FLOWS WITH KNOTTED CLOSED ORBITS

JOHN FRANKS AND MICHAEL C. SULLIVAN

ABSTRACT. We survey results concerning dynamics of flows on S^3 with special attention to the relationship between dynamical invariants and invariants of geometric topology.

1. Introduction

One of the key objects of study in the field of dynamical systems is the topological structure of the solutions of ordinary differential equations. Formally we can think of an action of the real numbers $\mathbf{R}$ on a manifold and one wants to study the topology of the orbits of this action. All such orbits are either injectively immersed copies of $\mathbf{R}$ or embedded copies of S^1. The latter are called *periodic orbits* since if we consider the flow

$$\phi_t : M \to M$$

where $t \in \mathbf{R}$ and M is the manifold on which the flow lives then a point $x \in M$ (sometimes called an initial condition) lies on a periodic orbit if and only if $\phi_{t_0}(x) = x$ for some $t_0 \in \mathbf{R}$. An orbit of a flow is compact if and only if it is periodic.

From a classical differential equations point of view (though probably not from the topologist's viewpoint) the most useful manifolds on which to investigate flows are Euclidean spaces $\mathbf{R}^n$, or perhaps the sphere S^n if it is particularly helpful to be in a compact setting.

An area where this classical setting intersects a rich topological environment is the study of flows on S^3 or $\mathbf{R}^3$ where periodic orbits can be knotted or linked. We can then ask what knot types occur as periodic orbits of a given flow (or of any flow) as well as numerous other natural questions relating the dynamics of a flow to the topology of the knots and links which occur as closed orbits of it.

2. Flows in Boxes

We begin by describing a very simple but remarkably rich construction (introduced by Birman and Williams [2]) to produce examples of flows with many knotted orbits. We will see that this is much more than a method of producing examples, however. Indeed, the behavior exhibited in these examples occurs quite generally in flows on S^3.

Date: January 7, 1999.

We want to build simple examples of flows on subsets of S^3. Here is a trivial example, but one which is important as a building block. Consider the cube, or box, $B = I \times I \times I$ where $I = [-1, 1]$ which we parameterize by (x, y, s) with $x, y, s \in I$. We consider the "flow" on B whose orbits are the line segments $(x, y, s), s \in I$. We orient the orbits in the direction of decreasing s. And we specify a constant speed of 2 so an orbit will enter a box and remain for one unit of time before exiting. Formally $\Phi_t(x, y, s) = (x, y, s - 2t)$. Strictly speaking this is not a flow because the orbits will exit B after a finite amount of time (both positively and negatively).

Conceptually we think of orbits entering B at the top $(s = 1)$, flowing downward and exiting at the bottom $(s = -1)$. On the sides of B the orbits lie in the boundary. This is not a very interesting example as it stands, but we can construct very interesting examples by using multiple copies of B as building blocks and attaching parts of the bottom faces of one cube to the top faces of another. It is better to refer to *entering* or *exiting* faces rather than top or bottom faces, because we will want to embed collections of attached cubes in $\mathbf{R}^3$ or S^3 in complicated ways for which "top" and "bottom" may not be useful descriptions.

When we attach part of the exiting face of one box to part of the entering face of another we will, of course assume that the orbit segments exiting the first at a point of attachment immediately enter the second forming a longer orbit segment crossing both boxes. We will only allow attachments of a very special form. In particular, we will make identifications using affine maps of faces which preserve the line segments in the faces parallel to the x and y axes. That is, a point $(x, y, -1)$ in the exiting region of box B_1 may be identified with its image under an affine map, $(\lambda x + a_0, \lambda^{-1} y + b_0, 1)$, in the entering region of box B_2 provided $|\lambda| > 1$ and the image under this affine map of the exiting region intersects the entering region as shown in Figure 1.

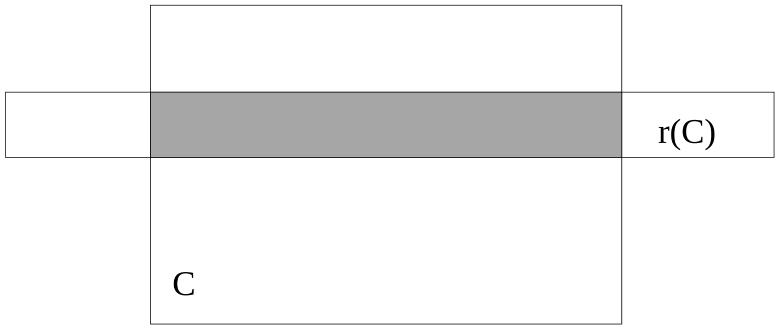

FIGURE 1. The affine attaching map r

In particular the affine map, which we will denote $r(x, y)$, must satisfy $r(x, y) = (\lambda x + a_0, \lambda^{-1} y + b_0)$ where

$$\begin{aligned} |\lambda + a_0| &> 1 \\ |-\lambda + a_0| &> 1 \\ |\lambda^{-1} + b_0| &< 1 \\ |-\lambda^{-1} + b_0| &< 1 \end{aligned}$$

so that if $C = I \times I$, the image under this map of C goes completely across C in the direction parallel to the x-axis, but in the direction parallel to the y-axis C goes completely across the image of C under this affine map.

Parts of one exiting region may be attached to more than one entering region of one or more than one box. The regions of attachment for all entering and exiting faces must be disjoint. See Figure 2. Different attachments are permitted to use different affine maps but they must all have the form above. In particular different choices of λ are permitted, but the coefficient of x (i.e. λ) must always have absolute value greater than one and, as a consequence, the coefficient of y (i.e. λ^{-1}) will have absolute value less than one.

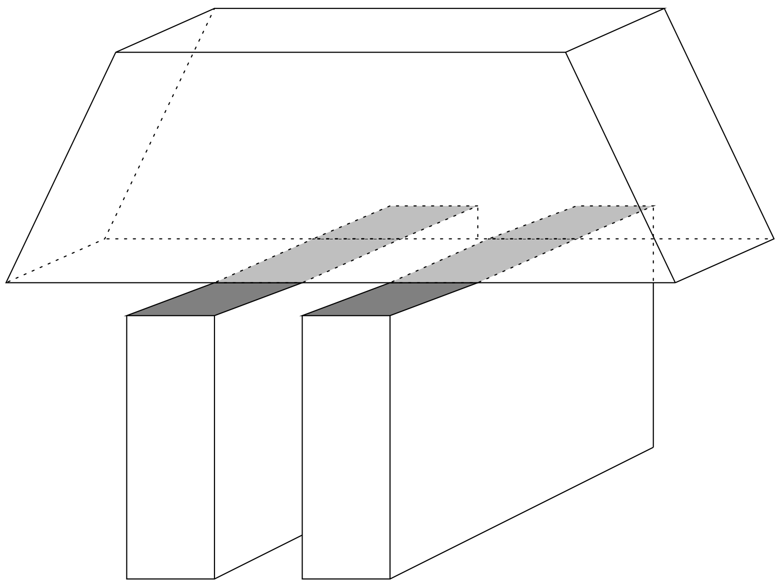

FIGURE 2. Attached boxes

In order to make interesting examples of flows on regions made up of boxes with these kinds of identifications we will need to make attachments that form a "cycle" of boxes, so it is possible for an orbit to leave a box B, pass through others, then return to B. Clearly this is necessary if we are to create a periodic orbit. It is not difficult to see that if we form a simple cycle of boxes like that in Figure 3 then there will be a single periodic orbit.

It is also not difficult to see that every other orbit in the boxes depicted in Figure 3 will eventually exit the boxes in one direction or the other.

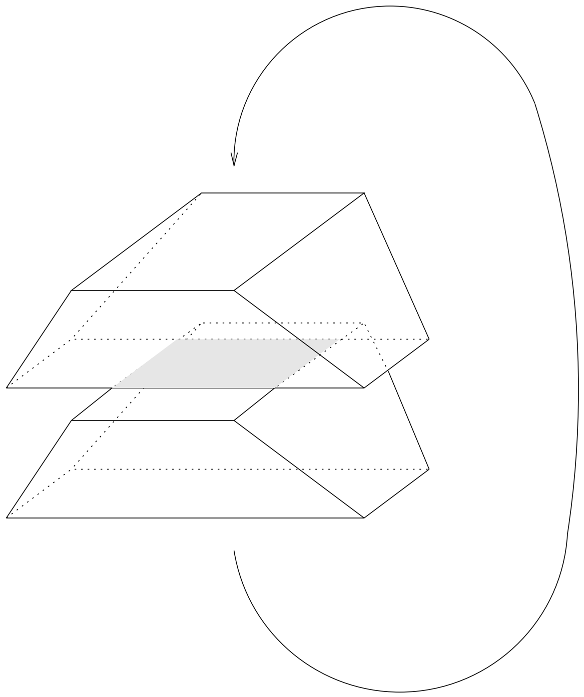

FIGURE 3. Two boxes, one closed orbit

The easiest way to demonstrate these facts is to consider the "return map" defined on (part of) one of the entering faces. Let C be one of these entering faces and let D_0 represent the subset of points in C which will return to C when the (positively oriented) orbit through them is followed, and denote by R_0 the points to which we return. See Figure 4. Then this return map $r : D_0 \to R_0$ is an affine map of the type described above for attachments. Indeed it is just the composition of the affine maps for the attachments through which an orbit passes.

From the form of this affine map it is easy to see that r has a unique fixed point, say (x_0, y_0). Moreover, the only points (x, y) for which $r^n(x, y)$ is defined for all $n > 0$ are those with $x = x_0$ while the only points for which $r^n(x, y)$ is defined for all $n < 0$ are those with $y = y_0$. Clearly the fixed point of r corresponds to a periodic orbit and any other periodic orbit (if there were any) would correspond to a periodic point of the map r.

We want now to investigate an analogous object but one which is more complicated than a simple cycle. Assume we have n boxes $B_1 \dots B_n$ with multiple attachments of the type described above. We will encode the box attachments in a matrix called the transition matrix A. We define it by setting A_{ij} equal to the number of attachments of the exiting face of B_i to

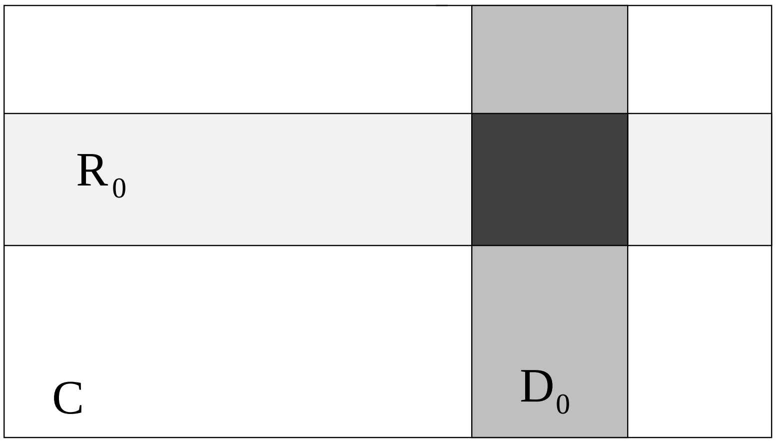

FIGURE 4

the entering face of B_j. The arrangement of these attachments is important, but for the moment we will ignore it. The boxes must be embedded in $\mathbf{R}^3$ in such a way as to realize the attachments and this requires that the embeddings be highly non-linear. In particular, it is perfectly legitimate for the exiting face of a box to be attached to the entering face of the same box.

Note that the transition matrix for our simple cycle in Figure 3 is

$$\mathbf{A} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

while the one for the attached boxes shown in Figure 5 is

$$\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$

Given a collection of boxes attached as described above with transition matrix A we want to investigate the collection of orbits which never exit the boxes in either direction. If we denote this set of points by Ω then it is easy to see that it is a compact set in the interior of the boxes which is invariant under the flow $\psi_t : \Omega \to \Omega$ whose orbits locally (in one box) are the arcs given by holding x and y fixed and varying t in the box co-ordinates. The orbits of this flow are parameterized so that $\psi_t(x, y, t_0) = (x, y, t_0 - t)$ whenever (x, y, t_0) and $(x, y, t_0 - t)$ are in the same box.

Definition 2.1. *The flow $\psi_t : \Omega \to \Omega$ described above will be called the* **boxed flow** *associated to the collection of embedded boxes.*

We will investigate this flow by considering a "return map" similar to the one described above. The name is perhaps not the best since typically points do not return under just one iteration of the map (unless the exiting face is attached to the entering face of the same box). The difference between this map and the return map discussed before is that now if there are n boxes in

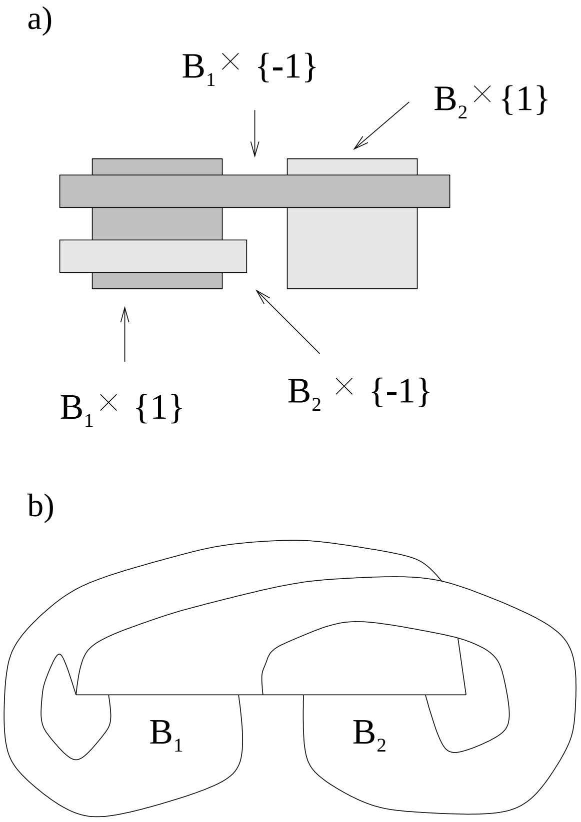

FIGURE 5. Another two-box flow, (a) the first return map, (b) side view of boxes.

our construct we will consider the return map r to be defined on a subset of the union of all entering faces and having values in a different subset of the union of all entering faces. More precisely, if x is in an entering face, say of box B_i, then $r(x)$ is defined provided the (positively oriented) orbit through x exits B_i in a region of attachment and immediately enters another box, say B_j. The value of $r(x)$ is the point in the entering face of B_j which is the first point of this box on the positive orbit segment starting at x.

The subset of an entering face on which r is defined will in general have several components each of which will be a rectangle entirely crossing that face in the direction parallel to the y-axis, like D_0 in Figure 4. The image of r in each entering face will in general consist of several components each of which will be a rectangle entirely crossing that face in the direction parallel to the x-axis, like R_0 in Figure 4. Since the transition matrix A was defined by letting A_{ij} equal the number of attachments of box B_i to B_j it is clear that A_{ij} of the components of the domain of r in the entering face of B_i will be mapped by r to A_{ij} of the components of the image of r in the entering face of B_j. We will denote these domain components by $D_{ij}(k)$ and the range components by $R_{ij}(k)$, where $1 \leq k \leq A_{ij}$. Note then that the restriction $r : D_{ij}(k) \to R_{ij}(k)$ is an affine map which expands the horizontal or x component and contracts the vertical or y component.

Let D denote the union of all the domain components, i.e. the full domain of r and let

$$\Lambda = \bigcap_{n=\infty}^{\infty} r^n(D).$$

Then $r : \Lambda \to \Lambda$ is a homeomorphism (assuming Λ is non-empty).

Each point in Λ has an associated forward and backward "itinerary." This is just the sequence of rectangles $D_{ij}(k)$ through which the trajectory of the point travels. More formally the the *forward itinerary* of a point $z \in \Lambda$ is the sequence $d_0, d_1, d_2, \ldots$ where d_n is the element of the collection $\{D_{ij}(k)\}$ which contains $r^n(z)$. The *backward itinerary* of a point $z \in \Lambda$ is defined similarly as the sequence $\ldots, d_{-2}, d_{-1}, d_0$ where d_n is the element of the collection $\{D_{ij}(k)\}$ which contains $r^n(z)$. Combining the two provides the *complete itinerary* $\ldots, d_{-2}, d_{-1}, d_0, d_1, d_2, \ldots$ of the point z.

Lemma 2.2. *Suppose $z = (x_0, y_0, 1) \in \Lambda$ is in the entering face of B_i. Then the set of $(x, y, 1) \in B_i$ with the same forward itinerary as z is*

$$\{(x_0, y, 1) \mid y \in [-1, 1]\},$$

i.e. it is the interval in the entering face of B_i which passes through z and is parallel to the y-axis. Likewise the set of $(x, y, 1) \in B_i$ with the same backward itinerary as z is $\{(x, y_0, 1) \mid x \in [-1, 1]\}$. Moreover, if two points in Λ have the same forward (backward) itinerary then the distance between the images of these points under Ψ_t tends to 0 as $t \to \infty$ ($t \to -\infty$). These line segments with the same forward (resp. backward) itinerary are called **local stable manifolds**, *(resp.* **local unstable manifolds***).*

Proof. The affine attaching maps described above contract each line segment parallel to the y-axis by a factor of λ^{-1} and expand each line segment parallel to the x-axis by a factor of λ. It follows that if $z_1 = (x_0, y, 1)$ the distance between $r^n(z)$ and $r^n(z_1)$ is $\lambda^{-n}|y - y_0|$. Hence $r^n(z_1)$ is defined for all $n \geq 0$ and z and z_1 have the same forward itinerary.

Conversely if $z_2 = (x, y, 1)$ is in the entering face of B_i and has the same forward itinerary as z then the distance between $r^n(z)$ and $r^n(z_2)$ is greater than $\lambda^n|x - x_0|$. Since this distance must be bounded independent of n we can conclude that $x = x_0$.

From the fact that $\|r^n(z) - r^n(z_1)\| \leq \lambda^{-n}|y - y_0|$ if follows that $\|\Psi(z) - \Psi(z_1)\|$ tends to 0 as $t \to \infty$.

A similar proof shows that points with backward itinerary equal to that of z are the interval specified and that the distance between such points tends to zero under Ψ_t as $t \to -\infty$. $\square$

We can also start with a possible itinerary and show that there are points which realize it. An *allowable sequence* $\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots$ of elements of the collection $\{D_{ij}(k)\}$ is one for which $r(d_n)$ intersects d_{n+1} for all $n \in \mathbf{Z}$. Forward and backward allowable sequences are defined similarly.

Clearly a necessary condition for a sequence to be a complete itinerary is that it be an allowable sequence, since if $r(d_n) \cap d_{n+1} = \emptyset$ there are no points whose trajectories enter one box through d_n and the next box through d_{n+1}. The next lemma asserts that this is also a sufficient condition.

Lemma 2.3. *Suppose that* $\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots$ *is an allowable sequence of elements of the collection* $\{D_{ij}(k)\}$ *with* $d_0 \subset B_i$. *Then the set of points in the entering face of* B_i *with forward itinerary* $d_0, d_1, d_2, \dots$ *is nonempty and consists of the line segment* $\{(x_0, y, 1) \mid y \in [-1,1]\}$, *for some fixed* x_0. *Similarly the set of points in the entering face of* B_i *with backward itinerary* $\dots, d_{-2}, d_{-1}, d_0$ *is a line segment of the form* $\{(x, y_0, 1) \mid x \in [-1,1]\}$ *for some* y_0. *And the set of points in the entering face of* B_i *with complete itinerary* $\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots$ *is the single point* $(x_0, y_0, 1)$.

Proof. By induction on n it is not difficult to see that

$$W_n = \bigcap_{i=0}^{n} r^{-i}(d_i) = \{z \mid r^i(z) \in d_i, 0 \le i \le n\}$$

is a rectangle in d_0 of the form $[a,b] \times [-1,1]$ whose width $b-a$ is $2\lambda^{-n-1}$. (The width of d_0 is $2\lambda^{-1}$.) It then follows that the set of points in the entering face of B_i with forward itinerary $d_0, d_1, d_2, \dots$ is $\bigcap_{n=0}^{\infty} W_n$ which is a line segment of the form $\{(x_0, y, 1) \mid y \in [-1,1]\}$. A similar argument shows that the set of points with backward itinerary $\dots, d_{-2}, d_{-1}, d_0$ is a line segment of the form $\{(x, y_0, 1) \mid x \in [-1,1]\}$ for some y_0. The set of points in in the entering face of B_i with complete itinerary $\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots$ must be the intersection of these two line segments, namely the point $(x_0, y_0, 1)$. $\square$

Clearly the itinerary sequences associated to a point in Λ give a great deal of information about the point. It has proven extremely valuable to abstract this concept and consider the "symbolic dynamics" associated with a transition matrix A which we describe in the next section.

It is interesting that (as we shall see) a complete topological description of Λ and the homeomorphism r depends only on the matrix A! As a consequence the topological type of the set of orbits which remain always in the boxes depends only on this matrix. But it is important to note that while A determines the abstract topological type of this space it says very little about the embedding of this space in $\mathbf{R}^3$ or S^3. It is really this embedding that is the subject of this article. And just as S^1 is a rather simple topological object with a very rich class of embeddings in S^3, it is also the case that as abstract topological objects the flows we have described above are well understood, but their embeddings in S^3 are far from understood.

The flows described in this section may seem to be of a rather special nature, and to some extent that is true. Nevertheless, the exhibit very typical behavior. In [12], for example, it is shown that for any smooth flow on $\mathbf{R}^3$ with a compact invariant set having positive topological entropy there is an invariant subset on which the flow is qualitatively equivalent to a boxed flow in a sense we define in the next section (topological equivalence). A definition of topological entropy can be found in [20]. It is a numerical invariant of the topological complexity of a flow.

3. Abstract Symbolic Dynamics

In this section we give an abstract "symbolic" description of a class of flows and then show that any boxed flow is topologically equivalent to such a flow. Of course, we must first specify our notion of equivalence.

Definition 3.1. *Two flows $\phi_t : X \to X$ and $\psi_t : Y \to Y$ are called* **topologically equivalent** *provided there is a homeomorphism $H : X \to Y$ carrying orbits of ϕ_t onto orbits of ψ_t. The flows are called* **topologically conjugate** *provided for every t and every $x \in X$ the equation $H(\phi_t(x)) = \psi_t(H(x))$ is satisfied.*

Both of these are equivalence relations. From a topological point of view topological equivalence is the more appropriate notion and we will focus largely on it. Topological conjugacy is generally too rigid for our purposes. For example, the numerical value of the period of each periodic orbit is easily seen to be an invariant of topological conjugacy, from which it follows that there are uncountably many different equivalence classes.

Let A be an $n \times n$ matrix of non-negative integers. We construct a finite graph Γ with oriented edges as follows. Γ has n vertices numbered 1 to n and Γ has A_{ij} oriented edges running from vertex i to vertex j. Each edge is assumed to have unit length. We will denote these oriented edges by $e_{ij}(k), 1 \leq k \leq A_{ij}$ and denote by E the set of all edges.

We want to consider the set of bi-infinite sequences of edges and give it a topology. If the finite set E is given the discrete topology and we give the product topology to

$$\prod_{i=-\infty}^{\infty} E = E^{\mathbf{Z}}$$

then this space is compact, totally disconnected and perfect (assuming E has more than one element) and hence can be shown to be homeomorphic to a Cantor set.

We are interested in a subset of this space of bi-infinite sequences of edges. We will denote by Σ_A the subset of elements $\{a_i\}_{i\in\mathbf{Z}}$ with the property that for each i the edge a_i of Γ ends in the vertex where a_{i+1} begins. In other words $a_i = e_{pq}(k)$ and $a_{i+1} = e_{qr}(k')$ for some choices of p, q, r, k and k'. Thus, subset Σ_A consists of precisely those sequences of oriented edges which could be traced out by an infinite continuous path on the graph which always respected the orientation.

The fact that E is a set of edges of a graph is only of heuristic importance and we could equally well consider Σ_A as bi-infinite sequences of "symbols" satisfying a finite collection of combinatorial rules about which symbols are allowed to follow which other symbols.

It is a straightforward exercise to show that Σ_A is a closed (and hence compact) subset of $E^{\mathbf{Z}}$.

Introducing some dynamics, we define the map

$$\sigma : \Sigma_A \to \Sigma_A$$

by $\sigma(\{a_i\}) = \{b_i\}$ where $b_i = a_{i-1}$. The map σ is called the *subshift of finite type* based on the matrix A. It is a shift because it simply shifts the bi-infinite sequences of symbols one place to the left. If any symbol were allowed to follow any other (i.e. If the matrix A is a nonzero 1×1 matrix and hence any symbol is allowed to follow any other symbol) then σ is called a *full* shift. The "finite type" part of the name refers to the fact that there are finitely many rules (encoded in the matrix A) about which symbols may follow which other symbols. It is straightforward to see that σ is continuous and we can immediately exhibit its inverse (shifting to the right), so it is a homeomorphism.

There is an alternate description of Σ_A which will be useful for us. Let Γ be the graph associated to A as described above with path metric assigning unit length to each edge. Let Σ be the space of maps from $\mathbf{R}$ to Γ which preserve arc length and take integers in $\mathbf{R}$ to vertices of Γ. We give this space the compact-open topology.

It is easy to see that the map from Σ to Σ_A obtained by associating to each path of Σ the bi-infinite sequence of edges through which it passes, is a homeomorphism. More precisely if $\alpha(t)$ is an element of Σ we define $h(\alpha(t))$ to be the sequence $\{a_i\}_{i \in \mathbf{Z}}$ with the property that for each i the edge of Γ containing $\alpha(i + 1/2)$ is a_i. Then h is a homeomorphism. Also it is clear that the shift map on Σ is given by $\sigma(\alpha)(t) = \alpha(t+1)$, or more precisely h is a topological conjugacy from $\sigma : \Sigma_A \to \Sigma_A$ to $\sigma : \Sigma \to \Sigma$.

The most thoroughly studied subshifts of finite type are those which are *irreducible*. A subshift of finite type is called *irreducible* if the graph Γ described above has the property that given any two vertices on it there is an oriented path joining them. In terms of the matrix A this is equivalent to asserting that there are no i and j such that the ij^{th} entry of A^n is zero for all $n > 0$.

A great deal is known about irreducible subshifts of finite type. See [18], for example. We mention only that any irreducible subshift of finite type has a dense orbit and has a dense set of periodic points.

Associated to a subshift of finite type $\sigma : \Sigma_A \to \Sigma_A$ we can construct an abstract flow. The construction we use is called the "mapping torus" by topologists, but unfortunately is called the "suspension" by dynamicists.

Let X_A be the quotient space of $\Sigma_A \times \mathbf{R}$ under the identification $(\sigma(x), s) \sim (x, s+1)$. Equivalently, let X_A be the quotient space of $\Sigma_A \times [0,1]$ under the identification $(x,1) \sim (\sigma(x),0)$. The first of these descriptions is better for describing the flow ϕ_t on X, because this flow is simply the quotient of the flow Φ_t defined on $\Sigma_A \times \mathbf{R}$ by $\Phi_t(x,s) = (x, s+t)$.

In terms of the description of the subshift of finite type as the space of maps from $\mathbf{R}$ to the graph Γ associated to the matrix A, this flow also has a nice description. Recall that the subshift map σ is defined on the space Σ of arc-length preserving paths mapping $\mathbf{R}$ to Γ which take integers to vertices by $\sigma(\alpha)(s) = \alpha(s+1)$. If we let Λ denote the space of arc-length preserving paths mapping $\mathbf{R}$ to Γ which are consistent with the orientations of the edges of Γ and give Λ the compact-open topology, then the mapping torus flow defined above is topologically conjugate to the flow on Λ given by $\Phi_t(\alpha)(s) = \alpha(s+t)$.

Because of the analogous properties for the subshift of finite type σ, whenever the matrix A is irreducible the flow ϕ_t on X_A has a dense orbit and a dense set of periodic orbits.

Proposition 3.2. *A boxed flow ψ_t with transition matrix A is topologically equivalent to the mapping torus flow of the subshift of finite type $\sigma : \Sigma_A \to \Sigma_A$.*

Proof. In fact something stronger is true. The boxed flow ψ_t with transition matrix A as described above and the mapping torus flow ϕ_t on X_A constructed from the subshift of finite type $\sigma : \Sigma_A \to \Sigma_A$, are topologically conjugate.

Recall from §2 that the ij entry of the transition matrix A equals the number of components of the attaching region of the exiting face of box B_i with the entering face of box B_j. Thus if we use the matrix A to construct a graph Γ and the corresponding subshift of finite type $\sigma : \Sigma_A \to \Sigma_A$, then the vertices of Γ are in one-to-one correspondence with the boxes $\{B_i\}$ and the edges from vertex i to vertex j are in one-to-one correspondence with the components of the attaching region of the exiting face of box B_i with the entering face of box B_j.

Then if $\Psi_t : \Omega \to \Omega$ is the boxed flow, this gives us a map $h : \Omega \to \Lambda$ which is a conjugacy from the boxed flow to the mapping torus flow for the subshift of finite type. It is defined by first considering the map $j : \Omega \to \Gamma$ which assigns to the point $(x,y,s) \in B_i$ the point on the edge of Γ corresponding to the region where the orbit of (x,y,s) exits B_i and whose distance from the beginning of that edge is $(1-s)/2$. That is, its position on that edge proportional to its distance from the entering region in B_i.

Then the conjugacy h is defined using the compact-open path space description of the mapping torus flow. To a point $z \in \Omega$ we first consider its orbit $\{\Psi_t(z)\}$ in Ω and then obtain a path in Γ. This path is arc-length preserving and hence an element of Λ the space of arc-length preserving infinite paths (which is given the compact-open topology). More precisely, $h(z) = \alpha(t)$ where $\alpha(t)$ is the path $j(\Psi_t(z))$ in Γ. It is immediate that

$$h(\Psi_s(z)) = \alpha(t+s) = \Phi_s(\alpha) = \Phi_s(h(z))$$

so h is a conjugacy.

It is straightforward to show that h is continuous. It is only necessary to check that h is one-to-one and onto, i.e. invertible, since Ω and Λ are both compact. Given an element $\alpha(t) \in \Lambda$ we wish to find $h^{-1}(\alpha)$. We can obtain the point $\alpha(0) \in \Gamma$, and from the edge of Γ we know the B_i which must contain $h^{-1}(\alpha)$. Also, if s is the distance of $\alpha(0)$ from the start of its edge, then $h^{-1}(\alpha) = (x_0, y_0, 1-2s)$ for some x_0 and y_0 in the co-ordinates of the box B_i. The fact that such an x_0 and y_0 exist and are unique can be shown as follows.

The fact that the edges of G are in one-to-one correspondence with the attaching regions of the exiting and entering faces of the boxes and the fact that each such attaching region is a subset of one of the domain components $\{D_{ij}(k)\}$ means that to the path $\alpha(t)$ we can associate an allowable sequence $\dots, d_{-2}, d_{-1}, d_0, d_1, d_2, \dots$, where each d_i is one of the domain components $\{D_{ij}(k)\}$ and they occur in the order in which the path $\alpha(t)$ crosses the corresponding edges of Γ.

According to Lemma 2.3 there is a unique point $z_0 = (x_0, y_0, 1)$ in the entering face of B_i with this itinerary. By construction $h(z_0)(t)$ is a path in G starting at the beginning of the edge containing $\alpha(0)$ and following the same sequence of edges as α. Then $h(z_0)(s) = \alpha(0)$ so $h(z_0)(t) = \alpha(t-s)$. From this it is clear that $h(\Psi_s(z)) = \alpha(t)$ and that $\Psi_s(z)$ is the unique point with this property. □

A corollary of this result and its proof is the following.

Proposition 3.3. *Orbits of a boxed flow ψ_t with transition matrix A are in one-to-one correspondence with orbits of the subshift of finite type $\sigma : \Sigma_A \to \Sigma_A$. In particular, periodic orbits of ψ_t correspond to periodic orbits of σ (or, equivalently, to periodic allowable sequences of symbols associated with Σ_A).*

Rather surprisingly, if we are only interested in classifying boxed flows up to topological equivalence there is a complete and easily computed answer given by the following theorem. We emphasize that this is only for classification up to topological equivalence; the analogous question for topological conjugacy is much more subtle (see [28] or [18]).

Theorem 3.4. *(Franks [11]) Suppose A and B are square, non-negative, integer matrices which are irreducible and Ψ and Φ are the mapping torus flows of the corresponding subshifts of finite type. Then necessary and sufficient conditions that these flows be topologically equivalent are that*

$$det(I - A) = det(I - B)$$

and

$$\mathbf{Z}^n/(I-A)\mathbf{Z}^n \cong \mathbf{Z}^m/(I-B)\mathbf{Z}^m,$$

where n and m are the size of A and B respectively.

Despite the fact that as abstract topological flows we can easily classify boxed flows up to topological equivalence, many problems remain if we try to understand how they are situated in three dimensional space. It is the pursuit of some understanding of these embeddings that the next section considers.

4. Templates

In the previous sections we described "boxed flows," a class of flows defined on subsets of $\mathbf{R}^3$ (but which could be easily extended to flows on all of $\mathbf{R}^3$). In this section we want to begin the study of how these subsets are embedded. In particular any closed orbit of a flow on $\mathbf{R}^3$ is a knot and any finite set of closed orbits is a link. One of our long range objectives is the study of the relationship between dynamics and the knot types of periodic orbits.

It turns out that the union of the boxes in a boxed flow is not as conceptually simple for describing a flow as another construct called a "template" which we now describe.

Definition 4.1. *The* **template** *associated with a boxed flow Φ_s corresponding to the union of boxes $X = \cup B_i$ is the quotient space L of X formed by collapsing each of the boxes B_i making up X to a rectangle by identifying any two points in B_i of the form (x, y_1, t) with (x, y_2, t). There is a well defined* **semi-flow** *ϕ_s induced by Φ_s because any two points which are identified are carried by Φ_s to two points which are identified for all the values of s for which $\Phi_s(x, y_i, t)$ is defined.*

Note that two points in X are identified in forming L if and only if they are in the same box B_i and they have the same forward itinerary. This is because we saw in Lemma 2.2 that the line segments we are collapsing to points are precisely the local stable manifolds consisting of points with the same forward itinerary.

The semi-flow ϕ_s induced by Φ_s is only "semi" because it cannot in general be inverted, i.e. it is a semi-group action of $\mathbf{R}$ rather than an action of $\mathbf{R}$ on the compact set of points on which ϕ_s is defined for all positive s.

Definition 4.2. *The* **dual template** *associated with a boxed flow Φ_s corresponding to the union of boxes $X = \cup B_i$ is the template of the inverse flow. Equivalently, it is the quotient space L of X formed by collapsing each of the boxes B_i making up X to a rectangle by identifying any two points in B_i of the form (x_1, y, t) with (x_2, y, t).*

It is important to note that a boxed flow, by definition, includes an embedding of its boxes in $\mathbf{R}^3$ or S^3. Thus if L is the template (or dual template) associated with a boxed flow then one can embed L in $\mathbf{R}^3$ or S^3 in such a way that each point z of L lies in the interval in $\cup B_i$ which is collapsed to form z. Any two embeddings of L with this property are isotopic and any of them will be called an embedding of L associated to the boxed flow.

The following result is a special case of a result of Birman and Williams [2]. It tells us that understanding the closed orbits of a boxed flow in $\mathbf{R}^3$ is equivalent to understanding the closed orbits of the associated embedded template semi-flow.

Theorem 4.3. *There is a one-to-one correspondence between the periodic orbits of a boxed flow and the periodic orbits of the induced semi-flow on the associated embedded template. Moreover, closed orbits paired by this correspondence are isotopic as embedded circles and finite sets of closed orbits matched by this correspondence are isotopic as links (embedded finite disjoint unions of circles).*

Proof. If $h : X \to L$ is the quotient map defining L then h is a semi-conjugacy. That is, $h \circ \Phi_t = \phi_t \circ h$ for all $t \geq 0$. It follows that if γ is a closed orbit of Φ_t then $h(\gamma)$ is a closed orbit of ϕ_t.

Conversely, if γ_0 is a closed orbit of ϕ_t then $h^{-1}(\gamma_0)$ is a bundle over γ_0 whose fiber is an interval. This bundle (either an annulus or a Möbius strip) is invariant under the flow Φ_t and this flow preserves and contracts the fibers. If t_0 is the period of γ then Φ_{t_0} is a contraction map of each fiber to itself and hence has a unique fixed point. The collection of these fixed points is a (unique) closed orbit γ of period t_0 for Φ which is mapped by h to γ_0.

If L has an embedding associated with the flow, then for any point z of γ_0 there is precisely one point $\alpha(z)$ in $h^{-1}(z) \cap \gamma$. Sliding the point $\alpha(z)$ along the interval $h^{-1}(z)$ to z, for all points $z \in \gamma_0$ simultaneously, defines an isotopy from γ to γ_0. □

5. Knots and Links

A *knot* k is an embedding of S^1 into S^3 (or $\mathbf{R}^3$), $k : S^1 \to S^3$. We are only interested in smooth embeddings. A knot may be given an orientation or preferred direction. We will always use a flow to induce an orientation on our knots. It is a common abuse of notation to use the same symbol for a knot k and its image $k(S^1) \subset S^3$. A *link* of n components is an embedding of n disjoint copies of S^1.

Two knots k_1 and k_2 (or two links) are *equivalent* if there is an isotopy of S^3 that takes k_1 to k_2. When we talk about a knot we almost always mean its equivalence class, or *knot type*. Detecting knot equivalence is the primary goal of knot theory.

In order to publish papers about knots, knot theorists have developed the *knot diagram*. This is just a projection of a knot or link into a plane such that any crossings are transverse. The crossings are then labeled as positive or negative. See Figure 6 for the convention used here. If a knot has a diagram with no crossings then it is called an *unknot* or less formally a trivial knot.

Suppose we have a knot diagram for k and we have parameterized the planar curve. If, using polar coordinates, $d\theta/dt > 0$ for all t, then we

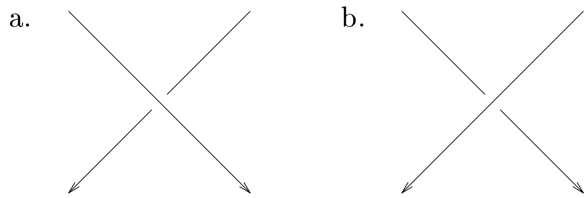

FIGURE 6. (a) a positive crossing, (b) a negative crossing

say that the diagram represents a *braiding* of k or that k is in *braid form.* Figure 7 shows two knot diagrams for the *figure-8 knot*, only one of which is braided. The reader is encouraged to demonstrate the equivalence. It is well known that any smooth knot or link is equivalent to a braid [6, Chapter 10]. In fact templates themselves can be braided, meaning that all the closed orbits are braided simultaneously [12]. If a knot k has a braid presentation in which all the crossings are positive then k is called a *positive braid.* Note: there exist knots which are not positive braids, but that can be presented by diagrams with only positive crossings [27].

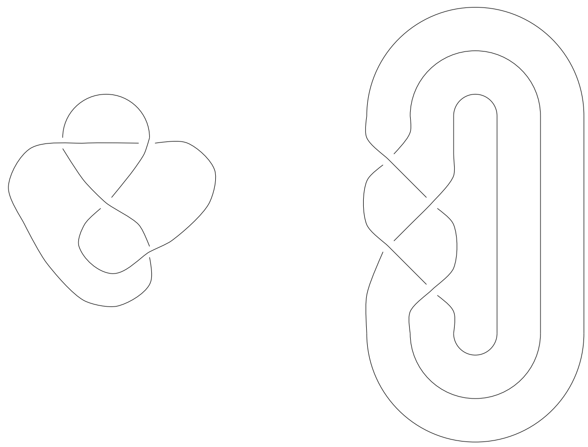

FIGURE 7. The figure-8 knot.

Though a practical algorithm for detecting knot equivalence is not known, there are many useful invariants. One of special interest to us is the *genus* of a knot or link. Every link forms the boundary of an embedded orientable surface, called a Seifert surface [6]. Abstractly we can attach a disk along each boundary component of a Seifert surface, that is along each component of the link, producing a closed surface. The genus of the link is then defined to be the minimum genus over all such closed surfaces.

If a knot is in braid form there are formulas which allow us to compute, or at least estimate, the genus.

Proposition 5.1. *(Bennequin [1]) Let k be a knot with genus g. Suppose k has a braid presentation on s strands with c_+ positive crossings and c_- negative crossings. Then $|c_+ - c_-| - s + 1 \leq 2g \leq |c_+ + c_-| - s + 1$.*

A similar result holds for links. If the braid presentation of k is positive, then we have $2g = c - s + 1$, $(c = c_+)$ a fact that was proved independently in [3].

Given a two-component link $k_1 \cup k_2$, the *linking number* of k_1 with k_2 is the sum of the signs of each crossing of k_1 under k_2 and is denoted $lk(k_1, k_2)$. The linking number is a link invariant. If the components of a link l can be separated by a 2-sphere that misses the link, then we say that l is a *split link*. For a two-component link $l = k_1 \cup k_2$, being a split link implies $lk(k_1, k_2) = 0$, though the converse is false.

The last item from knot theory we review is the notion of primeness. A knot $k \subset S^3$ is *composite* if there exists a smooth 2-sphere S^2 such that $S^2 \cap k$ is just two points p and q, and if γ is any arc on S^2 joining p to q then the knots

$$k_1 = \gamma \cup (k \cap \text{ outside of } S^2) \text{ and}$$

$$k_2 = \gamma \cup (k \cap \text{ inside of } S^2),$$

are each nontrivial, (i.e. not the unknot). We call k_1 and k_2 *factors* of k and write

$$k = k_1 \# k_2.$$

We call k the *connected sum* of k_1 and k_2. If a nontrivial knot is not composite, then it is *prime.*

Figure 8 gives an example. It shows how to factor the *square knot* into two trefoils. Trefoils are prime. It was shown by Schubert [6, Chapter 5] that any knot can be factored uniquely into primes, up to order. Note: the unknot serves as a unit.

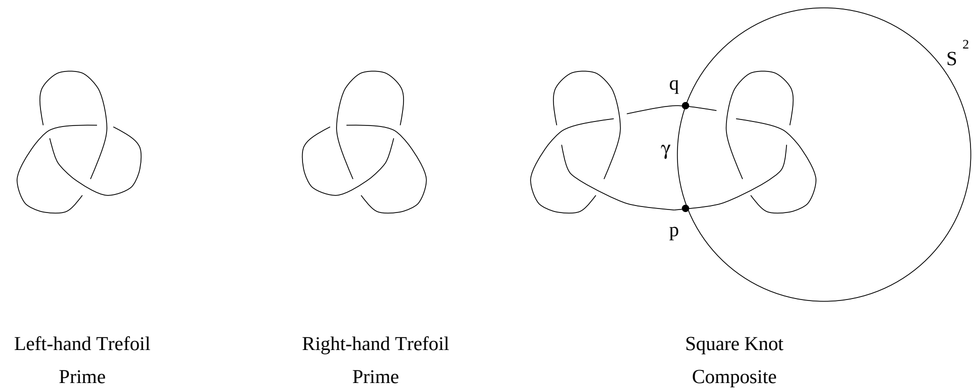

FIGURE 8. The square knot is the sum of two trefoils.

6. The Lorenz Template

Perhaps the simplest example of a template is the *Lorenz Template*. It is shown in Figure 9. It was developed by Williams as a naive model of the strange attractor apparently associated with the Lorenz equations, a 3×3 ODE used to study turbulent flows. See [2] and the references cited there. Although it remains unknown whether the Lorenz equations define a "Lorenz type" attractor, the existence of such attractors has been confirmed in various families of differential equations. For a brief history of this work see [20, §7.11.2] and the recent paper [8].

It is worth pointing out that boxed flows are saddle sets and not attractors. A boxed flow that would be modeled by the Lorenz template would have transition matrix

$$\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

The reader may wish to consider how one might embed two boxes in $\mathbf{R}^3$ with ends attached as prescribed by this matrix in such a way that the associated template is essentially the one shown in Figure 9. (In this figure the corners on the edges of the template have been trimmed a bit for simplicity.)

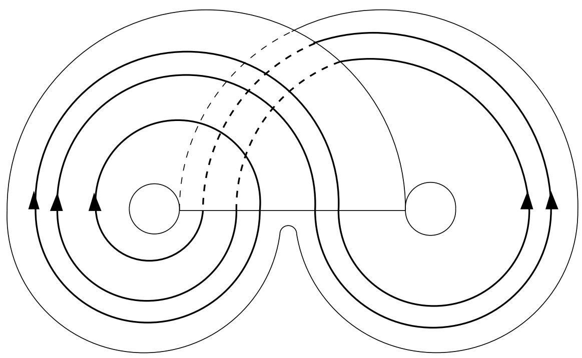

Figure 9. The Lorenz template with a trefoil orbit shown.

Here we shall study a little of what is known about the class of knots realized as periodic orbits of the Lorenz template's semi-flow. An important symbolic tool is the use of *words* to describe orbits. We label the right band x and the left band y. Then given a staring point on the branch line, we can write down a sequence of x's and y's to describe the path of the forward orbit of our point. If the orbit is periodic we shall just record the number of symbols needed to describe one circuit. Thus the orbit shown in Figure 9 is $xxyxy$ (and turns out to be a trefoil knot). Of course, any cyclic permutation of $xxyxy$ would give the same path.

The key fact is, there is a one-to-one correspondence between the equivalence classes of finite words in two symbols under cyclic permutations and the set of periodic orbits of the Lorenz template. Similar correspondences can be set up for any template and an appropriate symbol space. This follows from the theory of subshifts of finite type and, in particular, Proposition 3.3.

Knots which occur as periodic orbits in the Lorenz template are called *Lorenz knots.*

Proposition 6.1. *(1) All Lorenz knots are positive braids with full twists, (2) all torus knots (defined below) are Lorenz knots, (3) all Lorenz knots are prime, and (4) the only split Lorenz links have either the x or y loops as a component.*

We shall only prove 1, 2, and 4. Although Williams proved 3 directly in [29] it is now known that 1 implies 3; that is positive braids with a full twist are prime knots, as Williams himself had conjectured [7, 26]. It is generally believed that there are examples of positive braids with full twists that are not Lorenz knots, but this has not to our knowledge been written up.

The proof of Proposition 6.1 (and several others) uses what is known as the *belt trick.* Consider the strip with a "loop-de-loop" as show in Figure 10. As we pull the ends apart the "loop-de-loop" turns into a full twist.

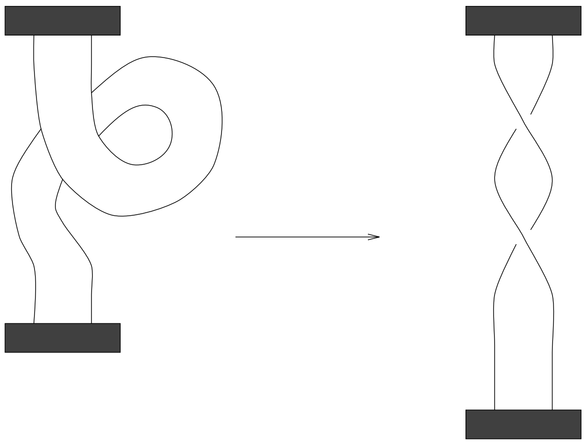

FIGURE 10. The belt trick.

Proof of 1. The full twist is revealed by doing *surgery* on the template. We delete the two segments of flow lines that meet at the center of the branch line and start at the branch line. Then deform the resulting new template. But this new template has just the same periodic orbits as does the Lorenz template. The surgery and deformation are shown in Figure 11. The deformation uses the belt and in the last step a move we shall call the *lamp shade trick.* It insures that $d\theta/dt > 0$ for all orbits except the x and y orbits (which are horizontal in Figure 11. But the result holds for these two orbits trivially.

Positivity is clear as only positive crossings can be realized. □

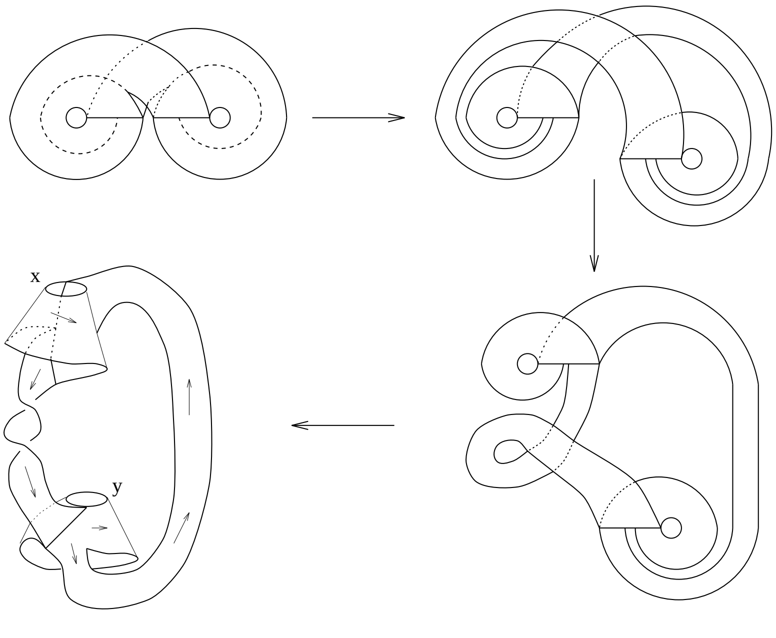

FIGURE 11. Surgery on a template

A *torus knot* is a knot that can be represented by a closed curve on a standardly embedded torus. A torus is standardly embedded if it is the boundary of an unknotted solid torus. All torus knots can be represented by a pair of relatively prime integers (p, q). The number p gives the number of times the knot wraps around the long way (longitudinally), while q gives the number of time the knot wraps around the torus the short way (meridionally).

A torus knot can always be presented so as to have all positive or all negative crossings. We are only interested in positive torus knots. It is not hard to check that $(p, q) = (q, p)$, but otherwise the (p, q) representation is one-to-one for positive torus knots and $p, q > 0$. We shall assume $p < q$.

Proof of 2. In Figure 12 we show how to place a (p, q) torus knot on a torus. We do not need the whole torus; the knot can be drawn in the split strip that wraps about the torus in Figure 12. The main branch of the strip has p strands of the (p, q) knot. This part of the strip makes $n > 0$ full twists around the torus, though the figure shows only one full twist. Then the strip splits in two with a strands staying of the side of the torus facing the reader and b strands making an extra full twist around the torus. Then the two strips come back together. One sees that $q = np + b$, and that there is a single knot as long as p and q are relatively prime.

Figure 13 shows how to place the split strip onto the Lorenz template so that any torus knot can be realized as a closed orbit of the template's semi-flow. Here we have used the belt trick $n + 1$ times. Note that in this figure too, only the $n = 1$ case is shown, but the reader should be able to see how to add on extra loops of the strip about the x band of the template. $\square$

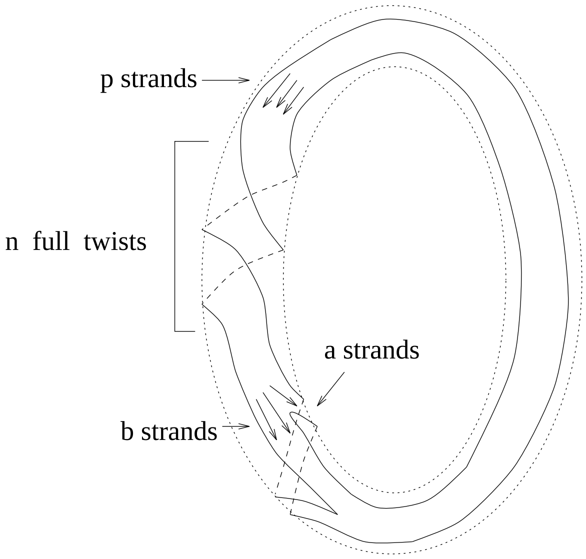

FIGURE 12. Torus knots

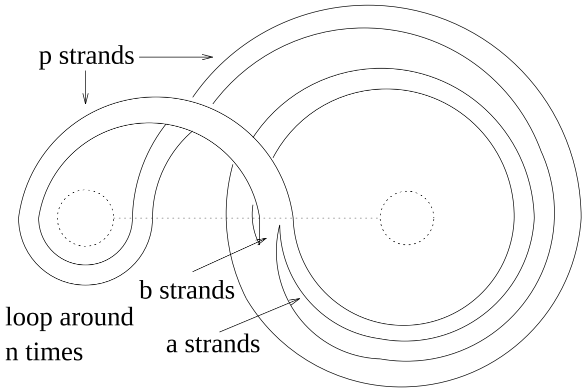

FIGURE 13. Torus knots are Lorenz

Proof of 4. Any two orbits whose words involve both x's and y's must cross, and all the crossings are positive. Thus, the linking number is not zero. $\square$

We conclude this section by stating one of the first general results that template theory has given to the study of flows and differential equations.

Theorem 6.2 (Franks & Williams). *A smooth flow in* $\mathbf{R}^3$ *or* S^3 *that has a compact invariant set and positive topological entropy has infinitely many distinct knot types among its closed orbits.*

The *invariant* set of a flow ϕ_t on a manifold M is $\bigcap_{-\infty<t<\infty} \phi_t(M)$. We shall not define topological entropy, but only state that it a standard device for measuring "mixing". The gist of the proof is to show that any such flow must have a part of its invariant set that can be modeled by an embedded Lorenz template. Proposition 5.1 is used to show that there are closed orbits of arbitrarily high genus and thus infinity many distinct knot types can be realized as closed orbits.

7. Lorenz-like Templates

By *Lorenz-like templates* we mean one of the templates depicted in Figure 14, where there are m half twists in the x branch and n half twists in the y branch. We shall denote such a template by $L(m,n)$. Thus $L(0,0)$ is the Lorenz template. Notice that $L(m,n) = L(n,m)$ via a 180^o rotation. We shall at times abuse our own notation and use the symbol $L(m,n)$ to represent the set of knots and links realizable by the closed orbits in the semi-flow on $L(m,n)$.

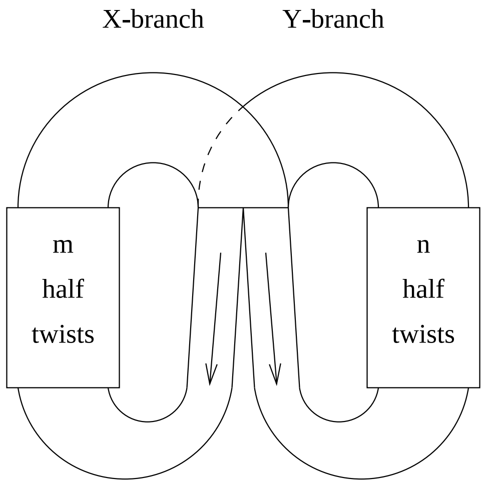

FIGURE 14. Lorenz-like Templates.

Proposition 7.1 (Williams). *Knots in $L(0,n)$ for $n \geq 0$ are positive braids with a full twist and hence are prime.*

Proof. For $n = 0$ we already have this result. For $n \geq 2$ Figure 15 shows how to manipulate a template $L(0,n)$ to see the full twist. This is actually a simple application of the lamp shade trick used in the proof of Proposition 6.1 (1). Positivity is clear. In [17] it shown that knots on $L(0,1)$ do indeed have a full twist presentation, but there does not appear to be a presentation of the template where all knots are simultaneously presented as positive braids with full twists. Also, the result in [26] shows that having a half twist in a positive braid is enough for primeness. ∎

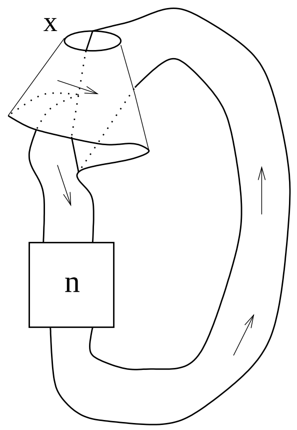

FIGURE 15. At least n half twists.

Proposition 7.2. *There is a lower bound on the genus of nontrivial knots in $L(0,n)$ of $g \geq (n-1)/2$, for $n \geq 0$.*

Proof. Referring again to Figure 15, suppose we have a knot with s strands. Notice that s is also just the number of y strands. If there is only one y strand, then we have an unknot. Each half twist forces $s(s-1)/2$ crossings. Suppose $s > 1$. Thus, since all the crossings are positive, Proposition 5.1 gives $g \geq (n-1)/2$. □

This shows, for example, that there can be no trefoils on $L(0,n)$ for $n \geq 3$, since the trefoil is a genus one knot.

Proposition 7.3. *As sets of knots, $L(0,n) \subset L(0,n-2)$, for all integers n.*

Proof. The proof is pictorial. Figure 16 shows how to place $L(0,n)$ into $L(0,n-2)$. Again we make use of the belt trick, paying careful attention to the sign of the new full twist. □

Proposition 7.4. *For $n < 0$, $L(0,n)$ contains composite knots.*

Proof. Figure 17 shows a composite knot in $L(0,-1)$, found by Williams [29]. Its word is $xxyyyyxxy$ and, as the reader can check, it is a connected sum of positive and negative trefoils. Figure 18 is of a composite in $L(0,-2)$ found in [22]. The knot type is the same as before. Its word is $xxxyxxxyyxyxxyyxy$. These examples and Proposition 7.3 give the result. □

Proposition 7.5. *As sets of knots $L(0,\pm 4)$ is a subset of $L(0,\pm 1)$ (respectively).*

Proof. Figure 19 gives the proof for the minus case. The surgery shown starts by cutting along the y orbit. This changes the invariant part of the

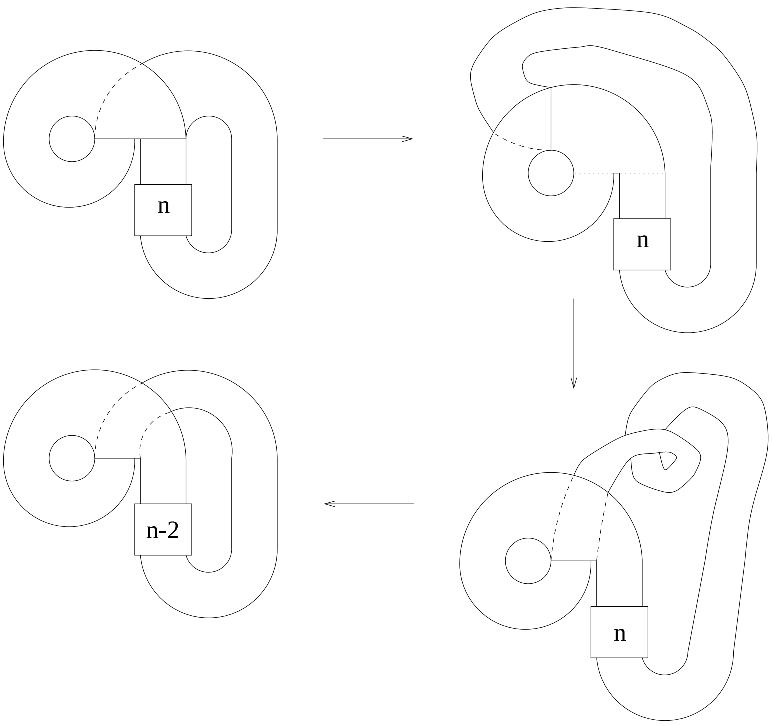

FIGURE 16. $L(0, n)$ is in $L(0, n-2)$.

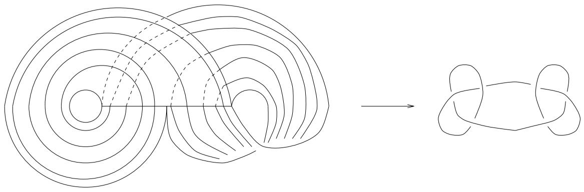

FIGURE 17. A composite knot in $L(0, -1)$.

semi-flow, but the knot type of the new y orbit is still that of the unknot. The linking number of the new y orbit with another closed orbit will be twice the linking number of the original y with the other orbit. So, the links have changed, but only a little. Any link in $L(0, -4)$ that does not use the y orbit is isotopic to a link in $L(0, -1)$. The plus case is similar.

The proof goes back to when one of us, as a seventh grader, was shown that when you cut a Möbius band down the middle you get a strip with four half twists in it. □

The following dichotomy has emerged. In the case $n \geq 0$ the templates are *positive*, that is all closed orbits are positive braids, while if $n < 0$ the templates are *mixed*, that is they have knots which are not positive braids.

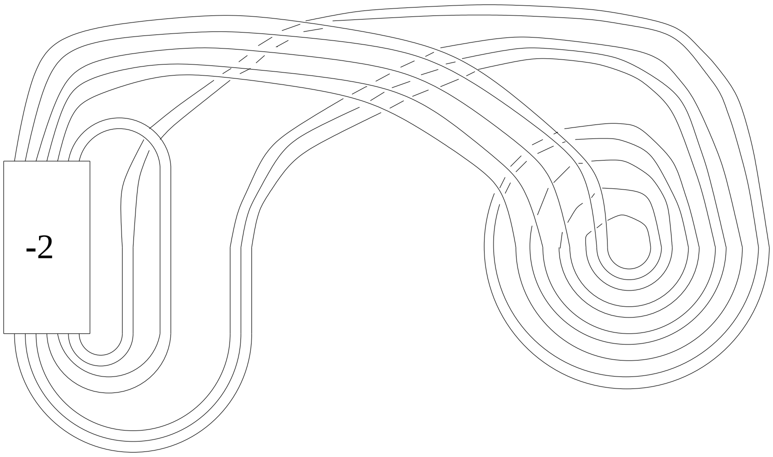

FIGURE 18. A composite knot in $L(-2, 0)$.

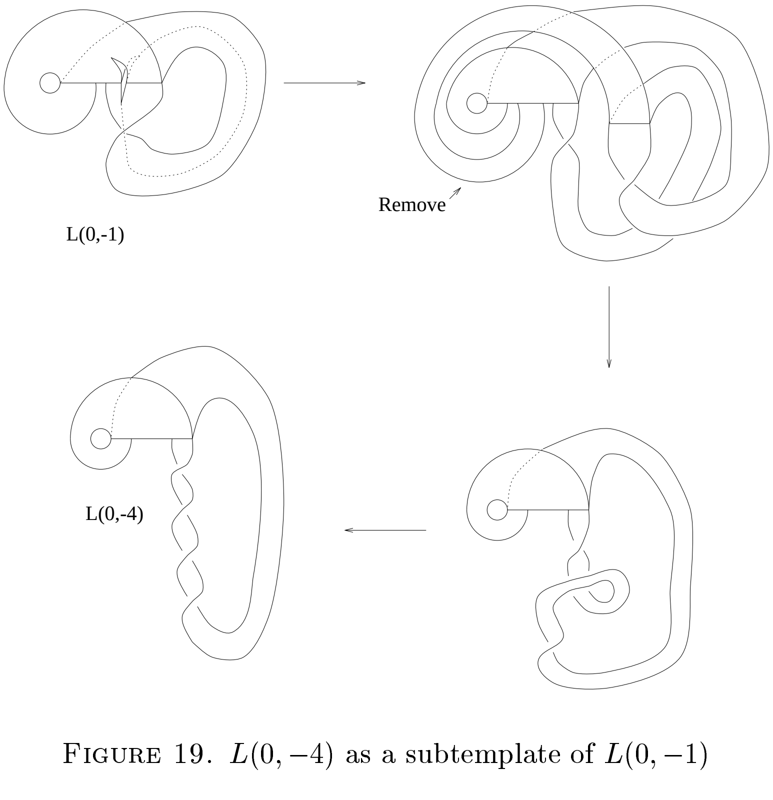

FIGURE 19. $L(0, -4)$ as a subtemplate of $L(0, -1)$

In the former case all the knots are prime, while in the latter case there are composite knots. In the next section we will see that this dichotomy is even stronger.

8. Universal Templates

A *universal template* is one which contains all knots and links. It was originally conjectured in [3] that such objects could not exist. However, Ghrist has recently shown that they do [13]. Here we will only outline the major steps needed to show this. More recently Ghrist has found a template that contains all other templates as subtemplates.

Step 1. Call the template shown in Figure 20 U. Define U_n to be the n fold cover of U shown in Figure 21. Ghrist shows that for all $n > 0$, U_n is a subtemplate of U. It is worth noting that his method does not involve visually constructing the isotopy as above. Instead Ghrist developed a symbolic formalism for orbit-to-orbit template embeddings to show that each U_n lives in U.

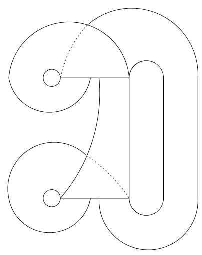

FIGURE 20. The template U

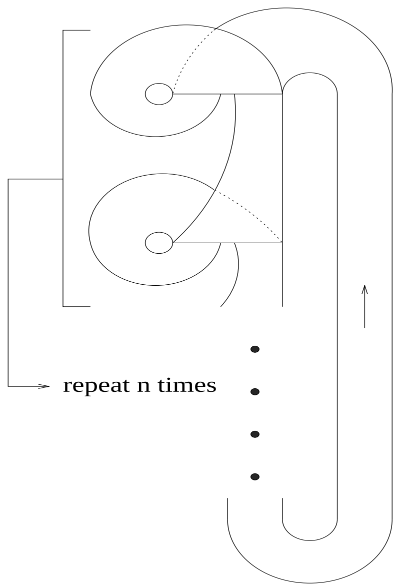

FIGURE 21. The templates U_n

Step 2. For every braid b there in an n such that $b \in U_n$. In Figure 22 we show how a negative crossing between the second and third strand can be realized on U_n. The positive crossing case is similar. By induction one can show that any braid with a single crossing is realized by a closed orbit on some U_n. We can thus construct b on a U_n by concatenation.

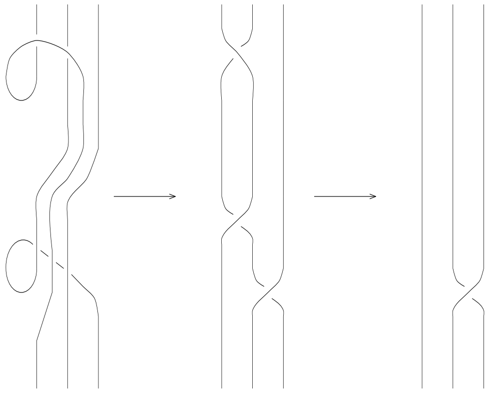

FIGURE 22. Any crossing can be realized on U_n

In [23] it is shown that U is a subtemplate of $L(0, -2)$. This, together with Propositions 7.3 and 7.5, shows that all the templates $L(0, n)$ for $n < 0$ are universal. While there are examples of positive templates with composite knots, *e.g.* $L(n, m)$ for n, $m > 1$, there always appears to be a limit on the number of prime factors. It is conjectured to be two for the example just cited. For examples (besides the Lorenz template) where such bounds have been confirmed, see [24].

Conjecture 8.1. *If T is a positive template, there is a number n such all periodic orbits of T have n or fewer prime factors.*

Holmes and Ghrist [14] have found differential equations and an open set of parameter values in which U arises. There are even electrical circuits governed by these equations. On the other hand, Holmes [16] has pointed out that positive templates arise naturally in the dynamics of forced oscillators, provided the hypothesis of hyperbolicity holds.

We remark that Ghrist's isotopic embedding of U_n in U is quite convoluted. For the figure-8 knot the number of strands in Ghrist's presentation runs into the millions. This is of some comfort to Sullivan and Williams, who had looked high and low for the figure-8 knot in this template without success. Still, it would be interesting to find minimum representations of a given knot or link in U. Anyone care to try?

The study of knotted periodic orbits and related phenomena is a growing and exciting field. The reader wanting to learn more may wish to consult the recently published book, *Knots and Links in Three-Dimensional Flows* [15].

References

[1] D. Bennequin. Entrelacement et equations de Pfaff. *Asterisque,* **107-108** (1983) pp. 87-161.

[2] J. Birman & R. Williams. Knotted Periodic Orbits in Dynamical Systems I: Lorenz Knots. *Topology,* **22** (1983) 47–82.

[3] J. Birman & R. Williams. Knotted Periodic Orbits in Dynamical Systems II: Knot Holders for Fibered Knots. *Contemporary Math.,* **20** (1983) 1–59.

[4] R. Bowen. One-dimensional hyperbolic sets for flows. *J. Differential Equations,* **12** (1972) 173–179.

[5] R. Bowen & J. Franks. Homology for zero dimensional basic sets. *Ann. of Math.,* **106** (1977) 73–92.

[6] G. Burde & H. Zieschang. *Knots,* Walter de Gruyter, New York, 1985.

[7] P. Cromwell. Positive braids are visually prime. *Proc. London Math. Soc.,* (3) **67** (1993) 384-424.

[8] F. Dumortier, H. Kokubu, & H. Oka. A degenerate singularity generating geometric Lorenz attractors. *Ergod. Th. & Dynam. Sys.*, **15** (1995), no. 5, 833–856.

[9] J. Franks. *Homology and Dynamical Systems,* AMS, 1980.

[10] J. Franks. Symbolic dynamics in flows on three-manifolds. *Tran. Amer. Math. Soc.*, Vol. 279, No. 1, (1983) 231–236.

[11] J. Franks. Flow equivalence of subshifts of finite type. *Ergod. Th. & Dynam. Sys.*, Vol. 4 (1984) 53–66.

[12] J. Franks & R. Williams. Entropy and Knots. *Trans. Amer. Math. Soc.*, **291** (1985) 241–253.

[13] R. W. Ghrist. Branched two-manifolds supporting all links. *Topology*, 36(2), 423–438, 1997.

[14] R. W. Ghrist. & P. Holmes. An ODE whose solutions contain all knots and links. To appear in *Intl. J. Bifurcation and Chaos*, 1996.

[15] R. W. Ghrist, P. Holmes & M. Sullivan. *Knots and Links in Three-Dimensional Flows: Lecture Notes in Mathematics*, Vol. 1654, Springer-Verlag, Berlin, 1997.

[16] P. Holmes. Knot and Orbit Genealogies in Nonlinear Oscillators. In *New Directions in Dynamical Systems*, pp 150-191, ed. T. Bedford and J. Swift, London Math. Soc. Lecture Notes 127, Cambridge University Press, 1988.

[17] P. Holmes & R. Williams. Knotted Periodic Orbits in Suspensions of Smale's Horseshoe: Torus Knots and Bifurcation Sequences. *Arch. Rational Mech. Anal.,* **90** (1985) 115–194.

[18] D. Lind & B. Marcus. *Symbolic Dynamics and Coding*, Cambridge Univ. Press, 1995.

[19] B. Parry & D. Sullivan. A topological invariant of flows on 1-dimensional spaces. *Topology*, **14** (1975) 297–299.

[20] C. Robinson. *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*, CRC Press Inc., Boca Raton FL, 1994.

[21] S. Smale. Differentiable dynamical systems. *Bull. Amer. Math. Soc.,* **73** (1967) 797–817.

[22] M. Sullivan. Prime decomposition of knots in Lorenz-like templates. *Journal of Knot Theory and its Ramifications,* Vol. 2, No. 4 (1993) 453–462.

[23] M. Sullivan. Composite knots in the figure-8 knot complement can have any number of prime factors. *Topology and its Applications,* 55 (1994) 261–272.

[24] M. Sullivan. The prime decomposition of knotted periodic orbits in dynamical systems. *The Journal of Knot Theory and its Ramifications,* Vol. 3 No. 1 (1994) 83–120.

[25] M. Sullivan. A zeta-function for flows with positive templates. *Topology and its Applications,* 66 (1995) 199–213.

[26] M. Sullivan. Positive braids with a half twist are prime. *The Journal of Knot Theory and its Ramifications,* Vol. 6, No. 3 (1997) 405–415.

[27] J. M. Van Buskirk. Positive knots have positive Conway polynomials. In *Knot Theory and Manifolds: Lecture Notes in Mathematics*, Vol. 1144, Ed. D. Rolfsen, Springer-Verlag, 1983.

[28] R. Williams. Classification of subshifts of finite type, *Ann. of Math.* **98** (1973), 120-153; Errata **99** (1974) 380–381.

[29] R. Williams. Lorenz Knots are Prime. *Ergod. Th. & Dynam. Sys.*, **4** (1983) 147–163.

Department of Mathematics, Northwestern University, Evanston, IL 60208-2730, USA

Department of Mathematics, Southern Illinois University, Carbondale, IL 62901-4408, USA