Centrum voor Wiskunde en Informatica

REPORT*RAPPORT*

MAS

Modelling, Analysis and Simulation

Modelling, Analysis and Simulation

MAS

Entropy-minimising models of surface diffeomorphisms relative to homoclinic and heteroclinic orbits

P.J. Collins

REPORT MAS-E0418 NOVEMBER 2004

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3703

Entropy-minimising models of surface diffeomorphisms relative to homoclinic and heteroclinic orbits

ABSTRACT
In the theory of surface diffeomorphisms relative to homoclinic and heteroclinic orbits, it is possible to compute a one-dimensional representative map for any irreducible isotopy class. The topological entropy of this graph representative is equal to the growth rate of the number of essential Nielsen classes of a given period, and hence is a lower bound for the topological entropy of the diffeomorphism. In this paper, we show that this entropy bound is the infemum of the topological entropies of diffeomorphisms in the isotopy class, and give necessary and sufficient conditions for the infemal entropy to be a minimum.

2000 Mathematics Subject Classification: 37E30; 37C27; 37B40; 37E15.
Keywords and Phrases: Surface diffeomorphism; Homoclinic orbit; Trellis; Topological entropy.
Note: This work was partially funded by Leverhulme Special Research Fellowship SRF/4/9900172.

Entropy-minimising models of surface diffeomorphisms relative to homoclinic and heteroclinic orbits

Pieter Collins *

Centrum voor Wiskunde en Informatica,
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Email: Pieter.Collins@cwi.nl

28 October

Abstract

In the theory of surface diffeomorphisms relative to homoclinic and heteroclinic orbits, it is possible to compute a one-dimensional representative map for any irreducible isotopy class. The topological entropy of this graph representative is equal to the growth rate of the number of essential Nielsen classes of a given period, and hence is a lower bound for the topological entropy of the diffeomorphism. In this paper, we show that this entropy bound is the infemum of the topological entropies of diffeomorphisms in the isotopy class, and give necessary and sufficient conditions for the infemal entropy to be a minimum.

Keywords: Surface diffeomorphism; Homoclinic orbit; Trellis; Topological entropy.
2000 Mathematics subject classification: 37E30; 37C27; 37B40; 37E15

1 Introduction

In the Nielsen-Thurston theory of surface diffeomorphisms (see [3]), it is possible to find a diffeomorphism in each isotopy class which minimises both the number of periodic orbits of each period, and the topological entropy (see [9, 7, 10, 2]). A constructive proof was given independently by Bestvina and Handel [1] and by Franks and Misiurewicz [8]; a one-dimensional representative of the diffeomorphism called a *train track* is computed and the Thurston-minimal model is easily constructed from this.

The case of diffeomorphisms relative to homoclinic orbits to a saddle point is more complicated. It is well-known that any diffeomorphism with a transverse homoclinic point has positive topological entropy, but it is easy to construct examples of such diffeomorphisms with arbitrarily small entropy. Therefore, it may not even be possible to find a diffeomorphism realising the lower bound of the entropy.

In this paper, we discuss the relation between the *Nielsen entropy*, which measures the growth rate of the number of essential Nielsen classes of periodic points, and the infemal entropy in the isotopy class. The Nielsen entropy is equal to the topological entropy of the *graph representative*, which can be computed using the algorithms described in [5]. The key step is to represent the isotopy class relative to the homoclinic/heteroclinic orbits using a *trellis*, which is a subset of the homoclinic tangle of the saddle periodic points. We show that under mild non-degeneracy conditions, the Nielsen entropy is equal to the infemum of the topological entropies in the isotopy

*This work was partially funded by Leverhulme Special Research Fellowship SRF/4/9900172.

class. This result can be considered an optimality result for the computational techniques of the trellis theory.

The proof that the Nielsen entropy is infemal involves the construction of uniformly-hyperbolic diffeomorphisms realising the entropy bound arbitrarily closely. In certain cases, it is possible to construct a uniformly-hyperbolic diffeomorphism realising the entropy bound. We show that the only obstruction to the construction of a uniformly hyperbolic diffeomorphisms realising the entropy bound is the existence of *almost wandering segments*; intervals of (un)stable manifold which must contain intersection points under some iterate of any representative diffeomorphism, but for which no lower bound on the number of iterates needed exists. We further show that if the Nielsen entropy is realised and the system is *irreducible*, then any isotopy which removes any of the intersection points of the trellis also reduces the Nielsen entropy. As well as uniformly-hyperbolic diffeomorphisms, we also consider the construction of pseudo-Anosov diffeomorphisms realising the entropy bound.

The main results of the paper are summarised in the following theorem.

Main Theorem *Let $[f;T]$ be a well-formed trellis type. Then for any $\epsilon > 0$ there exists $(\widehat{f},\widehat{T}) \in [f;T]$ such that $h_{top}(\widehat{f}) < h_{niel}[f;T] + \epsilon$. If there exists a diffeomorphism $\widehat{f}$ isotopic to f relative to T such that any $\widehat{f}$-extension of T is minimal, then there exists a uniformly-hyperbolic diffeomorphism $\widetilde{f}$ isotopic to f relative to T such that $h_{top}(\widetilde{f}) = h_{niel}[f;T]$. Further, if $[f;T]$ is irreducible and $(\widetilde{f},T) \in [f;T]$ such that $h_{top}(\widetilde{f}) = h_{niel}[f;T]$, then any $\widetilde{f}$-extension of T is minimal.*

The paper is organised as follows. In section 2 we provide an overview of the trellis theory used in the paper. In section 3, we prove some technical results on the geometry of a trellis and the associated dynamics near the periodic saddle points. The main results of the paper are contained in section 4, and are preceded by a number of illustrative examples.

2 An overview of trellis theory

In this section we introduce the notion of trellis and trellis type, and relate these to the well-known notion of homoclinic/heteroclinic orbit and tangle. We give a number of definitions which allow us to describe some of the basic properties of trellises. We then give some definitions concerning homotopy and isotopy classes of curves. Finally, we review the definitions of (thick) graphs, and show how to associate a graph representative to each irreducible trellis type. Most of these definitions are also given in [5]. Throughout this paper, we restrict to compact, orientable surfaces, possibly with boundary.

2.1 Tangles and Trellises

If P is an invariant set of periodic saddle points of a surface diffeomorphism f, we denote the stable and unstable sets of P by $W^S(f;P)$ and $W^U(f;P)$ respectively. A *tangle* W for a diffeomorphism f is a pair (W^U, W^S) where $W^U = W^U(f;P)$ and $W^S = W^S(f;P)$ for some invariant set of periodic saddle points P of f. If P consists of a single point, then W is a *homoclinic* tangle, otherwise W is a *heteroclinic* tangle. A *branch* of $W^{U/S}(f;P)$ is a component of $W^{U/S} \setminus P$, where we use the notation U/S for statements which hold for both stable and unstable manifolds.

Let f be a diffeomorphism of a surface M with a finite invariant set P of hyperbolic saddle points. A *trellis* for f is a pair $T = (T^U, T^S)$, where T^U and T^S are subsets of $W^U(f;P)$ and $W^S(f;P)$ respectively such that:

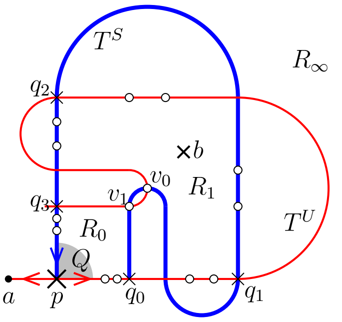

Figure 1: A trellis T_3 for a diffeomorphism f_3

(1). T^U and T^S are both the union of finitely many compact intervals with non-empty interiors,

(2). $f(T^U) \supset T^U$ and $f(T^S) \subset T^S$.

We write $(f;T)$ to denote the pair consisting of a diffeomorphism f and a trellis T for f. An *intersection point* of a tangle T is a point in $T^U \cap T^S$, and is a point of a periodic, homoclinic or heteroclinic orbit. We denote the set of intersections by T^V, and the set of periodic points of $T^U \cup T^S$ by T^P. An intersection point q is *transverse* or *tangential* according to whether T^U and T^S cross transversely or tangentially at q. A trellis is *transverse* if all intersection points are transverse.

We denote the closed interval in $T^{U/S}$ with endpoints a and b by $T^{U/S}[a,b]$, and the open interval by $T^{U/S}(a,b)$. An *arc* of $T^{U/S}$ is a closed subinterval of $T^{U/S}$ with endpoints in T^V. A *segment* of $T^{U/S}$ is an arc of $T^{U/S}$ with no topologically transverse intersection points in its interior. A segment S is *periodic* if $f^{\pm n}(S) \subset S$ for some n, or equivalently, if it contains a point of T^P.

An *endpoint* of $T^{U/S}$ is a point which is at the end of an interval of $T^{U/S}$; we denote the set of endpoints of $T^{U/S}$ by $\partial T^{U/S}$. An *end interval* of $T^{U/S}$ is a subinterval of $T^{U/S}$ with one endpoint in T^V and the other in $\partial T^{U/S}$. An intersection point is an *end intersection point* if is contained in some end interval, or equivalently, if it is not contained in the interior of any arc.

As long as no segment of T^U maps into an end interval of T^U under f^{-1}, and no segment of T^S maps into an end interval of T^S under f, the end intervals are dynamically irrelevant, and can be removed freely. If we know the orientation of the intersection at an endpoint v of $T^{U/S}$, then small end intervals may be freely introduced at v.

A *branch* of $T^{U/S}$ is the closure of the intersection of $T^{U/S}$ with a branch of $W^{U/S}$, and may be connected or disconnected. A connected branch is *trivial* if it contains no points of $T^V \setminus T^P$, so is either a single point or an end interval.

A trellis $T = (T^U, T^S)$ for a diffeomorphism f is *well-formed* if every component of $T^U \cup f(T^S)$ contains a point of T^P.

An *open region* of T is a component of $M \setminus (T^U \cup T^S)$. A *(closed) region* is the closure of an open region, and hence includes the stable and unstable boundary segments. A *(closed) domain* of T is a closed subset of M with boundary in $T^U \cup T^S$; note that we do not require domains to be simply-connected, though this is the case we will usually consider.

There are two special types of region which play an important role, namely *bigons* and *rectangles*. A *bigon* is a region which is a topological disc bounded by one stable and one unstable segment with internal angles less than π. A *rectangle* is a region which is a topological disc bounded by two stable and two unstable segments with internal angles less than π.

The type-3 trellis T_3 is depicted in figure 1. It is formed by subsets of the stable and unstable manifolds of the direct saddle fixed point (i.e. the saddle point with positive eigenvalues) of the

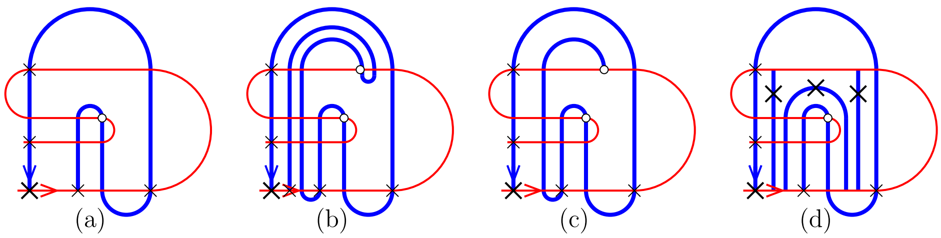

Figure 2: Supertrellises. (a) A trellis T_{3e}, (b) an iterate formed by taking $f^{-1}(T^S_{3e})$, (c) an extension which is not an iterate, and (f) a supertrellis which is not an extension.

Hénon map for certain parameter values. The stable and unstable sets are subsets of the stable and unstable manifolds of the saddle fixed point p; the stable set is drawn in a thicker line width than the unstable set. The branches of this trellis are connected and all the intersection points are transverse. There are ten regions, an unbounded region R_∞, five bigons, three rectangles, and a hexagon R_1. The points p, q_0 and q_3 are the end intersections. The points q_0, q_1, q_2 and q_3 are intersection points on a single homoclinic orbit. The orbits of the intersection points v_0 and v_1 are shown in white dots. One of the unstable branches of the trellis ends in an attracting fixed point, a and one of the stable branches contains only of $\{p\}$, and so is trivial. Both of these are therefore trivial branches.

A surface diffeomorphism f with a periodic saddle orbit has infinitely many trellises, which are partially ordered by inclusion. Taking a smaller trellis gives a *subtrellis*, and a larger trellis a *supertrellis*. Thus T is a subtrellis of $\widehat{T}$ if $T^{U/S} \subset \widehat{T}^{U/S}$, and T is a subtrellis of a tangle W if $T^{U/S} \subset W^{U/S}$. We say $\widehat{T}$ is an *f-supertrellis* of T and $\widehat{W}$ is an *f-supertangle* of T. Of particular importance are those supertrellises which can be obtained by iterating segments or branches. A trellis $\widehat{T}$ is an *f-iterate* of T if there exist positive integers n_u and n_s such that $\widehat{T}^U = f^{n_u}(T^U)$ and $\widehat{T}^S = f^{-n_s}(T^S)$. A trellis $\widehat{T}$ is an *f-extension* of T if there exists n such that $T^U \subset \widehat{T}^U \subset f^n(T^U)$ and $T^S \subset \widehat{T}^S \subset f^{-n}(T^S)$. An iterate/extension is a *stable iterate/extension* if $\widehat{T}^U = T^U$ and an *unstable iterate/extension* if $\widehat{T}^S = T^S$.

Given a diffeomorphism f with a trellis T, we can obtain a canonical map of pairs by *cutting* along unstable curves T^U. The topological pair obtained by cutting along the T^U is denoted $\mathcal{C}T = (\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$. The diffeomorphism f lifts to a map $\mathcal{C}f$ on $\mathcal{C}T$. Notice that the pair $\mathcal{C}T = (\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$ contains the pair $(M \setminus T^U, T^S \setminus T^U)$ as an open subset which is invariant under $\mathcal{C}f$. Indeed, $\mathcal{C}_{T^U}M$ can be regarded as a natural compactification of $M \setminus T^U$, and the homotopy properties of $(\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$ and $(M \setminus T^U, T^S \setminus T^U)$ are essentially the same.

2.2 Isotopies and trellis types

If T is a trellis, then by an *isotopy relative to* T, we mean an isotopy f_t for $t \in [0,1]$ such that T is a trellis for f_t for all t. We write $[f]_T$ for the isotopy class of f relative to T, and say f_0 and f_1 are *isotopic* relative to T. We say $(f_0; T_0)$ and $(f_1; T_1)$ are *conjugate* if there is a homeomorphism h such that $h \circ f_0 = f_1 \circ h$ and $h(T_0^{U/S}) = T_1^{U/S}$. These relations allow us to define the equivalence classes of trellis map which will be our primary object of study.

The *trellis mapping class* $([f]; T)$ is the set of all pairs $(\widetilde{f}; T)$ for which $\widetilde{f} \in [f]_T$. The *trellis type* $[f; T]$ is the set of all pairs $(\widehat{f}, \widehat{T})$ which are conjugate to some $(\widetilde{f}; T)$ with $\widetilde{f} \in [f]_T$. We consider $(f_0; T_0)$ and $(f_1; T_1)$ equivalent if $[f_0; T_0] = [f_1; T_1]$.

The most important dynamical feature of a trellis type is its *entropy*, denoted $h_{top}[f; T]$, and defined to be the infemum of the topological entropies of diffeomorphisms with a trellis in $[f; T]$;

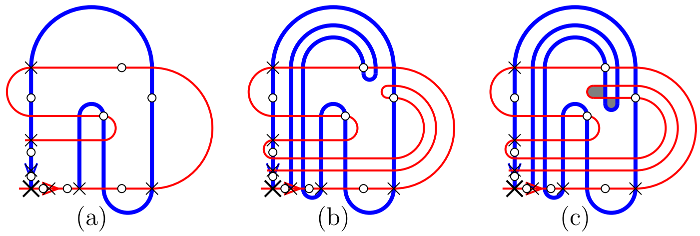

Figure 3: The trellis shown in (b) is a minimal supertrellis (indeed, a minimal iterate) of that in (a), but the trellis in (c) is not, since the shaded bigons do not contain a point in $\bigcup_{n\in\mathbb{Z}} \widetilde{f}^n(T^V)$.

that is

$$h_{top}[f;T] = \inf\{h_{top}(\widehat{f}) : (\widehat{f},\widehat{T}) \in [f;T]\} . \tag{1}$$

The *Nielsen entropy* of a trellis type $[f;T]$, denoted $h_{niel}[f;T]$ is the growth rate of the number of essential Nielsen classes of $\mathcal{C}f$; that is

$$h_{niel}[f;T] = \limsup_{n\to\infty} \frac{\log N^n(\mathcal{C}f)}{n} . \tag{2}$$

See [4] for the definition of the relative Nielsen numbers N^n.

A trellis mapping class $([f];T)$ has an *invariant-curve reduction* if there is a closed one-dimensional manifold C which is disjoint from $T^U \cup T^S$, invariant under some $\widetilde{f}$ with $(\widetilde{f};T) \in ([f];T)$ such that each component of the complement of C either contains a point of $T^U \cup T^S$, or has negative Euler characteristic. A trellis mapping class $([f];T)$ has an *attractor-repeller reduction* if there is a diffeomorphism $\widetilde{f}$ homotopic to f relative to T and a closed one-dimensional manifold C such that C divides M into subsets A and B with $\mathrm{cl}(\widetilde{f}(A)) \subset \mathrm{int}(A)$, and both A and B contain a point of T^P. If $([f]:T)$ has an invariant-curve reduction or an attractor-repeller reduction, then $([f];T)$ is *reducible*, otherwise $([f];T)$ is *irreducible*.

We say a trellis $\widehat{T}$ is an $[f]$-supertrellis of T if $\widehat{T}$ is a $\widehat{f}$-supertrellis of T for some $\widehat{f} \in [f]_T$. A trellis $\widetilde{T}$ is an $[f]$-minimal supertrellis of T if

(1). $\widetilde{T}$ is a trellis for a diffeomorphism $\widetilde{f} \in [f]_T$.

(2). every tangency of $\widetilde{T}$ is a point of $\bigcup_{n\in\mathbb{Z}} \widetilde{f}^n(T^V)$,

(3). every bigon of $\widetilde{T}$ contains a point of $\bigcup_{n\in\mathbb{Z}} \widetilde{f}^n(T^V)$, and either

(4). the end intersections of $\widetilde{T}^U$ lie in $f(T^S)$ and the end intersections of $\widetilde{T}^S$ lie in $f^{-1}(T^U)$, or

(5). $\widetilde{T}$ is a subtrellis of a trellis $\widehat{T}$ which satisfies conditions (1) to (4).

The purpose of condition (5) is to allow the condition (4) on the end intersections of $\widetilde{T}^U$ and $\widetilde{T}^S$ to be relaxed. A minimal supertrellis which is an extension or iterate is, respectively, a *minimal extension* or a *minimal iterate*.

The Nielsen entropy of a minimal supertrellis is the same as that of the original trellis type, as the following result (Theorem 5 of [5]) shows:

Theorem 2.1 (Nielsen entropy of minimal supertrellises) *Let $([f];T)$ be a well-formed trellis mapping class. If $(\widetilde{f};\widetilde{T})$ is a trellis mapping class such that $\widetilde{f} \in [f]_T$ and $\widetilde{T}$ is an $[f]$-minimal supertrellis of f, then $h_{niel}[\widetilde{f};\widetilde{T}] = h_{niel}[f;T]$.*

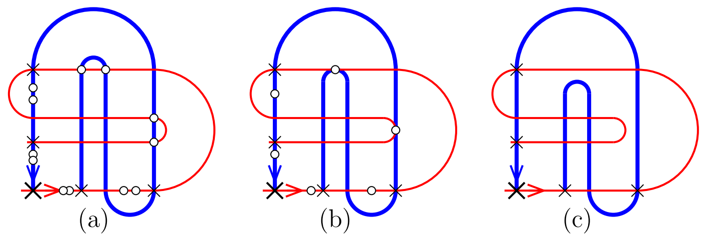

Figure 4: A pruning isotopy. The intersections marked with a circle in (a) are removed by the tangency at (b), but the end intersections are preserved, resulting in the trellis shown in (c).

Given an isotopy of diffeomorphisms f_t with a continuously-varying family of periodic orbits P_t, we can construct a continuously-varying family of trellises T_t with $T_t^{U/S} \subset W^{U/S}(f_t; P_t)$. Of particular importance is the case that the endpoints of each $T_t^{U/S}$ are intersections. In this case, the only changes in the trellis type occur when T_t has tangential intersections. An isotopy $(f_t; T_t)$ in which intersections are removed at tangencies, but not created, is a *pruning isotopy*. An example of a pruning isotopy is shown in figure 4.

The following theorem (Theorem 1 of [5]) shows that entropy decreases during pruning.

Theorem 2.2 (Entropy decreases during pruning) *Suppose $(f_t; T_t)$ is a pruning isotopy. Then $h_{niel}[f_0; T_0] \geqslant h_{niel}[f_1; T_1]$. Further, any essential Nielsen class for $[f_1; T_1]$ can be continued to an essential Nielsen class for $[f_0; T_0]$.*

2.3 Graph maps

Our main tool for computing and describing the dynamics of a trellis mapping class is to relate the trellis type to a graph map. This graph map can be used to compute the Nielsen entropy, as shown in Sections 4 and 5 of [5]. In this paper, we only need to use the graph map in Theorem 4.10 and Lemma 4.5, in which we consider changes in the Nielsen entropy.

We will mostly use the same terminology and notation for graphs as in [1], though many of the same concepts are given in [8]. A *graph* G is a one-dimensional CW-complex with vertices $V(G)$ and edges $E(G)$. Each edge has two orientations; the reverse of an directed edge e is denoted $\bar{e}$ and the initial vertex by $\imath(e)$. An *edge-path* is a list $e_1 \ldots e_n$ of directed edges of G such that $\imath(\bar{e}_i) = \imath(e_{i+1})$ for $1 \leqslant i < n$, and an *edge loop* is a cyclic list of directed edges also satisfying $\imath(\bar{e}_i) = \imath(e_{i+1})$. The *trivial* edge-path contains no edges and is denoted $\bullet$. The reverse of an edge-path $e_1 \ldots e_n$ is the edge-path $\bar{e}_n \ldots \bar{e}_1$ An edge-path $e_1 \ldots e_n$ *back-tracks* if $e_{i+1} = \bar{e}_i$ for some i, otherwise it is *tight*. A *graph map* g is a self-map of G taking each vertex to a vertex, and each directed edge e to an edge-path $e_1 \ldots e_k$ such that $\imath(e_1) = g(\imath(e))$ for all directed edges e, and $g(\bar{e}) = \overline{g(e)}$. The *derivative* map ∂g takes oriented edges to oriented edges or $\bullet$, with $\partial g(e_i) = e_j$ if $g(e_i) = e_j \ldots$ and $\partial g(e_i) = \bullet$ if $g(e_i) = \bullet$.

We will always consider a graph embedded in a surface by an embedding i. This induces a natural relation $\lhd$ on the oriented edges starting at each vertex v given by $e_1 \lhd e_2$ if e_2 is the edge immediately following e_1 in an anticlockwise loop at v. A pair of edges (e_1, e_2) is a *turn* in G at vertex v if $v = \imath(e_1) = \imath(e_2)$ and $e_1 \lhd e_2$, so e_2 immediately follows e_1 in the cyclic order at v.

An edge-path $\pi = p_1, \ldots, p_n$ is an *edge-loop* if $\imath(p_1) = \imath(\bar{p}_n)$. An edge-loop π is *peripheral* in G if $(p_i, \bar{p}_{i-1})$ is a turn in G for all i, where we define $p_0 = p_n$. The *peripheral subgraph* P of g is the maximal invariant subset of G consisting of a union of peripheral loops on which g acts as a homeomorphism. Edges of P are called *peripheral edges*.

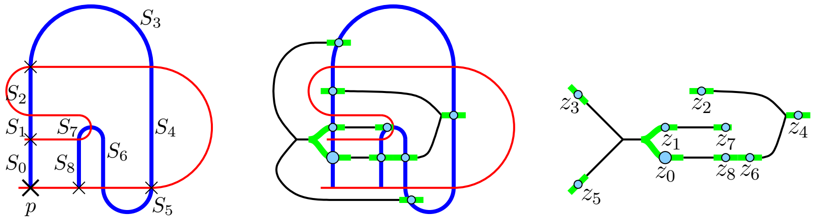

(a) (b) (c)

Figure 5: (a) The type-3 trellis T_3, (b) a controlled graph (G_3, W_3) compatible with the trellis T_3 embedded in $(\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$, and (c) the controlled graph $(G_3.W_3)$, with the control edge crossing the segment S_i of T_3^S labelled by z_i.

The *transition matrix* of a graph map g is the matrix $A = (a_{ij})$ where a_{ij} counts the number of times the undirected edge e_j appears in the image path of edge e_i. The largest eigenvalue of A is the *growth rate* λ of g, and the logarithm of the growth rate gives the topological entropy of g. A *length function* on G is a strictly-positive function $l : E(G) \longrightarrow \mathbb{R}$. The *length* of an edge-path $e_1e_2 \ldots e_n$ is defined to be $l(e_1e_2 \ldots e_n) = \sum_{i=1}^{n} l(e_i)$. If g is a graph map with topological entropy λ, then for any $\epsilon > 0$ there is a length function with $l(g(e)) < (\lambda + \epsilon)l(e)$ for any edge e.

Let (G, W) be a topological pair where G is a graph and W is a finite subset of G. The pair (G, W) is a *controlled graph*, and the edges of G containing points of W are called *control edges*. A graph map $g : (G, W) \longrightarrow (G, W)$ is a *controlled graph map* if $g(z)$ is a control edge whenever z is a control edge. A vertex of G which is the endpoint of a control edge is called a *control vertex*. All other vertices are called *free vertices*, and edges which are not control edges are called *free edges*.

To relate maps of pairs on different spaces, we need the notion of *exact homotopy equivalence*. A map of pairs $f : (X_1, Y_1) \longrightarrow (X_2, Y_2)$ is *exact* if $f^{-1}(Y_2) = Y_1$. Pairs (X_1, Y_1) and (X_2, Y_2) are *exact homotopy equivalent* if there are exact maps $p_1 : (X_1, Y_1) \longrightarrow (X_2, Y_2)$ and $p_2 : (X_2, Y_2) \longrightarrow (X_1, Y_1)$ such that $p_2 \circ p_1 \sim id$ and $p_1 \circ p_2 \sim id$. The map p_1 is an *exact homotopy equivalence*, and p_2 is the *exact homotopy inverse* of p_1. Maps of pairs $f_1 : (X_1, Y_1) \longrightarrow (X_1, Y_1)$ and $f_2 : (X_2, Y_2) \longrightarrow (X_2, Y_2)$ are *exact homotopy equivalent* if there is an exact homotopy equivalences $p_1 : (X_1, Y_1) \longrightarrow (X_2, Y_2)$ such that $p_1 \circ f_1 \sim f_2 \circ p_1$. Note that all homotopies are taken through maps of pairs which need not be exact.

A graph map representing the topology of a trellis via exact homotopy equivalence is called *compatible* with the trellis. To avoid unnecessary complications, we restrict to transverse trellises, and assume that the endpoints of T^U are not intersection points, but that the endpoints of T^S are (in other words, we assume $\partial T^U \cap T^S = \varnothing$ but $\partial T^S \subset T^U$). Let (G, W) be a topological pair where G is a graph and W is a finite subset of G. Then (G, W) is *compatible* with a transverse trellis T if (G, W) and $(\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$ are exact homotopy equivalent by an embedding $i : (G, W) \longrightarrow (\mathcal{C}_{T^U}M, \mathcal{C}_{T^U}T^S)$. A controlled graph map g of (G, W) is *compatible* with the trellis mapping class $([f]; T)$ if the embedding i is an exact homotopy equivalence between g and $\mathcal{C}f$, and for each pair (z, S), where z is a control edge crossing segment S, the relative orientation of $g(z)$ with the segment containing $f(S)$ is the same as that of $f(z)$ with $f(S)$. The controlled graph (G_3, W_3) compatible with the type-3 trellis T_3 is shown in figure 5.

There are many controlled graph maps compatible with a trellis mapping class $([f]; T)$. Graph maps $(g_1; G_1, W_1)$ and $(g_2; G_2, W_2)$ are *equivalent* if there is a homeomorphism $h : (G_1, W_1) \longrightarrow (G_2, W_2)$ such that $h \circ g_1 = g_2 \circ h$. A controlled graph map $(g; G, W)$ is *optimal* if for every turn $e_1 \lhd e_2$, either $\partial g(e_1) \neq \partial g(e_2)$, or at least one of e_1 or e_2 is a control edge. A controlled graph map $(g; G, W)$ is a *graph representative* of a transverse trellis type $[f; T]$ if $(g; G, W)$ is an optimal graph map which is compatible with $([f]; T)$. The following theorem (Theorem 7 of [5]) shows that every proper irreducible trellis type has a unique graph representative.

Theorem 2.3 (Existence and uniqueness of graph representatives) *Let $[f; T]$ be a transverse trellis type with no invariant curve reduction. Then $[f; T]$ has a unique graph representative $(g; G, W)$. Further, if $[f_0; T_0]$ and $[f_1; T_1]$ are trellis types, the graph representatives $(g_0; G_0, W_0)$ and $(g_1; G_1, W_1)$ are equivalent if and only if $[f_0; T_0] = [f_1; T_1]$.*

In particular, this result shows that the graph representative provides a convenient way of specifying a trellis type.

3 Quadrants, regular domains and alpha-chains

In this section we introduce two technical concepts which we require to analyse trellis mapping classes. Regular domains are a local concept concerning neighbourhoods of T^P, whereas alpha-chains capture the global dynamics.

3.1 Curves and homotopies

The main tool for studying the geometry, topology and dynamics associated with trellis maps will be to consider curves in topological pairs (X, Y), which we will usually take to be either the cut surface $\mathcal{C}T = (\mathcal{C}_{T^U} M, \mathcal{C}_{T^U} T^S)$ or a controlled graph (G, W). Since the geometry of $\mathcal{C}T$ can easily be seen from T, we will illustrate curves in $\mathcal{C}_{T^U} M$ by drawing T, and showing the curve embedded in $M \setminus T^U$, as this is easier to visualise.

A *curve* α is a map in the category of topological pairs, $\alpha : (I, J) \longrightarrow (X, Y)$, where I is a compact interval. The *path* of such a curve α is the set $\alpha(I)$. A curve $\alpha_2 : (I_2, J_2) \longrightarrow (X, Y)$ is a *reparameterisation* of $\alpha_1 : (I_1, J_1) \longrightarrow (X, Y)$ if there is an orientation-preserving homeomorphism $(I_1, J_1) \longrightarrow (I_2, J_2)$ with $h(J_1) = J_2$ such that $\alpha_1 = \alpha_2 \circ h$.

For the most part, we are only interested in curves up to homotopy or isotopy, and we always take homotopies and isotopies of curves through maps of pairs. We write $\alpha_0 \sim \alpha_1$ if α_0 is homotopic to α_1. For Nielsen theory we will always keep the endpoints fixed during the homotopy. For most other purposes, we only consider curves for which J contains the endpoints of I, so that the endpoints of the curves α_t lie in Y, but may move. If $\alpha : (I, J) \longrightarrow (X, Y)$ is a curve and J contains the endpoints of I we say *α has endpoints in Y*. By Theorem 3.1 of Epstein [6], if α_0 and α_1 are homotopic simple curves in an orientable surface, then they are isotopic.

Reparameterising a curve but may change the set J which maps into Y. This means that different parameterisations of the same path may not even be comparable under homotopy. However, we consider different parameterisations of the same curve as equivalent.

Definition 3.1 (Homotopy equivalence of curves) Curves $\alpha_1 : (I_1, J_1) \longrightarrow (X, Y)$ and $\alpha_2 : (I_2, J_2) \longrightarrow (X, Y)$ are *homotopy equivalent* if there is an orientation-preserving homeomorphism $h : (I_1, J_1) \longrightarrow (I_2, J_2)$ with $h(J_1) = J_2$ such that $\alpha_1 \sim \alpha_2 \circ h$ as curves $(I_1, J_1) \longrightarrow (X, Y)$. The homotopy may be taken relative to endpoints, as appropriate.

We can also define a weaker notion.

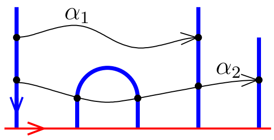

Figure 6: The exact curve α_2 tightens onto α_1.

Definition 3.2 (Tightening curves) Let $\alpha_1 : (I_1, J_1) \longrightarrow (X, Y)$ and $\alpha_2 : (I_2, J_2) \longrightarrow (X, Y)$ be exact curves. We say α_2 *tightens* to α_1 if there is an injective map $h : (I_1, J_1) \longrightarrow (I_2, J_2)$ such that $\alpha_1 \sim \alpha_2 \circ h$ as curves $(I_1, J_1) \longrightarrow (X, Y)$.

It is immediate that the tightening relation (α_2 tightens onto α_1) is reflexive and transitive. If we restrict to curves $\alpha : (I, J) \longrightarrow (X, Y)$ where J has finitely many components, it is easy to show that the tightening relation is also antisymmetric, so is a partial order on exact homotopy classes.

We are especially interested in iterates with minimal number of components intersecting Y.

Definition 3.3 (Minimal iterates of curves) Let $\alpha : (I, J) \longrightarrow (X, Y)$ be a simple curve with endpoints in Y, and $f : (X, Y) \longrightarrow (X, Y)$ be a map of pairs. Then a *minimal iterate* of α under f is a curve β which is homotopic to $f \circ \alpha$ relative to J and which minimises the number of components of I mapping into Y. We let $J' = \beta^{-1}(Y)$, and consider β as an exact curve $(I, J') \longrightarrow (X, Y)$. If we further require that an intersection is isolated whenever possible, then the curve β is well-defined up to equivalence, so we obtain a well-defined map $f_{\min}$ on equivalence classes of curves given by $f_{\min}[\alpha] = [\beta]$.

By $f^n_{\min}[\alpha]$ we mean $(f_{\min})^n[\alpha]$, and not $(f^n)_{\min}[\alpha]$, which typically has fewer intersections. An example of minimal and non-minimal iterates is given in figure 7. Taking repeated minimal iterates of curves gives an *alpha-chain*.

Definition 3.4 (Alpha-chain) A list $[\alpha_i]$, $0 \leqslant i \leqslant k$ of homotopy classes of exact curves with endpoints in Y is an *alpha-chain* if $f_{\min}[\alpha_i]$ tightens onto $[\alpha_{i+1}]$ for $0 \leqslant i < k$.

If $[\alpha_0]$ is an exact homotopy class of curve such that the initial endpoint of α_0 lies in a component Y_0 of Y for which $f^n(Y_0) \subset Y_0$, then taking repeated minimal iterates $[\alpha_i] = f^i_{\min}[\alpha_0]$ gives exact homotopy classes $[\alpha_{kn}]$ with initial endpoint in Y_0. We can consider how the component Y_i containing the first intersection of $[\alpha_i]$ with Y evolves with i. Of particular interest is when the component containing the first intersection with $[\alpha_{kn}]$ is eventually constant.

Definition 3.5 (Exit set) Let Y_0, Y_1 be components of Y such that $f^n(Y_0) \subset Y_0$, and $[\alpha] : ([0, 1], \{0, 1\}) \longrightarrow (X, Y)$ be a homotopy class of exact curves such that $\alpha(0) \in Y_0$ and $\alpha(1) \in Y_1$. Then Y_1 is an *exit set* for α if there is some minimal iterate $\alpha_n = f^n_{\min}[\alpha]$ and some $t \in (0, 1)$ such that $\alpha_n|_{[0,t]}$ is homotopy-equivalent to α_0.

In other words, Y_1 is an exit set if the initial part of some minimal iterate of α tightens onto α. A similar definition holds for curves with initial endpoint at an essential period-n Nielsen class of f, except that we take an exact homotopy class $[\alpha_0] : ([0, 1], \{1\}) \longrightarrow (X, Y)$, and perform isotopies fixing $\alpha(0)$. An example of an exit set is shown in figure 7(b).

We will also need the following result of plane topology, which can be proved using results of [6] and by passing to a universal cover.

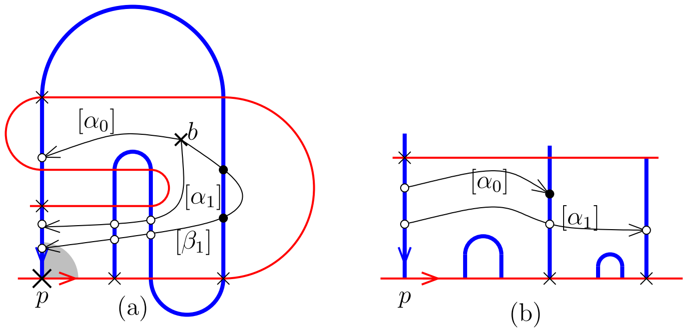

Figure 7: (a) The exact homotopy class $[\alpha_1]$ is a minimal iterate of $[\alpha_0]$, but $[\beta_1]$ is not, since it contains extra intersections marked with black dots. (b) The segment S_1 is an exit set for $[\alpha_0]$ starting at segment S_0, since the initial part of $[\alpha_1] = f_{\min}[\alpha_0|$ tightens onto $[\alpha_0]$.

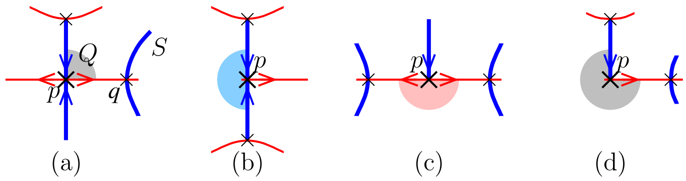

Figure 8: Sectors at p. (a) depicts a quadrant Q, (b) an attracting secant, (c) a repelling secant and (d) a coquadrant.

Theorem 3.6 *Let $\{\alpha_i : i = 0, \ldots, n-1\}$ and $\{\beta_i : i = 0, \ldots, n-1\}$ be sets of mutually disjoint simple curves with endpoints in T^S which are disjoint from T^U and have minimal intersections with T^S. Suppose that the homotopy classes of the α_i are mutually disjoint, and that $\alpha_i \sim \beta_i$ for $i = 0, \ldots, n-1$. Then there is a isotopy h_t such that $h(T^U) = T^U$, $h(T^S) = T^S$, $h_0 = id$ and $h_1 \circ \alpha_i = \beta_i$ for $i = 0, \ldots, n-1$.*

3.2 Quadrants, attractors and repeller

In general, we do not have much control over the geometry of a trellis. The exception is near a point of T^P, where the dynamics are conjugate to a linear map. It is important to consider how the nontrivial branches T divide the surface in a neighbourhood of T^P.

Definition 3.7 (Quadrant, secant and coquadrant) Let $T = (T^U, T^S)$ be a trellis such that the endpoints of $T^{U/S}$ are points of T^V. Then a *sector* Q of T at p is a local component of a region of T in a neighbourhood of p. A sector is a *quadrant* if the boundary of Q at p forms an angle less than π, in which case the boundary includes a piece of an unstable branch $T^U(Q)$ and a piece of a stable branch $T^S(Q)$. A sector is an *secant* if the boundary at p forms an angle exactly π, in which case the boundary includes an interval either T^U or T^S. We say the secant is *attracting* if this is an interval in T^S, and *repelling* if this is an interval of T^U. A sector is a *coquadrant* if the boundary of Q at p forms an angle greater than π.

Quadrants, secants and coquadrants are shown in figure 8. The type three trellis shown in figure 1 has a quadrant Q contained in the region R_0.

The *image* of a quadrant Q is the quadrant containing $f(Q)$; note that $T^{U/S}(f(Q))$ is the same branch as $f(T^{U/S}(Q))$. A stable segment S with endpoint $q \in T^U(Q)$ lies on the *Q-side* of T^U if locally S lies on the same side of $T^U(Q)$ at q as $T^S(Q)$ does at p. Similarly, an unstable segment U with endpoint $q \in T^S(Q)$ lies on the *Q-side* of T^S if locally U lies on the same side

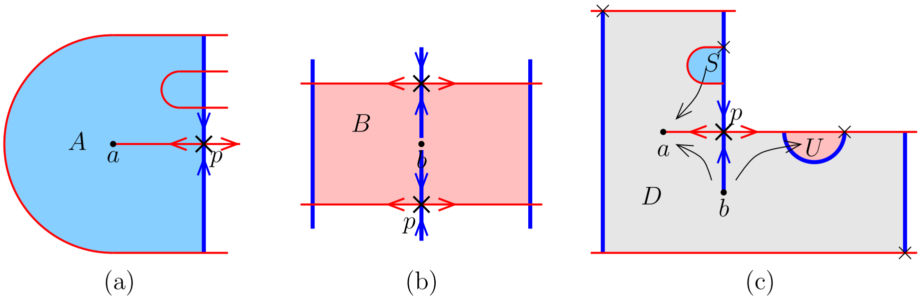

Figure 9: Attracting and repelling domains. (a) The domain A is an attracting domain with stable fixed point a, and (b) the region B is a repelling domain with unstable fixed point b. (c) A domain D with an attracting fixed point a and repelling fixed point b. Points enter D through the region S, which is a stable region, and leave D through U, which is an unstable region.

of $T^S(Q)$ at q as $T^U(Q)$ does at p. The segment S in figure 8(a) lies on the Q-side of T^U. We write $S(Q)$ for the first segment of T^S crossing the branch $T^U(Q)$ which lies on the Q-side of T^U.

As well as saddle periodic orbits, diffeomorphisms also have attracting and repelling periodic orbits. Clearly, a stable segment cannot be in the basin of attraction of an attracting orbit, and neither can an unstable segment be in the basin of a repelling orbit. However, it is possible for the interior of an entire domain to lie in the basin of an attracting or repelling periodic orbit. Such domains are called *stable* or *unstable* domains (or regions).

Let T be a trellis for a diffeomorphism f.

Definition 3.8 (Attracting and repelling domains) A domain D of T is *attracting* if there is a diffeomorphism $\widetilde{f} \in [f]_T$ and a period-n point $a \in \mathrm{int}(D)$ of $\widetilde{f}$ such that $\widetilde{f}^{in}(x) \to a$ as $i \to \infty$ for all $x \in \mathrm{int}(D)$. Similarly, a domain D of T is *repelling* if there is a diffeomorphism $\widetilde{f} \in [f]_T$ and a period-n point $b \in \mathrm{int}(D)$ of $\widetilde{f}$ such that $\widetilde{f}^{-in}(x) \to b$ as $i \to \infty$ for all $x \in \mathrm{int}(D)$. A domain D of T is *attracting/repelling* if there is a diffeomorphism $\widetilde{f} \in [f]_T$ such that D contains an attracting period-n point a and a repelling period-n point b such that all $x \in \mathrm{int}(R)$, either $\widetilde{f}^{in}(x) \to a$ or $\widetilde{f}^{-in}(x) \to b$ as $i \to \infty$.

Examples of attracting and repelling domains are shown in figure 9.

Definition 3.9 (Stable and unstable domains) A domain D of T is *stable* if there is a diffeomorphism $\widetilde{f}$ isotopic to f relative to T and a period-n orbit a of $\widetilde{f}$ such that $\widetilde{f}^{in}(x) \to a$ as $i \to \infty$ for all $x \in \mathrm{int}(D)$. Similarly, a domain D is *unstable* if there is a diffeomorphism $\widetilde{f}$ isotopic to f relative to T and a period-n orbit b of $\widetilde{f}$ such that $\widetilde{f}^{-in}(x) \to b$ as $i \to \infty$ for all $x \in \mathrm{int}(D)$.

Note that T^S is disjoint from the interior of A, T^U is disjoint from the interior of B, and there are no points of T^V in D; this is immediate from the definitions. It is immediate that the interior of an attracting or stable domain is disjoint from T^S, the interior of a repelling or unstable domain is disjoint from T^U, and the interior of an attracting/repelling domain D contains no points of T^V. Further, there exists a diffeomorphism $\widetilde{f}$ isotopic to f relative to T such that $\widetilde{T}^V$ contains no intersection points in the interior of such a domain for any $\widetilde{f}$-extension of T.

In contrast, some regions of T contain an intersection point for some $\widetilde{f}$-iterate of T for any $\widetilde{f} \in [f]_T$.

Definition 3.10 (Chaotic region) Let $([f];T)$ be a trellis mapping class. A region R is *chaotic* if for every diffeomorphism $\widehat{f} \in [f]_T$, there exists an integer n such that $\widehat{f}^n(T^U) \cap \widehat{f}^{-n}(T^S)$ contains a point in the interior of R.

An open segment which is in the basin of an attracting or repelling periodic orbit is a wandering set, and contains no intersection points. There are some nonwandering segments such that any finite iterate need not contain an intersection point. We call such segments *almost wandering.*

Definition 3.11 (Almost wandering segment) Let $([f];T)$ be a trellis mapping class. An open (unstable or stable) segment S is *wandering* if there is a diffeomorphism $\widetilde{f} \in [f]_T$ such that for all $n \in \mathbb{Z}$, $\widetilde{f}^n(S) \cap T^V = \varnothing$. An open segment is *almost wandering* if for all $n \in \mathbb{Z}$, there exists $\widetilde{f} \in [f]_T$ such that $\widetilde{f}^n(S) \cap T^V = \varnothing$, but for every diffeomorphism $\widetilde{f} \in [f]_T$, there exists $n \in \mathbb{Z}$ such that $\widetilde{f}^n(S) \cap T^V \neq \varnothing$.

In other words, an open segment is wandering if there exists a diffeomorphism $\widetilde{f}$ isotopic to f relative to T such that every $\widetilde{f}$-iterate of the segment contains no intersection points. An open segment is almost wandering if it is not wandering, but for every n, there exists $\widetilde{f}$ isotopic to f relative to T such that the nth $\widetilde{f}$-iterate of the segment contains no intersection points.

3.3 Regular quadrants

Any diffeomorphism is topologically conjugate to a linear hyperbolic diffeomorphism in a neighbourhood of a hyperbolic periodic saddle orbit. We therefore expect a diffeomorphism f with trellis T to behave in a fairly predictable way in a neighbourhood of a point of T^P. In particular, we know how a rectangle with sides parallel to the local stable and unstable foliations behaves. Unfortunately, for a given trellis mapping class, the neighbourhood of T^P on which we have hyperbolic behaviour may be arbitrarily small. To deal with this problem, we introduce the concept of a *regular domain*, which is a rectangular domain which behaves similarly to a sufficiently small rectangular neighbourhood of a periodic saddle point.

Definition 3.12 (Regular domain) Let f be a diffeomorphism with trellis T. A rectangular domain D of T is a *regular domain* for $(f;T)$ at a period-n quadrant Q of T if D has sides $T^U[p,q^u]$, $T^S[p,q^s]$, $T^U[q^s,r]$ and $T^S[q^u,r]$, such that

(1u) $f^n(T^U[p,q^u]) \subset T^U$ with $f^n(T^U[p,q^u]) \cap D = T^U[p,q^u]$,

(1s) $f^{-n}(T^S[p,q^s]) \subset T^S$ with $f^{-n}(T^S[p,q^s]) \cap D = T^S[p,q^s]$,

(2u) $f^{-n}(T^U[q^s,r]) \cap D = \varnothing$, and

(2s) $f^n(T^S[q^u,r]) \cap D = \varnothing$.

The sides $T^U[p,q^u]$ and $T^S[p,q^s]$ are called the *adjacent sides*, and the sides $T^U[q^s,r]$ and $T^S[q^u,r]$ are the *opposite sides* of D. Notice that the definition is invariant under time reversal if we interchange T^U and T^S. In general, we therefore only need to prove statements on the topology of a regular domain for either the unstable (u) or stable (s) case.

Definition 3.13 (Regular quadrant) A quadrant Q is *regular* if Q is contained in a regular domain D such that D contains no points of T^P.

A regular domain is shown in figure 10. We shall always denote the vertices of a regular domain by p, q^u, q^s and r as used in the definition. We now give a number of elementary properties of a regular domain.

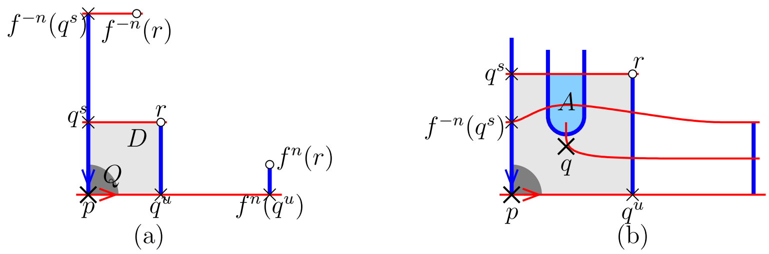

Figure 10: The domain D shown in (a) is a regular domain at a regular quadrant Q. The domain D shown in (b) is also a regular domain, but Q is not a regular quadrant, since the point $\widehat{p}$ is a point of T^P in D.

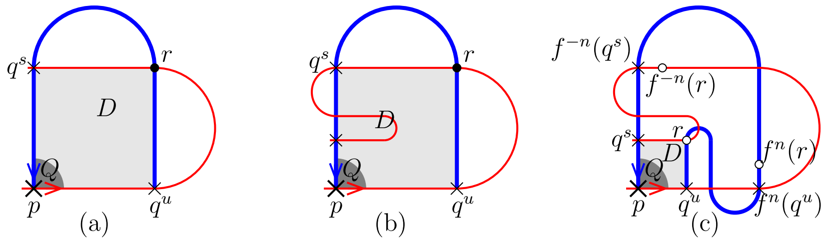

Figure 11: The domain D is a regular domain at quadrant Q in (c), but not in (a) or (b).

Lemma 3.14 *Let D be a regular domain for $(f;T)$ at Q. Then*

(a). If $f^{-(n+1)}(T^S[p,q^s])$ and $f^{-1}(T^S[q^u,r])$ are subsets of T^S, then $f^{-1}(D)$ is a regular domain for $(f;T)$ at Q.

(b). If $\widetilde{f}$ is isotopic to f relative to T, then D is a regular domain for $(\widetilde{f};T)$ at Q.

(c). If $\widehat{T}$ is an f-extension of T, then D is a regular domain for $(f;\widehat{T})$ at Q.

PROOF:

(a). Since $f^{-1}(T^U) \subset T^U$, $f^{-1}(T^S[p,q^s]) \subset f^{-(n+1)}(T^S[p,q^s]) \subset T^S$ and $f^{-1}(T^S[q^u,r]) \subset T^S$, the boundary of $f^{-1}(D)$ is composed of segments of T. On the adjacent sides we have $f^n(f^{-1}(T^U[p,q^u])) \cap f^{-1}(D) = f^{-1}(f^n(T^U[p,q^u]) \cap D) = f^{-1}(T^U[p,q^u])$ and $f^{-n}(f^{-1}(T^S[p,q^s])) \cap f^{-1}(D) = f^{-1}(f^{-n}(T^U[p,q^s]) \cap D) = f^{-1}(T^S[p,q^s])$. On the opposite sides, we have $f^n(f^{-1}(T^S[q^u,r])) \cap f^{-1}(D) = f^{-1}(f^n(T^S[q^u,r]) \cap D) = \varnothing$ and similarly $f^{-n}(f^{-1}(T^U[q^s,r])) \cap f^{-1}(D) = \varnothing$.

(b). Condition (1u) shows that $f^n(q^u) \in T^V$, and thus $\widetilde{f}^n(T^U[p,q^u]) = T^U[\widetilde{f}^n(p),\widetilde{f}^n(q^u)] = T^U[f^n(p),f^n(q^u)] = f^n(T^U[p,q^u])$, so condition (1u) is satisfied for $\widetilde{f}$. If $f^n(r) \in T^V$, then $\widetilde{f}^n(T^S[q^u,r]) = f^n(T^S[q^u,r])$, so condition (2s) is satisfied. Otherwise, $f^n(T^S[q^u,r]) \subset T^S(f^n(q^u),x)$ for some $x \in T^V$ such that $T^S(f^n(r),x)$ contains no points of T^V. Then since $T^S[f^n(q^u),f^n(r)] \cap D = f^n(T^S[q^u,r]) \cap D = \varnothing$, we have $T^S[f^n(q^u),x) \cap D = \varnothing$, and since $\widetilde{f}^n(T^S[q^u,r]) \subset T^S[f^n(q^u),x)$, we have $\widetilde{f}^n(T^S[q^u,r]) \cap D = \varnothing$, so condition (2s) is satisfied. A similar argument proves conditions (1s) and (2u).

(c). Since $\widehat{T}^U[p,q^u] = T^U[p,q^u]$, we have $f^n(\widehat{T}^U[p,q^u]) = f^n(T^U[p,q^u]) \subset T^U[p,q^u] \subset \widehat{T}^U[p,q^u]$, so condition (1u) is satisfied. Since $\widehat{T}^U[q^s,r] = T^U[q^s,r]$, we have $f^{-n}(\widehat{T}^U[q^s,r]) \cap D = f^{-n}(T^U[q^s,r]) \cap D = \varnothing$, so condition (2u) is satisfied. A similar argument proves conditions (1s) and (2s).

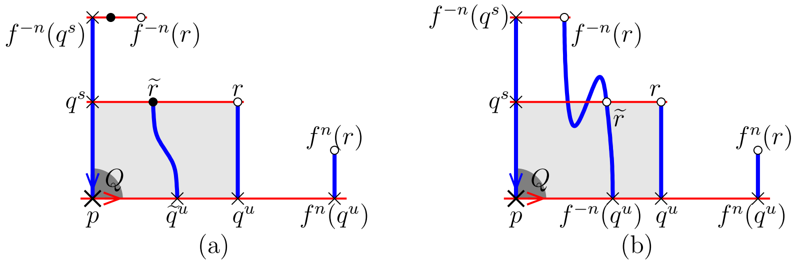

Figure 12: (a) The sub-domain $\widetilde{D}$ is a regular domain if D is a regular domain. (b) A regular domain D together with $f^{-n}(T^S[q^u,r]) = T^S[f^{-n}(q^u), f^{-n}(r)]$, which crosses $T^U[q^s,r]$ since $f^{-n}(r) \notin D$.

$\square$

Lemma 3.14(b) shows that being a regular domain or quadrant is a property of being a trellis mapping class.

We now show how to construct regular subdomains of a regular domain.

Lemma 3.15 *Let D be a regular domain for $(f;T)$ at a quadrant Q.*

(a). If $T^S[\tilde{q}^u,\tilde{r}] \subset D$ is an arc with endpoints $\tilde{q}^u \in T^U(p,q^u)$ and $\tilde{r} \in T^U(q^s,r)$, then the rectangular domain $\widetilde{D}$ with vertices at $\{p,\tilde{q}^u,q^s,\tilde{r}\}$ is a regular domain at Q.

Similarly, if $T^U[\tilde{q}^s,\tilde{r}] \subset D$ is an arc with endpoints $\tilde{q}^s \in T^S(p,q^s)$ and $\tilde{r} \in T^S(q^u,r)$, then the rectangular domain vertices at $\{p,q^u,\tilde{q}^s,\tilde{r}\}$ is a regular domain at Q.

(b). If $f^{-n}(T^S[q^u,r]) \subset T^S$, then $f^{-n}(T^S[q^u,r])$ contains a subinterval of the form $T^S[\tilde{q}^u,\tilde{r}] \subset D$ such that $\tilde{q}^u = f^{-n}(q^u) \in T^U(p,q^u)$ and $\tilde{r} \in T^U(q^s,r)$.

Similarly, if $f^n(T^U[q^s,r]) \subset T^U$, then $f^n(T^U[q^s,r])$ contains a subinterval of the form $T^U[\tilde{q}^s,\tilde{r}] \subset D$ such that $\tilde{q}^s = f^n(q^s) \in T^S(p,q^s)$ and $\tilde{r} \in T^S(q^u,r)$.

The constructions are illustrated in figure 12.

Proof:

(a). The only nontrivial step is to show $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap \widetilde{D} = \varnothing$. First, note that $f^n(\tilde{q}^u) \in T^U(\tilde{q}^u, f^n(q^u))$, so $f^n(\tilde{q}^u) \notin \widetilde{D}$. Further, $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap T^U[p,\tilde{q}^u] = f^n(T^S[\tilde{q}^u,\tilde{r}] \cap f^{-n}(T^U[p,\tilde{q}^u])) = f^n(\varnothing) = \varnothing$ and $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap T^U[q^s,\tilde{r}] = f^n(T^S[\tilde{q}^u,\tilde{r}] \cap f^{-n}(T^U[q^s,\tilde{r}])) \subset f^n(D \cap f^{-n}(T^U[q^s,\tilde{r}])) = \varnothing$, so $f^n(T^S[\tilde{q}^u,\tilde{r}^u])$ does not intersect the unstable boundary of $\widetilde{D}$. Additionally, $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap T^S[p,q^s] = f^n(T^s[\tilde{q}^u,\tilde{r}] \cap f^{-n}(T^S[p,q^s])) \subset f^n(T^S[\tilde{q}^u,\tilde{r}] \cap T^S[p,q^s]) = \varnothing$. It remains to show that $f^n(T^S[\tilde{q}^u,\tilde{r}])$ is disjoint from $T^S[\tilde{q}^u,\tilde{r}]$. We have already seen that $f^n(T^S[\tilde{q}^u,\tilde{r}])$ does not contain $\tilde{q}^u$ or $\tilde{r}$. Further, $f^n(\{\tilde{q}^u\}) \cap T^S[\tilde{q}^u,\tilde{r}] \subset f^n(T^U(p,q^u)) \cap T^S[\tilde{q}^u,\tilde{r}] = \{\tilde{q}\}$, but clearly $f^n(\tilde{q}^u) \neq \tilde{q}^u$, so $f^n(\tilde{q}^u) \notin T^S[\tilde{q}^u,\tilde{r}]$. Since $T^S[\tilde{q}^u,\tilde{r}]$ is an interval, this is enough to show that $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap T^S[\tilde{q}^u,\tilde{r}] = \varnothing$. Hence $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap \partial\widetilde{D} = \varnothing$, so $f^n(T^S[\tilde{q}^u,\tilde{r}]) \cap \widetilde{D} = \varnothing$. A similar analysis proves the statement for a curve $T^U[\tilde{q}^s,\tilde{r}]$.

(b). Since $f^n(Q) = Q$, f^n is orientation-preserving, so the orientation of the intersection at $f^{-n}(q^u)$ is the same as that at q^u. Further, $f^{-n}(T^S[q^u,r]) \cap T^U[p,q^u] = f^{-n}(T^S[q^u,r] \cap f^n(T^U[p,q^u])) = f^{-n}(T^S[q^u,r] \cap T^U[p,q^u]) = \{f^{-n}(q^u)\}$, so $f^{-n}(T^S(q^u,r])$ does not cross $T^U[p,q^u]$. Finally, $\{f^{-n}(r)\} \cap D \subset f^{-n}(T^U[q^s,r]) \cap D = \varnothing$, so $f^{-n}(r) \notin D$. Therefore, $f^{-n}(T^S[q^u,r])$ intersects $T^U[q^s,r]$, and we let $\tilde{r}$ be the first such intersection.

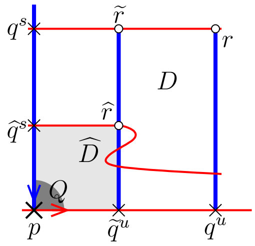

Figure 13: The domain $\widehat{D}$ is a regular domain for Q if D is a regular domain.

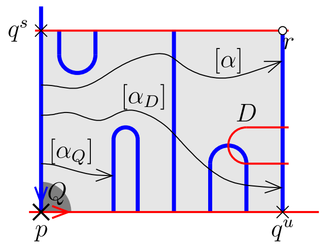

Figure 14: The exact homotopy classes $[\alpha_Q]$ and $[\alpha_D]$, and a homotopy class $[\alpha]$ crossing a regular domain D. The exact homotopy class $[\alpha_D]$ has minimal intersections with T^S and contains curves homotopic to $T^U[p, q^u]$.

□

In Figure 13 we show a regular domain D of $(f; T)$ at a quadrant Q. An application of Lemma 3.15 shows that the rectangular domain with vertices at p, $\widetilde{q}^u$, q^s and $\widetilde{r}$ is a regular domain for Q, and a further application of Lemma 3.15(a) shows that the rectangular domain $\widehat{D}$ with vertices $\{p, \widetilde{q}^u, \widehat{q}^s, \widehat{r}\}$ is also a regular domain. By taking iterates of T, we can ensure that Q has a regular domain contained in an arbitrarily small neighbourhood of p.

We are particularly interested in curves crossing a regular domain.

Definition 3.16 (Domain crossing curves) If D is a regular domain for a quadrant Q, we say that a homotopy class of curves $[\alpha]$ *crosses* D if $\alpha(I) \subset D$ with $\alpha(0) \in T^S[p, q^s]$ and $\alpha(1) \in T^S[q^u, r]$.

Let $[\alpha_D]$ be the exact homotopy class of curve in D with initial endpoint in the segment of $T^S[p, q^s]$ containing p, final endpoint in the segment of $T^S[q^u, r]$ containing q^u, which has minimal intersections with T^S and is homotopic to $T^U[p, q^u]$.

Let $[\alpha_Q]$ be the homotopy class of an exact curve $\alpha_Q : (I, \partial I) \longrightarrow \mathcal{C}T$ with initial endpoint in the segment of $T^S(Q)$ containing p, and the other in the segment S on the Q-side of $T^U(Q)$ containing the first intersection of T^S with $T^U(Q)$.

An exact homotopy class $[\alpha]$ crossing a regular domain D, together with the exact homotopy classes $[\alpha_D]$ and $[\alpha_Q]$ are shown in figure 14

Lemma 3.17 *Let Q be a regular period-n quadrant of $(f; T)$, and D be a regular domain for Q. Then*

(a). there exists k such that $f_{\min}^{kn}[\alpha_Q]$ tightens onto $[\alpha_D]$, and

(b). if $\widetilde{D}$ is any other regular domain for Q, then there exists k such that $f_{\min}^{kn}[\alpha_{\widetilde{D}}]$ tightens onto $[\alpha_D]$.

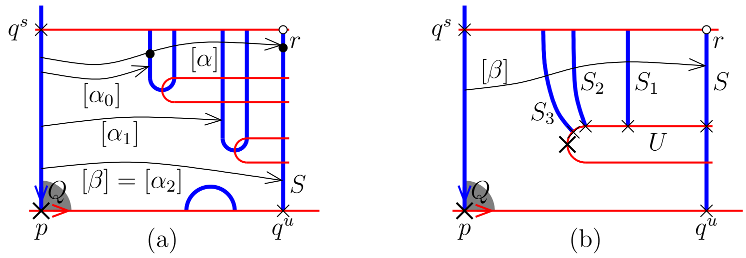

Figure 15: A regular domain at a quadrant Q. (a) $f^{kn}_{\min}[\alpha]$ tightens onto $[\alpha_i]$ for $i = 0, 1, 2$, and $f^n_{\min}[\alpha]$ tightens onto itself. (b) Backward minimal iterates S_i of segment S, iterates of one endpoint converging to a point of T^P.

PROOF:

(a). Let S be the segment containing the final endpoint of $[\alpha_Q]$, and q be the intersection of S with $T^U(Q)$. Then $q \in T^U[p, q^u]$, so there exists k such that $f^{kn}(q) \in T^U[q^u, f^n(q^u)]$. For this k, $f^{kn} \circ \alpha_Q$ is a curve homotopic to T^U relative endpoints, so leaves D through the segment of $T^S[q^u, r]$ containing q^u. Hence $f^{kn}_{\min}[\alpha_Q]$ tightens onto $[\alpha_D]$.

(b). If $\widetilde{q}^u \subset T^U[p, q^u]$, then $[\alpha_{\widetilde{D}}]$ tightens onto $[\alpha_D]$, since both homotopy classes contain curves lying along $T^U[p, q^u]$. Otherwise, there exists k such that $f^{kn}(\widetilde{q}^u) \in T^U[q^u, f^n(q^u)]$, and then $f^{kn}_{\min}[\alpha_{\widetilde{D}}]$ tightens onto $[\alpha_D]$.

□

Lemma 3.18 *Let D be a regular domain for a period-n quadrant Q such that D contains no point of T^P in its interior. If $[\alpha]$ is any homotopy class of curves crossing D, then there exists k such that $f^{kn}_{\min}[\alpha]$ tightens onto $[\alpha_D]$.*

PROOF: Let $[\alpha_i]$ be the exact homotopy class of the sub-curve of $f^{in}_{\min}[\alpha]$ from $T^S[p, q^s]$ to the first crossing with T^S. Then for sufficiently large i, the initial endpoint of $[\alpha_i]$ lies in the segment of $T^S[p, q^s]$ containing p. Since f is a homeomorphism and D is simply-connected, the final endpoint of $[\alpha_i]$ lies in a segment S for all i sufficiently large. We let $[\beta]$ be the exact homotopy class of curve from the segment of $T^S[p, q^s]$ containing p to the segment S, as shown in figure 15(a).

If $f^n(S) \subset S$, then S contains a point of T^P, a contradiction. Since the first crossing of $f^n_{\min}[\beta]$ with T^S lies in segment S, it must be the case that $f^{-n}_{\min}(S)$ has an essential crossing with β. We can therefore define segments S_i recursively by taking $S_0 = S$ and S_{i+1} to be the segment of $f^{-n}_{\min}(S_i)$ with an essential crossing with β, as shown in figure 15(b). Since f is a diffeomorphism, there are segments U_0 and U_1 of T^U which contain the endpoints of S_i for all sufficiently large i. The segments containing $f^{-n}(U_0)$ and $f^{-n}(U_1)$ cannot both lie outside of D, since then $f^{-n}_{\min}(S_i)$ would lie outside D. Hence there is an unstable segment U in the interior of D such that $f^{-n}(U) \subset U$, and so there is a point of T^P in U. Then U must be the segment of $T^U[p, q^u]$ containing p, and hence one endpoint of S lies in $T^U[p, q^u]$.

Therefore, the curve β is either homotopic to $T^U[p, q]$, or bounds a set B which contains a subset of T^U in its interior. However, we can isotope f so that B is repelling, and then the interior of B must contain a periodic point in T^U, a contradiction. If β is homotopic to $T^U[p, q]$, we can use a similar argument to that used in Lemma 3.17 to show that some minimal iterate of $[\beta]$, and hence some $f^{kn}_{\min}[\alpha]$, tightens onto $[\alpha_D]$. □

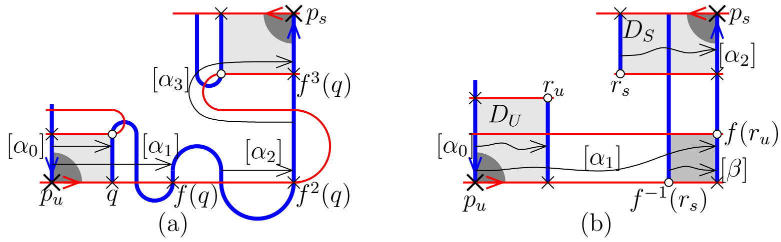

Figure 16: Trellises with regular domains D_U at p_u and D_S at p_s. (a) There is an alpha-chain from $[\alpha_{D_U}] = [\alpha_0]$ to $[\alpha_3]$ which tightens onto $[\alpha_{D_S}]$. (b) Since $f(D_U)$ and $f^{-1}(D_S)$ intersect to form a rectangular domain, there is an alpha-chain from $[\alpha_{D_U}] = [\alpha_0]$ to $[\alpha_2]$, which is the reverse of $[\alpha_{D_S}]$.

Definition 3.19 (Transitive trellis mapping class) Let Q_U and Q_S be regular quadrants contained in regular domains D_U and D_S, respectively. There is an alpha-chain *linking* Q_U to Q_S if there is an alpha-chain from $[\alpha_{D_U}]$ to $[\alpha_{D_S}]$. A trellis mapping class $([f];T)$ is *transitive* if every quadrant Q is regular, and for pair of quadrants Q_U and Q_S, there is an alpha-chain linking Q_U to Q_S.

An example of quadrants Q_U and Q_S for which there is a alpha-chain linking Q_U to Q_S is shown in figure 16(a). Note that if there is an alpha-chain from $[\alpha_{D_U}]$ to $[\alpha_{D_S}]$, then there is an alpha-chain from $[\alpha_{\widetilde{D}_U}]$ to $[\alpha_{\widetilde{D}_S}]$ for any regular domains $\widetilde{D}_U$ and $\widetilde{D}_S$ containing Q_U and Q_S. Transitivity of a trellis mapping class is an especially useful property since it implies that the trellis type has a transitive graph representative (see Lemma 4.9).

The following lemma gives a sufficient condition for the existence of an alpha-chain linking Q_U to Q_S.

Lemma 3.20 *Let D_U and D_S be regular domains for Q_U and Q_S with periods n_u and n_s, respectively.*

(a). If there exist k_u and k_s such that $f^{k_u n_u}(D_U) \cap f^{-k_s n_s n}(D_S)$ is a rectangular domain with unstable edges contained in $f^{k_u n_u}(\partial D_U) \cap T^U$ and stable edges in $f^{-k_s n_s}(\partial D_S) \cap T^S$, then there is an alpha-chain linking Q_U to Q_S.

(b). If there is an alpha-chain linking Q_U to Q_S for $(f;T)$ and $\widehat{T}$ is an f-extension of T, then there is an alpha-chain linking Q_U to Q_S for $(f;\widehat{T})$.

PROOF:

(a). Let β be a curve in $f^{k_u n_u}(D_U) \cap f^{-k_s n_s}(D_S)$ such that $[\beta]$ is the restriction of $[\alpha_{D_U}]$ to $f^{k_u n_u}(D_U) \cap f^{-k_s n_s}(D_S)$, as shown in figure 16(b). Then $f^{k_n n_u}[\alpha_{D_U}]$ tightens onto $[\beta]$, and since β is a curve crossing $f^{-k_s n_s} D_S$, we also have $f^{k_s n_s}_{\min}[\beta]$ tightens onto $[\alpha_{D_S}]$ by Lemma 3.17. Hence $f^{k_u n_u + k_s n_s}_{\min}[\alpha_{D_U}]$ tightens onto $[\alpha_{D_S}]$ as required.

(b). Let D_U and D_S be regular domains containing Q_U and Q_S respectively such that $f^m_{\min}[\alpha_{D_U}]$ tightens onto $[\alpha_{D_S}]$ for some m. By Lemma 3.14(c), D_U and D_S are still regular domains for $(f,\widehat{T})$, and the curves $[\widehat{\alpha}_{D_U}]$ and $[\widehat{\alpha}_{D_S}]$ which cross D_U and D_S with minimal intersections relative to $\widehat{T}$ tighten onto $[\alpha_{D_U}]$ and $[\alpha_{D_S}]$. Then $f^m_{\min}[\widehat{\alpha}_{D_U}]$ crosses D_S, and by Lemma 3.17(b), some $f^{m+k_s n_s}_{\min}[\widehat{\alpha}_{D_U}]$ tightens onto $[\widehat{\alpha}_{D_S}]$ for some k_s.

□

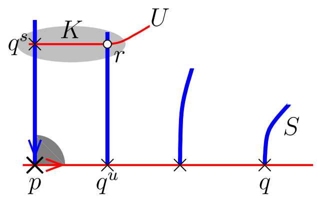

Figure 17: Backward iterates of S with endpoint q eventually form the opposite stable side of a regular domain with $q^s = f^{-kn}(q)$.

We now show that the hyperbolicity near T^P is enough to create intersections from which we can deduce regularity. Note that for this result we are concerned with f-extensions of T rather than $[f]$-minimal extensions.

Lemma 3.21 *Let T be a trellis for a diffeomorphism f, and let Q a quadrant of T. Let $q \in T^U(Q)$ be the endpoint of an unstable segment S on the Q-side of $T^S(Q)$, and let $q^s \in T^S(Q)$ be the endpoint of a stable segment U on the Q-side of $T^S(Q)$. Then there exists k such that $f^{-kn}(S)$ intersects U at a point r such that $\{p, f^{-kn}(q), q^s, r\}$ are the vertices of a regular domain for Q.*

Proof: By the Lambda lemma, $f^{-in}(S)$ limits on $W^S(Q)$ in the C^1 topology as $i \to \infty$. Take a neighbourhood K of q^s such that $f^{-n}(K) \cap T^S[p, q^s] = \varnothing$, and choose k such that $f^{-kn}(S)$ intersects U in at a point r in K such that $T^U[q^s, r] \subset K$, and the domain D with vertices at $\{p, f^{-kn}(q), q^s, r\}$ is a rectangle which does not intersect $f^{-n}(K)$, and such that $T^U[p, q] \cap D = T^U[p, f^{-kn}(q)]$, as shown in figure 17. Let $q^u = f^{-kn}(q)$, and $\widehat{T} = (T^U, f^{-nk}(T^S))$. Then $f^{-n}(T^U[q^s, r]) \cap D \subset f^{-n}(K) \cap D = \varnothing$, and $f^n(T^S[q^u, r])$ is contained in a segment with one endpoint at $f^n(q^u) \notin D$, so $f^n(T^S[q^u, r]) \cap D = \varnothing$. Hence D is a regular domain of $(f; \widehat{T})$. $\square$

3.4 Construction of minimal supertrellises

Most of the procedures we will use to construct diffeomorphisms in a given trellis mapping class rely on extending the original trellis and introducing new branches in a controlled way. The most important type of supertrellis is a minimal supertrellis, since the Nielsen entropy for a minimal supertrellis is the same as that of the original trellis mapping class.

We now show that we can introduce new stable and unstable branches at essential Nielsen classes of f and obtain a minimal supertrellis. The only difficulty here is on finding the correct initial segment of a branch; once this has been achieved, we can take minimal iterates. In the sequel, arithmetic on i will be assumed to be modulo n.

We can use theorem 3.6 to construct the minimal supertrellises. We say that a periodic point p of f *shadows* T^P if p is Nielsen equivalent to a periodic point of T^P; this means that it can be joined to a point of T^P by an exact curve α such that $f^n_{\min}[\alpha] = [\alpha]$.

Lemma 3.22 *Let $([f]; T)$ be a well-formed irreducible trellis mapping class. Then if p is an essential periodic point of $([f]; T)$ which does not shadow T^P, there is an $[f]$-minimal supertrellis of T for which p is contained in a region which can be chosen to be attracting, repelling, or both.*

Proof: Suppose that p is contained in a region R and has period n. We only consider the construction of a minimal supertrellis with new unstable curves; the construction of new stable

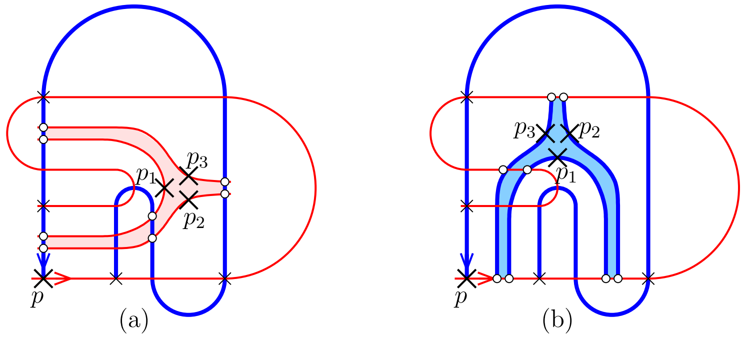

Figure 18: Supertrellises of the type-3 trellis shown in figure 1. The trellis shown in (a) is formed by introducing new stable branches and forming an attracting region, and that in (b) is formed by introducing new unstable branches and forming a repelling region.

curves follows by reversing time. Since p does not shadow T^P, there must be an exit segment S for P, and a simple curve γ from P to S such that the initial piece of $f_{\min}^{kn}[\gamma]$ tightens onto γ for some least k. Since p is essential, we must have $k > 1$. Let $[\gamma_i]$ be the initial sub-curve of $f_{\min}^{i}[\gamma]$ from $f^i(p)$ to a stable segment for $i = 0, \ldots, nk-1$. Since h is a homeomorphism, the curves γ_{jn+i} are cyclically ordered around $f^i(p)$, with γ_i being followed by $\gamma_{i+ln \bmod kn}$ for some l. We let $[\alpha_i]$ be the homotopy class $[\overline{\gamma}_i \cdot \gamma_{i+ln \bmod kn}]$. The homotopy classes $[\alpha_i]$ are distinct, and the representative curves α_i can be taken to be disjoint.

Let $[\beta_{i+1}] = f_{\min}[\alpha_i]$, and note that each β_i tightens onto α_i and has minimal intersections with T^S. Without loss of generality, we can isotope each β_i so that it contains α_i as a sub-curve. Since $f \circ \alpha_i$ and β_{i+1} are homotopic in curves for the trellis $(T^U, f(T^S))$, by theorem 3.6, we can find an isotopy h_t such that $h_t(T^U) = T^U$, $h_t(f(T^S)) = f(T^S)$, $h_0 = id$ and $h_1 \circ f \circ \alpha = \beta$. Then for the diffeomorphism $\widetilde{f} = h_1 \circ f$ we have $\widetilde{f}^{-1}(\beta(I)) = \alpha(I)$. Since the curves β_i have minimal intersections with T^S, the trellis $\widetilde{T}$ with $\widetilde{T}^U = T^U \cup \bigcup_{i=1}^{kn} \beta(I)$ and $\widetilde{T}^S = T^S$ is an $[f]$-minimal supertrellis of T. □

Lemma 3.23 *Let $([f];T)$ be a well-formed irreducible trellis mapping class, $p \in T^P$ and Q an attracting or repelling secant at p.*

(a). There there is a then there is an $[f]$-minimal supertrellis $\widetilde{T}$ of T for which Q is contained in an attracting or repelling region.

(b). If Q lies in a chaotic region of T, then there is an $[f]$-minimal supertrellis $\widetilde{T}$ of T for which there is a nontrivial branch of $\widetilde{T}^{U/S}$ through Q.

PROOF:

(a). Suppose Q is a repelling sector, and let R be the region containing Q. If R has only one stable boundary segment, then R is already a repelling region and we are done. Otherwise, let $[\gamma_0]$ be the exact homotopy class of curve in R joining two stable segments crossing $T^U(p)$, and define exact homotopy classes $[\gamma_i]$ recursively as follows. If $f_{\min}^n[\gamma_i]$ crosses R, we take $[\gamma_{i+1}]$ to be a subcurve $f_{\min}^n[\gamma_i]$ of crossing R.

Since R is not repelling, then there exists least i such that $f_{\min}^n[\gamma_i]$ tightens onto itself. We then take $\alpha = \gamma_i$, and α_i to be the sub-curve of $f_{\min}^i[\alpha_0]$ in the region containing $f^i(p)$ for $i = 0, \ldots, n-1$, and let $\beta_{i+1} = f_{\min}[\alpha_i]$. By homotoping the α_i and β_i if necessary, we

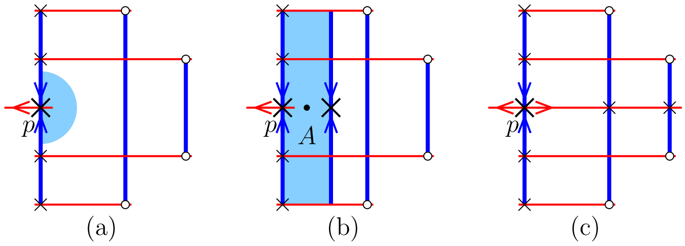

Figure 19: A trellis with a stable secant (a), and supertrellises. In (b) we create a new attracting region as in Lemma 3.23(a); in (c) we subdivide the sector into two quadrants as in Lemma 3.23(b).

can ensure that α_i is an initial sub-curve of β_i, and is tangent to T^U at $f^i(p)$. We can therefore find an isotopy h_t fixing T^U and $f(T^S)$ with $h_t(0) = id$ and $h_1 \circ f \circ \alpha = \beta$. The trellis $\widetilde{T}$ with $\widetilde{T}^U = T^U \cup \bigcup_{i=1}^{kn} \beta(I)$ and $\widetilde{T}^S = T^S$ is then an $[f]$-minimal supertrellis of T. The region $\widetilde{R}$ of $\widetilde{T}$ containing the secant Q is repelling since there can be no stable segment which enters $\widetilde{R}$.

(b). Suppose Q is a stable sector. Since Q is not contained in an attracting domain, there must be an exit segment S for p in the region containing Q. Let $[\alpha]$ be an exact curve from p to S, and let α_i be the initial piece of $f^i_{\min}[\alpha]$ for $i = 0, \ldots, n-1$. Let $\beta_{i+1} = f_{\min}[\alpha_i]$. Then since S is an exit segment for p, the initial piece of each $[\beta_i]$ tightens onto $[\alpha_i]$. By homotoping the α_i and β_i if necessary, we can ensure that α_i is an initial sub-curve of β_i, and is tangent to T^U at $f^i(p)$. We can therefore find an isotopy h_t fixing T^U and $f(T^S)$ with $h_t(0) = id$ and $h_1 \circ f \circ \alpha = \beta$. The trellis $\widetilde{T}$ with $\widetilde{T}^U = T^U \cup \bigcup_{i=1}^{kn} \beta(I)$ and $\widetilde{T}^S = T^S$ is then an $[f]$-minimal supertrellis of T.

□

We use these results to construct new branches at essential periodic orbits and trivial branches of a trellis. These cases are shown in figure 19. These branches are may used to create new attractors and repellers, as in (b), or to create a new branch dividing a sector into two quadrants as in (c).

We can similarly ensure that every coquadrant is contained in an attracting/repelling region.

Lemma 3.24 *Let $([f]; T)$ be a well-formed irreducible trellis mapping class, $p \in T^P$ and Q a coquadrant at p. There there is a then there is an $[f]$-minimal supertrellis $\widetilde{T}$ of T for which Q is contained in an attracting/repelling region.*

The proof follows that of Lemma 3.23(a).

4 Entropy-Minimising Diffeomorphisms

In this section, we show that the entropy bound obtained by the Nielsen entropy is sharp under the mild assumption that the trellis is well-formed. That is, the topological entropy of the graph representative g for a trellis type $[f; T]$ (which is the same as the Nielsen entropy $h_{niel}[f; T]$) is the infemum of the topological entropies of diffeomorphisms in the class.

We now give some examples which illustrate the hypotheses of the theorem. The following example shows that the hypothesis that the trellis be well-formed is necessary.

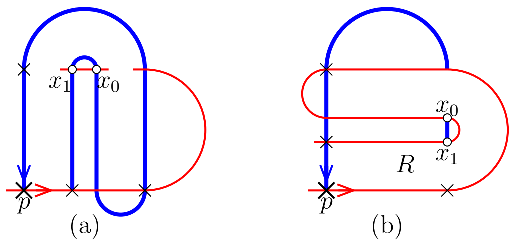

Figure 20: Two ill-formed trellis mapping classes. The trellis in (a) has Nielsen entropy $\log 2$, whereas the trellis in (b), which is the time reversal of that in (a), has Nielsen entropy zero.

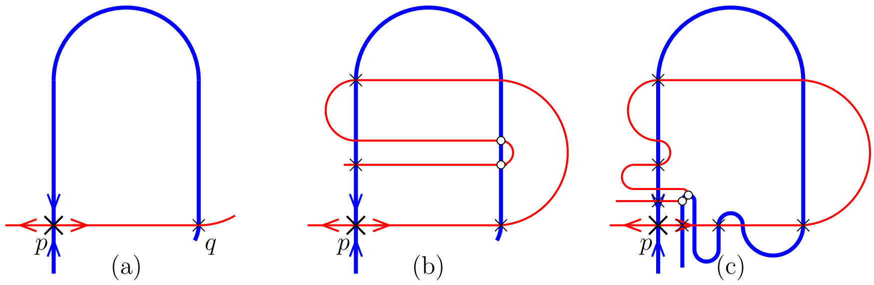

Figure 21: The trellis type $[f;T]$ shown in (a) has Nielsen entropy 0, but since T has a transverse homoclinic point, f has positive topological entropy. The trellis types $[f_2;T_2]$ and $[f_5;T_5]$ shown respectively in (b) and (c) are extensions of T with positive Nielsen entropy.

Example 4.1 The trellis mapping classes in figure 20 are not well-formed. The Nielsen entropy of the trellis mapping class in (a) is equal to $\log 2$, so any diffeomorphism in the trellis mapping class must have topological entropy at least $\log 2$. Since the Smale horseshoe map has this trellis type, the topological entropy of the trellis type is exactly $\log 2$. The trellis mapping class of figure 20(b) is conjugate to the time-reversal of the trellis mapping class in (a). Since the topological entropy of a diffeomorphism is the same as that of its inverse, any diffeomorphism in this trellis mapping class must have topological entropy at least $\log 2$. However, all the edges of the graph representative are control edges, so the Nielsen entropy is zero.

The above example illustrates that a trellis which is not well-formed may have Nielsen entropy strictly less than the topological entropy, and may even have different Nielsen entropy from its time-reversal. The following example shows that even if a trellis mapping class is well-formed, it is not necessarily true that the Nielsen entropy is realised.

Example 4.2 A trivial example is of a well-formed trellis type for which the Nielsen entropy equals the topological entropy, but no diffeomorphism realises the entropy bound, is the planar trellis type with a single transverse homoclinic intersection, as shown in figure 21(a). The Nielsen entropy of this trellis type is $h_{niel}[f;T]=0$, but every diffeomorphism with a transverse homoclinic point has strictly positive topological entropy. It is simple to construct trellis types with topological entropy arbitrarily close to zero; in figure 21(b) we show the trellis type $[f_2;T_2]$ for which $h_{niel}[f_2;T_2]\approx 0.693$, and in figure 21(c) we show the trellis type $[f_5;T_5]$ for which $h_{niel}[f_5;T_5]\approx 0.372$.

We now give a nontrivial example.

Example 4.3 A trellis type $[f;T]$ for which the Nielsen entropy is not realisable is shown in figure 22(a). Taking three backward minimal iterates of the segment S eventually yields a

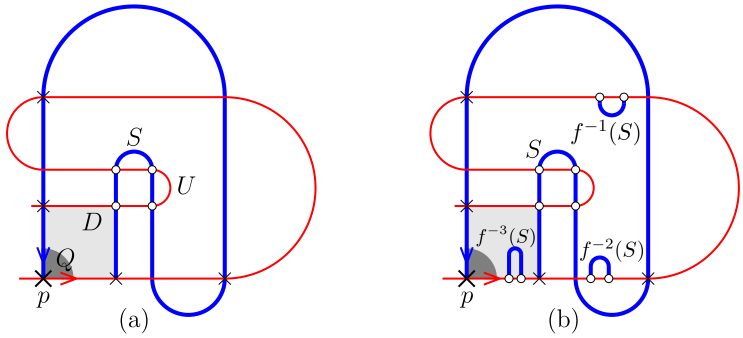

Figure 22: A trellis for which the Nielsen entropy is not realisable.

segment $f^{-3}(S)$ lying in the domain D containing quadrant Q as shown in figure 22(b). Hence, by the Lambda lemma, under any diffeomorphism f in the trellis mapping class, $f^{-n}(S)$ approaches T^S as $n \to \infty$, so contains an intersection with T^U for some n, even though any *minimal* backward iterate of S has no intersections with T^U. Similarly, $f^n(U)$ must intersect T^S for some n even though any minimal iterate does not. Once one of these extra intersections occurs, the Nielsen entropy of the resulting trellis type can be shown to increase. Hence the Nielsen entropy of $[f;\widehat{T}]$ is greater than that of $[f;T]$ for some f-extension $\widehat{T}$ of T, so $h_{top}(f) > h_{niel}[f;T]$.

The realisability of the entropy bound is closely related to the existence of a diffeomorphism for which every extension is a minimal extension. In the case where this infemum is realised, we show how to construct a minimal-entropy uniformly-hyperbolic diffeomorphism in the trellis mapping class. Otherwise, we show, for any $\epsilon > 0$, how to construct a diffeomorphism whose entropy is within ϵ of the Nielsen entropy.

4.1 Existence of entropy minimisers

To prove the existence of a diffeomorphism in a trellis mapping class whose topological entropy is the Nielsen entropy of the class, we reduce to the case for which every chaotic region is a rectangle. We then construct such a diffeomorphism for a particularly simple class of trellises.

Theorem 4.4 (Existence of entropy minimisers) *Let $([f];T)$ be a trellis mapping class. Suppose there is a diffeomorphism $\widehat{f}$ isotopic to f relative to T such that every $\widehat{f}$-extension of T is minimal. Then there is a uniformly-hyperbolic diffeomorphism $\widetilde{f}$ isotopic to f relative to T such that every extension of T by $\widetilde{f}$ is minimal, and $h_{top}(\widetilde{f}) = h_{niel}[f;T]$.*

Proof: By Theorem 2.1, any trellis type $[\widehat{f};\widehat{T}]$ with $\widehat{T} = (T^U, \widehat{f}^{-n}(T^S))$ has the same Nielsen entropy at $[f;T]$. By Lemma 3.21 we can therefore take a $\widehat{f}$-extension T_1 of T such that every quadrant of T_1 lies in a regular domain, and by irreducibility, we can ensure that every region of T_1 is a topological disc or annulus. By introducing new unstable curves as in Lemma 3.23(a), we can take an $[f_1]$-minimal supertrellis of T_2 of T_1 such that every secant of T_2 lies in an attracting or repelling region. We can further ensure that every coquadrant lies in an attracting/repelling region by Lemma 3.24.

Now suppose there is a chaotic region R of $([f_2];T_2)$ which is not a rectangle. Then the graph representative $(g_2;G_2,W_2)$ of $[f_2;T_2]$ has a peripheral loop or a valence-n vertex in R which corresponds to a boundary component or essential periodic orbit which does not shadow T^S. Introducing new stable curves for all such R as in Lemma 3.22 gives an $[f_2]$-minimal supertrellis T_3 for a diffeomorphism f_3 for which every chaotic region is a rectangle.

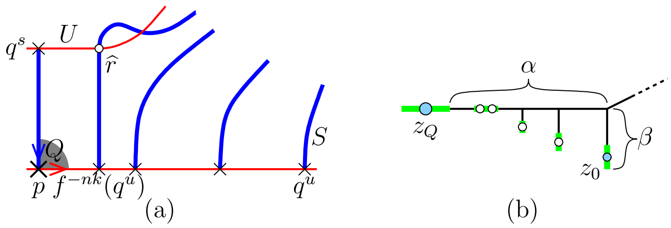

Figure 23: (a) A non-minimal iterate of a trellis to form a regular quadrant, and (b) the corresponding graph representative.

Taking an $[f_3]$-minimal iterate T_4 of T_3 with $T_4 = (T_3^U, f_4^{-1}(T_3^S))$ gives a trellis mapping class such that every domain D of $([f_4];T_4)$ with boundary in $T_4^U \cup f_4(T_4^S)$ containing a chaotic region is a rectangle. Foliate every domain D containing a chaotic region, foliate D and $f(D)$ by an unstable foliation $\mathcal{F}^{\mathcal{U}}$ parallel to T^U and a transverse stable foliation $\mathcal{F}^{\mathcal{S}}$ parallel to T^S. Isotope f_4 to obtain a diffeomorphism $\widetilde{f}$ which preserves the stable and unstable foliations, and for which all points of non-chaotic regions are in the basin of a stable or unstable periodic orbit. Let $(\widetilde{g};\widetilde{G},\widetilde{W})$ be a graph representative of $[f_4;T_4]$, and $\pi : \mathcal{C}T_4 \longrightarrow (\widetilde{G},\widetilde{W})$ be a deformation-retract which collapses each leaf of $\mathcal{F}^{\mathcal{S}}$ onto a point of $\widetilde{G}$. Then $\pi \circ \widetilde{f} = \widetilde{g} \circ \pi$ on every chaotic region R, so $h_{top}(\widetilde{f}) = h_{top}(\widetilde{g}) = h_{niel}[f;T]$. Further, it is clear that every $\widetilde{f}$-extension of T_4 is minimal, so every f-extension of T is minimal. □

4.2 Approximate entropy minimising diffeomorphisms

We now show that, for any $\epsilon > 0$, we can find a diffeomorphism $\widehat{f}$ isotopic to f relative to T which has entropy less than $h_{top}[f;T] + \epsilon$. We first consider how to perform non-minimal extensions without increasing the Nielsen entropy above $h_{niel}[f;T] + \epsilon$. To control the entropy bound, we need to consider the graph representative, which introduces some technical difficulties. To avoid having to deal directly with graph representative in the sequel, we prove a result, Lemma 4.5, which applies directly to trellis mapping classes.

Our final goal is to to construct a trellis mapping class which satisfies the conditions of Theorem 4.4. To do this, we may need to introduce new periodic points to the trellis to create attracting and repelling regions. We then iterate curves bounding an attractor or repeller into regular domains, and finally take non-minimal iterates to move bigon boundaries into attractors and repellers.

Lemma 4.5 *Let $([f];T)$ be a trellis mapping class and Q a period-n quadrant of $([f];T)$. Let S be the segment of T^S with endpoint q^u on the Q-side of $T^U(Q)$, and U be the segment of T^U with endpoint q^s on the Q-side of $T^S(Q)$. Then for any $\epsilon > 0$ there is a minimal stable extension $([\widehat{f}];\widehat{T})$ such that $h_{niel}[\widehat{f};\widehat{T}] < h_{niel}[f;T] + \epsilon$, and an integer k such that $\widehat{f}^{-kn}(S) \subset \widehat{T}^S$ and intersects U at a point $\widehat{r}$ such that $\{p, f^{-kn}(q^u), q^s, \widehat{r}\}$ are the vertices of a regular domain for Q.*

Proof: Choose λ and λ_ϵ such that $h_{niel}[f;T] < \lambda < \lambda_\epsilon < h_{niel}[f;T] + \epsilon$. Let $(g;G,W)$ be the graph representative of $[f;T]$ and l be a length function on G such that for all edges e of G, $l(g(e)) < \lambda l(e)$. Let z_Q be the control edge crossing $T^S[p,q^s]$, and z_0 be the control edge crossing S. Let α be an edge-path starting at z_Q and finishing at the first vertex v between z_Q and z_0, and let β be the edge-path from v to the end of z_0.

Since $l(g(\alpha))/\lambda_\epsilon < l(g(\alpha))/\lambda \leqslant l(\alpha)$, there exists k such that $l(g(\alpha))/\lambda_\epsilon + 2l(\beta)/\lambda_\epsilon^{nk} < l(\alpha)$. Let $\widetilde{f}$ be a diffeomorphism such that $\widetilde{f}^{1-kn}(T^S)$ is a minimal iterate of T^S, and set $\widetilde{T} =$

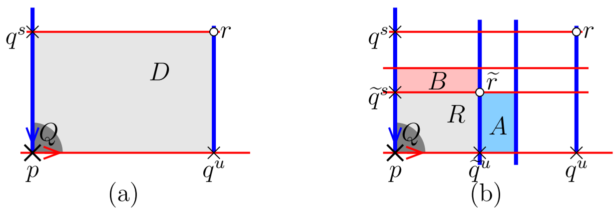

Figure 24: (a) A quadrant Q in a regular domain D, and (b) the same quadrant contained in a region R which is a regular domain with an attractor A across $T^S[\widetilde{q}^u, \widetilde{r}]$ and a repeller B across $T^U[\widetilde{q}^s, \widetilde{r}]$.

$(T^U, \widetilde{f}^{1-kn}(T^S))$. Then $[\widetilde{f}; \widetilde{T}]$ has a graph representative $(\widetilde{g}; \widetilde{G}, \widetilde{W})$ for which l extends to a length function with $l(\beta_i) = l(\beta)/\lambda_\epsilon^i$ for the ith backward iterate β_i of β. Let $\widehat{f}$ be a diffeomorphism which is isotopic to $\widetilde{f}$ relative to $\widetilde{T}$, for which $\widehat{f}^{-nk}(S)$ intersects U transversely at two points, including a point $\widehat{r}$, but otherwise has minimal intersections with T^U, as shown in figure 23. Let $\widehat{T} = (T^U, \widehat{f}^{-nk}(T^S))$, and $(\widehat{g}; \widehat{G}, \widehat{W})$ be the graph representative of $[\widehat{f}; \widehat{T}]$. The graph map $\widehat{g}$ maps α to $g(\alpha)$ and twice over β_{nk-1}, with total length $l(\widehat{g}(\alpha)) = l(g(\alpha)) + 2l(\beta)/\lambda_\epsilon^{nk-1}$, so $l(\widehat{g}(\alpha)) < \lambda_\epsilon l(\alpha)$ as required. Since the image of all other edges is unchanged, the growth rate of l under $\widehat{g}$ is less than λ_ϵ, so $h_{niel}[\widehat{f}, \widehat{T}] = h_{top}(\widehat{g}) < \log \lambda_\epsilon$. □

Before proving the main theorem we show how to ensure that every quadrant Q is contained in a regular domain whose opposite sides bound attractors or repellers.

Lemma 4.6 *Let $([f]; T)$ be a transitive trellis mapping class. Then there is a minimal supertrellis $([\widetilde{f}]; \widetilde{T})$ of $([f]; T)$ such that every quadrant Q is contained in a region R such that R is a regular domain for Q such that $\widetilde{T}^S[\widetilde{q}^u, \widetilde{r}]$ is a boundary arc of a stable domain and $\widetilde{T}^U[\widetilde{q}^s, \widetilde{r}]$ is the boundary of an unstable region.*

PROOF: Let P be an essential periodic orbit of $([f]; T)$ which does not shadow T^P. Construct an $[f]$-minimal supertrellis T_1 for a diffeomorphism f_1 as in Lemma 3.22 containing P in a stable region. Since $([f]; T)$ is transitive, for any regular domain D, the homotopy class $[\alpha_D]$ must have some minimal iterate $f^k_{\min}[\alpha_D]$ intersecting the new stable curves of T_1. Taking T_2 to be the trellis formed by taking $[f_1]$-minimal iterates of $T_1^S \setminus T^S$, we can ensure that every regular domain D is crosses by curves of T_2^S. Then every quadrant Q is contained in a region R of T_2 such that R is a regular domain for Q and $T_2^S[q^u, r]$ is a boundary arc of a stable domain, as shown in figure 24. A supertrellis T_3 such that every quadrant Q is contained in a region R of T_2 such that R is a regular domain for Q and $T_3^U[q^s, r]$ is a boundary arc of an unstable domain can be constructed in a similar way. □

Theorem 4.7 (Existence of approximate entropy minimisers) *Let $([f]; T)$ be a well-formed trellis mapping class. Then for every $\epsilon > 0$, there exists a diffeomorphism $\widehat{f} \in [f]_T$ such that $h_{top}(\widehat{f}) < h_{niel}[f; T] + \epsilon$.*

PROOF: We repeatedly construct trellis mapping classes $([f_i]; T_i)$ where T_{i+1} is a supertrellis of T_i and f_{i+1} is isotopic to f_i relative to T_i, at each stage ensuring that $h_{niel}[f_i; T_i] < h_{niel}[f; T] + \epsilon = \log \lambda_\epsilon$. Without loss of generality we assume $([f]; T)$ is irreducible, since if $([f]; T)$ is reducible, we can consider irreducible components. Take $([f_0]; T_0) = ([f], T)$. By irreducibility, there

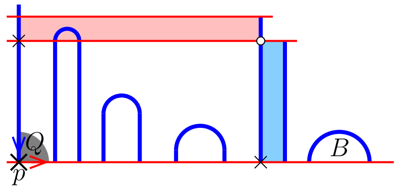

Figure 25: Backward iterates of B give an inner bigon in a repelling domain.

exists a minimal extension T_1 of T_0 such that for every pair of nontrivial branches $T_1^U[p_u, b_u]$ and $T_1^S[p_s, b_s]$, there are points $p_u = p_0, p_1, \ldots, p_n = p_s$ such that $T_1^U(p_u, b_u] \cap T_1^S(p_1) \neq \varnothing$, $T_1^U(p_i) \cap T_1^S(p_{i+1}) \neq \varnothing$ for $1 \leqslant i < n-1$ and $T_1^U(p_{n-1}) \cap T_1^S(p_s, b_s] \neq \varnothing$. By Lemma 4.5, there exists $f_2 \in [f_1]_{T_1}$ and an f_2-extension T_2 of T_1 such that every quadrant of T_2 is contained in a regular rectangular region.

We next construct a non-minimal f_3-extension T_3 of T_2 such that every nontrivial branch of T_3^U intersects every nontrivial branch of T_3^S as follows: Suppose Q_U, Q_S and Q are quadrants such that $T_3^U(Q_U)$ intersects $T_3^S(Q)$ and $T_3^U(Q)$ intersects $T_3^S(Q_S)$. Then can take minimal iterates of $T_3^U(Q_U)$ and $T_3^S(Q_S)$ until the closest intersections to the periodic point p at Q along $T_3^S(Q)$ and $T_3^U(Q)$ lie on $T^U(Q_U)$ and $T^S(Q_S)$, respectively. Then an application of Lemma 4.5 shows that there is a non-minimal extension such that $T^U(Q_U) \cap T^S(Q_S) \neq \varnothing$. By considering all pairs of quadrants Q_U and Q_S, we can ensure that every nontrivial branch of T_3^U intersects every nontrivial branch of T_3^S.

Now construct a non-minimal f_4-extension T_4 of T_3 such that there is an alpha-chain linking any quadrant Q_U of T_4 to any other quadrant Q_S. Since $T_3^S(Q_S)$ intersects $T_3^U(Q_U)$, by taking minimal backward iterates of $T_3^S(Q_S)$ we can ensure that $T^S(Q_S)$ crosses $U(Q_U)$. Then by Lemma 4.5, there is a non-minimal extension such that the branch $T^S(Q_S)$ crosses the regular region containing Q_U. Similarly, we can ensure $T^U(Q_U)$ crosses the regular region containing Q_S. A further application of Lemma 4.5 gives a non-minimal extension with an alpha-chain linking Q_U to Q_S. Repeating this construction for all quadrants gives the required trellis T_4.

Let T_5 be a minimal supertrellis of T_4 such that all secants and coquadrants of T_4 are contained in a non-chaotic region, as given by Lemma 3.23(a) and Lemma 3.24. If $([f_5]; T_5)$ has no attractors or no repellers, take a further minimal supertrellis T_6, with at least one attractor and one repeller, as given by Lemma 3.22. Since $([f_6]; T_6)$ is transitive, we can take a minimal extension T_7 by iterating the stable and unstable curves bounding stable and unstable regions such that all quadrants of T_7 are contained in a regular region with the opposite stable side bounding an attracting domain, and the opposite unstable side bounding a repelling domain.

Now, for any inner bigon B such that the stable boundary $\partial B \cap T^S$ is an almost wandering segment, we take backward iterates of $\partial B \cap T^S$ until $f^{-k}(B)$ is contained in a regular domain with stable and unstable regions across its opposite sides. By Lemma 4.5, we can take further backwards iterates until $f^{-n})(B)$ contains a single bigon in an unstable domain, as shown in figure 25. Applying this procedure to all bigons gives a trellis mapping class $([f_8]; T_8)$.

Finally, we take a minimal supertrellis T_9 such that every chaotic region is a rectangle by introducing new stable curves as in Lemma 3.22. Then by Theorem 4.4 there is a uniformly-hyperbolic diffeomorphism $\widehat{f} \in [f_9]_{T_9}$, and hence in $[f]_T$ such that $h_{top}(\widehat{f}) = h_{niel}[f_9; T_9] < h_{niel}[f; T] + \epsilon$ as required. □

We now prove the existence of pseudo-Anosov map which approximate the Nielsen entropy. The condition on attracting and repelling domains is to ensure that the trellis mapping class contains a pseudo-Anosov map, since pseudo-Anosov maps have no attracting or repelling do-

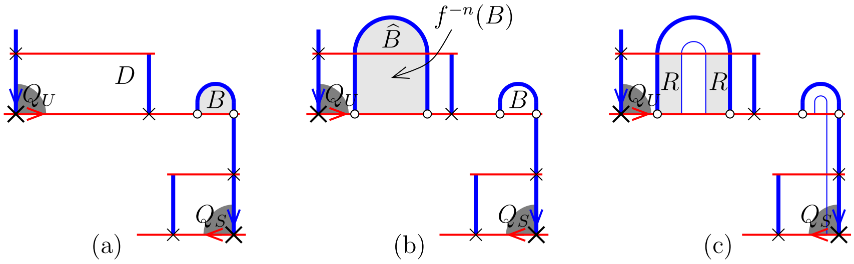

Figure 26: The domain with vertices q_0 and q_1 must have periodic orbits in the rectangles R_0 and R_1.

mains. The strategy is to create the periodic orbits which will give the one-prong singularities of the pseudo-Anosov map. Many of the steps of the proof mimic those the proof of Theorem 4.7.

Theorem 4.8 (Existence of pseudo-Anosov representatives) *Let $([f];T)$ be a trellis mapping class with no attractors or repellers. Then for any $\epsilon > 0$ there exists a pseudo-Anosov diffeomorphism $\widehat{f} \in [f]_T$ such that $h_{top}(\widehat{f}) < h_{top}[f;T] + \epsilon$.*

PROOF: We repeatedly construct trellis mapping classes $([f_i];T_i)$ where T_{i+1} is a supertrellis of T_i and f_{i+1} is isotopic to f_i relative to T_i.

If T has a coquadrant, take an extension T_1 such that the two trivial branches inside the coquadrant intersect in a single transverse homoclinic point. This does not affect the Nielsen entropy. If T has any other trivial branches, take a minimal extension T_2 such that these branches have an intersection point, which is possible by Lemma 3.23(b) since T has no attracting or repelling regions. As in the proof of Theorem 4.7, take an extension T_3 such that every branch of T_3 intersects every other. Further, we can ensure that every branch intersects every other with both orientations by constructing new intersections as adjacent pairs or triples with alternating orientations. Take a further extension T_4 such that $([f_4];T_4)$ is transitive.

Let B be an inner bigon of T_4, as shown in figure 26(a). By Lemma 4.5, we can find a stable extension such that $f^{-n}(\partial B \cap T^S)$ crosses the regular region D containing Q_U for some n. Further, by removing intersections by a pruning isotopy if necessary, we can ensure that $f^{-n}(\partial B \cap T^S)$ crosses D twice and gives a new inner bigon $\widehat{B}$, as shown in figure 26(b). A further application of Lemma 4.5 shows that we can find a stable extension T_4 such that $\partial\widehat{B} \cap T^U$ crosses some regular domain, as shown in figure 26(c). Since the regions R are mapped over by $[\alpha_{Q_U}]$ and map over $[\alpha_{Q_S}]$, they must contain a periodic orbit, since T is transitive. Applying this construction for every inner bigon of $([f_4];T_4)$ gives a non-minimal extension T_5 such that every bigon of T_4 is a domain of T_5 containing an essential periodic orbit of $([f_5];T_5)$.

Let $([f_6];T_6)$ be the trellis mapping class obtained by puncturing at a periodic orbit in every inner bigon of T^4 to give a surface $\widehat{M}$. Then $h_{niel}[f_6;T_6] = h_{niel}[f_5;T_5]$. Since T_4 is a subtrellis of T_6, we can take a trellis mapping class $([f_7];T_7) = ([f_6];T_4)$ in the surface M_6. Since every inner bigon of T_4 contains component of $\partial\widehat{M}$, the trellis mapping class $([f_7];T_7)$ has no inner bigons. The graph representative $(g_7;G_7,W_7)$ is locally injective except at cusps, so is efficient and hence is a train-track map for a pseudo-Anosov diffeomorphism $\widehat{f} \in ([f_7];T_7)$. Then $h_{top}(\widehat{f}) = h_{niel}[f_7;T_7] \leqslant h_{niel}[f_6;T_6] < h_{niel}[f;T] + \epsilon$ as required. □

4.3 Non-existence of entropy minimisers

In this section we show that if an irreducible trellis type has an entropy minimiser, then removing intersections results in a trellis type with strictly smaller Nielsen entropy. We first prove that a

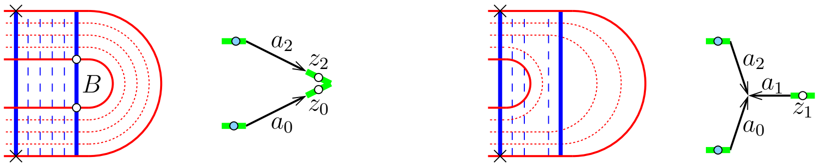

(a) (b)

Figure 27: Change in trellis and graph at a homoclinic bifurcation.

transitive trellis has a transitive graph representative.

Lemma 4.9 *Let* $([f];T)$ *be an transitive trellis mapping class. Then the graph representative* $(g;G,W)$ *of* $([f];T)$ *has a single transitive component with positive topological entropy.*

Proof: Consider (G,W) to be embedded as a homotopy retract of $\mathcal{C}T$ For every homotopy class $[\alpha]$ in $\mathcal{C}T$ with endpoints in T^S, let α be the representative curve in G. Let $\overline{G} = \bigcap_{n=0}^{\infty} g^n(G)$ be the essential subgraph of G, and let e be any expanding edge of G. By irreducibility, there is a quadrant Q contained in a regular domain D such that $g^n(e)$ tightens to $[\alpha_D]$ for some n.

We now consider preimages of edges. Again, let e be an expanding edge of G. There exists a tight curve β_0 in $\widetilde{G}$ with endpoints in W such that β_0 is homotopic to a subinterval of a branch $T^U[p,b]$. We then find a curve β_1 in $\widetilde{G}$ homotopic to a subinterval of a branch at $f^{-1}(p)$. Proceeding recursively gives an edge-path $\beta_n = \alpha_D$ where β_i is a sub-path of $g(\beta_{i+1})$ for $0 \leqslant i < n$. Therefore, there exists a quadrant Q contained in a regular domain D and an integer n such that $g^n(\alpha_D) \supset e$.

Now if Q_U and Q_S are any two regular quadrants, and are contained in regular domains D_U and D_S, respectively, there exists n such that $g^n(\alpha_{D_U})$ contains a sub-curve α_{D_S}. Since $(f;T)$ is irreducible, we can find N such that for any two quadrants and any $n \geqslant N$, the iterate $g^n(\alpha_{Q_U})$ contains a sub-path α_{Q_S}. Therefore, any such curve α generates the same graph component under iteration.

Combining these results shows that there exists N such that if e_1 and e_2 are any two edges of $\overline{G}$ and $n \geqslant N$, then $e_2 \subset g^n(e_1)$. □

The following result shows that any isotopy removing intersections results in a trellis mapping class with strictly smaller Nielsen entropy.

Theorem 4.10 *Let* $([f];T)$ *be a well-formed irreducible trellis mapping class, and suppose* f *is a uniformly-hyperbolic diffeomorphism such that* $h_{top}(f) = h_{niel}[f;T]$*. Then if* $([\overline{f}];\overline{T})$ *is a trellis mapping class which is obtained from* $([f];T)$ *by a pruning isotopy, then* $h_{niel}[\overline{f};\overline{T}] < h_{niel}[f;T]$.

Proof: By taking an f-supertrellis T_1 of T, we can ensure that every quadrant is regular, and contained in a rectangular region, that $([f];T_1)$ is transitive, and that that every inner bigon of $([f];T_1)$ is contained in a larger domain with the topology of figure 27(a). Clearly, $h_{niel}[f;T] \leqslant h_{niel}[f;T_1] \leqslant h_{top}(f)$, and hence we have $h_{niel}[f;T_1] = h_{top}(f)$. Further, if $(g_1;G_1,W_1)$ is the graph representative of $([f];T_1)$, then (G_1,W_1) has free edges a_0 and a_2 and control edges z_0 and z_2 as shown. Since $([f];T_1)$ is transitive, a_0 and a_1 lie in the transitive component of g_1 with positive entropy, and there is an edge e mapping $g_1(e) = \ldots a_0 z_0 \bar{z}_2 \bar{a}_2, \ldots$.

Now consider the effect of a pruning isotopy yielding a trellis mapping class $([f_2];T_2)$ with graph representative $(g_2;G_2,W_2)$ by removing one pair of orbits on the same inner bigons. The resulting trellis and graph locally have the topology of the right of figure 27(b). The control edges z_0 and z_2 are folded to give control edge z_1, and the edges a_0 and a_2 are partially folded together yielding a new edge a_1. The image of e is tightened from a path $\ldots a_0\bar{a}_1 z_1 \bar{z}_1 a_1 \bar{a}_2 \ldots$ to $g_2(e) = \ldots a_0 \bar{a}_2 \ldots$, which means that $h_{top}(g_2) < h_{top}(g_1)$. Hence the entropy of the graph representative of $([f_2]; T_2)$ is less than that of $([f]; T_1)$, so $h_{niel}[f_2; T_2] < h_{niel}[f; T_1] = h_{niel}[f; T]$. Since we can then prune $([f_2]; T_2)$ to obtain $([\overline{f}]; \overline{T})$, we have $h_{niel}[\overline{f}; \overline{T}] \leqslant h_{niel}[f_2; T_2]$, hence $h_{niel}[\overline{f}; \overline{T}] < h_{niel}[f; T]$ as required. □

We use this to prove that the sufficient condition given for the existence of entropy minimisers in Theorem 4.4 is necessary.

Theorem 4.11 (Non-existence of entropy minimisers) *Let $([f]; T)$ be an irreducible trellis mapping class. Suppose that for every diffeomorphism $\widehat{f}$ isotopic to f relative to T there is a $\widehat{f}$-extension of T which is not minimal. Then there does not exist a diffeomorphism isotopic to f relative to T whose topological entropy equals $h_{niel}[f; T]$.*

PROOF: Suppose there is an f-extension T_1 of T which is not minimal. Let T_2 be a transitive f-extension of T_1, and T_3 be a further extension such that every non-wandering segment of T crosses a regular domain of T_3. Then we have entropy bound $h_{top}(f) \geqslant h_{niel}[f; T_3]$.

Take T_4 to be an f_4-supertrellis of $([f]; T_3)$ which is $[f]$-minimal and for which the opposite sides of every regular region bound a stable or unstable region. Then every nonwandering segment of T enters a stable or unstable region of T_4 under iterates of f_4. Isotope to remove intersections of T_4 not in T_1 to obtain a trellis mapping class $([f_5]; T_5)$ such that T_5 is a supertrellis of T_1 and $([f_5]; T_5)$ satisfies the conditions of Theorem 4.4. Then $h_{niel}[f_3; T_3] = h_{niel}[f_4; T_4] \geqslant h_{niel}[f_5; T_5]$, but by Theorem 4.10, $h_{niel}([f_5]; T_5) > h_{niel}[f; T]$. Combining these inequalities we have $h_{top}(f) \geqslant h_{niel}[f_5; T_5] > h_{niel}[f; T]$ as required. □

The results of this section show that many fundamental properties of an irreducible trellis type $[f; T]$ depends on whether there is a diffeomorphism isotopic to f relative to T for which $h_{top}(f) = h_{niel}[f; T]$. If such a diffeomorphism exists, then the entropy of the trellis type is carried in a uniformly hyperbolic diffeomorphism, but is fragile in the sense that any pruning will reduce the Nielsen entropy, and any diffeomorphism $\widehat{f}$ for which some extension is non-minimal must have strictly greater topological entropy. If no such diffeomorphism exists, every diffeomorphism in the trellis type has an extension which is non-minimal, but pruning this extension may give a trellis type for which the entropy is still greater than the Nielsen entropy of $[f; T]$.

References

[1] Mladen Bestvina and Michael Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.

[2] Philip L. Boyland. Isotopy stability of dynamics on surfaces. In Marcy Barge and Krystyna Kuperberg, editors, *Geometry and Topology in Dynamics*, number 246 in Contemporary Mathmatics, pages 17–45, Providence, Rhode Island, 1999. American Mathematical Society.

[3] A. Casson and S. Bleiler. *Automorphisms of surfaces, after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, 1988.

[4] Pieter Collins. Relative periodic point theory. *Topology Appl.*, 115(1):97–114, 2001.

[5] Pieter Collins. Dynamics of surface diffeomorphisms relative to homoclinic and heteroclinic orbits. *Dynam. Sys.*, 19(1):1–39, 2004.

[6] David Epstein. Curves on 2-manifolds and isotopies. *Acta Math.*, 115:83–107, 1966.

[7] Albert Fathi. Homotopical stability of pseudo-anosov diffeomorphisms. *Ergodic Theory Dynamical Systems*, 10(2):387–294, 1990.

[8] John Franks and Michael Misiurewicz. Cycles for disk homeomorphisms and thick trees. In *Nielsen Theory and Dynamical Systems*, number 152 in Contemporary Mathematics, pages 69–139, 1993.

[9] Michael Handel. Global shadowing of pseudo-Anosov homeomorphisms. *Ergodic Theory Dynamical Systems*, 5(3):373–377, 1985.

[10] Boju Jiang. Estimation of the number of periodic orbits. *Pacific J. Math.*, 172:151–185, 1996.