2001.

[6] Adriano Da Silva and Christoph Kawan. Invariance entropy of hyperbolic control sets. *arXiv preprint arXiv:1408.2416*, 2014.

[7] Fulvio Forni and Rodolphe Sepulchre. A differential lyapunov framework for contraction analysis. *IEEE Trans. on Automatic Control*, 59(3):614–628, 2014.

[8] Fulvio Forni and Rodolphe Sepulchre. Differential dissipativity theory for dominance analysis. *IEEE Trans. on Automatic Control*, 2018.

[9] Boris Hasselblatt. Hyperbolic dynamical systems. In *Handbook of dynamical systems*, volume 1A, pages 239–319. Elsevier, 2002.

[10] Boris Hasselblatt. Introduction to hyperbolic dynamics and ergodic theory. In *Ergodic Theory and Negative Curvature*, pages 1–124. Springer, 2017.

[11] Boris Hasselblatt and Yakov Pesin. Partially hyperbolic dynamical systems. In *Handbook of dynamical systems*, volume 1B, pages 1–55. Elsevier, 2005.

[12] Richard D Hill. Inertia theory for simultaneously triangulable complex matrices. *Linear Algebra and its Applications*, 2(2):131–142, 1969.

[13] Morris W Hirsch and Charles C Pugh. Stable manifolds and hyperbolic sets. In *Global Analysis*, pages 133–163, 1970.

[14] Raphaël M Jungers, Amir Ali Ahmadi, Pablo A Parrilo, and Mardavij Roozbehani. A characterization of Lyapunov inequalities for stability of switched systems. *IEEE Trans. on Automatic Control*, 62(6):3062–3067, 2017.

[15] Christoph Kawan. Uniformly hyperbolic control theory. *Annual Reviews in Control*, 44:89–96, 2017.

[16] Jeffrey M Lee. *Manifolds and Differential Geometry*. Graduate studies in mathematics. American Mathematical Society, 2009.

[17] Pei-Dong Liu. Stability of orbit spaces of endomorphisms. *Manuscripta Mathematica*, 93(1):109–128, 1997.

[18] Winfried Lohmiller and Jean-Jacques E Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.

[19] Sheldon Newhouse. Cone-fields, domination, and hyperbolicity. *Modern dynamical systems and applications*, pages 419–432, 2004.

[20] George Osipenko. *Dynamical systems, graphs, and algorithms*. Springer, 2006.

[21] Clark Robinson. *Dynamical systems: stability, symbolic dynamics, and chaos (second edition)*. CRC press, 1998.

[22] Olga Taussky. A generalization of a theorem of Lyapunov. *SIAM Journal*, 9(4):640–643, 1961.