Block Additivity of $\mathbb{Z}_2$-Embeddings

Marcus Schaefer[1] and Daniel Štefankovič[2]

[1] School of Computing, DePaul University, Chicago, IL 60604
mschaefer@cs.depaul.edu
[2] Computer Science Department, University of Rochester, Rochester, NY 14627-0226
stefanko@cs.rochester.edu

Abstract. We study embeddings of graphs in surfaces up to $\mathbb{Z}_2$-homology. We introduce a notion of genus mod 2 and show that some basic results, most noteworthy block additivity, hold for $\mathbb{Z}_2$-genus. This has consequences for (potential) Hanani-Tutte theorems on arbitrary surfaces.

1 Introduction

A graph G *embeds* in a surface $\mathcal{S}$ if it can be drawn in $\mathcal{S}$ so that no pair of edges cross. In this paper we want to relax the embedding condition using $\mathbb{Z}_2$-homology, that is, we are only interested in the parity of the number of crossings between independent edges; in terms of algebraic topology we are studying the "mod 2 homology of the deleted product of the graph" [1]. We say a graph $\mathbb{Z}_2$-*embeds* in $\mathcal{S}$ if it can be drawn in $\mathcal{S}$ so that every pair of independent edges crosses evenly. This approach is inspired by two Hanani-Tutte theorems which, for the plane [2] and the projective plane [3], show that embeddability is equivalent to $\mathbb{Z}_2$-embeddability. For other surfaces, it is only known that $\mathbb{Z}_2$-embeddability is a necessary condition for embeddability. In this paper we want to lay the foundations for a study of $\mathbb{Z}_2$-embeddings of graphs in surfaces which may, at some point, lead to a proof of the Hanani-Tutte theorem for arbitrary surfaces. Our main result is that if we define the notion of $\mathbb{Z}_2$-genus as a homological invariant of $\mathbb{Z}_2$-embeddings, then block additivity holds for $\mathbb{Z}_2$-genus just as it does for the standard notion of genus (as proved by Battle, Harary, Kodama and Youngs for the orientable case, and by Stahl and Beineke in the non-orientable case, see [4, Section 4.4]). This implies that a counterexample to the Hanani-Tutte theorem on an arbitrary surface can be assumed to be 2-connected (Corollary 1).

2 $\mathbb{Z}_2$-Embeddings

2.1 Definition

In the introduction we defined a $\mathbb{Z}_2$-embedding in a surface $\mathcal{S}$, as a drawing of a graph G in which every pair of independent edges crosses evenly. In this section, we want to develop a more algebraic version of this definition, which separates the topology of the surface from the drawing. We start with the plane, and then add crosscaps and handles.

Pick an initial drawing D of $G = (V, E)$ in the plane. For edges $e, f \in E$ let $i_D(e, f)$ be the number of crossings of e and f in the drawing D. We want to extend the drawing to a surface $\mathcal{S}$ with c crosscaps. Since we only plan to keep track of the parity of the number of crossings of independent edges and hence we use $\mathbb{Z}_2$-homology. For each edge we have a vector $y_e \in \mathbb{Z}_2^c$ where $(y_e)_i = 1$ if e is pulled through the i-th crosscap an odd number of times (in a drawing, we can deform e so it passes through the i-th crosscap; this changes the crossing parity of e with any edge that passes through the i-th crosscap an odd number of times). We also allow changing the planar part of the drawing—for each edge we have a vector $x_e \in \mathbb{Z}_2^V$ where $(x_e)_v = 1$ indicates that we made an (e, v)-*move*, that is, we pull the edge e over v (this changes the crossing parity between e and any edge incident to v). We say that the initial drawing together with $\{x_e\}_{e\in E}$ and $\{y_e\}_{e\in E}$ is a $\mathbb{Z}_2$-*embedding* of G in $\mathcal{S}$ if for each pair of independent edges $e = \{u, v\}, f = \{s, t\}$ we have

$$i_D(e, f) + (x_e)_s + (x_e)_t + (x_f)_u + (x_f)_v + y_e^T y_f \equiv 0 \pmod 2. \tag{1}$$

All congruences in this paper are modulo 2, so to simplify notation, we drop (mod 2) from now on. See Figure 1 for a $\mathbb{Z}_2$-embedding of K_5 in the projective plane, illustrating the effect of crosscap- and (e, v)-moves. This definition is equivalent to the more intuitive definition given in the introduction (see, for example, Levow [5, Theorem 3]). We say that the drawing is *orientable* if $y_e^T y_e \equiv 0$ for every $e \in E$ (that is, every y_e has an even number of ones).

Handles can be dealt with in the following way: For each handle we have three coordinates in y_e and the possible settings for these coordinates are 000, 110, 101, and 011. We extend the definitions of $\mathbb{Z}_2$-embeddings and orientability given earlier to y_e containing handles. Note that each of the four vectors modeling an edge passing through handle has an even number of ones, so if the surface contains only handles, then any drawing on it is orientable by the earlier definition.

A handle and a crosscap are equivalent to three crosscaps (Dyck, see [6, Section 1.2.4]): in the $\mathbb{Z}_2$-homology this corresponds to the following bijection (we replace the $3 + 1$ coordinates in y_e by 3 coordinates in y_e):

$$\begin{aligned} &000\,0 \leftrightarrow 000, 000\,1 \leftrightarrow 111, 011\,0 \leftrightarrow 011, 011\,1 \leftrightarrow 100,\\ &110\,0 \leftrightarrow 110, 110\,1 \leftrightarrow 001, 101\,0 \leftrightarrow 101, 101\,1 \leftrightarrow 010, \end{aligned} \tag{2}$$

where the first three coordinates on the left hand sides correspond to the handle. Note that (2) preserves the parity of the number of ones (and thus the orientability of the drawing). Also note that (2) is linear (add vector 111 times the last coordinate to the vector of the first three coordinates) and hence preserves the dimension of the space generated by the $\{y_e\}_{e\in E}$.

Remark 1 ($\mathbb{Z}_2$-drawings). Call a drawing D of a graph G in the plane together with $\{x_e\}_{e\in E}$ and $\{y_e\}_{e\in E}$ a $\mathbb{Z}_2$-*drawing*, and define $i_{D,x,y}(e, f) := i_D(e, f) + (x_e)_s + (x_e)_t + (x_f)_u + (x_f)_v + y_e^T y_f$. With this notion of $\mathbb{Z}_2$-drawing, we can model drawings of graphs in a surface up to the $\mathbb{Z}_2$-homology we are interested in: If D is

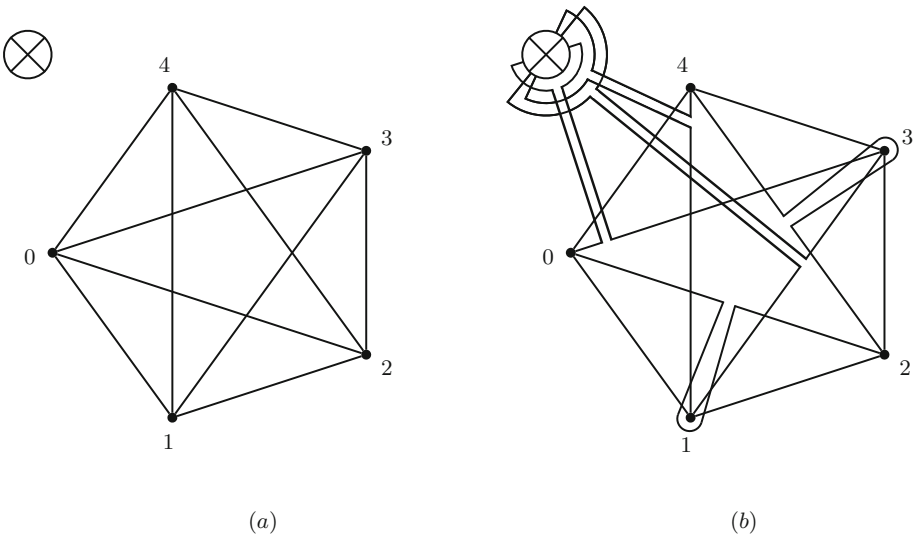

Fig. 1. (a) shows the initial drawing of $G = K_5$ in the projective plane. (b) shows a $\mathbb{Z}_2$-embedding of G in the projective plane with $(x_{02})_1 = (x_{24})_3 = 1$, $y_{03} = y_{14} = y_{24} = 1$ (dropping the subscript for the single crosscap) and all other values being zero.

a drawing of a graph G in some surface $\mathcal{S}$, then there is a $\mathbb{Z}_2$-drawing (D', x, y) of G in $\mathcal{S}$ so that $i_D(e, f) \equiv i_{D',x,y}(e, f)$ for every pair (e, f) of independent edges. As mentioned earlier, a result like this (with a slightly different model) was stated by Levow [5]. In the plane, algebraic topologists would phrase this as saying that any two drawings differ by a coboundary, or that they define the same cohomology class in the second symmetric cohomology, see, for example, [7, Section 4.6].

By the observations in Remark 1, any embedding in a surface $\mathcal{S}$ can be considered a $\mathbb{Z}_2$-embedding, so that having a $\mathbb{Z}_2$-embedding is a necessary condition for embeddability in a surface. Hanani-Tutte theorems state that this condition is also sufficient. As we mentioned earlier, this is only known for the plane [2] and the projective plane [3].

Remark 2 (Crosscaps versus Handles in $\mathbb{Z}_2$-Embeddings). Suppose that a surface $\mathcal{S}$ contains c crosscaps and h handles. By the classification theorem for surfaces, each handle is equivalent to two crosscaps, as long as $c > 0$. The same is true for $\mathbb{Z}_2$-embeddability: If $c > 0$, then $\mathbb{Z}_2$-embeddability in $\mathcal{S}$ is equivalent to $\mathbb{Z}_2$-embeddability in a surface with $c + 2h$ crosscaps—we apply (2) and convert each handle into 2 crosscaps; the transformation is possible because $c > 0$. If $c = 0$, then $\mathbb{Z}_2$-embeddability in $\mathcal{S}$ is equivalent—again using (2)—to $\mathbb{Z}_2$-embeddability in a surface with $2h + 1$ crosscaps where drawings are restricted to be orientable. The orientability ensures that when applying (2) to convert

crosscaps into handles, the single crosscap left at the end is not used by any edge and hence can be discarded.

With this terminology, we can now define $\mathbb{Z}_2$-homological variants of the genus and the Euler genus of a graph. We write $\mathbf{g}(G)$ and $\mathbf{eg}(G)$ for the traditional genus and Euler genus of G (following [4]).

Definition 1 ($\mathbb{Z}_2$-genus and $\mathbb{Z}_2$-Euler genus). *If a graph G has a $\mathbb{Z}_2$-embedding in an orientable surface with h handles, but not in any surface with fewer handles, we write $\mathbf{g}_0(G) = h$ and call h the* $\mathbb{Z}_2$-genus *of G. If G has a $\mathbb{Z}_2$-embedding in a surface $\mathcal{S}$ with c crosscaps and h handles, but not in any surface with a smaller value of $2h + c$, we write $\mathbf{eg}_0(G) = 2h + c$ and call $2h + c$ the* $\mathbb{Z}_2$ Euler genus *of G.*

By definition, we have $\mathbf{g}_0(G) \leq \mathbf{g}(G)$ and $\mathbf{eg}_0(G) \leq \mathbf{eg}(G)$, where $\mathbf{g}(G)$ is the genus of G and $\mathbf{eg}(G)$ is the Euler genus of G.

2.2 Basic Properties

We derive some basic properties of $\mathbb{Z}_2$-embeddings. We call two graphs G and H *disjoint* if $V(G) \cap V(H) = \emptyset$.

Lemma 1. *Let G be a graph $\mathbb{Z}_2$-embedded in a (possibly non-orientable) surface $\mathcal{S}$. Let C_1 and C_2 be two disjoint cycles in G. Then*

$$\sum_{e \in C_1, f \in C_2} y_e^T y_f \equiv 0. \tag{3}$$

Let $e_1 \in C_1$ and $e_2 \in C_2$. Suppose that all edges e in $C_1 \setminus \{e_1\}$ have $y_e = 0$ and all edges $f \in C_2 \setminus \{e_2\}$ have $y_f = 0$. Then

$$y_{e_1}^T y_{e_2} \equiv 0. \tag{4}$$

Proof. We have

$$\sum_{e \in C_1, f \in C_2} \left(i_D(e, f) + (x_e)_s + (x_e)_t + (x_f)_u + (x_f)_v + y_e^T y_f \right) \equiv 0, \tag{5}$$

where s, t are endpoints of f and u, v are endpoints of e (note that s,t,u and v vary over the terms in the sum). The equality in (5) follows from the disjointness of C_1 and C_2 (any $e \in C_1$ and $f \in C_2$ are independent and hence (1) has to be satisfied).

We have

$$\sum_{e \in C_1, f \in C_2} (i_D(e, f) + (x_e)_s + (x_e)_t + (x_f)_u + (x_f)_v) \equiv 0, \tag{6}$$

since (6) corresponds to a drawing in the plane (and two transversally intersecting cycles in the plane intersect evenly; the cycles have to intersect transversally since they are disjoint).

Combining (5) with (6) we obtain (3). Equation (4) is an immediate corollary (since terms in (3) other than $y_{e_1}^T y_{e_2}$ are zero). □

Lemma 1 suggests that orthogonality plays a role in understanding $\mathbb{Z}_2$-embeddings. We will need the following fact about vector spaces over finite fields (see [8, Section 2.3]).

Lemma 2. *Let A be a subspace of $\mathbb{Z}_2^t$. Then for $A^\perp := \{x \in \mathbb{Z}_2^t \mid (\forall y \in A)\ x^T y \equiv 0\}$ we have*

$$\dim A + \dim A^\perp = t.$$

Let A, B be subspaces of $\mathbb{Z}_2^t$ with $A \subseteq B$. Then for $A^{\perp B} := \{x \in B \mid (\forall y \in A)\ x^T y \equiv 0\} = A^\perp \cap B$ we have

$$\dim A + \dim A^{\perp B} = \dim B + \dim \operatorname{rad} B,$$

where $\operatorname{rad} B := B^{\perp B}$ is the radical *of B.*

The dimension of the $\mathbb{Z}_2$-embedding (the dimension of the space spanned by $\{y_e\}_{e\in E}$) is closely related to its $\mathbb{Z}_2$-genus. In Lemma 3 we extend this result to $\mathbb{Z}_2$-Euler genus.

Lemma 3. *Let G be a graph $\mathbb{Z}_2$-embedded in a (possibly non-orientable) surface $\mathcal{S}$. Assume that the drawing is orientable. Let d be the dimension (over $\mathbb{Z}_2$) of the vector space generated by the edge vectors $\{y_e\}_{e\in E}$. Then G can be $\mathbb{Z}_2$-embedded in an orientable surface of genus $\lfloor d/2 \rfloor$.*

Proof. In the light of Remark 2 we can assume that $\mathcal{S}$ has t crosscaps (and no handles). We are going to remove the crosscaps from $\mathcal{S}$ one by one. Let S be the space generated by the edge vectors $\{y_e\}_{e\in E}$. Let $d = \dim S$. Let $T = S^\perp$ be the space of vectors that are perpendicular to S. Assume that T contains a vector z such that $z \not\equiv 0$ and $z^T z \equiv 0$ (the computations are in $\mathbb{Z}_2$). Rearranging coordinates, if necessary, we can assume that $z_1 = 1$. To each y_e for which $(y_e)_1 = 1$ we add z. This transformation has the following properties:

- it is linear ($y \mapsto y + y_1 z$) and hence the dimension of S cannot increase,
- it preserves orientability (since $(y+z)^T(y+z) \equiv y^T y$),
- it preserves the parity of the number of crossings for every independent pair of edges e, f (since $(y_e+z)^T(y_f+z) \equiv y_e^T y_f$ and $(y_e+z)^T y_f \equiv y_e^T y_f$; here we use the fact that $T = S^\perp$).

After the transformation, the first crosscap is not used by any edge and hence we can remove it thus decreasing t. We repeat the crosscap removal process as long as such a z exists. We distinguish three cases depending on whether the process stops with $d \le t-2$, $d = t-1$ or $d = t$.

If $d \le t-2$, then T always contains $z \neq 0$ with $z^T z \equiv 0$ (since by a dimension argument there are two distinct vectors z_1, z_2 in $T\setminus\{0\}$ and then one of z_1, z_2, z_1+z_2 satisfies $z^T z \equiv 0$) and hence we can always remove a crosscap in this case. Therefore, the process ends up with either $d = t$ or $d = t-1$, $T = \langle z \rangle$, and $z^T z \equiv 1$. If $d = t$, then we convert the crosscaps back into $\lfloor d/2 \rfloor$ handles (if t is odd we end up with $(t-1)/2$ handles, if t is even we add an extra crosscap

and end up with $t/2$ handles, in both cases use Remark 2 on being able to drop a crosscap in an orientable embedding).

The final case to handle is $d = t - 1$. Let k be the number of ones in z. W.l.o.g. the first k coordinates of z are 1 and the rest are 0. Note that k is odd and that $z^T y_e \equiv 0$ for each e (that is, if we restrict our attention to the first k crosscaps, the drawing is orientable). Hence we can convert the first k crosscaps into $(k-1)/2$ handles. Then—as in the $d = t$ case—we convert the remaining $t - k$ crosscaps into $\lfloor (t-k)/2 \rfloor$ handles. In total, we have

$$(k-1)/2 + \lfloor (t-k)/2 \rfloor = \lfloor (t-1)/2 \rfloor = \lfloor d/2 \rfloor$$

handles. □

For non-orientable surfaces we have the following analogue of Lemma 3, replacing the notion of genus by Euler genus.

Lemma 4. *Let G be a graph $\mathbb{Z}_2$-embedded in a (possibly non-orientable) surface $\mathcal{S}$. Let d be the dimension (over $\mathbb{Z}_2$) of the vector space generated by the edge vectors $\{y_e\}_{e \in E}$. Then G can be $\mathbb{Z}_2$-embedded in a (possibly non-orientable) surface of Euler genus d.*

Proof. The proof is almost the same as the proof of Lemma 3. We first convert handles to crosscaps and work on a surface with crosscaps only. We again remove crosscaps one by one until we end up with $d = t$ or with $d = t - 1$, $T = \langle z \rangle$, and $z^T z \equiv 1$. In the case that $d = t$ we are done.

In the case $d = t - 1$ we assume, as in the proof of Lemma 3, that the first k coordinates of z are 1 and the rest are 0 and convert the first k crosscaps into $(k-1)/2$ handles. We leave the remaining crosscaps as they are. The Euler genus of the resulting surface is $2((k-1)/2) + t - k = t - 1 = d$. □

We end this section with a more complex move that allows us to zero out the labels of all edges in a spanning forest.

Lemma 5. *Suppose G is $\mathbb{Z}_2$-embedded on a surface $\mathcal{S}$, and F is a spanning forest of G. Then there is a $\mathbb{Z}_2$-embedding of G on $\mathcal{S}$ in which all edges of F are labeled with zero vectors.*

Proof. By Remark 2 we can assume that $\mathcal{S}$ is a surface with $c > 0$ crosscaps. Choose $z \in \mathbb{Z}_2^c$ and $v \in V$. Consider the following collection of moves: 1) add z to y_e for all e that are adjacent to v, and 2) for every f not adjacent to v and so that $y_f^T z \equiv 1$ perform an (f, v)-move. This collection of moves preserves the parity of the number of crossings between any pair of independent edges. Moreover, if z contains an even number of ones, the parity of the number of ones in no y-label is changed. Pick a root for each component of F, orient the edges of F away from the root, and process the edges in each component in a breadth-first traversal; for each edge e in this traversal, we turn its label into the zero vector, by performing the collection of moves above with $z = y_e$ and v the head of e. Note that this changes the label of e into the zero vector, without

affecting the labels of any edges that have already been processed (since F is a forest, and $y_f^T z \equiv 0$ for edges f already processed, because $y_f = 0$ for those edges). If the $\mathbb{Z}_2$-embedding was orientable to begin with, it remains so, since z is chosen from the set of existing labels, all of which contain an even number of ones originally and throughout the relabeling. □

3 Block Additivity mod $\mathbb{Z}_2$

As a warm-up we show the additivity of genus over connected components (a result that is nearly obvious for embeddings).

Lemma 6. *The $\mathbb{Z}_2$-genus of a graph is the sum of the $\mathbb{Z}_2$-genera of its connected components.*

Proof. Let G be a graph. Let $g := \mathbf{g}_0(G)$ be the $\mathbb{Z}_2$-genus of G. By Remark 2 we have an orientable drawing of G on the surface with $t := 2g + 1$ crosscaps. Assume that G is the disjoint union of G_1 and G_2. Let F_1 be a maximum spanning forest of G_1 and F_2 be a maximum spanning forest of G_2. We can assume (see Lemma 5) that the y_e-labels for edges e in F_1 and F_2 are zero.

Let e_1 be an edge in $G_1 - F_1$ and let e_2 be an edge in $G_2 - F_2$. Let C_1 be the unique cycle in $F_1 + e_1$ and let C_2 be the unique cycle in $F_2 + e_2$. Note that C_1 and C_2 are disjoint (since G_1 and G_2 are disjoint) and hence by Lemma 1 we have

$$y_{e_1}^T y_{e_2} \equiv 0; \tag{7}$$

that is, the vectors y_{e_1} and y_{e_2} are perpendicular. Let S_1 be the vector space generated by the y_e-labels on the edges in G_1 and let S_2 be the vector space generated by the y_e-labels on the edges in G_2. Then $S_1 \perp S_2$ and hence

$$\dim S_1 + \dim S_2 \leq t = 2g + 1. \tag{8}$$

By Lemma 3, we can $\mathbb{Z}_2$-embed G_i in an orientable surface with $\lfloor (\dim S_i)/2 \rfloor$ handles. Note

$$\lfloor (\dim S_1)/2 \rfloor + \lfloor (\dim S_2)/2 \rfloor \leq g$$

and hence $\mathbf{g}_0(G_1) + \mathbf{g}_0(G_2) \leq \mathbf{g}_0(G)$. □

Again, one also has the analogue of Lemma 6 for non-orientable surfaces.

Lemma 7. *The $\mathbb{Z}_2$-Euler genus of a graph is the sum of the $\mathbb{Z}_2$-Euler genera of its connected components.*

Proof. The proof is the same as the proof of Lemma 6 except the final part. We have

$$\dim S_1 + \dim S_2 \leq t, \tag{9}$$

where $t := \mathbf{eg}_0(G)$ is the $\mathbb{Z}_2$-Euler genus of G. By Lemma 4 we can draw G_i in a surface with $\dim S_i$ crosscaps. Hence $\mathbf{eg}_0(G_1) + \mathbf{eg}_0(G_2) \leq \mathbf{eg}_0(G)$. □

We are ready now to establish additivity of $\mathbb{Z}_2$-genus and $\mathbb{Z}_2$-Euler genus over 2-connected components (blocks).

Theorem 1. *The $\mathbb{Z}_2$-genus of a graph is the sum of the $\mathbb{Z}_2$-genera of its blocks. The $\mathbb{Z}_2$-Euler genus of a graph is the sum of the $\mathbb{Z}_2$-Euler genera of its blocks.*

Proof. There is a large shared part in the arguments for $\mathbb{Z}_2$-genus and $\mathbb{Z}_2$-Euler genus (only the initial setup and the final drawing step are different).

The initial setup for the $\mathbb{Z}_2$-genus case is the following. Let $G = (V, E)$ be a connected graph (we can assume this by Lemma 6) and let $g := \mathbf{g}_0(G)$ be the $\mathbb{Z}_2$-genus of G. Thus we have an orientable $\mathbb{Z}_2$-embedding of G on the surface with $t := 2g + 1$ crosscaps. Let B be the subspace of $\mathbb{Z}_2^t$ consisting of vectors with an even number of ones (we will keep our drawing orientable, that is, all the edge labels will be from B). Note that

$$\operatorname{rad} B = \{0\}, \tag{10}$$

since each vector in $B \setminus \{0\}$ has a zero and a one. Let $\hat{t} := \dim B = t - 1$.

The initial setup for the $\mathbb{Z}_2$-Euler genus case is the following. Let $G = (V, E)$ be a connected graph (Lemma 7), and let $g := \mathbf{eg}_0(G)$ be the $\mathbb{Z}_2$-Euler genus of G. Thus we have a $\mathbb{Z}_2$-embedding of G on the surface with t crosscaps. In this case we let $B := \mathbb{Z}_2^t$. (And we trivially have (10).) Let $\hat{t} := \dim B = t$.

Let v be a cut vertex of G. Let $G_1 = (V_1, E_1)$ be a block of G containing v and let $G_2 = (V_2, E_2)$ be the union of the remaining blocks (note that $V_1 \cap V_2 = \{v\}$). Let $T = (V, F)$ be a depth-first search (DFS) spanning tree of G with the exploration starting at v. We can assume $y_e = 0$ for the edges in F (see Lemma 5). The reason for taking a DFS spanning tree is that we will need the following property: if $e \in E \setminus F$ is not adjacent to v then the unique cycle in $F + e$ does not contain v.

Let S_i be the vector space generated by the y_e-labels of $e \in E_i$ that are not adjacent to v. Let Z_i be the vector space generated by the y_e-labels of $e \in E_i$ that are adjacent to v. See Figure 2. Our plan is to modify the y_e-labels of the edges adjacent to v (changing Z_1, Z_2) so that: 1) no new odd crossings between independent edges are introduced, and 2) after the modification $\dim(S_1 + Z_1) + \dim(S_2 + Z_2) \leq t$.

We can modify the y_e of an edge e adjacent to v by adding any vector in $O := (S_1 + S_2)^{\perp B}$. This does not change the parity of the number of crossings between independent pairs of edges (for f that are not adjacent to v we have $y_f^T(y_e + z) \equiv y_f^T y_e$; and f which are adjacent to v are not independent of e). Note that the modification in the $\mathbb{Z}_2$-genus case also preserves orientability (by choice of B).

We are going to modify the y_e-labels of edges adjacent to v (by adding vectors in O) as follows. For $i \in \{1, 2\}$ we do the following. Let $a := \dim(Z_i \cap S_i)$, $b := \dim(Z_i \cap (S_i + O))$, and $c := \dim(Z_i)$. Let $z_1, \ldots, z_c$ be a basis of Z_i such that 1) $z_1, \ldots, z_a$ is a basis of $Z_i \cap S_i$, 2) $z_1, \ldots, z_b$ is a basis of $Z_i \cap (S_i + O)$, and 3) $z_{a+1}, \ldots, z_b \in Z_i \cap O$. Such a basis can be constructed as follows: first apply the Steinitz exchange lemma on a basis of $Z_i \cap S_i$ and a basis of $Z_i \cap O$ (the

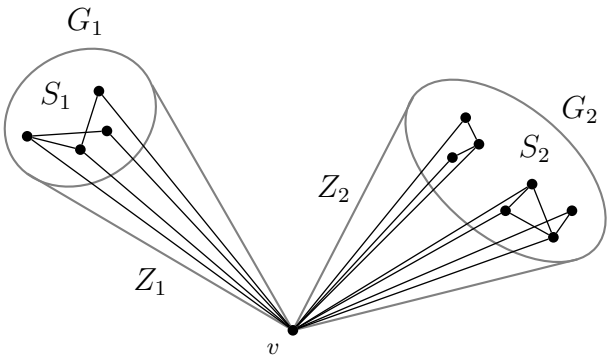

Fig. 2. Graph G with cutvertex v, block G_1 and union of remaining blocks G_2. Vector space Z_i (S_i) is generated by labels of edges in G_i (not) incident to v.

basis of $Z_i \cap S_i$ will be extended by vectors in the basis of $Z_i \cap O$ to a basis of $Z_i \cap (O + S_i)$); then apply the Steinitz exchange lemma on the resulting basis and a basis of Z_i. For each edge in G_i adjacent to v we relabel $y_e = \alpha_1 z_1 + \cdots + \alpha_c z_c$ by setting $\alpha_{a+1} = \alpha_{a+2} = \cdots = \alpha_b = 0$ (note that this corresponds to adding an element of O to y_e). After the modification (which also changed Z_i) we have that $z_1, \ldots, z_a$ is a basis of $Z_i \cap (S_i + O)$ and also a basis of $Z_i \cap S_i$. Thus we have

$$Z_i \cap (S_i + O) = Z_i \cap S_i. \tag{11}$$

Let $e_1 \in E_1 \setminus F$ and let $e_2 \in E_2 \setminus F$ be such that e_2 is not adjacent to v. Let C_1 be the unique cycle in $F + e_1$ and let C_2 be the unique cycle in $F + e_2$. Note that C_2 does not contain v (since F is a DFS spanning tree and e_2 is not adjacent to v). Thus C_1 and C_2 are disjoint (C_1 is in G_1, C_2 is in G_2 and $V_1 \cap V_2 = \{v\}$) and hence, by Lemma 1, we have

$$y_{e_1}^T y_{e_2} \equiv 0.$$

Thus $(Z_1 + S_1) \subseteq S_2^{\perp B}$ and since $O \subseteq S_2^{\perp B}$ (by the definition of O) we also have $(Z_1 + S_1 + O) \subseteq S_2^{\perp B}$. By symmetry we also have $(Z_2 + S_2 + O) \subseteq S_1^{\perp B}$ and hence, by Lemma 2 and (10), we obtain

$$\dim(Z_i + S_i + O) + \dim(S_{3-i}) \leq \hat{t}. \tag{12}$$

Thus we have (using $\dim(A + B) = \dim(A) + \dim(B) - \dim(A \cap B)$)

$$\dim(Z_i) + \dim(S_i + O) - \dim(Z_i \cap (S_i + O)) + \dim(S_{3-i}) \leq \hat{t},$$

and (11) yields

$$\dim(Z_i) + \dim(S_i + O) - \dim(Z_i \cap S_i) + \dim(S_{3-i}) \leq \hat{t}. \tag{13}$$

Adding (13) for $i = 1, 2$ and simplifying (again using $\dim(A + B) = \dim(A) + \dim(B) - \dim(A \cap B)$) we obtain

$$\dim(Z_1 + S_1) + \dim(Z_2 + S_2) + \dim(S_1 + O) + \dim(S_2 + O) \leq 2\hat{t}. \tag{14}$$

We have

$$\begin{aligned}
&\dim(S_1 + O) + \dim(S_2 + O) \\
&= \dim(S_1) + \dim(O) - \dim(S_1 \cap O) + \dim(S_2) + \dim(O) - \dim(S_2 \cap O) \\
&= \dim(S_1 + S_2) + \dim(S_1 \cap S_2) + 2\dim(O) - \dim(S_1 \cap O) - \dim(S_2 \cap O) \\
&= \hat{t} + \dim(S_1 \cap S_2) + \dim(O) - \dim(S_1 \cap O) - \dim(S_2 \cap O) \\
&\geq \hat{t} + \dim(S_1 \cap S_2 \cap O) + \dim(O) - \dim(S_1 \cap O) - \dim(S_2 \cap O) \\
&= \hat{t} + \dim(O) - \dim((S_1 + S_2) \cap O) \geq \hat{t},
\end{aligned}$$

where in the first, second, and fourth equality we used $\dim(A+B)+\dim(A\cap B) = \dim(A)+\dim(B)$; in the third equality we used $\dim(S_1+S_2)+\dim(O) = \hat{t}$ (which follows from the definition of O and Lemma 2); in the first and the last inequality we used the monotonicity of dimension,

Plugging $\dim(S_1 + O) + \dim(S_2 + O) \geq \hat{t}$ into (14) we obtain

$$\dim(Z_1 + S_1) + \dim(Z_2 + S_2) \leq \hat{t} \leq t. \tag{15}$$

The $\mathbb{Z}_2$-genus case of the lemma now follows from Lemma 3, using the argument from the proof of Lemma 6 (the final part after equation (8)) giving us $\mathbb{Z}_2$-embeddings of G_1 and G_2 on two surfaces which have g handles total. In the $\mathbb{Z}_2$-Euler genus case we apply Lemma 4, using the argument from the proof of Lemma 7 (the final part after equation (9)). □

The Hanani-Tutte theorem for surface $\mathcal{S}$ would—if true—state that if a graph has a $\mathbb{Z}_2$-embedding on surface $\mathcal{S}$, then it can be embedded on $\mathcal{S}$. Theorem 1 implies that in a search for counterexamples we can concentrate on 2-connected graphs. For the projective plane this result was obtained using much simpler means in [3, Lemma 2.4].

Corollary 1. *A minimal counterexample to the Hanani-Tutte theorem on any surface $\mathcal{S}$ is 2-connected.*

Proof. Suppose G is a minimal counterexample to Hanani-Tutte on some surface $\mathcal{S}$. If G is not connected, let $G_1, \ldots, G_k$, $k \geq 2$ be its connected components. If $\mathcal{S}$ is orientable, let g be the genus of $\mathcal{S}$. Then $g \geq \mathbf{g}_0(G) = \sum_{i=1}^{k} \mathbf{g}_0(G_i) = \sum_{i=1}^{k} \mathbf{g}(G_i) = \mathbf{g}(G)$, where the first equality is true by Lemma 6, and the second equality because G is minimal (and the third is a standard property of the genus of a graph). Therefore, $\mathbf{g}(G) \leq g$, so G can be embedded in $\mathcal{S}$, meaning it cannot be a counterexample. If $\mathcal{S}$ is non-orientable we can make essentially the same argument with Lemma 7 replacing Lemma 6, and Euler genus replacing genus. If G is connected, but not 2-connected, we repeat the same argument with Theorem 1 replacing the two lemmas, and using block additivity of (Euler) genus to conclude that $\sum_{i=1}^{k} \mathbf{g}(G_i) = \mathbf{g}(G)$ where the G_i are the blocks of G having cutvertex v in common. □

4 Questions

The (Euler) genus of a graph is an obvious upper bound on the $\mathbb{Z}_2$-(Euler) genus of a graph, but are they always the same?

Conjecture 1. The $\mathbb{Z}_2$-(Euler) genus of a graph equals its (Euler) genus.

The truth of this conjecture would imply the Hanani-Tutte theorem for arbitrary surfaces, so we have to leave the question open. The block additivity result from Section 3 implies that a minimal counterexample to the conjecture (if it exists) and, thereby, to the Hanani-Tutte theorem on an arbitrary surface, can be assumed to be 2-connected (since it cannot have a cut-vertex).

A much more modest goal than Conjecture 1 would be to bound the standard (Euler) genus in the $\mathbb{Z}_2$-(Euler) genus: Are there functions f and g so that $\mathbf{g}(G) \leq f(\mathbf{g}_0(G))$ and $\mathbf{eg}(G) \leq g(\mathbf{eg}_0(G))$?

In the absence of a proof of Conjecture 1, we can ask what other results for (Euler) genus also hold for $\mathbb{Z}_2$-(Euler) genus. For example, is the computation of the $\mathbb{Z}_2$-(Euler) genus **NP**-hard (as it is for (Euler) genus [9])? And is $\mathbb{Z}_2$-embeddability decidable in polynomial time for a fixed surface $\mathcal{S}$ (as it is for embeddability [10])? One could also try to extend the block additivity result: if G_1 and G_2 are two edge-disjoint graphs with $|V(G_1) \cap V(G_2)| = 2$, is it true that $|\,\mathbf{g}_0(G) - (\mathbf{g}_0(G_1) + \mathbf{g}_0(G_2))| \leq 1$? (This inequality is known to be true for the standard genus, a result by Decker, Glover, Huneke, and Stahl, see [4, Section 4.4]).

References

1. Sarkaria, K.S.: A one-dimensional Whitney trick and Kuratowski's graph planarity criterion. Israel J. Math. 73(1), 79–89 (1991)
2. Chojnacki (Haim Hanani), C.: Über wesentlich unplättbare Kurven im dreidimensionalen Raume. Fundamenta Mathematicae 23, 135–142 (1934)
3. Pelsmajer, M.J., Schaefer, M., Stasi, D.: Strong Hanani–Tutte on the projective plane. SIAM Journal on Discrete Mathematics 23(3), 1317–1323 (2009)
4. Mohar, B., Thomassen, C.: Graphs on surfaces. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore (2001)
5. Levow, R.B.: On Tutte's algebraic approach to the theory of crossing numbers. In: Hoffman, F., Levow, R.B., Thomas, R.S.D. (eds.) Proceedings of the Third Southeastern Conference on Combinatorics, Graph Theory and Computing, Boca Raton, Fla., Florida Atlantic University, pp. 315–314 (1972)
6. Stillwell, J.: Classical topology and combinatorial group theory. Second edn, 2nd edn. Graduate Texts in Mathematics, vol. 72. Springer, New York (1993)
7. de Longueville, M.: A course in topological combinatorics. In: Universitext, Springer, New York (2013)
8. Babai, L., Frankl, P.: Linear algebra methods in combinatorics with applications to geometry and computer science. Department of Computer Science, The University of Chicago (1992)
9. Thomassen, C.: The graph genus problem is NP-complete. J. Algorithms 10(4), 568–576 (1989)
10. Mohar, B.: A linear time algorithm for embedding graphs in an arbitrary surface. SIAM J. Discrete Math. 12(1), 6–26 (1999)