

Extremal Norms for Fiber-Bunched Cocycles

Jairo Bochi Eduardo Garibaldi

October 4, 2019

Abstract

In traditional Ergodic Optimization, one seeks to maximize Birkhoff averages. The most useful tool in this area is the celebrated Mañé Lemma, in its various forms. In this paper, we prove a non-commutative Mañé Lemma, suited to the problem of maximization of Lyapunov exponents of linear cocycles or, more generally, vector bundle automorphisms. More precisely, we provide conditions that ensure the existence of an extremal norm, that is, a Finsler norm with respect to which no vector can be expanded in a single iterate by a factor bigger than the maximal asymptotic expansion rate. These conditions are essentially irreducibility and sufficiently strong fiber-bunching. Therefore we extend the classic concept of Barabanov norm, which is used in the study of the joint spectral radius. We obtain several consequences, including sufficient conditions for the existence of Lyapunov maximizing sets.

1 Introduction

1.1 Extremal norms

Let \mathbb{E} be a d -dimensional real vector bundle over a compact metric space X , with projection map π . Let $T: X \rightarrow X$ be a homeomorphism. We say that Φ is an *automorphism of \mathbb{E} covering T* if the diagram

$$\begin{array}{ccc} \mathbb{E} & \xrightarrow{\Phi} & \mathbb{E} \\ \pi \downarrow & & \downarrow \pi \\ X & \xrightarrow{T} & X \end{array}$$

commutes and moreover the restriction of Φ to each fiber $\mathbb{E}_x := \pi^{-1}(x)$ is a linear isomorphism Φ_x onto the fiber \mathbb{E}_{T_x} . The set of such automorphisms is denoted $\text{Aut}(\mathbb{E}, T)$. The simplest situation is when the vector bundle is trivial, say $\mathbb{E} = X \times \mathbb{R}^d$. Then Φ takes the form

$$\Phi(x, u) = (T(x), F(x)u), \tag{1.1}$$

for some continuous map $F: X \rightarrow GL(d, \mathbb{R})$. The pair (T, F) is called a *(linear) cocycle*.

2010 *Mathematics subject classification*. 37H15, 37D20, 37D30, 15A60, 93D30

Acknowledgements. Bochi was partially supported by projects Fondecyt 1180371 and Conicyt PIA ACT172001. Garibaldi was partially supported by FAPESP's Thematic Project 2012/18780-0.

A *Finsler norm*¹ on \mathbb{E} is a continuous map $\|\cdot\|: \mathbb{E} \rightarrow \mathbb{R}$ whose restriction to each fiber \mathbb{E}_x is a norm. If L is a linear map from a fiber \mathbb{E}_x to another fiber \mathbb{E}_y , then we define the operator norm:

$$\|L\|_{y \leftarrow x} := \sup_{\substack{u \in \mathbb{E}_x \\ u \neq 0}} \frac{\|L(u)\|}{\|u\|}. \quad (1.2)$$

When no confusion is likely to arise we denote this simply by $\|L\|$.

Fix an automorphism Φ covering T and a Finsler norm $\|\cdot\|$. Given $x \in X$, the limit

$$\chi_1(\Phi, x) := \lim_{n \rightarrow +\infty} \frac{1}{n} \log \|\Phi_x^n\| = \lim_{n \rightarrow +\infty} \frac{1}{n} \log \|\Phi_{T^{n-1}x} \circ \cdots \circ \Phi_{Tx} \circ \Phi_x\|,$$

if it exists, is called the (*first*) *Lyapunov exponent* of Φ at the point x . The Lyapunov exponent is obviously independent of the choice of the Finsler norm. If μ is a T -invariant Borel probability measure for T , then the Lyapunov exponent $\chi_1(\Phi, x)$ exists for μ -almost every $x \in X$; this is a well-known consequence of Kingman's subadditive ergodic theorem; see e.g. [50]. Let us denote $\chi_1(\Phi, \mu) := \int \chi_1(\Phi, \cdot) d\mu$. If the measure μ is ergodic then $\chi_1(\Phi, x) = \chi_1(\Phi, \mu)$ for μ -almost every $x \in X$.

In this paper we are interested in the *maximal Lyapunov exponent*, defined as:

$$\beta(\Phi) := \sup_{\mu \in \mathcal{M}_T} \chi_1(\Phi, \mu), \quad (1.3)$$

where \mathcal{M}_T denotes the set of all T -invariant Borel probability measures. The supremum is always attained by an ergodic measure – this follows from upper semicontinuity of $\chi_1(\Phi, \cdot)$ with respect to the weak-star topology, and the fact that \mathcal{M}_T is a compact convex set whose extreme points are exactly the ergodic measures. Let us mention that the maximal Lyapunov exponent can also be characterized in more elementary terms as follows:

$$\beta(\Phi) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sup_{x \in X} \log \|\Phi_x^n\| = \sup_{x \in X} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \|\Phi_x^n\|. \quad (1.4)$$

(We use “limf” to denote a limit that is also an infimum.) These equalities follow from general results on “subadditive ergodic optimization”: see [55, Appendix A].

A trivial upper bound for the maximal Lyapunov exponent, which depends on the chosen Finsler norm, is given by:

$$\beta(\Phi) \leq \log \sup_{x \in X} \|\Phi_x\|. \quad (1.5)$$

If equality holds then $\|\cdot\|$ is called an *extremal norm* for Φ . More precisely, the norm is so “tight” that there is *no* vector $u \neq 0$ in \mathbb{E} whose expansion factor $\|\Phi(u)\|/\|u\|$ exceeds the maximum asymptotic expansion rate $e^{\beta(\Phi)}$. In particular, if $\beta(\Phi) \leq 0$ then the extremal norm is a (non-strict) Lyapunov function for Φ .

Extremal norms first appeared in the 1960 paper [64] by Rota and Strang, who considered the particular setting of one-step cocycles (details are given below), but apparently were not considered in our level of generality before.

¹Beware that other definitions of Finsler norms appear in the literature; here the main point is that the norm is not necessarily induced by inner products (i.e. “Riemannian”).

The existence of an extremal norm is far from automatic², and has strong consequences. In this paper we construct extremal norms for a large and natural class of vector bundle automorphisms.

1.2 Previous results

Consider the case of a 1-dimensional vector bundle \mathbb{E} , with an arbitrary Finsler norm $\|\cdot\|$. Given $\Phi \in \text{Aut}(\mathbb{E}, T)$, there exists a unique continuous function $f: X \rightarrow \mathbb{R}$ such that

$$u \in \mathbb{E}_x \quad \Rightarrow \quad \|\Phi(u)\|_{Tx} = e^{f(x)} \|u\|_x. \quad (1.6)$$

Note that in this case the maximal Lyapunov exponent $\beta(\Phi)$ equals:

$$\beta(f) := \sup_{\mu \in \mathcal{M}_T} \int f d\mu. \quad (1.7)$$

Any other Finsler norm $\|\cdot\|$ is of the form:

$$\|u\|_x = e^{h(x)} \|u\|_x,$$

for some continuous function $h: X \rightarrow \mathbb{R}$. Then $\|\cdot\|$ is an extremal norm if and only if h satisfies the ‘‘cohomological inequality’’:

$$f + h \circ T - h \leq \beta(f).$$

Such a function h is called a *subaction* for (T, f) . Existence of subactions can fail dramatically: see e.g. [24, §3] and [37, Appendix]. However, if the dynamics T is in some sense hyperbolic (e.g., a shift) and the function f is regular enough (e.g., Hölder) then subactions h do exist. Results of this type are sometimes called *Mañé lemmas*; see [34, 67, 33, 22, 23] for various versions and approaches, and see [14, Prop. 2.1] for a negative result. Important applications include [25, 32]. The study of invariant measures that attain that supremum in (1.7) is called *ergodic optimization*; we refer the reader to [42, 43, 37] for much more information. For a discussion of ergodic optimization in a more general context, including optimization of Lyapunov exponents, see [14].

When $\dim \mathbb{E} > 1$, commutativity is lost and much less is known. The most studied situation is the following one. Let $T: X \rightarrow X$ be the full shift on N symbols, defined on the space $X := \{0, 1, \dots, N-1\}^{\mathbb{Z}}$. Given a N -tuple (A_0, \dots, A_{N-1}) of invertible $d \times d$ matrices, let $F: X \rightarrow GL(d, \mathbb{R})$ be given by $F(x) = A_{x_0}$. We say that (T, F) is a *one-step cocycle*. Let Φ the associated automorphism (1.1). In that case, the quantity $e^{\beta(\Phi)}$ is known as the *joint spectral radius* of the set $\{A_0, \dots, A_{N-1}\}$.³ It was introduced by Rota and Strang [64].

If, for example, $N = 1$ and $A_0 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, then no extremal norm exists. However, if the set $\{A_0, \dots, A_{N-1}\}$ is *irreducible*, in the sense that there is no common invariant non-trivial subspace, then extremal norms $\|\cdot\|$ do exist, and can be taken so that $\|u\|_x$

²On the other hand, one can always construct ‘‘almost-extremal’’ norms, i.e., norms for which the inequality (1.5) is an approximate equality, and such norms can be taken Riemannian. Furthermore, it is possible to find a Riemannian norm with respect to which all the singular values of the linear maps Φ_x (and not only the first) are suitably controlled: see [14, Prop. 4.1].

³More generally, one could consider (possibly infinite) bounded sets of (possibly non-invertible) square matrices.

is independent of $x \in X$. Actually, Barabanov [10] proved that there exists a norm $\|\cdot\|$ on \mathbb{R}^d with the following stronger property:

$$\forall u \in \mathbb{R}^d, \quad \max_{i \in \{0, \dots, N-1\}} \|A_i u\| = e^{\beta(\Phi)} \|u\|. \quad (1.8)$$

For more information on the joint spectral radius and Barabanov norms, see [74, 44]. Further applications of extremal norms were obtained by Morris [54, 55].

Still in the setting of one-step cocycles, a modification of the concept of Barabanov norm was used in [18, 16] to study Lyapunov-maximizing and also Lyapunov-minimizing measures.

Extremal norms for certain locally constant cocycles over sofic shifts have been studied in the papers [58, 30].

The main purpose of this paper is to establish existence of extremal norms in a far more general setting.

1.3 The main result

We now describe the hypotheses on the automorphism Φ and the underlying dynamics T from which we will prove the existence of extremal norms. We first describe them informally, leaving the precise definitions for later sections.

First, we assume that $T: X \rightarrow X$ is a transitive *hyperbolic homeomorphism* of a compact metric space X . Hyperbolicity means that T has local stable and unstable sets with uniform exponential bounds, which satisfy a local product property. Examples include subshifts of finite type and Anosov diffeomorphisms.

Second, we assume that the vector bundle \mathbb{E} has a Hölder structure, and that the automorphism Φ respects this structure. In the case of trivial vector bundles, this means that the matrix function F in formula (1.1) is Hölder continuous.

Third, we assume that the automorphism Φ is *fiber-bunched*. In crude terms, this means that the non-conformality of the linear maps Φ_x is small when compared to the hyperbolicity rates of T . The precise condition involves the Hölder exponent of the automorphism, so that more regular automorphisms are allowed to be less conformal. In the case that T and Φ are differentiable, fiber-bunching means that the projectivization of Φ is a *partially hyperbolic* diffeomorphism.

Actually, for $d \geq 3$ we need to assume a stronger form of fiber-bunching.

Our last assumption is *irreducibility*, meaning that Φ admits no nontrivial regular subbundle, where regular means as regular as the automorphism itself. We remark that this condition is satisfied for typical fiber-bunched automorphisms: it holds on an open and dense subset of infinite codimension.

The main result of this paper is that *under the conditions above, extremal norms exist*. See Corollary 5.1 for a precise statement.

In the case where the base dynamics T is a subshift of finite type, we are able to improve our main result and obtain an extremal norm with a further property akin to the Barabanov property: see Subsection 5.3.

Classical Barabanov norms are usually non-Riemannian (that is, they do not come from inner products), and it is easy to produce examples⁴. On the other hand, in our setting, there is more flexibility as the norm is allowed to depend on the basepoint. So one could wonder if the Finsler extremal norms in our main result could be taken Riemannian. Unfortunately, that is not the case: we construct an explicit example in Appendix B.2.

⁴The pair of matrices (B.1) is one such example.

1.4 Consequences

As a consequence of our result on the existence of extremal norms, we can show that the maximal Lyapunov exponent is a locally Lipschitz function on the space of strongly bunched irreducible automorphisms (see Proposition 5.4 for a more precise statement), thus extending a result of Wirth [74].

We are also able to obtain several general properties of strongly bunched automorphisms Φ (not necessarily irreducible):

- Their growth obeys certain uniform bounds: see Theorem 4.6.
- They obey the *subordination principle*: if μ and ν are invariant probability measures such that ν is Lyapunov maximizing in the sense that $\chi_1(\Phi, \nu) = \beta(\Phi)$, and $\text{supp } \mu \subseteq \text{supp } \nu$, then μ is Lyapunov maximizing as well: see Theorem 7.1. This property is far from being tautological, even in the commutative setting; in fact it was introduced in this setting by Bousch [22].
- The maximal Lyapunov exponent $\beta(\Phi)$ can be approximated by Lyapunov exponents of measures supported on periodic orbits, and moreover the quality of this approximation is superpolynomial with respect to the period: see Theorem 7.2. This extends a result of Morris [54], who gave a quantitative version of the celebrated theorem of Berger–Wang [11].

We also introduce *Mather sets* in our context; these sets are the habitat of Lyapunov maximizing measures. We prove an important structural result on the existence of *dominated splittings* on the Mather sets, namely Theorem 6.5, which is an essential ingredient in the proof of the aforementioned Theorem 7.2.

1.5 Organization of the paper

In Section 2 we introduce the setting for our results, providing the definitions and properties of fiber-bunched automorphisms and related concepts.

In Section 3 we study irreducibility and related concepts.

In Section 4 we provide sufficient conditions for *relative product boundedness*, an intermediate property which is required for the existence of extremal norms.

The construction of extremal norms is given in Section 5, together with the construction of Barabanov-like norms for shifts and an application to the regularity of $\beta(\cdot)$.

In Section 6 we introduce Mather sets in a very general setting and, under the assumption of existence of an extremal norm, establish fine properties about them.

In Section 7 we collect several applications of our results.

Appendix A contains the proofs of several subsidiary results, therefore making the paper self-contained.

In Appendix B we exhibit some “pathological” examples, including an example that fits in the setting of our main results, but where no Riemannian extremal norm exists.

2 The fiber-bunched setting

In this section, we fix the basic setting for our theorems. Namely, we define and state the basic properties of Hölder vector bundles, intrinsically hyperbolic homeomorphisms, fiber-bunching, holonomies, and irreducibility. Our approach is influenced by

[20, 72, 46], and we tried to make it as general as possible. We also obtain some new regularity results that are essential for the main theorems of the paper. However, to make the presentation more fluid, we postpone most proofs to Appendix A.

2.1 The Hölder exponent

From now on, assume that (X, d) is a compact metric space. We also fix $\theta > 0$ such that the algebra of θ -Hölder functions on X is *normal*, that is, given any two disjoint compact subsets of X , there exists a θ -Hölder function that takes values in the interval $[0, 1]$ and equals 0 on one set and 1 on the other. This assumption is automatically satisfied if $\theta \leq 1$. If X is a Cantor set, then the assumption holds for any $\theta > 0$. Normality implies the existence of θ -Hölder partitions of unity: see e.g. [48, p. 221].

2.2 Hölder vector bundles

Let \mathbb{E} be a d -dimensional vector bundle over X . We recall the definition and fix the terminology. \mathbb{E} is a topological space endowed with a continuous map $\pi: \mathbb{E} \rightarrow X$ (called the *projection*), a cover of X by open sets U_i (called *coordinate neighborhoods*), and a family of homeomorphisms (called *coordinate maps*)

$$\psi_i: U_i \times \mathbb{R}^d \rightarrow \pi^{-1}(U_i) \quad \text{such that } \pi(\psi_i(x, u)) = x \text{ for all } (x, u) \in U_i \times \mathbb{R}^d,$$

which is required to have the following compatibility property: whenever $x \in U_i \cap U_j$, the map

$$g_{j \leftarrow i}(x) := [\psi_j(x, \cdot)]^{-1} \circ \psi_i(x, \cdot): \mathbb{R}^d \rightarrow \mathbb{R}^d$$

is linear. Therefore we obtain a family of continuous maps:

$$g_{j \leftarrow i}: U_i \cap U_j \rightarrow GL(d, \mathbb{R}), \tag{2.1}$$

which are called *coordinate transformations*. Moreover, each *fiber* $\mathbb{E}_x := \pi^{-1}(x)$ has a unique structure of d -dimensional vector space such that the maps

$$h_i(x) := \psi_i(x, \cdot): \mathbb{R}^d \rightarrow \mathbb{E}_x \tag{2.2}$$

become isomorphisms. Since X is assumed to be compact, we will from now on assume that the cover $\{U_i\}$ is finite.

We say that \mathbb{E} is a *θ -Hölder vector bundle* if the coordinate transformations (2.1) are locally θ -Hölder. By compactness, we can reduce the coordinate neighborhoods so that the coordinate transformations become (uniformly) θ -Hölder.

As mentioned in Subsection 1.1, a *Finsler norm* is a continuous function $\|\cdot\|$ on \mathbb{E} that restricts to a norm $\|\cdot\|_x$ on each fiber \mathbb{E}_x . A Finsler norm $\|\cdot\|$ is called *Riemannian* if each $\|\cdot\|_x$ is induced by an inner product $\langle \cdot, \cdot \rangle_x$. A Finsler norm $\|\cdot\|$ is called *θ -Hölder* if for every $u \in \mathbb{R}^d$ and every coordinate neighborhood, the function $x \in U_i \mapsto \|h_i(x)u\|$ is θ -Hölder. Every θ -Hölder vector bundle \mathbb{E} admits a θ -Hölder Riemannian norm; the proof is straightforward using a θ -Hölder partition of unity.

We will also need a way of “transporting” vectors from one fiber to another:

Proposition 2.1 (Transport maps). *Let \mathbb{E} be a θ -Hölder vector bundle. There exists a family of linear maps $I_{y \leftarrow x}: \mathbb{E}_x \rightarrow \mathbb{E}_y$ with the following properties:*

- (a) For every point $x \in X$, the linear map $I_{x \leftarrow x}$ equals the identity.
- (b) For every pair of indices i, j , the matrix-valued map

$$(x, y) \in U_i \times U_j \mapsto [h_j(y)]^{-1} \circ I_{y \leftarrow x} \circ h_i(x)$$

is θ -Hölder.

See Appendix A.1 for the proof of Proposition 2.1. The next propositions, also proved in Appendix A.1, give additional quantitative properties of the transport maps that will be useful in subsequent calculations. Recall that we agree to denote a norm and its induced operator norm by the same symbol, as in (1.2).

Proposition 2.2. *Let \mathbb{E} be a θ -Hölder vector bundle, endowed with a Finsler norm. Let $\{I_{y \leftarrow x}\}$ be the family of transport maps provided by Proposition 2.1. Then there is $C > 0$ such that for all points $x, y, z \in X$,*

$$\|I_{y \leftarrow z} \circ I_{z \leftarrow x} - I_{y \leftarrow x}\| \leq C \max\{d(x, z)^\theta, d(y, z)^\theta\},$$

Proposition 2.3. *Let \mathbb{E} be a θ -Hölder vector bundle, endowed with a Finsler norm $\|\cdot\|$. Let $\{I_{y \leftarrow x}\}$ be the family of transport maps provided by Proposition 2.1. Then the Finsler norm $\|\cdot\|$ is θ -Hölder if and only if there exists $C > 0$ such that for all points $x, y \in X$,*

$$\|I_{y \leftarrow x}\| - 1 \leq Cd(x, y)^\theta.$$

2.3 θ -Hölder bundle automorphisms

Assume that \mathbb{E} is a θ -Hölder vector bundle over the compact metric space X . Fix a θ -Hölder Riemannian norm on \mathbb{E} .

A map $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ is called an *endomorphism* of \mathbb{E} if there exists a continuous map $T: X \rightarrow X$ such that $\pi \circ \Phi = T \circ \pi$ (we say that Φ *covers* T) and for each $x \in X$, the restriction of Φ to the fiber \mathbb{E}_x is a linear map Φ_x to the fiber \mathbb{E}_{Tx} . If T is a homeomorphism and each Φ_x is an isomorphism then we say that Φ is an *automorphism*.

We say that the endomorphism Φ covering T is θ -Hölder if T is Lipschitz and the maps

$$x \in U_i \cap T^{-1}(U_j) \mapsto [h_j(Tx)]^{-1} \circ \Phi_x \circ h_i(x) \in GL(d, \mathbb{R})$$

are θ -Hölder.⁵ As an immediate consequence, the function $x \in X \mapsto \|\Phi_x\|$ is θ -Hölder.

We can characterize θ -Hölder automorphisms in terms of the transport maps from Proposition 2.1:

Proposition 2.4. *An endomorphism $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ covering a Lipschitz map T is θ -Hölder if and only if there exists $K > 0$ such that for all $x, y \in X$, we have*

$$\|I_{Ty \leftarrow Tx} \circ \Phi_x - \Phi_y \circ I_{y \leftarrow x}\| \leq Kd(x, y)^\theta.$$

A proof is provided in Appendix A.1.

Next, we want to topologize the set of θ -Hölder automorphisms.

⁵This is similar to the definition of θ -bounded vertical shear in [62].

Let $\text{End}^\theta(\mathbb{E}, T)$ denote the vector space of θ -Hölder endomorphisms $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ covering T . Define the C^0 norm:

$$\|\Phi\|_0 := \sup_{x \in X} \|\Phi_x\|. \quad (2.3)$$

The stronger θ -Hölder norm makes $\text{End}^\theta(\mathbb{E}, T)$ a Banach space:

$$\|\Phi\|_\theta := \max \left\{ \|\Phi\|_0, \sup_{x \neq y} \frac{\|I_{Ty \leftarrow Tx} \circ \Phi_x - \Phi_y \circ I_{y \leftarrow x}\|}{d(x, y)^\theta} \right\}. \quad (2.4)$$

The set $\text{Aut}^\theta(\mathbb{E}, T)$ of θ -Hölder automorphisms is a C^0 -open subset of $\text{End}^\theta(\mathbb{E}, T)$. Given $K \geq 1$, let:

$$\text{Aut}_K^\theta(\mathbb{E}, T) := \{\Phi \in \text{Aut}^\theta(\mathbb{E}, T) ; \|\Phi\|_\theta \leq K, \|\Phi^{-1}\|_\theta \leq K\}. \quad (2.5)$$

2.4 Hyperbolic homeomorphisms

The concept of hyperbolicity in differentiable dynamical systems was introduced by Anosov [3] and Smale [70]. Even without recourse to a differentiable structure, it is possible to define hyperbolicity (and to prove interesting theorems); this has been done in various ways: [26, 65, 2, 1, 4]. In this paper, we will use a minor variation of the definition of hyperbolic homeomorphism given by Sakai [66] (see Remark 2.5 below).

Recall that X is a compact metric space. Let $T: X \rightarrow X$ be a homeomorphism. Given $x \in X$ and $\varepsilon > 0$, we define the following sets:

- *local unstable set* $W_\varepsilon^u(x) := \{y \in X ; d(T^{-n}y, T^{-n}x) \leq \varepsilon \text{ for all } n \geq 0\}$;
- *local stable set* $W_\varepsilon^s(x) := \{y \in X ; d(T^n y, T^n x) \leq \varepsilon \text{ for all } n \geq 0\}$.

We say that T is a *hyperbolic homeomorphism* if the following axioms hold:

- T is bi-Lipschitz, i.e., both T and T^{-1} are Lipschitz;
- there exist a constant $\varepsilon_0 > 0$ and a pair of continuous positive functions λ_u, λ_s (called the *hyperbolicity exponents*) such that:

$$x \in X, x', x'' \in W_{\varepsilon_0}^u(x) \Rightarrow d(T^{-1}x', T^{-1}x'') \leq e^{-\lambda_u(x)} d(x', x''), \quad (2.6)$$

$$y \in X, y', y'' \in W_{\varepsilon_0}^s(y) \Rightarrow d(Ty', Ty'') \leq e^{-\lambda_s(y)} d(y', y''); \quad (2.7)$$

- there exists a constant $\varepsilon_1 \in (0, \varepsilon_0)$ such that for any pair of points $x, y \in X$ with $d(x, y) \leq 2\varepsilon_1$, the intersection $W_{\varepsilon_0}^u(x) \cap W_{\varepsilon_0}^s(y)$ contains exactly one point, denoted by $[x, y]$ and called *the bracket of x and y* , which depends continuously on x and y ;
- there exists a constant $C > 0$ such that:

$$x, y \in X, d(x, y) \leq 2\varepsilon_1 \Rightarrow \max \{d([x, y], x), d([x, y], y)\} \leq Cd(x, y). \quad (2.8)$$

Remark 2.5. Sakai [66] uses the terminology \mathcal{L} -hyperbolic homeomorphism, while Ruelle [65] uses *Smale spaces*. Modulo a change of metric, both definitions are equivalent to ours, and also to expansivity plus the shadowing property: see [66] and references cited there.

Let us also define other sets associated with T :

- *unstable set* $W^u(x) := \{y \in X ; d(T^{-n}y, T^{-n}x) \rightarrow 0 \text{ as } n \rightarrow +\infty\}$;
- *stable set* $W^s(x) := \{y \in X ; d(T^n y, T^n x) \rightarrow 0 \text{ as } n \rightarrow +\infty\}$;

If T is a hyperbolic homeomorphism then, as an immediate consequence of part (b) of the definition, for every $\varepsilon \in (0, \varepsilon_0]$ we have the following set relations:

$$W^u(x) = \bigcup_{n \geq 0} T^n(W_\varepsilon^u(T^{-n}x)), \quad W^s(x) = \bigcup_{n \geq 0} T^{-n}(W_\varepsilon^s(T^n x)). \quad (2.9)$$

The transverse regularity of the unstable and stable sets is a classical subject, and fine results about hyperbolicity rely on it: see [47, Chapter 19]. Nevertheless, we could not find a reference for the following property for hyperbolic homeomorphisms:

Proposition 2.6. *Let T be a hyperbolic homeomorphism. There exist constants $0 < \kappa_s \leq 1$ and $C > 0$ such that if $x, x', y, y' \in X$ satisfy (see Fig. 1):*

$$x' \in W_{\varepsilon_0}^u(x), \quad y' \in W_{\varepsilon_0}^u(y), \quad y \in W_{\varepsilon_0}^s(x), \quad y' \in W_{\varepsilon_0}^s(x'), \quad (2.10)$$

then:

$$d(y, y') \leq C d(x, x')^{\kappa_s}.$$




Figure 1: Four points in the configuration (2.10).

The proof, which includes an estimate for the constant κ_s , is given in Appendix A. If T is the restriction of a C^2 diffeomorphism to a basic hyperbolic basic set then a better estimate for κ_s is given in [68]. Even better regularity estimates can be obtained under various types of extra assumptions: see [59] and references therein.

2.5 Fiber-bunched automorphisms and their holonomies

We now discuss the notion of *fiber-bunching*. It was introduced in a setting very similar to ours by Bonatti, Gómez-Mont, and Viana [20], though related concepts can be traced back to Brin and Pesin [29] and Hirsch, Pugh, and Shub [41]. Earlier papers [20, 21] use a different terminology (“dominated cocycles”), but subsequently the term “fiber-bunched cocycles” prevailed: [7, 8, 46].

If L is a linear isomorphism between inner product spaces, we define its *bolicity*⁶ as

$$\text{bol}(L) := \|L\| \|L^{-1}\|, \quad (2.11)$$

which measures the lack of conformality of L (see Proposition 3.1).

Let \mathbb{E} be a θ -Hölder d -dimensional vector bundle over X . Assume that T is a hyperbolic homeomorphism, and that Φ is a θ -Hölder automorphism of \mathbb{E} covering T . We say that \mathbb{E} is *fiber-bunched* if there exists a Riemannian norm (sometimes called an *adapted norm*) such that for all $x \in X$,

$$\log \text{bol}(\Phi_x) < \min \{ \theta \lambda_u(x), \theta \lambda_s(x) \}, \quad (2.12)$$

where λ_u, λ_s are the hyperbolicity rates of T . By perturbing the adapted norm if necessary, we can assume it is also θ -Hölder.

Consider the subset of fiber-bunched automorphisms in the space $\text{Aut}^\theta(\mathbb{E}, T)$ of θ -Hölder automorphisms; then this set is open with respect to the C^0 norm (2.3), and therefore also open with respect to the stronger θ -Hölder norm (2.4).

Sometimes we need stronger bunching: we say that it is (η_u, η_s) -*bunched* for certain constants $\eta_u, \eta_s \in (0, \theta]$ if, for some adapted norm, and all $x \in X$,

$$\log \text{bol}(\Phi_x) < \min \{ \eta_u \lambda_u(x), \eta_s \lambda_s(x) \}. \quad (2.13)$$

Remark 2.7. We have used the *pointwise* definition of fiber-bunching; the more stringent notion of *absolute* fiber-bunching requires the same condition with *constant* hyperbolicity exponents λ_u, λ_s . Furthermore, our definition of fiber-bunching is *immediate* in the sense that it manifests itself in a single iterate; one can also define a notion of *eventual* fiber-bunching.

The most basic and fruitful consequence of fiber-bunching is the existence of certain unstable and stable holonomy maps. Like the transport maps from Proposition 2.1, unstable and stable holonomy maps provide a way of linearly transporting vectors from a fiber \mathbb{E}_x to another fiber \mathbb{E}_y (as long as the points x, y belong to the same unstable or stable set), but with several extra properties:

Proposition 2.8 (Holonomy maps). *Let $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$ be a fiber-bunched automorphism. For each $\star \in \{u, s\}$, there exist a unique family of linear maps $H_{y \leftarrow x}^\star : \mathbb{E}_x \rightarrow \mathbb{E}_y$, defined whenever $y \in W^\star(x)$, such that the following properties hold:*

- (a) $H_{x \leftarrow x}^\star = \text{id}$.
- (b) $H_{z \leftarrow y}^\star \circ H_{y \leftarrow x}^\star = H_{z \leftarrow x}^\star$.
- (c) $\Phi_y \circ H_{y \leftarrow x}^\star = H_{T y \leftarrow T x}^\star \circ \Phi_x$.
- (d) *There exists a constant $C > 0$ such that:*

$$y \in W_{\varepsilon_0}^\star(x) \quad \Rightarrow \quad \|H_{y \leftarrow x}^\star - I_{y \leftarrow x}\| \leq C d(x, y)^\theta. \quad (2.14)$$

- (e) *The following map is continuous:*

$$\begin{aligned} \{(u, y) \in \mathbb{E} \times X ; y \in W_{\varepsilon_0}^\star(\pi(u))\} &\rightarrow \mathbb{E} \\ (u, y) &\mapsto H_{y \leftarrow \pi(u)}^\star(u) \end{aligned} \quad (2.15)$$

⁶The term and the notation come from [60, 61]. In numerical analysis, the bolicity is called *condition number*.

Furthermore, the constant C in (2.14) works for all automorphisms in a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$, and the right-hand side in (2.15) depends continuously on the automorphism in that neighborhood.

The maps H^u and H^s are called *unstable* and *stable* holonomies, respectively. Properties (a) and (b) are called *groupoid properties*, and property (c) is called *equivariance*.

The stable holonomies are actually defined by the following formula:

$$H_{y \leftarrow x}^s := \lim_{n \rightarrow +\infty} (\Phi_y^n)^{-1} \circ I_{T^n y \leftarrow T^n x} \circ \Phi_x^n,$$

and unstable holonomies are defined likewise, taking $n \rightarrow -\infty$ instead. The proof of Proposition 2.8 consists essentially in proving uniform convergence in these formulas, and it turns out that fiber-bunching is the precise condition for this to work. Except for minor adjustments, the argument is the same as in [20, §1.4], [46, §4.1], but for completeness and convenience of the reader we present the proof in Appendix A.2.

Remark 2.9. Fiber-bunched automorphisms satisfy a non-commutative version of Walters' condition [22], namely:

$$\forall \varepsilon > 0 \exists \delta > 0 \text{ such that } \sup_{i \in [0, n]} d(T^i x, T^i y) < \delta \Rightarrow \|\Phi_y^n \circ I_{y \leftarrow x} - I_{T^n y \leftarrow T^n x} \circ \Phi_x^n\| < \varepsilon.$$

Indeed, consider $z := [x, y]$ and note the following identity:

$$\Phi_y^n = H_{T^n x \leftarrow T^n z}^s \circ H_{T^n z \leftarrow T^n x}^u \circ \Phi_x^n \circ H_{x \leftarrow z}^u \circ H_{z \leftarrow y}^s.$$

Then, using the continuity of the bracket and the regularity of the holonomies, it is straightforward to obtain the non-commutative Walters' condition.

We use the holonomies to define certain subsets of \mathbb{E} . For $\varepsilon > 0$, $u \in \mathbb{E}$, and $x = \pi(u)$, let:

$$\begin{aligned} \mathbb{W}_\varepsilon^u(u) &:= \{H_{y \leftarrow x}^u(u) ; y \in W_\varepsilon^u(x)\}, \\ \mathbb{W}^u(u) &:= \{H_{y \leftarrow x}^u(u) ; y \in W^u(x)\} = \bigcup_{n \geq 0} \Phi^n(\mathbb{W}_{\varepsilon 0}^u(\Phi^{-n}(u))), \end{aligned}$$

Analogously we define $\mathbb{W}_\varepsilon^s(u)$ and $\mathbb{W}^s(u)$. The sets \mathbb{W}^u (resp. \mathbb{W}^s) form a Φ -invariant partition of \mathbb{E} and project by π onto the sets W^u (resp. W^s).

Part (d) of Proposition 2.8 basically says that the “leaves” \mathbb{W}^u , \mathbb{W}^s are Hölder-continuous. We will need the transverse regularity of the holonomies:

Proposition 2.10. *Let $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$ be a fiber-bunched automorphism. There exist $\theta_s \in (0, \theta_{\kappa_s}]$ and $C > 0$ such that if $x, x', y, y' \in X$ satisfy conditions (2.10) as in Fig. 1 then:*

$$\|H_{y' \leftarrow y}^u \circ H_{y \leftarrow x}^s - H_{y' \leftarrow x'}^s \circ H_{x' \leftarrow x}^u\| \leq C d(x, x')^{\theta_s}.$$

Furthermore, the same constants θ_s and C work for every automorphism in a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$.

We were not able to find such a statement in the literature, so we provide a proof in Appendix A.3.

3 Invariant subbundles

3.1 Subbundles and rigidity

Let E be an inner product space of dimension d , and let $p \in \llbracket 1, d \rrbracket$. We denote by $\mathcal{G}_p(E)$ the p -th Grassmannian of E , i.e., the set of all p -dimensional subspaces of E . There are many metrics on this set that are “natural” in the sense that they are preserved by the action of orthogonal linear maps: see [63]. As shown in Appendix A.4, we can find one such metric d with the useful properties stated in the following propositions:

Proposition 3.1. *If $L: E \rightarrow F$ is a linear isomorphism between d -dimensional inner product spaces and $p < d$, then the induced map $\mathcal{G}_p(E) \rightarrow \mathcal{G}_p(F)$ is Lipschitz with a constant equal to the bolicity of L (2.11).*

Proposition 3.2. *If $L: E \rightarrow E$ is a linear isomorphism of a inner product space such that $\|L - \text{id}\| \leq \delta$ then the induced map on $\mathcal{G}_p(E)$ is $O(\delta)$ -close to the identity, provided δ is sufficiently small.*

Proposition 3.3. *The map that associates a p -tuple of linearly independent vectors to its span is locally Lipschitz.*

In particular, the metric d induces the usual topology on the Grassmannian.

Now consider a θ -Hölder d -dimensional vector bundle \mathbb{E} over X . For each $p \in \llbracket 1, d - 1 \rrbracket$, let $\mathcal{G}_p(\mathbb{E})$ denote the fiber bundle whose fiber over $x \in X$ is $\mathcal{G}_p(\mathbb{E}_x)$. As just explained, the fixed Riemannian norm on \mathbb{E} induces a distance on each fiber of this bundle.

Let \mathbb{F} be a continuous p -dimensional subbundle of \mathbb{E} . We say that \mathbb{F} is η -Hölder, for some $\eta \in (0, \theta]$, if there exists $C > 0$ such that for all sufficiently close points $x, y \in X$ we have:

$$d(\mathbb{F}_y, I_{y \leftarrow x}(\mathbb{F}_x)) \leq C d(x, y)^\eta, \quad (3.1)$$

where the d in the left hand side is the distance in $\mathcal{G}_p(\mathbb{E}_y)$. (Recall that $I_{y \leftarrow x}$ is an isomorphism when x and y are close enough.)

Let Φ be a fiber-bunched automorphism of \mathbb{E} . We say that a subbundle $\mathbb{F} \subseteq \mathbb{E}$ is Φ -invariant if for all $x \in X$, we have

$$\Phi_x(\mathbb{F}_x) = \mathbb{F}_{Tx}.$$

We say that \mathbb{F} is H^u -invariant (or W^u -saturated) if for all $x \in X$ and all $y \in W^u(x)$, we have

$$H_{y \leftarrow x}^u(\mathbb{F}_x) = \mathbb{F}_y.$$

We say that \mathbb{F} is η -Hölder along unstable sets, for some $\eta \in (0, \theta]$, if there exists $C > 0$ and $\varepsilon > 0$ such that the estimate (3.1) holds whenever $y \in W_\varepsilon^u(x)$. Equivalently, there exists $C \geq 0$ such that for all $x \in X$ and all $y \in W_{\varepsilon_0}^u(x)$, we have:

$$d(\mathbb{F}_y, H_{y \leftarrow x}^u(\mathbb{F}_x)) \leq C d(x, y)^\eta; \quad (3.2)$$

to see the equivalence, use θ -Hölderiness of the holonomy (2.14) and Proposition 3.2. H^s -invariance and η -Hölderiness along stable sets are defined analogously.

Proposition 3.4 (Rigidity). *Let $\eta \in (0, \theta]$. Suppose that Φ is (η, θ) -bunched. Let $\mathbb{F} \subseteq \mathbb{E}$ be a continuous Φ -invariant subbundle. If \mathbb{F} is η -Hölder along unstable sets then \mathbb{F} is H^u -invariant, and in particular \mathbb{F} is actually θ -Hölder along unstable sets.*

Proof. Since Φ is (η, θ) -bunched, there is a constant $r \in (0, 1)$ such that $\text{bol}(\Phi_x) < re^{\eta\lambda_u(x)}$ for every $x \in X$. Now fix $x \in X$ and $y \in W_{\varepsilon_0}^u(x)$. For each $n \geq 0$, let $x_n := T^{-n}x$ and $y_n := T^{-n}y$. Then:

$$\begin{aligned} d(\mathbb{F}_y, H_{y \leftarrow x}^u(\mathbb{F}_x)) &= d(\Phi_{y_n}^n(\mathbb{F}_{y_n}), \Phi_{y_n}^n(H_{y_n \leftarrow x_n}^u(\mathbb{F}_{x_n}))) \quad (\text{by } \Phi\text{-invariance of } \mathbb{F}) \\ &\leq \text{bol}(\Phi_{y_n}^n) d(\mathbb{F}_{y_n}, H_{y_n \leftarrow x_n}^u(\mathbb{F}_{x_n})) \quad (\text{by Proposition 3.1}) \\ &\leq C \text{bol}(\Phi_{y_n}^n) d(x_n, y_n)^\eta \quad (\text{by } \eta\text{-H\"olderness of } \mathbb{F}). \end{aligned}$$

On one hand, by submultiplicativity of bolicity, $\text{bol}(\Phi_{y_n}^n) \leq \text{bol}(\Phi_{y_1}) \cdots \text{bol}(\Phi_{y_n})$. On the other hand, using (2.6) recursively,

$$\begin{aligned} d(x_n, y_n) &= d(T^{-1}x_{n-1}, T^{-1}y_{n-1}) \leq e^{-\lambda_u(y_{n-1})} d(x_{n-1}, y_{n-1}) \leq \cdots \\ &\leq e^{-\lambda_u(y_0) - \cdots - \lambda_u(y_{n-1})} d(x, y). \end{aligned}$$

Combining these estimates, we have:

$$\begin{aligned} d(\mathbb{F}_y, H_{y \leftarrow x}^u(\mathbb{F}_x)) &\leq C \left[\prod_{j=1}^n \text{bol}(\Phi_{y_j}) \right] \left[\prod_{j=0}^{n-1} e^{-\eta\lambda_u(y_j)} \right] d(x, y)^\eta \\ &\leq CB^2 r^{n-2} d(x, y)^\eta, \end{aligned}$$

where $B \geq 1$ is the maximal bolicity. As $n \rightarrow \infty$, the right hand side tends to zero. So $\mathbb{F}_y = H_{y \leftarrow x}^u(\mathbb{F}_x)$, proving that the subbundle \mathbb{F} is H^u -invariant. Since (3.2) holds with $C = 0$, the subbundle \mathbb{F} is θ -H\"older along unstable sets. \square

Corollary 3.5. *Let Φ be a fiber-bunched automorphism of \mathbb{E} . Let $\mathbb{F} \subseteq \mathbb{E}$ be a Φ -invariant continuous subbundle. Then the following conditions are equivalent:*

- (a) \mathbb{F} is a θ -H\"older subbundle;
- (b) \mathbb{F} is both H^u - and H^s -invariant.

Proof. If condition (a) holds then \mathbb{F} is θ -H\"older along unstable sets, and so Proposition 3.4 (with $\eta = \theta$) guarantees that \mathbb{F} is H^u -invariant. By symmetry, \mathbb{F} is also H^s -invariant. That is, condition (b) holds. Conversely, assume that condition (b) holds, and consider a pair x, y of nearby points. Then the bracket $z := [x, y]$ is well-defined, and by property (2.8), it is $O(d(x, y))$ -close to either x or y . By hypothesis, $\mathbb{F}_y = H_{y \leftarrow z}^s \circ H_{z \leftarrow x}^u(\mathbb{F}_x)$. Using Proposition 2.2 and θ -H\"olderness of the holonomies (2.14), we see that $\|I_{y \leftarrow x} - H_{y \leftarrow z}^s \circ H_{z \leftarrow x}^u\| = O(d(x, y)^\theta)$. It follows from Proposition 3.2 that $d(\mathbb{F}_y, I_{y \leftarrow x}(\mathbb{F}_x)) = O(d(x, y)^\theta)$, i.e., condition (a) holds. \square

3.2 Irreducibility

The *trivial* subbundles of \mathbb{E} are the zero section and \mathbb{E} itself. A fiber-bunched automorphism Φ is called *reducible* if it has a nontrivial Φ -invariant subbundle \mathbb{F} satisfying either of the equivalent conditions of Corollary 3.5, and *irreducible* otherwise.

While the existence of continuous Φ -invariant subbundles is common, the existence of θ -H\"older ones is not. For example, if the automorphism admits a dominated splitting⁷, then the subbundles that form the splitting are Φ -invariant, continuous,

⁷See Subsection 6.3 for the definition and properties of dominated splittings

and actually Hölder, but usually with smaller Hölder exponent. Actually, the dominating bundle is H^u -invariant and so θ -Hölder along unstable sets, but usually not so well behaved along stable sets.

A precise formulation of the fact that reducibility is uncommon among fiber-bunched automorphisms is provided by Proposition A.6.

3.3 The strong bunching hypothesis

If $d = 2$ then ordinary fiber-bunching suffices for our main results, while if $d \geq 3$ we need Φ to be not only fiber-bunched, but (η_0, θ) -bunched, where η_0 is given by the following:

Lemma 3.6. *There exists $\eta_0 \in (0, \theta]$ that depends only on the hyperbolic homeomorphism T (or, more precisely, on its hyperbolicity exponents) and on the Hölder exponent θ such that if Φ is a (η_0, θ) -bunched automorphism then the associated regularity exponent θ_s from Proposition 2.10 satisfies:*

$$\theta_s \geq \eta_0.$$

For the proof (and an explicit value for η_0), see Appendix A.3.

Let us say that a θ -Hölder automorphism $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ covering T is *strongly bunched* if:

- the vector bundle \mathbb{E} has fibers of dimension $d \leq 2$, and Φ is fiber-bunched; **or**
- Φ is a (η_0, θ) -bunched automorphism.

The precise point of our proofs where we need strong bunching is for the validity of Theorem 3.7, explained in the next subsection.

3.4 Spannability

The following concept of *spannability* will play an important role in this paper; it is vaguely similar to the concept of accessibility in partially hyperbolic dynamics (see e.g. [57, §8.1]).

Let us say that a fiber-bunched automorphism Φ is *spannable* if for all $x, y \in X$, and all nonzero $u \in \mathbb{E}_x$, there exist:

- points $x_1, \dots, x_d \in W^u(x)$;
- integers $n_1, \dots, n_d \geq 0$ such that the points $y_i := T^{n_i} x_i$ all belong to $W^s(y)$;

with the property that the vectors $v_1, \dots, v_d \in \mathbb{E}_y$ defined by

$$v_i := H_{y \leftarrow y_i}^s \circ \Phi_{x_i}^{n_i} \circ H_{x_i \leftarrow x}^u(u) \tag{3.3}$$

form a basis for \mathbb{E}_y .

It is clear that every spannable automorphism is irreducible. The following important result provides a converse under extra assumptions:

Theorem 3.7 (Sufficient conditions for spannability). *Let T be a transitive hyperbolic homeomorphism. Let Φ be a strongly bunched irreducible automorphism covering T . Then Φ is spannable.*

In particular (see Proposition A.6), typical strongly bunched automorphisms are spannable (provided T is transitive).

It would be interesting to know whether or not strong bunching is really necessary for the validity of Theorem 3.7; see Remark 3.12 below for a possible approach to this question.

In order to prove the theorem, we need the following easy property of the unstable and stable sets for the base dynamics:

Lemma 3.8. *For every $x \in X$, the sets $\bigcup_{n \geq 0} W^u(T^n x)$ and $\bigcup_{n \geq 0} W^s(T^{-n} x)$ are dense in X .*

Proof. Let D be the set of points whose forward orbits are dense. Since T is transitive, D is itself dense. Moreover, D is W^s -saturated (i.e., it is a union of stable sets). By definition of hyperbolic homeomorphism, local stable and unstable sets whose base-points are sufficiently close always intersect. It follows that D intersects all unstable sets. This implies that for every $x \in X$, the set $\bigcup_{n \geq 0} W^u(T^n x)$ is dense. Applying this to T^{-1} we obtain that $\bigcup_{n \geq 0} W^s(T^{-n} x)$ is also dense. \square

Proof of Theorem 3.7. Fix a point $x \in X$ and a nonzero vector $u \in \mathbb{E}_x$. Let $\Lambda := \bigcup_{n \geq 0} W^u(T^n x)$, which by Lemma 3.8 is a dense subset of X . Define the following subsets of the vector bundle \mathbb{E} :

$$\mathbb{U} := \bigcup_{n \geq 0} W^u(\Phi^n(u)), \quad \mathbb{S} := \bigcup_{v \in \mathbb{U}} W^s(v), \quad \mathbb{F} := \text{span}(\mathbb{S}),$$

where the latter equation means that for each $y \in X$, the fiber $\mathbb{F}_y := \mathbb{E}_y \cap \mathbb{F}$ is the vector space spanned by $\mathbb{S}_y := \mathbb{E}_y \cap \mathbb{S}$. In order to prove the theorem, we need to show that $\mathbb{F} = \mathbb{E}$. Clearly,

- \mathbb{U} projects onto Λ , and is both forward- Φ -invariant (i.e., $\Phi(\mathbb{U}) \subseteq \mathbb{U}$) and W^u -saturated (i.e., it is a union of W^u sets);
- \mathbb{S} projects onto X , and is both forward- Φ -invariant and W^s -saturated; therefore \mathbb{F} has the same properties.

We claim that the function $y \in X \mapsto \dim \mathbb{F}_y$ has the following properties:

- (a) it is non-decreasing along orbits of T (i.e., $\dim \mathbb{F}_{Ty} \geq \dim \mathbb{F}_y$);
- (b) it is constant along W^s sets;
- (c) it is lower semicontinuous.

Indeed, properties (a) and (b) follow from the facts that \mathbb{F} is forward-invariant and W^s -saturated, respectively. In order to check property (c), fix an arbitrary point $y \in X$ and let $p := \dim \mathbb{F}_y$. Then there exist points $x_1, \dots, x_p \in W^u(x)$ and integers $n_1, \dots, n_p \geq 0$ such that the points $y_i := T^{n_i} x_i$ all belong to $W^s(y)$, and the vectors v_i given by formula (3.3) span \mathbb{F}_y . If y' is sufficiently close to y , then for each i we can find $y'_i \in W^u(y_i) \cap W^s(y')$ such that the holonomies $H_{y'_i \leftarrow y_i}^u$ and $H_{y' \leftarrow y'_i}^s$ are respectively close to the identity and $H_{y \leftarrow y_i}^s$. Then each vector $v'_i := H_{y' \leftarrow y'_i}^s \circ H_{y'_i \leftarrow y_i}^u \circ H_{y_i \leftarrow y}^s(v_i)$ is close to v_i , and so the span of $\{v'_1, \dots, v'_p\}$ has dimension p . Since each v'_i belongs to $\mathbb{F}_{y'}$, we conclude that $\dim \mathbb{F}_{y'} \geq p$, therefore proving the semicontinuity property (c).

Let C be the set of points $y \in X$ where $\dim \mathbb{F}_y$ attains its minimum. By the properties (a), (b), and (c) that we have just proved, the set C is nonempty, backwards-invariant (i.e., $T^{-1}(C) \subseteq C$), W^s -saturated, and closed. It follows from Lemma 3.8 that $C = X$. In other words, \mathbb{F} has constant dimension, say p . So \mathbb{F} is not only forward- Φ -invariant, but Φ -invariant.

Let θ_s be given by Proposition 2.10. We claim that \mathbb{F} is θ_s -Hölder along unstable sets, in the sense defined in Subsection 3.1. By compactness, it suffices to prove this claim on a neighborhood of each point $y \in X$. Take points $x_1, \dots, x_p \in W^u(x)$ and integers $n_1, \dots, n_p \geq 0$ such that the points $y_i := T^{n_i} x_i$ all belong to $W^s(y)$, and the vectors v_i given by formula (3.3) span \mathbb{F}_y . Take $k \geq 0$ large enough so that the points $T^k y_i$ all belong to $W^s_{\varepsilon_1}(T^k y)$, where ε_1 is constant from condition (c) in the definition of hyperbolic homeomorphism. If we prove that \mathbb{F} is θ_s -Hölder along unstable sets on a neighborhood of $T^k y$ then, by invariance, it will follow that \mathbb{F} is θ_s -Hölder along unstable sets on a neighborhood of y . So let us assume that $k = 0$, for simplicity of notation. Let $y' \in W^u_{\varepsilon_1}(y)$ be close to y . Then the brackets $y'_i := [y_i, y']$ are well-defined; see Fig. 2. We need to compare the following two subspaces of $\mathbb{E}_{y'}$:

$$\begin{aligned} \mathbb{F}_{y'} &= \text{span} \left\{ H_{y' \leftarrow y'_i}^s \circ \underbrace{H_{y'_i \leftarrow y_i}^u \circ H_{y_i \leftarrow y}^s}_{\textcircled{1}}(v_i) \right\}_{i=1}^p, \\ H_{y' \leftarrow y}^u(\mathbb{F}_y) &= \text{span} \left\{ H_{y' \leftarrow y'_i}^s \circ \underbrace{H_{y'_i \leftarrow y'}^s \circ H_{y' \leftarrow y}^u}_{\textcircled{2}}(v_i) \right\}_{i=1}^p. \end{aligned}$$

By Proposition 2.10, $\|\textcircled{1} - \textcircled{2}\| = O(d(y, y')^{\theta_s})$; moreover $\|H_{y' \leftarrow y'_i}^s\| = O(1)$. So, by Proposition 3.3, we conclude that $d(\mathbb{F}_{y'}, H_{y' \leftarrow y}^u(\mathbb{F}_y)) = O(d(y, y')^{\theta_s})$. This concludes the proof that \mathbb{F} is θ_s -Hölder along unstable sets. A fortiori, \mathbb{F} is continuous (since it is invariant under stable holonomies).




Figure 2: Proof of Theorem 3.7

The proof ends differently according to the dimension of \mathbb{E} . If $d = 1$ then $\mathbb{F} = \mathbb{E}$ and we are done.

Next, consider the case $d = 2$. Assume for a contradiction that $\mathbb{F} \neq \mathbb{E}$, i.e., that \mathbb{F} has 1-dimensional fibers. For each $y \in \Lambda$, the set \mathbb{U}_y contains a nonzero vector and therefore spans \mathbb{F}_y . Since Λ is dense in X and \mathbb{F} is continuous, we conclude that \mathbb{F}

is the closure of $\text{span}(\mathbb{U})$. In particular, \mathbb{F} is \mathbb{W}^u -saturated. Recalling that \mathbb{F} is also \mathbb{W}^s -saturated, we contradict irreducibility. This concludes the proof in the case $d = 2$.

Now consider the case $d \geq 3$. Then, by definition of strong bunching, Φ is (η_0, θ) -bunched. Recall from Lemma 3.6 that $\eta_0 \leq \theta_s$. So Proposition 3.4 (rigidity) applies and the regularity of the subbundle is upgraded: it is actually θ -Hölder along unstable sets. Irreducibility implies that $\mathbb{F} = \mathbb{E}$, thus concluding the proof. \square

We will use an apparently stronger, but equivalent form of spannability. Recall that $\varepsilon_1 > 0$ is one of the constants that appear in the definition of hyperbolic homeomorphism (Subsection 2.4).

Proposition 3.9 (Uniform spannability). *Suppose Φ is a spannable automorphism. Then there exist constants $\bar{n} \geq 0$ and $C_0 > 0$ with the following properties: For all points $x, y \in X$, and all unit vectors $u \in \mathbb{E}_x$, there exist:*

- points $x_1, \dots, x_d \in W_{\varepsilon_1}^u(x)$;
- integers $n_1, \dots, n_d \in \llbracket 0, \bar{n} \rrbracket$ such that the points $y_i := T^{n_i} x_i$ all belong to $W_{\varepsilon_1}^s(y)$;

with the property that the vectors $v_1, \dots, v_d \in \mathbb{E}_y$ defined by (3.3) form a basis for \mathbb{E}_y ; moreover, if $L: \mathbb{E}_y \rightarrow \mathbb{E}_y$ is a linear map that sends this basis to an orthonormal basis then $\|L\| < C_0$.

Proof. If $u \in \mathbb{E}$ is a nonzero vector, let $[u]$ denote its class in the projective bundle $\hat{\mathbb{E}} := \mathcal{G}_1(\mathbb{E})$. Let Φ be a spannable automorphism. Given $([u], y) \in \hat{\mathbb{E}} \times X$, consider $x = \pi(u)$ and let x_i, n_i, y_i , and v_i , where $i \in \llbracket 1, d \rrbracket$, be as in the definition of spannability. Note that if $([\tilde{u}], \tilde{y})$ belongs to a sufficiently small neighborhood of $([u], y)$ then we can find the corresponding data $(\tilde{x}_i, \tilde{n}_i, \tilde{y}_i, [\tilde{v}_i])$ close to $(x_i, n_i, y_i, [v_i])$ (so $\tilde{n}_i = n_i$) and actually depending continuously on $([\tilde{u}], \tilde{y})$. Since the space $\hat{\mathbb{E}} \times X$ is compact, we can cover it by finitely many such neighborhoods U_j . We can also assume that the sets U_j are compact.

Fix any set U_j and an element $([u], y) \in U_j$. Let $x = \pi(u)$ and let $(x_i, n_i, y_i, [v_i])$, $i \in \llbracket 1, d \rrbracket$ be the corresponding spannability data. For each $k \geq 0$, the pair $([\Phi^{-k}(u)], T^k y) \in \hat{\mathbb{E}} \times X$ has $(T^{-k} x_i, n_i + 2k, T^k y_i, [\Phi^k(v_i)])$, $i \in \llbracket 1, d \rrbracket$ as valid spannability data. By (2.9), if k is large enough then

$$T^{-k} x_i \in W_{\varepsilon_1}^u(T^{-k} x) \quad \text{and} \quad T^k y_i \in W_{\varepsilon_1}^s(T^k y) \quad \text{for each } i \in \llbracket 1, d \rrbracket.$$

By continuity of the spannability data on the compact set U_j , this conclusion holds provided k is bigger than some k_j . There are finitely many indices j to consider, so let us fix a definitive k bigger than all k_j 's. The compact sets $V_j := \{([\Phi^{-k}(u)], T^k y) ; ([u], y) \in U_j\}$ also cover the space $\hat{\mathbb{E}} \times X$. They provide the spannability data with the required uniformity properties. This proves the proposition. \square

Corollary 3.10. *Given a spannable automorphism $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$, we can choose $\bar{n} \geq 0$ and $C_0 > 0$ such that the the statement of Proposition 3.9 holds for all automorphisms in a C^0 -neighborhood of Φ in the space $\text{Aut}_K^\theta(\mathbb{E}, T)$.*

In particular, spannable automorphisms form a C^0 -open subset of $\text{Aut}_K^\theta(\mathbb{E}, T)$.

Proof. By part (e) of Proposition 2.8, holonomies depend continuously on the fiber-bunched automorphism Φ , with respect to the C^0 -norm. So, in the situation of Proposition 3.9, if we make a C^0 -perturbation of Φ (among θ -Hölder automorphisms) then the vectors v_1, \dots, v_d change little and therefore stay linearly independent. \square

Remark 3.11. Let us say that a automorphism is *topologically irreducible* if it admits no *continuous* proper invariant subbundle. As the proof of Theorem 3.7 shows, if a fiber-bunched automorphism over a transitive hyperbolic homeomorphism is topologically irreducible then it is spannable.

Remark 3.12. As Clark Butler has pointed out to us, if a fiber-bunched automorphism satisfies the *pinching-and-twisting* condition from [21, Def. 1.3], [7, Def. 1.2] then it is spannable. In other words, one can remove the strong bunching hypothesis from Theorem 3.7, provided one replaces irreducibility with the (strictly stronger) pinching-and-twisting condition.

Let us sketch the proof. Let $\mathbb{U} \subseteq \mathbb{F}$ be as in proof of Theorem 3.7. Let \mathbb{V} be the closure of the span of $\bigcup_{n \geq 0} \Phi^{-n}(\mathbb{U})$; then \mathbb{V} is Φ -invariant, \mathbb{W}^u -saturated, projects down on X , and is contained in \mathbb{F} . Let μ be the T -invariant probability measure on X with maximal entropy (other choices are possible). Let $\mathbb{P}\Phi$ be the projectivization of the automorphism Φ . Then $\mathbb{P}\Phi$ admits an *invariant u-state*, that is, an invariant measure \hat{m} that projects on μ and whose disintegration w.r.t. to this projection is μ -a.e. invariant under unstable holonomies: see [7, § 4.1]. By adapting the construction, we can ensure that the invariant u-state \hat{m} gives full weight to $\mathbb{P}\mathbb{V}$, and in particular to $\mathbb{P}\mathbb{F}$, which is a continuous \mathbb{W}^s -saturated invariant subbundle of $\mathbb{P}\mathbb{E}$. Under the pinching-and-twisting assumption, [21, Prop. 5.1] or [7, Prop. 5.1] say that such a situation is impossible unless $\mathbb{F} = \mathbb{E}$. (Actually in these papers T is a shift, but the proofs can be adapted to the general situation, or we can use a Markov partition.) Therefore Φ is spannable.

It is not clear how to relax the pinching-and-twisting hypothesis in the arguments from [21, 7]. Therefore we still lack an optimal criterion for spannability.

4 Bounding the growth

4.1 Relative product boundedness

A vector bundle automorphism Φ is called *product bounded* if

$$\sup_{n \geq 0} \sup_{x \in X} \|\Phi_x^n\| < \infty$$

for some and hence any Finsler norm on \mathbb{E} . This condition evidently implies that $\beta(\Phi) \leq 0$, i.e. the maximal Lyapunov exponent (1.3) is nonpositive. On the other hand, we say that Φ is *relatively product bounded* if $e^{-\beta(\Phi)}\Phi$ is product bounded, that is,

$$\sup_{n \geq 0} e^{-\beta(\Phi)n} \sup_{x \in X} \|(\Phi^n)_x\| < \infty.$$

Of course, if Φ has an extremal norm then it is relatively product bounded. The converse is true in the 1-step case, as noted by Rota and Strang [64]. But the converse is not true in general⁸; in fact it may fail even in dimension 1, as shown by Morris [53, Proposition 2]. In Morris' example, the dynamics is uniformly hyperbolic (actually a full shift), but the function is not Hölder.

In this paper, we need to prove relative product boundedness as an essential preliminary step in the construction of extremal norms. We will show the following:

Proposition 4.1. *Every spannable automorphism is relatively product bounded.*

The proof, which will occupy Subsections 4.2 and 4.3, is roughly as follows: first, we find pieces of \mathbb{W}^u sets of uniform size that stay relatively product bounded for a

⁸The naive attempt of defining an extremal norm by $\|u\| := \sup_{n \geq 0} e^{-\beta(\Phi)n} \|\Phi^n(u)\|$ does not necessarily work because continuity may fail.

long time (Lemma 4.3), then we use a compactness argument to find small pieces of \mathbb{W}^u sets that stay relatively product bounded forever (Lemma 4.4), and finally we use spannability to spread this property to the whole bundle (Lemma 4.5).

Let us close this subsection with some remarks about product boundedness and relative product boundedness.

It was shown by Blondel and Tsitsiklis [12] that the product boundedness of a pair of rational matrices is algorithmically undecidable.

A result of Coronel, Navas, and Ponce [35] states if T is a *minimal* homeomorphism (i.e. all its orbits are dense) and Φ and Φ^{-1} are both product bounded then there exists an *invariant* Riemannian norm.

It is easy to give examples of regular (e.g. Hölder) automorphisms that are not relatively product bounded: any cocycle constant equal to $(\begin{smallmatrix} 1 & \\ 0 & 1 \end{smallmatrix})$ will do. Here is a more interesting example:

Example 4.2. Let the base dynamics T be an irrational rotation of the circle \mathbb{R}/\mathbb{Z} , and consider the matrix-valued map $F(x) := (\begin{smallmatrix} 2 & 0 \\ 0 & 1/2 \end{smallmatrix}) R_{2\pi x}$, where R_θ denotes the rotation matrix by angle θ . As shown by Herman [40, p. 471–473], the $SL(2, \mathbb{R})$ -cocycle (T, F) has a positive Lyapunov exponent, but it is not uniformly hyperbolic. Therefore it cannot be relatively product bounded, because otherwise it would contradict a result of Morris [54, Theorem 2.1].

4.2 Existence of local unstable sets with relatively bounded orbits

Let Φ be a fiber-bunched automorphism in the set $\text{Aut}_K^\theta(\mathbb{E}, T)$. By the definition (2.5) of this set,

$$\|\Phi_x^{\pm 1}\| \leq K \quad \text{for all } x \in X. \quad (4.1)$$

By equicontinuity of local holonomies, there exists a constant $C_1 > 1$ such that:

$$\|H_{y \leftarrow x}^\star\| < C_1 \quad \text{for all } x \in X, \star \in \{u, s\}, y \in W_{\varepsilon_0}^\star(x). \quad (4.2)$$

Moreover, by Proposition 2.8, it is possible to choose a constant C_1 that works for all automorphisms in a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$.

Let \mathbb{E}^\times denote the complement of the zero section in \mathbb{E} . Recall that ε_0 comes from the definition of hyperbolic homeomorphism.

Lemma 4.3. *Let Φ be a fiber-bunched automorphism. Then there exists $\varepsilon_2 \in (0, \varepsilon_0)$, depending only on T , such that for every integer $m > 0$ there exists $u \in \mathbb{E}^\times$ with the following property:*

$$\sup_{n \in [1, m]} \sup_{v \in \mathbb{W}_{\varepsilon_2}^u(u)} e^{-\beta(\Phi)n} \|\Phi^n(v)\| \leq 2\|u\|.$$

Proof. Multiplying Φ by a nonzero constant, we can assume that $\beta(\Phi) = 0$. Let λ_u be the hyperbolicity exponent of T along unstable sets, and let

$$a := \sup_{x \in X} e^{-\lambda_u(x)} < 1. \quad (4.3)$$

Hyperbolicity implies:

$$\forall x \in X, \forall \varepsilon \in (0, \varepsilon_0], \quad T^{-1}(W_\varepsilon^u(x)) \subseteq W_{a\varepsilon}^u(T^{-1}x). \quad (4.4)$$

Let $\varepsilon_2 := (1 - a)\varepsilon_0$. In order to show that the conclusion of the lemma holds for this ε_2 , let us assume for a contradiction that there exists an integer $m > 0$ such that:

$$\forall u \in \mathbb{E}^\times \exists n = n(u) \in \llbracket 1, m \rrbracket \exists v = v(u) \in \mathbb{W}_{\varepsilon_2}^u(u) \text{ s.t. } \|\Phi^n(v)\| > 2\|u\|. \quad (4.5)$$

We recursively define sequences $(u_k), (v_k)$ in \mathbb{E}^\times and (n_k) in $\llbracket 1, m \rrbracket$ as follows: We choose $u_0 \in \mathbb{E}^\times$ arbitrarily. Assuming u_k was already defined, we let $n_k := n(u_k)$ and $v_k := v(u_k)$ be given by (4.5), and let $u_{k+1} := \Phi^{n_k}(v_k)$. Note that for each $k \geq 0$ we have $\|u_{k+1}\| > 2\|u_k\|$ and so $\|u_k\| \geq 2^k\|u_0\|$.

Now let $\ell_k := n_0 + n_1 + \dots + n_{k-1}$ (so $\ell_0 := 0$), and $w_k := \Phi^{-\ell_k}(v_k)$. We claim that each w_k belongs to $\mathbb{W}_{\varepsilon_0}^u(u_0)$; indeed:

$$\begin{aligned} v_k &\in \mathbb{W}_{\varepsilon_2}^u(u_k) && \Rightarrow \\ \Phi^{-n_{k-1}}(v_k) &\in \mathbb{W}_{a\varepsilon_2}^u(v_{k-1}) \subseteq \mathbb{W}_{(1+a)\varepsilon_2}^u(u_{k-1}) && \Rightarrow \\ \Phi^{-n_{k-2}-n_{k-1}}(v_k) &\in \mathbb{W}_{(a+a^2)\varepsilon_2}^u(v_{k-2}) \subseteq \mathbb{W}_{(1+a+a^2)\varepsilon_2}^u(u_{k-2}) && \Rightarrow \\ &\vdots && \\ w_k &= \Phi^{-n_0-\dots-n_{k-1}}(v_k) \in \mathbb{W}_{(a+\dots+a^k)\varepsilon_2}^u(v_0) \subseteq \mathbb{W}_{(1+a+\dots+a^k)\varepsilon_2}^u(u_0), \end{aligned}$$

proving the claim. In particular, by (4.2) we obtain $\|w_k\| \leq C_1\|u_0\|$. Since $v_k \in \mathbb{W}_{\varepsilon_2}^u(u_k)$, using (4.2) again we have $\|u_k\| \leq C_1\|v_k\|$. Therefore:

$$\frac{\|\Phi^{\ell_k}(w_k)\|}{\|w_k\|} = \frac{\|v_k\|}{\|w_k\|} \geq \frac{C_1^{-1}\|u_k\|}{C_1\|u_0\|} \geq C_1^{-2}2^k.$$

Since $\ell_k \leq mk$, using (1.4) we obtain

$$\beta(\Phi) = \limsup_{k \rightarrow \infty} \sup_{x \in X} \frac{\log \|\Phi_x^{\ell_k}\|}{\ell_k} \geq \limsup_{k \rightarrow \infty} \frac{\log(\|\Phi^{\ell_k} w_k\|/\|w_k\|)}{\ell_k} \geq \frac{\log 2}{m} > 0.$$

This contradiction concludes the proof. \square

The next lemma supersedes the previous one:

Lemma 4.4. *Let $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$ be a fiber-bunched automorphism. Then there exist a constant $C_2 > 1$ and a vector $u_* \in \mathbb{E}^\times$ such that*

$$\sup_{n \geq 0} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u_*)} e^{-\beta(\Phi)n} \|\Phi^n(v)\| \leq C_2\|u_*\|.$$

Moreover, the same constant C_2 works for all automorphisms in a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$.

Proof. Again, multiplying Φ by a nonzero constant (and increasing K if necessary), we can assume that $\beta(\Phi) = 0$.

Let ε_2 be given by Lemma 4.3. By the continuity of the bracket, there exists $\varepsilon_3 \in (0, \varepsilon_1)$ such that:

$$z_1, z_2 \in X, \text{ d}(z_1, z_2) < 2\varepsilon_3 \quad \Rightarrow \quad \text{d}([z_1, z_2], z_i) \leq \varepsilon_2.$$

For each integer $m \geq 1$, Lemma 4.4 provides $u_m \in \mathbb{E}^\times$, say with $\|u_m\| = 1$, such that for every $n \in \llbracket 1, m \rrbracket$ and every $v \in \mathbb{W}_{\varepsilon_2}^u(u_m)$ we have $\|\Phi^n(v)\| \leq 2$. Passing

to a subsequence if necessary, we assume that (u_m) converges to some \bar{u} , which has $\|\bar{u}\| = 1$. Let $x_m := \pi(u_m)$ and $\bar{x} := \pi(\bar{u})$.

We claim that

$$\sup_{n \geq 0} \sup_{\bar{v} \in \mathbb{W}_{\varepsilon_3}^u(\bar{u})} \|\Phi^n(\bar{v})\| \leq 2C_1. \quad (4.6)$$

Indeed, given $\bar{v} \in \mathbb{W}_{\varepsilon_3}^u(\bar{u})$ and $n \geq 0$, consider $\bar{y} := \pi(\bar{v})$. Since $x_m \rightarrow \bar{x}$, for every sufficiently large $m \geq n$ we have $d(x_m, \bar{y}) < 2\varepsilon_3$, and in particular $y_m := [x_m, \bar{y}]$ is well-defined and belongs to $W_{\varepsilon_2}^u(x_m)$. Let $v_m := H_{y_m \leftarrow x_m}^u(u_m)$ and $w_m := H_{\bar{y} \leftarrow y_m}^s(v_m)$ (see Fig. 3).




Figure 3: Proof of Lemma 4.4.

Then:

$$\|\Phi^n(w_m)\| = \|H_{T^m \bar{y} \leftarrow T^m y_m}^s(\Phi^n(v_m))\| \leq C_1 \|\Phi^n(v_m)\| \leq 2C_1.$$

As $m \rightarrow \infty$ (recall that n is fixed), we have $y_m \rightarrow [\bar{x}, \bar{y}] = \bar{y}$ and so:

$$w_m = H_{\bar{y} \leftarrow y_m}^s \circ H_{y_m \leftarrow x_m}^u(u_m) \rightarrow H_{\bar{y} \leftarrow \bar{x}}^u(\bar{u}) = \bar{v},$$

by continuity of holonomies. It follows that $\|\Phi^n(\bar{v})\| \leq 2C_1$, completing the proof of the claim (4.6).

Fix a constant $\ell > 0$ depending only on T such that $W_{\varepsilon_1}^u(T^\ell \bar{x}) \subseteq T^\ell(W_{\varepsilon_3}^u(\bar{x}))$. Let $u_* := \Phi^\ell(\bar{u})$. Then

$$\sup_{n \geq 0} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u_*)} \|\Phi^n(v)\| \leq \sup_{n \geq 0} \sup_{\bar{v} \in \mathbb{W}_{\varepsilon_3}^u(\bar{u})} \|\Phi^{n+\ell}(\bar{v})\| \leq 2C_1,$$

by (4.6). On the other hand, recalling (4.1), we have $\|u_*\| \geq K^{-\ell}\|u\| = K^{-\ell}$. So the vector u_* has the desired property with $C_2 := 2K^\ell C_1$, completing the proof of the lemma. \square

4.3 Proof of relative product boundedness

The next lemma uses spannability to spread local product boundedness from a local unstable set to the whole space:

Lemma 4.5. *Let $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$ be a spannable automorphism. There exists $C_3 > 1$ with the following properties. Suppose $u \in \mathbb{E}$ is a nonzero vector such the following quantity is finite:*

$$r := \frac{1}{\|u\|} \limsup_{n \rightarrow \infty} e^{-n\beta(\Phi)} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u)} \|\Phi^n(v)\|.$$

Then

$$\limsup_{n \rightarrow \infty} e^{-n\beta(\Phi)} \sup_{y \in X} \|\Phi_y^n\| \leq C_3 r.$$

Furthermore, the same constant C_3 works for all automorphisms in a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$.

Note that the hypothesis of Lemma 4.5 is non-void by Lemma 4.4, and that its conclusion implies that Φ is relatively product bounded. So Lemma 4.5 implies Proposition 4.1. However, the more technical statement of Lemma 4.5 is necessary for the construction of an extremal norm in the next section.

Proof of Lemma 4.5. It is sufficient to consider $\beta(\Phi) = 0$. Let \bar{n} and C_0 be the uniform spannability constants provided by Proposition 3.9. Fix a nonzero vector u for which the associated quantity r is finite, and without loss of generality, let us assume that $\|u\| = 1$. Let $r' > r$ be arbitrary. Then there exists n_* such that

$$\sup_{n \geq n_*} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u)} \|\Phi^n(v)\| \leq r'.$$

Consider arbitrary $y \in X$ and $w \in \mathbb{E}_y$. Apply Proposition 3.9 to the points $x := \pi(u)$ and y and the vector u , obtaining points $x_1, \dots, x_d \in W_{\varepsilon_1}^u(x)$ and times $n_1, \dots, n_d \in \llbracket 0, \bar{n} \rrbracket$ such that each point $y_i := T^{n_i} x_i$ belongs to $W_{\varepsilon_1}^s(y)$ and the vectors v_i defined by (3.3) form a basis for \mathbb{E}_y . Moreover, if we express w as a linear combination $a_1 v_1 + \dots + a_d v_d$, then the proposition also yields that $(\sum_i a_i^2)^{1/2} \leq C_0 \|w\|$. So each $|a_i| \leq C_0 \|w\|$. For each i and $n \geq n_*$, we have

$$\Phi^n(v_i) = \underbrace{H_{T^n y \leftarrow T^n y_i}^s}_{\textcircled{1}} \circ \underbrace{\Phi_{x_i}^{n_i+n} \circ H_{x_i \leftarrow x}^u}_{\textcircled{2}}(u).$$

We have $\|\textcircled{1}\| \leq C_1$ by (4.2), and $\|\textcircled{2}\| \leq r'$ by definition. Combining these estimates, we obtain:

$$\|\Phi^n(w)\| \leq dC_0 C_1 r' \|w\|,$$

that is, $\|\Phi_y^n\| \leq C_3 r'$, where $C_3 := dC_0 C_1$. So $\limsup_{n \rightarrow \infty} \sup_{y \in X} \|\Phi_y^n\|$ is bounded by $C_3 r'$, and actually by $C_3 r$, since $r' > r$ is arbitrary. This proves the desired inequality.

Now consider a C^0 -perturbation of Φ in the set $\text{Aut}_K^\theta(\mathbb{E}, T)$. By Corollary 3.10, this perturbation is also spannable, and we can use the same constants \bar{n} and C_0 . So the argument above applies verbatim for the perturbed automorphism. \square

4.4 Application: polynomial bounds

Let us give an application of what we have proved so far, namely that under the hypothesis of strong fiber-bunching, relative product boundedness fails at most by a polynomial factor. The reader anxious to see extremal norms may skip this subsection.

Theorem 4.6. *Let T be a transitive hyperbolic homeomorphism. Let $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ be strongly bunched automorphism covering T . Then there exists an integer $d' \in \llbracket 0, d-1 \rrbracket$ and $C > 0$ such that*

$$\|\Phi_x^n\| \leq C n^{d'} e^{n\beta(\Phi)} \quad \text{for all } x \in X \text{ and } n \geq 0.$$

For related results, see [46, Theorem 3.10], [44, §3.5–3.6].

Before proving this theorem, let us fix some terminology. Suppose \mathbb{E} is a θ -Hölder vector bundle over X , with a fixed θ -Hölder Riemannian norm, and that $\mathbb{F} \subseteq \mathbb{E}$ is a θ -Hölder subbundle. Let $\mathbb{F}^\perp \subseteq \mathbb{E}$ be the orthogonal complement subbundle, which is also θ -Hölder. Then the orthogonal projections

$$P: \mathbb{E} \rightarrow \mathbb{F} \quad \text{and} \quad Q: \mathbb{E} \rightarrow \mathbb{F}^\perp \tag{4.7}$$

are θ -Hölder endomorphisms covering id_X . Now suppose $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ is θ -Hölder automorphism covering T and that \mathbb{F} is Φ -invariant. Then there are two induced θ -Hölder automorphisms, both covering T , namely the obvious *restricted automorphism* $\Phi|_{\mathbb{F}}: \mathbb{F} \rightarrow \mathbb{F}$, and the *quotient automorphism* $\Phi/\mathbb{F}: \mathbb{F}^\perp \rightarrow \mathbb{F}^\perp$ defined by $\Phi/\mathbb{F} := (Q \circ \Phi)|_{\mathbb{F}^\perp}$. If the automorphism Φ is fiber-bunched (or strongly bunched) then so are $\Phi|_{\mathbb{F}}$ and Φ/\mathbb{F} .

Proof of Theorem 4.6. Let $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ be a strongly bunched automorphism. If Φ is irreducible then by Theorem 3.7 it is spannable, and by Proposition 4.1 it is relatively product bounded, hence our claim holds with $d' = 0$. In particular, the theorem holds when $d = 1$. Now suppose Φ is reducible, that is, there exists a θ -Hölder Φ -invariant nontrivial subbundle $\mathbb{F} \subset \mathbb{E}$. By induction on dimension, we can assume that the theorem holds for the restricted automorphism $\Phi|_{\mathbb{F}}$ and the quotient automorphism Φ/\mathbb{F} , that is, there are nonnegative integers $d_1 < \dim \mathbb{F}$ and $d_2 < d - \dim \mathbb{F}$ such that:

$$\|(\Phi|_{\mathbb{F}})_x^n\| = O\left(n^{d_1} e^{n\beta(\Phi|_{\mathbb{F}})}\right) \quad \text{and} \quad \|(\Phi/\mathbb{F})_x^n\| = O\left(n^{d_2} e^{n\beta(\Phi/\mathbb{F})}\right). \tag{4.8}$$

Note that, by the definitions of the automorphisms $\Phi|_{\mathbb{F}}$ and Φ/\mathbb{F} ,

$$\max\{\|(\Phi|_{\mathbb{F}})_x^n\|, \|(\Phi/\mathbb{F})_x^n\|\} \leq \|\Phi_x^n\| \quad \text{for all } x \in X \text{ and } n \geq 0,$$

and therefore

$$\max\{\beta(\Phi|_{\mathbb{F}}), \beta(\Phi/\mathbb{F})\} \leq \beta(\Phi). \tag{4.9}$$

Letting P and Q be the orthogonal projections (4.7), note the identity:

$$\Phi_x = (\Phi|_{\mathbb{F}})_x \circ P_x + P_{T^j x} \circ \Phi_x \circ Q_x + (\Phi/\mathbb{F})_x \circ Q_x.$$

More generally, for every $n \geq 1$, we have:

$$\Phi_x^n = (\Phi|_{\mathbb{F}})_x^n \circ P_x + \left[\sum_{j=0}^{n-1} (\Phi|_{\mathbb{F}})_{T^{j+1}x}^{n-j-1} \circ P_{T^{j+1}x} \circ \Phi_{T^j x} \circ (\Phi/\mathbb{F})_x^j \circ Q_x \right] + (\Phi/\mathbb{F})_x^n \circ Q_x, \tag{4.10}$$

which can be checked by induction. Using the bounds (4.8), it follows that:

$$\|\Phi_x^n\| = O\left(n^{d_1+d_2+1} e^{n \max\{\beta(\Phi|_{\mathbb{F}}), \beta(\Phi/\mathbb{F})\}}\right).$$

Noting that $d_1 + d_2 + 1 < d$ and recalling (4.9), we obtain the desired polynomial bound. \square

Incidentally, note that (4.10) implies that (4.9) is an equality, that is:

$$\beta(\Phi) = \max \{ \beta(\Phi|_{\mathbb{F}}), \beta(\Phi/\mathbb{F}) \}. \quad (4.11)$$

Actually, a more general fact holds: for any T -invariant ergodic probability measure μ ,

$$\chi_1(\Phi, \mu) = \max \{ \chi_1(\Phi|_{\mathbb{F}}, \mu), \chi_1(\Phi/\mathbb{F}, \mu) \}. \quad (4.12)$$

We will use this fact in Subsection 7.1. We were not able to find a precise reference for it, but it follows easily from the identity (4.10) together with an estimate such as [13, Lemma 12].

5 Construction of extremal norms

5.1 Extremal norms for spannable automorphisms

In this subsection we state and prove the central result of this paper, Theorem 5.2 below. Let us present a simple consequence first:

Corollary 5.1. *Let T be a transitive hyperbolic homeomorphism. Let Φ be a strongly bunched irreducible automorphism covering T . Then Φ admits an extremal norm.*

Here is the full statement of our result on extremal norms. Let θ_u be the exponent provided by applying Proposition 2.10 to Φ^{-1} .

Theorem 5.2. *Every spannable automorphism $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$ admits an extremal norm $\|\cdot\|$, which has the following additional properties:*

- (a) *there exists $C_4 > 1$ such that for every $u \in \mathbb{E}$,*

$$C_4^{-1} \|u\| \leq \|u\| \leq C_4 \|u\|; \quad (5.1)$$

- (b) *$\|\cdot\|$ is θ_u -Hölder, that is, there is a constant $C_5 > 0$ such that for all $x, x' \in X$,*

$$\left| \|I_{x' \leftarrow x}\| - 1 \right| \leq C_5 d(x, x')^{\theta_u}. \quad (5.2)$$

- (c) *$\|\cdot\|$ is $\bar{\theta}$ -Hölder along unstable sets with $\bar{\theta} := \max\{\theta, 1\}$, that is, there is a constant $C_6 > 0$ such that for all $x \in X$ and $x' \in W_{\varepsilon_0}^u(x)$,*

$$\left| \|H_{x' \leftarrow x}^u\| - 1 \right| \leq C_6 d(x, x')^{\bar{\theta}}; \quad (5.3)$$

Furthermore, for every sufficiently C^0 -small perturbation of the automorphism Φ in the set $\text{Aut}_K^\theta(\mathbb{E}, T)$, we can find an extremal norm that satisfies the properties above with the same constants θ_u, C_4, C_5, C_6 .

Combining the theorem above with Theorem 3.7 we immediately obtain Corollary 5.1.

Note that part (b) of the statement of Theorem 5.2 is compatible with the characterization of Hölderness of a norm given by Proposition 2.3. In summary, our extremal norm is Hölder, but perhaps with a smaller Hölder exponent than the original Φ .⁹

⁹A similar loss of exponent also appears in the first version of Mañé Lemma for Anosov diffeomorphisms, obtained by Lopes and Thieullen [51]. Later, Bousch [23] obtained a stronger Mañé Lemma without loss of exponent. However, it is unclear whether Bousch's strategy can be applied in our setting.

Nevertheless, part (c) says that the norm is more regular along unstable sets: there is no loss of exponent, and if $\theta < 1$ there is a gain.

Concerning the final part of the statement of Theorem 5.2, recall from Corollary 3.10 that the set of spannable automorphisms is a C^0 -open subset of the set $\text{Aut}_K^\theta(\mathbb{E}, T)$. So the theorem also says that our extremal norms vary in a bounded way if the automorphism is perturbed; this is useful to certain applications (see Subsection 5.2).

Before commencing the actual proof, let us establish an auxiliary fact:

Lemma 5.3. *Let $0 < a < 1$. Let $\bar{\theta} := \max\{\theta, 1\}$. Then there exists a $\bar{\theta}$ -Hölder function $\zeta: X \times X \rightarrow [0, 1]$ such that:*

$$\begin{aligned} \zeta(x, y) &= 1 & \text{if } d(x, y) \leq a\varepsilon_1; \\ \zeta(x, y) &= 0 & \text{if } d(x, y) \geq \varepsilon_1. \end{aligned}$$

Proof. If $\theta \leq 1$, let $f: [0, +\infty) \rightarrow [0, 1]$ be a non-increasing smooth function such that $f(a\varepsilon_1) = 1$ and $f(\varepsilon_1) = 0$. Then the function $\zeta(x, y) := f(d(x, y))$ meets our requirements.

If $\theta > 1$ then the existence of ζ is an immediate consequence of the fact that the algebra of θ -Hölder functions on $X \times X$ is normal (Lemma A.1). \square

Proof of Theorem 5.2. As in (4.3), let $a := \exp(-\min \lambda_u) \in (0, 1)$. Let ζ be given by Lemma 5.3. For each $u \in \mathbb{E}$, let

$$\|u\| := \limsup_{n \rightarrow \infty} e^{-\beta(\Phi)^n} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u)} \zeta(\pi(u), \pi(v)) \|\Phi^n(v)\|. \quad (5.4)$$

We will check that formula (5.4) defines an extremal norm with the additional properties stated in Theorem 5.2. To simplify writing, we assume from now on that $\beta(\Phi) = 0$.

Since $0 \leq \zeta \leq 1$ and Φ is relatively product bounded (thanks to Proposition 4.1), the quantity (5.4) is always finite, and therefore defines a seminorm on each fiber of \mathbb{E} .

Take arbitrary nonzero $u \in \mathbb{E}$. Since $\zeta(x, y) = 1$ whenever $y \in W_{a\varepsilon_1}^u(x)$, we have:

$$\limsup_{n \rightarrow \infty} \sup_{v \in \mathbb{W}_{a\varepsilon_1}^u(u)} \|\Phi^n(v)\| \leq \|u\|.$$

Recalling the hyperbolicity property (4.4), we have $\mathbb{W}_{a\varepsilon_1}^u(u) \supseteq \Phi^{-1}(\mathbb{W}_{\varepsilon_1}^u(\Phi(u)))$, and so:

$$\limsup_{n \rightarrow \infty} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(\Phi(u))} \|\Phi^n(v)\| \leq \|u\|. \quad (5.5)$$

So, letting $\tilde{u} := \Phi(u)$, we have:

$$\frac{1}{\|\tilde{u}\|} \limsup_{n \rightarrow \infty} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(\tilde{u})} \|\Phi^n(v)\| \leq \frac{\|u\|}{\|\tilde{u}\|} \leq K \frac{\|u\|}{\|u\|},$$

using the bound (4.1). This allows us to apply Lemma 4.5 to \tilde{u} and conclude that, for some constant $C_3 > 1$ that only depends on Φ ,

$$\limsup_{n \rightarrow \infty} \sup_{y \in X} \|\Phi_y^n\| \leq C_3 K \frac{\|u\|}{\|u\|}, \quad \text{for all } u \in \mathbb{E}^\times. \quad (5.6)$$

The left-hand side is at least 1; indeed by (1.4), for every $n > 0$ there exists $y \in X$ such that $\|\Phi_y^n\| \geq e^{n\beta(\Phi)} = 1$. Therefore:

$$\|u\| \geq K^{-1}C_3^{-1}\|u\|, \quad \text{for all } u \in \mathbb{E}. \quad (5.7)$$

In particular, the seminorm $\|\cdot\|$ is actually a norm.

Since $0 \leq \zeta \leq 1$ and $\zeta(x, y) = 0$ whenever $y \in W_{\varepsilon_1}^u(x)$, inequality (5.5) implies:

$$\|\Phi(u)\| \leq \|u\|, \quad \text{for all } u \in \mathbb{E},$$

that is, $\|\cdot\|$ is an extremal norm.

Now consider the vector $u_* \in \mathbb{E}^\times$ given by Lemma 4.4. It satisfies $\|u_*\| \leq C_2\|u_*\|$, where $C_2 > 1$ is a constant depending only on Φ . Applying (5.6) to this vector we obtain:

$$\limsup_{n \rightarrow \infty} \sup_{y \in X} \|\Phi_y^n\| \leq KC_2C_3.$$

Therefore, for all $u \in \mathbb{E}$,

$$\|u\| \leq \limsup_{n \rightarrow \infty} \sup_{v \in \mathbb{W}_{\varepsilon_1}^u} \|\Phi^n(v)\| \leq KC_2C_3 \sup_{v \in \mathbb{W}_{\varepsilon_1}^u} \|v\| \leq KC_1C_2C_3\|u\|,$$

where $C_1 > 1$ is the constant from (4.2). So, letting $C_4 := KC_1C_2C_3$ and recalling the lower bound (5.7), we obtain (5.1): the extremal norm is uniformly comparable to the original norm by a factor C_4 that works not only for Φ but also for its C^0 perturbations in $\text{Aut}_K^\theta(\mathbb{E}, T)$.

Before proving regularity properties of the extremal norm, let us establish a few auxiliary facts. For all $u \in \mathbb{E}$, $v \in \mathbb{W}_{\varepsilon_1}^u$, and $n \geq 0$, using (5.1), extremality, and (4.2), we obtain:

$$\|\Phi^n(v)\| \leq C_4\|\Phi^n(v)\| \leq C_4\|v\| \leq C_4^2\|v\| \leq C_1C_4^2\|u\| \leq C_1C_4^3\|u\|. \quad (5.8)$$

Fix a constant $b < 1$ sufficiently close to 1 so that:

$$d(x, y) \geq b\varepsilon_1 \quad \Rightarrow \quad \zeta(x, y) < \frac{1}{2}C_1^{-1}C_4^{-3}.$$

Then:

$$v \in \mathbb{W}_{\varepsilon_1}^u \setminus \mathbb{W}_{b\varepsilon_1}^u \quad \Rightarrow \quad \zeta(\pi(u), \pi(v))\|\Phi^n(v)\| \leq \frac{1}{2}\|u\|.$$

So vectors v outside $\mathbb{W}_{b\varepsilon_1}^u$ do not contribute in formula (5.4), which therefore can be rewritten as:

$$\|u\| := \limsup_{n \rightarrow \infty} \sup_{v \in \mathbb{W}_{b\varepsilon_1}^u} \zeta(\pi(u), \pi(v))\|\Phi^n(v)\|. \quad (5.9)$$

We will prove property (c) first, and use it later in the proof of property (b). In order to simplify writing, let us use the O notation to denote constants that depend only on Φ and can be taken uniform on a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$. In order to prove property (c), we need to show:

$$x \in X, x' \in W_{\varepsilon_0}^u(x), u \in \mathbb{E}_x, \|u\| = 1 \quad \Rightarrow \quad \|\|H_{x' \leftarrow x}^u\| - \|u\|\| = O(d(x, x')^{\bar{\theta}}).$$

It is sufficient to consider x' very close to x , so assume $d(x', x) \leq (1-b)\varepsilon_1$. Fix a unit vector $u \in \mathbb{E}_x$ and let $u' := H_{x' \leftarrow x}^u(u)$. For all $v \in \mathbb{W}_{\varepsilon_1}^u(u)$, and $n \geq 0$, using (5.8) and the fact that ζ is $\bar{\theta}$ -Hölder, we estimate:

$$\left| \zeta(x', \pi(v)) \|\Phi^n(v)\| - \zeta(x, \pi(v)) \|\Phi^n(v)\| \right| = O(d(x, x')^{\bar{\theta}}).$$

Noting that $\mathbb{W}_{b\varepsilon_1}^u(u') \subseteq \mathbb{W}_{\varepsilon_1}^u(u)$, we have:

$$\sup_{v \in \mathbb{W}_{b\varepsilon_1}^u(u')} \zeta(x', \pi(v)) \|\Phi^n(v)\| \leq \sup_{v \in \mathbb{W}_{\varepsilon_1}^u(u)} \zeta(x, \pi(v)) \|\Phi^n(v)\| + O(d(x, x')^{\bar{\theta}}) \|u\|.$$

Using (5.9) and (5.4), we obtain:

$$\|u'\| \leq \|u\| + O(d(x, x')^{\bar{\theta}}).$$

On the other hand, using $\mathbb{W}_{b\varepsilon_1}^u(u) \subseteq \mathbb{W}_{\varepsilon_1}^u(u')$, a similar argument shows that:

$$\|u\| \leq \|u'\| + O(d(x, x')^{\bar{\theta}}).$$

This completes the proof of property (c).

We are left to check θ_u -Hölderiness of the norm, that is, property (b). In fact, it is sufficient to prove θ_u -Hölderiness along stable sets, that is:

$$x \in X, x' \in W_{\varepsilon_0}^s(x), u \in \mathbb{E}_x, \|u\| = 1 \Rightarrow \|\|H_{x' \leftarrow x}^s(u)\| - \|u\|\| = O(d(x, x')^{\theta_u}). \quad (5.10)$$

Since we have already proven $\bar{\theta}$ -Hölderiness of the norm along unstable sets, and $\theta_u \leq \bar{\theta}$, property (b) will follow from (5.10): just mimic the proof of (b) \Rightarrow (a) in Corollary 3.5.

In order to prove (5.10), it is sufficient to consider x' very close to x . Fix a unit vector $u \in \mathbb{E}_x$ and let $u' := H_{x' \leftarrow x}^s(u)$. Consider arbitrary $v \in \mathbb{W}_{\varepsilon_1}^u(u)$, and write $y := \pi(v)$. Since $d(x', y) \leq \varepsilon_1 + d(x, x') < 2\varepsilon_1$, the bracket $[x', y] =: y'$ is well-defined. Let also $w := H_{y' \leftarrow y}^s(v)$, and $v' := H_{y' \leftarrow x'}^u(u')$: see Fig. 4.




Figure 4: Proof of property (b).

Then for each $n \geq 0$ we estimate:

$$\begin{aligned} |\zeta(x, y)\|\Phi^n(v)\| - \zeta(x', y')\|\Phi^n(v')\| &\leq \textcircled{1} + \textcircled{2} + \textcircled{3}, \quad \text{where} \\ \textcircled{1} &:= |\zeta(x, y) - \zeta(x', y')|\|\Phi^n(v)\|, \\ \textcircled{2} &:= \|\|\Phi^n(v)\| - \|\Phi^n(w)\|\|, \\ \textcircled{3} &:= \|\Phi^n(w - v')\|. \end{aligned}$$

In order to estimate $\textcircled{1}$, recall that by (5.8), $\|\Phi^n(v)\| = O(1)$. On the other hand, by Hölder-continuity of ζ ,

$$|\zeta(x, y) - \zeta(x', y')| = O(\max\{d(x, x')^{\bar{\theta}}, d(y, y')^{\bar{\theta}}\}).$$

Using Proposition 2.6 for T^{-1} , we have $d(y, y') = O(d(x, x')^{\kappa_u})$, where the exponent κ_u is at most 1. So:

$$|\zeta(x, y) - \zeta(x', y')| = O(d(x, x')^{\kappa_u \bar{\theta}}).$$

Note that $\kappa_u \bar{\theta} \geq \kappa_u \theta \geq \theta_u$, so the weaker estimate $\textcircled{1} = O(d(x, x')^{\theta_u})$ holds.

The next term is estimated as follows:

$$\begin{aligned} \textcircled{2} &= \|\|H_{T^n y' \leftarrow T^n y}^s \Phi^n(v)\| - \|\Phi^n(v)\|\| \\ &\leq \|\|H_{T^n y' \leftarrow T^n y}^s - I_{T^n y' \leftarrow T^n y}\| + \|I_{T^n y' \leftarrow T^n y} - 1\| \|\Phi^n(v)\|. \end{aligned}$$

Since $d(T^n y', T^n y) = o(1)$ (i.e., it tends to 0 as $n \rightarrow \infty$), using regularity of holonomies (2.14) and of the transport maps (Proposition 2.3) together with product boundedness (5.8), we conclude that $\textcircled{2} = o(1)$.

In order to estimate the last term, we use Proposition 2.10 applied to T^{-1} :

$$\textcircled{3} = O(\|v' - w\|) = O(\|H_{y' \leftarrow x'}^u \circ H_{x' \leftarrow x}^s - H_{y' \leftarrow y}^s \circ H_{y \leftarrow x}^u\|) = O(d(x, x')^{\theta_u}),$$

Summing the three estimates,

$$|\zeta(x, y)\|\Phi^n(v)\| - \zeta(x', y')\|\Phi^n(v')\| = O(d(x, x')^{\theta_u}) + o(1). \quad (5.11)$$

As in the proof of the previous property (c), we need to use the cutoff property (5.9) to conclude. If we are careful enough to take $d(x, x')$ sufficiently small then $d(y, y') = O(d(x, x')^{\kappa_u})$ is also small and therefore the following two implications are correct:

$$\begin{aligned} d(x, y) \leq b\varepsilon_1 &\Rightarrow d(x', y') \leq \varepsilon_1, \\ d(x', y') \leq b\varepsilon_1 &\Rightarrow d(x, y) \leq \varepsilon_1. \end{aligned}$$

That is,

$$\begin{aligned} v \in \mathbb{W}_{b\varepsilon_1}^u(u) &\Rightarrow v' \in \mathbb{W}_{\varepsilon_1}^u(u'), \\ v' \in \mathbb{W}_{b\varepsilon_1}^u(u') &\Rightarrow v \in \mathbb{W}_{\varepsilon_1}^u(u). \end{aligned}$$

Then, using (5.11), (5.9), and (5.4), we obtain:

$$\|\|u\| - \|u'\|\| = O(d(x, x')^{\theta_u}),$$

proving (5.10) and the theorem. \square

5.2 Application: Lipschitz continuity of the maximal Lyapunov exponent

As a simple application of Theorem 5.2, let us establish a local regularity result for the maximal Lyapunov exponent. A similar property for the joint spectral radius (under the assumption of irreducibility) was established by Wirth [74, Corol. 4.2], also using extremal norms; see also [49] for a more precise result.

Let \mathcal{S}_K denote the set of spannable automorphisms in $\text{Aut}_K^\theta(\mathbb{E}, T)$, which by Corollary 3.10 is relatively C^0 -open.

Proposition 5.4. *The maximal Lyapunov exponent $\beta(\cdot)$ is a locally Lipschitz function on the set \mathcal{S}_K , with respect to the C^0 -norm (2.3).*

Proof. Let $\Phi \in \mathcal{S}_K$. Let $\mathcal{U} \subset \mathcal{S}_K$ be a C^0 -neighborhood of Φ where Theorem 5.2 applies with uniform constants. Take any two automorphisms Φ_1 and Φ_2 in \mathcal{U} , and let $\|\cdot\|_1$ and $\|\cdot\|_2$ be the corresponding extremal norms provided by Theorem 5.2. Then, using the bound (5.1), we obtain:

$$\begin{aligned} e^{\beta(\Phi_2)} &\leq \sup_x \|\Phi_{2x}\|_1 \leq \sup_x \|\Phi_{1x}\|_1 + \sup_x \|\Phi_{1x} - \Phi_{2x}\|_1 \\ &\leq e^{\beta(\Phi_1)} + C_4 \sup_x \|\Phi_{1x} - \Phi_{2x}\| = e^{\beta(\Phi_1)} + C_4 \|\Phi_1 - \Phi_2\|_0, \end{aligned}$$

where $\|\cdot\|_0$ is the C^0 -norm (2.3). By symmetry, we obtain $|e^{\beta(\Phi_1)} - e^{\beta(\Phi_2)}| \leq C_4 \|\Phi_1 - \Phi_2\|_0$. This shows that the function $e^{\beta(\cdot)}$ is Lipschitz on the neighborhood \mathcal{U} , with respect to the C^0 -norm. Since the function $\beta(\cdot)$ is uniformly bounded on \mathcal{U} (and in the whole set $\text{Aut}_K^\theta(\mathbb{E}, T)$, in fact), it is Lipschitz as well. \square

Remark 5.5. For reducible automorphisms, it is clear that β is not locally Lipschitz: see e.g. [74, p. 27]. Nevertheless, β is continuous on the whole space of θ -Hölder automorphisms: indeed, upper semicontinuity is automatic from (1.4), while lower semicontinuity follows by a theorem of Kalinin [45, Theorem 1.4] that allows one to approximate β by the Lyapunov exponents of periodic orbits. Let us also remark that if T is no longer hyperbolic, then β becomes discontinuous with respect to the C^0 topology. For example, the cocycle from Example 4.2 can be C^0 -perturbed so that β drops to 0, as it follows e.g. from the result of [6].

5.3 Barabanov-like norms for linear cocycles over shifts

Let us consider subshifts of finite type, that is, X is the set of two-sided sequences $(x_n)_{n \in \mathbb{Z}}$ in an alphabet $\{0, 1, \dots, N-1\}$ whose neighboring pairs are those allowed by a fixed 0-1 matrix, and $T: X \rightarrow X$ is the (left) shift map. As usual, we consider on X the (ultra)metric

$$d(x, y) := e^{-\lambda k} \quad \text{where } k = \min\{|n|; x_n \neq y_n\}, \quad (5.12)$$

and $\lambda > 0$ is a fixed parameter. Then T is a hyperbolic homeomorphism. Indeed letting $\varepsilon_0 := e^{-\lambda}$, the corresponding local unstable and stable sets at $x = (x_n) \in X$ are:

$$\begin{aligned} W_{\text{loc}}^u(x) &:= W_{\varepsilon_0}^u(x) = \{(y_n) \in X; y_n = x_n \text{ for all } n \leq 0\}, \\ W_{\text{loc}}^s(x) &:= W_{\varepsilon_0}^s(x) = \{(y_n) \in X; y_n = x_n \text{ for all } n \geq 0\}. \end{aligned}$$

and so hyperbolicity property (b) holds with $\lambda_u = \lambda_s = \lambda$, property (c) holds with $2\varepsilon_1 = \varepsilon_0$, and property (d) holds with $C = 1$. Also note that Proposition 2.6 holds with $\kappa_s = 1 = C$.

We will consider θ -Hölder automorphisms covering the subshift $T: X \rightarrow X$. Since X is a Cantor set, every θ -Hölder vector bundle is trivial, i.e., θ -Hölder isomorphic to the product bundle. So we are actually dealing with θ -Hölder linear cocycles; nevertheless, we will keep using the vector bundle terminology.

Example 5.6. As mentioned in Subsection 1.2, the *one-step cocycle* determined by a N -tuple of matrices $(A_0, \dots, A_{N-1}) \in GL(d, \mathbb{R})^N$ is the pair (T, F) where T is the full shift on N symbols and $F: X \rightarrow GL(d, \mathbb{R})$ is given by $F(x) := A_{x_0}$. Let Φ the associated automorphism (1.1). Then $e^{\beta(\Phi)}$ is joint spectral radius of the set $\{A_0, \dots, A_{N-1}\}$. Since F is locally constant, it is θ -Hölder for any $\theta \in (0, +\infty)$. Choosing θ large enough, the automorphism Φ becomes fiber-bunched. (Alternatively, we can take $\theta = 1$, say, and then take the parameter λ large enough.) The holonomies are locally trivial:

$$\star \in \{u, s\}, y \in W_{\text{loc}}^*(x) \Rightarrow H_{y \leftarrow x}^* = \text{id}. \quad (5.13)$$

A useful generalization of one-step cocycles are the *sofic cocycles* from [17, § 5.1]; the same concept appears in [58] under the terminology *constrained switching systems*.

Let us present an improved version of Theorem 5.2 for subshifts of finite type. We obtain an extremal norm with an additional Barabanov-like property: given any vector $u \in \mathbb{E}$, there always exists a vector in its local unstable set $\mathbb{W}_{\text{loc}}^u(u) := \mathbb{W}_{\varepsilon_0}^u(u)$ whose expansion factor in a single iterate equals the maximum asymptotic expansion rate $e^{\beta(\Phi)}$. Furthermore, the norm is invariant under local unstable holonomies. Therefore, for the case of one-step cocycles, we reobtain the Barabanov property (1.8).

Theorem 5.7. *Let T be a two-sided subshift of finite type. Let \mathbb{E} be a d -dimensional θ -Hölder vector bundle. Let Φ be a spannable automorphism of \mathbb{E} covering T . Then*

$$\|u\| := \limsup_{n \rightarrow \infty} e^{-\beta(\Phi)n} \sup_{v \in \mathbb{W}_{\text{loc}}^u(u)} \|\Phi^n(v)\| \quad (5.14)$$

is a well-defined Barabanov norm on \mathbb{E} , namely, an extremal norm satisfying, for all $u \in \mathbb{E}$,

- (1) local H^u -invariance: $\|u\| = \|v\|$ for all $v \in \mathbb{W}_{\text{loc}}^u(u)$;
- (2) calibration: there exists $v \in \mathbb{W}_{\text{loc}}^u(u)$ such that $\|\Phi(v)\| = e^{\beta(\Phi)}\|v\|$.

Furthermore, $\|\cdot\|$ satisfies the other properties stated in Theorem 5.2.

Let us comment on the hypotheses. Given $\theta > 0$, Lemma 3.6 holds with the value $\eta_0 = \theta/3$; indeed, this follows from formula (A.24), recalling that $\lambda_u = \lambda_s = \lambda$, and noting that we can also take $\Lambda_u = \lambda$ in (A.7). Therefore θ -Hölder automorphism $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ covering T is strongly bunched if Φ is fiber-bunched and the fibers of \mathbb{E} have dimension $d \leq 2$, or Φ is a $(\theta/3, \theta)$ -bunched. In that case, it follows from Theorem 3.7 that Φ is spannable, provided it is irreducible and T is transitive.

Proof of Theorem 5.7. Since X is a Cantor set, we can simplify the construction in Theorem 5.2 and dispense with the bump function ζ from Lemma 5.3. Ultimately, we can replace the definition (5.4) of the extremal norm by the simpler formula (5.14). Note that in the latter formula we maximize over $\mathbb{W}_{\text{loc}}^u(u) := \mathbb{W}_{\varepsilon_0}^u(u)$ instead of $\mathbb{W}_{\varepsilon_1}^u(u) = \mathbb{W}_{\varepsilon_0/2}^u(u)$; this is possible because the metric on X is an ultrametric. It is

straightforward to check that the proof in Theorem 5.2 applies, with simplifications. It is immediate from its definition that the norm satisfies local H^u -invariance, that is property (1), which of course subsumes property (c) from Theorem 5.2.

We only left to check the calibration property (2). Given $u \in \mathbb{E}$, by definition, there exist sequences $n_i \nearrow \infty$ and $v_i \in \mathbb{W}_{\text{loc}}^u(u)$ such that

$$\|u\| = \lim_{i \rightarrow \infty} e^{-\beta(\Phi)n_i} \|\Phi^{n_i} v_i\|.$$

Denote $y_i = \pi(v_i)$. By compactness, we may suppose that $y_i \rightarrow y \in W_{\text{loc}}^u(x)$ and $v_i \rightarrow v \in \mathbb{W}_{\text{loc}}^u(u)$. For i large enough, $y'_i := T(y_i) \in W_{\text{loc}}^u(T(y))$. We can assume that this property is true for all i . Let $v'_i := \Phi(v_i)$. Thus, we have:

$$\begin{aligned} \|\Phi(v)\| &= \limsup_{n \rightarrow \infty} e^{-\beta(\Phi)n} \sup_{v' \in \mathbb{W}_{\text{loc}}^u(\Phi(v))} \|\Phi^n v'\| \\ &\geq \limsup_{i \rightarrow \infty} e^{-(n_i-1)\beta(\Phi)} \|\Phi^{n_i-1} v'_i\| \\ &= \lim_{i \rightarrow \infty} e^{-(n_i-1)\beta(\Phi)} \|\Phi^{n_i} v_i\| = e^{\beta(\Phi)} \|u\| = e^{\beta(\Phi)} \|v\|. \end{aligned}$$

By extremality, the inequality is actually an equality. This proves calibration. \square

6 Mather sets

In traditional ergodic optimization, that is, the optimization of Birkhoff averages (see [42, 43, 37]), a *maximizing set* is a closed subset such that an invariant probability is maximizing if and only if its support lies on this subset. The existence of such sets is guaranteed in any context where a Mañé Lemma holds. The *Mather set* is the smallest maximizing set: it is defined as the union of the supports of all maximizing measures. The nomenclature is borrowed from Lagrangian dynamics, where the concept of minimizing measures proved to be a useful generalization of the notion of action minimizing orbits: see [52]. There are other canonical maximizing sets, such as the Aubry set: see e.g. [37]. Some of these concepts have been already considered in the optimization of the top Lyapunov exponent: see [55, 38].

In this section, we are going to study some notions of Mather sets for continuous vector bundle automorphisms, not necessarily with any Hölder or hyperbolicity structures. Our approach was profoundly influenced by some works of Morris [54, 55].

Throughout the section, we assume that X is a compact metric space, $T: X \rightarrow X$ is a homeomorphism, \mathbb{E} is a d -dimensional vector bundle over X , and Φ is an automorphism covering T .

6.1 The first Mather set

By a *Lyapunov maximizing measure* we mean any T -invariant probability μ whose upper Lyapunov exponent $\chi_1(\Phi, \mu)$ equals $\beta(\Phi)$. Following Morris [55], we define the (*first*) *Mather set* $M(\Phi) \subseteq X$ as the union of the supports of all Lyapunov maximizing measures.

Proposition 6.1. *The Mather set $M(\Phi)$ is the support of some Lyapunov maximizing measure and, a fortiori, it is a nonempty, compact, and T -invariant set.*

Proof. The argument is quite standard, but we add it for completeness. For simplicity, write $M = M(\Phi)$. As explained at the introduction, at least one Lyapunov maximizing measure exists, so $M \neq \emptyset$. Given a countable basis $\{B_j\}_{j \in \mathbb{N}}$ for the topology of X , consider the subset of indices $J := \{j \in \mathbb{N} ; B_j \cap M \neq \emptyset\}$. For each $j \in J$, we assign a Lyapunov maximizing measure μ_j such that $B_j \cap \text{supp } \mu_j \neq \emptyset$, which means $\mu_j(B_j) > 0$. Define then $\mu := \sum_{j \in J} \alpha_j \mu_j$, where $\alpha_j > 0$ and $\sum_{j \in J} \alpha_j = 1$. As a convex combination of Lyapunov maximizing measures, μ is also Lyapunov maximizing. Now consider an arbitrary $k \in \mathbb{N}$ such that $B_k \cap \text{supp } \mu = \emptyset$. Then for all $j \in J$ we have $B_k \cap \text{supp } \mu_j = \emptyset$ and therefore $\mu_j(B_k) = 0$. This implies that $k \notin J$, and so $B_k \cap M = \emptyset$. We have shown that $X \setminus \text{supp } \mu \subseteq X \setminus M$, which yields $M(\Phi) = \text{supp } \mu$. \square

Proposition 6.2. *If Φ admits an extremal norm $\|\cdot\|$ then*

$$\|\Phi_x^n\| = e^{n\beta(\Phi)}, \quad \forall x \in M(\Phi), \forall n \geq 1,$$

and, in particular, every T -invariant probability measure whose support is contained in $M(\Phi)$ is Lyapunov-maximizing.

Proof. By extremality,

$$f_n(x) := \frac{1}{n} \log \|\Phi_x^n\| \leq \beta(\Phi), \quad \forall x \in X, \forall n \geq 1.$$

On the other hand, if μ is the Lyapunov maximizing measure with $\text{supp } \mu = M(\Phi)$ constructed in Proposition 6.1 then, by Kingman's subadditive theorem, $\inf_n \frac{1}{n} \int f_n d\mu = \beta(\Phi)$. It follows that $\frac{1}{n} \int f_n d\mu = \beta(\Phi)$ for each $n \geq 1$. Since the functions f_n are continuous, they must be identically equal to $\beta(\Phi)$ over $\text{supp } \mu = M(\Phi)$, as we wanted to show. \square

6.2 Mather sets of higher index

Up to here we have only considered the first Lyapunov exponent χ_1 , but now we will need to consider the full Lyapunov spectrum. Let us recall the definitions and main properties, referring to [5] for details.

If μ is a T -invariant probability measure then the *Lyapunov exponents* of the automorphism Φ with respect to μ are the numbers

$$\chi_1(\Phi, \mu) \geq \chi_2(\Phi, \mu) \geq \dots \geq \chi_d(\Phi, \mu) \tag{6.1}$$

uniquely defined by the following equations: for every $p \in \llbracket 1, d \rrbracket$,

$$\sum_{i=1}^p \chi_i(\Phi, \mu) = \chi_1(\Lambda^p \Phi, \mu), \tag{6.2}$$

where the automorphism $\Lambda^p \Phi: \Lambda^p \mathbb{E} \rightarrow \Lambda^p \mathbb{E}$ is the p -fold exterior power of the automorphism $\Phi: \mathbb{E} \rightarrow \mathbb{E}$.

Suppose μ is ergodic, and that λ is a Lyapunov exponent with respect to μ of *multiplicity* k , in the sense that it appears k times in the list (6.1). Then Oseledets' theorem says that for μ -a.e. $x \in X$, there exists a k -dimensional subspace $\mathbb{O}_x(\lambda)$ of the fiber \mathbb{E}_x , called a *Oseledets space*, such that:

$$u \in \mathbb{O}_x(\lambda) \setminus \{0\} \quad \text{if and only if} \quad \lim_{n \rightarrow \pm\infty} \frac{1}{n} \log \|\Phi^n(u)\| = \lambda.$$

Moreover, Oseledets spaces form a splitting of \mathbb{E}_x , depend measurably on the point x , and are Φ -equivariant.

We now consider other Mather sets that take multiplicity into account. Define a chain of sets

$$M(\Phi) = M_1(\Phi) \supseteq M_2(\Phi) \supseteq \cdots \supseteq M_d(\Phi) \quad (6.3)$$

as follows: the p -th Mather set $M_p(\Phi)$ is the union of the supports of all T -invariant probabilities μ whose p first Lyapunov exponents are all maximal, that is,

$$\chi_1(\Phi, \mu) = \chi_2(\Phi, \mu) = \cdots = \chi_p(\Phi, \mu) = \beta(\Phi).$$

Repeating the argument of the proof of Proposition 6.1, we see that if the set $M_p(\Phi)$ is nonempty then there exists a measure μ with p maximal Lyapunov exponents and whose support is exactly $M_p(\Phi)$; in particular, $M_p(\Phi)$ is compact and T -invariant.

The following properties follow immediately from the definition of Mather sets and relations (6.1) and (6.2):

Proposition 6.3. *For any $p \in \llbracket 1, d \rrbracket$, we have:*

$$\beta(\Lambda^p \Phi) \leq p\beta(\Phi) \quad \text{and} \quad M_p(\Phi) \subseteq M_1(\Lambda^p \Phi).$$

Furthermore, these two relations become equalities if and only if $M_p(\Phi) \neq \emptyset$.

Let us extend the chain of sets (6.3). Let $M_{d+1}(\Phi) := \emptyset$ and let $M_0(\Phi)$ be defined as the union of the supports of all T -invariant probability measures (which is also the support of one of them). So $M_0(\Phi)$ only depends on T , and is in fact the classical *minimal center of attraction* of T : see [69], [1, p. 164].

6.3 Dominated splittings over the Mather sets

We want to provide more information about the action of Φ on the fibers above the Mather sets, assuming the existence of an extremal norm. We will use the notion of dominated splitting, which is very useful in Differentiable Dynamics (see [19]). It appears in the celebrated book [41] as *relative pseudo hyperbolicity*. It also appears in ODE and Control Theory under the terminology *exponentially separated splitting* (see [31]), and is intimately related to the concept of Anosov representations in Geometric Group Theory (see [17]).

Let $Y \subseteq X$ be a nonempty T -invariant compact set, and let $\mathbb{E}_Y := \pi^{-1}(Y)$ be the restricted vector bundle. (Recall that $\pi: \mathbb{E} \rightarrow X$ denotes the bundle projection.) Let $\|\cdot\|$ be a Finsler norm on \mathbb{E} . Suppose that the bundle \mathbb{E}_Y splits as a direct sum $\mathbb{F} \oplus \mathbb{G}$ of two (continuous) subbundles whose fibers $\mathbb{F}_x, \mathbb{G}_x$ have constant dimensions, and are equivariant in the sense that $\Phi_x(\mathbb{F}_x) = \mathbb{F}_{Tx}$, $\Phi_x(\mathbb{G}_x) = \mathbb{G}_{Tx}$. We say that $\mathbb{F} \oplus \mathbb{G}$ is a *dominated splitting* with *dominating* bundle \mathbb{F} and *dominated* bundle \mathbb{G} if there are positive constants c and τ such that for each point $x \in Y$, if $u \in \mathbb{F}_x, v \in \mathbb{G}_x$ are unit vectors then

$$\|\Phi_x^n(v)\| \leq ce^{-\tau n} \|\Phi_x^n(u)\| \quad \text{for all } n \geq 0. \quad (6.4)$$

An equivalent definition is to say that there exists an *adapted norm* $\|\cdot\|$ for which relation (6.4) holds with $c = 1$ (and therefore only needs to be checked for $n = 1$): see [39]. Dominated splittings are unique given the dimensions: see [36, Prop. 2.2]. Continuity of the subbundles actually follows from the uniform estimates (6.4) and

therefore could be removed from the definition: see [36, Prop. 2.5]. Actually, if Φ is Hölder then the bundles of a dominated splitting are always Hölder (with a smaller exponent): see [36, Thm. 4.11]. Domination can be characterized in terms of existence of invariant cone fields: see [36, Thm. 2.6]. This implies strong robustness properties: see [36, Corol. 2.8].

We will use another criterion for the existence of dominated splittings, expressed in terms of singular values. Recall that if $L: E \rightarrow F$ is a linear map between d -dimensional inner product spaces, then the *singular values* $\sigma_1(L) \geq \dots \geq \sigma_d(L)$ are the eigenvalues of the symmetric operator $(L^*L)^{1/2}$. So $\sigma_1(L)$ coincides with the Euclidean operator norm $\|L\|$. Endowing the exterior power spaces with the induced inner products, the exterior powers of L have norm:

$$\|\wedge^p L\| = \sigma_1(L)\sigma_2(L)\cdots\sigma_p(L); \quad (6.5)$$

see e.g. [5, p. 120]. Another useful characterization of the singular values is:

$$\sigma_p(L) = \max_{V \in \mathcal{G}_p(E)} \min_{u \in V} \frac{\|Lu\|}{\|u\|}; \quad (6.6)$$

see e.g. [71, p. 68].

A theorem from [15] says that the domination is equivalent to a uniform exponential gap between singular values of the powers of Φ (computed with respect to a Riemannian norm fixed a priori). More precisely:

Theorem 6.4 (Bochi–Gourmelon). *The bundle \mathbb{E}_Y admits a dominated splitting with a dominating bundle of dimension p if and only if there exist positive constants c and τ such that*

$$\sigma_{p+1}(\Phi_x^n) \leq ce^{-\tau n} \sigma_p(\Phi_x^n) \quad \text{for all } x \in Y \text{ and } n \geq 0.$$

We now come back to the Mather sets:

Theorem 6.5. *Suppose Φ admits an extremal norm $\|\cdot\|$. Let $p \in \llbracket 1, d \rrbracket$. Suppose Y is a nonempty compact T -invariant set contained in $M_p(\Phi) \setminus M_{p+1}(\Phi)$. Then the restricted bundle \mathbb{E}_Y admits a dominated splitting $\mathbb{F} \oplus \mathbb{G}$ where the dominating bundle \mathbb{F} has fibers of dimension p and is calibrated in the sense that $\|\Phi(u)\| = e^{\beta(\Phi)}\|u\|$ for every $u \in \mathbb{F}$.*

In particular, if exactly one of the sets $M_p(\Phi) \setminus M_{p+1}(\Phi)$ is nonempty then we obtain a dominated splitting over the whole Mather set $M(\Phi)$.

Related results were previously obtained by Morris: [54, Theorem 2.1] produces a dominated splitting under the weaker assumption of relative product boundedness, but with the strong hypothesis that the set Y is minimal (i.e., all orbits in Y are dense). Assuming existence of an extremal norm, Morris also proves the calibration property of the dominating bundle in his Theorem 2.2.

For a complement to Theorem 6.5, see Proposition B.4 in Appendix B.2.

6.4 Proof of Theorem 6.5

Consider the set of vectors whose bi-infinite orbits under Φ are *calibrated* with respect to the extremal norm $\|\cdot\|$:

$$\mathbb{K} := \{u \in \mathbb{E} ; \|\Phi^n(u)\| = e^{n\beta(\Phi)}\|u\| \text{ for all } n \in \mathbb{Z}\}. \quad (6.7)$$

This is a closed, Φ -invariant subset of \mathbb{E} . Denote its fibers by $\mathbb{K}_x := \mathbb{E}_x \cap \mathbb{K}$.

Proposition 6.6. *There exists a T -invariant Borel set $R \subseteq M(\Phi)$ such that:*

- $\mu(M(\Phi) \setminus R) = 0$ for every T -invariant probability measure μ ;
- for all $x \in R$, the Oseledets space corresponding to the Lyapunov exponent $\beta(\Phi)$ exists and coincides with \mathbb{K}_x .

Proof. Without loss of generality, assume that $\beta(\Phi) = 0$. Let $R_0 \subseteq X$ be the Borel set of points that satisfy the conclusions of Oseledets theorem. For each $x \in R_0 \cap M(\Phi)$, the Oseledets space $\mathbb{O}_x = \mathbb{O}_x(0) \subseteq \mathbb{E}_x$ is well-defined and has positive dimension, say $p(x)$. Calibrated vectors have zero Lyapunov exponent, so $\mathbb{K}_x \subseteq \mathbb{O}_x$. Consider the unit balls on these Oseledets spaces, i.e., $\mathcal{B}_x := \{u \in \mathbb{O}_x ; \|u\| \leq 1\}$. Then $\Phi_x(\mathcal{B}_x) \subseteq \mathcal{B}_{Tx}$, and therefore the following function is non-positive:

$$\psi(x) := \log \frac{\text{vol}(\Phi_x(\mathcal{B}_x))}{\text{vol}(\mathcal{B}_{Tx})}.$$

Here vol means $p(x)$ -dimensional volume with respect to a fixed Riemannian norm on the bundle \mathbb{E} ; of course, the choice of this metric does not affect the function ψ . Then ψ is cohomologous to the function $\tilde{\psi}(x) := \log \det \Phi(x)|_{\mathbb{O}_x}$, where \det denotes the signless determinant induced by the Riemannian metric (see [5, p. 213]); indeed $\psi = \tilde{\psi} + \varphi - \varphi \circ T$ where $\varphi(x) := \log \text{vol}(\mathcal{B}_x)$. All these functions are Borel measurable and bounded. Let μ be any T -invariant probability measure supported on $M(\Phi)$, that is, any Lyapunov maximizing measure. As a consequence of Oseledets theorem, we have $\int \tilde{\psi} d\mu = 0$ (see [5, p. 214]). Since ψ is cohomologous to $\tilde{\psi}$, its integral is zero as well. But $\psi \leq 0$, so $\psi = 0$ μ -a.e. Let $R_1 := \{x \in R_0 \cap M(\Phi) ; \psi(x) = 0\}$ and $R := \bigcap_{n \in \mathbb{Z}} T^{-n}(R_1)$; then $\mu(R) = 1$. Noting that $\psi(x) = 0$ if and only if $\Phi(x)|_{\mathbb{O}_x}$ preserves $\|\cdot\|$, we see that if $x \in R$ then $\mathbb{O}_x \subseteq \mathbb{K}_x$. As remarked before, the reverse inclusion is automatic, so $\mathbb{O}_x = \mathbb{K}_x$ for every $x \in R$. Since set R has full measure with respect to any Lyapunov maximizing measure, the set $M(\Phi) \setminus R$ has zero measure with respect to any T -invariant probability measure, as we wanted to show. \square

Corollary 6.7. *Suppose Φ admits an extremal norm $\|\cdot\|$. For each $p \in \llbracket 1, d \rrbracket$ and $x \in M_p(\Phi)$, the set \mathbb{K}_x contains a vector space of dimension p .*

Proof. Let μ be a measure whose p first Lyapunov exponents equal $\beta(\Phi)$ and whose support equals $M_p(\Phi)$. Given $x \in M_p(\Phi)$, take a sequence of neighborhoods U_i converging to x . Since $\mu(U_i) > 0$, by Proposition 6.6 we can find $x_i \in U_i$ such that the Oseledets space corresponding to the Lyapunov exponent $\beta(\Phi)$ exists and coincides with \mathbb{K}_{x_i} . Moreover, these spaces have dimensions at least p . Passing to subsequences, we can assume that these dimensions are constant equal to some $q \geq p$, and that \mathbb{K}_{x_i} converges to some q -dimensional space V . As \mathbb{K} is a closed subset of \mathbb{E} , we conclude that $V \subseteq \mathbb{K}_x$, completing the proof. \square

Remark 6.8. It is not necessarily the case that \mathbb{K}_x is a subspace: see Example B.1 in Appendix B.1. On the other hand, if \mathbb{K}_x is a subspace, then by Corollary 6.7 its dimension is at least the number p such that $x \in M_p(\Phi) \setminus M_{p+1}(\Phi)$. However, it is not necessarily true that $\dim \mathbb{K}_x = p$: see Example B.2 in Appendix B.1.

Proof of Theorem 6.5. As usual, it is sufficient to consider $\beta(\Phi) = 0$.

In the case $p = d$, we have $Y \subseteq M_d(\Phi)$ and so by Corollary 6.7 the extremal norm is preserved along the bundle \mathbb{E}_Y . So the trivial splitting $\mathbb{E}_Y \oplus 0$ has the required properties. So let us suppose that $p < d$.

Fix a Riemannian norm $\|\cdot\|$ on \mathbb{E} . For each $x \in Y$, by Corollary 6.7 the fiber \mathbb{E}_x contains a p -dimensional subspace formed by vectors u such that for every $n \geq 0$, we have $\|\Phi^n(u)\| = \|u\|$ and therefore $c_1^{-1}\|u\| \leq \|\Phi^n(u)\| \leq c_1\|u\|$, for some constant $c_1 \geq 1$. Recalling the maxmin characterization of singular values (6.6), we conclude that:

$$c_1^{-1} \leq \sigma_p(\Phi_x^n) \leq \dots \leq \sigma_1(\Phi_x^n) \leq c_1 \quad \text{for all } x \in Y \text{ and } n \geq 0. \quad (6.8)$$

On the other hand, note that the set $M_{p+1}(\Phi|_{\mathbb{E}_Y})$ is contained in $Y \cap M_{p+1}(\Phi)$ and therefore is empty. So Proposition 6.3 yields

$$\beta(\Lambda^{p+1}(\Phi|_{\mathbb{E}_Y})) < (p+1)\beta(\Phi|_{\mathbb{E}_Y}) = 0.$$

Recalling that $\beta(\cdot)$ can also be characterized by (1.4), we conclude that there exist positive constants c_2 and τ such that for all $x \in Y$ and $n \geq 0$,

$$\|\Lambda^{p+1}\Phi_x^n\| \leq c_2 e^{-\tau n}.$$

So, using (6.5) and (6.8), we have:

$$\sigma_{p+1}(\Phi_x^n) = \frac{\|\Lambda^{p+1}\Phi_x^n\|}{\|\Lambda^p\Phi_x^n\|} \leq c_1^p c_2 e^{-\tau n}$$

and $\sigma_p(\Phi_x^n) \geq c_1^{-1}$. So we have a uniform exponential gap between the p -th and $p+1$ -th singular values. By Theorem 6.4, the bundle \mathbb{E}_Y admits a dominated splitting $\mathbb{F} \oplus \mathbb{G}$ with a dominating bundle \mathbb{F} of dimension p .

To conclude, we need to check that Φ preserves the extremal norm along the bundle \mathbb{F} . We will actually show that, in terms of notation (6.7), $\mathbb{K}_x = \mathbb{F}_x$ for every $x \in Y$. Since Φ is product bounded, domination implies that vectors in \mathbb{G} are uniformly contracted in the future, and therefore uniformly expanded in the past. Furthermore, any vector in $u \in \mathbb{E}_x \setminus \mathbb{F}_x$ is uniformly expanded in the past, since we can write $u = v + w$ with $v \in \mathbb{F}_x$, $w \in \mathbb{G}_x \setminus \{0\}$ and then

$$\|\Phi^{-n}(u)\| \geq \|\Phi^{-n}(w)\| \left(1 - \frac{\|\Phi^{-n}(v)\|}{\|\Phi^{-n}(w)\|}\right) \rightarrow \infty \quad \text{as } n \rightarrow +\infty.$$

In particular, vectors in $\mathbb{E}_x \setminus \mathbb{F}_x$ cannot be calibrated; that is, $\mathbb{K}_x \subseteq \mathbb{F}_x$. This inclusion cannot be strict, thanks to Corollary 6.7. So $\mathbb{K}_x = \mathbb{F}_x$, as claimed. \square

7 Further applications of extremal norms and Mather sets

7.1 Subordination

By definition, the Mather set $M(\Phi)$ contains the support of every Lyapunov maximizing measure. Let us see that the converse holds under the hypothesis of strong fiber-bunching, regardless of reducibility:

Theorem 7.1. *Let T be a transitive hyperbolic homeomorphism. Let Φ be strongly bunched automorphism covering T . Then every T -invariant probability measure whose support is contained in the Mather set $M(\Phi)$ is Lyapunov maximizing.*

Proof. Let $\Phi: \mathbb{E} \rightarrow \mathbb{E}$ be a strongly bunched automorphism and let ν be a T -invariant probability measure whose support is contained in $M(\Phi)$; we want to prove that $\chi_1(\Phi, \nu) = \beta(\Phi)$. By ergodic decomposition, it is sufficient to consider the case of ergodic ν .

If Φ is irreducible (which is certainly the case if $d = 1$) then Φ is spannable by Theorem 3.7, and so Φ admits an extremal norm by Theorem 5.2. Then Proposition 6.2 yields the desired conclusion.

From now on, assume that Φ is reducible, that is, there exists a θ -Hölder Φ -invariant nontrivial subbundle $\mathbb{F} \subset \mathbb{E}$. By induction on dimension, we can assume that the theorem holds for the restricted automorphism $\Phi|_{\mathbb{F}}$ and the quotient automorphism Φ/\mathbb{F} . Recall from (4.11) that $\beta(\Phi) = \max\{\beta(\Phi|_{\mathbb{F}}), \beta(\Phi/\mathbb{F})\}$.

As a first case, suppose that $\beta(\Phi) = \beta(\Phi|_{\mathbb{F}}) > \beta(\Phi/\mathbb{F})$. Then it follows from (4.12) that an ergodic measure μ is Lyapunov maximizing for Φ if and only if it is Lyapunov maximizing for $\Phi|_{\mathbb{F}}$. Therefore the Mather sets coincide: $M(\Phi) = M(\Phi|_{\mathbb{F}})$. The measure ν fixed at the beginning is supported on this set; so, by the induction hypothesis, it is Lyapunov maximizing for $\Phi|_{\mathbb{F}}$, that is, it is Lyapunov maximizing for Φ , as we wanted to show.

The second case where $\beta(\Phi) = \beta(\Phi/\mathbb{F}) > \beta(\Phi|_{\mathbb{F}})$ is entirely analogous.

In the last case, we have $\beta(\Phi) = \beta(\Phi|_{\mathbb{F}}) = \beta(\Phi/\mathbb{F})$. Then it follows from (4.12) that an ergodic measure μ is Lyapunov maximizing for Φ if and only if it is Lyapunov maximizing for $\Phi|_{\mathbb{F}}$ or for Φ/\mathbb{F} . Therefore $M(\Phi) = M(\Phi|_{\mathbb{F}}) \cup M(\Phi/\mathbb{F})$. So the measure ν fixed at the beginning has a support contained in the union of the two closed T -invariant sets $M(\Phi|_{\mathbb{F}})$ and $M(\Phi/\mathbb{F})$. By ergodicity, this support must be contained in one of the two sets. By the induction hypothesis, ν is Lyapunov maximizing for $\Phi|_{\mathbb{F}}$ or for Φ/\mathbb{F} . In either case, it is Lyapunov maximizing for Φ , as we wanted to show. \square

7.2 Lyapunov almost-maximizing periodic orbits of low period

Let Φ be a θ -Hölder automorphism covering a hyperbolic homeomorphism. For each integer $n \geq 1$, let

$$\beta_n(\Phi) := \max \{ \chi_1(\Phi, \mu) ; \mu \text{ is supported on a periodic orbit of period } \leq n \}.$$

This is a bounded non-decreasing sequence, and so it is convergent. Actually, the limit is:

$$\lim_{n \rightarrow \infty} \beta_n(\Phi) = \beta(\Phi). \tag{7.1}$$

Indeed, this follows from a much more general result of Kalinin [45, Theorem 1.4] on the approximation of Lyapunov exponents using measures supported on periodic orbits. In the case of one-step cocycles, formula (7.1) is known as the *Berger–Wang theorem*, and it was first proved in [11]. For other extensions of Berger–Wang theorem, see [56, 28].

It is quite possible that the limit (7.1) is attained for some finite n (and indeed this is expected to be the typical situation). On the other hand, in the worst-case scenario, what can we say about the speed of the approximation in formula (7.1)? A result of Morris [54] says that for one-step cocycles, this speed is always superpolynomial. Here we show that the same is true for strongly bunched automorphisms:

Theorem 7.2. *If Φ is a strongly bunched automorphism then for every $\tau > 0$,*

$$\beta(\Phi) - \beta_n(\Phi) = O(n^{-\tau}).$$

The first result of superpolynomial approximation was actually obtained in the context of ergodic optimization of Birkhoff averages by Bressaud and Quas [27], who also showed that this type of bound is essentially sharp. The key ingredient is a quantitative version of Anosov Closing Lemma, also due to Bressaud and Quas [27], which we state as follows:

Theorem 7.3 (Bressaud–Quas Closing Lemma). *Let $T: X \rightarrow X$ be a hyperbolic homeomorphism. Let $Y \subseteq X$ be a nonempty compact T -invariant set. Then for every $\tau > 0$ and every sufficiently large n , there exists a periodic orbit of period at most n entirely contained in the $n^{-\tau}$ -neighborhood of Y .*

This result is proved in [27] for the one-sided full shift; as remarked in that paper, one can use standard techniques to reduce to that case. Alternatively, one can prove Theorem 7.3 directly, and we do so in Appendix A.6.

Proof of Theorem 7.2. Recall from (4.11) that if Φ is reducible then we can replace it by either a restricted or a quotient automorphism with the same maximal Lyapunov exponent. Repeating this procedure a finite number of times, we eventually find a irreducible automorphism with the same maximal Lyapunov exponent; this induced automorphism will also be strongly bunched. So, without loss of generality, we assume that Φ is irreducible. By Theorem 3.7, Φ is spannable, and by Theorem 5.2, Φ admits a Hölder extremal norm $\|\cdot\|$. Let $p \in \llbracket 1, d \rrbracket$ be maximal such that the p -th Mather set $M_p(\Phi) =: Y$ is nonempty. By Theorem 6.5, the restricted bundle \mathbb{E}_Y admits a dominated splitting $\mathbb{F} \oplus \mathbb{G}$ where the dominating bundle \mathbb{F} has fibers of dimension p and is calibrated in the sense that for every $x \in Y$ and $u \in \mathbb{F}_x$, we have $\|\Phi(u)\| = e^{\beta(\Phi)} \|u\|$.

By robustness of dominated splittings [36, Corol. 2.8], there exists a closed neighborhood U of Y such that if $Z := \bigcap_{k \in \mathbb{Z}} T^{-k}(U) \supseteq Y$ is the maximal invariant set in this neighborhood, then the restricted bundle \mathbb{E}_Z over the compact invariant set admits a dominated splitting $\mathbb{F} \oplus \mathbb{G}$, extending the previously found dominated splitting on \mathbb{E}_Y . Recall that the bundles of a dominated splitting are Hölder-continuous [36, Thm. 4.11]. Furthermore, the extremal norm is also Hölder-continuous. It follows that there exist $\rho > 0$ and $C_0 > 0$ such that for every $x \in Z$ and every $u \in \mathbb{F}_x$,

$$e^{\beta(\Phi) - C_0 d(x, Y)^\rho} \|u\| \leq \|\Phi(u)\| \leq e^{\beta(\Phi)} \|u\|.$$

Fix $\tau > 0$. Let C be given by Theorem 7.3; so for all sufficiently large n , there exists a periodic orbit of period at most n supported on the $Cn^{-\tau}$ -neighborhood of Y , and so contained in Z . Let ν_n be the invariant probability measure supported on that orbit. The bound obtained before implies:

$$\chi_1(\Phi, \nu) \geq \beta(\Phi) - C_0(Cn^{-\tau})^\rho.$$

So $\beta(\Phi) - \beta_n(\Phi) = O(n^{-\rho\tau})$. Since $\tau > 0$ is arbitrary, the theorem is proved. \square

A Appendix: Proof of some technical results

A.1 Basic constructions on θ -Hölder bundles

Recall our assumption from Subsection 2.1 that the algebra of θ -Hölder functions on X is normal. Let us metrize the product $X \times X$ by $d((x, y), (x', y')) := \max\{d(x, y), d(x', y')\}$.

Lemma A.1. *The algebra of θ -Hölder functions on $X \times X$ is normal.*

Proof. Let $K_0, K_1 \subset X \times X$ be two disjoint nonempty compact sets. Let $\varepsilon > 0$ be a lower bound for the distance between a point in K_0 and a point in K_1 . Let $\{B_i\}$ be a finite cover of X by open sets of diameter less than ε . Let $\{\rho_i\}$ be a partition of unity subordinated to this cover and formed by θ -Hölder functions. Define a function $f: X \times X \rightarrow \mathbb{R}$ by:

$$f(x, y) := \sum_{\substack{(i,j) \text{ such that} \\ (B_i \times B_j) \cap K_1 \neq \emptyset}} \rho_i(x) \rho_j(y).$$

Then f is θ -Hölder, takes values in the interval $[0, 1]$, equals 0 on K_0 , and equals 1 on K_1 . This proves normality. \square

*Proof of Proposition 2.1 (existence of transport maps $I_{y \leftarrow x}$).*¹⁰ Consider the finite cover of $X \times X$ formed by the following open sets:

$$V_{k,\ell} := \begin{cases} U_k \times U_k & \text{if } k = \ell; \\ U_k \times U_\ell \setminus \Delta & \text{if } k \neq \ell, \end{cases}$$

where $\Delta \subseteq X \times X$ is the diagonal. Consider a partition of unity subordinate to this cover, composed of θ -Hölder functions $\rho_{k,\ell}$; its existence is a consequence of Lemma A.1. Given any pair of points $x, y \in X$, define a linear map from \mathbb{E}_x to \mathbb{E}_y by:

$$I_{y \leftarrow x} := \sum_{(k,\ell)} \rho_{k,\ell}(x, y) h_\ell(y) \circ [h_k(x)]^{-1},$$

where the sum is taken over the indices (k, ℓ) such that $V_{k,\ell} \ni (x, y)$. If $(x, y) \in U_i \times U_j$ then the matrix:

$$[h_j(y)]^{-1} \circ I_{y \leftarrow x} \circ h_i(x) = \sum_{(k,\ell)} \rho_{k,\ell}(x, y) g_{j \leftarrow \ell}(y) \circ g_{k \leftarrow i}(x)$$

is θ -Hölder continuous as a function of (x, y) , and equals the identity when $x = y$. \square

For the following proofs, it is convenient to fix another open cover $\{V_i\}$ of X such that $\bar{V}_i \subset U_i$ for each k . Note that for any Finsler norm $\|\cdot\|$ on \mathbb{E} , we have:

$$\max_i \sup_{x \in V_i} \max \{ \|h_i(x)\|, \|[h_i(x)]^{-1}\| \} < \infty, \quad (\text{A.1})$$

where these operators norms are relative to the norm $\|\cdot\|_x$ on \mathbb{E}_x and the Euclidean norm $\|\cdot\|_{\text{eucl}}$ on \mathbb{R}^d .

Proof of Proposition 2.2 (composition of transport maps). In order to prove the assertion, it is sufficient to consider triples of points x, y, z that are close enough so that they belong to a common coordinate neighborhood V_i . Consider the matrix

$$\tilde{I}_{y \leftarrow x} := [h_i(y)]^{-1} \circ I_{y \leftarrow x} \circ h_i(x), \quad (\text{A.2})$$

which by Proposition 2.1 is $O(d(x, y)^\theta)$ -close to the identity matrix. Using a similar notation for the other points, we have:

$$\|\tilde{I}_{z \leftarrow x} - \text{Id}\|_{\text{eucl}} = O(d(x, z)^\theta) \quad \text{and} \quad \|\tilde{I}_{y \leftarrow z} - \text{Id}\|_{\text{eucl}} = O(d(y, z)^\theta).$$

¹⁰A different construction that provides the additional property $(I_{y \leftarrow x})^{-1} = I_{x \leftarrow y}$ (for sufficiently close x, y) can be found in [46, p. 169]; however, we will not need that property.

Therefore:

$$\|\tilde{I}_{y \leftarrow z} \circ \tilde{I}_{z \leftarrow x} - \tilde{I}_{y \leftarrow x}\|_{\text{eucl}} = O(\max\{d(x, z)^\theta, d(y, z)^\theta\}).$$

Since $I_{y \leftarrow z} \circ I_{z \leftarrow x} - I_{y \leftarrow x} = h_i(y) \circ (\tilde{I}_{y \leftarrow z} \circ \tilde{I}_{z \leftarrow x} - \tilde{I}_{y \leftarrow x}) \circ [h_i(x)]^{-1}$, using the boundedness property (A.1) we obtain

$$\|I_{y \leftarrow z} \circ I_{z \leftarrow x} - I_{y \leftarrow x}\|_{\text{eucl}} = O(\max\{d(x, z)^\theta, d(y, z)^\theta\}),$$

as we wanted to show. \square

Proof of Proposition 2.3 (characterization of θ -Hölder norms). Let $\|\cdot\|$ be a θ -Hölder Finsler norm. In order to prove the desired estimate, it is sufficient to consider pairs of points x, y that are close enough so that they belong to a same set V_i . By definition, for every $u \in \mathbb{R}^d$, the map $x \in V_i \mapsto \|h_i(x)u\|$ is θ -Hölder, and so there is a constant $C > 0$ such that, for all $x, y \in V_i$,

$$\| \|h_i(y)u\| - \|h_i(x)u\| \| \leq C \|u\|_{\text{eucl}} d(x, y)^\theta. \quad (\text{A.3})$$

Using the boundedness property (A.1) and compactness of the unit sphere, we can find a uniform C so that the estimate above holds for every $u \in \mathbb{R}^d$.

Recall that that the matrix defined in (A.2) satisfies $\|\tilde{I}_{y \leftarrow x} - \text{Id}\|_{\text{eucl}} = O(d(x, y)^\theta)$. Now, given $v \in \mathbb{E}_x$, consider $u := [h_i(x)]^{-1}v$. Then:

$$\begin{aligned} \|I_{y \leftarrow x}v\| - \|v\| &= \|h_i(y)\tilde{I}_{y \leftarrow x}u\| - \|h_i(x)u\| \\ &\leq \|h_i(y)\tilde{I}_{y \leftarrow x}u\| - \|h_i(y)u\| + \| \|h_i(y)u\| - \|h_i(x)u\| \| \end{aligned}$$

Using (A.1) and (A.3), we conclude that $\|I_{y \leftarrow x}v\| - \|v\| = O(\|v\|d(x, y)^\theta)$, that is, $\|I_{y \leftarrow x}\| - 1 = O(d(x, y)^\theta)$, as claimed.

The proof of the converse is entirely analogous. \square

Proof of Proposition 2.4 (characterization of θ -Hölder endomorphisms). Suppose Φ is θ -Hölder. In order to prove the desired estimate, it is sufficient to consider pairs of points x, y that are close enough so that they belong to a same set $V_i \cap T^{-1}(V_j)$. Let $\tilde{\Phi}_x := [h_j(Tx)]^{-1} \circ \Phi_x \circ h_i(x)$ and similarly define $\tilde{\Phi}_y$. Let $\tilde{I}_{y \leftarrow x}$ be defined by (A.2), and similarly define $\tilde{I}_{T^y \leftarrow T^x}$; by Proposition 2.1 these matrix-valued maps are θ -Hölder as functions of (x, y) . So the map $(x, y) \mapsto \tilde{I}_{T^y \leftarrow T^x} \circ \tilde{\Phi}_x - \tilde{\Phi}_y \circ \tilde{I}_{y \leftarrow x}$ is also θ -Hölder, and since it vanishes on (x, x) we conclude that:

$$\|\tilde{I}_{T^y \leftarrow T^x} \circ \tilde{\Phi}_x - \tilde{\Phi}_y \circ \tilde{I}_{y \leftarrow x}\| = O(d(x, y)^\theta). \quad (\text{A.4})$$

Using the boundedness property (A.1) we obtain:

$$\|I_{T^y \leftarrow T^x} \circ \Phi_x - \Phi_y \circ I_{y \leftarrow x}\| = O(d(x, y)^\theta),$$

as desired.

Conversely, assume that such an estimate holds; then (A.4) follows from (A.1). By Proposition 2.1, the matrices $\tilde{I}_{y \leftarrow x}$ and $\tilde{I}_{T^y \leftarrow T^x}$ are $O(d(x, y)^\theta)$ -close to the identity. It follows that the matrices $\tilde{\Phi}_x$ and $\tilde{\Phi}_y$ are $O(d(x, y)^\theta)$ -close. This means that Φ is θ -Hölder. \square

A.2 Existence of holonomies

We begin with a straightforward estimate:

Lemma A.2. *Let $\Phi \in \text{Aut}_K^\theta(\mathbb{E}, T)$. For every $x, y \in X$ and $n \geq 0$ we have:*

$$\|(\Phi_y^n)^{-1}\| \|\Phi_x^n\| \leq \prod_{j=0}^{n-1} e^{K_1 d(T^j x, T^j y)^\theta} \text{bol}(\Phi_{T^j y}),$$

where K_1 depends only on K .

Proof. By submultiplicativity of norms and the definition of bolicity, we have:

$$\|(\Phi_y^n)^{-1}\| \|\Phi_x^n\| \leq \prod_{j=0}^{n-1} \|(\Phi_{T^j y})^{-1}\| \|\Phi_{T^j x}\| = \prod_{j=0}^{n-1} \|\Phi_{T^j y}\|^{-1} \|\Phi_{T^j x}\| \text{bol}(\Phi_{T^j y}),$$

and so the claimed inequality holds with K_1 being the θ -Hölder constant of $\log \|\Phi\|$. This constant can be estimated in terms of K , using (2.3). \square

Proof of Proposition 2.8 (existence of holonomies). By symmetry, it is sufficient to consider $\star = s$. The stable holonomy is defined as:

$$H_{y \leftarrow x}^s := \lim_{n \rightarrow +\infty} \underbrace{(\Phi_y^n)^{-1} \circ I_{T^n y \leftarrow T^n x} \circ \Phi_x^n}_{H_n}, \quad (\text{A.5})$$

where x and y are in a same stable set. Let us establish convergence. Assume first that $y \in W_{\varepsilon_0}^s(x)$. We have:

$$H_{n+1} - H_n = (\Phi_y^{n+1})^{-1} \circ \underbrace{(I_{T^{n+1} y \leftarrow T^{n+1} x} \circ \Phi_{T^n x} - \Phi_{T^n y} \circ I_{T^n y \leftarrow T^n x})}_{\Delta_n} \circ \Phi_x^n,$$

and so, using the definition (2.5) of the set $\text{Aut}_K^\theta(\mathbb{E}, T)$ and Lemma A.2,

$$\begin{aligned} \|H_{n+1} - H_n\| &\leq K \|\Delta_n\| \|(\Phi_y^n)^{-1}\| \|\Phi_x^n\| \\ &\leq K^2 d(T^n x, T^n y)^\theta \prod_{j=0}^{n-1} e^{K_1 d(T^j x, T^j y)^\theta} \text{bol}(\Phi_{T^j y}). \end{aligned}$$

By property (2.7) in the definition of hyperbolicity, for every $j \geq 0$ we have

$$d(T^j x, T^j y) \leq e^{-\lambda_s^{(j)}(y)} d(x, y), \quad \text{where} \quad \lambda_s^{(j)}(y) := \sum_{i=0}^{j-1} \lambda_s(T^i y).$$

Since λ_s is strictly positive, the series $\sum_{j=0}^{\infty} d(T^j x, T^j y)^\theta$ is convergent. Therefore

$$\|H_{n+1} - H_n\| \leq K_2 e^{-\theta \lambda_s^{(n)}(y)} \left(\prod_{j=0}^{n-1} \text{bol}(\Phi_{T^j y}) \right) d(x, y)^\theta,$$

where $K_2 > 0$ is another constant. Take a small constant $\eta > 0$ such that the fiber-bunching condition (2.12) still holds if the right hand side is multiplied by $1 - \eta$. In particular, $\text{bol}(\Phi_{T^j y}) < e^{(1-\eta)\theta \lambda_s(T^j y)}$ and so

$$\|H_{n+1} - H_n\| \leq K_2 e^{-\eta \theta \lambda_s^{(n)}(y)} d(x, y)^\theta, \quad (\text{A.6})$$

This establishes uniform exponential convergence in formula (A.5) when $y \in W_{\varepsilon_0}^s(x)$. Using (2.9) we see that convergence holds whenever $y \in W^s(x)$. The groupoid properties (a) and (b) are the equivariance property (c) are automatic from the definition. The Hölderness property (d) follows by summing (A.6) for $n = 0$ to ∞ , and noting that $H_0 = I_{y \leftarrow x}$. The joint continuity property (e) also follows from the uniformity of our estimates. Finally, if we consider a small C^0 perturbation of Φ in the set $\text{Aut}_K^\theta(\mathbb{E}, T)$, then we can use the same constants K_2 and η in (A.6), and so the remaining assertions of the proposition follow. \square

A.3 Regularity estimates

In this appendix, we prove Propositions 2.6 and 2.10. Before going into the proofs, let us state our estimates for Hölder exponents.

Since T is Lipschitz, we can find $\varepsilon_1 \in (0, \varepsilon_0)$ and a continuous strictly positive function Λ_u such that for all $x, x', x'' \in X$,

$$x', x'' \in W_{\varepsilon_1}^u(x) \quad \Rightarrow \quad \begin{cases} Tx', Tx'' \in W_{\varepsilon_0}^u(Tx), \\ d(Tx', Tx'') \leq e^{\Lambda_u(x)} d(x', x''). \end{cases} \quad (\text{A.7})$$

We will show that the conclusion of Proposition 2.6 holds for any κ_s in the range:

$$0 < \kappa_s < \inf_X \frac{\lambda_s + \lambda_u}{\lambda_s + \Lambda_u}. \quad (\text{A.8})$$

Let

$$\varphi(x) := \log \text{bol}(\Phi_x); \quad (\text{A.9})$$

We will show that the conclusion of Proposition 2.10 holds for any κ_s in the range:

$$\theta_s < \inf_X \frac{\theta \lambda_s - \varphi}{\lambda_s + \Lambda_u}. \quad (\text{A.10})$$

(Note that the numerator is positive by fiber-bunching.)

The idea of the proof of Proposition 2.6 is roughly as follows. Without loss of generality, we can assume that the “quadrilateral” in Fig. 1 has a “base” $d(x, x')$ much smaller than the two “legs” $d(x, x')$, $d(x, y)$. We take $N \geq 0$ as big as possible for which we can guarantee that the T^N -image of that quadrilateral has a base smaller than the legs. We estimate the “summit” $d(T^N y, T^N y')$ using the triangle inequality, and finally we iterate backwards to obtain the desired estimate for $d(y, y')$. The proof of Proposition 2.10 uses the same “there and back again” idea. Formal proofs follow.

Let us denote the Birkhoff sums of a function $f: X \rightarrow \mathbb{R}$ as:

$$f^{(n)} := f + f \circ T + \cdots + f \circ T^{n-1}, \quad f^{(0)} := 0.$$

Lemma A.3. *For any strictly positive continuous function f on X and any $a \in (0, 1)$, there exists $b(f, a) > 0$ such that for any $z \in X$ and any $n \geq 0$,*

$$z' \in W_{\varepsilon_0}^s(z) \cup T^{-n}(W_{\varepsilon_0}^u(T^n z)) \quad \Rightarrow \quad f^{(n)}(z') \geq a f^{(n)}(z) - b(f, a).$$

Proof. It is sufficient to consider the case $z' \in W_{\varepsilon_0}^s(z)$, since the case $z' \in T^{-n}(W_{\varepsilon_0}^u(T^n z))$ follows by reversing the time. By uniform continuity of f and uniform contraction on local stable sets, we can find an integer $k = k(f, a) \geq 0$ such that if $z' \in W_{\varepsilon_0}^s(z)$ then $f(T^j z') \geq a f(T^j z)$ for every $j \geq k$. Letting $b(f, a) := ak \sup_X f$, we obtain the desired conclusion. \square

We start the proofs of Propositions 2.6 and 2.10 with some estimates that are common to them. Fix κ_s and θ_s satisfying (A.8) and (A.10), respectively. Fix a number $a \in (0, 1)$ sufficiently close to 1 such that:

$$\kappa_s < \inf_X \frac{a(\lambda_s + \lambda_u)}{a\lambda_s + \Lambda_u}, \quad (\text{A.11})$$

$$\theta_s < \inf_X \frac{a(\theta\lambda_s - \varphi)}{a\lambda_s + \Lambda_u}. \quad (\text{A.12})$$

Fix four points x, x', y, y' satisfying (2.10). Let $\delta := d(x, x')$. Note that to prove Propositions 2.6 and 2.10, it is sufficient to consider δ smaller than a fixed positive constant, say ε_1 from (A.7). Let N be the largest nonnegative integer such that:

$$a\lambda_s^{(N)}(x) + \Lambda_u^{(N)}(x) < \log(\varepsilon_1/\delta); \quad (\text{A.13})$$

Then:

$$a\lambda_s^{(N)}(x) + \Lambda_u^{(N)}(x) \geq \log(\varepsilon_1/\delta) - c, \quad (\text{A.14})$$

for some constant c , namely $c := \sup_X (a\lambda_s + \Lambda_u)$. In particular, assuming that δ is small enough, N will be large and so the following inequality will hold:

$$e^{a\lambda_s^{(N)}(x)} > 2 + e^{b(\lambda_s, a)} \quad (\text{A.15})$$

(where b comes from Lemma A.3).

Using (A.7) and (A.13), one checks by induction that the following chain of inequalities hold for each $n \in \llbracket 0, N \rrbracket$:

$$d(T^n x, T^n x') \leq \delta e^{\Lambda_u^{(n)}(x)} \leq \varepsilon_1 e^{-a\lambda_s^{(n)}(x)} \leq \varepsilon_1. \quad (\text{A.16})$$

This gives estimates for the base of the ‘‘quadrilateral’’ obtained as the T^n -image of that of Fig. 1. Let us estimate the other sides; the ‘‘legs’’ are:

$$d(T^n x, T^n y) \leq \varepsilon_0 e^{-\lambda_s^{(n)}(x)}, \quad (\text{A.17})$$

$$d(T^n x', T^n y') \leq \varepsilon_0 e^{-\lambda_s^{(n)}(x')} \leq \varepsilon_0 e^{-a\lambda_s^{(n)}(x) + b(\lambda_s, a)}, \quad (\text{A.18})$$

where in the last inequality we used Lemma A.3. Therefore the ‘‘summit’’ is:

$$d(T^n y, T^n y') \leq d(T^n x, T^n y) + d(T^n x, T^n x') + d(T^n x', T^n y') \quad (\text{A.19})$$

$$\leq (2 + e^{b(\lambda_s, a)}) \varepsilon_0 e^{-a\lambda_s^{(n)}(x)} \quad (\text{A.20})$$

$$\leq \varepsilon_0, \quad (\text{A.21})$$

where in the last step we used assumption (A.15).

We estimate the base of the original quadrilateral by iterating backwards:

$$\begin{aligned} d(y, y') &\leq e^{-\lambda_u^{(N)}(y)} d(T^N y, T^N y') && \text{(by (A.21))} \\ &= O(e^{-a\lambda_u^{(N)}(x)} d(T^N y, T^N y')) && \text{(by Lemma A.3)} \\ &= O(e^{-a[\lambda_s^{(N)}(x) + \lambda_u^{(N)}(x)]}) && \text{(by (A.20))} \\ &= O(e^{-\kappa_s[a\lambda_s^{(N)}(x) + \Lambda_u^{(N)}(x)]}) && \text{(by (A.11))} \\ &= O(\delta^{\kappa_s}) && \text{(by (A.14)).} \end{aligned}$$

This proves Proposition 2.6.

We proceed to the proof of Proposition 2.10. In what follows, the constants implicit in O can be taken uniform on a C^0 -neighborhood of Φ in $\text{Aut}_K^\theta(\mathbb{E}, T)$. For $n \geq 0$, define

$$\Gamma_n := H_{T^n y' \leftarrow T^n y}^u \circ H_{T^n y \leftarrow T^n x}^s - H_{T^n y' \leftarrow T^n x'}^s \circ H_{T^n x' \leftarrow T^n x}^u.$$

Let us estimate the norm of these linear maps. First,

$$\|\Gamma_0\| \leq \|H_{y' \leftarrow y}^u \circ H_{y \leftarrow x}^s - I_{y' \leftarrow x}\| + \|H_{y' \leftarrow x'}^s \circ H_{x' \leftarrow x}^u - I_{y' \leftarrow x}\| =: \textcircled{1} + \textcircled{2}.$$

We estimate the first term:

$$\textcircled{1} \leq \|H_{y' \leftarrow y}^u \circ (H_{y \leftarrow x}^s - I_{y \leftarrow x})\| + \|(H_{y' \leftarrow y}^u - I_{y' \leftarrow y}) \circ I_{y \leftarrow x}\| + \|I_{y' \leftarrow y} \circ I_{y \leftarrow x} - I_{y' \leftarrow x}\|.$$

Using (2.14) and Proposition 2.2, we conclude that

$$\textcircled{1} = O(\max\{d(x, y)^\theta, d(y, y')^\theta\}).$$

An analogous reasoning yields:

$$\textcircled{2} = O(\max\{d(x, x')^\theta, d(x', y')^\theta\}).$$

So we obtain:

$$\|\Gamma_0\| = O(\max\{d(x, x')^\theta, d(y, y')^\theta, d(x, y)^\theta, d(x', y')^\theta\}).$$

Any of these four distances, say $d(y, y')$, is less than the sum of the other three; so:

$$\|\Gamma_0\| = O(\max\{d(x, x')^\theta, d(x, y)^\theta, d(x', y')^\theta\}).$$

Now if $n \in \llbracket 0, N \rrbracket$, the corresponding quadrilateral has sides are bounded by ε_0 (estimates (A.16)–(A.21)), and the exact same argument yields:

$$\|\Gamma_n\| = O(\max\{d(T^n x, T^n x')^\theta, d(T^n x, T^n y)^\theta, d(T^n x', T^n y')^\theta\}).$$

Then, using estimates (A.16), (A.17), and (A.18) we obtain:

$$\|\Gamma_n\| = O(e^{-\theta a \lambda_s^{(n)}(x)}). \quad (\text{A.22})$$

As a consequence of estimates (A.16) and (A.18), if $j \in \llbracket 0, N \rrbracket$ then $d(T^j y', T^j x) = O(e^{-a \lambda_s^{(j)}(x)})$. So it follows from Lemma A.2 that

$$\|(\Phi_{y'}^n)^{-1}\| \|\Phi_x^n\| = O(e^{\varphi^{(n)}(y')}), \quad \text{for } n \in \llbracket 0, N \rrbracket. \quad (\text{A.23})$$

Now we want to iterate backwards to obtain a finer estimate for $\|\Gamma_0\|$. By the groupoid properties of holonomies, $\Gamma_0 = (\Phi_{y'}^n)^{-1} \circ \Gamma_n \circ \Phi_x^n$. Therefore:

$$\begin{aligned} \|\Gamma_0\| &= O(e^{\varphi^{(N)}(y')} \|\Gamma_N\|) && \text{(by (A.23))} \\ &= O(e^{-\theta a \lambda_s^{(N)}(x) + a \varphi^{(N)}(x)}) && \text{(by (A.22) and Lemma A.3)} \\ &= O(e^{-\theta_s [a \lambda_s^{(N)}(x) + \Lambda_u^{(N)}(x)]}) && \text{(by (A.12))} \\ &= O(\delta^{\theta_s}) && \text{(by (A.14)).} \end{aligned}$$

Proposition 2.10 is proved.

Proof of Lemma 3.6 (the strong bunching constant). Fix any positive

$$\eta_0 \leq \inf_X \frac{\theta \lambda_s}{\lambda_s + \lambda_u + \Lambda_u}. \quad (\text{A.24})$$

Suppose Φ is a (η_0, θ) -bunched automorphism. This means that the function φ defined by (A.9) is less than $\eta_0 \lambda_u$. Therefore we have pointwise inequalities:

$$\frac{\theta \lambda_s - \varphi}{\lambda_s + \Lambda_u} > \frac{\theta \lambda_s - \eta_0 \lambda_u}{\lambda_s + \Lambda_u} \geq \eta_0.$$

So there exists $\theta_s \geq \eta_0$ that satisfies (A.10). \square

A.4 The metric on the Grassmannian

Let E be an inner product space of dimension d . If $V_1, V_2 \subseteq E$ are subspaces of the same dimension $p > 0$, we define:

$$d(V_1, V_2) := \inf_{F_1, F_2} \|F_1 - F_2\|, \quad (\text{A.25})$$

where each F_i runs over all linear isomorphisms $F_i: \mathbb{R}^p \rightarrow V_i$ such that $\|F_i^{-1}\| \leq 1$. (We consider \mathbb{R}^p endowed with the canonical inner product, and $\|\cdot\|$ always denotes the operator norm.)

Proposition A.4. *d is a metric on the Grassmannian $\mathcal{G}_p(E)$.*

Proof. Symmetry and the triangular inequality are trivially satisfied, so let us check non-degeneracy. Suppose $V_1 \neq V_2 \in \mathcal{G}_p(E)$. Take a unit vector v_1 in V_1 but not in V_2 . Then there exists $\delta > 0$ such that $\|v_1 - v_2\| \geq \delta$ for every $v_2 \in V_2$. For each $i \in \{1, 2\}$, let $F_i: \mathbb{R}^p \rightarrow V_i$ be a linear isomorphism such that $\|F_i^{-1}\| \leq 1$. Then:

$$\|F_1 - F_2\| \geq \frac{\|v_1 - F_2(F_1^{-1}(v_1))\|}{\|F_1^{-1}(v_1)\|} \geq \|v_1 - F_2(F_1^{-1}(v_1))\| \geq \delta.$$

This shows that $d(V_1, V_2) \geq \delta > 0$. \square

Proof of Proposition 3.1 (linearly induced maps are Lipschitz). Consider a linear isomorphism $L: E \rightarrow F$ between d -dimensional inner product spaces. For each $i \in \{1, 2\}$, let $V_i \in \mathcal{G}_p(E)$, and let $W_i := L(V_i)$. Let $F_i: \mathbb{R}^p \rightarrow V_i$ be a linear isomorphism such that $\|F_i^{-1}\| \leq 1$. Define $G_i: \mathbb{R}^p \rightarrow W_i$ by $G_i := \|L^{-1}\| L \circ F_i$. Then G_i is a linear isomorphism and $\|G_i^{-1}\| \leq 1$. So

$$d(W_1, W_2) \leq \|G_1 - G_2\| = \|L^{-1}\| \|L \circ F_1 - L \circ F_2\| \leq \text{bol}(L) \|F_1 - F_2\|.$$

Taking infimum over the F_i 's, we obtain $d(W_1, W_2) \leq \text{bol}(L) d(V_1, V_2)$. This proves that the map induced by L has Lipschitz constant $\text{bol}(L)$. \square

Proof of Proposition 3.2 (maps close to the identity). Suppose $L: E \rightarrow E$ satisfies $\|L - \text{id}\| \leq \delta \leq \frac{1}{2}$. Note that $\|L^{-1}\| \leq (1 - \delta)^{-1}$. Fix an arbitrary $V \in \mathcal{G}_p(E)$. Let $F_1: \mathbb{R}^p \rightarrow V$ be an isometry, and let $F_2 := (1 - \delta)^{-1} L \circ F_1$. Then $\|F_1^{-1}\| = 1$, $\|F_2^{-1}\| \leq 1$, and so

$$d(V, LV) \leq \|F_1 - F_2\| = \|\text{id} - (1 - \delta)^{-1} L\| \leq \|\text{id} - L\| + \frac{\delta}{1 - \delta} \|L\| \leq \frac{2\delta}{1 - \delta} \leq 4\delta. \quad \square$$

Proof of Proposition 3.3 (span is locally Lipschitz). This is an easy consequence of the definition (A.25), and details are left to the reader. \square

Remark A.5. It can be shown that our metric (A.25) coincides with the metric used in [17, § A.1].

A.5 Typical fiber-bunched automorphisms are irreducible

Recall from Subsection 2.3 that $\text{End}^\theta(\mathbb{E}, T)$ denotes the vector space of θ -Hölder endomorphisms, which becomes a Banach space with the θ -Hölder norm (2.4). The set $\text{Aut}^\theta(\mathbb{E}, T)$ of θ -Hölder automorphisms and the subset $\mathcal{B} \subset \text{Aut}^\theta(\mathbb{E}, T)$ of fiber-bunched automorphisms are both open subsets of $\text{End}^\theta(\mathbb{E}, T)$ (actually they are C^0 -open).

A subset of a Banach space is said to be of *infinite codimension* if it is locally contained in the union of finitely many closed submanifolds of arbitrarily large codimension.

Proposition A.6. *Suppose X is infinite, $T: X \rightarrow X$ is a transitive hyperbolic homeomorphism, and \mathbb{E} is a θ -Hölder vector bundle over X . Then there exists an open and dense subset \mathcal{I} of the set $\mathcal{B} \subset \text{Aut}^\theta(\mathbb{E}, T)$ of fiber-bunched automorphisms such that every $\Phi \in \mathcal{I}$ is irreducible. Furthermore, the set $\mathcal{B} \setminus \mathcal{I}$ has infinite codimension.*

The proof is an obvious adaptation of arguments from [20, 72], so we will make it concise.

Proof. As a consequence of shadowing and expansivity, the hyperbolic homeomorphism T has infinitely many periodic points (see e.g. [1, p. 228]). Select one of these, say a point p of period k , and a homoclinic point q associated to p . If Φ is reducible then it admits a non-trivial θ -Hölder Φ -invariant subbundle \mathbb{F} which by Corollary 3.5 is both H^u - and H^s -invariant. Then the subspace $\mathbb{F}_p \subseteq \mathbb{E}_p$ is invariant under two different linear maps, namely Φ_p^k and $H_{p \leftarrow q}^u \circ H_{q \leftarrow p}^u$. On the other hand, we claim that the property that these two maps admit a common nontrivial invariant subspace is *atypical* in the space $\text{Aut}^\theta(\mathbb{E}, T)$; more precisely, it has positive codimension and, a fortiori, empty interior.

First note that the property that an element of $GL(d, \mathbb{R})$ admits infinitely many invariant subspaces is atypical (because it implies the existence of a complex eigenvalue of geometric multiplicity bigger than 1). So for typical $\Phi \in \text{Aut}^\theta(\mathbb{E}, T)$, the collection of Φ_p^k -invariant subspaces is finite.

On the other hand, choose a closed neighborhood U of q that is disjoint from the future and past iterates of q . If we perturb the automorphism in this neighborhood (or rather in $\pi^{-1}(U)$) then the maps Φ_p^k and $H_{q \leftarrow p}^u$ are unaffected, but $H_{p \leftarrow q}^s$ changes, and actually any small perturbation of $H_{p \leftarrow q}^s$ can be realized with a perturbation of Φ supported in U . In particular, with a well-chosen perturbation, the composition of holonomies sends each of the (finitely many) Φ_p^k -invariant nontrivial subspaces of \mathbb{E}_p into something transverse to it. Such an automorphism Φ cannot be reducible. This shows that irreducibility has dense interior in $\text{Aut}^\theta(\mathbb{E}, T)$.

The argument actually shows that if Φ is reducible then it must satisfy infinitely many independent conditions of positive codimension, at least one for each homoclinic orbit. Therefore reducibility has infinite codimension. See [72, §4] for full details. \square

A.6 Bressaud–Quas Closing Lemma

Here we will prove Theorem 7.3. Though our formulation is different, the key ideas come from [27].

Let $f: Y \rightarrow Y$ be any homeomorphism of a compact metric space (Y, d) . For $\varepsilon > 0$, an (ε, d, f) -pseudorbit is a string of points $(x_0, x_1, \dots, x_{n-1})$ such that $d(f(x_i), x_{i+1}) < \varepsilon$ for every $i \in \llbracket 0, n-2 \rrbracket$. If additionally $d(f(x_{n-1}), x_0) < \varepsilon$ then we say that the pseudo-orbit is *periodic*, with *period* n ; in that case indices can be taken as integers mod n instead. Let $R(\varepsilon, d, f)$ denote the minimal period of a periodic (ε, d, f) -pseudorbit. Note that:

$$R(\varepsilon, d, f) \leq nR(\varepsilon, d, f^n) \quad \text{for every } n \geq 1. \quad (\text{A.26})$$

A set $E \subseteq Y$ is called (ε, d) -separated if $d(x, y) \geq \varepsilon$ for every pair of distinct points $x, y \in E$. Let $S(\varepsilon, d)$ be the maximal cardinality of a (ε, d) -separated set.

Define a sequence of metrics by:

$$d_{n,f}(x, y) := \max_{i \in \llbracket 0, n-1 \rrbracket} d(f^i(x), f^i(y)).$$

Lemma A.7. *Let $\varepsilon > 0$. Suppose that $R(\varepsilon, d, f) > m > 0$. Then:*

$$\log m \leq \log S\left(\frac{\varepsilon}{2}, d\right) - \frac{1}{m} \log S(\varepsilon, d_{m,f}) + 1.$$

Proof. Let E be a $(d, \frac{\varepsilon}{2})$ -separated set of maximal cardinality. Note that the d -balls of radius $\frac{\varepsilon}{2}$ and centers at the points of E cover Y , because otherwise we could enlarge E by adding any point not covered.

Let F be a $(d_{m,f}, \varepsilon)$ -separated set of maximal cardinality. For each $y \in F$, choose a m -tuple (x_0, \dots, x_{m-1}) of points in E such that $d(x_j, f^j(y)) < \frac{\varepsilon}{2}$ for each $j \in \llbracket 0, m-1 \rrbracket$. First, we claim that these x_j 's are all distinct. Indeed, if $x_j = x_k$ with $j < k$, then $d(f^j(y), f^k(y)) < \varepsilon$, so $(f^i(y))_{i \in \llbracket j, k-1 \rrbracket}$ is a periodic (ε, d) -pseudorbit of period $k - j \leq m - 1 < R(\varepsilon, d)$, contradiction.

Second, we claim that if $y \neq y' \in F$ then the corresponding m -tuples (x_0, \dots, x_{m-1}) and (x'_0, \dots, x'_{m-1}) are distinct. Indeed, if the two m -tuples coincide then for each $j \in \llbracket 0, m-1 \rrbracket$ we have $d(f^j(y), f^j(y')) < \varepsilon$. This means that $d_{m,f}(y, y') < \varepsilon$. Since the set F is $(d_{m,f}, \varepsilon)$ -separated, we conclude that $y = y'$.

Third, we claim that if $y \neq y' \in F$ then the sets $\{x_0, \dots, x_{m-1}\}$ and $\{x'_0, \dots, x'_{m-1}\}$ are distinct. Indeed, if the two sets coincide then $x'_i = x_{\sigma(i)}$ for some permutation σ of $\llbracket 0, m-1 \rrbracket$. By the previous claim, this permutation is not the identity; therefore there exists $\ell \in \llbracket 0, m-2 \rrbracket$ such that $k := \sigma(\ell) > \sigma(\ell+1) =: j$. Then $d(f^\ell(y'), f^k(y)) < \varepsilon$ and $d(f^{\ell+1}(y'), f^j(y)) < \varepsilon$. Therefore $(f^j(y), f^{j+1}(y), \dots, f^{k-1}(y), f^\ell(y'))$ is a periodic (ε, d) -pseudorbit of period $k - j + 1 \leq m < R(\varepsilon, d)$, contradiction.

We conclude that the number of elements of the set F cannot exceed the number of subsets of the set E with exactly m elements, that is,

$$|F| \leq \binom{|E|}{m} \leq \frac{|E|^m}{m!} \leq \left(\frac{e|E|}{m}\right)^m.$$

Taking log's, recalling that $|E| = S(d, \frac{\varepsilon}{2})$ and $|F| = S(d_{m,f}, \frac{\varepsilon}{2})$, and rearranging, we obtain the inequality stated in the lemma. \square

Recall that a homeomorphism $f: Y \rightarrow Y$ is called *expansive* if there is a uniform separation between every pair of distinct orbits. In that case, the topological entropy $h_{\text{top}}(f)$ is finite; furthermore, for every sufficiently small $\varepsilon > 0$, the limit

$$\lim_{n \rightarrow \infty} \frac{\log S(\varepsilon, d_{n,f})}{n} \text{ exists and equals } h_{\text{top}}(f); \quad (\text{A.27})$$

see [73, p. 174, 177].

Lemma A.8. *If f is expansive then for every sufficiently small $\varepsilon > 0$ we have:*

$$\lim_{n \rightarrow \infty} \frac{\log R(\varepsilon, d_{n,f}, f)}{n} = 0.$$

Proof. Fix a small $\varepsilon > 0$ and a large integer n . Write $R_n := R(\varepsilon, d_{n,f}, f)$ and $m_n := \lfloor (R_n - 1)/n \rfloor$. Assume that $m_n > 0$, otherwise $(\log R_n)/n$ is already small. Using (A.26), we have $R(\varepsilon, d_{n,f}, f^{m_n}) > m_n$. Applying Lemma A.7, we obtain that

$$\frac{\log m_n}{n} \leq \frac{\log S(\varepsilon/2, d_{n,f})}{n} - \frac{\log S(\varepsilon, d_{nm_n, f})}{nm_n} + \frac{1}{n}.$$

By (A.27), the right-hand side is small: the first two terms essentially cancel each other. It follows that $(\log R_n)/n$ is small. \square

Proof of Theorem 7.3. Given the hyperbolic homeomorphism T and the compact T -invariant set $Y \neq \emptyset$, let f be the restriction of T to Y . Hyperbolic homeomorphisms are expansive (recall Remark 2.5), so f is expansive as well.

Fix $\varepsilon > 0$ small enough so that Lemma A.8 applies. Note that if $(x_i)_{i \in \mathbb{Z}/k\mathbb{Z}}$ is a periodic (ε, d_n, f) -pseudorbit then, letting $y_i := f^{\lfloor n/2 \rfloor}(x_i)$, we have, for all $i \in \mathbb{Z}/k\mathbb{Z}$,

$$\max_{j \in [-\lfloor n/2 \rfloor, \lfloor n/2 \rfloor - 1]} d(f^{j+1}(y_i), f^j(y_{i+1})) < \varepsilon.$$

Hyperbolicity implies that $d(f(y_i), y_{i+1}) < C e^{-\lambda n} \varepsilon =: \varepsilon_n$, where C and λ are positive constants. That is, $(y_i)_{i \in \mathbb{Z}/k\mathbb{Z}}$ is a periodic (ε_n, d, f) -pseudorbit. So $R(\varepsilon, d_{f,n}, f) \geq R(\varepsilon_n, d, f) =: N_n$. In particular,

$$\frac{\log N_n}{n} \quad \text{and} \quad \frac{\log N_n}{\log \varepsilon_n^{-1}} \quad \text{also tend to 0 as } n \rightarrow \infty.$$

Therefore, for any given $\tau > 0$, if n is large enough then $\varepsilon_n < N_n^{-\tau}$. By definition, there exists a periodic (ε_n, d, T) -pseudorbit of period N_n in the set Y . By the Lipschitz shadowing lemma [66, Thm. 2], there exist a periodic orbit for T of period N_n within distance $O(\varepsilon_n) = O(N_n^{-\tau})$. This proves the theorem. \square

B Appendix: Examples

Here we present examples that show some of the limits of our results.

B.1 Examples of sets of calibrated vectors with exceptional behavior

The following two examples show that the set \mathbb{K} of calibrated vectors defined by (6.7) can have exceptional fibers, justifying Remark 6.8.

Example B.1. Let $T: X \rightarrow X$ be a hyperbolic homeomorphism having a fixed point x_0 . Let f be a non-negative Hölder function vanishing only at x_0 . Consider the cocycle

$$A(x) := \begin{pmatrix} 1 & 0 \\ 0 & e^{-f(x)} \end{pmatrix}.$$

Then the corresponding automorphism Φ on the trivial bundle $\mathbb{E} := X \times \mathbb{R}^2$ has $\beta(\Phi) = 0$, and its Mather sets are $M_1(\Phi) = X$ and $M_2(\Phi) = \{x_0\}$. The max norm $\|(u_1, u_2)\| := \max\{|u_1|, |u_2|\}$ is extremal. Consider the corresponding set \mathbb{K} of calibrated vectors, defined by (6.7). If $x \in W^u(x_0) \setminus \{x_0\}$ then the fiber

$$\mathbb{K}_x = \{(u_1, u_2) \in \mathbb{R}^2 ; |u_2| \leq e^{\sum_{n=1}^{\infty} f(T^{-n}x)} |u_1|\}$$

is not a subspace.

Example B.2. Suppose $T: X \rightarrow X$ is a homeomorphism admitting two nonempty compact invariant sets X_1, X_2 such that:

- each X_i equals the support of some T -invariant probability measure μ_i ;
- $X_1 \cup X_2 = X$;
- $X_1 \cap X_2 = \{x_0\} \cup \{T^n y_0 ; n \in \mathbb{Z}\}$ where x_0 is a fixed point and $y_0 \neq x_0$ is an homoclinic point.

Let f be a non-negative continuous function vanishing only at $X_1 \cap X_2$. Consider the cocycle

$$A(x) := \begin{pmatrix} 1 & 0 \\ 0 & e^{-f(x)} \end{pmatrix} \quad \text{if } x \in X_1, \quad A(x) := \begin{pmatrix} e^{-f(x)} & 0 \\ 0 & 1 \end{pmatrix} \quad \text{if } x \in X_2.$$

Then the corresponding automorphism Φ on the trivial bundle $\mathbb{E} := X \times \mathbb{R}^2$ has $\beta(\Phi) = 0$, and its Mather sets are $M_1(\Phi) = X$ and $M_2(\Phi) = \{x_0\}$. The Euclidean norm is extremal. Consider the corresponding set \mathbb{K} of calibrated vectors, defined by (6.7). Then $\mathbb{K}_{y_0} = \mathbb{R}^2$ despite the fact that $y_0 \in M_1(\Phi) \setminus M_2(\Phi)$.

Remark B.3. One may contend that “correctly” defined Mather sets should not lie in the base X , but instead in the bundle \mathbb{E} , or in its projectivization $\hat{\mathbb{E}}$. Fix a norm $\|\cdot\|$ on \mathbb{E} and define a function $f: \hat{\mathbb{E}} \rightarrow \mathbb{R}$ by $f([u]) := \log(\|\Phi u\|/\|u\|)$. Let $\hat{\mathbb{M}}$ be the union of the supports of all probability measures on $\hat{\mathbb{E}}$ that are invariant under the automorphism $\hat{\Phi}$ and that maximize the integral of the function f . Let $\mathbb{M} := \{u \in \mathbb{E} ; u = 0 \text{ or } [u] \in \hat{\mathbb{M}}\}$. This is a closed subset of \mathbb{E} that projects down on the Mather set $M(\Phi) \subseteq X$. Given an extremal norm, it is clear that the fibers of \mathbb{M} are calibrated, i.e. $\mathbb{M}_x \subseteq \mathbb{K}_x$ for every $x \in M(\Phi)$. A stronger property actually holds: $\text{span}(\mathbb{M}_x) \subseteq \mathbb{K}_x$ for every $x \in M(\Phi)$; we omit the proof. However, \mathbb{M}_x may fail to be a subspace. Indeed, in Example B.2 the set \mathbb{M}_{y_0} is a union of two lines.

B.2 On Riemannian extremal norms

After having established the existence of extremal *Finsler* norms (under appropriate hypotheses), one naturally wonders about the existence of extremal *Riemannian* norms. Let us begin with a weak positive result:

Proposition B.4. *In the situation of Theorem 6.5, there exists a Riemannian norm $\|\cdot\|'$ such that for all $x \in Y$, the spaces \mathbb{F}_x and \mathbb{G}_x are orthogonal, and*

$$\|\Phi(v)\|' = e^{\beta(\Phi)} > \|\Phi(w)\|' \quad \text{for all unit vectors } v \in \mathbb{F}_x, w \in \mathbb{G}_x.$$

Proof. As usual, assume $\beta(\Phi) = 0$. For each $x \in Y$, consider the restriction of the extremal norm to the space \mathbb{F}_x , and let $\mathcal{B}_x \subseteq \mathbb{F}_x$ be the unit ball. Let \mathcal{E}_x be the John ellipsoid of \mathcal{B}_x , namely the unique ellipsoid of maximal volume contained in \mathcal{B}_x (see e.g. [9]). This field of ellipsoids is continuous, since finding the John ellipsoid is a continuous operation (as a consequence of its uniqueness). Consider the Riemannian norm on the bundle \mathbb{F} whose unit balls are the \mathcal{E}_x 's. Since $\Phi_x(\mathcal{B}_x) = \mathcal{B}_{T_x}$ and the John ellipsoid is equivariant with respect to linear isomorphisms, we obtain $\Phi_x(\mathcal{E}_x) = \mathcal{E}_{T_x}$. This means that the Riemannian norm just constructed on the bundle \mathbb{F} is preserved by Φ .¹¹

In the bundle \mathbb{G} , we use the standard construction of Lyapunov norms (see e.g. [47, p. 667]). Fix a small positive ε , and for each $x \in M_p$ and $w \in \mathbb{G}_x$, let:

$$\|w\|'_x := \left(\sum_{n=0}^{\infty} e^{2\varepsilon n} \|\Phi_x^n(w)\|^2 \right)^{1/2}.$$

As a consequence of domination, the series converges exponentially, so the formula yields a well-defined continuous Riemannian norm on the bundle \mathbb{G} . It is immediate that $\|\Phi_x(w)\|' \leq e^{-\varepsilon} \|w\|'$, so the norm along \mathbb{G} is uniformly contracted. Finally, we extend the Riemannian norm to the fibers \mathbb{E}_x for $x \in Y$ by declaring \mathbb{F}_x and \mathbb{G}_x to be orthogonal. This completes the construction. \square

The Riemannian norm $\|\cdot\|'$ provided by Proposition B.4 is extremal over the restricted subbundle $\mathbb{E}_Y = \pi^{-1}(Y)$. Can one extend this Riemannian norm to the whole bundle, keeping it extremal? The answer is no, as we will see next.

We will present an example of an irreducible fiber-bunched automorphism in dimension 2 that admits no *Riemannian* extremal norm.

Consider the following two matrices:

$$A_0 := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad A_1 := \begin{pmatrix} 0.8 & -0.1 \\ 0.8 & 0.1 \end{pmatrix}. \quad (\text{B.1})$$

Let (T, F) be the corresponding one-step cocycle (see Example 5.6), and let Φ be the corresponding automorphism of the trivial vector bundle $\mathbb{E} = X \times \mathbb{R}^2$. Consider a Hölder exponent $\theta = 1$, and take the parameter λ in the metric (5.12) large enough so that Φ becomes fiber-bunched. Consider the fixed point $p = (p_n)$ where each $p_n := 0$. Since $F(p) = A_0$ has non-real eigenvalues, the automorphism Φ is irreducible: there can be no nontrivial Φ -invariant subbundle. Then Theorem 5.2 yields the existence of an extremal norm. Actually, the max norm in \mathbb{R}^2 , defined by $\|(u_1, u_2)\| := \max\{|u_1|, |u_2|\}$, is an extremal norm. Indeed, the operator norms of our two matrices are:

$$\|A_0\| = 1, \quad \|A_1\| = 0.9.$$

Since the spectral radius of A_0 is 1, it follows that $\beta(\Phi) = 0$, and so $\|\cdot\|$ is a (constant) extremal norm, as claimed. Also note that δ_p is the unique Lyapunov-maximizing measure.

Proposition B.5. *The automorphism Φ admits no Riemannian extremal norm.*

¹¹Incidentally, note that if the John ellipsoid were monotonic with respect to set inclusion, then we could use it to “Riemannize” any given Finsler extremal norm. However, monotonicity fails: consider for instance a pair of rectangles as in Fig. 5.

Proof. Assume for a contradiction that Φ admits a Riemannian extremal norm $\{\|\cdot\|_x\}_{x \in X}$. Since $e^{\beta(\Phi)} = 1$ is the spectral radius of $F(p) = A_0$, we must have:

$$\|A_0\|_{p \leftarrow p} \leq 1,$$

in the operator norm notation (1.2). This means that if $D \subset \mathbb{R}^2$ denotes the unit ball in the norm $\|\cdot\|_p$, we have $A_0(D) \subseteq D$. Since the norm is assumed to be Riemannian, D is a (filled) ellipse, and since A_0 is a rotation, this ellipse must be a disk. Rescaling the norm if necessary, we can assume that D is the unit disk. Equivalently, $\|\cdot\|_{p \leftarrow p}$ is the usual Euclidean operator norm, which for emphasis we will write $\|\cdot\|_{\text{eucl}}$.

Consider the homoclinic point $q := (\dots, 0, 0, 1, 0, 0, \dots)$, i.e., the sequence that has a unique symbol 1 at position 0. Note that for any $k > 0$ we have the identity:

$$H_{p \leftarrow q}^s \circ H_{q \leftarrow p}^u = \Phi_p^{-k} \circ H_{p \leftarrow T^k q}^s \circ \Phi_{T^{-k} q}^{2k} \circ H_{T^{-k} q \leftarrow p}^u \circ \Phi_p^{-k}. \quad (\text{B.2})$$

In particular, taking $k = 1$, by triviality of local holonomies (5.13) we obtain:

$$\begin{aligned} H_{p \leftarrow q}^s \circ H_{q \leftarrow p}^u &= A_0^{-1} \circ \text{id} \circ A_1 A_0 \circ \text{id} \circ A_0^{-1} \\ &= A_0^{-1} A_1. \end{aligned}$$

Using that A_0 preserves Euclidean norm, applying the extremal Riemannian norm to (B.2):

$$\|A_1\|_{\text{eucl}} = \|A_0^{-1} A_1\|_{\text{eucl}} = \|H_{p \leftarrow q}^s \circ H_{q \leftarrow p}^u\| \leq \underbrace{\|\Phi_p^{-k}\|}_{=1} \underbrace{\|H_{p \leftarrow T^k q}^s\|}_{\rightarrow 1} \underbrace{\|\Phi_{T^{-k} q}^{2k}\|}_{\leq 1} \underbrace{\|H_{T^{-k} q \leftarrow p}^u\|}_{\rightarrow 1} \underbrace{\|\Phi_p^{-k}\|}_{=1}.$$

Taking $k \rightarrow \infty$ we obtain that $\|A_1\|_{\text{eucl}} \leq 1$. This is a contradiction: actually $\|A_1\|_{\text{eucl}} = 0.8\sqrt{2} > 1$ (see Fig. 5). \square




Figure 5: The unit balls with respect to the two norms and their images under A_1 .

Let us comment on other properties of our example. We claim that there are perturbations $\tilde{\Phi}$ of Φ for which the measure δ_p ceases to be Lyapunov-maximizing; so the “locking property” (*verrouillage*) is not satisfied. Indeed, let $k \gg 1$ be an integer, let $m := 4k + 2$, and let \tilde{A}_0 be rotation matrix of angle $\frac{\pi}{2} - \frac{\pi}{4m}$. Then:

$$\tilde{A}_0^m A_1 = \begin{pmatrix} -0.8\sqrt{2} & 0 \\ 0 & -0.1\sqrt{2} \end{pmatrix}.$$

Consider the associated one-step cocycle \tilde{F} , and the associated automorphism $\tilde{\Phi}$. Then the probability measure $\tilde{\mu}$ supported on the orbit of the periodic point

$$\tilde{p} := (\dots, 1, \underbrace{0, \dots, 0}_m, 1, \underbrace{0, \dots, 0}_m, \dots), \quad \text{i.e., } \tilde{p}_n = 1 \text{ iff } m + 1 \text{ divides } n,$$

has Lyapunov exponent

$$\chi_1(\tilde{\Phi}, \tilde{\mu}) = \frac{\log(0.8\sqrt{2})}{m + 1} > 0 = \chi_1(\tilde{\Phi}, \delta_p),$$

showing that δ_p was “unlocked”. Therefore the argument of the proof of Proposition B.5 does not apply to the perturbation $\tilde{\Phi}$, and it is possible that these perturbations $\tilde{\Phi}$ admit Riemannian extremal norms (though there is no obvious candidate). So the main property of our example Φ , namely not to possess Riemannian extremal norms, may be fragile. Going beyond this specific example, we ask:

Question B.6. *Let $T: X \rightarrow X$ be a hyperbolic automorphism. Let \mathbb{E} be a 2-dimensional θ -Hölder vector bundle over X . Let $\mathcal{B} \subset \text{Aut}^\theta(\mathbb{E}, T)$ be the set of fiber-bunched irreducible automorphisms, endowed with the θ -Hölder topology. Let $\mathcal{R} \subset \mathcal{B}$ be the subset of automorphisms that admit a Riemannian extremal norm. Is \mathcal{R} dense in \mathcal{B} ? Is the interior of \mathcal{R} dense in \mathcal{B} ?*

Acknowledgements. We are very much indebted to Rafael Potrie for numerous illuminating and influential conversations. We also thank Clark Butler and Kiho Park for interesting discussions, and the referee for corrections and suggestions.

References

- [1] AKIN, E. – *The general topology of dynamical systems*. Graduate Studies in Mathematics, 1. American Mathematical Society, Providence, RI, 1993. (Cited on pp. 8, 33, and 46.)
- [2] ALEKSEEV, V.M.; YAKOBSON, M.V. – Symbolic dynamics and hyperbolic dynamic systems. *Phys. Rep.* 75 (1981), no. 5, 287–325. (Cited on p. 8.)
- [3] ANOSOV, D.V. – Roughness of geodesic flows on compact Riemannian manifolds of negative curvature. (Russian) *Dokl. Akad. Nauk SSSR* 145 (1962), 707–709. (Cited on p. 8.)
- [4] AOKI, N.; HIRAIDE, K. – *Topological theory of dynamical systems. Recent advances*. North-Holland Mathematical Library, 52. North-Holland Publishing Co., Amsterdam, 1994. (Cited on p. 8.)
- [5] ARNOLD, L. – *Random dynamical systems*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 1998. (Cited on pp. 32, 34, and 35.)
- [6] AVILA, A.; BOCHI, J. – A uniform dichotomy for generic $\text{SL}(2, \mathbb{R})$ cocycles over a minimal base. *Bull. Soc. Math. France* 135 (2007), no. 3, 407–417. (Cited on p. 29.)
- [7] AVILA, A.; VIANA, M. – Simplicity of Lyapunov spectra: a sufficient criterion. *Port. Math.* 64 (2007), no. 3, 311–376. (Cited on pp. 9 and 18.)
- [8] AVILA, A.; VIANA, M. – Extremal Lyapunov exponents: an invariance principle and applications. *Invent. Math.* 181 (2010), no. 1, 115–189. (Cited on p. 9.)
- [9] BALL, K. – An elementary introduction to modern convex geometry. *Flavors of geometry*, 1–58, Math. Sci. Res. Inst. Publ., 31, Cambridge Univ. Press, Cambridge, 1997. (Cited on p. 50.)

- [10] BARABANOV, N.E. – On the Lyapunov exponent of discrete inclusions. I. *Automat. Remote Control* 49 (1988), no. 2, part 1, 152–157. (Cited on p. 4.)
- [11] BERGER, M.A.; WANG, Y. – Bounded semigroups of matrices. *Linear Algebra Appl.* 166 (1992), 21–27. (Cited on pp. 5 and 37.)
- [12] BLONDEL, V.D.; TSITSIKLIS, J.N. – The boundedness of all products of a pair of matrices is undecidable. *Systems Control Lett.* 41 (2000), no. 2, 135–140. (Cited on p. 19.)
- [13] BOCHI, J. – *The multiplicative ergodic theorem of Oseledets*. Note available at www.mat.uc.cl/~jairo.bochi/docs/oseledets.pdf (2008). (Cited on p. 24.)
- [14] BOCHI, J. – Ergodic optimization of Birkhoff averages and Lyapunov exponents. *Proc. Int. Cong. of Math. – 2018 Rio de Janeiro*, Vol. 2, 1821–1842. (Cited on p. 3.)
- [15] BOCHI, J.; GOURMELON, N. – Some characterizations of domination. *Math. Z.* 263 (2009), no. 1, 221–231. (Cited on p. 34.)
- [16] BOCHI, J.; MORRIS, I.D. – Continuity properties of the lower spectral radius. *Proc. Lond. Math. Soc.* 110 (2015), no. 2, 477–509. (Cited on p. 4.)
- [17] BOCHI, J.; POTRIE, R.; SAMBARINO, A. – Anosov representations and dominated splittings. *J. Eur. Math. Soc.* 21 (2019), no. 11, 3343–3414. (Cited on pp. 30, 33, and 46.)
- [18] BOCHI, J.; RAMS, M. – The entropy of Lyapunov-optimizing measures of some matrix cocycles. *J. Mod. Dyn.* 10 (2016), 255–286. (Cited on p. 4.)
- [19] BONATTI, C.; DÍAZ, L.J.; VIANA, M. – *Dynamics beyond uniform hyperbolicity. A global geometric and probabilistic perspective*. Encyclopaedia of Mathematical Sciences, 102. Mathematical Physics, III. Springer-Verlag, Berlin, 2005. (Cited on p. 33.)
- [20] BONATTI, C.; GÓMEZ-MONT, X.; VIANA, M. – Généricité d’exposants de Lyapunov non-nuls pour des produits déterministes de matrices. *Ann. Inst. H. Poincaré Anal. Non Linéaire* 20 (2003), no. 4, 579–624. (Cited on pp. 6, 9, 11, and 46.)
- [21] BONATTI, C.; VIANA, M. – Lyapunov exponents with multiplicity 1 for deterministic products of matrices. *Ergodic Theory Dynam. Systems* 24 (2004), no. 5, 1295–1330. (Cited on pp. 9 and 18.)
- [22] BOUSCH, T. – La condition de Walters. *Ann. Sci. École Norm. Sup.* 34 (2001), no. 2, 287–311. (Cited on pp. 3, 5, and 11.)
- [23] BOUSCH, T. – Le lemme de Mañé-Conze-Guivarc’h pour les systèmes amphidynamiques rectifiables. *Ann. Fac. Sci. Toulouse Math.* 20 (2011), no. 1, 114. (Cited on pp. 3 and 24.)
- [24] BOUSCH, T.; JENKINSON, O. – Cohomology classes of dynamically non-negative C^k functions. *Invent. Math.* 148 (2002), no. 1, 207–217. (Cited on p. 3.)
- [25] BOUSCH, T.; MAIRESSE, J. – Asymptotic height optimization for topical IFS, Tetris heaps, and the finiteness conjecture. *J. Amer. Math. Soc.* 15 (2002), no. 1, 77–111. (Cited on p. 3.)
- [26] BOWEN, R. – Periodic points and measures for Axiom A diffeomorphisms. *Trans. Amer. Math. Soc.* 154 (1971), 377–397. (Cited on p. 8.)
- [27] BRESSAUD, X.; QUAS, A. – Rate of approximation of minimizing measures. *Nonlinearity* 20 (2007), no. 4, 845–853. (Cited on pp. 38 and 47.)
- [28] BREUILLARD, E.; FUJIWARA, K. – On the joint spectral radius for isometries of non-positively curved spaces and uniform growth. Preprint [arXiv:1804.00748](https://arxiv.org/abs/1804.00748) (Cited on p. 37.)
- [29] BRIN, M.I.; PESIN, J.A.B. – Partially hyperbolic dynamical systems. (Russian) *Izv. Akad. Nauk SSSR Ser. Mat.* 38 (1974), 170–212. (Cited on p. 9.)
- [30] CICONE, A.; GUGLIELMI, N.; PROTASOV, V. YU. – Linear switched dynamical systems on graphs. *Nonlinear Anal. Hybrid Syst.* 29 (2018), 165–186. (Cited on p. 4.)
- [31] COLONIUS, F.; KLIEMANN, W. – *The dynamics of control*. Systems & Control: Foundations & Applications. Birkhäuser Boston, Inc., Boston, MA, 2000. (Cited on p. 33.)
- [32] CONTRERAS, G. – Ground states are generically a periodic orbit. *Invent. Math.* 205 (2016), no. 2, 383–412. (Cited on p. 3.)
- [33] CONTRERAS, G.; LOPES, A.O.; THIEULLEN, P. – Lyapunov minimizing measures for expanding maps of the circle. *Ergodic Theory Dynam. Systems* 21 (2001), no. 5, 1379–1409. (Cited on p. 3.)

- [34] CONZE, J.P.; GUIVARC'H, Y. – Croissance des sommes ergodiques et principe variationnel. Unpublished manuscript, circa 1993. (Cited on p. 3.)
- [35] CORONEL, D.; NAVAS, A.; PONCE, M. – On bounded cocycles of isometries over minimal dynamics. *J. Mod. Dyn.* 7 (2013), no. 1, 45–74. (Cited on p. 19.)
- [36] CROVISIER, S.; POTRIE, R. – Introduction to partially hyperbolic dynamics. Notes ICTP, 2015. (Cited on pp. 33, 34, and 38.)
- [37] GARIBALDI, E. – *Ergodic optimization in the expanding case: concepts, tools and applications*. SpringerBriefs in Mathematics. Springer, Cham, 2017. (Cited on pp. 3 and 31.)
- [38] GARIBALDI, E.; GOMES, J.T.A. – Aubry set for asymptotically sub-additive potentials. *Stoch. Dyn.* 16 (2016), no. 2, 1660009, 13 pp. (Cited on p. 31.)
- [39] GOURMELON, N. – Adapted metrics for dominated splittings. *Ergodic Theory Dynam. Systems* 27 (2007), no. 6, 1839–1849. (Cited on p. 33.)
- [40] HERMAN, M.-R. – Une méthode pour minorer les exposants de Lyapounov et quelques exemples montrant le caractère local d'un théorème d'Arnol'd et de Moser sur le tore de dimension 2. *Comment. Math. Helv.* 58 (1983), no. 3, 453–502. (Cited on p. 19.)
- [41] HIRSCH, M.W.; PUGH, C.C.; SHUB, M. – *Invariant manifolds*. Lecture Notes in Mathematics, Vol. 583. Springer-Verlag, Berlin-New York, 1977. (Cited on pp. 9 and 33.)
- [42] JENKINSON, O. – Ergodic optimization. *Discrete Contin. Dyn. Syst.* 15 (2006), no. 1, 197–224. (Cited on pp. 3 and 31.)
- [43] JENKINSON, O. – Ergodic optimization in dynamical systems. *Ergodic Theory Dynam. Systems* 39 (2019), no. 10, 2593–2618. (Cited on pp. 3 and 31.)
- [44] JUNGERS, R.M. – *The joint spectral radius: theory and applications*. Lecture Notes in Control and Information Sciences, 385. Springer-Verlag, Berlin, 2009. (Cited on pp. 4 and 23.)
- [45] KALININ, B. – Livšic theorem for matrix cocycles. *Ann. of Math.* 173 (2011), no. 2, 1025–1042. (Cited on pp. 29 and 37.)
- [46] KALININ, B.; SADOVSKAYA, V. – Cocycles with one exponent over partially hyperbolic systems. *Geom. Dedicata* 167 (2013), 167–188. (Cited on pp. 6, 9, 11, 23, and 39.)
- [47] KATOK, A.; HASSELBLATT, B. – *Introduction to the modern theory of dynamical systems*. Encyclopedia of Mathematics and its Applications, 54. Cambridge University Press, Cambridge, 1995. (Cited on pp. 9 and 50.)
- [48] KATZNELSON, Y. – *An introduction to harmonic analysis*. 2nd corrected edition. Dover Publications, Inc., New York, 1976. (Cited on p. 6.)
- [49] KOZYAKIN, V. – An explicit Lipschitz constant for the joint spectral radius. *Linear Algebra Appl.* 433 (2010), no. 1, 12–18. (Cited on p. 29.)
- [50] KRENGEL, U. – *Ergodic theorems*. de Gruyter Studies in Mathematics, 6. Walter de Gruyter & Co., Berlin, 1985. (Cited on p. 2.)
- [51] LOPES, A.O.; THIEULLEN, P. – Sub-actions for Anosov diffeomorphisms. Geometric methods in dynamics. II. *Astérisque* No. 287 (2003), xix, 135–146. (Cited on p. 24.)
- [52] MATHER, J.N. – Action minimizing invariant measures for positive definite Lagrangian systems. *Math. Z.* 207 (1991), no. 2, 169–207. (Cited on p. 31.)
- [53] MORRIS, I.D. – A sufficient condition for the subordination principle in ergodic optimization. *Bull. Lond. Math. Soc.* 39 (2007), no. 2, 214–220. (Cited on p. 18.)
- [54] MORRIS, I.D. – A rapidly-converging lower bound for the joint spectral radius via multiplicative ergodic theory. *Adv. Math.* 225 (2010), no. 6, 3425–3445. (Cited on pp. 4, 5, 19, 31, 34, and 37.)
- [55] MORRIS, I.D. – Mather sets for sequences of matrices and applications to the study of joint spectral radii. *Proc. Lond. Math. Soc.* (3) 107 (2013), no. 1, 121–150. (Cited on pp. 2, 4, and 31.)
- [56] OREGÓN-REYES, E. – A new inequality about matrix products and a Berger-Wang formula. Preprint [arXiv:1710.00639](https://arxiv.org/abs/1710.00639) (Cited on p. 37.)
- [57] PESIN, YA. – *Lectures on partial hyperbolicity and stable ergodicity*. Zurich Lectures in Advanced Mathematics. European Mathematical Society (EMS), Zürich, 2004. (Cited on p. 14.)

- [58] PHILIPPE, M.; ESSICK, R.; DULLERUD, G.E.; JUNGERS, R.M. – Stability of discrete-time switching systems with constrained switching sequences. *Automatica* 72 (2016), 242–250. (Cited on pp. 4 and 30.)
- [59] PINTO, A.A.; RAND, D.A. – Smoothness of holonomies for codimension 1 hyperbolic dynamics. *Bull. London Math. Soc.* 34 (2002), no. 3, 341–352. (Cited on p. 9.)
- [60] PUGH, C. – On arbitrary sequences of isomorphisms in $R^m \rightarrow R^m$. *Trans. Amer. Math. Soc.* 184 (1973), 387–400. (Cited on p. 10.)
- [61] PUGH, C.; SHUB, M.; WILKINSON, A. – Hölder foliations. *Duke Math. J.* 86 (1997), no. 3, 517–546. Correction: *ibid.* 105 (2000), no. 1, 105–106. (Cited on p. 10.)
- [62] PUGH, C.; SHUB, M.; WILKINSON, A. – Hölder foliations, revisited. *J. Mod. Dyn.* 6 (2012), no. 1, 79–120. (Cited on p. 7.)
- [63] QIU, L.; ZHANG, Y.; LI, C.-K. – Unitarily invariant metrics on the Grassmann space. *SIAM J. Matrix Anal. Appl.* 27 (2005), no. 2, 507–531. (Cited on p. 12.)
- [64] ROTA, G.-C.; STRANG, G. – A note on the joint spectral radius. *Indag. Math.* 22 (1960), 379–381. (Cited on pp. 2, 3, and 18.)
- [65] RUELLE, D. – *Thermodynamic formalism*. Encyclopedia of Mathematics and its Applications, 5. Addison-Wesley Publishing Co., Reading, Mass., 1978. (Cited on p. 8.)
- [66] SAKAI, K. – Shadowing properties of \mathcal{L} -hyperbolic homeomorphisms. *Topology Appl.* 112 (2001), no. 3, 229–243. (Cited on pp. 8 and 48.)
- [67] SAVCHENKO, S.V. – Homological inequalities for finite topological Markov chains. *Funct. Anal. Appl.* 33 (1999), no. 3, 236–238. (Cited on p. 3.)
- [68] SCHMELING, J.; SIEGMUND-SCHULTZE, R. – Hölder continuity of the holonomy maps for hyperbolic basic sets I. *Ergodic theory and related topics, III (Güstrow, 1990)*, 174–191, Lecture Notes in Math., 1514, Springer, Berlin, 1992. (Cited on p. 9.)
- [69] SIGMUND, K. – On minimal centers of attraction and generic points. *J. Reine Angew. Math.* 295 (1977), 72–79. (Cited on p. 33.)
- [70] SMALE, S. – Differentiable dynamical systems. *Bull. Amer. Math. Soc.* 73 (1967), 747–817. (Cited on p. 8.)
- [71] STEWART, G.W. – *Matrix algorithms. Vol. I: Basic decompositions*. SIAM, Philadelphia, PA, 1998. (Cited on p. 34.)
- [72] VIANA, M. – Almost all cocycles over any hyperbolic system have nonvanishing Lyapunov exponents. *Ann. of Math.* 167 (2008), no. 2, 643–680. (Cited on pp. 6 and 46.)
- [73] WALTERS, P. – *An introduction to ergodic theory*. Graduate Texts in Mathematics, 79. Springer-Verlag, New York-Berlin, 1982. (Cited on p. 48.)
- [74] WIRTH, F. – The generalized spectral radius and extremal norms. *Linear Algebra Appl.* 342 (2002), 17–40. (Cited on pp. 4, 5, and 29.)

Jairo Bochi
 Facultad de Matemáticas, Pontificia Universidad Católica de Chile
www.mat.uc.cl/~jairo.bochi
jairo.bochi@mat.uc.cl

Eduardo Garibaldi
 IMECC, Unicamp
www.ime.unicamp.br/~garibaldi
garibaldi@ime.unicamp.br