HAMBURGER BEITRÄGE ZUR MATHEMATIK

Heft 298

Embedding graphs in surfaces: MacLane's theorem for higher genus

Henning Bruhn & Reinhard Diestel

February 2008

Embedding graphs in surfaces: MacLane's theorem for higher genus*

Henning Bruhn and Reinhard Diestel

Abstract

Given a closed surface S, we characterise the graphs embeddable in S by an algebraic condition asserting the existence of a sparse generating set for their cycle space. When S is the sphere, the condition defaults to MacLane's planarity criterion.

1 Introduction

MacLane's well-known planarity criterion [9, 5] characterises the finite planar graphs in terms of their cycle space. As the (unoriented) *cycle space* $\mathcal{C}(G)$ of a graph G we take the $\mathbb{Z}_2$-vector space generated by the edge sets of cycles in G, with symmetric difference as addition. Its elements are those sets $F \subseteq E(G)$ such that every vertex of G is incident with an even number of edges in F. Call a family $\mathcal{F}$ of sets $F \subseteq E(G)$ *sparse* if every edge of G lies in at most two members of $\mathcal{F}$.

MacLane's planarity criterion can then be stated as follows:

MacLane's Theorem. *A finite graph is planar if and only if its cycle space is generated by some sparse family of (edge sets of) cycles.*

In this paper we generalise MacLane's theorem to embeddability criteria for arbitrary closed surfaces.

Our approach is motivated by simplicial homology, as follows. Let a connected graph G be embedded in a closed surface S of minimum Euler genus $\varepsilon := 2-\chi(S)$. Then S can be viewed as the underlying space of a 2-dimensional CW-complex C with 1-skeleton G. Its first homology group $Z_1(C;\mathbb{Z}_2)/B_1(C;\mathbb{Z}_2)$ is $\mathbb{Z}_2^\varepsilon$, the direct product of ε copies of $\mathbb{Z}_2$.

In graph theoretic language this means that the subspace $\mathcal{B}$ $(= B_1(C;\mathbb{Z}_2))$ spanned in $\mathcal{C}(G)$ $(= Z_1(C;\mathbb{Z}_2))$ by the set of face boundaries of G in S has codimension ε in $\mathcal{C}(G)$. Now the set of face boundaries is a sparse set of cycles. Thus, if G embeds in a surface of small Euler genus, at most ε, then G has a sparse set of cycles spanning a large subspace in $\mathcal{C}(G)$, one of codimension at most ε.

MacLane's theorem says that, for $\varepsilon = 0$, the converse implication holds too: if G has a sparse set of cycles whose span in $\mathcal{C}(G)$ has codimension at most $\varepsilon = 0$, then G embeds in the (unique) surface of Euler genus at most $\varepsilon = 0$, the

*This paper is an extended version of [4] providing more background and discussion.

sphere. Our initial aim, then, would be to prove this converse implication for arbitrary ε.

This naive extension soon runs into difficulties, and indeed is not true. In Section 3 we discuss the obstructions encountered as they arise, and modify our naive conjecture accordingly. The result will be a collection of theorems, presented in Section 4, which each characterise embeddability in a given surface, or in a surface of given Euler genus, by a condition akin to MacLane's planarity criterion that is both necessary and sufficient. All proofs are given in Section 5.

Some previous work in this direction can be found in the literature. Lefschetz [8] characterises the graphs that are embeddable in a given surface so that every face is bounded by a cycle. His theorem for orientable surfaces will follow from Theorem 5 (i). Lefschetz's theorem for non-orientable surfaces, stated in [8] without a formal proof, is incorrect; our Theorem 5 (ii) corrects and strengthens his result. Mohar [10] starts out from the necessary condition discussed earlier for embeddability in a surface of Euler genus at most ε, namely, that the graph must have a sparse set of cycles whose span in its cycle space has codimension at most ε. Unlike our plan here, Mohar does not strengthen this condition to one that is also sufficient, but establishes how much it implies as it is; the (best possible) result is that it implies embeddability in a surface of Euler genus at most 2ε. Širáň and Škoviera [12, 13] investigate when a given family of closed walks in a graph G can appear as face boundaries in an embedding of G in some surface, not necessarily of small genus (as will be our aim). Their work extends our discussions in Section 3 and provides an interesting backdrop for our proofs in Section 5, some of which use techniques they developed. We shall also use techniques of Edmonds [6], who studies embeddability in arbitrary surfaces in terms of duality.

2 General definitions and background

All graphs we consider are finite. Our notation will be that of [5], except that instead of 'multigraph' we say 'graph'. (Thus, our graphs may have loops and multiple edges, and degrees and connectivity are defined as they are in [5] for multigraphs. In particular, 2-connected graphs cannot have loops.) In the statements of some of our results we do not allow loops, but only to avoid unnecessary complication in our terminlogy: those theorems can be applied to graphs with loops by subdividing (and thereby eliminating) these.

The set of edges of a graph $G = (V, E)$ incident with a given vertex v is denoted by $E(v)$. When W is a walk in G, we denote the subgraph of G that consists of the edges on W and their incident vertices by $G[W]$; note that this need not be an induced subgraph of G. The (*unoriented*) *edge space* of G is the $\mathbb{Z}_2$ vector space of all functions $E \to \mathbb{Z}_2$ under pointwise addition. We usually write these as subsets of E, so vector addition becomes symmetric difference of edge sets. The (*unoriented*) *cycle space* $\mathcal{C}(G)$ of G is the subspace of $\mathcal{E}(G)$ generated by *circuits*, the edge sets of cycles.

A triple (e, u, v) consisting of an edge $e = uv$ together with its ends listed in a specific order is an *oriented edge*. The two oriented edges corresponding to e are its two *orientations*, denoted by $\vec{e}$ and $\overleftarrow{e}$. Thus, $\{\vec{e}, \overleftarrow{e}\} = \{(e, u, v), (e, v, u)\}$, but we cannot generally say which is which. Given a set E of edges, we write $\vec{E}$ for the set of their orientations, two for every edge in E.

The *oriented edge space* $\vec{\mathcal{E}}(G)$ of $G = (V, E)$ is the real vector space of all functions $\phi\colon \vec{E} \to \mathbb{R}$ satisfying $\phi(\overleftarrow{e}) = -\phi(\vec{e})$ for all $\vec{e} \in \vec{E}$. When $v_0 \dots v_{k-1}v_0$ is a cycle and $e_i := v_i v_{i+1}$ (with $v_k := v_0$), the function mapping the oriented edges (e_i, v_i, v_{i+1}) to 1, their inverses (e_i, v_{i+1}, v_i) to -1, and every other oriented edge to 0, is an *oriented circuit*. The *oriented cycle space* $\vec{\mathcal{C}}(G)$ is the subspace of $\vec{\mathcal{E}}(G)$ generated by the oriented circuits.

If G is connected and has n vertices and m edges, its oriented and its unoriented cycle space both have dimension

$$\dim \mathcal{C}(G) = \dim \vec{\mathcal{C}}(G) = m - n + 1. \tag{1}$$

A (closed) *surface* is a compact connected 2-manifold without boundary. It is *orientable* if it admits a triangulation whose 2-simplices (triangles) can be compatibly oriented. Equivalent conditions are that every triangulation has this property, and that the surface does not contain a Möbius strip [2].

An *n-dimensional CW-complex*, or *n-complex*, is a finite set C of open balls $B^i_j \subseteq \mathbb{R}^i$ with $i \leq n$, called *i-cells*, that have disjoint closures and whose union is made into a topological space $|C|$ as follows. The union C^0 of all 0-cells (which are singletons, so C^0 is just a set of points) carries the discrete topology. Assume now that the union of all i-cells with $i \leq k < n$, the *k-skeleton* C^k of C, has been given a topology, and denote this space by $|C^k|$. For every $(k+1)$-cell $B^{k+1}_j \in C$ choose a continuous *attachment map* $f_j : \partial B^{k+1}_j \to |C^k|$ from its boundary ∂B^{k+1}_j in $\mathbb{R}^{k+1}$ to $|C^k|$. Then give $|C^{k+1}|$ the quotient topology of the (disjoint) union of $|C^k|$ with all the closures of the B^{k+1}_j obtained by identifying every $x \in \partial B^{k+1}_j$ with $f_j(x)$.

Every graph G is a 1-complex, with vertices as 0-cells and edges as 1-cells. A topological embedding of G in another space S is a *2-cell-embedding* if G is the 1-skeleton of a 2-complex C such that the embedding of G in S extends to a homeomorphism $\varphi\colon |C| \to S$. The images under φ of the 2-cells of C are the *faces* of G in S. If S is a surface, their attachment maps define closed walks in G. These walks are unique up to cyclic shifts and orientation, a difference we shall often ignore. We thus have one such walk (with two orientations) assigned to each face, and call this family the (unique) *family of facial walks*. If W is the facial walk of some face f, then φ maps the subgraph $G[W]$ onto the frontier of f in S, and we call $G[W]$ the *boundary* of the face f.

Given a surface S, consider any 2-cell-embedding of any graph in S. Let n be its number of vertices, m its number of edges, and ℓ its number of faces in S. Euler's theorem tells us that $n - m + \ell$ is equal to a constant $\chi(S)$ depending only on S (not on the graph), the Euler characteristic of S. The *Euler genus* $\varepsilon(S)$ of S is defined as the number $2 - \chi(S)$. Euler's theorem then takes the following form, which we refer to as *Euler's formula*:

$$\varepsilon(S) = m - n - \ell + 2. \tag{2}$$

Given a graph G, let $\varepsilon = \varepsilon(G)$ be minimum such that G has a topological embedding φ in a surface of Euler genus at most ε. This ε is the *Euler genus* of G, and any such φ is a *genus-embedding* of G. Every connected graph has a genus-embedding that is a 2-cell-embedding [11, p. 95]. If G has components $G_1, \dots, G_n$, then $\varepsilon(G) = \varepsilon(G_1) + \cdots + \varepsilon(G_n)$, a fact referred to as *genus additivity* [11]. (The same is true for blocks rather than components, but we do not need this.)

We say that a family $\mathcal{W}$ of walks *covers* a subgraph H of G (often given in terms of its edge set) if every edge of H lies on some walk of $\mathcal{W}$. It covers an edge e *k times* if $k = \sum_{W \in \mathcal{W}} k_W(e)$, where $k_W(e)$ is the number of occurrences of e on W (irrespective of the direction in which W traverses e). $\mathcal{W}$ is a *double cover* of G if it covers every edge of G exactly twice. A walk is *non-trivial* if it contains an edge.

Given a walk W in G, we write $c(W)\colon E(G) \to \mathbb{Z}_2$ for the function that assigns to every edge e the number of times that W traverses e (in either direction), taken mod 2. Informally, we think of $c(W)$ as its support, the set of edges that appear an odd number of times in W. The *dimension* of a family $\mathcal{W}$ of walks, $\dim \mathcal{W}$, is the dimension of the subspace spanned in $\mathcal{E}(G)$ by the functions (or sets) $c(W)$ with $W \in \mathcal{W}$. If the walks are closed, their $c(W)$ lie in $\mathcal{C}(G)$; then the *codimension* of $\mathcal{W}$ in $\mathcal{C}(G)$ is the number $\dim \mathcal{C}(G) - \dim \mathcal{W}$.

Taking the natural orientation of W into account, we write $\vec{c}(W)$ for the function that assigns to every $\vec{e} \in \vec{E}$ the number of times that W traverses e in the direction of $\vec{e}$ minus the number of times that W traverses e in the direction of $\overleftarrow{e}$, and assigns 0 to any $\vec{e}$ with e not on W. The *oriented dimension* of a family $\mathcal{W}$ of walks, $\overrightarrow{\dim}\, \mathcal{W}$, is the dimension of the subspace of $\vec{\mathcal{E}}(G)$ spanned by the functions $\vec{c}(W)$ with $W \in \mathcal{W}$. If the walks are closed, their $\vec{c}(W)$ lie in $\vec{\mathcal{C}}(G)$; then the *codimension* of $\mathcal{W}$ in $\vec{\mathcal{C}}(G)$ is the number $\dim \vec{\mathcal{C}}(G) - \overrightarrow{\dim}\, \mathcal{W}$.

3 Reconstructing a surface

MacLane's theorem offers a necessary and sufficient condition for embeddability in a fixed surface, the sphere. Our aim is to find a similar condition characterising embeddability in an arbitrary but fixed surface S.

To illustrate what we mean by 'similar', let us think of MacLane's theorem as listing some properties of the facial cycles of a plane graph—sparseness and generating the entire cycle space—which, together, imply the following: that whenever we have *any* collection of cycles with these properties and attach a 2-cell to each of them, the 2-complex obtained is homeomorphic to the sphere. (This, indeed, is the outline of the standard topological proof of MacLane's theorem.)

For an arbitrary surface S, we are thus looking for a similar list of properties shared by the facial cycles of all graphs suitably embedded in S (with a genus-embedding, say) such that, given any family of cycles with these properties in a graph G, attaching a 2-cell along every cycle in this family turns G into a copy of S. One of those properties should be sparseness: if more than two 2-cells meet in an edge, the complex obtained will not be a surface. Following the homological approach outlined in the introduction, we might complement this by requiring that our cycles span a large enough subspace of the cycle space of G:

Naive Conjecture. *A graph G embeds in a surface S if and only if G has a sparse set of cycles whose span in $\mathcal{C}(G)$ has codimension at most $\varepsilon(S)$ in $\mathcal{C}(G)$.*

Notice that this conjecture can be true only if embeddability in a surface S depends only on $\varepsilon(S)$. For $\varepsilon = 0$ this is not an issue, since the sphere is the only surface with $\varepsilon = 0$. For even $\varepsilon > 0$, however, there are two surfaces of Euler genus ε—one orientable and one non-orientable—and the corresponding classes of graphs embeddable in them do not coincide. (Indeed, large projective-planar grids have unbounded orientable genus [3], while K_7 can be embedded in the torus but not in the Klein bottle [7].) Our first aim, therefore, will be to characterise embeddability not in a given surface S, but in 'some' surface of given Euler genus—in other words, to characterise the graphs of given Euler genus.

Another flaw in the Naive Conjecture is its reference to cycles: for surfaces other than the sphere, even genus-embeddings of 2-connected graphs can have facial walks that are not cycles. (For example, we can embed the graph G of Figure 2 in the torus by running the edge $e = uv$ along a handle added to the sphere to join two triangular faces containing u and v, respectively. Then e lies on the boundary of only one face, whose facial walk contains it twice and therefore is not a cycle. Zha [14] constructed for every surface S other than the sphere and the projective plane a 2-connected graph that has a genus-embedding in S but no embedding whose facial walks are all cycles.)

With these two modifications, our conjecture might become the following:

Revised Conjecture. *For every integer $\varepsilon \geq 0$, a graph G embeds in a surface of Euler genus at most ε if and only if it has a family of closed walks that covers every edge at most twice and whose codimension in $\mathcal{C}(G)$ is at most ε.*

However, as noticed by various authors [8, 10, 13], this is still not true: our list of properties of facial cycles—so far, sparseness and large dimension—needs a further addition.

To illustrate this, consider the plane graph A_1 shown in Figure 1. Let G be obtained from A_1 by identifying the vertices u and v. This graph G is one of the 35 forbidden minors that characterise embeddability in the projective plane (Archdeacon [1]), so $\varepsilon(G) \geq 2$.

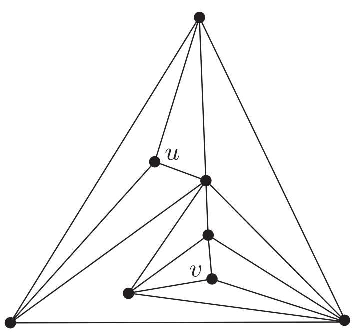

Figure 1: Identify u and v to obtain a graph G with $\varepsilon(G) \geq 2$

Let $\mathcal{W}$ denote the family of facial walks of A_1. The subspace it spans in $\mathcal{C}(G)$ is the cycle space of A_1. By (1), and since G has one vertex less than A_1 but the same number of edges, we deduce that

$$\dim \mathcal{W} = \dim \mathcal{C}(A_1) = \dim \mathcal{C}(G) - 1.$$

By the Revised Conjecture for $\varepsilon = 1$, this implies that G can be embedded in the projective plane—which it cannot.

To see what went wrong, let us form the 2-complex obtained by pasting a 2-cell on every walk in $\mathcal{W}$: the complex that 'should' be the projective plane but is not. The solution to the paradox is that this complex is not a surface at all: it is the pseudosurface obtained from a sphere by identifying two points.

To rule out this type of counterexample we could require that, for every vertex v, no proper subset of those of our given walks that pass through v can combine to a flat neighbourhood of v when we attach 2-cells to these walks. Since the facial walks in any 2-cell embedding of a graph have this property, it would certainly be an acceptable addition to our list. (In MacLane's theorem no such requirement is needed, because it follows; we shall prove this after stating Theorem 1 below.) If $\mathcal{W}$ is a double cover of G as in the example, or at least a cover, this flatness condition is sufficient and, indeed, the additional requirement we shall impose in Section 4 will then reduce to this. However if $\mathcal{W}$ does not cover G we need to be yet more careful. Our next example shows why.

Consider the plane graph A_2 shown in solid lines in Figure 2. Let G be obtained from A_2 by adding the edge uv. This graph G is another of Archdeacon's 35 forbidden minors for the projective plane, so again $\varepsilon(G) \geq 2$.

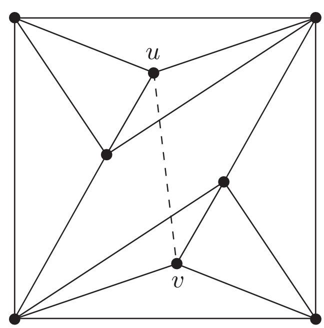

Figure 2: Add the edge uv to obtain a graph G with $\varepsilon(G) \geq 2$

As before, the subspace $\mathcal{W}$ spanned in $\mathcal{C}(G)$ by the facial walks of A_2 is the cycle space of A_2. By (1), and since G has one more edge than A_2 but the same number of vertices, we deduce that

$$\dim \mathcal{W} = \dim \mathcal{C}(A_2) = \dim \mathcal{C}(G) - 1.$$

By the Revised Conjecture for $\varepsilon = 1$, this implies that G can be embedded in the projective plane—which it cannot.

Again let us see what goes wrong in the 2-complex C formed by attaching a 2-cell to each walk in $\mathcal{W}$. While we obtain from the walks in $\mathcal{W}$ a flat neighbourhood around each vertex, it is impossible to extend $\mathcal{W}$ so as to include the edge uv without producing a non-flat neighbourhood at u and at v: the only way to do this is to add a walk just along uv and back, and pasting a disc onto this walk will add a second sphere, touching the sphere of $|C|$ in u and v.

To rule out counterexamples such as this, we could simply require that our family $\mathcal{W}$ should cover all the edges of G. This would be an acceptable addition to our list of requirements on $\mathcal{W}$ in that every collection of facial walks of an embedded graph satisfies it. However it would be against the spirit of MacLane's theorem—that marrying a purely graph-theoretical sparseness condition on $\mathcal{W}$ to a purely algebraic richness condition can yield a characterisation of planarity. Requiring that $\mathcal{W}$ cover the edges of G would spoil this dichotomy by adding a graph-theoretical richness condition.

Our solution to this dilemma will be to strengthen the requirement of 'sparseness at vertices' on $\mathcal{W}$, discussed after our first example, as follows: we shall require that no subfamily $\mathcal{U}$ of $\mathcal{W}$ (plus discs) shall form a flat neighbourhood of a vertex v unless it covers all the edges at v, regardless of whether they lie on walks of $\mathcal{W} \setminus \mathcal{U}$ or not. This condition will rule out both the above examples, while the graph-theoretical richness requirement it entails—that $\mathcal{U}$ must cover $E(v)$—will at least be kept local.

There is one more twist. Although, as we shall prove, our sparseness condition is now strong enough to ensure that graphs with a sparse family of closed walks of codimension at most ε embed in a surface of Euler genus at most ε, it is not true that pasting a disc on each of those walks will yield such as surface. For example, consider in a graph drawn on the sphere two vertices that lie on a common face boundary W. Identifying these two vertices into a new vertex v turns the sphere into a pseudosurface S on which the old facial walks still bound discs, so attaching discs to the walks after identification yields this pseudosurface. But those facial walks still form a sparse family: any non-empty subfamily summing to zero at v must contain W, but then it contains edges from both of the 'two disjoint disc neighbourhoods' of v on S and hence contains all the facial walks through v and thus covers $E(v)$.

Fortunately, there is a way to dissolve such singularities: rather than pasting a disc on W along its original orientation (which will result in a pseudosurface), we change the orientation of half of W, reversing it as shown in Figure 3 on one of its two segments between its two visits to v. (Recall that, after identification, W passes through v twice.) As the reader may verify, this alteration dissolves the singularity of our pseudosurface, turning it into a projective plane. In general, we shall prove that all singularities that can arise from pasting discs on a sparse family of closed walks can be dissolved in this way.

The ideas discussed so far will enable us to characterise, for any given ε, the graphs embeddable in either the orientable or the non-orientable surface of Euler genus ε (Theorem 1). To distinguish between the two, we shall have to refine our sparseness condition at vertices once more, and make use of the oriented cycle space. The key observation is that, given a 2-cell embedding of a graph G in a surface S, the facial walks—suitably oriented—will sum to zero in $\vec{\mathcal{C}}(G)$ when S is orientable, but will never sum to zero if S is non-orientable. With this observation suitably implemented, we shall finally be able to derive our desired MacLane-type characterisation of the graphs embeddable in a given surface, Theorem 5.

4 Statement of results

Recall that a family $\mathcal{F}$ of subsets of $E(G)$ is *sparse* if every edge of G lies in at most two members of $\mathcal{F}$. Similarly, we shall call a family of walks *sparse at* an edge e if it covers e at most twice. In view of our discussion in Section 3, we now wish to supplement this by a sparseness requirement at vertices.

Given a family $\mathcal{W}$ of walks and a vertex v, let us call a non-empty subfamily $\mathcal{U}$ of the walks in $\mathcal{W}$ through v a *cluster at* v if $\sum_{W \in \mathcal{U}} c(W) \cap E(v) = \emptyset$ but $\mathcal{U}$ fails to cover $E(v)$. We say that $\mathcal{W}$ is *sparse* if it is sparse at all edges and does not have a cluster at any vertex. For families of edge sets rather than walks we retain our earlier notion of sparseness, meaning sparseness at edges.

We can now state our first extension of MacLane's theorem. It can be read as a characterisation of the graphs of given Euler genus:

Theorem 1. *For every integer $\varepsilon \geq 0$, a graph G can be embedded in some surface of Euler genus at most ε if and only if there is a sparse family of closed walks in G whose codimension in $\mathcal{C}(G)$ is at most ε.*

For $\varepsilon = 0$, Theorem 1 implies MacLane's theorem. This is not immediately obvious: one has to show that a sparse family $\mathcal{B}$ of edge sets of cycles generating $\mathcal{C}(G)$ (as in MacLane's theorem) must be sparse also as a family of walks, i.e., that it does not have any clusters. We may assume that G is 2-connected. Suppose that $\mathcal{B}$ has a cluster at a vertex v. Thus, there is a non-empty subfamily $\mathcal{F}$ of $\mathcal{B}$ whose edges at v sum to zero but which fails to cover some other edge vw at v. Choose $\mathcal{F}$ minimal, and pick an edge uv from a cycle in $\mathcal{F}$. As G is 2-connected, $G - v$ contains a u–w path P; then $C = uPwvu$ is a cycle. We claim that no set $\mathcal{B}' \subseteq \mathcal{B}$ can sum to C, contradicting the choice of $\mathcal{B}$. Indeed, since $\mathcal{B}$ is sparse and $\mathcal{F}$ sums to zero at v, every edge in $D := E(v) \cap \bigcup \mathcal{F}$ lies on exactly two cycles in $\mathcal{F}$ but not on any cycle in $\mathcal{B} \setminus \mathcal{F}$. The set of cycles in $\mathcal{B}'$ with an edge in D, therefore, is precisely $\mathcal{B}' \cap \mathcal{F}$. In particular, if $uv \in \sum \mathcal{B}'$ then $uv \in E' := \sum(\mathcal{B}' \cap \mathcal{F})$. Since every cycle in $\mathcal{B}' \cap \mathcal{F}$ has two edges in D, we know that $|E' \cap D|$ is even. Hence if $uv \in \sum \mathcal{B}'$, there must be another edge $e \neq uv$ in $E' \cap D = (\sum \mathcal{B}') \cap D$. This edge cannot be $vw \notin D$, so it does no lie on C. Thus, $\sum \mathcal{B}'$ differs from C either in uv or in e, i.e. $\sum \mathcal{B}' \neq C$ as claimed.

The forward implication of Theorem 1 is well known, and its proof will not be hard. In our proof of the backward implication we shall take a detour via 'locally sparse' families of walks, which we define next. (We shall also need this concept again to state and prove our second main result, Theorem 5 below.) In order to keep our terminology simple we shall now ban loops; this will be easy to undo when we later prove Theorem 1.

Let $W = v_1 e_1 \dots v_n e_n v_1$ be a closed walk in a loopless graph G, where the v_i are vertices and the e_i are edges. For a vertex v we call a subsequence $e_{j-1} v_j e_j$

of W with $v_j = v$ (where $e_0 := e_n$) a *pass* of W *through* the vertex v. Extending our earlier notation for walks, we write $c(e_{j-1}v_je_j) := \{e_{j-1}, e_j\}$ if $e_{j-1} \neq e_j$, and $c(e_{j-1}v_je_j) := \emptyset$ if $e_{j-1} = e_j$. In order keep track of how often a given walk passes through a given vertex, we shall consider the *family of all passes of W through v*, the family $(e_{j-1}v_je_j)_{j\in J}$ where $J = \{j : v_j = v, 1 \leq j \leq n\}$. Similarly, if $\mathcal{W} = (W_i)_{i\in I}$ is a family of walks then the *family of all passes of $\mathcal{W}$ through v* is the family $\mathcal{A}(\mathcal{W}, v) := (p_{ij})_{i\in I, j\in J_i}$ where, for each i, $(p_{ij})_{j\in J_i}$ is the family of all passes of W_i through v. Let us call a non-empty subfamily $\mathcal{F} \subseteq \mathcal{A}(\mathcal{W}, v)$ a *local cluster at v* if $\sum_{p\in\mathcal{F}} c(p) = \emptyset$ but $\mathcal{F}$ fails to cover $E(v)$. We say that $\mathcal{W}$ is *locally sparse* if $\mathcal{W}$ is sparse at all edges and has no local cluster at any vertex. Note that any locally sparse family of closed walks in G is sparse, since for every vertex v and every closed walk W we have $c(W) \cap E(v) = \sum_{p\in\mathcal{A}((W),v)} c(p)$.

The following equivalence, whose implication (ii)→(i) will be a lemma in our proof of the backward implication of Theorem 1, is weaker than that implication in that it requires local sparseness rather than just sparseness in (ii). But it is also stronger, in that it allows us to make our *given* walks into face boundaries.

Lemma 2. *Let $G = (V, E)$ be a loopless connected graph, $\mathcal{W}$ a family of closed walks in G, and $\varepsilon \geq 0$ an integer. Then the following two statements are equivalent:*

(i) *There is a surface S of Euler genus at most ε in which G can be 2-cell-embedded so that $\mathcal{W}$ is a subfamily of the family of facial walks.*

(ii) *There is a locally sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$ and includes $\mathcal{W}$.*

In order to make Lemma 2 usable for the proof of Theorem 1, we next have to address the task of turning a sparse family $\mathcal{W}$ of closed walks into a locally sparse family $\mathcal{W}'$ without changing its codimension in $\mathcal{C}(G)$. In fact, we shall be able to do much more: we shall obtain $\mathcal{W}'$ from $\mathcal{W}$ by merely changing the order in which a walk traverses its edges. This is not unremarkable: it means, for example, that by merely changing the order in which the offending boundary walk W in the example discussed at the end of Section 3 traverses its edges we can turn the resulting pseudosurface into a surface.

To do this formally, consider any family $\mathcal{W}$ of closed walks in G. Call a family $\mathcal{W}' = (\,W' : W \in \mathcal{W}\,)$ of closed walks *similar* to $\mathcal{W}$ if, for every $e \in E(G)$ and every $W \in \mathcal{W}$, the edge e occurs on W' as often as it does on W. Thus if $\mathcal{W}'$ is similar to $\mathcal{W}$ then $G[W'] = G[W]$ and $c(W') = c(W)$ for every $W \in \mathcal{W}$, and in particular $\dim \mathcal{W}' = \dim \mathcal{W}$. Note that although a family similar to a locally sparse family need not itself be locally sparse (which indeed is our reason for defining similarity), a family similar to a sparse family will always be sparse.

Our next step, then, will be to prove the following lemma:

Lemma 3. *For every sparse family $\mathcal{W}$ of closed walks in a connected loopless graph G there exists a locally sparse family $\mathcal{W}'$ similar to $\mathcal{W}$.*

Using Lemmas 2 and 3 it will be easy to prove the following equivalence, a more explicit version of Theorem 1:

Theorem 4. *Let G be a connected graph, $\mathcal{W}$ a family of closed walks in G, and $\varepsilon \geq 0$ an integer. Then the following statements are equivalent:*

(i) *There is a surface of Euler genus at most ε in which G can be 2-cell-embedded so that the family of facial walks has a subfamily similar to $\mathcal{W}$.*

(ii) *There is a sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$ and includes $\mathcal{W}$.*

While Theorem 1 characterises the graphs of given Euler genus, our initial aim was to characterise the graphs embeddable in a given surface S. This will be achieved by the following theorem, which is our main result.

Theorem 5. *Let S be any surface, and let ε denote its Euler genus. Let G be any loopless graph, and let k denote the number of its components.*

(i) *If S is orientable, then G can be embedded in S if and only if G has a double cover by a locally sparse family $\mathcal{W}$ of closed walks whose oriented dimension is at most $|\mathcal{W}|-k$ and which has codimension at most ε in $\vec{\mathcal{C}}(G)$.*

(ii) *If G is connected and S is not orientable, then G can be embedded in S if and only if there is a sparse family $\mathcal{W}$ of closed walks in G whose codimension in $\vec{\mathcal{C}}(G)$ is at most $\varepsilon - 1$.*

We conjecture that 'locally sparse' cannot be replaced by 'sparse' in (i). And we remark that the connectivity requirement in (ii) cannot be dropped. Indeed, consider a graph G consisting of k disjoint copies of a graph that can be embedded in the projective plane but not in the sphere. By (ii), G can be covered by a sparse family of closed walks that has codimension 0 in $\vec{\mathcal{C}}(G)$. However, G cannot be embedded in any surface of Euler genus less than k.

5 The proofs

Let $\mathcal{W}$ be a family of closed walks in a loopless graph G that is sparse at edges. Recall that, for each vertex $v \in G$, we denoted by $\mathcal{A}(\mathcal{W}, v)$ the family of all passes of $\mathcal{W}$ through v. As a tool for our proofs, let us define for every vertex v an auxiliary graph $H = H(\mathcal{W}, v)$ with vertex set $\mathcal{A}(\mathcal{W}, v)$. Its edge set will be a subset of $E(G)$, with incidences defined as follows. Whenever two distinct vertices p, q of H (i.e., passes that are distinct as family members—they may be equal as triples) share an edge $e \in G$, we let e be an edge of H joining p and q. If $\mathcal{W}$ contains a pass $p = eve$, we let e be a loop at p. Clearly, H has maximum degree at most 2, since a pass evf can be incident only with the edges e and f. (For example, if there are three edges e, f, g at v in G, and $\mathcal{W}$ contains the passes evf, fvg, gve, then these three passes and the three edges e, f, g form a triangle in H. As another example, if $\mathcal{W}$ has two passes consisting of the triple evf, or one pass evf and another pass fve, then these two passes are joined by the pair $\{e, f\}$ of double edges in H and have no other incident edge.) If $\mathcal{W}$ is a double cover of G, then every $H(\mathcal{W}, v)$ is 2-regular.

Note that if $\mathcal{W}$ covers $E(v)$, then $\mathcal{W}$ has a local cluster at v if and only if $H = H(\mathcal{W}, v)$ contains a non-spanning cycle. Thus, $\mathcal{W}$ is locally sparse if and only if (it is sparse at edges and) each of the graphs $H(\mathcal{W}, v)$ is either a forest—possibly empty—or, if $\mathcal{W}$ covers $E(v)$, a single cycle.

We begin with a lemma which says that sparse double covers by closed walks[1] are nearly independent: that $\dim \mathcal{W} = |\mathcal{W}| - 1$. We shall need this for the family of face boundaries in the proof of (i)→(ii) of Theorem 4, and again for arbitrary sparse families in the proof of Theorem 5.

Lemma 6. *Let $G = (V, E)$ be a connected graph, and let $\mathcal{W}$ be a sparse family of non-trivial walks.*

(i) *For every non-empty subfamily $\mathcal{U}$ of $\mathcal{W}$ that is not a double cover of G, the family $(\, c(U) : U \in \mathcal{U}\,)$ is linearly independent in $\mathcal{C}(G)$.*

(ii) *If $\mathcal{W}$ is a double cover then* $\dim \mathcal{W} = |\mathcal{W}| - 1$.

Proof. It suffices to prove (i), since this implies that $\dim \mathcal{W} \geq |\mathcal{W}| - 1$: then (ii) follows, since $\mathcal{W}$ covers every edge twice and hence $\sum_{W \in \mathcal{W}} c(W) = \emptyset$.

For a proof of (i), let $\mathcal{U}$ be given as stated. Suppose the assertion fails; then $\mathcal{U}$ has a non-empty subfamily $\mathcal{U}' \subseteq \mathcal{U}$ such that $\sum_{U \in \mathcal{U}'} c(U) = \emptyset$. Then any edge covered by $\mathcal{U}'$ is covered by it twice, so as $\mathcal{U}$ is not a double cover there exists an edge not covered by $\mathcal{U}'$. On the other hand, since $\mathcal{U}'$ is non-empty and its walks are non-trivial, $\mathcal{U}'$ covers some edge of G. Since G is connected, it therefore has a vertex v that is incident both with an edge that is covered by $\mathcal{U}'$ and an edge that is not. Denote by $\mathcal{U}'(v)$ the non-empty family of all walks in $\mathcal{U}'$ containing v. As

$$\sum_{U \in \mathcal{U}'(v)} c(U) \cap E(v) \subseteq \sum_{U \in \mathcal{U}'} c(U) = \emptyset,$$

and as $\mathcal{U}'(v)$ does not cover $E(v)$, $\mathcal{U}'(v)$ is a cluster at v, contradicting that $\mathcal{W}$ is sparse. □

Next, we show that locally sparse families extend to double covers. It is possible to deduce this from results of Širáň and Škoviera [13], but for simplicity we sketch a direct proof.

Lemma 7. *Let G be a loopless graph and $\mathcal{W}$ a locally sparse family of closed walks in G. Then $\mathcal{W}$ can be extended to a locally sparse double cover $\mathcal{W}'$ of G by closed walks.*

Proof. Let $\mathcal{W}' \supseteq \mathcal{W}$ be a maximal family of closed walks that is locally sparse. We show that $\mathcal{W}'$ is a double cover.

Suppose not. Let F be the set of edges in G not covered twice. Our aim is to find a closed walk W in (V, F) such that $\mathcal{W}'' := \mathcal{W}' \cup \{W\}$ is again locally sparse; this will contradict our maximal choice of $\mathcal{W}'$.

For every vertex v incident with an edge in F, consider the auxiliary graph $H(v) := H(\mathcal{W}', v)$ defined at the start of this section. Let us show that $H(v)$ is a (possibly empty) forest. Suppose not, and let U be the vertex set of a cycle

[1] Indeed by any edge sets without clusters: our proof of Lemma 6 will not use the fact that $\mathcal{W}$ is a family of walks.

in $H(v)$. Then $\sum_{u\in U} c(u) = 0$. By assumption v is incident with an edge $f \in F$, which thus lies in at most one pass of $\mathcal{W}'$ through v. As this pass has degree at most 1 in $H(v)$ it cannot be in U, which implies that U, as a family of passes, does not cover $E(v)$. Then, however, U is a local cluster at v—a contradiction to our assumption that $\mathcal{W}'$ is locally sparse.

The components of $H(v)$, therefore, are paths. The edges of these paths are precisely the edges at v which $\mathcal{W}'$ covers twice, those in $E(v) \setminus F$. For every such path P put $\partial P := \sum_{p\in V(P)} c(p)$; this is a set of two edges in $F \cap E(v)$, and all these 2-sets are disjoint. Let $C(v)$ be a cycle on $F \cap E(v)$ as its vertex set such that $E(C(v)) \supseteq \{\, \partial P : P \text{ is a component of } H(v)\,\}$. Call the edges in this last set *red*, and the other edges of $C(v)$ *green*. (We allow $C(v)$ to be a loop or to consist of two parallel edges.) Call the number of green edges of $C(v)$ incident with a given vertex f of $C(v)$ the *green degree* of f in $C(v)$.

The green degree in $C(v)$ of an edge $f \in F \cap E(v)$ equals $2-k$, where $k \in \{0,1\}$ is the number of times that $\mathcal{W}'$ covers f. (3)

To construct our additional walk W in (V, F), we start by picking a vertex v_0 of G that is incident with an edge $f_0 \in F$. Then $H(v_0)$ and $C(v_0)$ are defined. Let us construct a maximal walk $W = v_0 f_0 v_1 f_1 \ldots f_{n-1} v_n$ in (V, F) such that $f_{i-1}f_i$ is a green edge of $C(v_i)$ and these green edges are distinct for different i. To ensure that we do not use a green edge again, let us delete the green edges as we construct W inductively, $f_{i-1}f_i$ at the time we add $f_{i-1}v_i f_i$ to W. Note, for $i = 1, \ldots, n-1$ inductively, that assertion (3) still holds for f_{i-1} and f_i at v_i with $W_i := v_0 f_0 \ldots f_i v_{i+1}$ added to $\mathcal{W}'$ and the green edges $f_{j-1}f_j$ deleted for all j with $1 \le j \le i$. This implies that when W gets to v_i via f_{i-1}, there is still a green edge $f_{i-1}f$ in $C(v_i)$ at f_{i-1} at that time, so W can continue and leave v_i via $f =: f_i$—unless $v_i = v_0$ and $f = f_0$, for which the extended assertion (3) does not hold (and was not proved above). Hence when our construction of W terminates we have $v_n = v_0$, and f_{n-1} is joined to f_0 by a green edge of $C(v_0)$. Thus, W is indeed a closed walk, and $\mathcal{W}'' := \mathcal{W} \cup \{W\}$ is again sparse at edges.

It remains to show that $\mathcal{W}''$ has no local clusters at vertices. The passes of W through a vertex v are all triples evf such that ef is a green edge of $C(v)$. Adding these passes as new vertices to $H(v)$, with adjacencies as defined before, turns $H(v)$ into a graph $H'(v)$ that is either a single cycle containing all of $E(v)$ (if W 'traverses' every green edge of $C(v)$) or a disconnected graph whose components are still paths: $H'(v)$ cannot contain cycles other than a Hamilton cycle, because $C(v)$ is a single cycle. Therefore, as any family $\mathcal{F}$ of passes of $\mathcal{W}''$ through v with $\sum_{p\in\mathcal{F}} c(p) = \emptyset$ induces a cycle in $H'(v)$, this can happen only when $\mathcal{F}$ covers $E(v)$. Thus, $\mathcal{W}''$ is again locally sparse, contradicting the maximal choice of $\mathcal{W}'$. □

We remark that Lemma 7 remains true if we replace 'locally sparse' with 'sparse', but we will not need this.

Proof of Lemma 2. (i)→(ii) Extend $\mathcal{W}$ to the family $\mathcal{W}'$ of all the facial walks of G in S. Since S is locally homeomorphic to the plane, $\mathcal{W}'$ covers every edge of G twice, and elementary topological arguments show that $\mathcal{W}'$ cannot have a local cluster at any vertex. Hence $\dim \mathcal{W}' = |\mathcal{W}'| - 1$ by Lemma 6 (ii). Using (1) and Euler's formula (2), we deduce that

$$\dim \mathcal{C}(G) - \varepsilon = |E(G)| - |V(G)| + 1 - \varepsilon \leq |\mathcal{W}'| - 1 = \dim \mathcal{W}'$$

as desired.

(ii)→(i) Replacing $\mathcal{W}$ with the extension of $\mathcal{W}$ whose existence is asserted in (ii), we may assume that $\mathcal{W}$ itself is locally sparse and has codimension at most ε in $\mathcal{C}(G)$. Extending $\mathcal{W}$ by Lemma 7 if necessary, we may further assume that $\mathcal{W}$ is a double cover of G.

Let C be the 2-dimensional CW-complex obtained as follow. We start with G as its 1-skeleton. As the 2-cells we take disjoint open discs $D_W \subseteq \mathbb{R}^2$, one for each walk $W \in \mathcal{W}$, divide the boundary of D_W into as many segments as W is long, and map consecutive segments homeomorphically to consecutive edges in W.

In order for $S := |C|$ to be a surface, we have to check that every point has an open neighbourhood that is homeomorphic to $\mathbb{R}^2$. For points in the interior of 2-cells or edges, this is clear; recall that $\mathcal{W}$ is a double cover. Now consider a vertex v of G. Define $H(v)$ as earlier. Since $\mathcal{W}$ is a double cover, $H(v)$ is now 2-regular, and since $\mathcal{W}$ has no local cluster at v it contains no cycle properly. Hence, $H(v)$ is a cycle. For each pass $p = evf \in V(H(v))$ we let $D(p)$ be a closed disc whose interior lies inside a disc D_W such that p is a pass of W, choosing each $D(p)$ so that its boundary contains v and intersects W in one segment contained in $e \cup f$ and meeting both e and f. These discs $D(p)$ can clearly be chosen with disjoint interiors for different p. Using the elementary fact that the union of two closed discs intersecting in a common segment of their boundaries is again a disc, one easily shows inductively that the interior of the union of all the discs $D(p)$ is an open disc, and hence homeomorphic to $\mathbb{R}^2$. This completes the proof that S is a surface.

Since C is finite, S is compact. Since G is connected, so is S. Finally, Euler's formula (2) applied to C, together with (1), the trivial inequality of Lemma 6 (ii), and our assumption that $\mathcal{W}$ has codimension at most ε in $\mathcal{C}(G)$, yields

$$\begin{aligned}
\varepsilon(S) &= 2 - (|V(G)| - |E(G) + |\mathcal{W}|) \\
&= (|E(G)| - |V(G)] + 1) - (|\mathcal{W}| - 1) \\
&= \dim \mathcal{C}(G) - \dim \mathcal{W} \\
&\leq \varepsilon .
\end{aligned}$$

Thus, (i) is proved. □

We need an easy technical lemma relating $\overrightarrow{\dim}\, \mathcal{W}$ to $\dim \mathcal{W}$.

Lemma 8. *Let $G = (V, E)$ be a connected graph, and let $\mathcal{W} = (W_1, \ldots, W_n)$ be a sparse family of non-trivial walks.*

(i) $\overrightarrow{\dim}\, \mathcal{W} \geq \dim \mathcal{W}$.[2]

(ii) *If $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$ then there exist $\mu_i \in \{1, -1\}$ such that $\sum_{i=1}^{n} \mu_i \vec{c}(W_i) = 0$.*

[2]This is true regardless of whether $\mathcal{W}$ is sparse. But the special case proved here is all we need.

Proof. Assertation (i) will follow at once from the following claim:

If there exist $\lambda_1, \dots, \lambda_n \in \mathbb{R} \setminus \{0\}$ *such that* $\sum_{i=1}^n \lambda_i \vec{c}(W_i) = 0$ *in* $\vec{\mathcal{E}}(G)$, *then there are also* $\mu_1, \dots, \mu_n \in \{-1, 1\}$ *such that* $\sum_{i=1}^n \mu_i \vec{c}(W_i) = 0$. (4)

Indeed, whenever two walks W_i, W_j share an edge e, we have $|\lambda_i| = |\lambda_j|$ because $\mathcal{W}$ is sparse at e. Let H be the graph on $\{1, \dots, n\}$ in which ij is an edge whenever W_i and W_j share an edge. Then the values of $|\lambda_i|$ coincide for all i in a common component C of H, and letting $\mu_j := \lambda_j / \lambda_i$ for some fixed i and all j in C satisfies (4).

Let us now prove (ii). If $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$ there are $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ not all zero such that $\sum_{i=1}^n \lambda_i \vec{c}(W_i) = 0$. By (4), we may assume the λ_i to be in $\{1, 0, -1\}$. Applying Lemma 6 (i) to the subfamily $\mathcal{U}$ of the W_i with $\lambda_i \neq 0$ we see that the λ_i are in fact all non-zero, as desired. □

Next, let us prove Lemma 3, our tool for turning a sparse family of walks into a locally sparse one without changing the edge sets of its walks. The proof employs a trick from surface surgery to dissolve singularities, which we learnt from Edmonds [6]. In fact, we prove a slightly stronger statement:

Lemma 9. *For every sparse family* $\mathcal{W}$ *of closed walks in a connected loopless graph* G *there exists a locally sparse family* $\mathcal{W}'$ *similar to* $\mathcal{W}$. *If* $\mathcal{W}$ *is not locally sparse, then* $\mathcal{W}'$ *can be chosen so that* $\overrightarrow{\dim}\, \mathcal{W}' = |\mathcal{W}'|$.

Proof. For families $\mathcal{W}'$ of closed walks, define $\gamma(\mathcal{W}') := \sum_{v \in V(G)} \gamma_{\mathcal{W}'}(v)$ where $\gamma_{\mathcal{W}'}(v)$ denotes the number of components of $H(\mathcal{W}', v)$. Assuming that $\mathcal{W}$ is not locally sparse, we will construct a family $\mathcal{W}'$ similar to $\mathcal{W}$ such that $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$; we will further ensure that $\overrightarrow{\dim}\, \mathcal{W}' = |\mathcal{W}'|$. Since $\gamma(\mathcal{W})$ is bounded below by 0, this will prove the lemma.

Let us construct $\mathcal{W}'$. As $\mathcal{W}$ is not locally sparse, there must exist a local cluster at some vertex v. Seen in $H := H(\mathcal{W}, v)$ this local cluster forms a cycle. Since $\mathcal{W}$ is sparse, one of the vertices of C must be a pass $p = eve'$ of a walk $W \in \mathcal{W}$ which also contains a pass $q = fvf'$ that is a vertex in another component $D \neq C$ of H. Choose these passes so that W has a subwalk $ve' \dots fv$ not containing e or f'. Let W' be the closed walk obtained from W by reversing this subwalk (Figure 3), and let $\mathcal{W}'$ be obtained from $\mathcal{W}$ by replacing W with W'. Clearly, W' is again a closed walk, and $\mathcal{W}'$ is similar to $\mathcal{W}$.

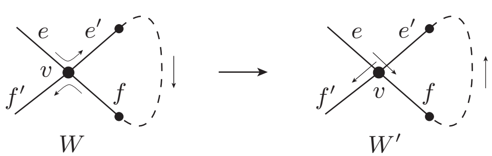

Figure 3: Turning W into W' by reversing the segment $ve' \dots fv$

Let us show that $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$. For vertices $u \neq v$ of G we have $H(\mathcal{W}', u) = H(\mathcal{W}, u)$, so $\gamma_{\mathcal{W}'}(u) = \gamma_{\mathcal{W}}(u)$. At v, however, we have $\gamma_{\mathcal{W}'}(v) < \gamma_{\mathcal{W}}(v)$, so $\gamma(\mathcal{W}') < \gamma(\mathcal{W})$. Indeed, $H' := H(\mathcal{W}', v)$ arises from H by the replacement of $p = eve' \in V(C)$ and $q = fvf' \in V(D)$ with two new vertices, $p' := evf$ and $q' := e'vf'$, and redefining the incidences for the edges $e, f, e', f' \in E(H) = E(H')$ accordingly. As one easily checks (see Figure 4), this has the effect of merging the components C and D of H into one new component, leaving the other components of H intact. Thus, the components of H' are those of H other than C and D, plus one new component arising from $(C-p)\cup(D-q)$ by adding the new vertex p' incident with e and f and the new vertex q' incident with e' and f' (leaving the other incidences of e, e', f, f' in H' as they were in H).

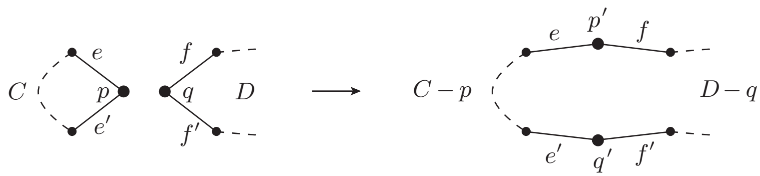

Figure 4: Merging the components C and D of H to form H'

It remains to show that $\vec{\dim}\,\mathcal{W}' = |\mathcal{W}'|$. First note that, if $C = e_1 \dots e_m$ where $e = e_1$ and $e' = e_m$ then $fe_1 \dots e_m f'$ is a subpath of C', the new component that arose from merging C and D.

Suppose now that $\vec{\dim}\,\mathcal{W}' < |\mathcal{W}'|$. Then for all $U \in \mathcal{W}'$ there are $\mu_U \in \{1, -1\}$ such that $\sum_{U\in\mathcal{W}'} \mu_U \vec{c}(U) = 0$ (Lemma 8 (ii)), and we may assume that $\mu_{W'} = 1$. Reversing the orientation of each $U \in \mathcal{W}'$ with $\mu_U = -1$ we obtain $\sum_{U\in\mathcal{W}'} \vec{c}(U) = 0$. Since the orientation of W' has not changed, $p' = e_1 vf$ and $q' = e_m vf'$ are still subwalks of W'. The orientations of the walks in $\mathcal{W}'$ induce orientations on the passes at v; therefore $\sum_{U\in\mathcal{W}'} \vec{c}(U) = 0$ implies that $\sum_{r\in V(C')} \vec{c}(r) = 0$, the passes r being interpreted as subwalks. Hence as $p' = e_1 vf \in V(C')$, each of the passes $e_{i+1} v e_i$ is traversed by some walk in $\mathcal{W}'$ in this order: e_{i+1} towards v, and e_i away from v ($i = 1, \dots, m-1$). In particular, e_m is traversed towards v in the pass $e_m v e_{m-1} \neq q'$. However, this is also the case in q'. As $\mathcal{W}'$ is sparse at edges, this implies $\sum_{r\in V(C')} \vec{c}(r) \neq 0$, a contradiction. □

Proof of Theorem 4. Denote by $\dot{G}$ the loopless graph obtained from G by subdividing every loop once. Note that there is an obvious isomorphism $\mathcal{C}(G) \doteq \mathcal{C}(\dot{G})$, and in particular, the two spaces have the same dimension.

To prove the implication (i)→(ii), consider an embedding of G as in (i). The embedding of G immediately induces an embedding of $\dot{G}$, so that there is a 1-1 correspondence between the facial walks $\dot{\mathcal{U}}$ of the embedding of $\dot{G}$ and the facial walks $\mathcal{U}$ of the embedding of G. Applying Lemma 2 to $\dot{\mathcal{U}}$, which is a double cover, we see that $\dot{\mathcal{U}}$ is locally sparse and of codimension $\le \varepsilon$ in $\mathcal{C}(\dot{G})$. Then the same holds for $\mathcal{U}$ with respect to $\mathcal{C}(G)$. Replacing in $\mathcal{U}$ the subfamily of $\mathcal{U}$ similar to $\mathcal{W}$ with $\mathcal{W}$ preserves both the sparseness of $\mathcal{U}$ and its dimension, so (ii) follows.

For a proof of the implication (ii)→(i), let $\mathcal{W}' \supseteq \mathcal{W}$ be the sparse family of codimension $\le \varepsilon$ in $\mathcal{C}(G)$ provided by (ii). Then the subdivided walks $\dot{\mathcal{W}}'$ in $\dot{G}$ are still sparse and have codimension $\le \varepsilon$ in $\mathcal{C}(\dot{G})$. We use Lemma 9 to turn $\dot{\mathcal{W}}'$ into a locally sparse family $\dot{\mathcal{W}}''$ similar to $\mathcal{W}'$, which, by Lemma 2, is a subfamily of the family $\dot{\mathcal{U}}$ of facial walks of an embedding of $\dot{G}$ in a surface of Euler genus at most ε. If each walk W in $\dot{\mathcal{U}}$ is a subdivision of a walk in G then the embedding of $\dot{G}$ induces one of G in which $\mathcal{W}$ is similar to a subfamily of the facial walks, since $\dot{\mathcal{U}}$ contains $\dot{\mathcal{W}}'' \sim \dot{\mathcal{W}}'$. This can fail only if W contains a pass eve through a subdividing vertex v. If it does, let f be the other edge of $\dot{G}$ at v. Then the subfamily $\mathcal{F} = \{eve\}$ of $\dot{\mathcal{U}}$ satisfies $\sum_{p \in \mathcal{F}} c(p) = 0$, but fails to cover f. Thus the local cluster $\mathcal{F}$ at v contradicts that $\mathcal{U}$ is locally sparse. □

Theorem 4 immediately implies Theorem 1 for connected graphs. To complete the proof of Theorem 1, it remains to reduce the disconnected to the connected case.

Proof of Theorem 1. For the forward direction, let G and ε be such that G embeds in a surface of Euler genus at most ε. Our aim is to find a certain family of closed walks of codimension at most ε, so there is no loss of generality in choosing ε minimum, i.e., in assuming that $\varepsilon = \varepsilon(G)$. Let $G_1, \ldots, G_n$ be the components of G. For each $i = 1, \ldots, n$ choose a genus-embedding $G_i \hookrightarrow S_i$. These can be chosen to be 2-cell-embeddings, and by genus additivity we have $\varepsilon_1 + \cdots + \varepsilon_n = \varepsilon$ for $\varepsilon_i := \varepsilon(S_i) = \varepsilon(G_i)$. For each i let $\mathcal{W}_i$ be the family of facial walks of G_i in S_i. By Theorem 4, the $\mathcal{W}_i$ are sparse and have codimension at most ε_i in $\mathcal{C}(G_i)$: as $\mathcal{W}_i$ already covers every edge of G_i twice, it cannot be extended to a larger sparse family. Since the G_i are vertex-disjoint, $\mathcal{W} := \mathcal{W}_1 \cup \cdots \cup \mathcal{W}_n$ is again sparse, and it has codimension at most $\varepsilon_1 + \cdots + \varepsilon_n = \varepsilon$ in $\mathcal{C}(G)$, since $\mathcal{C}(G)$ is the direct sum of the spaces $\mathcal{C}(G_i)$.

For a proof of the backward direction, let $\mathcal{W}$ be a sparse family of closed walks in G that has codimension at most ε in $\mathcal{C}(G)$. If G has components $G_1, \ldots, G_k$, say, write $\mathcal{W}_i$ for the subfamily of walks contained in G_i, and ε_i for the codimension of $\mathcal{W}_i$ in $\mathcal{C}(G_i)$. Then $\varepsilon(G_i) \le \varepsilon_i$, by (ii)→(i) of Theorem 4. Moreover, $\sum_{i=1}^{k} \varepsilon_i \le \varepsilon$, since $\mathcal{C}(G)$ is the direct sum of the spaces $\mathcal{C}(G_i)$. Hence, by genus additivity,

$$\varepsilon(G) = \sum_{i=1}^{k} \varepsilon(G_i) \le \sum_{i=1}^{k} \varepsilon_i \le \varepsilon\,.$$

Thus, G can be embedded in a surface of Euler genus at most ε. □

We finally come to the proof of Theorem 5. We need another easy lemma.

Lemma 10. *Let G be a loopless and connected graph. If $\mathcal{W}$ is the family of facial walks of an embedding of G in a surface S, then S is orientable if and only if* $\overrightarrow{\dim}\, \mathcal{W} < |\mathcal{W}|$.

Proof. If $\mathcal{W}$ is the family of facial walks of an embedding of G in S, insert a new vertex in every face and join it to all the vertices on the boundary of that face. This yields a triangulation of S. If S is orientable, we can orient the 2-simplices of this complex C (i.e., the newly created triangles) compatibly, so that every edge receives opposite orientations from the orientations of the two 2-simplices containing it. Then the 2-simplices triangulating a given face induce orientations on the edges of its boundary walk $W \in \mathcal{W}$ that either all coincide with their orientations induced by W or are all opposite to them. Let $\lambda_W := 1$ or $\lambda_W := -1$ accordingly. Then $\sum_{W\in\mathcal{W}} \lambda_W \vec{c}(W) = 0$, showing that $\vec{\dim}\,\mathcal{W} < |\mathcal{W}|$.

Conversely, if $\vec{\dim}\,\mathcal{W} < |\mathcal{W}|$ then, by Lemma 8 (ii), there are $\mu_W \in \{1, -1\}$, $W \in \mathcal{W}$, so that $\sum_{W\in\mathcal{W}} \mu_W \vec{c}(W) = 0$. Reversing the orientation of every W with $\mu_W = -1$ yields $\sum_{W\in\mathcal{W}} \vec{c}(W) = 0$. These new orientations of the boundary walks W therefore extend to compatible orientations of the 2-simplices of C, showing that S is orientable. □

Proof of Theorem 5. (i) We assume that G is connected; the general case then follows as in the proof of Theorem 1.[3] Suppose first that G can be embedded in S. Replacing S with a surface of smaller oriented genus if necessary, we may assume that this is a 2-cell embedding. (Any such replacement reduces ε, so this assumption entails no loss of generality.) By Lemma 2, the family $\mathcal{W}$ of facial walks is locally sparse and has codimension at most ε in $\mathcal{C}(G)$. Its codimension in $\vec{\mathcal{C}}(G)$ is no greater, since $\vec{\dim}\,\mathcal{W} \geq \dim \mathcal{W}$ by Lemma 8 (i), and $\dim \vec{\mathcal{C}}(G) = \dim \mathcal{C}(G)$ by (1). It remains to show that $\vec{\dim}\,\mathcal{W} \leq |\mathcal{W}| - 1$, which follows from Lemma 10.

For the converse implication of (i), Lemmas 6 (ii) and 8 (i) and our assumption about $\vec{\dim}\,\mathcal{W}$ give

$$\dim \mathcal{W} \leq \vec{\dim}\,\mathcal{W} \leq |\mathcal{W}| - 1 = \dim \mathcal{W},$$

with equality. By (1), then, also the codimension of $\mathcal{W}$ is the same in $\mathcal{C}(G)$ as in $\vec{\mathcal{C}}(G)$, at most ε. By (ii)→(i) of Lemma 2, there exists a surface S' with $\varepsilon' := \varepsilon(S') \leq \varepsilon$ in which G has a 2-cell-embedding with $\mathcal{W} =: (W_1, \ldots, W_n)$ as the family of facial walks. By Lemma 10, S' is orientable. Adding $(\varepsilon - \varepsilon')/2$ handles turns S' into a copy of S with G embedded in it, as desired.

(ii) For the forward implication let $\mathcal{W}$ be the family of facial walks of the given embedding. By Lemma 2, $\mathcal{W}$ is sparse. By Lemma 10, $\vec{\dim}\,\mathcal{W} = |\mathcal{W}|$. By (1) and (2), the codimension of $\mathcal{W}$ in $\vec{\mathcal{C}}(G)$ is $\varepsilon - 1$.

For the backward implication in (ii), let us assume first that the (unoriented) codimension of $\mathcal{W}$ in $\mathcal{C}(G)$ is also at most $\varepsilon - 1$. By Theorem 1, we can embed G in a surface S' of Euler genus $\varepsilon' \leq \varepsilon - 1$. The addition of $\varepsilon - \varepsilon' \geq 1$ crosscaps turns S' into a copy of S with G embedded in it.

We may therefore assume that $\mathcal{W}$ has codimension at least ε in $\mathcal{C}(G)$. Let us show that $\mathcal{W}$ is a double cover of G. If not, then Lemmas 8 (i) and 6 (i) imply

$$|\mathcal{W}| \geq \vec{\dim}\,\mathcal{W} \geq \dim \mathcal{W} = |\mathcal{W}|$$

with equality, so $\vec{\dim}\,\mathcal{W} = \dim \mathcal{W}$. By (1), this contradicts our assumption that the codimensions of $\mathcal{W}$ in $\mathcal{C}(G)$ and $\vec{\mathcal{C}}(G)$ differ. Moreover, by assumption and Lemma 6 we have

$$\dim \mathcal{C}(G) - \varepsilon \geq \dim \mathcal{W} \geq |\mathcal{W}| - 1 \geq \vec{\dim}\,\mathcal{W} - 1 \geq \dim \vec{\mathcal{C}}(G) - \varepsilon.$$

[3]Use the additivity of oriented genus rather than of Euler genus.

By (1), we have equality throughout. In particular, $\mathcal{W}$ has codimension exactly ε in $\mathcal{C}(G)$, and $\overrightarrow{\dim}\, \mathcal{W} = |\mathcal{W}|$. By Lemma 9 there is a locally sparse family $\mathcal{W}'$ similar to $\mathcal{W}$ such that $\overrightarrow{\dim}\, \mathcal{W}' = |\mathcal{W}'|$. Since $\mathcal{W}'$, like $\mathcal{W}$, is a double cover, $\mathcal{W}'$ is by Lemma 2 the family of facial walks of an embedding of G in a surface S' of Euler genus $\varepsilon' \leq \varepsilon$. By Lemma 10, S' is not orientable. Adding $\varepsilon - \varepsilon'$ crosscaps we turn S' into a copy of S with G embedded in it. □

Acknowledgement

We thank Bojan Mohar, Jozef Širáň and Martin Škoviera for updating us on what was known on the subject. We further thank a thoughtful referee of [4] for pointing out that the notion of sparseness at vertices (local or global), which we had earlier defined to imply that the family of walks should cover the graph, would benefit from making this requirement only locally. (Compare the discussion at the end of Section 3.)

References

[1] D. Archdeacon, *A Kuratowski theorem for the projective plane*, J. Graph Theory **5** (1981), 243–246.

[2] M.A. Armstrong, *Basic topology*, Springer-Verlag, 1983.

[3] L. Auslander, T.A. Brown, and J.W.T. Youngs, *The imbedding of graphs in manifolds*, J. Math. and Mech. **12** (1963), 629–634.

[4] H. Bruhn and R. Diestel, *MacLane's theorem for arbitrary surfaces*, to appear in *J. Combin. Theory (Series B)*.

[5] R. Diestel, *Graph Theory* (3rd edition), Springer-Verlag, 2005, Electronic edition available at: `http://www.math.uni-hamburg.de/home/diestel/books/graph.theory`.

[6] J. Edmonds, *On the surface duality of linear graphs*, J. Res. Natl. Bur. Stand., Sect. B **69** (1965), 121–123.

[7] P. Franklin, *A six color problem*, J. Math. Massachusetts **13** (1934), 363–369.

[8] S. Lefschetz, *Applications of algebraic topology*, Springer-Verlag, 1975.

[9] S. MacLane, *A combinatorial condition for planar graphs*, Fund. Math. **28** (1937), 22–32.

[10] B. Mohar, *Generalization of MacLane's characterization of planar graphs*, Congr. Numerantium **56** (1987), 231–236.

[11] B. Mohar and C. Thomassen, *Graphs on surfaces*, Johns Hopkins, 2001.

[12] J. Širáň and M. Škoviera, *Oriented relative embeddings of graphs*, Applicationes Math. **XIX** (1987), 589–597.

[13] ———, *Relative embeddings of graphs on closed surfaces*, Math. Nachr. **136** (1988), 275–284.

[14] X. Zha, *On minimum-genus embeddings*, Disc. Math. **149** (1996), 261–278.

Henning Bruhn
Reinhard Diestel
Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany

Version 11.2.2008