DE GRUYTER Adv. Geom. 2017; 17 (4):423–434

Bruce Solomon*

Central figure-8 cross-cuts make surfaces cylindrical

DOI 10.1515/advgeom-2017-0030. Received 5 December, 2015

Abstract: Let M be a complete connected C^2-surface in $\mathbb{R}^3$ in general position, intersecting some plane along a clean figure-8 (a loop with total curvature zero) and such that all compact intersections with planes have central symmetry. We prove that M is a (geometric) cylinder over some central figure-8. On the way, we establish interesting facts about centrally symmetric loops in the plane; for instance, a clean loop with even rotation number $2k$ can never be central unless it passes through its center exactly twice and $k = 0$.

Keywords: Surface, cross-cut, figure-8, central symmetry, cylinder.

2010 Mathematics Subject Classification: 53A04

Communicated by: P. Eberlein

1 Introduction

A set $X \subset \mathbb{R}^{n+1}$ has a *center* $c \in \mathbb{R}^{n+1}$ (or has *central symmetry*, or is *central*) if the c-fixing reflection $x \mapsto 2c - x$ maps X to itself. What can one say about a set $X \subset \mathbb{R}^{n+1}$ that meets every hyperplane along a central set?

When P is a hyperplane, we (for now) call $X \cap P$ a *cross-cut* of X. Later we define *cross-cut* more narrowly. Are cross-cuts of central sets always central? Not generally, unless they go through the center. A cube in $\mathbb{R}^3$ is central, but a plane that severs a small corner cuts it along a triangle, which is not. Do central cross-cuts make a set central? Not in $\mathbb{R}^2$. For instance, all cross-cuts of a plane triangle are (trivially) central, but again, no triangle is central.

When $n + 1 > 2$, however, no such counterexample comes to mind. And in the presence of convexity, central cross-cuts can force *more* than just centrality. For example, when $n + 1 > 2$ and $K \subset \mathbb{R}^{n+1}$ is a convex body, central cross-cuts make K *ellipsoidal*. The most general known formulation of this fact was proven by Olovjanischnikoff [7], who relaxed restrictions (e.g. on smoothness) in earlier results by Brunn and Blaschke (see [2, §44, §84] and [1]). Burton gives a nice statement of Olovjanischnikoff's result in [3, Lemma 3]. In [8], we drew a similar conclusion for (not necessarily convex) hypersurfaces of revolution in $\mathbb{R}^{n+1}$. If their compact convex cross-cuts are central, they must be quadrics: ellipsoids, hyperboloids, paraboloids, or circular cylinders. Using that fact, we got a broader result in [9]: When a complete immersed surface in $\mathbb{R}^3$ has a connected compact transverse cross-cut, and all such cross-cuts are central, uniformly convex ovals, the surface is either a central cylinder or a tubular quadric.

None of these results, however, manages to exploit centrality of cross-cuts without also requiring their convexity. Here for the first time, we drop the convexity requirement, replacing it with a very different, albeit special, alternative. We consider surfaces in $\mathbb{R}^3$ whose cross-cuts are *clean* (meaning they never visit a point twice *and* do so tangent to the same line) and have total geodesic curvature zero. The latter condition makes them *figure-8's* up to regular homotopy. Our main result, Theorem 3.6, says that a surface with this property must be a central cylinder. Section 3 of our paper presents the main arguments in its proof.

Section 2 (which we find interesting on its own) is devoted mainly to the proof of a simple but critical ingredient: *Any clean, central figure-8 must visit its center exactly twice*. The key role this plays in Section 3 is explained in the paragraphs immediately below our statement of Theorem 3.6. In proving the supporting fact,

***Corresponding author: Bruce Solomon,** Math Department, Indiana University, Bloomington IN 47405, USA, email: solomon@indiana.edu

however, we get the general theory summarized in Proposition 2.16, which says, in part, that a clean central loop must either have *odd* rotation index, and avoid its center entirely, or else be a figure-8 (index zero) that visits its center precisely twice.

Simple examples—the unit circle traced twice, for instance, or the loop in Figure 2—show that such statements fail for loops that are not cleanly immersed. The reasoning in both §2 and §3 would simplify considerably if we had not needed to assume and exploit general position arguments to exclude unlikely "pathologies" like those.

2 Reparametrization and symmetry

Definition 2.1 (Central symmetry). An immersion $F: M \to \mathbb{R}^{n+1}$ is *central* if its image has central symmetry.

Definition 2.2 (Reparametrization). When $\alpha, \beta: \mathbb{S}^1 \to \mathbb{R}^2$ are immersed loops, we say that β *reparametrizes* α when $\beta = \alpha \circ \varphi$ for some diffeomorphism $\varphi: \mathbb{S}^1 \to \mathbb{S}^1$. It *preserves* or *reverses* orientation when φ preserves or reverses the orientation of the circle.

The following fact will pester us: Two immersed loops with the same image do not always reparametrize each other, even if they visit each point equally often. Centrality does not mitigate this inconvenient truth, as discussed with regard to Figure 2 below.

Examples 2.3. The unit circle $\mathbb{S}^1 \subset \mathbb{C}$ is central about the origin. If we parametrize it as usual by $t \mapsto e^{it}$, reflection through the origin produces the orientation-*preserving* reparametrization $t \mapsto -e^{it}$. Contrastingly, consider the figure-8 parametrized by $t \mapsto (\cos t, \sin 2t)$. While likewise central about the origin, reflection through the origin induces the orientation-*reversing* reparametrization $t \mapsto (\cos t, -\sin 2t)$. See Figure 1.

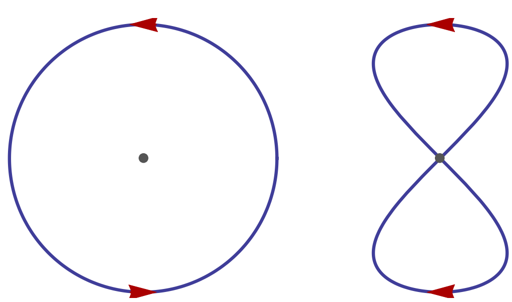

Figure 1: Reflection through the origin reparametrizes both the unit circle and figure-8, preserving orientation on the former, but reversing it on the latter.

In Figure 2 however, we depict a smooth, origin-central immersion $\alpha: \mathbb{S}^1 \to \mathbb{R}^2$ whose reflection $-\alpha$ does *not* reparametrize α. To see this, orient the open sets U_L, U_0 and U_R there in the standard way (so that integration of $dx \wedge dy$ gives a positive result), and observe that $\alpha = \partial(U_0 + U_R - U_L)$ and $-\alpha = \partial(U_0 - U_R + U_L)$. The oriented domains bounded by α and $-\alpha$ are neither equal nor opposite, so $-\alpha$ neither preserves nor reverses the orientation of α. It evidently does not, therefore, reparametrize.

We can exclude the behavior depicted in Figure 2 by requiring our loops to be *cleanly immersed*:

Definition 2.4 (Double-points/clean loops). A point p in the image of an immersed curve α is a *double-point* when its preimage contains two or more points. When it contains *exactly* two, we call it *simple*.

An immersion $\alpha: \mathbb{S}^1 \to \mathbb{R}^2$ has *clean* double-points (or *is clean*) if, whenever $t_1, t_2 \in \mathbb{S}^1$ are distinct preimages of a single point in $\mathbb{R}^2$, they have respective neighborhoods U_1 and U_2 whose images $\alpha(U_1)$ and $\alpha(U_2)$ intersect transversally.

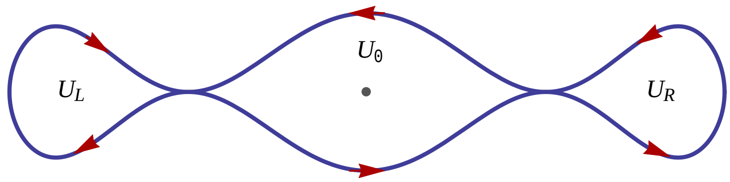

Figure 2: This immersed loop α winds counter-clockwise around U_R and U_0, but clockwise around U_L. Reflection through the center preserves its image, but $-\alpha$ is not a reparametrization of α.

Remark 2.5. Though more familiar, a *general position* assumption (like the one we use early in §3 below) would be more restrictive than that of *clean double-points* for loops in $\mathbb{R}^2$. The latter lets a loop pass through a single point three or more times as long as no two velocity vectors there are dependent. General position would prohibit triple intersections.

The double-points in Example 2.3 are *not* clean, and we shall see that when two loops with the same image do *not* reparametrize each other, they *must* have unclean double-points. Indeed, the principal result of this section, Proposition 2.11, and the main facts leading up to it all fail without cleanness, as reconsideration of Figure 2 will quickly reveal.

We denote the intrinsic distance between points s_1, s_2 in the sphere of any dimension (including the circle) by $\varphi(s_1, s_2)$. We write κ_g for the *geodesic* (i.e. *signed*) *curvature* along a loop $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$. It is given by

$$\kappa_g(t) := \frac{\det(\alpha', \alpha'')}{|\alpha'|^3}. \tag{2.1}$$

Observation 2.6. *Suppose that we have a C^2 unit-speed loop $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$. If $\bar\kappa := \max_{\mathbb{S}^1} |\kappa_g|$, and $\alpha^{-1}(p)$ contains distinct inputs $s_1, s_2 \in \mathbb{S}^1$ for some $p \in \mathbb{R}^2$, then $\varphi(s_1, s_2) \geq \pi/\bar\kappa$.*

Proof. Let A be either arc in $\mathbb{S}^1$ joining α_1 to α_2. Then some point $s_0 \in A$ maximizes $|\alpha(s) - p|^2$ on A, and $\dot\alpha(s_0)$ is then *perpendicular* to $\alpha(s_0) - p$. Choose $\mathbf{0} \neq \mathbf{v} \in \mathbb{R}^2$ with $\mathbf{v} \cdot (\alpha(s_0) - p) = 0$. Then $\mathbf{v} \cdot \alpha(s)$ attains at least one local extremum on each component of $A \setminus \{s_0\}$, say at points s_- and s_+ respectively. Since $\dot\alpha$ must be *parallel* to $\alpha(s_0) - p$ at these points, α maps the intervals (s_-, s_0) and (s_0, s_+) both to arcs with total absolute curvature at least $\pi/2$. As α has unit speed, we may then deduce that

$$\varphi(s_-, s_+)\bar\kappa \geq \int_{s_-}^{s_+} |\kappa_g(s)|\, ds \geq \pi. \qquad \square$$

In general, a C^2-immersion $\mathbb{S}^1 \to \mathbb{R}^2$ can have infinitely many double points, even without retracing any open arc along its image. Not so for clean immersions:

Lemma 2.7. *A clean C^2-immersion $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ has at most finitely many double-points, and at any double-point p, there is an $\varepsilon = \varepsilon(p) > 0$ for which $\alpha^{-1}(B(p, \varepsilon))$ is a finite union of embedded arcs passing through p with pairwise distinct tangent lines.*

Proof. With no loss of generality, assume that α has unit speed. Set $\bar\kappa := \max_{\mathbb{S}^1} |\kappa_g|$ as in the preceding Observation. Suppose (toward a contradiction) that α had infinitely many double points. Since $\mathbb{S}^1$ and $\alpha(\mathbb{S}^1)$ are both compact, that would imply the existence of a cluster point $p \in \alpha(\mathbb{S}^1)$, along with convergent sequences $(s_n), (s_n') \subset \mathbb{S}^1$ with

$$s_n \neq s_n' \quad \text{and} \quad \alpha(s_n) = \alpha(s_n') \to p$$

and yet $\alpha(s_n) \neq p$ for all $n \in \mathbb{N}$. Observation 2.6 ensures $|s_n - s_n'| > \pi/\bar\kappa$, so the respective limits s and s' of these sequences must obey that same estimate. In particular, $s \neq s'$. By continuity, however, $\alpha(s) = \alpha(s')$, which forces the collinearity of

$$\frac{\alpha(s_n) - \alpha(s)}{s_n - s} \quad \text{and} \quad \frac{\alpha(s_n') - \alpha(s')}{s_n' - s'}$$

for each n. Letting $n \to \infty$, we see that $\dot\alpha(s)$ and $\dot\alpha(s')$ must also be collinear. This contradicts our assumption of clean double-points. So α has at most finitely many double-points.

To prove the remaining claim, recall that Observation 2.6 puts a lower bound on the distance between any two points in $\alpha^{-1}(p)$, so the compactness of $\mathbb{S}^1$ makes $\alpha^{-1}(p)$ finite. By definition of *immersion*, the inverse function theorem then yields the asserted $\varepsilon(p) > 0$, while that of *clean* makes tangent lines pairwise distinct at p. □

Lemma 2.8. *Suppose that $\alpha, \beta\colon \mathbb{S}^1 \to \mathbb{R}^2$ are clean unit-speed C^2 loops with the same image. Suppose that $p = \beta(t_0)$ is a double-point of α, and that $\varepsilon > 0$ is small enough to make $\alpha^{-1}(B(p,\varepsilon))$ a union of embedded arcs with distinct tangent lines at p, as provided by Lemma 2.7. Then one such arc contains $\beta(t_0 - \delta, t_0 + \delta)$ for all sufficiently small $\delta > 0$.*

Proof. Take $\varepsilon > 0$ small enough to satisfy the hypothesis of Lemma 2.7, and let $A_1, A_2, \ldots, A_k$ denote the (distinct) arcs whose union then constitutes $\alpha^{-1}(B(p,\varepsilon))$. Define $I_n := (t_0 - \frac{1}{n}, t_0 + \frac{1}{n})$ for $n \in \mathbb{N}$. When n is large, β embeds I_n, and since β and α have the same image, $\beta(I_n)$ must then lie in the *union* of the A_i's.

If for every such n, we could find $t_n, t_n' \neq t$ in I_n with $\beta(t_n)$ and $\beta(t_n')$ in *different* A_i's, we could renumber the A_i's and pass to a subsequence to arrange $\beta(t_n) \in A_1$ and $\beta(t_n') \in A_2$ for all large n. But $\lim_{n\to\infty} t_n = \lim_{n\to\infty} t_n' = t_0$, and β is differentiable, so computing $\dot\beta(t_0)$ on the two different sequences would give the same result, forcing the tangent lines to A_1 and A_2 at $p = \beta(t_0)$ to agree. This would contradict the last assertion of Lemma 2.7. So $\beta(I_n)$ must stay in one A_i, as claimed. □

Definition 2.9. By the *lift* of an immersed unit-speed arc $\alpha\colon (a,b) \to \mathbb{R}^2$, we mean the arc parametrized by $s \mapsto (\alpha(s), \dot\alpha(s))$ in the unit tangent bundle $\mathbb{R}^2 \times \mathbb{S}^1$.

Using Lemma 2.7, the reader will easily verify

Observation 2.10. *If $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ is a cleanly immersed loop, its lift is embedded. The lift of any reparametrization $\alpha \circ \varphi$ either reparametrizes that of α, or never meets it, depending on whether φ preserves or reverses orientation respectively.*

We can now prove the fact that puts the main results of this section in easy reach.

Proposition 2.11. *Suppose that $\alpha, \beta\colon \mathbb{S}^1 \to \mathbb{R}^2$ are clean, unit-speed C^2-immersions with the same image. Then β reparametrizes α, and the two loops have the same orientation if and only if their lifts meet.*

Proof. By Observation 2.10, α and β have embedded lifts. If they *meet* above $\beta(b)$ for some $b \in \mathbb{S}^1$, then Lemma 2.8 provides an $a \in \mathbb{S}^1$ and a $\delta > 0$ such that $(a-\delta, a+\delta)$ and $(b-\delta, b+\delta)$ lift, via α and β respectively, to the same arc in $\mathbb{R}^2 \times \mathbb{S}^1$. The lifts of α and β therefore meet along a set relatively open in the image of each. The coincidence set is also closed (trivially), so the two lifts coincide entirely, manifesting (via the Inverse Function Theorem) an invertibly C^1 transition map between them.

The identity map on $\mathbb{R}^2 \times \mathbb{S}^1$ then induces a diffeomorphism between the circles parametrizing α and β, allowing us to read β as a reparametrization of α. Orientation is preserved, for the lifts would otherwise be completely disjoint by Observation 2.10.

If the lifts *are* completely disjoint, then (since clean immersions have at most finitely many double-points by Lemma 2.7) we can find a point $p \in \alpha(\mathbb{S}^1)$ with a single pre-image $\{t\} = \alpha^{-1}(p)$. Then α and β share a unique tangent line at p. If their lifts do not meet, β must lift to $(p, -\dot\alpha(t_0))$ above p. The lift of any orientation-*reversing* reparametrization β' of β thus *meets* that of α above p, making β' an orientation-*preserving* reparametrization of α by what we have already proven. So β itself reverses orientation, as claimed. □

We will use the proposition just proven mainly in the form of an immediate

Corollary 2.12. *Any clean* central *C^2-loop is reparametrized by its central symmetry.*

As Figure 1 shows, the reparametrization induced by a central symmetry of a clean loop may preserve or reverse orientation. The two possibilities have starkly different geometric implications, however. To see that, we will need Corollary 2.14 below — a further consequence of Proposition 2.11 — which requires this

Definition 2.13. The *centroid* (center of mass) $\mu(\alpha)$ of a C^1-loop $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^{n+1}$ with length L is the mean value of α relative to an arc-length parameter s:

$$\mu(\alpha) := \frac{1}{L}\int_{\mathbb{S}^1} \alpha(s)\, ds.$$

Note that the centroid of a loop with central symmetry may not coincide with its center of symmetry. For example, take the circles $(x \pm 1)^2 + y^2 = 1$, and parametrize their union, starting at $\mathbf{0}$, by tracing clockwise around the right-hand lobe, then counterclockwise around the left, and finally, clockwise around the right again. The origin will be a center of symmetry, but the centroid lies at $(1/3, 0)$.

Clean loops never exhibit that kind of discrepancy:

Corollary 2.14. *For a clean, central C^2-loop, the center of symmetry and the centroid coincide.*

Proof. Suppose that $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ is a clean C^2-loop with center of symmetry at $\mathbf{c} \in \mathbb{R}^2$, and length L. View it as a unit-speed L-periodic immersion $\mathbb{R} \to \mathbb{R}^2$. As we assume α to be clean, Proposition 2.11 provides a diffeomorphism φ of the circle (which lifts to $\mathbb{R}$) such that $2\mathbf{c} - \alpha = \alpha \circ \varphi$. By the chain rule and constancy of speed (which is preserved by the reflection), we must also have $|\varphi'| \equiv 1$. If we denote the unit-speed parameter for α by s, then $u = \varphi(s)$ gives a unit-speed parameter for its reflection $2\mathbf{c} - \alpha$, whose centroid is then clearly

$$2\mathbf{c} - \mu(\alpha) = \frac{1}{L}\int_0^L 2\mathbf{c} - \alpha(s)\, ds = \frac{1}{L}\int_0^L \alpha \circ \varphi(s)\, ds = \frac{1}{L}\int_0^L \alpha \circ \varphi(s)|\varphi'(s)|\, ds = \frac{1}{L}\int_0^L \alpha(u)\, du = \mu(\alpha).$$

Thus $\mu(\alpha) = \mathbf{c}$, as claimed. □

Definition 2.15. When an immersed C^1-loop $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ is central with respect to $\mathbf{c} \in \mathbb{R}^2$, we call the line segment joining $\alpha(t)$ to its reflection $2\mathbf{c} - \alpha(t)$ a *diameter* of α. If we can parametrize α so that

$$2\mathbf{c} - \alpha(t) = \alpha(t + \pi) \quad \text{for all } t \in \mathbb{S}^1 \tag{2.2}$$

(intertwining reflection through $\mathbf{c}$ with the antipodal map on $\mathbb{S}^1$), we say that α is *diameter-central*.

Diameter-central loops are obviously central, but the converse is false, as shown by the central figure-8 in Figure 1. Careful consideration of that picture reveals that the figure-8 is not diameter-central.

Proposition 2.16. *Suppose that $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ is a clean, central C^2-loop. Then either*

A) *the symmetry preserves orientation, in which case α is regularly homotopic to $e^{(2k+1)\theta}$ for some $k \in \mathbb{Z}$, avoids its center, and is diameter-central,*

   or

B) *the symmetry reverses orientation, in which case α is regularly homotopic to the figure-8, has a simple double-point at its center, and is not diameter-central.*

Proof. We can assume that α is centered at the origin $\mathbf{0}$ and (after applying a homothety giving it length 2π) has unit speed. Corollary 2.12 then says that $-\alpha = \alpha \circ \varphi$ for some diffeomorphism $\varphi\colon \mathbb{S}^1 \to \mathbb{S}^1$. By the chain rule, our unit speed assumption forces $|\varphi'| \equiv 1$, making φ an *isometry* of $\mathbb{S}^1$. An isometry either rotates $\mathbb{S}^1$ or reflects it across a diameter, preserving or reversing orientation respectively.

When we view $\mathbb{S}^1 \approx \mathbb{R}/2\pi$ as an additive group, rotation takes the form $\varphi(t) = t + l$ for some $l \in \mathbb{S}^1$. So if the symmetry preserves orientation, we get $-\alpha(t) = \alpha(t + l)$ for all t. Since α is not constant, we may assume that $0 < |l| \le \pi$. Iterating the symmetry then gives $\alpha(t + 2l) = \alpha(t)$, and hence $\dot\alpha(t + 2l) = \dot\alpha(t)$. Having clean double points, however, obstructs this pair of identities for $0 < |l| < \pi$. So in the orientation-preserving case, we must have $|l| = \pi$, which makes α diameter-central, as Conclusion A asserts.

A diameter-central loop has parallel tangent lines at $\alpha(t+\pi)$ and $\alpha(t)$, as seen by differentiating (2.2). We assume clean double-points, so this forces $\alpha(t + \pi) \ne \alpha(t)$ for all $t \in \mathbb{S}^1$. But we just saw that $\alpha(t + \pi) = -\alpha(t)$ for all $t \in \mathbb{S}^1$. So in the orientation-preserving case, our loop must avoid the origin — its center — as claimed by (A).

In the orientation-*reversing* case, by contrast, α is reparametrized by an isometry $\varphi\colon \mathbb{S}^1 \to \mathbb{S}^1$ that reflects across a diameter, fixing two antipodal points that we can assume to be $t = 0$ and $t = \pi$. In this case, for all $t \in \mathbb{S}^1$ we have $\varphi(t) = -t$, and thus

$$-\alpha(t) = \alpha(-t). \tag{2.3}$$

A central symmetry fixes only its center, however, forcing α to map both $t = 0$ and $t = \pi$ to the origin. In fact, the origin must be a *simple* double-point, as (B) claims. For, any central loop has parallel tangent lines at the ends of diameters, and when (2.3) holds, that means parallel tangent lines at $\alpha(t)$ and $\alpha(-t)$ for every $t \in \mathbb{S}^1$. If we had $\alpha(t) = \alpha(-t)$ for some t *not* fixed by φ, we would breach our clean double-points assumption.

It remains to verify the claims about regular homotopy. As is well-known (see e.g. [10] or [6, Proposition 2.1.6]), the regular homotopy class of an immersed plane loop $\alpha\colon \mathbb{S}^1 \to \mathbb{R}^2$ is classified by its *rotation index* — the degree $\omega_\alpha \in \mathbb{Z}$ of its unit tangent map $\alpha'/|\alpha'|\colon \mathbb{S}^1 \to \mathbb{S}^1$, which we may compute by integrating the geodesic curvature (2.1) along α:

$$\omega_\alpha = \frac{1}{2\pi}\int_0^{2\pi} \kappa_g(t)\,dt. \tag{2.4}$$

Consider first the orientation-preserving case. There, as we have seen, α is diameter-central: $\alpha(t+\pi) = -\alpha(t)$ for all t. It follows trivially that velocity and acceleration change sign too when we rotate the input by π. As easily seen from formula (2.1), however, this makes κ_g *even* on the circle: $\kappa_g(t+\pi) = \kappa_g(t)$. So when orientation is preserved, the total signed curvature of α is twice that along the arc $\alpha(0,\pi)$. At the same time, we have $\dot\alpha(\pi) = -\dot\alpha(0)$, forcing the unit tangent $\dot\alpha/|\dot\alpha|$ to traverse an odd number of semicircles as t varies from 0 to π. So

$$\int_0^{2\pi} \kappa_g(t)\,dt = 2\int_0^{\pi} \kappa_g(t)\,dt = 2(2k+1)\pi \quad \text{for some } k \in \mathbb{Z}.$$

By (2.4) we then have $\omega_\alpha = 2k+1$, an odd integer, as claimed.

In the orientation-*reversing* case, identity (2.3) replaces the diameter-central condition above. Differentiate that identity twice and use (2.4) to see that κ is now an odd function on the circle: $\kappa(-t) = -\kappa(t)$. The integral of an odd function vanishes, so (2.4) now yields $\omega_\alpha = 0$, making α regularly homotopic to the figure-8, as stated. This completes our proof. □

Corollary 2.17. *A clean plane C^2-loop with* even, non-zero *rotation index* cannot *have central symmetry.*

3 Main result

Supposing M^n is a smooth manifold, we now take up our motivating question: *What can we say about a complete, proper immersion $F\colon M^n \to \mathbb{R}^{n+1}$ when $F(M)$ has central intersections with an open set of hyperplanes?*

To address this, we introduce some notation. We write $u_p^\perp$ for the hyperplane containing $p \in \mathbb{R}^{n+1}$ and normal to $u \in \mathbb{S}^n$. When p is the origin, we simply write $u^\perp$. These hyperplanes are, respectively, zero sets of the affine functions u_p^* and u^* given by

$$u_p^*(x) = u\cdot(x-p), \qquad u^*(x) = u\cdot x$$

When using this notation, we always assume u to be a *unit* vector. As above, $\varphi(u,v) := \arccos(u\cdot v)$ denotes the angular distance between unit vectors u and v.

When $a > 0$ and $P = u_p^\perp$, we write P_a for the a-neighborhood of P:

$$P_a := \{q \in \mathbb{R}^{n+1} : |u_p^*(q)| < a\}. \tag{3.1}$$

Call $\nu \in \mathbb{S}^n$ a *unit normal* to an immersion $F\colon M^n \to \mathbb{R}^{n+1}$ at a point $x \in M$ if ν is orthogonal to the hyperplane $dF(T_xM)$ in $\mathbb{R}^{n+1}$.

We can then say that F has *general position* if, whenever $y \in \mathbb{R}^{n+1}$ and $\nu_1, \nu_2, \ldots, \nu_k$ are unit normals to F at distinct points in $F^{-1}(y)$, we have

$$\nu_1 \wedge \nu_2 \wedge \cdots \wedge \nu_k \neq \mathbf{0}. \tag{3.2}$$

If this holds when we extend F to $M \cup P$ for some hyperplane $P \subset \mathbb{R}^{n+1}$ via the inclusion map on P, we say that *F and P are in general position*. Note that in this case the restriction of F to M must itself have general position. When (3.2) holds for $k = 2$ (i.e., whenever ν_1, ν_2 are unit normals to F at distinct points of $F^{-1}(y)$), we get weaker conditions that we respectively express by saying *F has transverse self-intersections*, or *P meets F transversally*.

Transversality alone makes $F^{-1}(P)$ an embedded hypersurface in M; see [5, p.22]. General position guarantees more: when $n = 2$, for instance, it is not hard to see that it makes all double-points of $P \cap F(M)$ *clean* as specified in Definition 2.4.

We want to focus on the case where F and P have general position and the *compact* components of $F^{-1}(P)$ map to sets with central symmetry. Two definitions will help:

Definition 3.1 (Cross-cut). When a hyperplane $P \subset \mathbb{R}^{n+1}$ meets an immersion $F\colon M^n \to \mathbb{R}^{n+1}$ transversally, a *cross-cut of F relative to P* is a compact component $\Gamma \subset F^{-1}(P)$. We also call its image $F(\Gamma)$ a cross-cut; context will signal which meaning applies.

We call Γ a *clean* cross-cut when P and F are in general position.

The transversality assumption in Definition 3.1 ensures that the tangential projection on M of the unit normal u to P, i.e. $u - (u \cdot \nu)\nu$, is a *non-vanishing* transverse vectorfield along Γ. (The choice of unit normal ν to F is obviously irrelevant here.) Cross-cuts are thus *orientable* in M. A routine differential topology exercise then yields the existence of what we shall call a *good tubular coordinate neighborhood* U of a cross-cut Γ — a neighborhood that F maps to a tube foliated by cross-cuts diffeomorphic to Γ, each a level set of the height function u_p^*:

Definition 3.2 (Good tubular patch). Suppose that $\Gamma \subset M$ is a cross-cut of F relative to a hyperplane $P = u_p^\perp$. By a *good tubular coordinate neighborhood* (or *good tubular patch*) for Γ, we mean a pair (U, ψ), where $U \subset M$ is the image of an embedding $\psi\colon \Gamma \times [-a, a] \to M$ for some $a > 0$, and ψ has these three properties for all $(\theta, h) \in \Gamma \times [-a, a]$:

a) $\psi(x, 0) = x$ for all $x \in \Gamma$,
b) $(u_p^* \circ F \circ \psi)(\theta, h) = h$ and
c) $d(u_p^* \circ F \circ \psi) \neq \mathbf{0}$.

Property (b) means that for each $h \in [-a, a]$, the composition $F \circ \psi$ maps $\Gamma \times \{h\}$ into the plane $u_p^* \equiv h$. Property (c) makes F transverse to these same planes, so that $\psi(\Gamma \times \{h\})$ is a cross-cut of F for each $h \in [-a, a]$.

As mentioned above, the existence of a good tubular neighborhood of a cross-cut is routine. When a cross-cut is clean, we can guarantee that nearby cross-cuts are likewise clean:

Lemma 3.3. *Suppose that $\Gamma \subset M$ is a clean cross-cut relative to $P = u_p^\perp$, and that (U, ψ) is a good tubular patch for Γ. Then there is an $\varepsilon > 0$ for which $|q - p| < \varepsilon$ and $\varphi(v, u) < \varepsilon$ together ensure that $F^{-1}(v_q^\perp) \cap U$ is again a clean cross-cut, and is regularly homotopic to Γ.*

Proof. Define the map

$$\mathcal{F}\colon U \times \mathbb{R}^{n+1} \times \mathbb{S}^n \to \mathbb{R} \times \mathbb{R}^{n+1} \times \mathbb{S}^n \quad \text{via} \quad \mathcal{F}(x, q, v) = ((F(x) - q) \cdot v, q, v).$$

Property (c) in Definition 3.2 makes $d\mathcal{F}$ surjective at each point of $\mathcal{F}^{-1}(0, p, u) = \Gamma \times \{p\} \times \{u\}$, and lower-semicontinuity of rank then makes $d\mathcal{F}$ surjective on some neighborhood of $\Gamma \times p \times u$. If we denote ε-neighborhoods of p and u in $\mathbb{R}^{n+1}$ and $\mathbb{S}^n$ respectively by $B_\varepsilon(p)$ and $B_\varepsilon(u)$, the Implicit Function Theorem and compactness of Γ then make it straightforward to deduce that for some $\varepsilon > 0$, the $\mathcal{F}$-preimage of $(-\varepsilon, \varepsilon) \times B_\varepsilon(p) \times B_\varepsilon(u)$ is foliated by preimages $\mathcal{F}^{-1}(h, q, v)$, all regularly homotopic to $\mathcal{F}^{-1}(0, p, u) = \Gamma \times \{p\} \times \{u\}$. It follows that $\mathcal{F}^{-1}(\{0\} \times B_\varepsilon(p) \times B_\varepsilon(u))$ is likewise foliated. Since $\mathcal{F}^{-1}(0, q, v) = (F^{-1}(v_q^\perp) \cap U) \times \{q\} \times \{v\}$, this shows that

$|q - p| < \varepsilon$ and $\varphi(v, u) < \varepsilon$ together ensure, for every such q and v, that $F^{-1}(v_q^\perp) \cap U$ is a cross-cut regularly homotopic to Γ.

Finally, by making $\varepsilon > 0$ smaller still if necessary, we can guarantee that these cross-cuts are all clean too. Otherwise, we could find convergent sequences $(q_k) \to p$ and $(v_k) \to u$ for which each corresponding cross-cut $(F(x) - v_k) \cdot q_k \equiv 0$ in U was not clean. Condition (3.2) would then have to fail at some point y_k in each of these cross-cuts. Condition (3.2) is continuous in all variables, however, F is C^1, and U is compact. Passing to a subsequence, we could then take a limit as $k \to \infty$ and force a contradiction to our assumption that Γ itself was clean. □

With these differential-topological facts in hand, we now turn to the case of interest: where (the images of) all cross-cuts have central symmetry.

Definition 3.4 (CX). An immersion $F: M \to \mathbb{R}^{n+1}$ has the *central cross-cut property* (abbreviated CX) when

a) at least one clean cross-cut exists, and
b) the image of every clean cross-cut has *central symmetry*.

Note that CX is an *affine*-invariant property — not just a geometric one: if F has CX, and A is an affine isomorphism of $\mathbb{R}^{n+1}$, then $A \circ F$ has CX too.

In $\mathbb{R}^3$, circular cylinders and spheres have CX; they represent the only two kinds of examples we know:

- *Central cylinders:* If an immersion with a cross-cut is preserved by a line of translations *and* by a central reflection, we call it a *central cylinder*. Central cylinders clearly have CX, since every cross-cut is a translate of one through the center.
- *Tubular quadrics:* When a non-degenerate quadric hypersurface in $\mathbb{R}^{n+1}$ is affinely equivalent to a locus of the form $x_1^2 + x_2^2 + \ldots x_n^2 \pm x_{n+1}^2 = c \in \mathbb{R}$, it will always have compact and transverse, hence ellipsoidal (and thus central) cross-cuts. We call these hypersurfaces *tubular quadrics*. Note that in $\mathbb{R}^3$, all non-degenerate quadrics are tubular.

We suspect that these two classes exhaust all possibilities:

Conjecture 3.5. *A complete immersion $F: M^n \to \mathbb{R}^{n+1}$ with CX must either be a central cylinder, or a tubular quadric.*

In previous papers, we confirmed weakened versions of this conjecture, proving it

- for C^1 hypersurfaces of revolution (SO(n) symmetry) in $\mathbb{R}^{n+1}$ [8], and then, using that result,
- for C^2 surfaces in $\mathbb{R}^3$ whose cross-cuts are *convex* as well as central [9][1].

Here we add another case to this list: roughly, that of a complete surface in $\mathbb{R}^3$ with CX *and* for which some clean cross-cut is a figure-8. To make this precise, we first note that on any complete immersed C^2-surface with CX in $\mathbb{R}^3$, every clean cross-cut is a (clean) central plane C^2-loop. By Proposition 2.16, each of these loops is either regularly homotopic to a figure-8, or has odd rotation index.

The rotation index of a figure-8 is zero, and here (as sketched in our introduction) we verify Conjecture 3.5 for immersions with figure-8 cross-cuts. Since cross-cuts of quadrics cannot be figure-8's, such immersions must be cylindrical:

Theorem 3.6 (Main Result). *If $F: M \to \mathbb{R}^3$ is a complete C^2-immersion with CX, and some plane in general position with F cuts it along a clean figure-8, then $F(M)$ is a central cylinder.*

The figure-8 assumption is decisive for a simple reason: When a plane P cuts a surface with CX transversally along a figure-8 centered at $c \in \mathbb{R}^3$, and we tilt P slightly about c to get nearby cross-cuts, the latter remain centered at c.

1 This has recently been extended to hypersurfaces in $\mathbb{R}^n$ by M. Alper Gur in his Ph.D. thesis *Hypersurfaces with central convex cross-sections* (arXiv:1605.02862 [math.DG]).

Without the figure-8 assumption, this fails. Indeed, consider the unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$. It clearly has CX. Take $u \in \mathbb{S}^2$, $0 < \lambda < 1$, and set $c := \lambda u$. The plane $u_c^\perp$ will cut $\mathbb{S}^2$ along a circle centered at c. For any $v \in \mathbb{S}^2$ near u, however, the cross-cut $v_c^\perp \cap \mathbb{S}^2$ is clearly centered on the line spanned by λv (see Figure 3). So for $v \neq u$, the center moves.

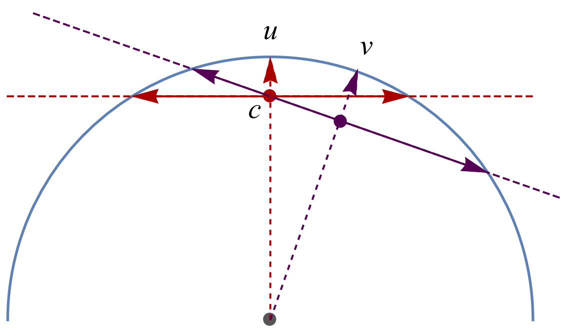

Figure 3: Cross-cuts on a sphere, via $u_c^\perp$ and $v_c^\perp$. Both hyperplanes contain c, but only one of the cross-cuts (red) is centered at c.

The center *cannot* move in this way when cross-cuts are figure-8's, as we make precise in Lemma 3.8 shortly below, using the notion of *central curve* of a good tubular patch:

Definition 3.7 (Central curve). Suppose that Γ is a cross-cut for an immersion $F\colon M \to \mathbb{R}^{n+1}$ relative to a hyperplane $P = u_p^\perp$. Let (U, ψ) denote a good tubular patch for Γ as in Definition 3.2, so that F maps $\psi(\Gamma, h)$ into the hyperplane $u_c^* \equiv h$ for each $h \in [-a, a]$. The *central curve* of the patch is the map $\mu\colon [-a, a] \to \mathbb{R}^{n+1}$ sending any $h \in [-a, a]$ to the centroid $\mu(h)$ of $F(\psi(\Gamma, h))$.

When F is C^k, the central curve of a good tubular patch is clearly C^k too. It is also *immersed*, since Condition (b) from Definition 3.2 yields $u_p^*(\mu(h)) = h$, and hence $\dot{\mu}(h) \cdot u \geq 1$.

We are about to formulate the advantage offered by figure-8 cross-cuts. Notation is as above: $F\colon M^2 \to \mathbb{R}^3$ is a proper C^2-immersion, $u \in \mathbb{S}^2$ and $c \in \mathbb{R}^3$ are fixed. We have a clean cross-cut $\Gamma \subset F^{-1}(u_c^\perp)$, for which $U \subset M$ is a good tubular neighborhood (Definition 3.2), and $\mu\colon [-a, a] \to \mathbb{R}^3$ is its central curve.

Lemma 3.8. *Suppose that F has CX, $F(\Gamma)$ is a figure-8, and $\varepsilon > 0$. If $\Gamma_{h,v} := U \cap F^{-1}(v_{\mu(h)}^\perp)$ is a clean cross-cut, regularly homotopic to Γ whenever $|h| < \varepsilon$ and $v \in \mathbb{S}^2$ with $\varphi(v, u) < \varepsilon$, then $F(\Gamma_{h,v})$ is a figure-8 with central symmetry about $\mu(h)$ for any such h and v.*

Proof. Lemma 3.3 says that for all sufficiently small $|h|$, the cross-cut $\Gamma_{h,u}$ (cut by the plane at signed height h above $u_c^\perp$) is, like $\Gamma_{0,u} = \Gamma$ itself, clean and regularly homotopic to Γ. For simplicity, we can assume this holds for all $|h| \leq a$. (If not, re-define our good tubular patch using a smaller $a > 0$.)

In this case, $F(\Gamma_{h,u})$ is a clean figure-8 for every $|h| \leq a$, and its center, by Proposition 2.16(b), is a simple double-point. The central curve μ of the patch thus consists entirely of simple double-points of F.

In particular, if we fix any $h \in (-\varepsilon, \varepsilon)$, then $F^{-1}(\mu(h)) \cap U$ is a pair $\{x_1, x_2\}$, and as an immersion, F *embeds* disjoint neighborhoods $U_1 \supset x_1$ and $U_2 \supset x_2$ in such a way that, in the ball $B_{h,r}$ centered at $\mu(h)$ with sufficiently small radius $r > 0$, we have

$$F(U) \cap B_{r,h} \subset F(U_1) \cup F(U_2).$$

Further, since F has general position, we can make $r > 0$ small enough to ensure that in $B_{r,h}$, the sheets $F(U_1)$ and $F(U_2)$ meet along a segment of the central curve and nowhere else.

Now, as long as $\varphi(v, u) < \varepsilon$, the nearby cross-cut $\Gamma_{h,v}$ is, by assumption, another clean cross-cut in U, regularly homotopic to $\Gamma = \Gamma_{0,u}$. Immersion preserves regular homotopy, so for all such v the nearby images

$F(\Gamma_{h,v})$, like $F(\Gamma)$, are clean figure-8's — and they are central, since F has CX. We just need to show that they stay centered, like $F(\Gamma_{h,u})$, at $\mu(h)$.

To see that they are, note that the planes $v^{\perp}_{\mu(h)}$ all cut the central curve μ transversally at $\mu(h)$ since the cross-cuts they form are all clean. So by shrinking $r > 0$ further if needed, we can ensure that in $B_{r,h}$, each of these planes cuts the central curve *only* at $\mu(h)$.

It follows that $\mu(h)$ is the *unique* double-point that $F(\Gamma_{h,v})$ has in $B_{r,h}$. Since $\Gamma_{h,v}$ varies smoothly with v, its centroid — and center of symmetry by Corollary 2.14 — varies smoothly too. So for v sufficiently near u, the center of the figure-8 $F(\Gamma_{h,v})$ must stay in $B_{r,h}$. As seen above, however, that center is a simple double-point, and we have just noted that for every v in question, the *only* double-point of $F(\Gamma_{h,v})$ in $B_{r,h}$ is $\mu(h)$. When $\varphi(v,u) < \varepsilon$, the center of $F(\Gamma_{h,v})$ is therefore trapped at $\mu(h)$. As this holds whenever $|h| < \varepsilon$, we have proven the lemma. □

We will prove our main result (Theorem 3.6) by combining this lemma with the Local Axis Lemma below, which shows that when F has CX, and centers of tilted cross-cuts stay (locally) on the central curve as in the lemma above, the central curve is locally straight. Note that it makes no figure-8 assumption. This lemma will quickly yield a local version of the main result in Corollary 3.10.

The construction behind the Local Axis Lemma is easier to describe than to execute, as it requires several estimates to ensure the argument goes through in the small piece of the tube where its hypotheses prevail. To help the reader see past the technical details, here is a sketch:

We have a short tube T whose every (not-too-steep) cross-cut is symmetric about some point on its central curve $\mu = \mu(h)$, and we want to prove this makes μ straight. Any point $p \in T$ that lies on a symmetric cross-cut reflects through the center $\mu(\zeta)$ of that cross-cut to another point $q \in T$. By varying the ζ a bit (and varying the cross-cut through p with it) $q = q(\zeta)$ varies too, tracing out a short arc on T. Done carefully, this forces $\dot{q}(\zeta) = 2\dot{\mu}(\zeta)$, so that the tangent plane to T at q (and symmetrically at p too) contains $\dot{\mu}(\zeta)$. Running the same argument for nearby values of ζ and p, we deduce that *every tangent plane of T near p contains every tangent line of the central curve near $\mu(\zeta)$*. If $\dot{\mu}$ were to span more than one line, those lines would therefore all lie in a single plane, and every tangent plane of T near p would then be parallel to that one plane. This forces T itself to lie, locally, in a plane. But we can take p to be maximally distant from $\mu(\zeta)$ in the original cross-cut. Since the distance to $\mu(z)$ cannot attain an interior max on an open subset of a plane, we get a contradiction unless the central curve to has just one tangent line (locally), making it locally straight, as needed.

To make all this precise, we now write (as earlier) $a > 0$ and $P = u^{\perp}_{c}$ for a fixed (but arbitrary) scalar and plane respectively, with P_a denoting the a-neighborhood of P. We have a clean cross-cut $\Gamma \subset F^{-1}(P)$, and a good tubular patch $\psi\colon \Gamma \times [-a, a] \to U \subset M$ around Γ, so that $F(\partial U) \subset \partial P_a$. Without loss of generality, we assume that $c = \mu(0)$, the initial value of the central curve μ of $F(U)$.

Lemma 3.9 (Local axis lemma). *Suppose that $\varepsilon > 0$, $0 < b < a$ and $F^{-1}(v^{\perp}_{\mu(t)}) \cap U$ is a boundaryless clean cross-cut whose image is central about $\mu(t)$ whenever $\varphi(u,v) < \varepsilon$ and $|t| < b$. Then μ maps $[-b, b]$ to a line segment.*

Proof. We may identify $\Gamma \approx \mathbb{S}^1$, and simplify notation accordingly by using coordinates from the domain of our good tubular patch so that, for instance, $F(\theta, h)$ really means $F(\psi(\theta, h))$.

Fix an arbitrary $\zeta \in (-b, b)$, and choose $\theta_0 \in \mathbb{S}^1$ so that $p_0 := F(\theta_0, \zeta)$ *maximizes* $|F(\theta,\zeta)|^2$ *on* $F(\Gamma, \zeta)$:

$$|p_0|^2 = |F(\theta_0, \zeta)|^2 \geq |F(\theta, \zeta)|^2 \quad \text{for all } \theta \in \mathbb{S}^1.$$

To prove the lemma, we will first need to show that $(\theta_0, \zeta) \subset U$ has a neighborhood with certain favorable attributes. For that, note that $|F(\theta, s) - \mu(h)|$ is continuous on the set of triples $(\theta, s, h) \in \Gamma \times [-a, a]^2$, and that $|p_0 - \mu(\zeta)| = 2r$ for some $r > 0$. So by making $\eta > 0$ small enough, we can ensure two properties:

i) $|\zeta \pm \eta| < b$,
ii) $|\theta - \theta_0|, |s - \zeta|, |h - \zeta| < \eta \Rightarrow |F(\theta, s) - \mu(h)| > r$

Now for any (θ, s, h) in the η-neighborhood of (θ_0, ζ, ζ) defined by (ii) above, consider the unit vector

$$w = w(\theta, s, h) := \frac{F(\theta, s) - \mu(h)}{|F(\theta, s) - \mu(h)|}.$$

Combining (b) from Definition 3.2 with (ii), we then have

$$|u \cdot w| = \frac{|s-h|}{|F(\theta,s)-\mu(h)|} \le \frac{|s-h|}{r}. \tag{3.3}$$

Subtract the w-component from u and normalize to construct a unit vector normal to w:

$$v := \frac{u-(u\cdot w)w}{|u-(u\cdot w)w|}. \tag{3.4}$$

By design, the plane $v^{\perp}_{\mu(h)}$ now contains $F(\theta,s)$. We shall want it to cut $F(U)$ along a central loop, and our hypotheses certify that, provided we can show that $\varphi := \varphi(u,v) < \varepsilon$. To do so, combine (ii) with the triangle inequality to deduce $|s-h| < 2\eta$, and hence

$$\sin^2\varphi = 1-\cos^2\varphi = 1-(u\cdot v)^2 = (u\cdot w)^2 \le \left|\frac{s-h}{r}\right|^2 \le \frac{4\eta^2}{r^2}.$$

Since $\varphi < \varepsilon$ when $\sin\varphi < \sin\varepsilon$, this yields the bound we seek if we require, along with (i) and (ii) above, that

iii) $0 < \eta < \frac{1}{2} r \sin\varepsilon$.

Together, restrictions (i), (ii), and (iii) on $\eta > 0$ now leverage our hypotheses to ensure that for v given by (3.4), the plane $v^{\perp}_{\mu(h)}$ contains both $F(\theta,s)$ and $\mu(h)$, and cuts $F(U)$ along a loop with central symmetry about $\mu(h)$.

We can now make the main geometric argument for our lemma. Consider the mapping that sends (θ,s,h) to the reflection of $F(\theta,s) \in F(U)$ through $\mu(h)$:

$$(\theta,s,h) \longmapsto 2\mu(h) - F(\theta,s) \tag{3.5}$$

Our hypotheses guarantee that for all small enough $|\tau| > 0$, the arc parametrized by $\beta(\tau) := 2\mu(h+\tau) - F(\theta,s)$ stays in $F(U)$. Trivially, its initial velocity is $2\dot\mu(h)$, which cannot vanish because $u^*(\dot\mu(h)) = 1$, by Condition (b) from Definition 3.2. This proves: *If $\eta > 0$ satisfies* (i), (ii), *and* (iii) *above, then for all* (θ,s,h) *with* $|\theta-\theta_0|, |s-\zeta|, |h-\zeta| < \eta$, *the plane tangent to $F(U)$ at $2\mu(h)-F(\theta,s)$ contains $\dot\mu(h) \neq \mathbf{0}$.*

It follows immediately that whenever $|t| < \eta$, each tangent plane to $F(U)$ in a neighborhood of $p_0 = F(\theta_0,\zeta)$ contains both $\dot\mu(\zeta)$ and $\dot\mu(\zeta+t)$. From this, we can deduce constancy of $\dot\mu$ near ζ:

Indeed, we would otherwise have $\dot\mu(\zeta+t) \neq \dot\mu(\zeta)$ for some $t \in (-\eta,\eta)$, and since they have the same u-component, by (b) from Definition 3.2, inequality means independence. Since p_0 has a neighborhood in $F(U)$ where every tangent plane contains — hence is spanned by — these same two non-zero vectors, independence forces constancy of the unit normal to $F(U)$ near p_0. A neighborhood of p_0 in $F(U)$ then lies in a plane — a plane cutting $u^{\perp}_{p_0}$ along a line. The cross-cut parametrized by $F(\ \cdot\ ,\zeta)$ must contain a segment of that line, with $p_0 = F(\theta_0,\zeta)$ in its interior. But we maximized $|F(\theta,\zeta)|^2$ at θ_0, and $x \mapsto |x|^2$ is strictly convex; it *cannot* reach a local max on the interior of a segment. We have thus contradicted the possibility that $\dot\mu(\zeta+t) \neq \dot\mu(\zeta)$ for any $|t| < \eta$. It follows that $\dot\mu \equiv \dot\mu(\zeta)$ on a neighborhood of ζ.

Because $\zeta \in (-b,b)$ was arbitrary, however, this (and continuity of $\dot\mu$) yields local constancy of $\dot\mu$ on subset of $[-b,b]$ that is simultaneously non-empty, open, and closed. The conclusion of our lemma follows at once. □

Corollary 3.10 (Local cylinder). *Under the assumptions of Lemma 3.9, $F(U)$ is a central cylinder.*

Proof. By Lemma 3.9, the central curve of $F(U)$ — what we shall henceforth call its *axis* — is a line segment parallel to $v := \dot\mu(0)$. The corollary follows easily from one additional

Claim. Every tangent plane to $F(U)$ contains v.

As $F(U)$ is closed, it suffices to prove this for an arbitrary point $p \in F(U)$ *not* lying on its axis. Let $h := u^*_c(p)$ denote the signed height of p above $u^{\perp}_c$ ($c = \mu(0)$). The cross-cut of $F(U)$ parallel to $u^{\perp}_c$ at height h is central about $\mu(h)$, so both p and its reflection $q := 2\mu(h) - p$ lie in $F(U) \cap u^{\perp}_{\mu(h)}$. Like p, of course, q avoids the axis of $F(U)$. Our assumptions say that slightly tilted cross-cuts of $F(U)$ are *also* central about the axis, and provided $|t| > 0$ is sufficiently small, some such cross-cut contains both q and $\mu(h+t)$. It follows that the arc

$t \mapsto 2\mu(h+t) - q$ lies in $F(U)$ for all sufficiently small $|t|$. The resulting differentiable arc passes through p when $t = 0$, with initial velocity $2\dot{\mu}(h) = 2v$, so v is tangent to $F(U)$ at p, as our Claim proposes.

The corollary quickly follows: $F(U)$ is everywhere tangent to the constant vectorfield v, so it is foliated by line segments parallel to v. This makes it a (generalized) cylinder, and it has a compact central cross-cut, so it is central too. □

By chaining together intermediate results from above, we can quickly prove our main theorem. We restate it here for the reader's convenience.

Theorem 3.6. *If $F\colon M \to \mathbb{R}^3$ is a complete C^2-immersion with* CX, *and some plane in general position with F cuts it along a clean figure-8, then $F(M)$ is a central cylinder.*

Proof of the Main Result. We are assuming that for some plane $P = u_c^\perp$ in general position relative to F, a clean cross-cut $\Gamma \subset F^{-1}(P)$ maps to a clean figure-8 $\gamma := F(\Gamma) \subset P$. As discussed in connection with Definition 3.2, that puts Γ in the image of a good tubular patch (U, ψ).

Lemma 3.3 now provides some $0 < \varepsilon < a$ for which every cross-cut $F^{-1}(v_{\mu(h)}^\perp) \cap U$ is a clean figure-8 when $|h| < \varepsilon$ and $\varphi(v, u) < \varepsilon$. As above, $\mu\colon [-a, a] \to \mathbb{R}^3$ here denotes the central curve of $F(U)$. Lemma 3.8 now certifies that each of these cross-cuts has central symmetry about $\mu(h)$, and Corollary 3.10 (with $b = \varepsilon$) then shows that, within the ε-neighborhood P_ε of P, the image $F(U)$ is a central cylinder.

A simple open/closed argument now shows that $F(M)$ is the complete extension of that cylinder. Indeed, call a scalar $a > 0$ *reachable* if there exists a good tubular patch $\psi\colon \Gamma \times [-a, a] \to M$ whose image is mapped by F to a central cylinder. Given what we have just proven, we know that

$$a^* := \sup\{a > 0\colon a \text{ is reachable}\} \geq \varepsilon > 0.$$

Our theorem amounts to the assertion $a^* = \infty$, which we can now establish by contradiction. For if $a^* < \infty$, the completeness and smoothness of F would let us construct a *maximal* good tubular patch $\psi\colon \Gamma \times [-a^*, a^*] \to M$, with $F \circ \psi$ mapping $\Gamma \times [-a^*, a^*]$ to a central cylinder with boundary in ∂P_{a^*}. The two loops bounding this cylinder would clearly be clean figure-8's. By applying the argument above, however, we could deduce that their preimages in M each have good tubular neighborhoods mapping to central cylinders via F. Our supposedly maximal good tubular patch could then be extended slightly at each boundary component, violating the maximality of a^*. Thus, a^* *cannot* be finite; we have $a^* = \infty$ which gives our theorem. □

References

[1] W. Blaschke, Über affine Geometrie XXII: Bestimmung der Flächen mit zentrischen ebenen Schnitten. *Leipz. Ber.* **70** (1919), 336–340. JFM 47.0661.01 (Also in: W. Blaschke, *Gesammelte Werke. Band 4: Affine Differentialgeometrie: Differentialgeometrie der Kreis- und Kugelgruppen*. Thales-Verlag, Essen, 1985)

[2] W. Blaschke, *Vorlesungen über Differentialgeometrie und geometrische Grundlagen von Einsteins Relativitätstheorie. II*. Springer 1923. JFM 49.0499.01

[3] G. R. Burton, Some characterisations of the ellipsoid. *Israel J. Math.* **28** (1977), 339–349. MR0513965 Zbl 0368.52004

[4] M. Ghomi, B. Solomon, Skew loops and quadric surfaces. *Comment. Math. Helv.* **77** (2002), 767–782. MR1949113 Zbl 1029.53002

[5] M. W. Hirsch, *Differential topology*. Springer 1976. MR0448362 Zbl 0356.57001

[6] W. Klingenberg, *A course in differential geometry*. Springer 1978. MR0474045 Zbl 0366.53001

[7] S. Olovjanischnikoff, Über eine kennzeichnende Eigenschaft des Ellipsoides. (Russian with German summary) *Uchenye Zapiski Leningrad State Univ., Math. Ser.* **83(12)** (1941), 114–128. MR0017553 Zbl 0061.37709

[8] B. Solomon, Central cross-sections make surfaces of revolution quadric. *Amer. Math. Monthly* **116** (2009), 351–355. MR2503320 Zbl 1229.53007

[9] B. Solomon, Surfaces with central convex cross-sections. *Comment. Math. Helv.* **87** (2012), 243–270. MR2914849 Zbl 1255.53007

[10] H. Whitney, On regular closed curves in the plane. *Compositio Math.* **4** (1937), 276–284. MR1556973 Zbl 0016.13804 JFM 63.0647.01