The Second Hull of a Knotted Curve

Jason Cantarella,[1,*] Greg Kuperberg,[2,†] Robert B. Kusner,[3,‡] and John M. Sullivan[4,§]

[1]*Department of Mathematics, University of Georgia, Athens, GA 30602*
[2]*Department of Mathematics, University of California, Davis, CA 95616*
[3]*Department of Mathematics, University of Massachusetts, Amherst, MA 01003*
[4]*Department of Mathematics, University of Illinois, Urbana, IL 61801*

(Dated: May 5, 2000; revised April 4, 2002)

The convex hull of a set K in space consists of points which are, in a certain sense, "surrounded" by K. When K is a closed curve, we define its higher hulls, consisting of points which are "multiply surrounded" by the curve. Our main theorem shows that if a curve is knotted then it has a nonempty second hull. This provides a new proof of the Fáry/Milnor theorem that every knotted curve has total curvature at least 4π.

A space curve must loop around at least twice to become knotted. This intuitive idea was captured in the celebrated Fáry/Milnor theorem, which says the total curvature of a knotted curve K is at least 4π. We prove a stronger result: that there is a certain region of space doubly enclosed, in a precise sense, by K. We call this region the second hull of K.

The convex hull of a connected set K is characterized by the fact that every plane through a point in the hull must interesect K. If K is a closed curve, then a generic plane intersects K an even number of times, so the convex hull is the set of points through which every plane cuts K twice. Extending this, we introduce a definition which measures the degree to which points are enclosed by K: a point p is said to be in the n^{th} *hull* of K if every plane through p intersects K at least $2n$ times.

A space curve is convex and planar if and only if it has the *two-piece property*: every plane cuts it into at most two pieces. This has been generalized by Kuiper (among others; see [CC97]) in the study of tight and taut submanifolds. It was Milnor [Mil50] who first observed that for a knotted curve, there are planes in every direction which cut it four times.

Our main theorem strengthens this to find points through which every plane cuts the knot four times. More precisely, it says that any curve K with a knotted inscribed polygon (in particular any knotted curved of finite total curvature) has a nonempty second hull, as shown for a trefoil in Figure 1. Furthermore, in a certain sense, this hull must extend across the knot: Using quadrisecants, we get an alternate approach to second hulls, and in particular we show that the second hull of a generic prime knot K is not contained in any embedded normal tube around K.

We use some ideas from integral geometry to show that the total curvature of a space curve with nonempty n^{th} hull is at least $2\pi n$. Thus, our theorem provides an alternate proof of the Fáry/Milnor theorem.

We do not yet know any topological knot type which is guaranteed to have a nonempty third hull, though we suspect that this is true for the (3,4)–torus knot (see Figure 5). It would be interesting to find, for $n \geq 3$, computable topological invariants of knots which imply the existence of nonempty n^{th} hulls.

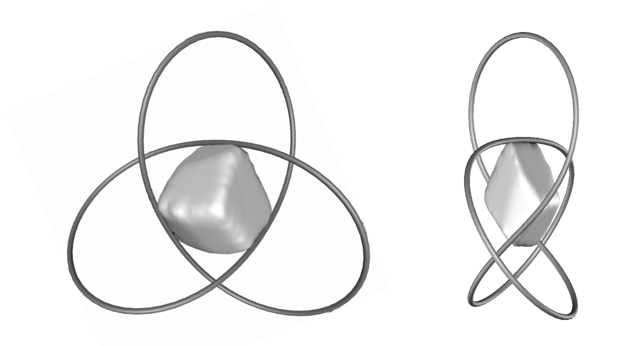

Figure 1: This figure shows the second hull of a trefoil knot, the shaded region in the center of the knot, from two perspectives. Every plane which intersects second hull must cut the knot at least four times. The region shown is only a numerical approximation to the true second hull of this knot; we expect the true second hull to be smoother. (The trefoil shown here is a critical point—and presumed minimizer—for the Möbius-invariant knot energy of Freedman, He and Wang [FHW94].)

Our definition of higher hulls is appropriate only for curves, that is, for one-dimensional submanifolds of $\mathbb{R}^d$. Although we can imagine similar definitions for higher-dimensional submanifolds, we have had no reason yet to investigate these. For 0-dimensional submanifolds, that is, finite point sets $X \subset \mathbb{R}^d$, Cole, Sharir and Yap [CSY87] introduced the notion of the k-*hull* of X, which is the result of removing from $\mathbb{R}^d$ all half-spaces containing fewer than k points of X. Thus it is the set of points p for which every closed half-space with p in the boundary contains at least k points of X. Although we did not know of their work when we first formulated our definition, our higher hulls for curves are similar in spirit.

1. DEFINITIONS AND CONVENTIONS

By a closed curve in $\mathbb{R}^d$, we mean a continuous map of a circle, up to reparametrization. Subject to the conventions introduced below about the counting of intersections, it is easy to define the n^{th} hull.

*Email: cantarel@math.uga.edu
†Email: greg@math.ucdavis.edu
‡Email: kusner@math.umass.edu
§Email: jms@math.uiuc.edu

Definition. If K is an oriented closed curve in $\mathbb{R}^d$, its n^{th} *hull* $h_n(K)$ is the set of points $p \in \mathbb{R}^d$ such that K cuts every hyperplane through p at least n times in each direction.

In this definition, we basically count the intersections of K with a plane P by counting the number of components of the pullback of $K \cap P$ to the domain circle. If the intersections are transverse, then (thinking of P as horizontal) we find equal numbers of upward and downward intersections; the total number of intersections (if nonzero) equals the number of components of $K \smallsetminus P$ (again counted in the domain).

To handle nontransverse intersections, we adopt the following conventions. First, if $K \subset P$, we say K cuts P once in either direction. If $K \cap P$ has infinitely many components, then we say K cuts P infinitely often. Otherwise, each connected component of the intersection is preceded and followed by open arcs in K, with each lying to one side of P. An *upward* intersection will mean a component of $K \cap P$ preceded by an arc below P *or* followed by an arc above P. (Similarly, a *downward* intersection will mean a component preceded by an arc above P or followed by an arc below P.) A glancing intersection, preceded and followed by arcs on the same side of P, thus counts twice, as both an upward and a downward intersection.

With these conventions, it is easy to see that an arc whose endpoints are not in a plane P intersects P an even number of times if the endpoints are on the same side of P and an odd number of times otherwise. For a closed curve, there are always equal numbers of intersections in each direction.

We will later apply the same definition to disconnected curves L, like links. In that context, of course, the number of intersections of L with P cannot simply be counted as the number of components of $L \smallsetminus P$.

2. THE SECOND HULL OF A KNOTTED CURVE

In order to prove that any knot has a second hull, we start with some basic lemmas about n^{th} hulls.

Lemma 1. *For any curve K, each connected component of its n^{th} hull $h_n(K)$ is convex. Furthermore, $h_n(K) \subset h_{n-1}(K)$.*

Proof. Let C be a connected component of $h_n(K)$. If q is a point in its convex hull, then by connectedness, any plane P through q must meet C at some point p. Thus P is also a plane through $p \in C \subset h_n(K)$, so P meets K at least $2n$ times. By definition, then, $q \in h_n(K)$. The second part of the statement then follows from the definition of h_n. □

A useful lemma allows us to simplify a given curve, with control on its higher hulls.

Lemma 2. *Let K^* be the result of replacing some subarc Γ of K by the straight segment S connecting its endpoints p, q. No plane cuts K^* more often than it cuts K, so $h_n(K^*) \subset h_n(K)$.*

Proof. Let P be any plane. If $K^* \cap P$ has infinitely many components then so does $K \cap P$. So we may assume $K^* \cap P$ has only a finite number of components. The generic case is when $p, q \notin P$. If they are on the same side of P, then S misses P and there is nothing to prove. Otherwise, S has a unique intersection with P, but Γ must also have cut P at least once.

The other cases require us to invoke our conventions for counting intersections, and thus require a bit more care. First suppose $p, q \in P$. If $K^* \subset P$, this counts as two intersections, but by our conventions K cuts P a nonzero even number of times, hence at least twice. Otherwise, let A be the component of $K^* \cap P$ containing S. It counts as one or two intersections with P, depending on whether points just before and after A are on different sides of P or the same side. But the corresponding arc of K then cuts P a nonzero odd or even number of times respectively, hence at least as often as A does.

Next, suppose $p \in P$ and q is above P. Let A be the component of $K \cap P$ containing p, and consider the arc of K from a point r just before A to q. This arc cuts P a nonzero odd or even number of times if r is below or above P, but then the corresponding arc in K^* intersects P once or twice, respectively.

In any case, K^* has at most as many intersections as K did with the arbitrary plane P. The final statement of the lemma then follows from the definition of h_n. □

We will need to locate regions of space which can be said to contain the knottedness of K.

Definition. We say that an (open) half-space H *essentially contains* a knot K if it contains K, except for possibly a single point or a single unknotted arc.

Note that if H essentially contains K, then K cuts ∂H at most twice. Thus, we can *clip K to H*, replacing the single arc (if any) outside H by the segment in ∂H connecting its endpoints. The resulting knot $\mathrm{cl}_H(K)$ is isotopic to K, and is contained in $\overline{H}$.

Our key lemma now shows that essential half-spaces remain essential under clipping:

Lemma 3. *Let $K^* := cl_H(K)$ be the clipping of a knot K to some half-space H which essentially contains it. Then any half-space H' which essentially contains K will also essentially contain K^*.*

Proof. We must show that either H' contains K^*, or the single arc of this new knot outside H' is unknotted. If H' contains K then, since K^* lies within the convex hull of K, H' must contain K^*, and we are done.

Otherwise, consider the single arc of K outside H', and count how many of its endpoints are outside H. The different cases are illustrated in Figure 2. If both endpoints are outside H, again we find H' contains K^*. If neither is, the unknotted arc of K outside H' has an unknotted subarc outside H, which gets replaced with an isotopic segment in ∂H. So the arc of K^* outside H' is unknotted, as desired.

Finally, we must consider the case where one endpoint of the arc is within H and one is outside H. This implies that the half-spaces are not nested, but divide space into four quadrants, with K winding once around their intersection line. The

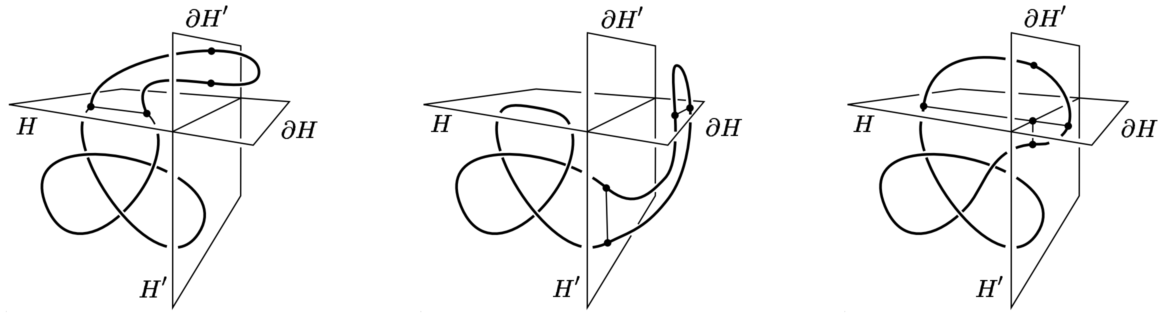

Figure 2: This figure shows the three cases of interest in the proof of Lemma 3. On the left, we see the case where both intersections of K with $\partial H'$ occur outside H. In the center, we see the case where both intersections of K with $\partial H'$ occur *inside* H. On the right, we see the case where K intersects $\partial H'$ once inside and once outside H. In each figure, the segments in ∂H and $\partial H'$ replacing unknotted subarcs of K when K is clipped to these planes are drawn in lightly. The half-space H is below the horizontal plane ∂H, while the half-space H' is to the left of the vertical plane $\partial H'$.

knot K is a connect sum of four arcs, in these quadrants. The sum of the two arcs outside H is trivial, as is the sum of the two outside H'. So in fact all the knotting happens in the one quadrant $H \cap H'$. But now, the arc of K^* outside H' consists of the original unknotted arc in one quadrant, plus a straight segment within ∂H; thus it is unknotted. ☐

For the proof of our theorem, we will need to decompose a knot into its prime pieces, but only when this can be done with a flat plane.

Definition. We say that a knot K in $\mathbb{R}^3$ is *geometrically composite* if there is a plane P (cutting K twice) which decomposes K into two nontrivial connect summands. Otherwise, we say K is *geometrically prime*. Note that a (topologically) prime knot is necessarily geometrically prime.

Main Theorem. *If K is any space curve with a knotted inscribed polygon, then its second hull $h_2(K)$ is nonempty.*

Proof. We can replace K by the knotted inscribed polygon, and by Lemma 2 its second hull can only shrink. If K is now geometrically composite, we can replace it by one of its two summands, and by Lemma 2 again its second hull will only shrink. Because the polygonal K is tame, we need only do this a finite number of times, and so we may assume that K is geometrically prime.

Let $A := \bigcap_H \overline{H}$ be the intersection of the closures of all the half-spaces H essentially containing K. We claim that any plane P through a point p in A cuts K at least four times. If not, P would cut K twice. Because K is geometrically prime, P cannot split K into two nontrivial summands. So P must be the boundary of an essential half-space H. We will construct a parallel half-space H' strictly contained in H which still essentially contains K. Since $p \notin \overline{H'}$ and $p \in A$, this is a contradiction.

If K is not contained in H, there are two edges of K which start in H but end outside H (possibly on P). Let v and w be their endpoints in H, and Γ be the subarc of K between v and w in H. If $K \subset H$, let $\Gamma = K$. In either case, Γ is a compact subset of H, and so is contained in a smaller parallel half-space H' which essentially contains K. This completes the proof of the claim.

Since planes through A cut K four times, $A \subset h_2(K) \subset h_1(K)$, and we can write

$$A = \bigcap_H \big(\overline{H} \cap h_1(K)\big).$$

This exhibits A as an infinite intersection of compact sets. To show it is nonempty, it suffices to prove that any finite intersection is nonempty (using the finite intersection property for compact sets). So suppose $H_1, \ldots, H_N$ are half-spaces essentially containing K. Clipping the knot successively to these half-spaces, we define $K_0 = K$ and $K_i = \mathrm{cl}_{H_i}(K_{i-1})$. Since K is essentially contained in every H_j, this is also true inductively for each of the K_i, using Lemma 3. Thus K_N, a knot isotopic to K, is contained in

$$\bigcap_1^N \overline{H_i} \cap h_1(K),$$

showing this to be nonempty, as desired. ☐

Note that our sets $\overline{H} \cap h_1(K)$ are convex. Therefore, we could appeal to Helly's theorem instead of the finite intersection property: as soon as any four of these compact convex sets have nonempty intersection, they all do. But the proof given above is no simpler for the special case $N = 4$.

We expect that the set A used in the proof is always a connected component of $h_2(K)$. This essential piece $A \subset h_2(K)$ is the intersection of all halfspaces essentially containing K, just as $h_1(K)$ is the intersection of all halfspaces containing K. We note also that the only place we used the assumption about polygons was to show that bounding planes of A must cut the knot four times, and thus that all of A, and not merely its (possibly empty) interior, is in $h_2(K)$.

Using Milnor's definition of total curvature for arbitrary curves, our theorem applies to any knot of finite total curvature:

Corollary. *A knotted space curve of finite total curvature has nonempty second hull.*

Proof. A lemma of Milnor [Mil50] shows that any curve of finite total curvature has an isotopic inscribed polygon, so our main theorem applies. ☐

3. LINKS

If we apply our definition of n^{th} hull to links L, that is, to unions of closed curves in space, then it is no longer necessarily true that $h_1(L)$ is the entire convex hull of L. However, Lemma 1 shows that $h_1(L)$ contains the convex hull of each component. In Figure 3 we show an example where both inclusions are strict.

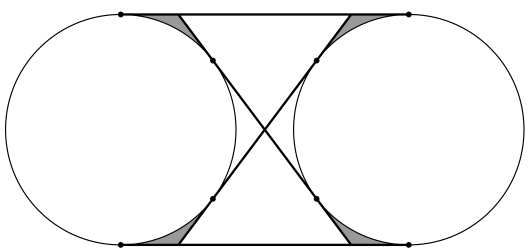

Figure 3: This figure shows two round circles forming an unlink L in the plane. The convex hull of L is bounded by a stadium curve. Its first hull $h_1(L)$, however, consists only of the hulls of the two components along with the four shaded triangular regions. These are bounded by the four lines tangent to both circles.

For links we can easily find points in higher hulls by intersecting the hulls of different components. The next lemma is immediate from the definitions:

Lemma 4. *If A and B are links, then*

$$h_m(A) \cap h_n(B) \subset h_{m+n}(A \cup B)$$

for all $m, n > 0$. ☐

This allows us to extend our theorem to the case of links.

Corollary. *Any nontrivial link (of finite total curvature) has a nonempty second hull.*

Proof. If the link contains a knotted component, we use the theorem. Otherwise it must contain two components A and B whose convex hulls intersect, and we use the last lemma. ☐

For a Hopf link whose two components are plane curves, Lemma 4 is sharp: the second hull consists exactly of the segment along which the convex hulls of the components intersect. This shows that the second hull of a link can have zero volume. We conjecture that this never happens for knots.

4. CONE ANGLES

We originally became interested in second hulls while attempting to improve lower bounds on the ropelength of knots [CKS02]. We wanted to find points in space from which a knot has large cone angle, or visual angle.

Lemma 5. *If $p \in h_n(K)$ then the cone from p to K has cone angle at least $2\pi n$.*

Proof. Consider the radial projection K' of K to the unit sphere around p. By the definition of h_n, K' intersects each great circle at least $2n$ times. But the length of a spherical curve, using integral geometry as in the proof of Fenchel's theorem, equals π times its average number of intersections with great circles (see [Mil50]). Thus the length of K', which is the cone angle of K at p, is at least $2\pi n$. ☐

Corollary. *If K is a nontrivial knot or link (of finite total curvature), there is some point p from which K has cone angle at least 4π.* ☐

For our application to ropelength, however, we needed to know more. If K has thickness τ, that is, if a normal tube of radius τ around K is embedded, then the cone point p can be chosen outside this tube. We will see below, using quadrisecants, that for generic prime knots, the second hull does extend outside this thick tube, as desired.

However, a different argument gives 4π cone points for arbitrary links: Brian White has a version of the monotonicity theorem for minimal surfaces with boundary [Whi96]. As he points out in work with Ekholm and Wienholtz [EWW02, Thm. 1.3], it shows directly that if we span K with a minimal surface, and p is a point where this surface intersects itself, then the cone angle from p to K is at least 4π. In fact, this result was originally proved in 1983 by Gromov [Gro83, Thm. 8.2.A].

The set of points from which the cone angle to a curve K is at least 2π has been called the *visual hull*, and is a superset of the convex hull. (See [Gro83, Gag80].) We might call the set where cone angle is at least 4π the *second visual hull*; as we have mentioned, this set contains our second hull. Gromov's result, reproved by our corollary above, is that any knot has a nonempty second visual hull.

We have shown [CKS02] that if K is a thick knot, then the self-intersection curve of a minimal spanning disk must go outside the tube of radius τ, as shown in Figure 4, giving us what we needed: a point far from the knot with large cone angle. We note that self-intersection points of a minimal surface need not be in the second hull of its boundary curve, as demonstrated by the second example in Figure 4, though they are necessarily in the second visual hull.

5. BRIDGE NUMBER AND TOTAL CURVATURE

We can define the *hull number* of a link to be the largest n for which the n^{th} hull is nonempty (and the hull number of a link type to be the minimum over all representatives). Perhaps sufficiently complicated link types have hull number greater than two, but we know of no way to prove this. We can get an easy upper bound:

Proposition. *The hull number of any link is bounded above by its bridge number.*

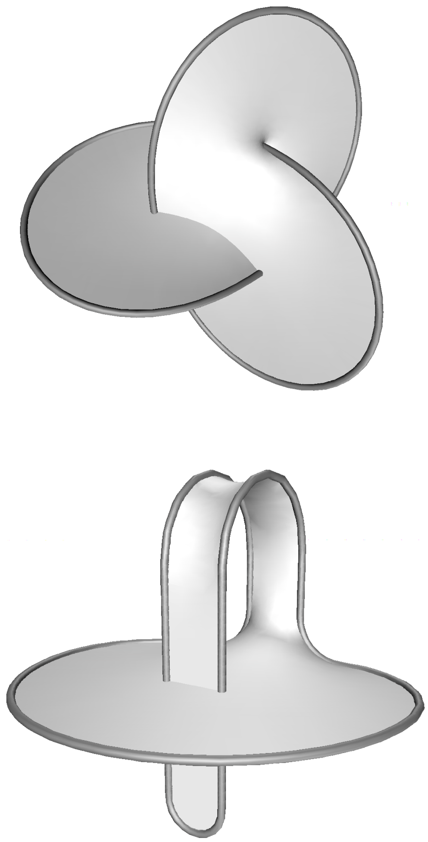

Figure 4: This figure shows two immersed minimal disks, one bounded by a trefoil knot and the other bounded by an unknot. By White's version of the monotonicity theorem for minimal surfaces, the cone angle to the boundary from any point on the line of self-intersection of either disk is at least 4π. In the trefoil above, this line is contained within the second hull. For the unknotted curve below, the second hull is empty, and so the line of self-intersections is in the second visual hull, but not in the second hull.

Proof. To say L is a link of bridge number n means there is some height function with just n minima and n maxima. Thus L will have empty $(n+1)^{\text{st}}$ hull, because it cuts no horizontal plane more than $2n$ times. □

Note that we don't always have equality, at least for composite links. For instance, the connect sum of two 2-bridge links (like Hopf links or trefoil knots) has bridge number three. But obvious configurations, where each summand is in its 2-bridge presentation and they are well separated vertically, have empty third hull. Perhaps any prime 3-bridge knot has nonempty third hull as in Figure 5: certainly there are planes in every direction which are cut six times by the knot. To extend our proof to this case, we might need to consider the notion of thin position for knots, as introduced by Gabai [Gab87] and studied by Thompson [Tho97].

Milnor's version [Mil50] of the Fáry/Milnor theorem says that the total curvature of a knot is greater than 2π times its bridge number; by our proposition, the same is true replacing bridge number by hull number. Since our main theorem says

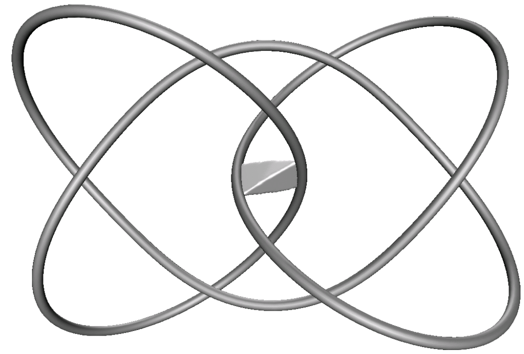

Figure 5: This presumed Möbius-energy-minimizing configuration (from [KS97]) of the $(3,4)$-torus knot 8_{19}, has a small but nonempty third hull. We do not know whether every configuration of 8_{19} has nonempty third hull. This knot and the button knot 8_{18} are the first prime 3-bridge knots in Rolfsen's table [Rol76]. As in Figure 1, we show a numerical approximation of the third hull, and we expect that the true hull is somewhat smoother.

hull number is at least two for a nontrivial knot, this provides an alternate proof of Fáry/Milnor.

The following lemma, which also appears in [EWW02], when combined with Lemma 5, gives yet another proof of Fáry/Milnor.

Lemma 6. *The cone angle of a knot K, from any point, is at most the total curvature of K.*

Proof. Applying Gauss-Bonnet to the cone itself, we see that the cone angle equals the total geodesic curvature of K in the cone, which is no greater than the total curvature of K in space. □

Fáry's proof [Fár49] of the Fáry/Milnor theorem showed that the total curvature of a space curve is the average total curvature of its planar projections, and that any knot projection has total curvature at least 4π. In fact, his proof of the latter statement actually showed that any knot projection has nonempty second hull.

When we first started to think about second hulls, we rederived this result, not knowing Fáry's proof. We hoped that perhaps a curve whose projections all had nonempty second hulls would necessarily also have one, but never succeeded in this line of argument. The next lemma does give a sort of converse, relating hulls and projections.

Lemma 7. *If K is a closed curve in $\mathbb{R}^d$, and Π is any orthogonal projection to $\mathbb{R}^{d-1}$, then $\Pi(h_n(K)) \subset h_n(\Pi(K))$.*

Proof. Call the projection direction "vertical". If P is a hyperplane in $\mathbb{R}^{d-1}$ through $p \in \Pi(h_n(K))$, then the vertical hyperplane $\Pi^{-1}P$ in $\mathbb{R}^d$ passes through a point in $h_n(K)$, so K cuts it $2n$ times. Thus $\Pi(K)$ cuts P at least $2n$ times. □

This lemma explains why, for instance, the projected third hull in Figure 5 must lie entirely within the central bigon of the knot projection.

6. QUADRISECANTS AND SECOND HULLS

Any knot or link in space has a quadrisecant, that is, a straight line which intersects the link four times. This was first proved by Pannwitz [Pan33] for generic polygonal links; the full result is due to Kuperberg [Kup94]. Quadrisecants can be used to prove a thickened version of our theorem. In particular, the middle segment of a quadrisecant for K often lies in the second hull of K.

In the case of links, if there are two components A and B with nonzero linking number, Pannwitz [Pan33] (and later Morton and Mond [MM82]) showed that there is an $ABAB$ quadrisecant, meaning one where the components are seen in that order along the secant line. The entire mid-segment then lies in the intersection of the convex hulls of A and B, and thus in $h_2(A \cup B)$. For the thickened version, simply note that this mid-segment starts inside the tube around A and ends inside the disjoint one around B, so it must leave the tubes altogether in the middle.

For knots, we can compare the linear ordering of the four intersection points (along the quadrisecant line) with their circular ordering (along the knot). Viewed from a point on the mid-segment of a quadrisecant, the four intersection points lie (two each) at the north and south poles of the visual sphere. There are two possibilities, which we call $NNSS$ and $NSNS$. The mid-segment of an $NSNS$ quadrisecant for K will be contained in the second hull of K: if we project K to the unit sphere around a point on the mid-segment, the projected curve cuts every great circle four times, since it visits the poles in the order $NSNS$. Thus the following conjecture would give an alternate proof of our main theorem.

Conjecture. *Any nontrivial knot has an $NSNS$ quadrisecant.*

Elizabeth Denne has already made good progress towards proving this conjecture as part of her doctoral dissertation at the University of Illinois. She makes use of ideas from Pannwitz [Pan33], Kuperberg [Kup94] and Schmitz [Sch98].

For generic polynomial curves in prime knot classes, we can prove a thickened version of our theorem using Kuperberg's notion of an essential (or "topologically nontrivial") quadrisecant [Kup94].

Definition. A secant S of K is *trivial* if its endpoints are on the same component of K, dividing it into arcs X and Y, and if one of the two circles $S \cup X$, $S \cup Y$ spans a disk whose interior avoids K. The disk may intersect S and itself. A quadrisecant is *essential* if neither the mid-segment nor either end-segment is trivial.

We start by proving a lemma about secants of unknots.

Lemma 8. *Let K be an unknot. Any secant of K is trivial on either side.*

Proof. Suppose that S is a secant of K. Pick either one of the arcs of K joining the endpoints of S. We will show that the closed curve C formed by joining this arc to S bounds a disk disjoint from K. We begin the disk by constructing an embedded ribbon bounded by C and a parallel curve, C', chosen so that the linking number of K and C' is zero. By the definition of linking number, this implies that C' is homologous to zero in the complement of K. Since K is the unknot, C' is therefore *homotopic* to zero in the complement of K. Thus, C' bounds a disk disjoint from K. Joining this disk to the ribbon completes the proof. □

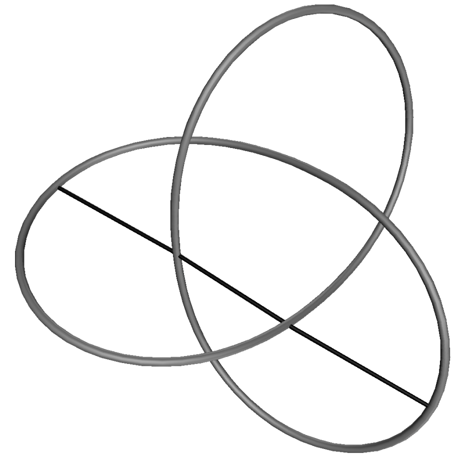

Figure 6: This figure shows an essential quadrisecant of a particular configuration of a trefoil knot. By Lemma 9, the mid-segment of this quadrisecant is contained within the second hull of K. In fact, this quadrisecant is $NSNS$ as well, so our earlier argument gives an alternate proof that the mid-segment is in the second hull of K. Since this is the same configuration of trefoil knot used in Figure 1, it is interesting to compare this picture with the image of the entire second hull.

Lemma 9. *Let K be a geometrically prime knot. The mid-segment of any essential quadrisecant for K is contained in the second hull of K.*

Proof. Let Q be a quadrisecant of K; if its mid-segment is not contained in the second hull, there is a plane P that meets the mid-segment of Q and meets K only twice. Because K is geometrically prime, it is trivial on one side of P. Let K^* be the unknotted arc on this side, completed with a straight line segment in P. By Lemma 8, the end-segment S of Q on this side is trivial with respect to K^*. That is, S together with one arc of K^* bounds a disk D whose boundary is disjoint from P.

By design, D does not intersect K^*, but it may cross P and intersect $K \smallsetminus K^*$. We claim that we can simplify D so that it does not cross P. Thus the secant S of Q is inessential with respect to K.

To check the claim, consider D as a map from the standard disk to $\mathbb{R}^3$. In general position, the inverse image of P is a collection of circles, and we can cut off the disk at the outermost circles. We know D is interior-disjoint from all of K^*, including the line segment lying in P. But D may intersect P in a complicated loop L that encircles the pair of points $K \cap P$ many times. In this case, after cutting D along this circle, we attach not a flat disk in P but instead a bigger disk that avoids K. □

Lemma 10. *The mid-segment of an essential quadrisecant for K leaves any embedded normal tube around K.*

Proof. If not, it would be isotopic within the tube to part of the core curve, that is, to one arc of K, contradicting the definition of essential. □

The result of [Kup94], that a generic polynomial link has an essential quadrisecant, combined with the last two lemmas, proves the thickened version of our theorem, for generic prime knots:

Theorem. *If K is a generic polynomial representative of a prime knot class, and T is an open embedded normal tube around K, then $h_2(K) \smallsetminus T$ is nonempty.* □

7. OTHER AMBIENT SPACES

We have been assuming that our knots lie in Euclidean space $\mathbb{R}^3$. However, our notion of second hull is not a metric notion but rather a projective notion, depending only on incidence relations of planes.

It follows that our definition applies in hyperbolic space $\mathbb{H}^3$, and that any knotted curve in $\mathbb{H}^3$ has a nonempty second hull: Simply embed hyperbolic space, in Klein's projective model, in a ball in $\mathbb{R}^3$. Then the hyperbolic second hull of any curve is the Euclidean second hull of its image in the model.

Recently, two proofs have been given showing that the Fáry/Milnor theorem holds for knots in any Hadamard manifold, that is, any simply connected manifold of nonpositive curvature. Alexander and Bishop [AB98] find a sequence of inscribed polygons limiting to a quadruply covered segment, while Schmitz [Sch98] comes close to constructing a quadrisecant. Arguments like either of these could be used to give alternate proofs of our theorem in Euclidean space.

We thus suspect that there should be some notion of second hull for curves in Hadamard manifolds. The problem is that there is no natural way to extend our definition, because there is no analog of a plane in this general context. If some reasonable definition for second hull can be found so that Lemma 2 remains true, then there should be no problem proving that knots have nonempty second hulls.

Acknowledgments

We gratefully acknowledge helpful conversations with many colleagues, including Colin Adams, Stephanie Alexander, Dick Bishop, Elizabeth Denne, Herbert Edelsbrunner, Mike Gage and Frank Morgan. This work has been partially supported by the NSF through grants DMS-99-02397 (to Cantarella), DMS-00-72342 (to Kuperberg), DMS-00-76085 (to Kusner) and DMS-00-71520 (to Sullivan).

[AB98] Stephanie B. Alexander and Richard L. Bishop, *The Fáry–Milnor theorem in Hadamard manifolds*, Proc. Amer. Math. Soc. **126**:11 (1998), 3427–3436.

[CC97] Thomas E. Cecil and Shiing-shen Chern (eds.), *Tight and taut submanifolds*, Cambridge University Press, Cambridge, 1997.

[CKS02] Jason Cantarella, Robert B. Kusner, and John M. Sullivan, *On the minimum ropelength of knots and links*, Invent. Math. (2002), to appear, arXiv:math.GT/0103224.

[CSY87] Richard Cole, Micha Sharir, and Chee-K. Yap, *On k-hulls and related problems*, SIAM J. Comput. **16**:1 (1987), 61–77.

[EWW02] Tobias Ekholm, Brian White, and Daniel Wienholtz, *Embeddedness of minimal surfaces with total boundary curvature at most 4π*, Ann. of Math. **155**:1 (2002), to appear.

[Fár49] István Fáry, *Sur la courbure totale d'une courbe gauche faisant un nœud*, Bull. Soc. Math. France **77** (1949), 128–138.

[FHW94] Michael H. Freedman, Zheng-Xu He, and Zhenghan Wang, *On the Möbius energy of knots and unknots*, Ann. of Math. **139**:1 (1994), 1–50.

[Gab87] David Gabai, *Foliations and the topology of 3-manifolds, III*, J. Differential Geom. **26**:3 (1987), 479–536.

[Gag80] Michael E. Gage, *A proof of Gehring's linked spheres conjecture*, Duke Math. J. **47**:3 (1980), 615–620.

[Gro83] Mikhael Gromov, *Filling Riemannian manifolds*, J. Diff. Geom. **18** (1983), 1–147.

[KS97] Robert B. Kusner and John M. Sullivan, *Möbius energies for knots and links, surfaces and submanifolds*, Geometric Topology (Willam H. Kazez, ed.), Amer. Math. Soc./Internat. Press, 1997, pp. 570–604, Proceedings of the Georgia International Topology Conference, August 1993.

[Kup94] Greg Kuperberg, *Quadrisecants of knots and links*, Journal of Knot Theory and its Ramifications **3**:1 (1994), 41–50, arXiv:math.GT/9712205.

[Mil50] John W. Milnor, *On the total curvature of knots*, Ann. of Math. **52**:2 (1950), 248–257.

[MM82] Hugh R. Morton and David M. Q. Mond, *Closed curves with no quadrisecants*, Topology **21**:3 (1982), 235–243.

[Pan33] Erika Pannwitz, *Eine elementargeometrishe Eigenschaft von Verschlingungen und Knoten*, Math. Ann. **108** (1933), 629–672.

[Rol76] Dale Rolfsen, *Knots and links*, Publish or Perish, Wilmington, 1976.

[Sch98] Carsten Schmitz, *The theorem of Fáry and Milnor for Hadamard manifolds*, Geom. Dedicata **71**:1 (1998), 83–90.

[Tho97] Abigail Thompson, *Thin position and bridge number for knots in the 3-sphere*, Topology **36**:2 (1997), 505–507.

[Whi96] Brian White, *Half of Enneper's surface minimizes area*, Geometric Analysis and the Calculus of Variations (Jürgen Jost, ed.), Internat. Press, 1996, pp. 361–367.