

52670
53042



Tightness for Smooth and Polyhedral Immersions of the Projective Plane with One Handle

DAVIDE P. CERVONE

ABSTRACT. The recent discovery that there is a tight polyhedral immersion of the projective plane with one handle, while there is no smooth tight immersion of the same surface, provides a rare example in low dimensions of a significant difference between smooth and polyhedral surfaces. In this paper the author shows that the obstruction to smoothing the polyhedral model is not local in nature, and describes some of the ways in which the proof of the nonexistence of the smooth tight surface does not carry over to the polyhedral case.

1. Introduction

A longstanding open problem in the study of tight surfaces centered around a question posed by Nicolaas Kuiper [1961] asking whether the surface with Euler characteristic -1 (a real projective plane with one handle) could be tightly immersed in three-space. Kuiper had established that all other surfaces admitted tight immersions in space except for the Klein bottle and the real projective plane, which do not. More than thirty years passed before François Haab [1992] proved that, for smooth surfaces, no such immersion exists. In light of this result and the failure of the attempts to find a polyhedral counterexample, it seemed only a matter of time before a corresponding proof would be found for the polyhedral case as well. Surprisingly, a polyhedral tight immersion of this surface *does* exist, as shown recently in [Cervone 1994]. Although the smooth and polyhedral theories differ substantially for surfaces in high-dimensional spaces, they correspond quite closely in low dimensions; the case of the real projective plane with one handle is important in that it represents one of only a handful of low-dimensional examples where the theories differ in a significant way (see Section 3).

In this article, we compare the smooth and polyhedral behaviors of the projective plane with one handle, and try to illuminate some of the reasons why they differ. The subject is approached from two different directions: first, we

analyze the polyhedral example in detail, especially the potential smoothability of specific configurations within it, and find that the obstruction to smoothing this model is not local in nature. Second, we outline the basic components of Haab's proof, and discuss why this proof does not carry over directly to the polyhedral case.

2. The Polyhedral Tight Immersion

We begin by presenting a tight polyhedral immersion of the real projective plane with one handle [Cervone 1994]. The model has 13 vertices, and their mapping into space is given in Figure 1 along with the 28 triangular faces, and a view of the surface from above. The self-intersection can be seen where faces meet without a heavy black line.




Figure 1. The tight polyhedral projective plane with one handle, including its vertex mapping and face list. The point of view is approximately $(-30, -10, 60)$. The x -axis and y -axis are parallel to the edges ef and fg . See <http://www.geom.umn.edu/locate/rp2-handle> for interactive three-dimensional viewing.

The surface is composed of a central core (an immersed real projective plane with two disks removed) surrounded by a cylinder formed by removing two disks from the convex envelope (the surface of the convex hull) of the central core. The core is based on the projective plane described by Kuiper [1961], who gave level sets for an immersed smooth surface with one maximum, one minimum, and one saddle point. The level sets of the polyhedron given here correspond closely to those described by Kuiper. In the complete surface, the maxima and minima are removed and replaced by a tube connecting the top to the bottom (the outer cylinder). The two curves where the core joins the tube are called *top cycles* (for a complete definition and more information about top cycles, see [Banchoff and Kühnel 1997] in this volume).

Any immersion of a surface of odd Euler characteristic must have a triple point [Banchoff 1974]; this model has exactly one, and it is visible at the center of the figure. Since the central core is essentially a projective plane with one triple point, we would expect the double curve to form three loops meeting at the triple point: six doubly covered lines emanate from the triple point, and since in an immersed surface the double locus forms closed curves, these six lines must be joined pairwise by the double curve, thus forming three loops (see Figure 2). This is indeed the case in our polyhedral model.




Figure 2. A triple point is formed by the intersection of three sheets, where six arcs of double points (dashed lines) meet. In an immersed projective plane with only one triple point, these arcs are joined pairwise by the double curve (dotted lines).

To verify that this is in fact a projective plane with one handle, we compute its Euler characteristic: the object has 13 vertices, 42 edges, and 28 faces, so its Euler characteristic is $V - E + F = -1$ as expected. A polyhedral surface is immersed if, and only if, the star of each of its vertices is embedded. An explicit check of the vertices of this model reveals that each star is embedded (the self-intersection curve does not pass through any vertex), so it is indeed an immersion.

There are several equivalent definitions of tightness (see [Banchoff and Kühnel 1997] in this volume). A geometric definition that applies to both smooth and polyhedral surfaces is that an immersion $f: M \rightarrow \mathbb{R}^3$ is *tight* if it has the *two-piece property*, namely that the preimage of any half-space is connected in M , or in other words any plane cuts the surface into at most two parts. For polyhedral surfaces, this provides a simple characterization of tight immersions [Banchoff and Kühnel 1997, Lemma 1.4.1]:

LEMMA 2.1. *A polyhedral immersion is tight if, and only if,*

- (i) *the one-skeleton of the convex hull is contained in the surface;*
- (ii) *every vertex that is not a vertex of the convex hull lies in the relative interior of some subset of its adjacent vertices; and*
- (iii) *the self-intersection set does not include any vertex of the convex hull of the surface.*

The final condition is required for immersions since such a vertex could be cut off, thereby dividing the surface into at least three pieces (two copies of the vertex, plus the rest of the surface).

We can use this lemma to check that the polyhedron given above is, in fact, tight: first, note that the convex envelope has seven vertices: a, c, d, e, f, g , and h (although b is on the convex envelope, it is not a vertex of it) and that the last eight faces listed in Figure 1 contain all the edges of the convex envelope. Of the remaining six vertices, two lie on the straight line segment between two neighbors (b lies on the segment ac , and k on the segment ag), two lie within a triangle formed by three neighbors (i lies within triangle beh , and j within bfg), and two lie inside tetrahedra formed by four neighbors (l lies within $acfh$, and m within $cfil$). Thus the surface is tight, as claimed.

3. Tight Smoothings of Polyhedral Surfaces

In the previous section we presented a tight polyhedral model of the real projective plane with one handle, and Haab [1992] provided a proof that no smooth immersion of this surface exists. An important question to ask is: Why can't this polyhedral model be smoothed tightly? Given a polyhedral surface, a *tight smoothing* is a tight smooth surface of the same topological type lying within an ε -neighborhood of the polyhedron, for some small ε . Frequently, there exists such a surface for arbitrary ε and these form a continuous deformation from the polyhedral to the smooth surface that is tight at every step. Such a deformation need not always exist, however. For example, in 5-space, there are essentially only two substantially embedded tight immersions of the projective plane: the Veronese surface (a smooth embedding), and the canonical embedding of the six-vertex polyhedral real projective plane (see [Banchoff and Kühnel 1997, Theorem 1.3.6 and Corollary 1.4.12] in this volume). In one sense, the Veronese surface is the obvious tight smoothing of the polyhedral projective plane, but

there is no continuous deformation by tight surfaces from one to the other, since no intermediate tight surfaces exist. Although no examples like this are known in three-space, we will not require a deformation by tight surfaces for our purposes, but only a smooth tight surface sufficiently near the polyhedral one.

One feature that might interfere with smoothability would be the presence of exotic saddle points. These are saddles where there is no direction for which the orthogonal projection of a neighborhood of the saddle is one-to-one, as in Figure 3. Such a saddle is not possible in a smooth surface, since the implicit function theorem guarantees that such a direction always exists. Exotic saddles play a crucial role in one of the few other low-dimensional examples of a difference between the smooth and polyhedral theories, namely the existence of an embedded polyhedral torus with a height function having exactly three critical points, but no such smooth embedding [Banchoff and Takens 1975]. Taking our cue from this, we look for exotic saddles in the polyhedral surface given in Section 2. Checking each of its vertices, however, we find that it contains no exotic saddles, since every vertex has a direction where the orthogonal projection of its star in that direction is a one-to-one mapping.




Figure 3. A polyhedral monkey saddle of the exotic type. There is no direction for which the orthogonal projection is a one-to-one mapping.

A second feature that might interfere with smoothability would be vertices that are not “generic” enough, i.e., ones that can not be perturbed without destroying tightness. For example, vertex k lies on the boundary of the convex hull of its neighbors (it is on the segment ag), and if it were raised in the z direction (for example) it no longer would be in the relative interior of any subset of its neighbors, and the model would fail to be tight (see lemma 2.1). Vertex b has a similar problem, and vertices i and j , which lie on faces of the convex hulls of their neighbors, have directions in which they can not be perturbed without losing tightness. The basic observation is that if a vertex lies on the convex envelope of its neighbors, then it can not be perturbed in some directions, nor can some of its neighbors, without losing tightness. On the other hand, if a vertex lies in the strict interior of the convex hull of its neighbors, then it can be moved slightly in all directions, and so can its neighbors, without destroying

tightness at that vertex. Ideally, then, we would like the model to be in *general position* (i.e., no three vertices lie on a line, and no four vertices lie in a plane).

It is possible to move vertices i , j , and k slightly so that they *do* lie in the interior of the convex hull of their neighbors while maintaining the tightness of the model (move i in the positive x direction, j in the negative x direction, and k in the negative z direction slightly). Since b lies on the convex envelope of the surface, it must be treated with a bit more care. The solution in this case is to move b in the negative y direction; this makes it a vertex of the convex hull of the surface, and so condition *ii* of lemma 2.1 no longer applies. Once this is done, the vertices that form the convex hull of the surface also can be perturbed, provided they stay within the planes of the top cycles to which they belong. Thus all the vertices of the model can be “jiggled” slightly without losing tightness. This is not quite general position, however, since the top cycles must remain in a plane; but these planes can be jiggled slightly as well, provided that all the vertices in the plane are moved together. We will call this *nearly general position*.

One may ask whether the model can be put into truly general position, and indeed it can. The crucial condition is that, to be in general position, the top cycles must be triangular since each top cycle must lie in a plane. The top cycles in this model are formed by quadrilaterals; however, the polyhedron can be modified so that its top cycles *are* triangular. One way to do this is to remove the convex envelope (leaving only the core) and then place large triangles in the planes of the top cycles so that the top cycles lie inside them. Triangulate the annular regions between each triangle and the top cycle contained within it, then move each triangle slightly in the direction perpendicular to its plane but away from the central core (the annular region will become a funnel-shaped “flange”). Finally, add a new cylinder connecting the two large triangles, to replace the convex envelope that was removed at the outset. Provided the new triangles are large enough and the distance they were moved is small enough, the resulting surface will be a tight immersion of the real projective plane with one handle, and its top cycles will be triangular. The vertices that previously formed the top cycles will now be inside the convex hull of the surface, and can be moved slightly to put the entire surface into general position.

Once in general position, every vertex can be perturbed slightly without damaging tightness, including those that form the top cycles. Note that this is in sharp contrast to the smooth situation, where the top cycles are less stable: small changes to a top cycle can easily destroy tightness.

We have seen that the model presented in Section 2 is not initially in general or even nearly general position. The reason for this is twofold: first, it shows that such unusual configurations are possible and not just contrived, and second, it makes checking tightness simpler, since it is easier to check that a vertex lies on a line segment or in a plane than it is to see that it lies inside a tetrahedron. It is important to note, however, that the presence of such configurations is *not* an obstacle to smoothing the polyhedral surface, since the surface can be put

into nearly general position with only minor adjustments, or general position with more extensive changes, without losing tightness.

For additional features that might interfere with smoothability, we turn to an algorithm developed in [Kühnel and Pinkall 1985] that will smooth a tight polyhedral immersion to produce a tight smooth immersion of the same surface that agrees with the polyhedron everywhere but in an ε -neighborhood of the edges. The algorithm does not apply to all polyhedra, however, but only to ones with certain properties, and the conditions that the polyhedral surface must meet are rather strict. Given a vertex of a polyhedral immersion, consider a small sphere centered at the vertex; its intersection with the surface forms a spherical polygon, and a neighborhood of the vertex is the cone over such a polygon. If this cone is convex, then the vertex is called *convex*. A vertex is a *standard saddle* if (i) it is 4-valent; (ii) all the angles of the faces at the vertex are strictly less than π ; and (iii) there are no local support planes through the vertex.

The conditions of the smoothing algorithm require that the nonconvex vertices of the surface be either 3-valent or standard saddles. In particular, no vertex interior to the convex hull can be more than 4-valent, and all the vertices that lie on the convex envelope either must be convex or as simple as possible, i.e., only 3-valent. On the other hand, the algorithm also allows nontriangular faces, and indeed, nonconvex and even non-simply-connected faces, so these restrictions aren't so severe as they at first appear.

The polyhedral model presented in the previous section does not satisfy these conditions, nor can it be modified to satisfy them, otherwise its smoothing would represent a counterexample to Haab's theorem. This leads us to ask: Which vertices cause trouble, and what is the obstruction to smoothing them? All the vertices except m fail the valence conditions, and vertices b , k , i , and j have facets with angles of π or greater (note: coplanar faces that share an edge form a single facet, so abk and kbg form one facet $abgk$ that has an angle of π at k since k lies on the segment between a and g). The latter problem can be resolved by putting the vertices into nearly general position as discussed above, since then all the interior facets are triangles, and so have angles strictly less than π . The remaining difficulty is the large number of edges at each vertex.

One approach to the valence problem is to try to split up a vertex of high valence into several vertices of lesser valence. The trick is to do it while still maintaining tightness. Since the model is in nearly general position, if the new vertices all are in a small neighborhood of the original vertex, this will not disrupt the tightness at its neighbors, so it is only necessary to check that the new vertices satisfy condition *ii* of lemma 2.1.

As an example, consider vertex a , a vertex of the convex hull of the surface that is nonconvex and 7-valent. It can be broken into three vertices, as shown in Figure 4. One of these, a_1 , is convex, and the other two, a_2 and a_3 , are 3-valent. Note that the introduction of nontriangular faces is allowed, and indeed




Figure 4. The nonconvex vertex a can be broken into three vertices, one convex, the other two 3-valent, all lying in the plane of the top cycle containing a .




Figure 5. The nonconvex vertex f can be broken into five vertices, one convex, two 3-valent, and the other two standard saddles.

is crucial to obtaining the correct valences at a_2 and a_3 . This modification maintains tightness, since a_1 is now a vertex of the convex hull, while a_2 and a_3 are in the relative interiors of their neighbors (a_2 is interior to the triangle formed by a_1 , a_3 and the vertex to the left of the original vertex a , while a_3 is interior to a_1 , a_2 and the vertex to the right of a).

We will call a vertex *locally smoothable* if it can be replaced by a collection of new vertices such that the resulting polyhedral model is still tight, and all the new vertices satisfy the smoothing criterion of Kühnel and Pinkall. Thus vertex a is locally smoothable. Vertices c , d and e likewise are locally smoothable using a similar decomposition. A local smoothing configuration for vertex f is given in Figure 5. Again, nontriangular faces are crucial to the construction, and care must be taken at the 4-valent vertices to assure that the angles are less than π . Vertex g can be handled similarly.

Vertex b presents a more difficult challenge. It can be modified as shown in Figure 6. Here, rather than pushing b outward so that it becomes a vertex of the convex envelope, we pull it into the interior and subdivide it. It is important that b_1 , b_2 , b_3 and g lie in a plane, with the angle at b_3 less than π . Vertex b_1 should be placed on the back side of the plane containing a , c and g , while b_2 and b_3 should be placed in the interior of the tetrahedron b_1gij . A similar, though slightly more complex, construction is possible at vertex h .

Vertices i and k are interior to the convex hull and are 5-valent; they can be treated in much the same way that b was above. Vertex m already satisfies the smoothing conditions. Vertex j can be split into two vertices, one 5-valent, the




Figure 6. Vertex b can be pulled into the interior of the convex hull (middle) and then split into three vertices (right), all of which are standard saddles.

other 4-valent (Figure 7, middle) provided j_1, j_2, b and g are coplanar; j_1 can then be split in a way similar to how b was modified above (Figure 7, right). An extra vertex j_5 is added between j_2 and h to divert from j_2 an edge generated while subdividing j_1 . This requires j_1, j_2, j_3 and j_5 to be coplanar, as well as j_1, j_3, j_4 , and f . One can arrange that these planarity conditions are satisfied while still placing each new vertex in the interior of the tetrahedron formed by its four neighbors. Thus j is locally smoothable. Finally, l can be split into two 5-valent vertices that can be handled in a similar fashion.




Figure 7. The 6-valent vertex j can be broken into two vertices (middle), one a standard saddle, the other 5-valent; the latter can in turn be broken into three standard saddles (right). An extra vertex is added near j_2 in order to divert an edge that would have increased the valence of j_2 .

We see, then, that *every* vertex of the polyhedral model is locally smoothable; but this does *not* mean that the model itself is smoothable. The reason is that we can not carry out all the modifications described above simultaneously without them interfering with each other. For example, the modification to b adds edges to a and c , which will disrupt the carefully planned valence configurations at those vertices. Some coordination among these modifications may be possible, however. For example, two new edges generated at neighboring vertices may form triangles that can be combined into a single planar quadrilateral, thus maintaining the proper valence at each vertex (the author currently is investigating such possibilities).

Although we can not carry out all the modifications concurrently, the fact that each vertex is locally smoothable has an important implication, namely that the obstruction to smoothing this polyhedral surface is not a local one, isolated at some unusual vertex (since neighborhoods of all the vertices look like smoothable patches that could become part of a smooth tight surface). Rather the obstruction is a global one, intrinsic to the surface itself. What this global property is remains an open question at this point.

4. Haab's Proof for Smooth Surfaces

Haab's proof [1992] of the fact that there is no smooth immersion of the real projective plane with one handle requires considerable machinery, of which we outline some of the highlights.

His basic idea is to consider mappings of surfaces into the plane, and to determine when these can be "factored" into an immersion of the surface in space followed by a projection into the plane. One of the important features of such projections are their fold curves. For smooth surfaces, the fold curves for almost all projections are formed by a collection of images of circles that are smooth immersions except at a finite number of cuspidal points.

The fold curves for tight immersions have a specific geometric form: one component is convex, and the others are locally concave (with respect to the image of the projected surface) and contain all the cusps (see Figure 8). Haab computes strict bounds on the number of components that can exist in the fold set for a tight immersion of a given surface. He defines a degree on each of the fold curves, and shows that it is 2 for the convex component, and strictly negative for the others, and that moreover, the sum of the absolute values of the degrees is equal to 4 minus the Euler characteristic of the surface. This provides a key connection between the fold curves and tightness.

Haab then considers height functions on surfaces with boundary and classifies their saddle points according to whether passing the saddle point changes the number of components in the level set of the height function. He uses a generalization of the Morse inequalities to show that the number of saddles that do not change the number of components is equivalent mod 2 to the genus of the




Figure 8. The fold curves for a tight immersion of the surface with Euler characteristic -3 , after [Haab 1992]. The saddles that change the number of components in the level set are shown by dashed lines, and those that don't are shown by dotted lines. Note that the type of saddle changes only at points where there is a doubly tangent line to the fold set.

surface. This allows him to prove an actual bound (in terms of the genus of the surface) on the number of such saddle points for any height function. Haab then considers the fold sets of projections of tight surfaces, and determines that the type of saddle that occurs (for height functions in directions perpendicular to the direction of projection) on the fold set for a given projection can change only at points where the tangent line to the fold curve is tangent to the fold set at more than one point (see Figure 8).

He applies this information to the case of the projective plane with one handle, and concludes that, for a tight immersion, the type of saddle is constant on each component of the fold set, and that the fold set has exactly three components. He uses the fold curves and the top cycles to decompose the surface into disjoint regions, and shows that one of these contains a cycle that separates it, but on which the Gauss map is constant. He considers a height function in the direction of the Gauss map, and, after showing that the curve bounds a region on which the Gauss map is *not* constant, deduces that the height function has a local extremum inside that region. The initial decomposition of the surface into regions ensures that this point is not a global extremum for the height function, which contradicts the fact that the immersion is tight, since a plane perpendicular to the direction of the height function and just below the lower extremum will cut the surface into three disjoint pieces.

5. Some Differences Between the Smooth and Polyhedral Theories

Haab relies heavily on the smoothness of the surface in his proof, but not all of his results carry over to the polyhedral situation. There are several important differences that arise in the polyhedral case.

First, for smooth surfaces, the fold curves of almost all projections form disjoint components, so it is possible to count the number of components accurately. In the polyhedral case, the analogue of the fold curves are formed by fold edges (ones where the two triangles sharing that edge both project onto the same side of the edge), but an arbitrarily large number of fold edges may come together at a single vertex of a polyhedral surface. Thus it is not always possible to determine a canonical way to divide the fold edges into fold curves, and the number of components may change with different divisions into curves. In the polyhedral model presented here the projection onto the xy -plane has either one or two components (see Figure 9) depending on how the choice is made at vertex h where four fold edges meet. This is not the three predicted by Haab for the smooth case, so already a difference has emerged between the two types of surfaces. It is interesting to note that the conditions of Kühnel and Pinkall's smoothing algorithm [1985] guarantee that there are at most two fold edges at each vertex, so for these polyhedra, the fold curves can be separated into components without ambiguity.




Figure 9. The fold curves for two projections of the tight polyhedral projective plane with a handle; each can be broken into fold curves in more than one way. The projection onto the xy plane (left) can be broken into one or two components depending on the decision made at the topmost vertex, h . The other projection (right) can be broken into as many as six components; it has three interior vertices where more than two fold edges meet. There are projections for which the fold curves are even more complicated.

Second, the idea of a cusp and of locally convex curves is harder to formulate in the polyhedral case. One might begin by identifying analogous polyhedral structures, such as those shown in Figure 10, but this becomes more complicated when more than two fold edges meet at one vertex.




Figure 10. Polyhedral analogues of the fold and cusp (left). The neighborhood of a vertex can wrap around the vertex an arbitrary number of times (right), though only a single revolution is shown here. The fold edges themselves do not contain enough information to distinguish between the two folds shown.

The problem is compounded by the fact that a star can wind around a vertex an arbitrary number of times with no visible effect on the angle between the fold edges. This makes computing the degree of a fold curve more complicated than in the smooth case (where small loops would be present). In the polyhedral case, the fold curve itself does not contain enough information to distinguish between the two situations shown in Figure 10, so any analog of Haab's fold-curve degree would have to be more complex, and would probably involve looking at the strip neighborhoods of the fold curves. (Again, note that the conditions of Kühnel and Pinkall's smoothing algorithm do not allow for such aberrant behavior.)

For smooth tight surfaces, almost every point of the (nonconvex) fold curves is a saddle for some direction, but this is no longer the case for polyhedral surfaces, where only vertices can be saddles. Moreover, these saddles can easily be of higher degree (i.e., monkey saddles or worse), and can be the exotic type that have no smooth counterpart (Figure 3). Such saddles are stable under small changes of direction, which is not the case for smooth surfaces. This complicates the issue of determining the "type" of each saddle, and thus the type of each fold component, which is crucial to Haab's argument. Haab's results concerning when the type of a fold curve can change would require additional work in the polyhedral case, since the idea of bitangent lines to the fold curves does not have a direct analog.

Finally, Haab uses the fact that the top cycles in a smooth tight immersion lie in distinct planes to break the central projective plane into regions, one of which is an annulus. In the polyhedral case, the top cycles can share vertices or even edges, making such a decomposition harder to formulate. For example,




Figure 11. A surface with three top cycles that share edges in pairs.

the Császár torus [Császár 1949], which is tight, has two triangular top cycles that share a vertex. To generate an example where the top cycles share edges, begin with a hexagonal prism and remove every other rectangular face. Place the prism within a large triangular prism whose faces are parallel to the ones removed from the hexagonal prism. Finally, replace each rectangular face of the triangular prism with four trapezoidal faces that connect its boundary to the boundary of the corresponding removed face of the hexagonal prism. The resulting embedded surface is tight (see Figure 11); it has three rectangular top cycles (the boundaries of the rectangular faces of the outer prism) that share edges in pairs: each edge joining the top and bottom triangles of the prism belongs to two distinct top cycles.

6. Conclusion

The existence of a tight polyhedral immersion of the real projective plane with one handle, but no smooth one, provides an opportunity to study some of the details of how these two types of surfaces differ even in low dimensions. One question that remains is how unique is this situation? Are there other tight polyhedral immersions of surfaces that can not be tightly smoothed? The paper [Cervone 1996] presents several tight polyhedral immersions to which Kühnel and Pinkall's smoothing algorithm does not apply, and for which no corresponding smooth tight immersion is known; are these models also examples of this same phenomena? Or is there a more general smoothing algorithm that will handle these cases? The example presented here shows that no smoothing algorithm will work for *every* tight polyhedral surface, and an understanding of just why that is so should provide insight into both the smooth and polyhedral worlds.

References

- [Banchoff 1974] T. F. Banchoff, “Triple points and surgery of immersed surfaces”, *Proc. Amer. Math. Soc.* **46** (1974), 407–413.
- [Banchoff and Kühnel 1997] T. F. Banchoff and W. Kühnel, “Tight submanifolds, smooth and polyhedral”, pp. 51–118 in *Tight and Taut Submanifolds*, edited by T. E. Cecil and S.-s. Chern, Cambridge U. Press, 1997.
- [Banchoff and Takens 1975] T. F. Banchoff and F. Takens, “Height functions on surfaces with three critical points”, *Illinois J. Math.* **19** (1975), 325–335.
- [Cervone 1994] D. P. Cervone, “A tight polyhedral immersion of the real projective plane with one handle”, 1994. Available at <http://www.geom.umn.edu/locate/rp2-handle>.
- [Cervone 1996] D. P. Cervone, “Tight immersions of simplicial surfaces in three space”, *Topology* **35**:4 (1996), 863–873.
- [Császár 1949] Á. Császár, “A polyhedron without diagonals”, *Acta Univ. Szeged. Sect. Sci. Math.* **13** (1949), 140–142.
- [Haab 1992] F. Haab, “Immersiones tendues de surfaces dans E^3 ”, *Comment. Math. Helv.* **67**:2 (1992), 182–202.
- [Kühnel and Pinkall 1985] W. Kühnel and U. Pinkall, “Tight smoothing of some polyhedral surfaces”, pp. 227–239 in *Global differential geometry and global analysis* (Berlin, 1984), edited by D. Ferus et al., Lecture Notes in Math. **1156**, Springer, Berlin, 1985.
- [Kuiper 1961] N. H. Kuiper, “Convex immersions of closed surfaces in E^3 . Nonorientable closed surfaces in E^3 with minimal total absolute Gauss-curvature”, *Comment. Math. Helv.* **35** (1961), 85–92.

DAVIDE P. CERVONE
 MATHEMATICS DEPARTMENT
 UNION COLLEGE
 SCHENECTADY, NY 12308
 dpvc@union.edu

