that the folds will enclose larger and larger domains. Their areas, however, stay bounded since Γ is an invariant curve on the cylinder. Therefore those additional areas must tend to zero. But this can only happen if Γ has a point of self-intersection, which contradicts its embeddedness.

I would like to thank Patrice Le Calvez for drawing my attention to the fact that Birkhoff's Theorem is not true without the area-preserving assumption, as well as Martin Beibel (from the Institute for Mathematical Stochastics, University of Freiburg) for reading and commenting on a preliminary version. This proof was presented in one of those evening sessions during the Dynamical Systems meeting in Oberwolfach (1997), and I thank everyone in the audience for attending.

2. BIRKHOFF'S THEOREM

We consider a C^1-diffeomorphism $\phi \colon \mathbf{S}^1 \times \mathbf{R} \to \mathbf{S}^1 \times \mathbf{R}$ of the two-dimensional cylinder; for the sake of simplicity, we keep the same notation for a lift of ϕ to $\mathbf{R}^2$ with coordinates x, y.

DEFINITION. We say that ϕ is a *monotone twist mapping* if the following three conditions hold:

- $\phi^*(dx \wedge dy) = dx \wedge dy$, i.e. ϕ preserves area and orientation.
- $\pi_y \circ \phi(x, y) \to \pm\infty$ as $y \to \pm\infty$, i.e. ϕ preserves the ends of the cylinder.
- $|\partial(\pi_x \circ \phi)/\partial y| \geq \delta > 0$, i.e. ϕ satisfies a uniform monotone twist condition.

According to the sign of $\partial(\pi_x \circ \phi)/\partial y$, we call ϕ a *positive*, respectively *negative*, monotone twist mapping.

The uniformity of the twist condition has the following geometric interpretation ("cone condition"). Let ϕ be a positive monotone twist map, and denote by v_x the vertical $\{x\} \times \mathbf{R}$. Then the image $\phi(v_x)$ crosses the vertical through $\phi(x, y)$ in positive direction and stays outside a cone around it with centre $\phi(x, y)$, whose angle depends only on the twist constant δ; see Figure 3.

Note that if ϕ is a positive monotone twist mapping then its inverse ϕ^{-1} is a negative monotone twist mapping.

For the statement of the theorem, recall that a closed continuous curve is embedded if it is homeomorphic to $\mathbf{S}^1$; in particular, it cannot have a point of self-intersection.
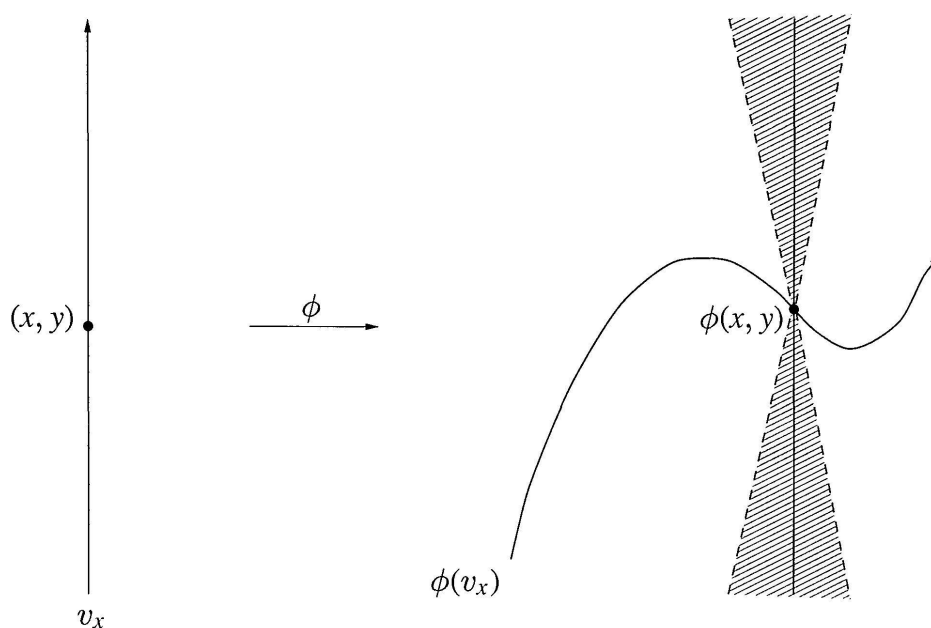
FIGURE 3

The "cone condition"

THEOREM (Birkhoff). *Let ϕ be a monotone twist mapping on $\mathbf{S}^1 \times \mathbf{R}$, and Γ a closed, embedded, homotopically nontrivial curve in $\mathbf{S}^1 \times \mathbf{R}$ such that $\phi(\Gamma) = \Gamma$.*

Then Γ is the graph of a Lipschitz continuous function on $\mathbf{S}^1$. Moreover, the Lipschitz constant can be bounded in terms of the twist constant δ.

The proof of Birkhoff's Theorem will take up the rest of this section. We assume that the monotone twist map ϕ possesses an embedded invariant curve Γ which is not a graph. From this we will conclude that Γ has a point of self-intersection, which contradicts the assumptions. The Lipschitz property will be proved at the very end.

I. SET-UP

We lift everything to $\mathbf{R}^2$ and keep the same notation. Fix a parametrization $\gamma \colon \mathbf{R} \to \mathbf{R}^2$ of Γ such that $\gamma(t + 1) = \gamma(t) + (1, 0)$. This equips Γ with an order inherited from $\mathbf{R}$, and we can say whether a point on Γ comes before or after another one. That Γ is not a graph means that the continuous function $f = \pi_x \circ \gamma \colon \mathbf{R} \to \mathbf{R}$ is not injective.

LEMMA 1. *We have one of the following two cases (or both):*
- *There are $d < e$ such that $f(d) = f(e)$ and $f(t) > f(d)$ for all $t \in (d, e)$;*
- *f is constant on some nontrivial interval.*

Proof. Since f is not injective there are $a < b$ with $f(a) = f(b) = h$. If f is not constant on $[a, b]$ then $m = \min_{[a,b]} f < h$ or $M = \max_{[a,b]} f > h$. In the first case, we set $d = \max\{t < a \mid f(t) = m\}$ and $e = \min\{t > a \mid f(t) = m\}$; then $f(d) = f(e) = m$ and $f(t) > m$ for $t \in (d, e)$. In the second case, we put $c = \min\{t > a \mid f(t) = M\}$ and set $d = \max\{t < c \mid f(t) = h\}$ and $e = \min\{t > c \mid f(t) = h\}$; then $f(d) = f(e) = h$ and $f(t) > h$ for $t \in (d, e)$. Note that all numbers are well-defined because f is continuous and $f(t) \to \pm\infty$ as $t \to \pm\infty$. $\square$

II. THE FIRST CASE

Let us deal with the first case from Lemma 1, and denote by v_x the vertical $\{x\} \times \mathbf{R}$. By construction, the points $D_0 = \gamma(d)$ and $E_0 = \gamma(e) = (x_0, y_0)$ lie on the same vertical v_{x_0}. Moreover, the part of Γ between D_0 and E_0, together with the part of the vertical v_{x_0} between E_0 and D_0, forms an embedded simply closed curve. By the Jordan-Schoenflies Theorem, this curve bounds a domain in $\mathbf{R}^2$ which we call Ω_0.

There are two alternatives: either D_0 lies above E_0 on v_{x_0}, i.e. $\pi_y(D_0) > \pi_y(E_0)$, or below. In the first case, we choose ϕ or ϕ^{-1} in such a way that we obtain a positive monotone twist map; the second alternative requires a negative twist map. Without loss of generality, we assume that D_0 lies above E_0 and ϕ is a positive monotone twist mapping.

We set $x_1 = \pi_x(\phi(E_0))$ and consider the intersection points of $\phi^{-1}(v_{x_1})$ and Γ; E_0 is one of them. Let A_0 be the first intersection point of $\phi^{-1}(v_{x_1})$ and Γ before D_0 (with respect to the order on Γ). See Figure 4 by way of illustration.

LEMMA 2. *The point A_0 is well-defined.*

Proof. The curve $y \mapsto \phi^{-1}(x_1, y)$ separates the plane into two domains and its second coordinate tends to $\pm\infty$ as $y \to \pm\infty$. The point $D_0 \in \Gamma$ lies in one of the two domains, more precisely, in $\phi^{-1}((x_1, +\infty) \times \mathbf{R})$ because ϕ^{-1} is a negative monotone twist map and D_0 lies above E_0.

Recall that Γ is parametrized by γ such that $\gamma(t + 1) = \gamma(t) + (1, 0)$. Therefore one of the points $\gamma(d - k) = D_0 - (k, 0)$ with $k \geq 1$ lies in the other domain $\phi^{-1}((-\infty, x_1) \times \mathbf{R})$. Since Γ is homotopically nontrivial, $\gamma|_{[d-k,d]}$ is a connecting path between them. Hence Γ must intersect $\phi^{-1}(v_{x_1})$.

Finally, we claim that there is a first intersection point on Γ before D_0; this will be our A_0. If not, there is a sequence of intersection points between $\phi^{-1}(v_{x_1})$ and Γ accumulating at D_0, and so, by continuity, $D_0 \in v_{x_0}$ belongs

also to $\phi^{-1}(v_{x_1})$. But then $\phi(v_{x_0}) \cap v_{x_1}$ contains two points, in contradiction to the twist property.   □

Let us define the pre-image $\phi^{-1}(E_1)$ of $E_1 = (x_1, y_1) \in v_{x_1}$ to be the last intersection point of Γ and $\phi^{-1}(v_{x_1})$ before A_0 (with respect to the natural order on $\phi^{-1}(v_{x_1})$ inherited from that on v_{x_1}). $\phi^{-1}(E_1)$ is different from A_0, since otherwise it would be a point of self-intersection for Γ, which is excluded by our assumption that Γ is embedded. Of course, it may happen that $\phi^{-1}(E_1)$ and E_0 are one and the same point on Γ, but in general $\phi^{-1}(E_1)$ comes after E_0.
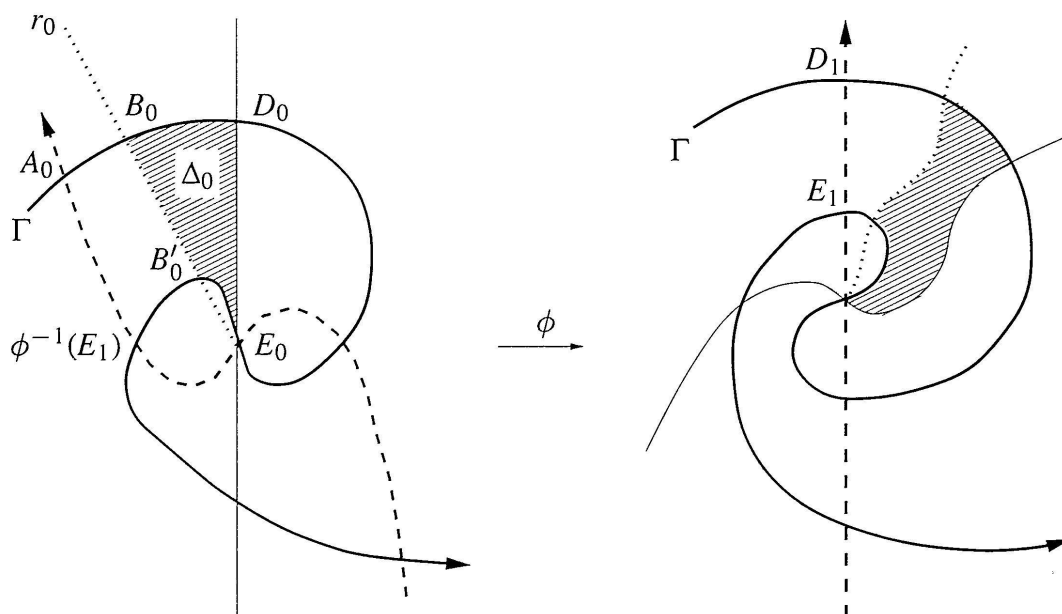

FIGURE 4

The first step of the iteration procedure

Again, the part of Γ between A_0 and $\phi^{-1}(E_1)$, together with that of $\phi^{-1}(v_{x_1})$ between $\phi^{-1}(E_1)$ and A_0, bounds a domain; its image under ϕ will be denoted by Ω_1. The vertical segment between E_0 and D_0 lies completely in $\phi^{-1}(\Omega_1)$ and divides it into two domains, Ω_0 and $\phi^{-1}(\Omega_1) \setminus \Omega_0$.

IIa. APPLYING ϕ ONCE

Now we apply ϕ to the whole picture. $\phi^{-1}(v_{x_1})$ will be mapped onto the vertical v_{x_1} through $D_1 = \phi(A_0)$ and E_1, where D_1 lies above E_1 because ϕ preserves the orientation. If we just look at the part of Γ between D_1 and E_1 and that of v_{x_1} between E_1 and D_1, we are in the same topological situation as before – together, they enclose the domain Ω_1. It does not matter that the

part of Γ may curl and intersect v_{x_1} again. What is important, however, is the fact that the area of the new Ω_1 has increased:

$$|\Omega_1| = |\Omega_0| + |\phi^{-1}(\Omega_1) \setminus \Omega_0|$$

We need an estimate from below for that additional area. To do so, we choose a ray r_0, centred at E_0 and pointing into the second quadrant, such that $\phi^{-1}(v_{x_1})$ does not intersect the open half cone between r_0 and $\{x_0\} \times [y_0, +\infty)$; see Figure 4. That this is possible follows from the above-mentioned "cone condition" for a monotone twist map. We point out that the angle of the corresponding half cone can be chosen independent of the base point on Γ.

We define B_0 to be the first intersection point of r_0 and Γ before D_0 (with respect to the order on Γ), and B_0' to be the last intersection point of Γ and r_0 before B_0 (with respect to the natural order on r_0). The existence of B_0 and B_0' is guaranteed by the same reasoning as in the proof of Lemma 2. Moreover, B_0' is different from B_0 because, otherwise, Γ would have a self-intersection. Note that it is possible that $B_0' = E_0$.

We call Δ_0 the domain bounded by the parts of Γ between B_0 and D_0, and E_0 and B_0', as well as r_0 between B_0' and B_0, and v_{x_0} between D_0 and E_0. Then we have

$$|\Omega_1| \geq |\Omega_0| + |\Delta_0|.$$

IIb.   APPLYING ϕ MANY TIMES

Now we iterate the above procedure. For this, we set $x_2 = \pi_x(\phi(E_1))$ and define A_1 and $\phi^{-1}(E_2)$ as intersection points of $\phi^{-1}(v_{x_2})$ and Γ in a completely analogous way as before. After one application of ϕ, we obtain a new domain Ω_2 whose area can be estimated by

$$|\Omega_2| \geq |\Omega_1| + |\Delta_1| \geq |\Omega_0| + |\Delta_0| + |\Delta_1|.$$

After n iterations, we obtain

$$|\Omega_n| \geq |\Omega_0| + \sum_{k=0}^{n-1} |\Delta_k|.$$

Note that $\phi^n(\Gamma) = \Gamma$ is fixed for all n and contained in some strip $\mathbf{R} \times [-R, R]$. Let us call L the horizontal diameter of the "fundamental part" $\gamma|_{[0,1]}$ of Γ. Then $\sup_{n \geq 0} |\Omega_n| \leq 2R \cdot L$, and hence

$$|\Delta_n| \to 0$$

as $n \to \infty$.

IIc. THE GRAPH PROPERTY

From the previous discussion, we will now derive that Γ must have a self-intersection, which contradicts the assumption that Γ is embedded. We define the points B_n, B_n', D_n and E_n on Γ exactly as before. Call Γ_n the part of Γ between B_n and D_n, and Γ_n' that between E_n and B_n' (which may reduce to the single point $E_n = B_n'$). We distinguish two cases.

If $\operatorname{dist}(\Gamma_n, \Gamma_n') \to 0$, then there are points $C_n \in \Gamma_n$ and $C_n' \in \Gamma_n'$ such that $\operatorname{dist}(C_n, C_n') \to 0$, and we may assume that all of them lie in $[0, 1] \times \mathbf{R}$. This means that (on subsequences) C_n and C_n' converge to one and the same point on Γ. This is a point of self-intersection, because the part of Γ between C_n and C_n' is always part of the boundary of a domain whose area is at least $|\Omega_0|$.

Ignoring subsequences, the other case is when $\operatorname{dist}(\Gamma_n, \Gamma_n') \geq \epsilon > 0$. Then we can put an open ball of diameter ϵ between Γ_n and Γ_n'. The area of Δ_n is at least that of the ball, intersected with the half cone between the rays from E_n through D_n (the upper part of the vertical v_{x_n}), and from E_n through B_n (which is r_n). Consider, in general, the area of the intersection of a half cone with a ball whose centre lies inside that half cone; this area becomes smallest if we put the centre of the ball right at the corner. In our situation, the crucial point is that the angle at the corner E_n is fixed for all n. Therefore the area of the above disk segment is a lower bound for all $|\Delta_n|$. But this contradicts $|\Delta_n| \to 0$, so this case cannot happen.

Thus our assumption that Γ is not a graph leads to a contradiction.

IId. THE LIPSCHITZ PROPERTY

We want to show that Γ is the graph of a Lipschitz function, whose Lipschitz constant can be estimated in terms of the twist constant δ. Pick any point P on Γ, and consider the ray r_P constructed in the same way for P, as r_0 had been constructed for E_0 in Section IIa. In particular, the angle between r_P and the vertical through P depends only on δ. If Γ intersects r_P in a second point different from P, then the pre-image of the vertical through $\phi(P)$ must intersect Γ in a second point, too; see Figure 5. This follows from the same arguments as in the proof of Lemma 2. But now one application of ϕ shows that the vertical through $\phi(P)$ intersects Γ in at least two points, which is impossible since Γ is a graph. Therefore Γ cannot intersect any of the r_P's, hence it is a Lipschitz graph with Lipschitz constant only depending on δ.
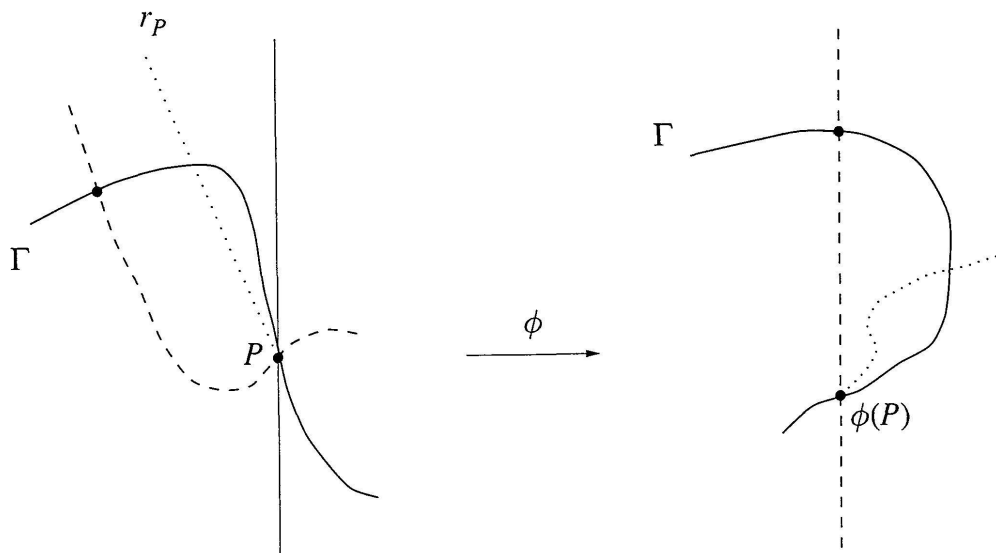
FIGURE 5

Why Γ must be a Lipschitz graph

III. THE SECOND CASE

Finally, the same remark can be applied in the second of the two cases from Lemma 1 where Γ contains a whole vertical interval. For we may take P to be the midpoint of that interval and apply ϕ once – the vertical through $\phi(P)$ will intersect Γ in two isolated points D_0 and E_0, and we are back in the first situation we already dealt with.

The proof of the theorem is complete.

3. CONCLUDING REMARKS

For the sake of clarity, we did not prove the most general result that can be obtained by our method. Here we just indicate possible generalizations.

First of all, our proof does not require the monotone twist condition but only a sort of "cone condition on Γ". Namely, what we really need is the requirement that all (pre-)images of verticals lie outside certain cones centred at points on Γ; we do not use the much more restrictive fact that they are graphs. (This subtle point might be the reason why we have not succeeded in proving a well-known generalization of Birkhoff's Theorem to boundaries of invariant annuli [Fa, He, KH] by our method.)