

The Degenerate Crossing Number and Higher-Genus Embeddings

Marcus Schaefer Daniel Štefankovič

Abstract. If a graph embeds in a surface with k crosscaps, does it always have an embedding in the same surface in which every edge passes through each crosscap at most once? This well-known open problem can be restated using crossing numbers: the degenerate crossing number, $\text{dcr}(G)$, of G equals the smallest number k so that G has an embedding in a surface with k crosscaps in which every edge passes through each crosscap at most once. The genus crossing number, $\text{gcr}(G)$, of G equals the smallest number k so that G has an embedding in a surface with k crosscaps. The question then becomes whether $\text{dcr}(G) = \text{gcr}(G)$, and it is in this form that it was first asked by Mohar.

We show that $\text{dcr}(G) \leq 3\text{gcr}(G)$, and $\text{dcr}(G) = \text{gcr}(G)$ as long as $\text{dcr}(G) \leq 3$. We can separate dcr and gcr for (single-vertex) graphs with embedding schemes, but it is not clear whether the separating example can be extended into separations on simple graphs. We also show that if a graph can be embedded in a surface with crosscaps, then it has an embedding in that surface in which every edge passes through each crosscap at most twice. This implies that dcr is **NP**-complete.

Finally, we extend some of these results to the orientable case (and bundled crossing numbers).

Keywords: degenerate crossing number, non-orientable genus, genus crossing number, bundled crossing number.

1 Introduction

When defining the (standard) crossing number of a graph (a perilous activity, see [18]), one typically requires that at most two edges cross in any point. If $k > 2$ edges cross in a single point, these edges can be perturbed slightly to create $\binom{k}{2}$ crossings of pairs of edges, so multiple crossings in a single point can always be avoided if crossings are counted pairwise. Günter Rote and M. Sharir (according to Pach and Tóth [16]) asked “what happens if multiple crossings are counted only *once*”. This led Pach and Tóth to introduce the degenerate crossing number: we allow drawings which are *degenerate* in the sense that more than two edges are allowed to cross in a single point (but which are otherwise standard, in particular, edges have to actually cross, not touch, and self-crossings are not allowed). The *degenerate crossing number* of the drawing is the number of crossing points in the drawing. The *degenerate crossing number*, $\text{dcr}(G)$, of a graph G is the smallest degenerate crossing number of any degenerate drawing of G . Some papers (e.g. [1]) restrict drawings to be simple, that is, every two edges intersect at most once; to distinguish this variant from dcr we call it the *simple degenerate crossing number*, $\text{dcr}^*(G)$.¹

If we relax the definition of the degenerate crossing number to allow self-crossings of edges, we obtain the *genus crossing number*, $\text{gcr}(G)$, which was introduced by Mohar [14]. By definition, $\text{gcr}(G) \leq \text{dcr}(G)$. Mohar conjectured that $\text{gcr}(G) = \text{dcr}(G)$ for all G . Equality of these two numbers would be particularly interesting, since, as Mohar observes, $\text{gcr}(G) = \tilde{\gamma}(G)$, where $\tilde{\gamma}(G)$ is the *non-orientable genus*, also known as the *minimum crosscap number*, of G , the smallest number k so that G can be embedded on a surface with k crosscaps (we allow the special case of $k = 0$ for planar graphs). Each crossing of multiple edges can be replaced by a crosscap and vice versa, since edges have to cross (and may not touch) in a crossing point. Similarly, $\text{dcr}(G)$ can be viewed (as we did in the abstract) as the smallest number k so that G has an embedding on a surface with k crosscaps so that every edge passes through each crosscap at most once. An edge not being allowed to pass through a crosscap more than once corresponds to prohibiting self-crossings in degenerate drawings in the plane. We view crosscaps as geometric, rather than purely topological objects, a view which we believe makes sense in graph drawing, where we need to visualize objects concretely. There is a more topological way to make precise the notion of “passing through a crosscap at most once”. Mohar [14] uses a “planarizing system of disjoint 1-sided curves”, abbreviated PD1S, where a system of disjoint curves is *planarizing* if cutting along all the curves leaves one with a sphere with holes. He observes that the genus crossing number of a graph is at most k if and only if there is an embedding of the graph in a surface with k crosscaps, and a PD1S system in that surface so that every curve in the system crosses every edge of the graph at most once.

We do not yet know, whether $\text{gcr}(G) = \text{dcr}(G)$ in general, but we can separate the two crossing numbers, if we are allowed to equip graphs with a rotation system (a fixed rotation at each vertex) or an embedding scheme (a rotation system, and a signature for each edge). In that case, there are graphs for which gcr is 3, but dcr is 4 as we will see in Theorem 7.

1.1 Visualizing Graphs in Higher-Order Surfaces

The exact relationship between gcr and dcr has consequences for visualizing graphs embeddable in higher-order surfaces in the plane. Typically, such graphs are visualized using a (canonical) polygonal schema. There are polynomial-time algorithms for this task, e.g., see [8] for orientable

¹The term *simple* has also been used to refer to drawings in which every two edges cross at most once; the difference is that a shared endpoint counts as an intersection, but not a crossing. An example in the entry on degenerate crossing number in [18] shows that it matters whether dcr^* is defined so as to allow crossings between adjacent edges or not.

surfaces, also see [7, 13, 9]. Many visualization algorithms start by contracting the graph to a single-vertex graph with an embedding scheme; for these algorithms, the example in Theorem 7 shows that edges can be forced to use the same topological feature more than once.

On the other hand, we can show that $\text{dcr}(G) \leq 3 \text{gr}(G)$, so any graph embeddable in a surface with k crosscaps can be embedded in a surface with at most $3k$ crosscaps so that every edge passes through each crosscap at most once. We will establish this in Theorem 5. If we allow an edge to pass through each crosscap just twice, it turns out that every graph can then be embedded in a surface with $\tilde{\gamma}(G) = \text{gr}(G)$ crosscaps (Theorem 10).

In Section 6 we will see that similar results can be obtained for the orientable case, which is related to the bundled crossing number.

1.2 Related Results

Pach and Tóth [16] showed that $\text{dcr}(G) < |E(G)|$. For the simple degenerate crossing number, a crossing lemma is known: $\text{dcr}^*(G) \geq c \cdot |E(G)|^3 / |V(G)|^2$ for $|E(G)| \geq 4|V(G)|$ (and some constant $c > 0$). This was shown by Ackerman and Pinchasi [1], improving an earlier result by Pach and Tóth. We should also mention work by Harborth [11], who may have been the first to study multiple crossings in drawings. His goal was to maximize the number of multiple crossings involving many edges. Let us call a degenerate crossing an *m-fold crossing* if it involves m edges. Harborth showed that K_{2m} can be drawn with two m -fold crossings; he conjectured that K_{2m} cannot be drawn with three or more m -fold crossings.

2 Terminology and Tools

2.1 Curves in Surfaces

A *surface* is a compact 2-manifold, with or without boundary. By the classification theorem of surfaces, every surface without boundary is homeomorphic to a sphere with h handles or k crosscaps [15, Theorem 3.1.3], so surfaces come in two *orientability* types: a surface S with h handles is called *orientable* with (*orientable*) *genus* h , or $\gamma(S) = h$; and a surface with $k \geq 1$ crosscaps is *non-orientable* with *non-orientable genus* k , or $\tilde{\gamma}(S) = k$, for short. We will also work with the *Euler genus*, $\text{eg}(S)$ of a surface, which is $2h$ for orientable surfaces, and k for non-orientable ones. For example, if $\text{eg}(S) = 2$, then S is either a torus, or a Klein bottle. If S is a surface with ℓ boundary components, the classification theorem still applies except that we need to replace the initial sphere with a sphere with ℓ holes [20, Section 1.3].

We start with some basic terminology for (simple) closed curves on a non-orientable surface S . A closed curve C is called *non-separating* if $S - C$ consists of a single component. Otherwise, C is *separating*. If it is separating, it can be *contractible* (one of the two pieces is homeomorphic to a disk) or *surface-separating* (also known as *splitting*). The *sidedness* of a closed curve is the number of sides it has: it is either *one-sided* (its neighborhood is a Möbius strip) or *two-sided* (its neighborhood is an annulus). A closed curve C in a non-orientable surface is *orienting* if $S - C$ is orientable.²

The following lemma is a well-known consequence of the classification of surfaces with boundary.

²There seems to be no standard name for curves of this type in the literature. Bojan Mohar suggests “orienting”; in the conference version of this paper, we used “maximal” because of the characterization given in Lemma 4.

Lemma 1 *If two non-separating simple, closed curves C and D on a surface S are such that $S - C$ and $S - D$ have the same orientability type, and C and D have the same sidedness (both one-sided or both two-sided), then there is a homeomorphism of S taking C to D .*

For the proof we use one more consequence of the classification of surfaces: up to a homeomorphism, a surface is determined by its Euler genus, its orientability type and the number of boundary components.

Proof: Suppose we have a simple, non-separating curve C on a surface S with Euler genus $\text{eg}(S)$. If we cut S along C , then the Euler genus of $S - C$ is $\text{eg}(S) - 2$ if C is two-sided, and $\text{eg}(S) - 1$ if C is one-sided (this follows from Euler's formula, see, for example, the proof of Lemma 3.1.4 in [15]).

We conclude that $S - C$ and $S - D$ have the same Euler genus. By assumption, they also have the same orientability type, and the same number of holes (one if C and D are one-sided, two if they are two-sided). The classification theorem then implies that they are homeomorphic, which allows us to construct the required homeomorphism of S mapping C to D . \square

A surface can contain only a small number of different types of disjoint simple, closed curves. The following lemma makes this precise.

Lemma 2 (Malnič, Mohar [15, Proposition 4.2.7]) *If G is a graph embedded in a surface S , and \mathcal{P} is a collection of internally disjoint paths between vertices a and b (where $a = b$ is allowed), so that no two of the paths bound a disk in S , then*

$$|\mathcal{P}| \leq \begin{cases} 3 \text{eg}(S) - 2 & \text{if } \text{eg}(S) \geq 2 \\ \text{eg}(S) + 1 & \text{otherwise.} \end{cases}$$

Remark 1 *We are interested in the case where $a = b$ and there are no surface separating paths; a better upper bound for that case would improve the upper bounds in Theorem 5 and Theorem 12.*

We conclude this section by characterizing separating and orienting curves in terms of how often they pass through crosscaps.

Lemma 3 *A simple, closed curve on a non-orientable surface is separating if and only if it passes through every crosscap an even number of times.*

In the proof we refer to pushing a curve off a crosscap; we do this as follows: instead of having the curve pass through the crosscap, we sever it just before and after it does so, and reconnect the two severed ends by routing the curve along one side of the crosscap; we do this for any strand of the curve passing through the crosscap. Figure 3(b)-(c) illustrates pushing a single strand of a curve off a crosscap.

Proof: Suppose C is a simple, closed curve that passes through every crosscap an even number of times. Let p be a point on C . If C is non-separating, then there must be a way to connect p to itself by a closed curve D that starts at p on one side of C and ends in p . Then C and D intersect once. Since C passes through each crosscap an even number of times, we can push it off every crosscap without changing the parity of crossing between C and D , so they still cross oddly. We can then push D off each crosscap it passes through. This again does not affect the parity of crossing, since C no longer passes through any crosscap. We end up with two closed curves in the plane (not necessarily simple) that cross an odd number of times, which is impossible. We conclude that C must be separating.

For the other direction, let us assume that C is separating, and color the two parts of the surface red and green. Along the boundary of each crosscap, the two colors must alternate. Since the crosscap connects opposite regions along its boundary, they must have the same color. From that it follows that C must pass through each crosscap an even number of times. \square

Any non-orientable surface can be made orientable by cutting through each crosscap once. The following lemma shows that (up to parity) this captures orienting curves precisely.

Lemma 4 *A simple, closed curve on a non-orientable surface is orienting if and only if it passes through every crosscap an odd number of times.*

Proof: Suppose C is a simple, closed curve on a non-orientable surface S such that C passes through some crosscap c an even number of times. We draw a simple, closed curve D : start close to c , pass through c once, and then follow the boundary of c until returning to the starting point. Curve D is one-sided, and it crosses C an even number of times (possibly not at all). Let p_1, \dots, p_{2k} be the crossings between C and D as one encounters them along a traversal of C , starting at an arbitrary point. We now remove crossings of D and C one pair p_{2i}, p_{2i+1} at a time. We move the crossing p_{2i} along C in the original direction of traversal towards p_{2i+1} until p_{2i} and p_{2i+1} are arbitrarily close (we extend D by closely following C on either side). We then erase D in a neighborhood of p_{2i} and p_{2i+1} and reconnect the severed ends on both sides of C , removing two crossings between C and D . This process turns D into a collection (possibly a singleton) of simple, closed curves, none of them crossing C . The overall parity of crossing through crosscaps does not change, so at least one of the resulting curves must be one-sided. But this curve remains in $S - C$, so C cannot be orienting.

For the other direction, we assume that C passes through every crosscap an odd number of times. If $S - C$ is not orientable, it must contain a one-sided curve D . Then D is also a one-sided curve in S , disjoint from C . Push D off of all crosscaps it passes through; since D is one-sided this requires an odd number of pushes, so C and D cross an odd number of times at this point. We can now push C off of any crosscaps it passes through, without changing the parity of crossing between C and D , giving us two closed curves in the plane that cross oddly, which is a contradiction, showing that $S - C$ is orientable. \square

2.2 Drawings and Embeddings

In a *drawing* of a graph in a surface S , or the plane, every vertex is drawn as a distinct point of S , and edges are realized as (typically simple) curves connecting their endpoints. We assume that edges do not pass through vertices and edges intersect at most finitely often. If two edges e and f intersect in a point, they either cross (intersect transversally) or touch. We do not allow edges in our drawings to touch. We also generally assume that at most two edges cross in a point; if more than two edges cross in a point, we call this a *multiple crossing*, and the drawing *degenerate*. If a graph can be drawn in a surface without any crossings, we say that the graph can be *embedded* in the surface, and we refer to the drawing as the *embedded graph* or the *embedding*.

We require some tools from topological graph theory which are used to describe embeddings in orientable and non-orientable surfaces. A *rotation* at a vertex is a cyclic permutation of the ends of edges incident on the vertex. Since we will often work with single-vertex graphs, we implicitly direct edges so we can distinguish their two ends at the same vertex. A *rotation system*, ρ , of a graph prescribes a rotation for every vertex. We say a drawing of a graph G in an orientable surface

realizes a given rotation system ρ if the clockwise order of ends around each vertex corresponds to the cyclic permutation prescribed by ρ .³

On non-orientable surfaces, we also prescribe, for every edge, its *signature*, which is a number in $\{-1, 1\}$. A rotation system ρ and signature λ together form an *embedding scheme* (ρ, λ) of a graph. With a bit of machinery, one can define what it means for a graph drawing to realize an embedding scheme purely topologically. Instead we opt to work in a more geometric model. To that end we localize the crosscaps and remove a point from the sphere to obtain a drawing in a plane with geometrically drawn crosscaps; see Figure 1 for an example.⁴

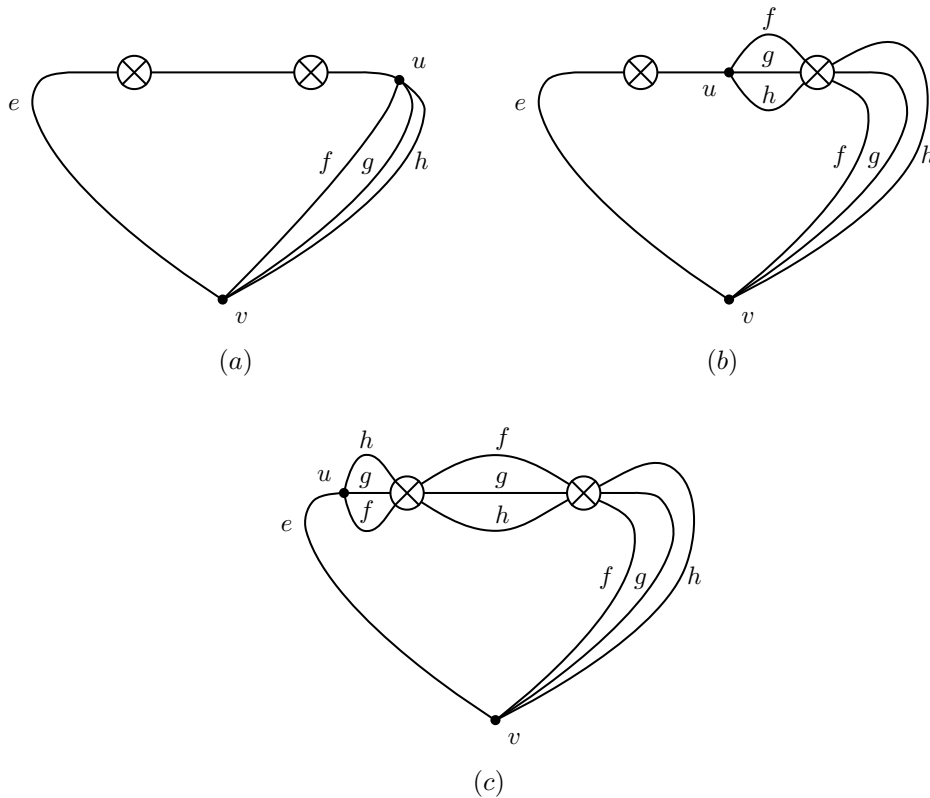


Figure 1: Three embeddings of a graph in a Klein bottle obtained from each other by moving u along a to v . (a) All edges have signature 1. (b) All edges have signature -1 . (c) All edges have signature 1 again.

In this geometrized crosscap model we say that a drawing (or embedding) of a graph G realizes an embedding scheme (ρ, λ) , if the rotation at each vertex is as prescribed by ρ , and the signature of an edge is 1 if and only if that edge crosses through crosscaps an even number of times. To

³Rotation systems are often defined as rotation systems of given embeddings or drawings. That approach is more intuitive, but it disguises the fact that rotation systems are purely combinatorial objects that exist independently of an embedding or drawing of the graph.

⁴To localize the crosscaps we take a planarizing set of disjoint one-sided curves and map them to the boundaries of disks which we fill with our geometric crosscaps.

shorten statements we will also speak of a drawing (or embedding) of (G, ρ, λ) meaning a drawing (or embedding) of G which realizes (ρ, λ) . With this definition, a cycle in a drawing of (G, ρ, λ) is two-sided if the signature of its edges multiply to 1, otherwise, it is one-sided.

The three embeddings pictured in Figure 1 are all the same embedding up to a homeomorphism of the surface; the crosscaps were localized in different ways. To capture this we call two embedding schemes for the same graph *equivalent* if they can be obtained from each other by *flipping* a vertex, which means reversing the rotation at the vertex (inverting the cyclic permutation of edge ends at the vertex), and changing the signature of each end incident to the vertex (as a result, a loop at the vertex gets changed twice, so overall its signature does not change). The three embedding schemes shown in Figure 1 are equivalent. The geometrized crosscap model allows us to identify an outer face (which will occasionally be useful in constructions).

Typical operations on graphs such as removing or adding a vertex or edge, and contracting an edge are easily performed purely combinatorially on the embedding scheme. For details, see [15, Section 3.3]. If in Figure 1(c) we contract edge e entirely, we obtain the single-vertex graph pictured in Figure 2. As long as we do not change the rotation at v we can later move back from the single-vertex graph to the graph in Figure 1(c) by *uncontracting* edge e . In other words, we can split v into u and v and reattach the ends to reconstruct the original rotation system, and assign edge e a signature of 1.

The *Euler genus* $\text{eg}(G, \rho, \lambda)$ of a graph G with embedding scheme (ρ, λ) is defined as $2 - |V(G)| + |E(G)| - |F(G, \rho, \lambda)|$, where $|F(G, \rho, \lambda)|$ is the number of faces in the embedding scheme. This is a purely combinatorial parameter, which can be computed from G and (ρ, λ) without surfaces being involved. By the Euler-Poincaré formula, $\text{eg}(G, \rho, \lambda)$ is a lower bound on the Euler genus of a non-orientable surface in which (G, ρ, λ) has an embedding.

Theorem 2 (Euler-Poincaré) *If (G, ρ, λ) can be embedded in a surface with k crosscaps, then $k \geq \text{eg}(G, \rho, \lambda)$.*

The theorem states that $\text{gcr}(G, \rho, \lambda) \geq \text{eg}(G, \rho, \lambda)$, and it is tempting to assume that the two values are equal, but that is not actually true; take, for example, a single vertex with two two-sided edges alternating at the vertex. The Euler genus of this graph is 2, while it requires 3 crosscaps to realize. We clarify the relationship between gcr and eg in Lemma 9.

2.3 Working with Single-Vertex Graphs

Arguments and algorithms for graph embeddings can often be simplified by replacing an embedded graph with a single-vertex graph with embedding scheme. This is often done for visualizing embeddings of graphs in higher-genus surfaces in the plane (see Section 1.1). In a single-vertex graph every edge is a loop, hence a closed curve, and we can talk about its sidedness, which then directly corresponds to its signature: a one-sided loop has signature -1 , and a two-sided loop signature 1. Figure 2 shows an embedding of a single-vertex graph obtained from the graph shown in Figure 1 by contracting edge e . In this case, the two crosscaps are not necessary, since the graph has a drawing realizing the same embedding scheme in the plane (or sphere).

The ends of two loops e and f at the same vertex can either be *parallel*, that is, in order $eeff$, or *alternate*, that is in order $efef$, in which case they must cross or use topology. The loops in Figure 2 are all pairwise parallel.

A loop separates the rotation at a vertex into two halves. By orienting the loop, we can distinguish between the two halves. Hence an oriented loop encloses a specific half of the rotation, which we call a *wedge*. For example, in Figure 2 the loop f encloses two wedges (depending on

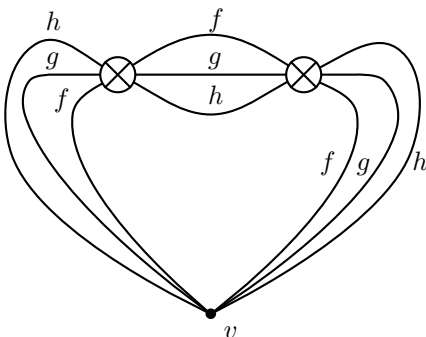


Figure 2: Embedding of a single-vertex graph in a Klein bottle.

how we orient it), one of them is empty, the other is $ghhg$. *Reversing* a wedge means reversing the rotation of the ends occurring in that wedge. If the rotation at v is $afghahgf$ (imagine adding a loop a to the rotation of v in Figure 2), and we are looking at the wedge fgh enclosed by a , then reversing that wedge results in the rotation $ahgfahgf$ at v .

For a graph G with embedding scheme (ρ, λ) , we define $\text{gcr}(G, \rho, \lambda)$ as the smallest number k so that G has an embedding realizing (ρ, λ) on a surface with k crosscaps. Similarly, $\text{dcr}(G, \rho, \lambda)$ is the smallest number k so that G has an embedding realizing (ρ, λ) on a surface with k crosscaps, and so that every edge passes through every crosscap at most once (since our crosscaps are geometrically localized, this is well-defined).

The following tool from topological graph drawing is helpful in simplifying some of the arguments in this section.

Theorem 3 (Weak Hanani-Tutte Theorem for Surfaces [4, 17]) *If G is drawn in a surface so that every pair of edges crosses an even number of times, then G has an embedding on the same surface with an equivalent embedding scheme.*

The next lemma shows that as far as gcr and Euler genus are concerned, we can replace a graph with a graph on a single vertex equipped with an embedding scheme. For dcr , we can do so for upper bounds only.

Lemma 5 *For every graph G there is an embedding scheme (ρ, λ) and a single-vertex graph G' with embedding scheme (ρ', λ') so that $\text{gcr}(G) = \text{gcr}(G, \rho, \lambda) = \text{gcr}(G', \rho', \lambda')$, and $\text{dcr}(G) = \text{dcr}(G, \rho, \lambda) \leq \text{dcr}(G', \rho', \lambda')$. Moreover, $|E(G')| \leq |E(G)|$, and (G, ρ, λ) can be derived from (G', ρ', λ') by uncontracting edges and deleting edges.*

Proof: Fix an embedding of G on a surface S with $k = \text{gcr}(G)$ crosscaps; let (ρ, λ) be the embedding scheme for that embedding, so $\text{gcr}(G) = \text{gcr}(G, \rho, \lambda)$. Choose a spanning forest F of G . Contract edges of F , merging rotations in the embedding scheme at vertices that are identified and updating signatures of edges. If there is more than one vertex left, let ζ be a two-sided curve connecting two of the vertices and so that ζ intersects the drawing only finitely often. We add ζ as an edge to the graph, and then contract that edge as above (since ζ is two-sided, we do not need to flip one of the rotations). Continue doing that until we obtain a drawing of a single-vertex graph G' with embedding scheme (ρ', λ') . When we added an edge along a curve ζ , we introduced

crossings into the drawing. But we then contracted the edge to a point, and all other edges incident to endpoints of the edge being contracted were loops, so the parity of crossing between no pairs of edges in G is changed (after the contraction). In particular, every two edges in the drawing of G' cross evenly. Hence, by Theorem 3, there is an embedding realizing (G', ρ', λ') in the surface S .

Hence, $\text{gcr}(G', \rho', \lambda') \leq k = \text{gcr}(G)$. By undoing the operations which turned (G, ρ, λ) into (G', ρ', λ') , namely by uncontracting contracted edges, and deleting edges we added to connect the graph, we can turn an embedding of (G', ρ', λ') into an embedding of (G, ρ, λ) on the same surface, so it follows that $\text{gcr}(G, \rho, \lambda) \leq \text{gcr}(G', \rho', \lambda')$, so $\text{gcr}(G) = \text{gcr}(G', \rho', \lambda')$. By the same argument, we can show that $\text{dcr}(G) = \text{dcr}(G, \rho, \lambda) \leq \text{dcr}(G', \rho', \lambda')$.

Since every loop of G' originates as an edge in G , we have $|E(G')| \leq |E(G)|$. □

Note that we do not claim that $\text{dcr}(G) \geq \text{dcr}(G', \rho, \lambda)$, the construction we used may force an edge through a crosscap multiple times, so dcr can increase. Lemma 5 allows us to replace a graph with a single-vertex graph when showing that dcr can be bounded in gcr .

We close this section with some results on orientable embeddings. An embedding of a graph G in a surface is *orientable* if all cycles in G are two-sided. If the embedding realizes an embedding scheme (ρ, λ) this is equivalent to saying that multiplying the signatures of edges along each cycle of G , one always gets 1. If G is a single-vertex graph, then its embedding is orientable, if all loops have signature 1. (On orientable surfaces, all embeddings are orientable.)

In some of the arguments we will talk about pushing an edge over and off a crosscap. Pushing an edge *over a crosscap* is done by creating a narrow band between the edge and the crosscap, rerouting the edge in that band so it approaches the crosscap, and then passing it through the crosscap. Pushing an edge *off a crosscap* means severing the edge where it passes through the crosscap and reconnecting any severed ends by routing around the crosscap. The effect of either move is to change the parity of crossing between the edge and every edge that passes through the crosscap an odd number of times; the parity of crossing between all other edge pairs remains unchanged. The move may introduce self-intersections of an edge, but those can always be removed locally. See Figure 3 for an illustration, and [17] for more information (including a figure illustrating how to remove a self-intersection).

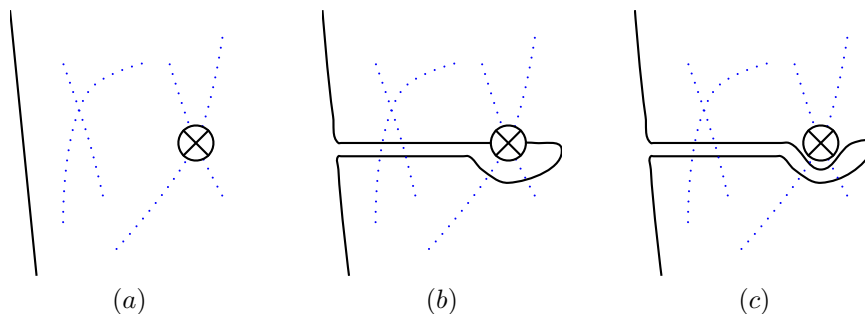


Figure 3: (a) Edge e and crosscap. (b) pushing e over the crosscap. (c) Pushing e off the crosscap.

It is well-known that with respect to surface homeomorphisms, a handle is equivalent to two crosscaps in the presence of another crosscap [15, Proposition 3.1.2]: a graph embeddable on a surface with h handles can be embedded on a surface with $2h + 1$ crosscaps so that every cycle is two-sided. The following lemma shows that the odd number of crosscaps is not accidental when we are restricted to orientable embeddings

Lemma 6 *Suppose k is minimal so that the connected graph G with rotation ρ has an orientable embedding on a surface with k crosscaps. Then either $k = 0$, or $k \geq 3$ and k is odd.*

Proof: Fix an orientable embedding of G realizing (ρ, λ) , for some signature λ , in a surface with k crosscaps, where k is minimal. We can assume that G is a single-vertex graph (contract edges of a spanning tree, this leaves the embedding orientable, so $\lambda(e) = 1$ for all loops e now). Suppose k is even. Let c be one of the crosscaps. For any edge that passes oddly through c , push that edge over all crosscaps. Note that pushing an edge over all crosscaps does not change the parity of crossing between any pair of edges since the number of crosscaps is even and every edge initially crosses through an even number of crosscaps oddly, and this remains true. At the end of this operation we have a drawing of G in which every pair of edges crosses an even number of times, and all edges pass through c an even number of times. We can then push all edges off of c , again maintaining that every pair of edges crosses evenly. Now, by Theorem 3, (G, ρ, λ) has an orientable embedding in the surface with $k - 1$ crosscaps, so k cannot have been even if it was minimal. If $k = 1$, then an orientable embedding on the projective plane implies that the graph is planar (since every edge passes through the single crosscap an even number of times). \square

One consequence of Lemma 6 is that an orientable embedding in a non-orientable surface is never maximal in the sense that we can always add a one-sided loop to it.

Corollary 4 *If a single-vertex graph (G, ρ) has an orientable embedding on a non-orientable surface with $k \geq 2$ crosscaps, we can add a one-sided loop into its embedding scheme, without changing the surface. If the orientable embedding of (G, ρ) contains at least one non-separating loop, we can assume that the one-sided loop we add is not orienting.*

Proof: Let $k' \leq k$ be minimal so that (G, ρ) has an orientable embedding on the surface with k' crosscaps; let λ be the signature of the embedding, so $\lambda(e) = 1$ for all loops e . If $k' = 0$, then we can add two crosscaps and a loop close to the vertex so that the loop passes through exactly one of the crosscaps; since $k \geq 2$ this is sufficient. Otherwise, by Lemma 6 we can assume that k' is odd and at least 3. To G add a loop with its ends consecutive in the rotation and next to a non-separating loop. Now push this loop once over each crosscap. Since all other loops are two-sided, every pair of edges crosses evenly, so by Theorem 3 the graph embeds in the surface with an equivalent embedding scheme. The loop we added is one-sided. (Lemma 7 shows that the loop is orienting, but we do not use that here.) If the loop is not orienting, we are done. Otherwise, the loop is orienting. We inserted the loop next to a non-separating loop f . Then f must pass through at least one crosscap an odd number of times by Lemma 3. We can now move the end of the one-sided loop along f and reattach it to the vertex. The resulting loop is still one-sided, and no longer orienting, since it passes an even number of times through the crosscap through which f passed an odd number of times (using Lemma 4). \square

Adding a one-sided loop (with consecutive ends) to an orientable embedding forces the loop to be orienting (as long as the embedding is minimal).

Lemma 7 *Let (G, ρ, λ) be a single-vertex graph embedded in a non-orientable surface with $\text{eg}(G, \rho, \lambda)$ crosscaps. If G contains exactly one one-sided loop and the ends of that loop are consecutive in ρ , then that loop is orienting.*

This seems intuitive: removing the one-sided loop leaves us with an orientable embedding scheme, but that does not automatically make the loop itself orienting; for that, we need to verify that the remaining surface is orientable.

Proof: Let e be the single one-sided loop in G , and let (G, ρ, λ) be embedded on a surface Σ with $k = \text{eg}(G, \rho, \lambda)$ crosscaps. By Theorem 2, k is minimal.

If we remove e from G we get a graph $G' = G - e$ with an induced embedding scheme (ρ', λ') on the same surface Σ . Since e is one-sided, and its ends are consecutive, removing e does not change the number of faces in the embedding, so $\text{eg}(G', \rho', \lambda') = k - 1$. Let k' be the smallest number of crosscaps so that (G', ρ', λ') can be embedded on a surface Σ' with k' crosscaps. By Lemma 6, k' is 0, or an odd number at least 3. The same argument as in the proof of Corollary 4 shows that (G, ρ, λ) can be embedded in the same surface Σ' , so $k \leq k'$, by minimality of k . In particular, $k' \geq 3$, since $k' = 0$ implies $k = 0$, which is not possible, since e is one-sided.

Assume that e is not orienting. Then, by Lemma 1, there is an embedding of (G, ρ, λ) in Σ in which e passes through a single cross-cap c , and does so once. Every other edge starts and ends on the same side of e , so all other edges pass through c an even number of times. We push all edges passing through c other than e off of c ; this gives us a drawing of (G', ρ', λ') in a surface with $k - 1$ crosscaps (we no longer need c). Since every pair of edges crosses an even number of times, Theorem 3 yields an embedding of (G', ρ', λ') in a surface with $k - 1$ crosscaps. This implies $k' \leq k - 1$ which contradicts $k \leq k'$. Hence e is orienting. \square

3 Removing Self-Crossings

Theorem 5 $\text{dcr}(H) \leq 3 \text{gcr}(H)$. *The bound remains true if a rotation system is fixed: $\text{dcr}(H, \rho) \leq 3 \text{gcr}(H, \rho)$; with a fixed embedding scheme $\text{dcr}(H, \rho, \lambda) \leq 6 \text{gcr}(H, \rho, \lambda)$.*

In other words, a degenerate drawing with self-crossings can be cleaned of self-crossings at the expense of increasing the number of degenerate crossings by a factor of three. We make use of the following lemma.

Lemma 8 *Let G be a single-vertex graph with rotation ρ and signature λ , then $\text{dcr}(G, \rho) \leq |E(G)|$, and $\text{dcr}(G, \rho, \lambda) \leq 2|E(G)|$.*

Proof: We use induction on $|E(G)|$. If $|E(G)| = 0$, there is nothing to show, so G has at least one loop. Pick a loop e whose ends at v are *closest* in the sense, that no other edge begins and ends in the wedge formed by the two ends of e . If we can, we pick e one-sided. Suppose e is one-sided. Let (G', ρ', λ') be obtained from (G, ρ, λ) by reversing the wedge formed by e , and removing e . This changes the signature of every edge enclosed by the wedge, since every edge has at most one end inside the wedge. By induction $\text{dcr}(G', \rho') \leq |E(G')|$ and $\text{dcr}(G', \rho', \lambda') \leq 2|E(G')|$. We can now add a crosscap close to v and pass all edges in the former wedge through that crosscap, reattaching them to v in their original order. This also reestablishes the original signatures of edges in G . Finally, we add back e in its proper place in the rotation, passing it through the crosscap once. By construction, $\text{dcr}(G, \rho) \leq 1 + \text{dcr}(G', \rho') \leq |E(G)|$ and $\text{dcr}(G, \rho, \lambda) \leq 1 + \text{dcr}(G', \rho', \lambda') \leq 2|E(G)|$.

If there is no closest, one-sided loop, e must be two-sided. If we only need to maintain the rotation system, we redefine $\lambda(e) = -1$ and proceed as above, proving that $\text{dcr}(G, \rho) \leq |E(G)|$. If we need to maintain the signature, define $\tilde{\lambda}$ to be λ , except $\tilde{\lambda}(e) = -1$. Arguing as in the previous case, we get $\text{dcr}(G, \rho, \tilde{\lambda}) \leq 1 + \text{dcr}(G', \rho', \tilde{\lambda}')$. Now add one additional crosscap passing only edge e through it, making it two-sided again. This shows that $\text{dcr}(G, \rho, \lambda) \leq 1 + \text{dcr}(G, \rho, \tilde{\lambda}) \leq 2 + \text{dcr}(G', \rho', \tilde{\lambda}') \leq 2|E(G)|$. \square

Proof of Theorem 5: Let H be a graph with $\text{gcr}(H) = k$. Fix an embedding of H on a surface S with k crosscaps. By Lemma 5, there is a graph G on a single vertex v with an embedding

scheme (ρ, λ) so that $\text{gcr}(H) = \text{gcr}(G, \rho, \lambda)$ and $\text{dcr}(H) \leq \text{dcr}(G, \rho, \lambda)$ and an embedding of H can be obtained from an embedding of G by uncontracting (and possible deleting) edges. We show the result by induction on $|E(G)| \leq |E(H)|$.

If $|E(G)| \leq 3k$, then the result follows from Lemma 8 (even if rotation or embedding scheme of H are fixed). So we can assume that $|E(G)| > 3k = 3\text{eg}(S)$. Lemma 2 implies that in this case there are two loops e and f so that $e \cup f$ bounds a disk (e and f are homotopic). Remove the disk (with any loops it may contain) from the surface, and identify e and f . Since this removes at least one edge from G we can apply induction to the resulting graph G' . From G' we can reconstruct an embedding of G by splitting e and f into two loops and reinserting the disk. Any loops in the disk which are not homotopic to e and f can be drawn close to v (so they do not use any crosscaps that e and f may be using). Any loops parallel to e and f use the same crosscaps as e and f , so in the resulting drawing no edge uses any crosscap more than once (note that any such loops have the same signature as e and f , since e and f bound a disk). \square

Since the proof works with single-vertex graphs with embedding schemes, the separation of gcr and dcr for those types of graphs (Theorem 7) implies that the proof approach in Theorem 5 will not yield $\text{gcr} = \text{dcr}$, but we can prove equality for small values.

Theorem 6 *If $\text{dcr}(G) \leq 3$, then $\text{gcr}(G) = \text{dcr}(G)$.*

For graphs with embedding scheme, this result is sharp, as Theorem 7 shows.

Proof: Since $\text{gcr}(G) \leq \text{dcr}(G)$ it is sufficient to show that if $\text{gcr}(G) \leq 2$, then $\text{dcr}(G) \leq \text{gcr}(G)$. By Lemma 5 it is sufficient to prove the result for single-vertex graphs with embedding scheme: for G there is a single-vertex graph G' and an embedding scheme (ρ', λ') so that $\text{dcr}(G) \leq \text{dcr}(G', \rho', \lambda')$ and $\text{gcr}(G', \rho', \lambda') = \text{gcr}(G)$, so establishing $\text{dcr}(G', \rho', \lambda') \leq \text{gcr}(G', \rho', \lambda')$ will prove the result.

If $\text{gcr}(G', \rho', \lambda') = 0$, there is nothing to prove. If $\text{gcr}(G', \rho', \lambda') = 1$ all loops are either two-sided and contractible, or one-sided. Pick a closest loop e (in the sense defined in Lemma 8: every edge has at most one end in the wedge formed by e). If e is one-sided, we can proceed as in Lemma 8, cutting along e , reversing the wedge enclosed by e and changing the signature of all edges in the wedge. The resulting graph is embedded in a plane (e is orienting, since it is one-sided and there is only one crosscap), and we can add back e so that it, and the edges it encloses cross through the crosscap exactly once. If e is two-sided, the ends of e must be consecutive. We can then remove e from the drawing, inductively draw the remaining graph, and add e back locally without using any crosscaps. If $\text{gcr}(G', \rho', \lambda') = 2$, there may be two-sided loops which are not contractible. However, if there is a closest one-sided loop, or a closest two-sided loop which is contractible, we can proceed as in the case of a single crosscap. Hence, all closest loops are two-sided, and either separating, or orienting. Suppose there is a one-sided loop f . Then the wedge enclosed by f must contain both ends of another loop e . Pick e so it is closest (within the wedge formed by f). Now e cannot be orienting, since the ends of an orienting loop alternate with the ends of a one-sided loop in the rotation. Hence e is separating. But then anything starting inside the wedge formed by e must end within the wedge as well, so since e was chosen to be closest, its ends have to be consecutive in the rotation. We can then remove e , inductively draw the remaining graph, and add e back into the rotation without using any additional crosscaps. We conclude that there is no one-sided loop f , so all loops are two-sided. By Lemma 6, the graph is planar in this case. \square

A closer look at the proof of Theorem 5 and Theorem 6 show that they are purely combinatorial, and the bounds can be implemented algorithmically.

4 Separating dcr and gcr with Embedding Schemes

Theorem 7 *There is a single-vertex graph G with embedding scheme (ρ, λ) for which $3 = \text{gcr}(G, \rho, \lambda) < \text{dcr}(G, \rho, \lambda) = 4$.*

Rephrased topologically: there is an embedding of a graph in a surface with three crosscaps, for which there is no PD1S system in which every curve intersects every edge of the graph at most once. The graph, as pictured in Figure 4, can be made into a pseudotriangulation (a vertex can appear more than once on the boundary of a pseudotriangle) by adding a single edge, turning it into a counterexample to the remark after Conjecture 3.4 in Mohar's paper [14].

Proof: See the graph pictured in Figure 4(a). The single vertex is drawn as the outer cycle, to make the picture easier to read. So there are 5 loop edges e_1, \dots, e_5 in this graph, the rotation at v is $e_1, e_2, e_3, e_4, e_5, e_3, e_2, e_1, e_4, e_5$, and the signatures are as in the embedding: $\lambda(e_1) = \lambda(e_3) = \lambda(e_4) = \lambda(e_5) = 1$ and $\lambda(e_2) = -1$. The drawing of G in Figure 4(a) shows that $\text{gcr}(G, \rho, \lambda) \leq 3$. If $\text{gcr}(G, \rho, \lambda) \leq 2$ were true, then e_2 would have to pass through exactly one of the two crosscaps oddly, say \otimes_1 . Since the ends of e_4 and e_5 alternate with the ends of e_2 , both e_4 and e_5 must also pass through \otimes_1 oddly. Since e_4 and e_5 are two-sided, they must then also pass through \otimes_2 oddly. But then e_4 and e_5 would be parallel (in the sense that their ends do not alternate), contradicting the fact that their ends alternate in the rotation. Hence, $\text{gcr}(G, \rho, \lambda) = 3$. The embedding in Figure 4(b) shows that $\text{dcr}(G, \rho, \lambda) \leq 4$, so we are left with the proof that $\text{dcr}(G, \rho, \lambda) \geq 4$. Suppose, for a contradiction, that G can be realized on a surface with three crosscaps so that every edge passes through each crosscap at most once, and the embedding scheme is (ρ, λ) , as specified in Figure 4(a). Then each edge in $\{e_1, e_3, e_4, e_5\}$ passes through an even number of crosscaps (we use that an edge can pass through a crosscap at most once). None of these edges can be separating (since they would all separate ends of other edges in the rotation), so they each pass through two crosscaps. Edge e_2 has signature -1 so it passes through an odd number of crosscaps. It cannot pass through all three crosscaps, since then all other edges would be parallel to it (as each would share two crosscaps with e_2), but the ends of e_2 alternate with the ends of e_4 and e_5 . Hence, e_2 passes through exactly one crosscap, say \otimes_1 . Since e_3 is parallel to e_2 , it must then pass through \otimes_2 and \otimes_3 . Now e_4 and e_5 alternate ends with both e_2 and e_3 , so one of them, say e_4 , by symmetry, passes through \otimes_1 and \otimes_2 and e_5 passes through \otimes_1 and \otimes_3 .

Edge	\otimes_1	\otimes_2	\otimes_3
e_2	1	0	0
e_3	0	1	1
e_4	1	1	0
e_5	1	0	1
e_1	0	1	1

Now e_1 is parallel to e_2 and e_3 and passes through two crosscaps, which must therefore be \otimes_2 and \otimes_3 . Now suppose there were such a drawing. Since edges pass through crosscaps at most once, we can think of crosscaps as vertices. But then, there is a path from an end of e_1 to an end of e_3 which passes through \otimes_2 and \otimes_3 but not through \otimes_1 . That path now separates the two ends of e_2 , since e_2 may only pass through \otimes_1 , so there is no way to connect the two ends of e_2 in the assumed drawing. \square

How important is the signature to Theorem 7; at a first glance, very important: if we do not fix the signature, then the graph (G, ρ) from the theorem can even be embedded in a surface with two

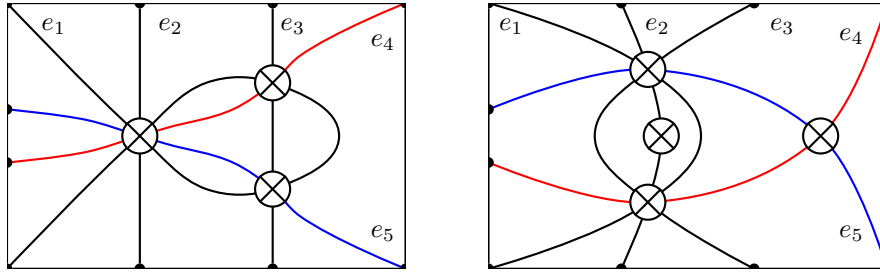


Figure 4: Graph G with rotation displayed as outer cycle. (a) G embedded in a surface with three crosscaps, requiring e_1 to pass through one crosscap twice. (b) G embedded in a surface with four crosscaps, each edge passing through each crosscap at most once.

crosscaps, and every edge passing through every crosscap at most once. However, the signature can be enforced using additional edges: create eight parallel copies of each edge in G , so that the resulting graph is still embeddable in a surface with three crosscaps. By Lemma 2 there must be at least two edges in each group that bound a disk, so the order of the ends of those two edges in the rotation determines the signature. This shows that the following result is true.

Corollary 8 *There is a single-vertex graph G with rotation system ρ for which $3 = \text{gcr}(G, \rho) < \text{dcr}(G, \rho) = 4$.*

Question 9 *Can the construction in Theorem 7 be used to construct for every n a single-vertex graph G with embedding scheme (ρ, λ) so that $n \leq \text{gcr}(G, \rho, \lambda) \leq (3/4) \text{dcr}(G, \rho, \lambda)$?*

5 Nice Embeddings of Higher Genus Graphs

In this section we consider relaxing the restriction on how often each edge may pass through each crosscap. It turns out that increasing the limit to two is sufficient.

Theorem 10 *If a graph is embeddable in a non-orientable surface, then it can be embedded—with an equivalent embedding scheme—in that surface so that every edge passes through each crosscap at most twice.*

This means, G always has a nearly degenerate drawing in the plane with at most $\text{gcr}(G)$ crossings, and in which each edge has at most $\text{gcr}(G)$ self-crossings. In topological language, following Mohar [14], the theorem states that there is a planarizing system of disjoint one-sided curves each of which intersects every edge of the graph at most twice.

By Theorem 7, the result is tight if the graph is given with an embedding scheme (which has to be maintained), even if the graph consists of a single vertex.

We will concentrate the proof in a more technical lemma, which may be of interest in its own right. The proof works with the Euler genus, $\text{eg}(G, \rho, \lambda)$ of an embedded single-vertex graph, which, by our earlier definition, is $1 + |E| - |F|$, where $|E|$ is the number of edges of G and $|F|$ the number of faces in the embedding scheme (ρ, λ) (a purely combinatorial property of the embedding scheme). It is tempting to assume that $\text{gcr}(G, \rho, \lambda) = \text{eg}(G, \rho, \lambda)$, but that is not actually true; take, for example, a single vertex with two two-sided edges alternating at the vertex. The Euler

genus of this graph is 2, while it requires 3 crosscaps to realize. The following lemma clarifies the relationship between gcr and eg .

Recall that we call an embedding scheme (ρ, λ) of a graph orientable, if $\lambda(e) = 1$ for all $e \in E(G)$.

Lemma 9 *If G is a single-vertex graph with embedding scheme (ρ, λ) , then it has an embedding in a surface with $\text{eg}(G, \rho, \lambda)$ crosscaps in which every edge uses every crosscap at most twice, unless (ρ, λ) is orientable, in which case such an embedding exists in a surface with $\text{eg}(G, \rho, \lambda) + 1$ crosscaps.*

The proof of this lemma can be viewed as a (more sophisticated) extension of the proof of Theorem 6. Since we allow edges to cross through a crosscap twice, the construction becomes simpler, in that we can process one-sided loops, even if they are not closest. The new ingredient needed is a technique for dealing with separating loops. Consider, for example, the embedding scheme described by $\rho = efefghgh$, and $\lambda(f) = -1$, and $\lambda(e) = \lambda(g) = \lambda(h) = 1$, as illustrated in Figure 5. The Euler genus of this graph is 3, and e is a separating loop, splitting the graph into two parts, one of Euler genus 1, and the other of Euler genus 2. The problem now is that the part of Euler genus 2 is orientable, and hence needs 3 crosscaps to realize by itself. Hence, some care is needed when merging drawings in this case.

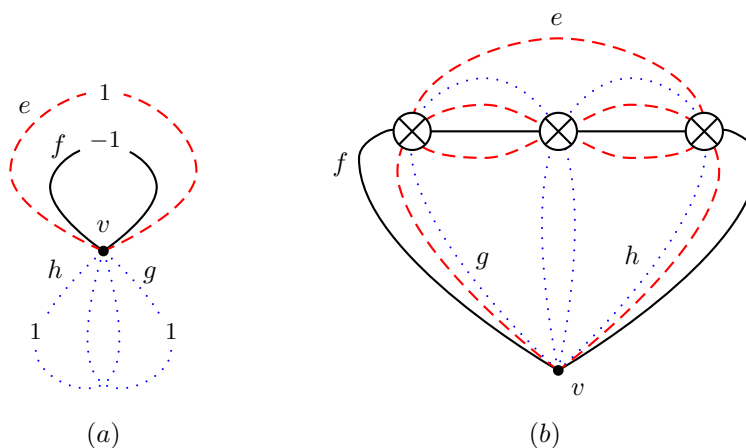


Figure 5: (a) Embedding scheme with Euler genus 3; edges are e (red/dashed), f (black), and g, h (blue/dotted). (b) Actual embedding of the same scheme on a surface with three crosscaps, in which every edge passes through every crosscap at most twice.

There is a point in the construction, where we need to ensure that a particular edge passes through a crosscap. The following lemma makes that possible.

Lemma 10 *If an embedding of a graph in a non-orientable surface contains edges which do not pass through a crosscap, we can modify the embedding (without changing the embedding scheme or the surface) so that each of those edges passes through a crosscap, and does so twice. The modification of the embedding does not increase the number of times any other edge passes through a crosscap.*

Proof: Suppose e is an edge that does not pass through a crosscap. Let ζ be a curve connecting an arbitrary point on e to a crosscap and so that ζ crosses the fewest number of edges. Suppose ζ crosses some edge f . Then f cannot pass through a crosscap, since otherwise we would have routed ζ along f up to the crosscap, reducing the number of crossings. We now push e and all edges crossed by ζ along ζ and through the crosscap. This ensures that e passes through a crosscap, and does so twice. We repeat this, until all edges pass through a crosscap. Edges which already passed through a crosscap are not affected by the redrawing. \square

Proof of Lemma 9: We prove the result by induction on $|E(G)|$ and, the number of edges being equal, the number of two-sided edges.

Case 1: There is a Contractible Loop. If the embedding of G contains a contractible loop, pick a contractible loop e forming a minimal wedge (in the sense that the number of ends in the wedge is minimal). The ends of e must then be consecutive, and we can remove e , apply induction, and reintroduce e without using any crosscap, and without affecting eg of the embedding (we removed one edge and one face).

Case 2: There is a Separating Loop. Suppose next that the embedding of G contains a separating loop e . Since we already eliminated contractible loops, e is non-contractible. We cut the surface along e , splitting v into two vertices v_1 and v_2 (each with its induced embedding scheme). As a result, we obtain embedded graphs (G_i, ρ_i, λ_i) , $i = 1, 2$ where $|E(G_1)| + |E(G_2)| + 1 = |E(G)|$. Since e is not contractible, we must have $|E(G_i)| \geq 1$, and therefore $|E(G_i)| \leq |E(G)| - 2$, $i = 1, 2$.

If (G_i, ρ_i, λ_i) is not orientable, we can inductively embed it in a surface with $\text{eg}(G_i, \rho_i, \lambda_i)$ crosscaps so that every edge passes through each crosscap at most twice. We can also ensure that the position in the rotation at v_i where v_{3-i} was attached lies on the outer face (without passing through a crosscap): take a point close to v_i which can be connected to v_i by a curve (without passing through a crosscap) ending where v_{3-i} was attached, and project that point to infinity.

If (G_i, ρ_i, λ_i) is orientable, let us add to it an edge e' in the empty wedge where e used to be in the rotation, but with e' being a one-sided curve. Since we added one edge, and did not increase the number of faces (note that the ends of e' are consecutive in $(G'_i, \rho'_i, \lambda'_i)$), we have $\text{eg}(G'_i, \rho'_i, \lambda'_i) = \text{eg}(G_i, \rho_i, \lambda_i) + 1$. Now $|E(G'_i)| = |E(G_i)| + 1 < |E(G)|$, so we can apply induction. Since $(G'_i, \rho'_i, \lambda'_i)$ is not orientable, we obtain a drawing of $(G'_i, \rho'_i, \lambda'_i)$ in a surface with $\text{eg}(G'_i, \rho'_i, \lambda'_i) = \text{eg}(G_i, \rho_i, \lambda_i) + 1$ crosscaps in which every edge passes through each crosscap at most twice. Edge e' has consecutive ends and it is the only one-sided loop in $(G'_i, \rho'_i, \lambda'_i)$, so it is orienting by Lemma 7. By Lemma 4, e' passes through every crosscap an odd number of times. Since it passes through every crosscap at most twice, e' passes through every crosscap exactly once. Projecting an appropriate point to infinity, we can ensure that the empty wedge between the two ends of e' lies on the outer face (without passing through a crosscap). We claim that we can now add e back into the rotation in the drawing of $(G'_i, \rho'_i, \lambda'_i)$ so that e' is on one side of e and the edges of (G_i, ρ_i, λ_i) on the other side: starting at one end of e , follow e' once through each crosscap; then in the outer face double-back to the first crosscap e entered and traverse the whole sequence of crosscaps again, finally following e' back to v_i . Note that e uses every crosscap exactly twice (edge e in Figure 5(b) illustrates this weaving construction). Call the resulting graph $(G''_i, \rho''_i, \lambda''_i)$.

When combining the drawings, we distinguish three cases based on how many of the (G_i, ρ_i, λ_i) are orientable.

Case 2.1: Neither of the (G_i, ρ_i, λ_i) is orientable. In this case, we simply join the two drawings of the (G_i, ρ_i, λ_i) at the appropriate point in the rotation at v_1 and v_2 . This is possible, since we ensured that these points were on the outer face; since they remain on the outer face after the

joining, we can add back e in its original position (without using any crosscaps). In this case, we used $\sum_{i=1}^2 \text{eg}(G_i, \rho_i, \lambda_i) = \text{eg}(G, \rho, \lambda)$ crosscaps, which is what we had to show.

Case 2.2: Exactly one of the (G_i, ρ_i, λ_i) is orientable. Without loss of generality, let us assume that (G_1, ρ_1, λ_1) is orientable, while (G_2, ρ_2, λ_2) is not. We have an embedding of (G_2, ρ_2, λ_2) using $\text{eg}(G_2, \rho_2, \lambda_2)$ crosscaps, and an embedding of $(G_1'', \rho_1'', \lambda_1'')$ using $\text{eg}(G_1, \rho_1, \lambda_1) + 1$ crosscaps. When combining these embeddings, we need to save one crosscap. We proceed as follows: join the drawings of $(G_1'', \rho_1'', \lambda_1'')$ and (G_2, ρ_2, λ_2) so that v_1 and v_2 are combined at the point where they were split. Let f be a loop in G_2 whose end is now next to e' in G_1'' . By Lemma 10, we can assume that f uses a crosscap.

Starting at $v_1 (= v_2)$ follow f until it is about to enter a crosscap for the first time. Remove that crosscap, severing edges passing through it. Take half the ends, starting with the end next to f , and route them along f to v_1 and then along e' (which we remove at this point) through all the crosscaps used by the drawing of (G_1, ρ_1, λ_1) . After passing through the last crosscap, we can route the edges back to the removed crosscap (since we are in the outer face). Since (G_1, ρ_1, λ_1) was orientable, it uses an odd number of crosscaps, so the rerouted edges now have the right order to reconnect to the severed ends at the removed crosscap. We have obtained a drawing of the original (G, ρ, λ) using $\sum_{i=1}^2 \text{eg}(G_i, \rho_i, \lambda_i) = \text{eg}(G, \rho, \lambda)$ crosscaps so that every edge passes through each crosscap at most twice. This is the case illustrated in Figure 5.

Case 2.3: Both of the (G_i, ρ_i, λ_i) are orientable. In this case, we have embeddings of $(G_i'', \rho_i'', \lambda_i'')$, $i = 1, 2$, using at most $\text{eg}(G_i, \rho_i, \lambda_i) + 1$ crosscaps and so that every edge passes through each crosscap at most twice. We can merge the two embeddings using the same construction described in the previous case. Removing the extra e' -edge, and the duplicate e , we obtain an embedding of (G, ρ, λ) using $1 + \sum_{i=1}^2 \text{eg}(G_i, \rho_i, \lambda_i) = \text{eg}(G, \rho, \lambda) + 1$ crosscaps, which meets the required bound in this case. This case is illustrated in Figure 6.

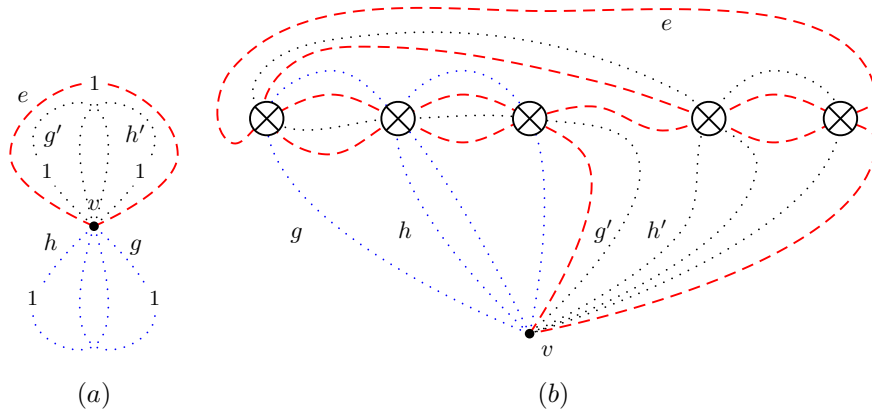


Figure 6: (a) Orientable embedding scheme with Euler genus 4; edges are e (red/dashed), g, h (blue/dotted), and g', h' (black/dotted). (b) Actual embedding of same scheme on surface with five crosscaps, in which every edge passes through every crosscap at most twice.

Case 3: No Separating Loops. We can therefore assume that (G, ρ, λ) contains no separating

loops (contractible or otherwise), so all two-sided loops are non-separating.

Case 3.1: There is a One-Sided, Non-Orienting Loop. If (G, ρ, λ) contains a one-sided loop e which is not orienting (not necessarily closest), we proceed similarly to how we did in Lemma 8, since in this case we did not increase the number of crosscaps: reverse the wedge of edges enclosed by e , change the signature of all edges occurring within the wedge exactly once, and remove e . The resulting embedding scheme (G', ρ', λ') has one edge less than G and the same number of faces (since the operation we performed corresponds to cutting the surface along e and removing the resulting hole), so $\text{eg}(G', \rho', \lambda') = \text{eg}(G, \rho, \lambda) - 1$. Since e is not orienting, (G', ρ', λ') is not orientable, and we can inductively find an embedding of (G', ρ', λ') in a surface with $\text{eg}(G', \rho', \lambda')$ crosscaps in which every edge passes through each crosscap at most twice.

We can then turn the embedding of G' into an embedding of G by reintroducing a crosscap close to the reversed wedge, and passing all edges of the wedge as well as e through it. This reestablishes the original embedding scheme. Moreover, each edge passes through the new crosscap at most twice (which happens if both its ends belong to the reversed wedge), and we used $\text{eg}(G', \rho', \lambda') + 1 = \text{eg}(G, \rho, \lambda)$ crosscaps.

Case 3.2: All One-Sided Loops are Orienting. If all loops are one-sided, and therefore (by assumption) orienting, their ends alternate pairwise in the rotation at the vertex. If we separate the rotation at the vertex into two halves, each half contains one end of every loop, in the same order (if the two ends of a loop form a wedge containing fewer than $|E(G)| - 1$ ends, one edge must be absent from that wedge entirely, making it parallel to the loop, which is not possible). The graph then trivially embeds in the projective plane, with each loop passing through the crosscap once. Hence, there is at least one two-sided loop.

Suppose there is both a one-sided (and necessarily orienting) and a two-sided loop. Then there must be a place in the rotation where an end of a one-sided (orienting) loop e is immediately followed by an end of a two-sided (necessarily not orienting) loop f . We create a new graph (G', ρ', λ') from (G, ρ, λ) by *sliding* e along f : that is, we detach the end of e next to f and move it along f until it reattaches to the rotation. Call that new edge e' . Note that $\text{eg}(G', \rho', \lambda') = \text{eg}(G, \rho, \lambda)$ and $|E(G)| = |E(G')|$ (since e turned into e'). The new edge e' is one-sided, and no longer orienting (we argue this as in Corollary 4, we need that the two-sided loop f is not separating). Hence we can again proceed as in the case where we have a one-sided loop which is not orienting. As we do this, orient e' so that the ends of f do *not* lie in the wedge enclosed by e' . We then obtain an embedding of (G', ρ', λ') on a surface with $\text{eg}(G', \rho', \lambda') = \text{eg}(G, \rho, \lambda)$ in which every edge passes through every crosscap at most twice. Moreover, since we worked with e' in the inductive step, we know that e' passes through a single crosscap, once, and that f does not pass through that crosscap. We can now slide e' back along f to turn the embedding into an embedding of (G, ρ, λ) . Since the crosscaps used by f are disjoint from the crosscap used by e' , e uses each crosscap at most twice.

Case 3.3: All Loops are Two-Sided. We are left with the case that all loops in (G, ρ, λ) are two-sided. We proceed similarly as we did in case 3.2, but the accounting is a bit different. Pick a closest loop e .

Consecutive with one of the ends of e , we add a new, one-sided loop f , giving us a graph $(\tilde{G}, \tilde{\rho}, \tilde{\lambda})$. This changes the number of edges, but not the number of faces, since the added loop is one-sided, and its ends are consecutive. Hence, $\text{eg}(\tilde{G}, \tilde{\rho}, \tilde{\lambda}) = \text{eg}(G, \rho, \lambda) + 1$. We now slide e along f to obtain (G', ρ', λ') with a new edge e' which is both one-sided and non-orienting (we argue this as we did earlier). One end of e' now lies inside the loop enclosed by f . Since the sliding move

does not change the number of faces or edges, we have $\text{eg}(G', \rho', \lambda') = \text{eg}(\tilde{G}, \tilde{\rho}, \tilde{\lambda})$.

We now remove e' , reverse the order of ends inside the wedge enclosed by e' and change the signature of all edges ending exactly once inside the wedge. Let (G'', ρ'', λ'') be the resulting graph. Then $|E(G'')| = |E(G)|$ and the number of two-sided edges has decreased (here we use that e' is not orienting, and (G, ρ, λ) is orientable). We can therefore inductively find an embedding of (G'', ρ'', λ'') on a surface with $\text{eg}(G'', \rho'', \lambda'')$ crosscaps so that every edge passes through every crosscap at most twice. Since $|E(G'')| = |E(G')| - 1$, and the number of faces did not change, we have $\text{eg}(G'', \rho'', \lambda'') \leq \text{eg}(G', \rho', \lambda') - 1 = \text{eg}(G, \rho, \lambda)$.

As we did before we now reintroduce e' and reestablish the embedding scheme (G', ρ', λ') by adding a single crosscap. Edges enclosed by the wedge formed by e' as well as e' itself will pass through the new crosscap exactly once (e was chosen to be closest; e' remained so, since the only additional end inside its wedge is one end of f).

This gives us an embedding of (G', ρ', λ') on a surface with $\text{eg}(G, \rho, \lambda) + 1$ crosscaps in which every edge passes through every crosscap at most twice. We slide the end of e' which neighbors f back along f , obtaining an embedding of $(\tilde{G}, \tilde{\rho}, \tilde{\lambda})$. In this embedding f is orienting (by Lemma 7), so it passes through each crosscap exactly once. This means that e passes through every crosscap once, except for the newly added one, through which it passes twice. Overall we have used $\text{eg}(G, \rho, \lambda) + 1$ crosscaps, which is what we had to establish in this case.

This completes the construction. \square

Proof of Theorem 10: Fix an embedding of a graph G on a surface with $k = \text{gcr}(G) = \tilde{\gamma}(G)$ crosscaps. By Lemma 5 there is a single-vertex graph G' with embedding scheme (ρ', λ') so that $\text{gcr}(G) = \text{gcr}(G', \rho', \lambda')$. It is sufficient to prove the result for G' , since an embedding of G' with embedding scheme (ρ', λ') can be turned back into an embedding of G by uncontracting and deleting edges (which may have been added to connect G). Since the uncontractions can be done close to the single vertex of G' , this does not affect how often edges pass through any crosscap. Hence, we can assume that G is given as a graph on a single vertex v with embedding scheme (ρ, λ) .

Since (G, ρ, λ) describes an embedding on the surface with k crosscaps, $\text{eg}(G, \rho, \lambda) \leq k$. If (G, ρ, λ) is not orientable, then the result follows immediately from Lemma 9. If (G, ρ, λ) is orientable, we apply Corollary 4 to extend (G, ρ, λ) to an embedding scheme (G', ρ', λ') which still embeds in the same surface, and is no longer orientable. Since $\text{eg}(G, \rho, \lambda) \leq \text{eg}(G', \rho', \lambda') \leq k$, and (G', ρ', λ') is not orientable, Lemma 9 gives us an embedding of (G', ρ', λ') , and thereby (G, ρ, λ) , in a surface with k crosscaps, in which every edge passes through each crosscap at most twice, completing the proof. \square

The proof of Theorem 10 is entirely combinatorial, so it can be made algorithmic.

The Complexity of the Degenerate Crossing Number

Corollary 11 *Determining the degenerate crossing number is NP-complete, even for cubic graphs.*

Proof: The problem lies in NP (since every edge passes through each crosscap at most once, we can guess the embedding). On the other hand, Thomassen [21, 15] showed that the non-orientable genus problem is NP-complete, even for cubic graphs. For a given cubic graph G , let G' be the result of replacing each edge of G with a path of length $2|E(G)|$, and attaching a (local, planar) gadget to each vertex of degree 2, to ensure that G' is cubic. If G has non-orientable genus at most k , then, by Theorem 10, there is an embedding in which every edge passes through each of

the crosscaps at most twice. Since we can assume that $k \leq |E(G)|$ (e.g. [16]), this implies that G' can be embedded so that every edge passes through each crosscap at most once. In other words, the degenerate crossing number of G' equals the non-orientable genus of G , showing that dcr is **NP**-complete. \square

6 Embeddings on Orientable Surfaces

What can we say about orientable surfaces? If we insist on every edge passing through each handle at most once, we can prove an upper bound as in Theorem 5. Just as in the non-orientable case, there is a crossing number view of this problem.

A *bundled crossing* in a drawing is a pseudodisk in which each of one bunch of parallel arcs crosses each of another bunch of parallel arcs. Given a drawing D of a graph, the *bundled crossing number*, $\text{bc}(D)$, is the smallest number of bundled crossings that are pairwise disjoint and cover all crossings of the graph. The *bundled crossing number*, $\text{bc}(G)$, of a graph G is the smallest $\text{bc}(D)$ for any good drawing D of G . If we allow drawings with self-crossings and multiple crossings, we get a variant denoted $\text{bc}'(G)$. In that variant, each bundled crossing can be viewed as a handle in an orientable surface, so $\text{bc}'(G) = \gamma(G)$ [2], where $\gamma(G)$ is the orientable genus of G .

The bundled crossing numbers $\text{bc}(G)$ and $\text{bc}'(G)$ are the orientable analogues of $\text{dcr}(G)$ and $\text{gcr}(G)$, and there has been some research on these, see [2, 10, 5, 3, 18]. For example, it is known that $\text{bc}(G)$ is **NP**-complete [10].

More strikingly, it is known that bc and bc' differ: $\text{bc}(K_6) > \text{bc}'(K_6) = 1$ as was shown by Alam, Fink, and Pupyrev [2].

By adapting and combining Lemma 5 and Theorem 5, we can show that the two bundled crossing numbers are within a small factor of each other.

Theorem 12 $\text{bc}(H) \leq 6 \text{bc}'(H)$.

In the proof, we think of handles as geometric objects, just like we did for crosscaps. We will use \boxplus to visualize such a handle.

Proof: We start with an embedding of H in an orientable surface of genus $\text{bc}'(H) = \gamma(H)$. As in Lemma 5, we can contract edges (possibly along edges we add to connect components) to obtain an embedding of a single vertex graph G in the same surface, so that an embedding of H can be recovered from an embedding of G by uncontracting and deleting edges.

We claim that G can be embedded on a surface with $|E(G)|$ handles (with its original embedding scheme) so that every edge passes through each handle at most once. Assuming the claim, we can deal with the case $|E(G)| \leq 6 \text{bc}'(G)$. In this case, $\text{bc}(H) \leq \text{bc}(G) \leq \text{bc}'(G) = \text{bc}'(H)$.

We prove the claim by induction on $|E(G)|$. There is nothing to prove for $|E(G)| = 0$, so we can assume that there is an edge. Pick e so that its ends are closest, that is, e alternates with every edge starting inside its wedge. We remove e from G and find the required embedding of $G - e$ inductively. We can now add back e into its proper place in the rotation. To avoid crossings, we use one handle that allows e to pass over all the edges inside its wedge; by choice of e , all those edges, as well as e itself, pass through the handle once. This proves the claim.

We are left with the case that $|E(G)| > 6 \text{bc}'(G)$. We show that in this case G' has an embedding in a surface with at most $6 \text{bc}'(G)$ handles (without changing the embedding scheme), and so that every edge passes through each handle at most once. Since $|E(G)| > 6 \text{bc}'(G) = 6 \gamma(S) = 3 \text{eg}(S)$, by Lemma 2 there must be two loops e and f that are homotopic; that is, they bound a disk. As in the proof of Theorem 5 we remove the disk and identify e and f . This reduces the number of

edges by at least one, so we can inductively obtain an embedding of the reduced graph into an orientable surface of genus $6 \text{bc}'(G)$ in which every edge uses every handle at most once. We can then undo the identification of e and f and embed the disk between them, without changing the surface. Loops lying inside the disk bounded by e and f must be contractible and can be embedded close to the vertex without using any handles, and e and f use the same handles, in particular, they use each handle at most once. \square

Since $\text{bc}'(H) = \gamma(H)$ we have $\text{bc}(H)/6 \leq \gamma(H) \leq \text{bc}(H)$, so $\text{bc}(H)$ and $\gamma(H)$ are within a constant factor of each other. This implies that approximation algorithms for the (Euler) genus, such as methods developed in [6, 12] are relevant for finding drawings with few bundled crossings. It also explains why finding good algorithms for the bundled crossing number is hard; any such algorithm also approximates the genus of the graph, which is known to be a difficult problem. The same conclusion applies to the degenerate crossing number (which has received less attention from an algorithmic point of view, than the bundled crossing number).

We next show that Theorem 10 can be used to establish a bound on how often edges have to pass through a handle in an orientable surface. This follows the philosophy suggested in [19] that looking at a problem on non-orientable surfaces can simplify a problem, *and* allow conclusions for orientable surfaces by restricting embeddings to orientable embeddings.

Theorem 13 *If a graph is embeddable in an orientable surface, then it can be embedded in that surface so that every edge passes through each handle at most four times.*

We start with a puzzle. Points p_1, \dots, p_n and q_n, \dots, q_1 lie on two horizontal lines, the p -points on the top, the q -points on the bottom line, in the given order (from left to right). We want to connect p_i to q_i , for all $1 \leq i \leq n$, without any crossings, including the horizontal lines. How many handles do we need? Say for $n = 5$?

Lemma 11 *The connection puzzle can be solved using at most $n/2$ handles and so that every curve passes through each handle at most twice.*

Proof: Figure 7 shows the proof idea: we lead q_1, \dots, q_n in parallel to the leftmost handle. We pass q_1 through the handle vertically to connect to p_1 . The remaining edges pass through the handle left to right. Immediately after that, $q_2 p_2$ doubles back to pass through the first handle vertically to connect to p_2 . We continue in this way: we pass $q_{2i-1} p_{2i-1}$ through the i -th handle vertically, connecting it to p_{2i-1} ; the remaining edges pass through the handle left to right; immediately after that $q_{2i} p_{2i}$ doubles back to pass through the i -th handle vertically to connect to p_{2i} . Each handle connects two more points. If n is odd, after the last handle, only $q_n p_n$ is left, which we can then connect to p_n directly. \square

Proof of Theorem 13: Suppose G is embedded in an orientable surface of (orientable) genus g . By Dyck's theorem, G has an orientable embedding on a non-orientable surface with $2g + 1$ crosscaps [15, Proposition 3.1.2]. We can add a one-sided loop e with consecutive ends to the rotation. By the analysis in Corollary 4, the resulting graph is still embeddable on the same surface, and e is orienting. Theorem 10 then gives us an embedding in which every loop passes through each crosscap at most twice. Since e is orienting, it passes through each crosscap exactly once. We now cut the surface along e . This gives us an orientable surface with a single hole, see Figure 8 (for $g = 1$). We group the edges which used to pass through each crosscap into $2g + 1$ bands. The solution to the connection puzzle, Lemma 11, shows how to embed the bands in a

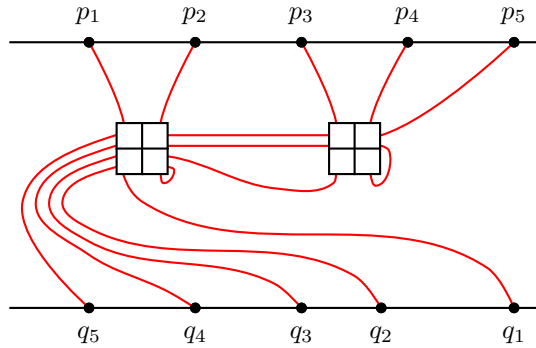


Figure 7: Connecting $p_i q_i$ without crossings and passing through each handle at most twice. Handles are represented geometrically as \boxplus .




Figure 8: An orienting loop passing through three crosscaps (*left*). After cutting the surface along the orienting loop with three twisted bands (*middle*). After untwisting the twisted bands (*right*).

surface with at most $(2g + 1)/2$ handles, which is at most g . Each band passes through each handle at most twice, and each band (corresponding to edges passing through a crosscap) contains each edge at most twice, so every edge passes through each handle at most four times in the resulting drawing. \square

7 Open Questions

The main open question which remains is whether $\text{dcr}(G) = \text{gcr}(G)$ for all graphs G ; one could weaken this question in various ways, and, for example ask whether $\text{dcr}(G) \leq \text{gcr}(G) + c$ for some constant c ? Another approach would be to ask whether $\text{dcr}(G) = \text{gcr}(G)$ if we allow a limited number of self-crossings along each edge. Theorem 10 implies that $\text{gcr}(G)$ self-crossings along each edge are sufficient, but can a constant bound be achieved?

For orientable embeddings, we showed that every edge needs to pass through every handle at most four times. A handle can be viewed as consisting of two parts: an overpass and an underpass (the left/right and top/bottom sides of \boxplus). In this view, Theorem 13 shows that every graph has a minimum genus embedding in which every edge uses every overpass and every underpass at

most twice. If $\text{dcr}(G) = \text{gcr}(G)$ for some graph G , then that bound improves: every edge uses every overpass and every underpass at most once. Since Alam, Fink, and Pupyrev [2] showed that $\text{bc}(K_6) > \text{bc}'(K_6)$, this would be the best we can expect. Conceivably the bound can be proved without establishing $\text{dcr}(G) = \text{gcr}(G)$.

We conclude with a computational question: how hard is it to find λ which minimizes $\text{eg}(G, \rho, \lambda)$ for a given single-vertex rotation system (G, ρ) ?

Acknowledgments

We would like to thank Bojan Mohar for suggesting the question, and giving us detailed feedback on earlier drafts of this paper. Thanks also go to Jan Kynčl for carefully working through more recent versions of the paper. And thanks to Niloufar Fuladi for spotting an error in Lemma 7. We are also grateful for many helpful comments by the anonymous reviewers.

References

- [1] E. Ackerman and R. Pinchasi. On the degenerate crossing number. *Discrete Comput. Geom.*, 49(3):695–702, 2013. doi:10.1007/s00454-013-9493-1.
- [2] M. J. Alam, M. Fink, and S. Pupyrev. The bundled crossing number. In *Graph drawing and network visualization*, volume 9801 of *Lecture Notes in Comput. Sci.*, pages 399–412. Springer, Cham, 2016. doi:10.1007/978-3-319-50106-2.
- [3] A. Arroyo and S. Felsner. Approximating the bundled crossing number. *ArXiv e-prints*, 2021. arXiv:2109.14892 (last accessed 10/4/2021). arXiv:2109.14892.
- [4] G. Cairns and Y. Nikolayevsky. Bounds for generalized thrackles. *Discrete Comput. Geom.*, 23(2):191–206, 2000. doi:10.1007/PL00009495.
- [5] S. Chaplick, T. C. van Dijk, M. Kryven, J.-w. Park, A. Ravsky, and A. Wolff. Bundled crossings revisited. *J. Graph Algorithms Appl.*, 24(4):621–655, 2020. doi:10.7155/jgaa.00535.
- [6] C. Chekuri and A. Sidiropoulos. Approximation algorithms for Euler genus and related problems. *SIAM J. Comput.*, 47(4):1610–1643, 2018. doi:10.1137/14099228X.
- [7] T. K. Dey and H. Schipper. A new technique to compute polygonal schema for 2-manifolds with application to null-homotopy detection. *Discrete Comput. Geom.*, 14(1):93–110, 1995. doi:10.1007/BF02570697.
- [8] C. A. Duncan, M. T. Goodrich, and S. G. Kobourov. Planar drawings of higher-genus graphs. *J. Graph Algorithms Appl.*, 15(1):7–32, 2011. doi:10.7155/jgaa.00215.
- [9] J. Erickson and S. Har-Peled. Optimally cutting a surface into a disk. *Discrete Comput. Geom.*, 31(1):37–59, 2004. ACM Symposium on Computational Geometry (Barcelona, 2002). doi:10.1007/s00454-003-2948-z.
- [10] M. Fink, J. Hershberger, S. Suri, and K. Verbeek. Bundled crossings in embedded graphs. In *LATIN 2016: theoretical informatics*, volume 9644 of *Lecture Notes in Comput. Sci.*, pages 454–468. Springer, Berlin, 2016. doi:10.1007/978-3-662-49529-2_34.

- [11] H. Harborth. Drawings of graphs and multiple crossings. In *Graph theory with applications to algorithms and computer science (Kalamazoo, Mich., 1984)*, Wiley-Intersci. Publ., pages 413–421. Wiley, New York, 1985.
- [12] K.-i. Kawarabayashi and A. Sidiropoulos. Polylogarithmic approximation for Euler genus on bounded degree graphs. In *STOC'19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 164–175. ACM, New York, 2019. doi:[10.1145/3313276.3316409](https://doi.org/10.1145/3313276.3316409).
- [13] F. Lazarus, M. Pocchiola, G. Vegter, and A. Verroust. Computing a canonical polygonal schema of an orientable triangulated surface. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry (SCG-01)*, pages 80–89, New York, June 3–5 2001. ACM Press. doi:[10.1145/378583.378630](https://doi.org/10.1145/378583.378630).
- [14] B. Mohar. The genus crossing number. *Ars Math. Contemp.*, 2(2):157–162, 2009. doi:[10.26493/1855-3974.21.157](https://doi.org/10.26493/1855-3974.21.157).
- [15] B. Mohar and C. Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.
- [16] J. Pach and G. Tóth. Degenerate crossing numbers. *Discrete Comput. Geom.*, 41(3):376–384, 2009. doi:[10.1007/s00454-009-9141-y](https://doi.org/10.1007/s00454-009-9141-y).
- [17] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings on surfaces. *European J. Combin.*, 30(7):1704–1717, 2009. doi:[10.1016/j.ejc.2009.03.002](https://doi.org/10.1016/j.ejc.2009.03.002).
- [18] M. Schaefer. The graph crossing number and its variants: A survey. *The Electronic Journal of Combinatorics*, 20:1–113, 2017. Dynamic Survey, #DS21. doi:[10.37236/2713](https://doi.org/10.37236/2713).
- [19] M. Schaefer and D. Štefankovič. Block additivity of \mathbb{Z}_2 -embeddings. In *Graph drawing*, volume 8242 of *Lecture Notes in Comput. Sci.*, pages 185–195. Springer, Cham, 2013. doi:[10.1007/978-3-319-03841-4_17](https://doi.org/10.1007/978-3-319-03841-4_17).
- [20] J. Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993. doi:[10.1007/978-1-4612-4372-4](https://doi.org/10.1007/978-1-4612-4372-4).
- [21] C. Thomassen. The genus problem for cubic graphs. *J. Combin. Theory Ser. B*, 69(1):52–58, 1997. doi:[10.1006/jctb.1996.1721](https://doi.org/10.1006/jctb.1996.1721).