Removing Even Crossings on Surfaces

Michael J. Pelsmajer
Department of Applied Mathematics
Illinois Institute of Technology
Chicago, Illinois 60616, USA
pelsmajer@iit.edu

Marcus Schaefer
Department of Computer Science
DePaul University
Chicago, Illinois 60604, USA
mschaefer@cs.depaul.edu

Daniel Štefankovič
Computer Science Department
University of Rochester
Rochester, NY 14627-0226
stefanko@cs.rochester.edu

January 22, 2009

Abstract

In this paper we investigate how certain results related to the Hanani-Tutte theorem can be extended from the plane to surfaces. We give a simple topological proof that the weak Hanani-Tutte theorem is true on arbitrary surfaces, both orientable and nonorientable. We apply these results and the proof techniques to obtain new and old results about generalized thrackles, including that every bipartite generalized thrackle on a surface S can be embedded on S. We also extend to arbitrary surfaces a result of Pach and Tóth that allows the redrawing of a graph so as to remove all crossings with even edges. From this we can conclude that $\mathrm{cr}_S(G)$, the *crossing number* of a graph G on surface S, is bounded by $2\,\mathrm{ocr}_S(G)^2$, where $\mathrm{ocr}_S(G)$ is the *odd crossing number* of G on surface S. Finally, we prove that $\mathrm{ocr}_S(G) = \mathrm{cr}_S(G)$ whenever $\mathrm{ocr}_S(G) \leq 2$, for any surface S.

1 Introduction

We continue the investigation of the Hanani-Tutte theorem and its close relatives begun in "Removing Even Crossings" [14], this time aiming for analogues on orientable and nonorientable surfaces. The theorem of Hanani and Tutte states that every drawing in the plane of a nonplanar graph contains two non-adjacent edges which cross an odd number of times.[1] There are several proofs of this theorem [5, 17, 6, 7, 16, 10] starting with the original papers by Hanani and

[1] We make the usual assumptions on drawings of graphs: any two edges intersect at most finitely often, their intersection points are either shared endpoints or crossings, and no more than two edges cross at a point [11, page 230]. For a detailed discussion, see [16].

Tutte. Kleitman's proof [7] is particularly short and elegant. All of these proofs invoke Kuratowski's theorem and then verify the result for subdivisions of $K_{3,3}$ and K_5. This approach seems hopeless for higher surfaces (the complete list of excluded minors for the torus is not known).[2] In [14] we gave a new proof of the Hanani-Tutte theorem in the plane which avoids Kuratowski's theorem and uses elementary topological methods only.

Here we show that these methods easily extend from the plane to surfaces. We prove the weak version of the Hanani-Tutte theorem for arbitrary surfaces: If a graph is drawn on a surface on which it cannot be embedded, then there must be two edges in the drawing that cross an odd number of times. Cairns and Nikolayevsky showed this result for orientable surfaces using homology theory [2]; our proof is entirely elementary, and algorithmic. (See Section 3.)

Cairns and Nikolayevsky proved the previous result to apply it to generalized thrackles, graphs that can be drawn so that every pair of edges intersects an odd number of times (a common endpoint of two edges counts as an intersection). Our topological approach handles generalized thrackles very naturally, and we can simplify and improve some of their central results. We show that a bipartite graph is a generalized thrackle on surface S if and only if it is embeddable on that surface; Cairns and Nikolayevsky proved this for orientable surfaces [2]. We also extend their result that a graph is a generalized thrackle on an orientable surface S if and only if it has a parity embedding on the surface obtained from S by adding a crosscap [3] to a statement that applies when S is nonorientable. (See Section 4.)

While this might leave the impression that results related to the Hanani-Tutte theorem easily generalize to arbitrary surfaces, this is not actually the case. Consider the following theorem from [14], which is at the core of our proof of the (strong) Hanani-Tutte theorem (for the plane). An *even* edge in a drawing is an edge that crosses every other edge an even number of times (including the possibility that it does not cross it at all).

Theorem 1.1 (Pelsmajer, Schaefer, Štefankovič)**.** *If D is a drawing of G in the plane, and E_0 is the set of even edges in D, then G can be drawn in the plane so that no edge in E_0 is involved in any crossings and there are no new pairs of edges that cross an odd number of times.*

This theorem is a strengthening of the following result of Pach and Tóth [12, Theorem 1].

Theorem 1.2 (Pach, Tóth)**.** *If D is a drawing of G in the plane, and E_0 is the set of even edges in D, then G can be drawn in the plane so that no edge in E_0 is involved in a crossing.*

In Section 5 we will see that this result of Pach and Tóth remains true on arbitrary surfaces, whereas Theorem 1.1 cannot even be extended to the projective plane or the torus. This means that some of the stronger consequences

[2]With one exception: we were recently able to establish the Hanani-Tutte theorem for the projective plane using excluded minors [13].

we could derive from Theorem 1.1 in the plane might fail or need a different approach for other surfaces. For example, we know that the crossing number of a graph, $\mathrm{cr}(G)$, and its odd crossing number, $\mathrm{ocr}(G)$, agree in the plane for $\mathrm{ocr}(G) \leq 3$ [14]; the proof used Theorem 1.1. We can still show that $\mathrm{cr}_S(G) = \mathrm{ocr}_S(G)$ if $\mathrm{ocr}_S(G) \leq 2$, but the proof is more intricate. On the other hand, our extension of Theorem 1.2 immediately yields that

$$\mathrm{cr}_S(G) \leq 2\,\mathrm{ocr}_S(G)^2$$

is true on arbitrary surface S, generalizing a result of Pach and Tóth for the plane [12], where $\mathrm{cr}_S(G)$ is the crossing number of G on S and $\mathrm{ocr}_S(G)$ is the odd crossing number of G on S.

2 Graphs and Surfaces

We begin with a short review of graphs and surfaces; see [9] or [4] for more background. In this paper, a *surface* is a connected, compact 2-manifold[3] without boundary unless we say otherwise.[4]

By the classification theorem for surfaces, an orientable surface is homeomorphic to a sphere with a number of handles attached, and a nonorientable surface is homeomorphic to a sphere with a number of crosscaps. That number is the *(orientable) genus* or *nonorientable genus* of the surface. (If there are more than two crosscaps, two crosscaps can be exchanged for one handle.) We say that a surface S has *smaller genus* than S' if S' can be obtained from S by adding handles and/or crosscaps.

Consider a closed curve C on a surface S. C is *S-separating* or simply *separating* if $S - C$ has two components. Otherwise C is *nonseparating*, and $S - C$ is connected. If C is *contractible* in S, that is, continuously deformable within S to a point, then C is S-separating. If S is nonorientable then C may be either *one-sided* or *two-sided*, depending on whether C passes through crosscaps an odd or even number of times. (If S is orientable then C must be two-sided.)

We may cut a surface S along a noncontractible closed curve C, temporarily attach a disk to each boundary component of $S - C$, then contract each disk to a point. If C is S-separating this yields two surfaces of smaller genus. If C is a nonseparating curve, we obtain one surface S' of smaller genus, which we call the *C-reduced surface*. If S is orientable, then S' is the orientable surface with genus one less than the genus of S. If S and S' are nonorientable then S' has nonorientable genus either one or two less than S, depending on whether C is one-sided or two-sided, as this determines whether one or two disks are attached after the cut. It can also happen that S is nonorientable and S' is orientable: for example, if S is the projective plane then there is no surface of smaller nonorientable genus, so this is the only possibility.

[3] A *2-manifold* is a Hausdorff topological space that is everywhere locally homeomorphic to a disk.

[4] Note that the plane $\mathbb{R}^2$ is not a surface by this definition, but embeddability in the plane and embeddability in the sphere are equivalent.

The following lemma helps us find nonseparating curves.

Lemma 2.1. *Suppose that C and C' are closed curves on a surface S. If C and C' cross an odd number of times, then both are nonseparating curves.*

Proof. Suppose that C is separating. Then $S - C$ has two components S_1, S_2, each with C as its boundary. If we trace the curve C', it must switch between S_1 and S_2 each time it crosses C, and never otherwise. Hence there must be an even number of switches, contradicting the assumption that C and C' cross oddly. Hence, neither C nor C' can be separating. ■

Next, we show how to adapt terminology for embedded graphs to graphs drawn on surfaces. We will phrase our definitions so that they work for multigraphs as well, since loops and multiple edges often arise in our proofs. Consider a drawing D of a multigraph G on a surface S. The (local) *rotation* at a vertex v is the cyclic order of the edges at v (for a loop we need to distinguish its two ends at v), and the *rotation system* of D is the collection of rotations. If S is orientable we choose rotations to be clockwise with respect to a fixed side of S; otherwise, we arbitrarily select the rotation to be clockwise to one side of S locally near v. The clockwise direction with respect to a local side of S is called the *local orientation* at the vertex. The local orientation can be transferred along the surface locally from one vertex u to another v along an edge $e = uv$;[5] the orientation at the endpoints will either agree, or—if they are clockwise with respect to opposite sides of S (locally along e)—will be reversed. An *embedding scheme*[6] is a rotation system plus a *signature* $\lambda : E \to \{-1, 1\}$ such that for each edge $e = uv$, we have $\lambda(e) = 1$ if the sense of clockwise rotation at u and v agree along e, and $\lambda(e) = -1$ otherwise. Note that when S is orientable we chose the rotation system so that every edge has signature 1.

We often create self-intersections when redrawing an edge, which we then remove by redrawing (as shown in Figure 1).

The crossing breaks the drawing of e into three curves, and we can think of each as having its own signature, with the crossing point as a vertex. Note that the signature of the entire drawing of e is the product of these three signatures, whether we use the original drawing or the redrawing. Therefore we can remove self-intersections from edges, while maintaining the embedding scheme.

A key step in many of our proofs is the contraction of an edge $e = uv$ by moving v towards u, and finally identifying v with u (see Figure 2). Performing

[5] This is reasonably straightforward and well-known for multigraph embeddings, but less so for drawings allowing self-intersections of edges. Since S is compact and the drawing of e is the continuous image of $[0, 1]$ (also compact) with 0 mapped to u and 1 mapped to v, we can specify a finite set of open disks in S that cover e, and partition $[0, 1]$ into subintervals, each of which is mapped into a disk. Each disk is orientable so the sense of clockwise is consistent there, and thus the local orientation can be transferred along each portion of the drawing of e corresponding to the image of one of the subintervals of $[0, 1]$.

[6] So called because it is commonly used for embedding multigraphs (and for creating a surface on which a multigraph can be embedded). However, our version also applies to multigraph drawings with edge crossings.

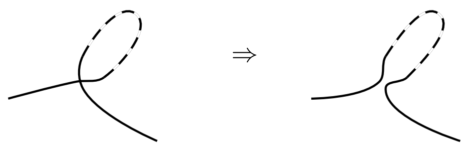

Figure 1: Removing a self-intersection; illustration taken from [14].

this contraction on a multigraph with an embedding scheme naturally induces an embedding scheme for the resulting multigraph as follows.

For a multigraph G drawn on a surface S and an edge $e = uv$, consider G' obtained by contracting edge e by moving v along e towards u and identifying u and v obtaining a new vertex u' (only the edge e is eliminated). We give u' the natural rotation suggested in Figure 2. If $\lambda(e) = 1$, the orientation at u and v agrees along e, so we use the same orientation for the rotation at u' as well. If $\lambda(e) = -1$, let the local orientation at u' agree with the one used at u; we must also flip the signature of each edge that was incident to v.

Our main tool connecting contraction and embedding schemes is the following observation.

Lemma 2.2. *Let S be a surface. Suppose we are given a multigraph G together with an embedding scheme. Let G' be obtained from G by contracting an edge $e = uv$ towards u and modifying the embedding scheme as described above. If G' can be embedded on S realizing this embedding scheme, then G can be embedded in S realizing its original embedding scheme.*

3 The Weak Hanani-Tutte Theorem on a Surface

In this section we want to show that the weak Hanani-Tutte theorem is true for arbitrary surfaces. This result is known for orientable surfaces, with a short and elegant proof using homology theory [2, Lemma 3]. In the spirit of our earlier paper we present an elementary, topological proof. We begin with a simple observation.

Proposition 3.1. *Let G be a multigraph with a single vertex v, drawn on a surface other than the sphere, so that all edges are even (loops). Then either G contains an edge e that is a nonseparating curve, or else we can draw a new nonseparating curve through v that crosses each edge of G an even number of times. (In the latter case, we can add a new edge e to G, drawn as the new curve, so that all edges are even.)*

Proof. Since the surface is not the sphere it must contain a nonseparating curve C, and we may assume that $v \notin C$. If C crosses no edge of G an odd number of times then we can use the curve C itself: deform a small segment of it (without crossing over the vertex) to approach v between two consecutive edges in the rotation, so the curve eventually contains v. Otherwise there is a loop e in G that crosses C an odd number of times. By Lemma 2.1, e is nonseparating. ■

We are now ready to state and prove the weak Hanani-Tutte theorem for arbitrary surfaces. Part of the argument is similar to our earlier proof of the weak Hanani-Tutte theorem for the plane [14]; new ideas are needed to deal with the case in which all edges in the graph are loops. We will also exploit these new ideas in later sections.

Theorem 3.2. *If G can be drawn on a surface S so that all its edges are even, then G can be embedded on that surface, i.e. drawn crossing-free, without changing the embedding scheme.*

Cairns and Nikolayevsky's proof [2] also preserves the embedding scheme.

Proof. Fix a drawing D of G on some surface S. The proof will be a double induction over the genus of the surface (using our definition of "smaller genus") and the number of vertices of G. To make the inductive step work, we prove the following slightly stronger statement:

> If D is a drawing of a multigraph G on a surface so that any pair of edges crosses an even number of times in D, then G can be drawn without crossings on that surface without changing the embedding scheme.

If D contains an even non-loop edge uv we proceed as in a previous paper [14, Theorem 1.1]: contract the edge uv by pulling v towards u as shown in Figure 2.

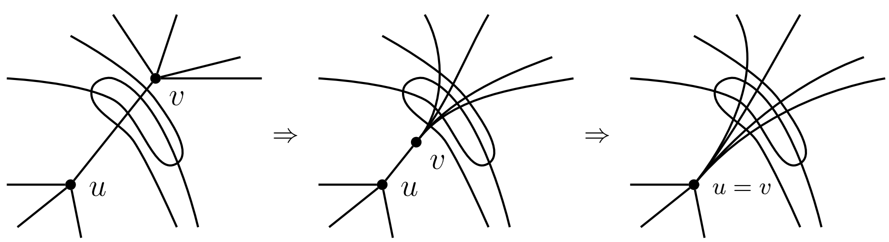

Figure 2: Pulling an endpoint (left to middle) and contracting the edge (middle to right); illustration taken from [14].

In the new drawing the edges incident to v remain even, since uv was an even edge. We might have introduced self-intersections by contracting uv, but these self-intersections are easily eliminated as shown earlier in Figure 1.

Finally, we contract uv to a single vertex u and join the rotations at the two vertices appropriately (as explained in Section 2). Call the new multigraph G'. By inductive assumption, there is a drawing of G' on the surface S without crossings with the embedding scheme unchanged. In such a drawing we can split the vertex u into vertices u and v and reintroduce the edge between the two vertices without introducing any crossings (Lemma 2.2).

Once we have contracted all non-loop edges, we are left with one or more vertices depending on whether G was connected or not. Suppose there are at least two vertices u and v left. Draw a curve C connecting u and v. Move v towards u along C, contracting C as if it were an edge in the procedure we described earlier, and join the rotations of u and v as before. Note that all edges remain even, since all edges attached to v are loops. Call the new multigraph G'. By induction, we can draw G' on S without crossings and embedding scheme unchanged. In this drawing we can split u and v to get a drawing of the original graph.

This leaves us with the case that G consists of a single vertex u with loops. If the surface S is the sphere, we can redraw the loops without any crossings: Since the loops are even, their ends at u cannot alternate. Pick a loop e whose ends are closest in the rotation, then the ends must be consecutive in the rotation at u. Recursively redraw $G - e$ without crossings, then connect the ends of e without introducing crossings.

We can therefore assume that S is not the sphere. By Proposition 3.1 there is an edge e in G—or one can be added to G so as to not create any odd crossings—which is drawn as a nonseparating curve. Next, we want to redraw G so that e is crossing-free. Since e is even, the crossings with e can be partitioned into pairs such that each pair involves the same edge (other than e). Now erase all crossings with e, and for each pair, on each side of e, draw a curve alongside e to connect the severed ends. (See Figure 3 for an example.)

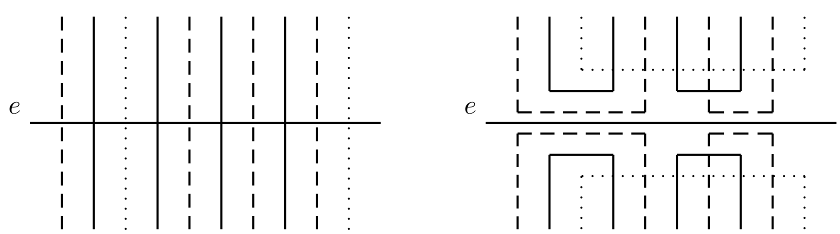

Figure 3: Eliminating crossings with e.

Note that this procedure does not change the parity of the number of crossings between any pair of edges, although it does lead to “curves” with multiple components, only one of which contains v. However, since e is nonseparating, this can be fixed: Within $S - e$, we can deform a small segment of a component so that it approaches another component of the same curve, without passing through any vertices. Then we apply the local redrawing move shown in Figure 4 to combine the two components. Since we avoid the vertex when redrawing, the number of crossings between each pair of edges remains even. Repeatedly doing this for all components makes every edge into a single closed curve again, none of which crosses e.

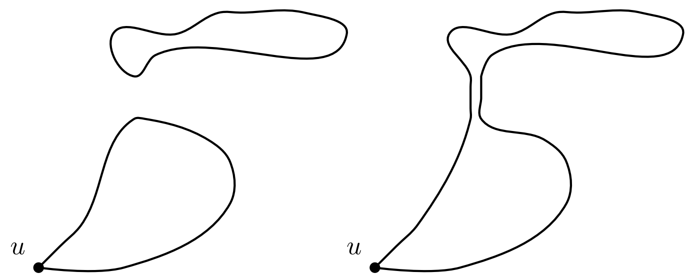

Figure 4: Reconnecting a closed component to the main curve.

The rotation at a vertex can be written as a cyclic word, where each edge end is represented by a letter in the cyclic word, in clockwise order; in particular, the rotation at u can be written as ew_1ew_2, where w_1, w_2 are words.

Now draw a closed curve C along e, through u, and remove e from the graph, and create the C-reduced surface S'. After cutting along C, we get a multigraph on two vertices u_1, u_2 with rotations w_1, w_2, respectively. If C is a two-sided curve, we apply induction to draw this graph on S' without crossings with the same embedding scheme; we then cut a hole close to each u_i so that it touches u_i between the first and last end in w_i, and glue together the hole boundaries to recover G drawn on S with the original embedding scheme.

However if C is a one-sided curve, contracting the attached disk will bring u_1 and u_2 together, so we must flip the local orientation at one of these vertices first, say u_1, and the signature of every edge in its rotation needs to be flipped as well. (If S' is orientable then we should continue to flip local orientations, and signatures of incident edges, until we have an overall consistent orientation on S', i.e., such that every signature is 1.) The rotation at the new vertex will be $w_1w_2^R$ or $w_1^Rw_2$ (the R means “reversed”). Now we can apply induction to the new single-vertex graph on S' to obtain a crossing-free drawing. We then split the vertex back into u_1, u_2 and cut a hole with u_1, u_2 on its boundary. If we let C_1 denote part of the boundary from u_1 to u_2 and let C_2 denote the other part of the boundary from u_2 to u_1, then identifying C_1 and C_2 recovers the original surface S and the original graph G, still with no crossings. ■

4 Generalized Thrackles

A graph is a *thrackle* if it can be drawn such that any pair of edges intersects exactly once, where a common endpoint of two edges counts as an intersection of these two edges. A *generalized thrackle* is a graph that can be drawn such that any pair of edges intersects an odd number of times (again counting endpoints). Generalized thrackles were introduced by Woodall in 1972 [1, p. 359–363].

Cairns and Nikolayevsky used the weak Hanani-Tutte theorem to prove the following result on generalized thrackles [2]. They used that result to get tighter bounds on the number of edges of thrackles and generalized thrackles.

Theorem 4.1 (Cairns, Nikolayevsky)**.** *Let G be bipartite. Then G is a generalized thrackle on an orientable surface if and only if G can be embedded on that surface.*

The special case of the sphere was first proved by Lovász, Pach, and Szegedy [8]: A bipartite graph is a generalized thrackle if and only if it is planar.

With the weak Hanani-Tutte theorem for arbitrary surfaces, we can generalize Cairns and Nikolayevsky's result to arbitrary surfaces. The following remark shows that the distinction between crossings and intersections in the definition of generalized thrackles is not essential.

Remark 4.2. Two edges *cross* (rather than intersect) if they intersect at a point which is not an endpoint of either edge. We can easily see that a graph is a generalized thrackle if and only if it can be drawn such that any pair of edges crosses (rather than intersects) an odd number of times: Simply add a *twist* near every vertex; that is, redraw the graph locally near each vertex so every pair of edges incident to the vertex cross an odd number of times as illustrated in Figure 5. This suffices since for each intersection of two edges at an endpoint there is one new crossing, and all other intersections are crossings. The desired redrawing can always be done: We can assume that, locally, the edges leaving the vertex look as shown on the left of Figure 5; that is, the edges are straight lines leaving upwards. We now reverse the rotation at the vertex and reconnect. (Cairns and Nikolayevsky use an equivalent idea.)

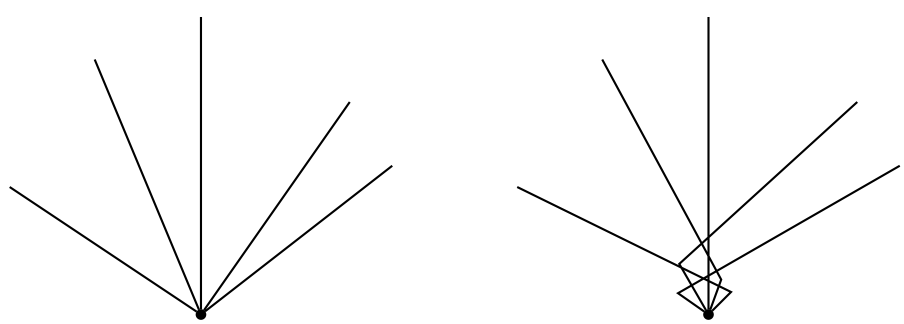

Figure 5: Flipping the parity of crossings at a vertex.

The following lemma is the main ingredient of our proof; it establishes the link between generalized thrackles and the weak Hanani-Tutte theorem.

Lemma 4.3. *If D is a drawing of a bipartite graph G on some surface S, then we can find a drawing D' of G on S such that any two edges cross oddly in D*

if and only if they cross evenly in D'. In other words, we can flip the crossing parity of all pairs of edges.

The lemma cannot be extended to all graphs: it is well-known that there are planar graphs that are not generalized thrackles (e.g. the wheel W_4, see [2]).

For the proof we make use of a simple move used, for example, by Tutte in his paper on crossing numbers [17]; for more recent uses, see [12, 3]. For any $e \in E(G), v \in V(G)$, the *e, v-move* deforms e so that a portion of e is moved close to v, then across v, without having e pass over any other vertices. The effect is to change the parity of crossing between e and every edge incident to v. Note that if e is incident to v, the self-intersections can be easily removed as before.

Proof. Fix a drawing D of the bipartite graph $G = (V_0 \cup V_1, E)$ with edges between $V_0 = \{v_0, \ldots, v_n\}$ and V_1. Let $i_D(e, f)$ be the parity of the number of crossings between edges e and f in the drawing D. Now for each v_i, perform an e, v_j-move for all $j > i$ and every $e \in E$ incident to v_i yielding the drawing D'. Since every edge is adjacent to some vertex in V_0 we conclude that

$$i_{D'}(e, f) = \begin{cases} i_D(e, f) & \text{if } e \text{ and } f \text{ share an endpoint in } V_0 \\ 1 - i_D(e, f) & \text{otherwise.} \end{cases}$$

Now apply the parity-flipping operation described in Remark 4.2 to each vertex in V_0 only. The resulting drawing D'' fulfills $i_{D''}(e, f) = 1 - i_D(e, f)$ for all edges e and f which is what was required. ■

If we are given a generalized thrackle G on a surface S, G can be drawn on S so that every pair of edges crosses oddly (Remark 4.2). Then G can be redrawn so that every pair of edges crosses evenly (Lemma 4.3). Hence G can be embedded in S (Theorem 3.2). On the other hand, if G can be embedded in S, then every pair of edges in the embedding crosses an even number of times, and we can redraw G so that every pair of edges crosses an odd number of times (Lemma 4.3), which is equivalent to G being a generalized thrackle (Remark 4.2). This completes an easy topological proof of the following theorem.

Theorem 4.4. *Let G be bipartite. Then G is a generalized thrackle on a surface if and only if G can be embedded on that surface.*

In a subsequent paper, Cairns and Nikolayevsky study generalized thrackles of non-bipartite graphs on orientable surfaces, using the following definition. A *parity embedding* of a graph is a drawing without crossings such that even cycles are two-sided curves and odd cycles are one-sided curves. They prove [3, Theorem 2]:

Theorem 4.5 (Cairns, Nikolayevsky)**.** *G is a generalized thrackle on an orientable surface S if and only if G has a parity embedding on the (nonorientable) surface obtained by adding a crosscap to S.*

Using our generalization of the Hanani-Tutte theorem that applies to nonorientable surfaces, we can give a fairly simple proof of this result, which originally had quite a lengthy proof. Our methods work just as well when S is nonorientable, once we see how to extend the definition of parity embedding appropriately. Let S' be the surface obtained by adding the crosscap X to the orientable surface S. Observe that a closed curve in S' is one-sided if and only if it passes through X an odd number of times. Therefore a parity embedding in S' is an embedding in which the parity of cycle length equals the parity of the number of times it passes through X. Now, for *any* nonorientable surface with a specified crosscap X, we define an *X-parity embedding* to be an embedding in which a cycle is odd if and only if it passes through X an odd number of times. (For an orientable surface with one added crosscap X, parity embedding is identical to X-parity embedding.)

Now we can state our result, which generalizes Theorem 4.5.

Theorem 4.6. *G is a generalized thrackle on a surface S if and only if G has an X-parity embedding on the surface obtained by adding a crosscap X to S, with the same embedding scheme.*

Proof. Consider a generalized thrackle on a surface S. By Remark 4.2 we can redraw G so that every two edges cross oddly. Add a crosscap to S. We detour every edge through the crosscap as follows: First redraw a portion of each edge to be near the crosscap, without passing through any vertices. Then make the curve go halfway around the boundary of the crosscap and then through the crosscap. The first part of the detour does not change the parity of crossing between any two edges, and the second part adds exactly one crossing between each pair of edges. (See Figure 6.) Since a cycle of k edges now uses X exactly k times, we obtain an X-parity embedding of G if we can remove all edge crossings while maintaining, for each edge, the parity of the number of times it uses X.

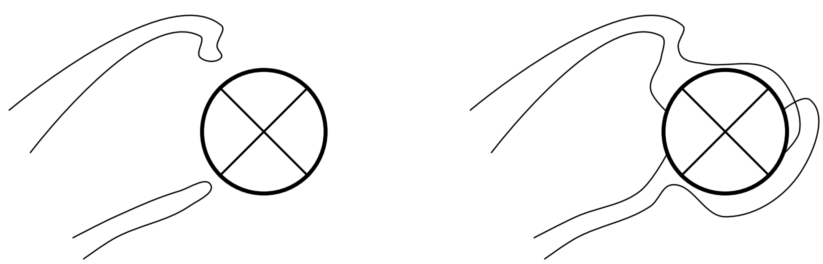

Figure 6: Detour through the crosscap

We cannot immediately apply Theorem 3.2 to remove crossings, since the proof of that theorem, while maintaining the embedding scheme, does not ensure that X is recovered as desired (or at all). Hence, before applying Theorem 3.2 we draw a one-sided closed curve D that lies along the boundary of X. As the proof proceeds, D is replaced by a set of curves $\mathcal{D}$ with ends at vertices of G.

We also strengthen the induction hypothesis of the proof of Theorem 3.2 so that (1) the rotation at each vertex, including all ends of $\mathcal{D}$, is preserved, and (2) the parity of the total number of crossings between $\mathcal{D}$ and an edge e is preserved, for every edge e.

Before contracting an edge e of G from v to u, we process the crossings between curves of $\mathcal{D}$ with e, in order along e starting with the closest to u: move a portion of the curve near the crossing from $\mathcal{D}$ along e to u, then modify $\mathcal{D}$ by splitting the curve in two where it intersects u; then add the two new curve ends to the rotation at u (see Figure 7). This only adds crossings between curves of $\mathcal{D}$ and edges that cross e (two crossings at a time), so the desired parity of total crossings between curves of $\mathcal{D}$ and edges is preserved. After e is contracted and induction is applied to recover e, we can recover $\mathcal{D}$ with the original number of crossings between $\mathcal{D}$ and e. This finishes the case of edge-contractions.

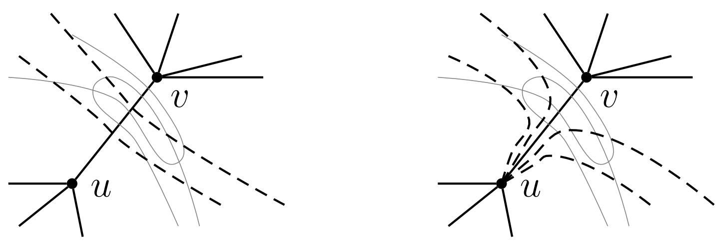

Figure 7: Rerouting curves in $\mathcal{D}$ (dashed) before contracting $e = uv$ (black). Initial situation *(left)*, adding new ends to u *(right)*. The parity of crossings of the other curves (gray) with all curves in $\mathcal{D}$ does not change.

The same method works when we join two single-vertex components of G, so let us consider the case that G has only one vertex, u. If S is the sphere, then parity of crossing between any two loops is entirely determined by the rotation at u; since each curve of $\mathcal{D}$ is drawn as a loop, this means that the redrawing cannot change parity of crossing between any curve in $\mathcal{D}$ and any edge in G. It remains to consider the case that G has a single vertex and S is not the plane.

As in the proof of Theorem 3.2, let C be a nonseparating curve that lies along an edge e. When edges are redrawn so that they do not cross e, every edge's parity of crossing with $\mathcal{D}$ is preserved. When we cut the surface along C, we also cut any curves of $\mathcal{D}$ at every crossing with C, obtaining a new set of curves $\mathcal{D}'$. When the newly attached disk or disks are contracted to a vertex, each end of a curve that was on C goes to the appropriate place in the rotation of a vertex of the new graph. After applying induction, as we recover the original graph and surface, $\mathcal{D}$ is naturally recovered as well.

This completes one direction of the proof.

Now suppose we have an X-parity embedding on the modified surface. Let

E_1 be the set of edges that pass through X an odd number of times, and let $E_0 = E(G) - E_1$.

Claim: $V(G)$ has a bipartition A, B such that $G - E_1$ is an A, B-bigraph and every edge of E_1 has either both endpoints in A or both in B. To see this, contract every edge in E_1, and let H be the resulting multigraph, with $E(H) = E_0$. If H contains an odd cycle, then there must be some cycle C in G that is contracted to the odd cycle. Then the length of C and the number of edges it has in E_1 differ by an odd number, making C violate the X-parity condition. Hence H is bipartite, with some bipartition A', B'. Let $A \subset V(G)$ be the vertices that are contracted to vertices in A', and define B' likewise. This yields A, B as desired.

Each time an edge uses X, detour the edge halfway around the boundary of X, redrawing so that no edge passes through X. Remove X; the new drawing D is now drawn on S. Observe that each pair $e, f \in E(G)$ crosses oddly if and only if $\{e, f\} \subseteq E_1$. We wish to redraw such that every pair of distinct edges crosses oddly—because then by Remark 4.2 we get a generalized thrackle and we are done.

We make use of the e, v-move introduced earlier; recall that it deforms e so it passes over v just once (and over no other vertex).

Step 1: For each pair of distinct vertices $u, v \in A$, choose one vertex, say u, and perform the e, v-move for every edge e incident to u. (This generalizes the procedure used in Lemma 4.3.)

Step 2: Twist the rotation at each $u \in A$.

As each pair u, v is considered in Step 1, parity changes between each pair of edges e, f where e is incident to u and f is incident to v. Thus after Steps 1 and 2, parity changes between e and f once per pair $u, v \in A$ (not necessarily distinct) such that u is an endpoint of e and v is an endpoint of f. Therefore parity of crossings changes an even number of times if either e or f have an even number of endpoints in A, i.e., if e or f is in E_1. For $e, f \in E_0$, parity changes exactly once. So after Steps 1 and 2, pairs $e, f \in E(G)$ cross oddly if both are in E_0 or if both are in E_1, and they cross evenly if one is in E_0 and the other is in E_1.

Step 3: Do the e, v-move for all $e \in E_1$ and $v \in A$.

As a result each $e \in E_1$ changes parity once with each edge in E_0 and 0, 2, or 4 times with each edge in E_1. Therefore, after Step 3 all pairs of edges cross oddly, which is what we had to show. ■

5 Removing Even Crossings on Arbitrary Surfaces

In Section 3 we gave a proof of the weak Hanani-Tutte theorem for surfaces. In the plane we know that a stronger result is true [12]:

Theorem 5.1 (Pach, Tóth)**.** *If D is a drawing of G in the plane, and E_0 is the set of even edges in D, then G can be drawn in the plane so that no edge in E_0 is involved in any crossings.*

Pach and Tóth applied their result to establish a relationship between two different notions of crossing numbers. The *crossing number*, $\mathrm{cr}(G)$, of a graph G is the smallest number of crossings in a drawing of G. The *odd crossing number*, $\mathrm{ocr}(G)$, is the smallest number of pairs of edges that cross oddly in a drawing of G. By definition $\mathrm{ocr}(G) \leq \mathrm{cr}(G)$; however there are graphs for which the two numbers differ [15]. On the other hand, Pach and Tóth showed that $\mathrm{cr}(G) \leq 2\,\mathrm{ocr}(G)^2$.

The redrawing procedure used in the proof of Theorem 5.1 can increase the odd crossing number, and therefore probably will not lead to better bounds of $\mathrm{cr}(G)$ in terms of $\mathrm{ocr}(G)$ (a linear bound is suspected). In a previous paper [14] we showed that Theorem 5.1 can be strengthened to avoid an increase in the odd crossing number in a strong sense:

Theorem 5.2 (Pelsmajer, Schaefer, Štefankovič)**.** *If D is a drawing of G in the plane, and E_0 is the set of even edges in D, then G can be drawn in the plane so that no edge in E_0 is involved in any crossings and there are no new pairs of edges that cross an odd number of times.*

As a consequence, we were able to show that crossing number and odd crossing number are the same when they are at most 3. Our goal is to extend as many results as possible from the plane to arbitrary surfaces. The following example, however, shows that the stronger Theorem 5.2 fails on any surface other than the sphere.

Example 5.3. We claim that for any surface S other than the sphere, there is a simple graph drawn on S with an even edge e, such that any redrawing with e free of crossings will have a new pair of edges that cross oddly. We describe an example for the torus and the projective plane only, but these can easily be modified for all "larger" surfaces, by adding edges to the graph that use up each extra crosscap and handle. We use the same graph for both the torus and the projective plane. We describe the graph as a single-vertex multigraph with fixed rotation system and later show how to turn it into a simple graph without rotation system.

Consider the multigraph G shown in Figure 8, drawn in the torus (left) and the projective plane (right). We first discuss the torus case. G consists of an even loop e and two pairs of loops, one pair with its ends inside of e (f_1, f_2 in Figure 8) and the other with ends outside of e (g_1 and g_2). The two loops in each pair alternate ends at the vertex, but do not cross each other. Each loop with ends inside of e crosses exactly one of the loops with ends outside of e oddly (and the other loop not at all). For a contradiction, assume that G is drawn with the given rotation at the vertex, so that no new pairs of edges cross oddly, and e is crossing-free. Then e is either contractible or nonseparating (there are no separating curves on the torus). If e is contractible, then one of the two pairs

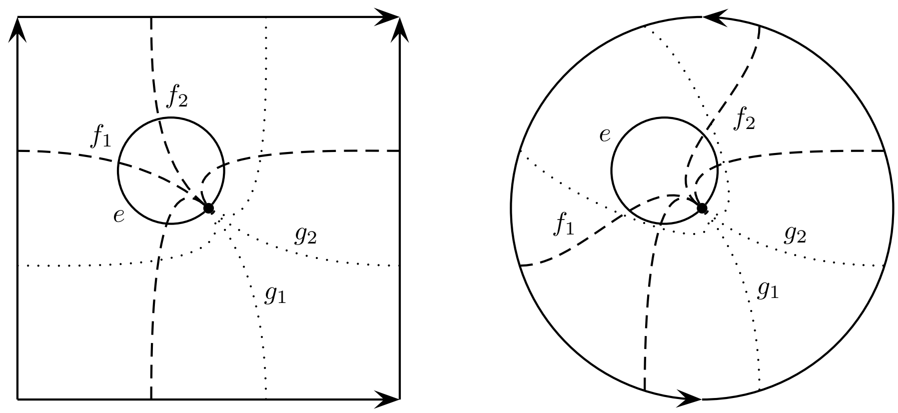

Figure 8: An example showing that Theorem 5.2 is not true for the torus (*left*) and the projective plane (*right*).

has to lie entirely within e, i.e. be embedded in a disk (a sphere with a hole). Consequently, the two loops in that pair must cross oddly, which they did not do before. If e is nonseparating, both pairs of loops are embedded in a region homeomorphic to an annulus (a sphere with two holes), which forces the loops in each pair to cross each other oddly. Again this changes which pairs of edges cross oddly.

The argument for the projective plane is nearly identical. As shown in the right part of Figure 8, G is drawn so that e is even and f_1 and f_2 each cross both g_1 and g_2 oddly. Consider a drawing of G in the projective plane in which e is crossing-free. If e is contractible then one pair of edges must be drawn in the interior of e, which is homeomorphic to a disk. Consequently, the two loops in that pair must cross oddly, giving us a new odd-crossing pair of edges. If e is not contractible it must be nonseparating, in which case both pairs of loops are embedded in a region homeomorphic to a disk, with the same result.

From this example we can create a simple graph F without rotation system for which Theorem 5.2 on the torus or projective plane fails. Erase the drawing in a small ball containing the vertex of G, and draw the wheel W_{10} there without crossings, such that its interior has ten 3-faces and one vertex of degree 10. The five edges of G have ten ends, which we can extend to meet each of the other vertices of W_{10}, without creating any crossings.

In the drawing of the resulting graph F the edges belonging to the W_{10} are free of crossings. Consider a drawing of F in which W_{10} and e are crossing-free. If W_{10} is embedded as in the original drawing (i.e. with ten triangular faces), then we can follow the multigraph argument above. Otherwise, one of the triangles of W_{10} does not bound a face, either because it is nonseparating or because the region it encloses contains the rest of W_{10}. In either case, the rest of F must be drawn in an annulus or a disk, which forces new odd pairs.

In these particular examples, it is easy to redraw the graph so that even edges become crossing-free and the odd crossing number does not increase. If this were true in general, then it would immediately follow that every graph has an ocr-optimal drawing in which every edge not involved in an odd crossing is actually crossing-free.

Question 5.4. Let D be a drawing of a graph G on a surface S, and let E_0 be the set of even edges in D. Is it always possible to redraw G in S so that no edge in E_0 is involved in any crossings and the odd crossing number of the drawing does not increase?

The original result of Pach and Tóth *is* true for arbitrary surfaces.

Theorem 5.5. *If D is a drawing of a graph G on some surface S, and E_0 is the set of even edges in D, then G can be drawn in S so that no edge in E_0 is involved in any crossings.*

Proof. Fix a drawing D of G in S. The proof will be a double induction over the genus of the surface S and the number of vertices of G. As usual we keep track of the embedding scheme. For the induction, we prove the following slightly stronger statement:

> If D is a drawing of a multigraph G on a surface S with E_0 the set of even edges, then G can be drawn in S so that edges of E_0 are crossing-free and the embedding scheme of the drawing is the same as that of D.

As in the proof of Theorem 3.2 we contract even edges while this is possible, maintaining the embedding scheme. In this way we obtain a graph in which all non-loop edges are odd, i.e. each crosses at least one other edge an odd number of times. We continue by contracting the odd edges as well. The important observation is that since all even edges are loops at this point, even edges remain even.

We are left with one or more vertices depending on whether G was connected or not; if there are multiple vertices, we merge the vertices into a single vertex as described in the proof of Theorem 3.2.

We are left with the case of a graph with a single vertex u and a bouquet of loops, some odd and some even. If the drawing contains a nonseparating even loop e, we proceed as in the proof of Theorem 3.2: we remove all crossings with e, cut along e and reduce the surface to a new surface, S', of smaller genus, apply induction to the new multigraph on that surface, and then we can recover the original graph on the original surface with the original embedding scheme.

If the drawing does not contain a nonseparating even loop, all even loops are separating. Consider the rotation at u and an even loop e at u. Since the surface is split by e into two pieces, any other loop f at u—which must cross e an even number of times—must begin and end on the same side of e. Hence the ends of e and f cannot alternate at u. This lets us redraw all the even loops in a very small neighborhood of u, without changing the rotation at u, so that no two of them cross each other. Because the rotation is unchanged, each even loop will still cross each odd loop an even number of times.

Consider an even loop e at u. If an odd loop f crosses e, we simply remove the segment between the first and last crossing with e and directly connect the two ends alongside the boundary of e (we argued earlier that the two ends are not separated by e). Repeating this finishes the proof. ∎

Observe that the proof potentially increases the odd crossing number of the drawing, both when contracting odd edges, as well as in the last step when reconnecting the ends of odd loops. Nevertheless, it is good enough to extend the result by Pach and Tóth that $\mathrm{cr}(G) \leq 2\,\mathrm{ocr}(G)^2$ in the plane to any surface, using essentially the same proof given by Pach and Tóth for the planar case [12].

Corollary 5.6. *For any surface S we have*

$$\mathrm{cr}_S(G) \leq 2\,\mathrm{ocr}_S(G)^2.$$

Proof. Let D be an ocr-optimal drawing of $G = (V, E)$ on surface S, i.e. a drawing realizing $\mathrm{ocr}_S(G)$. Let E_0 be the set of even edges in D. Using Theorem 5.5 we can obtain a drawing of G in which all edges of E_0 are free of crossings. In other words, only the edges in $E - E_0$ are involved in crossings, and there are at most $2\,\mathrm{ocr}_S(G)$ of them. Erase all of the edges in $E - E_0$ and redraw them so as to minimize the number of crossings between them. If any pair of edges crosses more than once, then it is easy to redraw with fewer crossings overall. Hence no pair of edges in $E - E_0$ crosses more than once, so the new drawing has crossing number at most $\binom{2\,\mathrm{ocr}_S(G)}{2} \leq 2\,\mathrm{ocr}_S(G)^2$. ∎

Remark 5.7. If the surface S is the sphere, then the proof of Theorem 5.5 can be simplified even further by removing the induction on the genus, giving a really simple proof of the fact that

$$\mathrm{cr}(G) \leq 2\,\mathrm{ocr}(G)^2$$

in the plane.

Even though Theorem 5.1 cannot be strengthened to an analogue of Theorem 5.2, it does allow us to derive a result on small crossing numbers; it is weaker than our result for the plane, where we could show that $\mathrm{ocr}(G) = \mathrm{cr}(G)$ whenever $\mathrm{ocr}(G) \leq 3$.

Theorem 5.8. *If G is a graph on a surface S with $\mathrm{ocr}_S(G) \leq 2$, then $\mathrm{ocr}(G) = \mathrm{cr}(G)$.*

Proof. We will prove by induction that a multigraph G with $\mathrm{ocr}(G) \leq 2$ can be redrawn with the same embedding scheme and at most $\mathrm{ocr}(G)$ crossings.

Fix an ocr-optimal drawing of G on S where $\mathrm{ocr}(G) \leq 2$, with the minimum possible number of odd edges. If there are no odd crossings, we apply Theorem 3.2. If there is an even non-loop, we contract it, apply induction, then uncontract. We deal with any non-separating even loop as in the proof of Theorem 5.5: we free it of all crossings, reduce the surface, apply induction, and then recover the original surface with the original multigraph drawn on it. In both cases no new crossings are added when recovering the original graph; thus we may assume that every non-loop is odd, and every even loop is separating.

Let G' be the subgraph consisting of all non-loops (a subset of the odd edges). Since each odd crossing involves two edges, G' has at most four edges.

If G' has a vertex v of degree 1, and $e = uv$ is the edge incident to v in G', then we modify G (not just G') by contracting v nearly to u along e. Since every other edge incident to v is a loop, this creates no odd crossings while removing all crossings from e, so we have a new drawing with smaller ocr, a contradiction. If, instead, v is incident to exactly two edges e, f in G', we apply the same procedure, which transfers all crossings with e to f. Edges that used to cross e oddly may now cross f oddly, but there can be no other new odd pairs, so ocr has not increased. However, the number of odd edges has decreased, which again contradicts the choice of drawing. Thus we can assume that G' has no vertices of degree 1 or 2.

If G' has a vertex of degree greater than 2, then its component must contain 3 or 4 edges and exactly two vertices; hence there is at most one component with edges of G'. If G has more than one component then we take one which is a single vertex with loops, and move it to join another component (as in previous proofs), which creates no odd crossings; we apply induction and separate the graph into the original graphs without adding crossings. Hence, we can assume that G is connected, with at most two vertices.

Consider the rotation at a vertex u. If there is a loop e whose ends are consecutive at u, then we can draw $G - e$ by induction with at most $\mathrm{ocr}(G)$ crossings and redraw e close to u without crossings. We proceed similarly if the rotation contains a word of the form efe where e is an odd loop: Since e is part of an odd pair, removing e lowers ocr, so applying induction to $G - e$ yields a drawing with at most $\mathrm{ocr}(G) - 1$ crossings; we can then draw e near u with exactly one additional crossing (with f). Thus we may assume that neither of these situations occur.

Suppose that e is an even loop at a vertex u. Since e separates the surface into two sides and e is even, any edge that starts on one side of e must end on the same side of e. In particular, if we write the rotation at u as a cyclic word $w_e e w'_e e$ (where w_e, w'_e are words over $E(G)$), then no edge appears in both w_e and w'_e. Similarly, G cannot have non-loop edges appear in both w_e and w'_e since their other ends would have to be on opposite sides of e; however, both ends must connect to the same vertex $v \neq u$ which cannot lie on both sides of e. Thus, if any non-loop appears in w_e, then all non-loops appear in w_e.

Now, let e be an even loop at u such that $m_e = \min(|w_e|, |w'_e|)$ is as small as possible; we may assume that $|w_e| = m_e$. Then no even edges appear in w_e, since w_e would have to contain both ends, contradicting the choice of e. If any non-loop edges end in w_e then they all do, so $m_e \geq 3$; if only loops end in w_e then we also have $m_e \geq 3$, since otherwise either e or one odd loop ending in w_e would have consecutive ends at u. Now choose f to be an even loop at u (possibly $f = e$) such that w_f is a subword of w'_e and $|w_f|$ is minimized. Then no even edges may appear in w_f (as we argued for w_e), and $|w_f| \geq m_e$ by the choice of e.

Using $|w_f| \geq m_e = |w_e| \geq 3$, we can finish this case. If G has two vertices then all non-loop edges appear together in w_e or w'_e, so there is at most one odd loop which can contribute to the other word; hence either $|w_e| \leq 2$ or $|w_f| \leq 2$, a contradiction. If u is the only vertex, then $|w_f|, |w_e| \geq 3$ implies that w_e and w_f each involve more than one odd loop; then previous arguments (no even ends in w_e or w_f, both ends of each odd edge are on the same side of e, and no consecutive ends) show that $w_e = abab$ and $w_f = cdcd$ where a, b, c, d are odd loops. After removing all odd edges, Theorem 3.2 lets us embed the remaining graph with the same rotation at u; then we can add a, b, c, d drawn near u with just two crossings (a crossing b and c crossing d). Clearly $\mathrm{ocr} \geq 2$ when there are more than two odd edges, so this suffices.

Thus we may assume that G has no even loops. If G has two vertices, with an odd loop e at a vertex u and exactly three non-loop edges, then the rotation at u must contain either ee or efe as a subword (for some non-loop edge f), which we already ruled out. If G consists of three or four edges between vertices u and v, then $\mathrm{ocr} \geq 2$ so it suffices to give a drawing with at most two crossings, and for any given rotations at u and v, this is easy. So we may assume that G has just one vertex, u.

Using the fact that the rotation at u has no subwords of the form ee or efe, one can show that the rotation has the form *abcabc*, *abcdabcd*, *abcadbcd*, or *abcadcbd*. If S is the sphere, then $\mathrm{ocr}(G)$ is the number of pairs of edges whose ends alternate in the rotation; in each of these four cases this implies that $\mathrm{ocr}(G) \geq 3$. Therefore S is not the sphere, and the surface has either a crosscap or a handle. Using a crosscap or a handle, the graph with rotation *abcabc* can be drawn without crossings, and in the other three cases the graph can be drawn such that there are no crossings between edges from the set $\{a, d\}$ and edges from $\{b, c\}$, and so that each pair a, d and b, c crosses at most once. The drawing has at most 2 crossings; since there are more than two edges, $\mathrm{ocr} \geq 2$ and we are done. ∎

References

[1] Unsolved problems. In *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972)*, pages 351–363. Inst. Math. Appl., Southend, 1972. Chairman: P. Erdős.

[2] Grant Cairns and Yury Nikolayevsky. Bounds for generalized thrackles. *Discrete Comput. Geom.*, 23(2):191–206, 2000.

[3] Grant Cairns and Yury Nikolayevsky. Generalized thrackle drawings of non-bipartite graphs. *Discrete Comput. Geom.*, 41(1):119–134, 2009.

[4] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.

[5] Chaim Chojnacki (Haim Hanani). Über wesentlich unplättbare Kurven im drei-dimensionalen Raume. *Fundamenta Mathematicae*, 23:135–142, 1934.

[6] Heiko Harborth. Parity of numbers of crossings for complete n-partite graphs. *Math. Slovaca*, 26(2):77–95, 1976.

[7] Daniel J. Kleitman. A note on the parity of the number of crossings of a graph. *J. Combinatorial Theory Ser. B*, 21(1):88–89, 1976.

[8] L. Lovász, J. Pach, and M. Szegedy. On Conway's thrackle conjecture. *Discrete Comput. Geom.*, 18(4):369–376, 1997.

[9] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.

[10] Serguei Norine. Pfaffian graphs, T-joins and crossing numbers. *Combinatorica*, 28(1):89–98, 2008.

[11] János Pach. Geometric graph theory. In *Handbook of discrete and computational geometry, 2nd edition*, CRC Press Ser. Discrete Math. Appl., pages 219–238. CRC, Boca Raton, FL, 2004.

[12] János Pach and Géza Tóth. Which crossing number is it anyway? *J. Combin. Theory Ser. B*, 80(2):225–246, 2000.

[13] Michael J. Pelsmajer, Marcus Schaefer, and Despina Stasi. Hanani-Tutte on the projective plane. Technical Report TR08-003, DePaul University, April 2008.

[14] Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Removing even crossings. *J. Combin. Theory Ser. B*, 97(4):489–500, 2007.

[15] Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Odd crossing number is not crossing number. In Patrick Healy and Nikola S. Nikolov, editors, *Graph Drawing*, volume 3843 of *Lecture Notes in Computer Science*, pages 386–396. Springer, 2005.

[16] László A. Székely. A successful concept for measuring non-planarity of graphs: the crossing number. *Discrete Math.*, 276(1-3):331–352, 2004. 6th International Conference on Graph Theory.

[17] William T. Tutte. Toward a theory of crossing numbers. *J. Combinatorial Theory*, 8:45–53, 1970.