A BASIC TO C^1 CONNECTING

LAN WEN AND ZHIHONG XIA

ABSTRACT. We give a proof for a basic C^1 perturbation theorem we assumed in a proof of a C^1 connecting lemma.

1. Introduction

We give a proof for a basic C^1 perturbation theorem we assumed in a proof of a C^1 connecting lemma. Let M be a compact manifold without boundary, and $f\colon M \to M$ be a diffeomorphism. Denote $\mathrm{Diff}^1(M)$ the set of diffeomorphisms of M, endowed with the C^1 topology. The problem of C^1 connecting two orbits was raised by Pugh in [P2]. It asks whether the positive orbit of a point p and the negative orbot of another point q can get connected by a C^1 perturbation if $\omega(p)$ intersects $\alpha(q)$, that is, if they are nearly connected at the first place. An affirmative answer to this would be called a C^1 connecting lemma. The C^1 connecting problem is of a fundamental importance and many authors have made important contributions to this problem. For a more complete introduction to the C^1 connecting problem the reader is refered to [H] or [W–X]. A surprising breakthrough came recently with Hayashi [H] who established a general C^1 connecting lemma, which has played a crucial role in proving the C^1 stability conjecture of Palis and Smale for flows (See [H] and [W2]). For a simpler proof of the C^1 connecting lemma the reader is refered to [W–X].

In the proof of the C^1 connecting lemma given in [W–X], we assumed a basic C^1 perturbation theorem, which can be extracted from the work of Liao, Pugh and Robinson (see [L], [P1], and [P–R]) on the C^1 closing lemma. This is Theorem A below, which forms a cornerstone in [W–X] for proving the C^1 connecting lemma. The statement of Theorem A is technical and needs some introductions, and hence is postponed to §2. Due to its importance there is a need for a proof for Theorem A on its own right. In this paper we give such a proof for Theorem A.

We came to the statement of Theorem A before we could prove it. Shantao Liao and Clark Robinson then confirmed to us that Theorem

Both authors are supported by the National Science Foundation and the Chinese National Natural Science Foundation.

Both authors are supported in part by National Science Foundation.

A was stated right, which encouraged us very much. Charles Pugh read through a preprint of the present paper and gave us many good suggestions. Let us take this opportunity to thank them for all of their help.

In §2 we introduce Theorem A. In §3 we introduce Theorem B, which is an alternate formulation of Theorem A using ellipsoids, and is easier to prove than Theorem A. In §4 we describe the main ideas of the proof of Theorem B. In §5 we pvove Theorem B.

2. An introduction for Theorem A

We introduce Theorem A in this section. Its formulation uses a geometrical notion called ε-kernel transition, which is due to Mai [M] and is the basic pattern for the C^1 perturbations constructed below. A detailed introduction for ε-kernel transitions can be found in [W–X]. This way of constructing perturbations actually appeared very early ([P1]). It is just the notion of ε-kernel transition that appeared relatively late ([M], [W1]). For convenience we quote in this section some of the relevant material from [W–X]. First we define ε-kernel lifts which serve as the basic elements of our C^1 perturbations.

Let $B \subset \mathbb{R}^m$ be a closed ball with radius r and let $\varepsilon > 0$. We denote εB the ball of the same center and of radius εr. We call εB the *ε-kernel* of B. Thus the number ε here tells a relative ratio but not an absolute size. For any x and y in the interior of B, there is a (in fact many) C^∞ diffeomorphism $h\colon \mathbb{R}^m \to \mathbb{R}^m$ that is identity outside B, while takes x to y. If x and y are in εB, we call such an h an *ε-kernel lift* that lifts x to y, supported on B. The following simple but foundamental lemma tells how ε controls the first devivatives of $h - id$ for certain ε-kernel lifts h. The formal formulation of this fact with the proof on manifolds can be found in [P-R, Theorem 6.1].

Lemma 2.1. *For any $\beta > 0$, there is an $\varepsilon > 0$ such that for any closed ball B in $\mathbb{R}^m$, and any x and y in εB, there is an ε-kernel lift h that lifts x to y, supported on B, such that all partial derivatives of $h - id$ have absolute values less than β.*

Proof. The proof is easy and hence omitted. There is a proof for this lemma in [W–X].

Roughly, the number ε controls the size of the first derivatives of $h - id$. Note that the radius r of B is not mentioned in the statement of lemma 2.1, which clearly controls the C^0 size of $h - id$. Therefore the ε-kernel lift h can be defined to be C^1 close to the identity if both ε and r are small, and the composition $h \circ f$ hence gives a C^1 perturbation of f. The C^1 perturbations used in this paper will be a composition of f

with a finitely many this kind of ε-kernel lifts with disjoint supports. By virtue of Lemma 2.1, we will not mention the ε-kernel lift h explicitly, but only mention the ball B and the two points $x,\ y \in \varepsilon B$. Whenever such B, x, and y are specified, we can put on a suitable ε-kernel lift h at any time. In this way we define ε-kernel avoiding transitions now, which are the basic patterns of C^1 perturbations used below. Let $V_0,\ V_1,\ \cdots,\ V_n,\ \cdots$, be a sequence of m-dimensional inner product spaces, and $T_n\colon V_n \to V_{n-1},\ n = 1, 2, \cdots$, be a sequence of linear isomorphisms. Let $\varepsilon > 0,\ x, y \in V_0,\ L \in \mathbb{N},\ Q \subset V_0$, and $G \subset V_0$ be given. By an *ε-kernel avoiding transition* of $\{T_n\}$ from x to y of length L, contained in Q, avoiding G we mean $L+1$ points $c_n \in V_n$, $0 \le n \le L$, together with L balls $B_n \subset V_n,\ 0 \le n \le L-1$, such that

(1) $c_0 = y,\ c_L = F_L^{-1}(x)$, where $F_n = T_1 \circ T_2 \circ \cdots \circ T_n$.
(2) $c_n \in \varepsilon B_n, T_{n+1}(c_{n+1}) \in \varepsilon B_n, 0 \le n \le L-1$.
(3) $B_n \subset F_n^{-1}(Q), 0 \le n \le L-1$.
(4) $B_n \cap F_n^{-1}(G) = \emptyset, 0 \le n \le L-1$.

Roughly, a transition of length L consists of $L+1$ points that form a pseudo orbit with L jumps. The associated L balls make these L jumps ε-kernel lifts. Of course every pseudo orbit like this would be an ε-kernel transition with respect to certain balls. The containing set Q and the avoidence set G then serve as constraints put on the transition. Note that the terminologies defined here are abreviated ones. Such an ε-kernel transition actually is from $F_L^{-1}(x)$ to y, and is contained in the tube $\bigcup_{n=1}^{L} F_n^{-1}(Q)$, and is avoiding a set of orbital arcs $\bigcup_{n=1}^{L}(G)$.

It would be interesting to see what these V_n and T_n have to do with M and f. Applied to the manifold via some standard linearization along a finite orbit of length L, these V_n, $n = 0, 1, \cdots, L-1$, simply correspond to disjoint neighborhoods of the iterates along a backward orbit of f, and these T_n simply correspond to f itself. Thus the transition transits a point from one orbit of f to another via L lifts which form a pseudo orbit. If Q is small (which bounds the C^0 size of the perturbation), and if ε is small too, then the transition gives a C^1 small perturbation. These details are not concerned with us in this paper however, because Theorem A we are going to prove is in the framework of V_n and T_n. For more illustrations about these details the reader is refered to [W–X].

Using the notion of ε-kernel transition we now formulate Theorem A. Let V be an m-dimensional inner product space and $e = (e_1, e_2, \cdots, e_m)$ be an orthonormal basis o V. An *e-box* Q of center $x \in V$ and of certain size $(\lambda_1, \lambda_2, \cdots, \lambda_m)$ is defined as

$$Q = \{y \in V|\ |y_i - x_i| \le \lambda_i,\ 1 \le i \le m\},$$

where x_i and y_i are coordinates of x and y, respecting the basis e. For $\alpha > 0$, define

$$\alpha Q = \{y \in V|\ |y_i - x_i| \le \alpha\lambda_i,\ 1 \le i \le m\}.$$

We say that a box Q' is *of type* Q, if

$$Q' = z + \alpha Q$$

for some $z \in V$ and some $\alpha > 0$.

Theorem A. *For any sequence of isomorphisms* $T_n \colon V_n \to V_{n-1},\ n = 1, 2. \cdots$, *there is an orthonormal basis* $e = (e_1, e_2, \cdots, e_m)$ *in* V_0 *such that for any* $\varepsilon > 0$, *and any* $0 < \alpha < 1$, *there is an e-box* A *and an integer* $L \in \mathbb{N}$ *such that for any e-box* Q *of type* A *and any two points* $x, y \in \alpha Q$, *there is an* ε*-kernel transition* $c_0, c_1, \cdots, c_L; B_0, B_1, \cdots, B_{L-1}$ *of* $\{T_n\}$ *from* x *to* y *of length* L, *contained in* Q. *Moreover, the radius of* B_n *is less than or equal to half of the distance between* $\partial(F_n^{-1}(D))$ *and* $\partial(F_n^{-1}(\alpha D))$.

This theorem constitutes a basis for the proof of the C^1 connecting lemma given in [W–X]. It asserts the existence of an ε-kernel transition that satisfies certain requirments. In particular, Theorem A requires that the support balls should be uniformly small in ratio as the last sentence of Theorem A claims. More precisely, in addition to that the ball B_n should be contained in $F_n^{-1}(Q)$, the last sentence of Theorem A requires that the ball B_n should be also small enough relative to the parallelepiped $F_n^{-1}(Q)$ that, via a parallel translation, it can be inserted into the gap between the two parallelepipeds $F_n^{-1}(Q)$ and $F_n^{-1}(\alpha Q)$. This turs out to be crucial to the proof of the C^1 connecting lemma given in [W–X].

There is a beautiful proof for the C^1 closing lemma by Mai [M], which is then generalized to some non-invertible maps by Wen [W1]. The proof is fairly simple. It is not based on Theorem A, because the radii of the balls used there do not have to satisfy the last requirement of Theorem A.

3. An alternate formulation of Theorem A using ellipsoids

We give an alternate formulation of Theorem A using ellipsoids in this section. This is Theorem B below. Theorem A is easier to use, but Theorem B is easier to prove. It is Theorem B that we prove in this paper.

It is well know that every linear isomorphism T factors as PQ where Q is orthogonal and P is positive definite symmetric. Hence the image of a round ball under T is an ellipsoid. If E is an ellipsoid in V_0, for a number $\alpha > 0$, we denote αE the ellipsoid with the same center as E and with a multiple α. More precisely, if x is the center of E, then αE is defined as $\alpha(E - x) + x$. If $\alpha < 1$, we say that αE is the α-kernel of E. We say that an ellipsoid E' is *of type* E, if

$$E' = z + \alpha E$$

for some $z \in V_0$ and some $\alpha > 0$. For an ellipsoid E of center x and a box D of center y we write $E \leq D$ if $E - x \subset D - y$.

The definition of ε-kernel avoiding transition for a sequence of isomorphisms $\{T_n\}$ can be reformulated in terms of ellipsoids. That is, for $n \geq 1$, we map the round balls B_n in V_n by their corresponding linear isomorphisms F_n to get ellipsoids Z_n in V_0, and also map the points c_n by F_n to get some points a_n in V_0. The sequence of linear isomorphisms $\{T_n\}_{n=1}^{\infty}$ gives a sequence of m-dimensional ellipsoids $\{E_n\}_{n=1}^{\infty}$ (up to type) with the subscripts running from 1 to ∞. For $n = 0$, we do not have the corresponding linear isomorphism T_0 in the definition, and a natural definition for the ellipsoid E_0 would be simply B_0 itself. This would give a constraint that the ellipsoid E_0 in this sequence should be always a round ball. For a technical reason below we will not take this constraint but allow the general case that E_0 is an m-dimensional ellipsoid in V_0 of any type. This gives the following definition for ε-kernel avoiding transition, in terms of ellipsoids. Let $\{E_n\}_{n=0}^{\infty}$ be a sequence of ellipsoids in V_0, and let $\varepsilon > 0$, $x, y \in V_0$, $L \in \mathbb{N}$, $Q \subset V_0$, $G \subset V_0$ be given. By an *ε-kernel transition* of $\{E_n\}$ from x to y of length L, contained in Q, avoiding G we mean $L+1$ points a_n in V_0, $0 \leq n \leq L$, together with L ellipsoids Z_n of type E_n in V_0, $0 \leq n \leq L-1$, such that

(1) $a_0 = y,\ a_L = x.$
(2) $a_n, a_{n+1} \in \varepsilon Z_n,\ \ 0 \leq n \leq L-1.$
(3) $Z_n \subset Q,\ \ 0 \leq n \leq L-1.$
(4) $Z_n \cap G = \emptyset,\ \ 0 \leq n \leq L-1.$

If, in addition, the following condition

(5) a_n is the center of Z_n for all $0 \leq n \leq L-1$

is satisfied, then we call this transition *centerwise*.

At first glance the transition so defined via ellipsoids may sound messy, because the ellipsoids $Z_0, Z_1, ..., Z_{L-1}$ are not disjoint, and perhaps overlap very much. But the point is that the way they overlap is irrelevant to us. The balls B_n will be mutually disjoint on the manifold anyway. What concerns us are just two things: every point a_n should be in the ε-kernel of the ellipsoid of the same subscript (i.e. $a_n \in \varepsilon Z_n$), and should also be in the ε-kernel of the ellipsoid of the previous subscript (i.e. $a_n \in \varepsilon Z_{n-1}$). In the centerwise case which is the only case we consider in this paper (one needs to consider the general non-centerwise case in the proof of the C^1 connecting lemma) the situation is even cleaner, and we just concentrate on one thing: the center of every ellipsoid should be in the ε-kernel of the previous ellipsoid. Thus when we look at Z_n, we only need to locate a_{n+1} from its ε-kernel, and not to care about all the other ellipsoids. This beautiful idea which has proved very effective is due to Mai [M].

Using this alternate definition of ε-kernel transition we can reformulate Theorem A into the following Theorem B.

Theorem B. *For any sequence of m-ellipsoids $\{E_n\}_{n=0}^{\infty}$ in V_0, there is an orthonormal basis $e = (e_1, e_2, ..., e_m)$ of V_0 such that for any $\varepsilon > 0$, and any $0 < \alpha < 1$, there is an e-box A and an integer L such that for any e-box Q of type A and any two points $x, y \in \alpha Q$, there is a centerwise ε-kernel transition $Z_0, Z_1, \cdots, Z_L$ of $\{E_n\}$ from x to y of length L, contained in Q. Moreover, $Z_n \leq (1-\alpha)Q$ for all $0 \leq n \leq L-1$.*

Clearly, Theorem B implies Theorem A, and we will prove Theorem B instead in §5.

4. The main ideas of the proof of Theorem B

Before proceeding to the proof of Theorem B, let us single out some of the main ideas involved in the proof. They are geometrically simple, but are the keys to the whole proof. First we make a remark on terminologies. We are going to find a centerwise transition from x to y. Note that x corresponds to the subscript L and y corresponds to the subscript 0. We construct the transition backwards. That is, we start with $y = c_0$. First we find the ellipsoid Z_0 centered at y, then we locate the point c_1 from its ε-kernel εZ_0. (Sometimes we say informally that this makes a "progress" from c_0 to c_1.) Then we find the ellipsoid Z_1 centered at c_1, and locate the point c_2 from εZ_1, etc. Thus we go from the subscript 0 to the subscript L. In other words, the thansition is *from* the subscript L *to* the subscript 0 according to the definition (this definition is natural because it represents on the manifold a transition that bring $f^{-L}(x)$ to y gradually), but we will search for these ellipsoids Z_n and these points c_n *from* the subscript 0 *to* the subscript L. This is what we mean "backwards". To avoid confusion we call below an ε-kernel transition from x to y an ε-kernel *movement* from y to x. They are exactly the same collection of points and ellipsoids. The only difference is the way of using the word "from" and "to".

Now we describe some of the main ideas involved in the proof of Theorem B.

(A). Let C be an m-cube of center a and of half-size $\delta > 0$. For any m-ellipsoid E, the length of every greatest axis (An ellipsoid may have infinitely many greatest axes) of the largest ellipsoid Y contained in C that has type E and center a is no less than δ (This is because C contains a round ball of radius δ). Hence, if b is an end point of such a greatest axis of Y, then the point $a' = a + \varepsilon(b - a)$ is in εY. This means, in the direction u of a greatest axis of Y, we can always get a progress of ε times the half-size of the bounding cube, via an ε-kernel lift. Of course, in other directions, say a direction that is orthogonal to u, we may not get such a progress. This is clear through the following 2-dimensional figure, where the ellipsoid is thin (Here and below, we use the word "thin" in the sense of type but not in an absolute sense. That is, "thin" means that the ratio of the greatest axis over the least axis is large), and we can not get in $e^{\perp}$-direction such a progress.

By replacing the bounding square C with a rectangle A, we can still get such a progress as long as the direction of the axes of the ellipsoid is correct. More precisely, no matter how thin an ellipsoid Y is, *as long as the direction u of the greatest axis of Y is exactly parallel to the longer side of a rectangle* A, and as long as A is sufficiently thin (here we use the word "thin" for rectangles in the same relative sense as for ellipsoids), we can still get in the direction $e^{\perp}$ a progress ε-proportional to the half of the shorter side of A. This is illustrated in the following figure, where we first draw two parallel lines tangent to Y at the two end points of the shorter axis of Y, then close up the infinite strip to form a rectangle A so that A contains Y.

However, if the direction of the greatest axis u is not parallel to the direction e of the longer side of A, the situation gets much more delicate. In fact, no matter how approximately parallel u is to e (as long as not exactly), there are some thin ellipsoids E such that no matter how thin the rectangle A is (even if infinitely thin as two parallel lines), we can not get such a proportional progress in $e^{\perp}$ direction. This is illustrated in the following figure, where the ellipsoid type E is so thin that the enclosed ellipsoid Y looks like an interval that touches the two parallel lines already, and hence stretching the longer sides of A no longer helps. This is the most delicate point we have to handle in the proof of Theorem A below in §5. This seems to be also the main difficulty discussed analytically in [P–R] on arbitrary sequence of isomorphisms. The next observation suggests a way to handle this delicate point geometrically.

Here comes the crucial observation:

(C). For any orthogonal basis $(e,\ e^{\perp})$ and any ellipsoid type E (no matter how thin it is and whatever directions of its axes are in), by adjusting the shape of the rectangle A, though we might not be able to get a progress of ε times half of the shorter side of A towards the $e^{\perp}$ direction as illustrated above, we can move to *some* direction so that, *projecting to* $e^{\perp}$, we do get such a progress. In the above figure, the direction from a to the tangency τ is such a direction. (The rectangle A should of course be thin enough so that the tangencies appear on the longer sides, but not the shorter sides. We can guarantee this by first drawing two parallel infinite lines in e-direction. Then the largest ellipsoid Y of type E that is contained in the infinite strip must be tangent to the two lines. Then we close up the rectangle to get the desired shape of A.) Briefly, by adjusting the shape of A, though we might not be able to get a desired progress in $e^{\perp}$ direction, but projectively, we can! This suggests that, to move from a point y to a point x, we can first try to move to the line e_m-axis$+x$ as the following figure shows, where e_m is a limit direction of the greatset axises of E_n.

The price paid will be that, when we get to a point z on that line, z may become a lot farther from x than y is. But according to (A), e_m is approximately the direction in which we have the strongest movability. It is easy to see that we can get near x from z by following a collection of moves of fashion (A).

(D). There is a final detail to be taken care of. That is, since the greatest axes of the ellipsoids E_n are only approximately, but not exactly, in the direction e_m, by the expected collection of moves of fashion (A) we can get only near to the point x, but probably not exactly x. And, after choosing subsequences, the type of the ellipsoids left may all be extremely thin so that it would be unclear if we can use them to move onto the (even nearby) point x. This difficulty is solved by leaving the first ellipsoid E_0 spared. That is, we leave aside E_0 for the most part of the proof, and using the sequence $\{E_n\}_{n=1}^{\infty}$ to make all the above expected moves to get from y to a point w near x. Then we make a parallel translation by adding the vector $x - w$ that takes the whole movement into a new movement. This means all the $L+1$ points and all the L associated balls in the movement are each shifted over by the vector $x - w$. It is clear that a parallel translation of an ε-kernel movement is still an ε-kernel movement, and this new movement will be from the point $u = y + (x - w)$ onto the point $w + (x - w)$ which is x. Since E_0 has fixed ballicity and u can be arbitrarily near y, using E_0 we may get from y to u as a preliminary (or, the first) move.

5. The proof of Theorem B

In this section we prove Theorem B. First we state an elementrary lemma related to observation (C) of §4. While observation (C) is for the case of dimension 2, the following lemma treats the case of general dimensions. The lemma is almost self-evident, but really exhibits the geometrical core of the proof for Theorem A. Let V be an m-dimensional inner product space, and

$$V = W \oplus W^{\perp}$$

an orthogonal splitting with W 1-dimensional. Denote

$$\pi : V \to W^{\perp}$$

the orthogonal projection. Let E be an m-dimensinal ellipsoid in V. Then πE is an $(m-1)$-dimensional ellipsoid in $W^{\perp}$. Let Z' be any $(m-1)$-ellipsoid in $W^{\perp}$ of type πE, and A' be any rectangular $(m-1)$-box in $W^{\perp}$ that has the same center as that of Z' and contains Z'. $\pi^{-1}(A')$ is an infinite rectangular cylinder in V. If c' is the center of A', then $\pi^{-1}(c')$ is the central axis of this infinite cylinder.

Lemma 5.1. *Let a be a point on $\pi^{-1}(c')$ and P be the largest m-ellipsoid contained in $\pi^{-1}(A')$ that has type E and is centered at a. Then $Z' \subset \pi(P)$.*

Proof. Since Z' is of type πE, there must be an m-ellipsoid Z of type E such that $\pi Z = Z'$. Hence $Z \subset \pi^{-1}(Z') \subset \pi^{-1}(A')$. Via a parallel translation along the line $\pi^{-1}(c')$ if necessary, we may assume Z has center a. Then $Z \subset P$ because P is the largest m-ellipsoid with these properties. Taking projections proves the lemma.

We remark that it is important here to have a full preimage of A' under π, that is, to have an infinite cylinder. If we use, instead of the infinite cylinder $\pi^{-1}A'$, a finite cylinder Σ that has the same base A', but is not sufficiently long in W-direction, then the largest ellipsoid P contained in Σ that has type Z and is centered at a may not at all have the property $Z' \subset \pi(P)$. This is actually the observation we made in (B) of §4. Anyway let us draw a figure again to illustrate this.

Now we prove Theorem B. The proof goes by induction. First we prove Theorem B for the case $m = 1$. In this case ellipsoids and boxes all reduce to intervals. Given a sequence of intervals $\{E_n\}$. Up to the type (we use E_n only up to their types), it is only one interval. We take a unit vector e_1 as the orthonormal basis. There are two unit vectors in V_0. We just take either one of them. Let $\varepsilon > 0$, $0 < \alpha < 1$ be given. The ellipsoid A will be just an interval too. The number L can be taken as

$$L = [4\alpha/(\varepsilon(1-\alpha)] + 1,$$

where $[\cdot]$ denotes the integer part.

We verify that A and L correspond to ε and α correctly. Let Q be any interval, and x, y be any two points in αQ. Without loss of generality we may assume that Q has length 2, and x, y are the two points on Q which are $1-\alpha$ from the ends of Q. We may also assume that x is on the e_1-direction from y. Let

$$a_0 = y,$$

$$a_1 = a_0 + \varepsilon(1-\alpha)e_1,$$

$$\ldots .$$

And let

$$Z_n = [a_n - (1-\alpha),\ a_n + (1-\alpha)].$$

Then it is easy to verify that this gives a centerwise ε-kernel movement from y to x, which is just a centerwise ε-kernel transition from x to y, that satisfies all the requirements of Theorem B. We only remark that the last move may deal with some tip, hence may not be exactly as given by the formulas.

Now we assume that the theorem is proved for dimension $m-1$. We prove it for dimension m.

Let $\{E_n\}_{n=0}^{\infty}$ be a sequence of m-dimensional ellipsoids in V_0. Let $u(E_n)$ be a unit vector along one of the greatest axes of E_n. Let e_m be a limit point of $\{u(E_n)\}$. This is a unit vector. Write

$$V_0 = [e_m] \oplus [e_m]^{\perp},$$

where $[e_m]$ denotes the one dimensional subspace of V_0 spanned by e_m, and $[e_m]^{\perp}$ denotes the orthogonal complement of $[e_m]$. Denote

$$\pi : V_0 \to [e_m]^{\perp}$$

the orthogonal projection.

Note that $\{\pi E_n\}_{n=0}^{\infty}$ is a sequence of $(m-1)$-dimensional ellipsoids in $[e_m]^{\perp}$. As discussed above, we first leave aside E_0. For the sequence $\{\pi E_n\}_{n=1}^{\infty}$, by induction hypothesis, there is an orthonormal basis $e' = (e_1, \ldots, e_{m-1})$ in $[e_m]^{\perp}$ with the properties stated in Theorem B. We remark that the first ellipsoid E_1 in this sequence may not be a round ball at all, and this is the reason why we have defined the ε-kernel transition of a sequence of ellipsoids in the way that the first ellipsoid could not be a round ball. Let

$$e = (e_1, \dots, e_{m-1}, e_m).$$

Then e is an orthonormal basis in V_0. We prove that this is a desired basis to us.

Let $\varepsilon > 0$ and $0 < \alpha < 1$ be given. We need to find an e-box A in V_0 and an integer L that satisfy the conditions stated in Theorem B. By induction hypothesis, corresponding to the sequence of $(m-1)$-dimensional ellipsoids $\{\pi E_n\}_{n=1}^{\infty}$ and the obtained orthonormal basis e' in $[e_m]^{\perp}$, for the same two numbers $\varepsilon > 0$ and $0 < \alpha < 1$, there is an e'-box A' and an interger L' such that for any e'-box Q' of type A', and any two points $x', y' \in \alpha Q'$, there is a centerwise ε-kernel transition $Z'_0, Z'_1, \cdots, Z'_{L'}$ of $\{\pi E_n\}_{n=1}^{\infty}$ from x' to y' of length L', contained in Q'. Moreover, $Z'_n \leq (1-\alpha)Q'$ for $1 \leq n \leq L'$. Since the $(m-1)$-box A' will be used only up to its type, we may assume that A' is centered at the origin. We might have also assumed that A' has a kind of normalized size, but this does not seem to simplify the notation much. So we just denote the size of A' as $\lambda_1, \cdots, \lambda_{m-1}$, and denote

$$\mu = \min\{\lambda_1, \cdots, \lambda_{m-1}\}.$$

Consider $\pi^{-1}A'$. This is an infinite cylinder with base A'. For each $n = 1, 2, \cdots$, denote Y_n the largest m-dimensional ellipsoid contained in $\pi^{-1}A'$ that has type E_n and is centered at origin (the same center as A'). Choose $d_n > 0$ so large that the m-dimensional e-box $A' \times [-d_n, d_n]$ contains Y_n. Let

$$\lambda_m = \max\{\lambda_1, \cdots, \lambda_{m-1}, \quad \mu + \Sigma_{n=1}^{L'} d_n\}.$$

Note that μ is still the minimum of $\{\lambda_1, \lambda_2, \cdots, \lambda_m\}$. Also note that $\lambda_m \geq d_n$ for all $n = 1, \cdots, L'$. Then let

$$A = [\lambda_1, \cdots, \lambda_{m-1}, \lambda_m].$$

This determines the desired e-box A. To get the desired integer L, we write

$$\delta = \mu(1-\alpha)/4,$$

and let

$$k = [\lambda_m/(\varepsilon\delta)] + 1.$$

For any finite subsequence $n_1, n_2, \cdots, n_k$, we will denote below

$$c_{n_1} = \varepsilon\delta u_{n_1},$$

$$c_{n_i} = c_{n_{i-1}} + \varepsilon\delta u_{n_i}, i = 2, \cdots, k.$$

That is, we start with the origin, and move a distance of $\varepsilon\delta$ in u_{n_i} directions succesively. Since e_m is a limit point of $u_n = u(E_n)$, and since k is fixed, we can choose a finite subsequence

$$n_1, \cdots, n_k$$

after $L' + 1$ such that u_{n_i} is so close to e_m that c_{n_i} is in the $\varepsilon\delta/b_0$-neighborhood of the point

$$c^0_{n_i} = i\varepsilon\delta e_m, i = 1, \cdots, k,$$

where b_0 denotes the bolicity of E_0 (i.e. the ratio of the greatest axis over the least axis for E_0). Then let

$$L = n_k.$$

Note that A and L are constructed out of $\{E_n\}$, ε, and α only.

Now we verify that A and L as chosen satisfy the requirements stated in Theorem B for the given sequence $\{E_n\}$, and for the given two numbers ε and α.

Let Q be any e-box of type A, and let x and y be any two points in αQ. We need to verify that there is a centerwise ε-kernel transition $Z_0, Z_1, \cdots, Z_{L-1}$ of $\{E_n\}_{n=0}^{\infty}$ from x to y of length L, contained in Q. Moreover, the ellipsoids Z_n have to satisfy the inequality $Z_n \le (1-\alpha)Q$ for $0 \le n \le L-1$. It is easy to see that if Q has this property, then any e-box of type Q has the same property too. Thus we may simply assume that Q is just A itself. Moreover, we may assume that the two points x and y are furthest apart in αA. In particular, we may assume that $y \in [e_m]^{\perp} \times \{-\alpha\lambda_m\}$ and $x \in [e_m]^{\perp} \times \{\alpha\lambda_m\}$. Let

$$\pi^* : V_0 \to [e_m]^{\perp} \times \{-\alpha\lambda_m\}$$

denote the orthogonal projection (We might have used π to denote π^* by an abuse of notation, but let us use π^* any way), and let

$$x^* = \pi^*(x), \quad \text{and} \quad A^* = \pi^*(A).$$

Now y and x^* are both in αA^*. But A^* and A' differ only by a translation, and $\{\pi^* E_n\}$ and $\{\pi E_n\}$ differ only by a translation too, so A^* has the same property stated in the above induction hypothesis with respect to $\{\pi^* E_n\}$ as A' does with respect to $\{\pi E_n\}$. That is, there is a centerwise ε-kernel transition $Z^*_1, \cdots, Z^*_{L'}$ of $\{\pi^* E_n\}_{n=1}^{\infty}$ from x^* to y of length L', supported in A^*. Moreover, $Z^*_n \le (1-\alpha)A^*$ for all $n = 1, \cdots, L'$. It will be more convenient below to call it a movement, because then we can say it is from y to x^*, rather than from x^* to y.

Now we lift the $(m-1)$-dimensional ε-kernel movement $Z^*_1, \cdots, Z^*_{L'}$ in the hyperplane $[e_m]^{\perp} \times \{-\alpha\lambda_m\}$ from y to x^* to an m-dimensional ε-kernel movement $P_1, \cdots, P_{L'}$ in V_0 from y to a point z on the line

$(\pi^*)^{-1}(x^*) = (\pi^*)^{-1}(x)$. Here we use the letter P but not Z because, as observed in (D) of §4, we will need to make a translation later to get our desired transition. The letter Z is reserved for that. As remarked before, we do it backwards. Let $a_1^*, \cdots, a_{L'+1}^*$ be the centers of $Z_1^*, \cdots, Z_{L'+1}^*$, respectively. Thus $a_1^* = y$, $a_{L'+1}^* = x^*$, and $a_{n+1}^* \in \varepsilon Z_n^*$ for all $n = 1, \cdots, L'$.

We first lift Z_1^*. This is an $(m-1)$-dimensional ellipsoid of type $\pi^* E_1$. Moreover, $Z_1^* \leq (1-\alpha)A^*$, which means $Z_1^* \subset (1-\alpha)A^* + a_1^*$. Then, as before, $(\pi^*)^{-1}((1-\alpha)A^* + a_1^*)$ is an infinite cylinder with base $(1-\alpha)A^* + a_1^*$, and $(1-\alpha)Y_1 + a_1^*$ is the largest m-dimensional ellipsoid contained in $(\pi^*)^{-1}((1-\alpha)A^* + a_1^*)$ that has type E_1 and is centered at a_1^*. Then we let

$$p_1 = a_1^* = y, \quad P_1 = (1-\alpha)Y_1 + p_1.$$

Here we use the letter p instead of the letter a also because we reserve the letter a for the final transition. By Lemma 5.1,

$$Z_1^* \subset \pi^*(P_1).$$

This is a crucial fact to what follows.

Note that the size of P_1 is well controled. In fact, by the choice of d_1, we have

$$P_1 \subset ((1-\alpha)A') \times [-(1-\alpha)d_1,\ (1-\alpha)d_1] + a_1^*.$$

Since $\lambda_m \geq d_1$, we have

$$P_1 \subset (1-\alpha)A + a_1^*,$$

that is,

$$P_1 \leq (1-\alpha)A.$$

Having lifted Z_1^* to P_1, we now lift the point a_2^* to a point p_2, which has to be in the ε-kernel of P_1. This is easy now. Since $a_2^* \in \varepsilon Z_1^*$, and since $Z_1^* \subset \pi^*(P_1)$ (Which is the crucial fact observed above), there is indeed a (actually an interval of) point $p_2 \in \varepsilon P_1$ such that $\pi^*(p_2) = a_2^*$.

Then we proceed as before. That is, let

$$P_2 = (1-\alpha)Y_2 + p_2.$$

Then similar arguments yield that

$$Z_2^* \subset \pi^*(P_2),$$

and

$$P_2 \leq (1-\alpha)A.$$

Then since $a_3^* \in \varepsilon Z_2^*$, there is a point $p_3 \in \varepsilon P_2$ such that

$$\pi^*(p_3) = a_3^*.$$

Then we let

$$P_3 = (1-\alpha)Y_3 + p_3,$$

and so on. In this way we get a lifted centerwise m-dimensional ε-kernel movement from y to a point $z \in V_0$ with

$$\pi x = \pi z.$$

As observed in (C) of §4, the price paid is that z may become rather far away from x. But the distance is well controlled. In fact, by the choice of λ_m,

$$\|z - x\| \le 2\lambda_m,$$

and

$$d(z, \partial A) \le (1-\alpha)\mu.$$

Here in the first inequality the upper bound $2\lambda_m$ could be somewhat reduced. But this is all right for our purpose. The second inequality holds because λ_m was chosen so that the total moves $\Sigma(1-\alpha)d_n$ in e_m-direction (even plus $(1-\alpha)\mu$) is no more than the $(1-\alpha)$-gap of A in e_m-direction. The room left in e_m-direction, which is $(1-\alpha)\mu$, seems to be narrow. But this is enough for the following steps.

In what follows, we construct a centerwise ε-kernel movement from z to a point w near x, using the ellipsoids $E_{n_1}, \cdots, E_{n_k}$. This corresponds to the collection of moves of fashion (A) discussed in §4. But this is easy now. Let C be the e-cube of size δ and of center origin, let

$$p_{n_i} = c_{n_i} + z,$$

and let P_{n_i} be the largest ellipsoid of type E_{n_i} contained in the cube $C + p_{n_i}$. Since the broken line $c_{n_1} \cdots c_{n_k}$ has been chosen so close to e_m-axis, it is easy to check that this gives a desired centerwise ε-kernel movement of length $L - L' - 1$ from z to a point w in the $\varepsilon\delta/b_0$ neighborhood of x. Here the ε-kernel lift for an integer n which is strictly between n_i and n_{i+1} is understood as the trivial ε-kernel lift, i.e., no lift. Combined with the movement $P_1, \cdots, P_{L'}$ obtained before, this gives a movement from y to w.

As observed in (D) of §4, there is a final step to be taken care of. That is, we need to translate the whole obtained movement $p_1, \cdots, p_L; P_1, \cdots, P_{L-1}$ from y to w into a movement from $u = y + x - w$ to $w + x - w$ which is x. This is simply done by letting

$$a_n = p_n + x - w, \quad n = 1, 2, \cdots, L,$$

and

$$Z_n = P_n + x - w, \quad n = 1, 2, \cdots, L-1.$$

Now let $a_0 = y$, and let Z_0 be the largest ellipsoid of type E_0 contained in the δ-cube $C + y$. It is easy to see that

$$a_0, a_1, \cdots, a_L; \quad Z_0, Z_1, \cdots, Z_{L-1}$$

finally gives a desired movement from y to x, or, what is the same, a desired transition from x to y. This completes the induction process, and proves Theorem B.

References

[H] S. Hayashi, *Connecting invariant manifolds and the solution of the C^1 stability and Ω-stability conjectures for flows*, Preprint, 1995.

[L] S. T. Liao, *An extension of the C^1 closing lemma*, Acta Sci. Natur. Univ. Pekinensis, 1979, No. 2, 1-41.

[M] J. Mai, *A simpler proof of C^1 closing lemma*, Scientia Sinica, **10** (1986), 1021-1031.

[P1] C. Pugh, *The closing lemma*, Amer. J. Math., **89**(1967), 956-1009.

[P2] C. Pugh, *Against the C^2 closing lemma*, Journal of Differential Equations, **17**(1975), 435-443.

[P-R] C. Pugh & C. Robinson, *The C^1 closing lemma, including Hamiltonians*, Ergod. Th. & Dynam. Sys., **3**(1983), 261-313.

[S] S. Smale, *Differentiable dynamical systems*, Bull. Amer. Math. Soc., **73**(1967), 747-817.

[W1] L. Wen, *The C^1 closing lemma for non-singular endomorphisms*, Ergod. Th. & Dynam. Sys., **11**(1991), 393-412.

[W2] L. Wen, *On the C^1 stability conjecture for flows*, Preprint.

[W–X] L. Wen & Z. Xia, *A simpler proof for the C^1 connecting lemma*, Preprint.

[X] Z. Xia, *Homoclinic points in symplectic and volume-preserving diffeomorphisms*, Commun. Math. Phys. **177**(1996), 435-449.

Department of Mathematics, Peking University, Beijing, 100871 China
E-mail address: `lwen@bepc2.ihep.ac.cn`

Department of Mathematics, Northwestern University, Evanston, IL 60208 USA
E-mail address: `xia@math.nwu.edu`