

Maximum genus embeddings of dense eulerian graphs with specified faces

M. N. Ellingham^{a,1}, Joanna A. Ellis-Monaghan^b

^a*Department of Mathematics, 1326 Stevenson Center, Vanderbilt University, Nashville, Tennessee 37240*

^b*Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Science Park 105-107, 1098 XH Amsterdam, the Netherlands*

Abstract

We give a density condition for when, subject to a necessary parity condition, an eulerian graph or digraph may be cellularly embedded in an orientable surface so that it has exactly two faces, each bounded by an euler circuit, one of which may be specified in advance. More generally, suppose that every vertex in an n -vertex eulerian digraph (loops and multiple arcs allowed) has at least $(4n + 2)/5$ neighbors, and specify any decomposition of the arcs into disjoint directed circuits (closed trails). We show that such a digraph has an orientable embedding in which the given circuits are facial walks and there are exactly one or two other faces. This embedding then has maximum genus relative to the given circuits being facial walks. When there is only one other face, it is necessarily bounded by an euler circuit. Consequently, if the numbers of vertices and edges have the same parity, a sufficiently dense digraph D with a given directed euler circuit T has an orientable embedding with exactly two faces, each bounded by an euler circuit, one of which is T . These results for digraphs give analogous results for graphs as immediate corollaries. The main theorem encompasses several special cases in the literature, such as when the digraph is a tournament.

Keywords: Eulerian, bi-eulerian, orientable graph embedding, maximum genus, circuit decomposition, digraph, directed embedding

2000 MSC: 05C10, 05C45

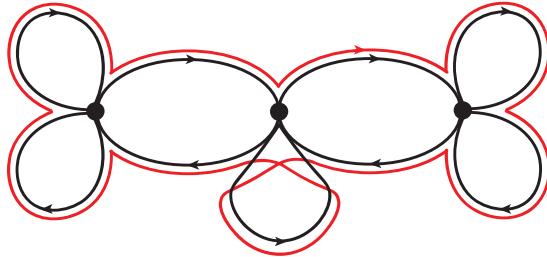
1. Introduction

When does a graph or digraph have an orientable *bi-eulerian* embedding? This is, when can it be cellularly embedded in a orientable surface so that it has exactly two faces, each bounded by an euler circuit, such as shown in Figure 1? Is it possible to specify one of the euler circuits in advance?

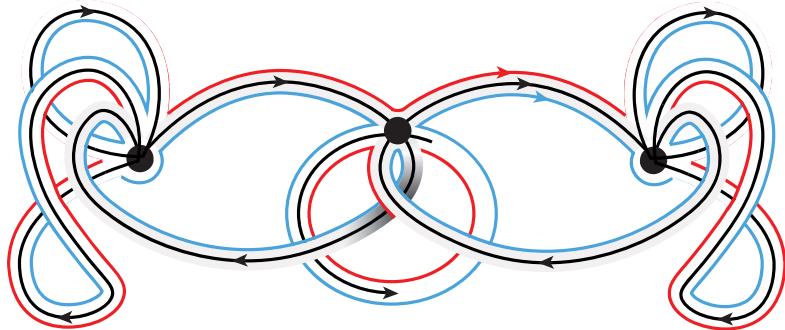
Email addresses: mark.ellingham@vanderbilt.edu (M. N. Ellingham), jellismonaghan@gmail.com (Joanna A. Ellis-Monaghan)

URL: <https://math.vanderbilt.edu/ellingmn/> (M. N. Ellingham),
<https://sites.google.com/site/joellismonaghan/> (Joanna A. Ellis-Monaghan)

¹Supported by Simons Foundation awards 429625 and MPS-TSM-00002760



(a) A digraph with a specified directed euler circuit.



(b) A directed embedding of the digraph with two faces.

Figure 1: A bi-eulerian orientable embedding of a digraph, with the specified euler circuit as one face (red), and one other euler circuit face (blue). The embedding is represented as a ribbon graph, while sewing a disc into each facial walk gives the surface.

More generally, when is it possible to specify an arbitrary circuit decomposition of the edges or arcs and complete it to an embedding with just one more face, necessarily bounded by an euler circuit? Finding such a face achieves a maximum genus orientable embedding having the circuits in a given decomposition as facial walks. This leads to the general question of determining the maximum genus of an embedding relative to a given circuit decomposition. Beyond topological graph theory, these questions arise in surprisingly diverse settings, including DNA self-assembly, Steiner triple systems, and latin squares.

Our main result, given in Theorem 4.1, answers these questions for sufficiently dense graphs and digraphs. We prove that if every vertex of an n -vertex digraph has at least $(4n+2)/5$ neighbors, and a necessary parity condition holds, then it is indeed always possible to achieve these special embeddings of maximum orientable genus. Our proof is constructive, so gives an algorithmic process for finding the desired embedding. Since the edges of an eulerian graph can always be directed to give an eulerian digraph, the stronger digraph results lead immediately to corollaries with analogous results for graphs.

Such consequences include Corollary 4.5, which says that if G is a sufficiently dense eulerian graph with euler circuit T , then there is a 2-face-colorable orientable embedding of G with T as the unique face of one color and with exactly one or two faces of the other

color, depending on whether $|V(G)| + |E(G)|$ is even or odd, respectively. This embedding has maximum genus among all 2-face-colorable orientable embeddings of G . When $|V(G)| + |E(G)|$ is even, then this is an orientable bi-eulerian embedding of G with T as a specified face, and the embedding has maximum genus among all orientable embeddings of G . However, in general not every maximum genus embedding of such a G is bi-eulerian.

Our original motivation for studying bi-eulerian embeddings came from an applied problem in DNA self-assembly posed by Jonoska, Seeman and Wu [22], which required a special closed walk in a graph for a DNA strand to follow. The presence of such walks corresponds to the existence of an orientable *edge-outer* embedding, namely an embedding with a special *outer* face whose boundary uses every edge at least once. The paper [22] proves the existence of orientable edge-outer embeddings, while [11] provides a short algorithm to find them and shows that finding an optimal embedding (with shortest outer face) is NP-hard. Neither [22] nor [11], however, provides any control over the number or sizes of the non-outer faces.

The startlingly simple (to state!) and intriguing questions in the first paragraph emerged from this application. Although determining the size of optimal edge-outer faces is hard in general, for eulerian graphs any optimal edge-outer face is necessarily bounded by an euler circuit. Thus, we seek to control the remaining faces in an optimal edge-outer embedding of an eulerian graph by specifying them in advance with a circuit decomposition. Of particular interest are bi-eulerian orientable embeddings, particularly when one of the circuits is specified in advance.

While our original motivation was DNA self-assembly, these and some closely related questions have also received considerable attention in various other special settings.

In [9], Edmonds proved that every eulerian graph has an bi-eulerian embedding in some surface, but noted that his main theorem was not sufficient to determine the orientability of the embedding. A result of Kotzig [23] implies that in a 4-regular graph every euler circuit, or more generally every circuit decomposition, can be extended to an embedding by adding an additional face bounded by an euler circuit; again the orientability is not determined. In [10, Section 6] we generalized both of these results to characterize when a circuit decomposition of an eulerian graph can be extended to a nonorientable embedding by adding an additional euler circuit face.

In this paper we focus on the more challenging orientable situation. Conditions for the existence of an orientable bi-eulerian embedding where one euler circuit face is specified in advance were given by Bonnington, Conder, Morton, and McKenna [3] for regular tournaments, and by the current authors [10], based on the degrees of vertices modulo 4. A series of papers, [15, 18, 19, 20] authored by Griggs and Širáň and sometimes also Erskine, Grannell, McCourt, or Psomas, discusses completing a triangular decomposition to an embedding by adding an euler circuit, for situations arising from a structures in design theory (Steiner triple systems, latin squares or symmetric 3-configurations).

More details of these, and other, prior results are given in Section 3, after reviewing necessary definitions and basic results in Section 2. We then state our main result and its various consequences and related results in Section 4, deferring the rather technical proof of the main result until Section 6. Section 5 contains the lemmas, some of independent interest, used for the proof of the main theorem. We conclude in Section 7 with some closing remarks.

An overview of some of the results in this paper has appeared in [12].

2. Terminology

We briefly recall some definitions and terminology from [10].

2.1. Representation of graphs and digraphs

Our graphs and digraphs allow multiple edge or arcs, and loops, and we often need to specify a particular end of an edge or arc. Therefore we define graphs and digraphs using half-edges and half-arcs, similarly to Fleischner [17].

A graph is a quadruple $G = (V, E^*, \psi, \omega)$ where V (vertices) and E^* (half-edges) are disjoint sets, $\psi : E^* \rightarrow V$ describes the incidence of each half-edge with a vertex, and $\omega : E^* \rightarrow E^*$ is a fixed-point-free involution that maps each half-edge to another half-edge. An edge is an unordered pair $\{h, \omega(h)\}$ where $h \in E^*$, and we let E denote the set of edges. We use $E^*(v)$ to represent $\psi^{-1}(v)$, the set of half-edges incident with a given vertex v . To specify a particular graph G we write $V(G)$, $E^*(G)$, ψ_G , ω_G , $E(G)$, and $E_G^*(v)$. The minimum degree of G is denoted $\delta(G)$.

Similarly, a digraph is a quintuple $D = (V, A^+, A^-, \psi, \omega)$ where V (vertices), A^+ (outgoing half-arcs) and A^- (incoming half-arcs) are disjoint sets, and if $A^* = A^+ \cup A^-$ then $\psi : A^* \rightarrow V$ describes incidences of half-arcs with vertices, and $\omega : A^* \rightarrow A^*$ is an involution that maps each element of A^+ to an element of A^- and vice versa. An arc of D is an ordered pair $(g, \omega(g))$ where $g \in A^+$, and we let A denote the set of arcs. We use $A^+(v)$, $A^-(v)$, and $A^*(v)$ to represent the set of elements of A^+ , A^- , and A^* , respectively, incident with a vertex v . To specify a particular digraph D we write $V(D)$, $A^+(D)$, $A^-(D)$, ψ_D , $A^*(D)$, ω_D , $A(D)$, $A_D^+(v)$, $A_D^-(v)$, and $A_D^*(v)$. Two distinct vertices u, v of D are *adjacent* if there is an arc (in either direction) with endvertices u and v .

We warn the reader that we will use either the framework just described or standard definitions as convenient. We follow the terminology of [28] unless otherwise noted.

A graph, walk, digraph, or directed walk is *nontrivial* if it has at least one edge or arc.

A *walk* of length ℓ in a graph G is a sequence $W = v_0g_1h_1v_1g_2h_2v_2 \dots v_{\ell-1}g_{\ell}h_{\ell}v_{\ell}$ where $v_i \in V(G)$, $g_i, h_i \in E^*(G)$, $\psi(g_i) = v_{i-1}$, $\psi(h_i) = v_i$ and $\omega(g_i) = h_i$ for all i with $1 \leq i \leq \ell$. A *directed walk* in a digraph D is similar to a walk, but with conditions $g_i \in A^+(D)$ and $h_i \in A^-(D)$ for all i with $1 \leq i \leq \ell$. Special types of walks include *closed walks* ($v_0 = v_{\ell}$), *paths* (no repeated vertices), *cycles* (closed walks that use every edge of a 2-regular subgraph exactly once), *trails* (no repeated edges), and *circuits* (closed trails). When it does not cause confusion, we do not distinguish between a closed (directed) walk and an equivalence class of closed (directed) walks under cyclic shifts, and we indicate such an equivalence class with parentheses, as $(v_0g_1h_1v_1 \dots v_{\ell-1}g_{\ell}h_{\ell})$. If $W = v_0g_1 \dots v_{\ell}$ and $W' = v'_0g'_1 \dots v'_{\ell'}$ are two (directed) walks with $v_{\ell} = v'_0$, then the (directed) walk $v_0g_1 \dots (v_{\ell} = v'_0)g'_1 \dots v'_{\ell'}$ obtained by concatenating W and W' is denoted by $W \cdot W'$.

An *euler circuit* in a graph is a circuit that contains all the edges and vertices. A graph with an euler circuit is *eulerian*, which is equivalent to being connected with every vertex of even degree. Directed euler circuits and eulerian digraphs are defined similarly; an eulerian

digraph is connected with every vertex having indegree equal to its outdegree. Every eulerian graph has at least one *circuit decomposition*, a collection of circuits using every edge exactly once. Similarly, every eulerian digraph has at least one *directed circuit decomposition*, using every arc exactly once.

2.2. Embeddings and directed embeddings

All graph embeddings in this paper are cellular. We assume that the reader is familiar with cellular embeddings of graphs in surfaces and their various combinatorial representations; standard references are [14, 21, 24]. An embedding can be represented by giving its collection of facial boundary walks, and we often do not distinguish between a face of an embedding and the corresponding facial walk. An *oriented* embedding is an embedding in an orientable surface with a specific global clockwise orientation. Our embeddings will mostly be oriented and represented combinatorially by rotation systems specifying the clockwise order of half-edges or half-arcs around each vertex, and we use standard tracing procedures to determine the faces. The *Euler genus* γ of a surface Σ is the genus if Σ is nonorientable and twice the genus if Σ is orientable, and Euler's formula for a connected graph G embedded in Σ is $|V| - |E| + |F| = 2 - \gamma$.

If we embed a digraph D so that every face is bounded by a directed closed walk of D (respecting the directions of the arcs in D), we call this a *directed embedding* of D . In an oriented directed embedding, each face is either a *proface*, with facial walk directed clockwise, or an *antiface*, with facial walk directed anticlockwise. Directed embeddings have the following basic properties.

Observation 2.1 (Bonnington, Conder, Morton and McKenna [3]).

- (a) An embedding of a digraph is a directed embedding if and only if at each vertex the half-arcs, taken in rotational order, alternate between entering and leaving the vertex.
- (b) Therefore, a connected digraph with a directed embedding must be eulerian, because at each vertex the numbers of entering and leaving half-arcs are equal.
- (c) An eulerian digraph always has at least one oriented directed embedding, given by any rotation system in which the half-arcs alternate in direction around each vertex.
- (d) An orientable directed embedding is always 2-face-colorable, where (for a specific orientation of the surface) the profaces and the antifaces form the two color classes.

Part (d) above can be strengthened, as follows.

Lemma 2.2 ([10, Lemma 2.2]). *A directed embedding is orientable if and only if it is 2-face-colorable.*

An embedding (directed embedding) of a graph (digraph) is *bi-eulerian* if it has two faces, each bounded by an euler circuit (directed euler circuit). Bi-eulerian embeddings of graphs may be orientable or nonorientable, but bi-eulerian directed embeddings of digraphs are always orientable, by Lemma 2.2.

The term *bi-eulerian* or *bieulerian* has occasionally been used in the literature, with a variety of meanings. For example, Xuong [29, p. 218] used it for embeddings with a single

face covering every edge twice, Fleischner [17, p. VI.17] used it for special eulerian digraphs, and Chen and Fang [4] used it for embeddings of bipartite eulerian graphs. We trust that our reuse of this term will cause no confusion.

2.3. Relative embeddings and upper relative embeddings

Suppose G is a connected graph and \mathcal{W} is a collection of closed walks in G . If \mathcal{W}' is another collection of closed walks in G such that \mathcal{W} and \mathcal{W}' together form the facial walks of an embedding Φ , then we say Φ is an *embedding of G relative to \mathcal{W}* . We say the faces bounded by elements of \mathcal{W} and of \mathcal{W}' are *specified* and *new* faces, respectively. If Φ is orientable, \mathcal{W} is a circuit decomposition \mathcal{C} of G , and there are only one or two new faces, then we say that Φ is an *upper embedding of G relative to \mathcal{C}* . (We require \mathcal{W} to be a circuit decomposition to avoid complications that arise in more general situations.) In particular, if we complete \mathcal{C} to an orientable embedding by adding an euler circuit, then we have an upper relative embedding. If an upper embedding of G relative to \mathcal{C} exists, it has maximum genus over all orientable embeddings of G relative to \mathcal{C} .

We can also apply these concepts to collections of directed walks, directed circuit decompositions, and directed embeddings of digraphs. The following observation relies on Lemma 2.2 for (a) and on evenness of the Euler genus of an orientable embedding for (b). Part (b) tells us whether an upper relative embedding will have just one new face or two new faces.

Observation 2.3. Suppose \mathcal{C} is a directed circuit decomposition of an eulerian digraph D and Φ is a directed embedding of D relative to \mathcal{C} .

- (a) Then Φ is orientable, and the surface can be oriented so that \mathcal{C} is the collection of profaces.
- (b) The number of new faces of Φ has the same parity as $|V(D)| + |A(D)| + |\mathcal{C}|$.

We will use the following result on the existence of oriented directed embeddings relative to a directed circuit decomposition.

Lemma 2.4. *Suppose \mathcal{C} is a directed circuit decomposition of an eulerian digraph D . Then there exists an oriented directed embedding of D in which the profaces are precisely the elements of \mathcal{C} .*

Proof. Since \mathcal{C} is a directed circuit decomposition of D , every half-arc occurs exactly once in \mathcal{C} . For each $h \in A_D^-$, let $\gamma(h) \in A_D^+$ be the half-arc that follows h in some element of \mathcal{C} . Then $\gamma(A^-(v)) = A^+(v)$ for all $v \in V(D)$.

We construct a rotation at each vertex v . Suppose v has $\text{indeg}(v) = \text{outdeg}(v) = d$, and $A^-(v) = \{h_0, h_1, \dots, h_{d-1}\}$. Let $g_i = \gamma(h_i)$ for $0 \leq i \leq d-1$, and define the clockwise rotation at v to be $(g_0 h_0 g_1 h_1 \dots g_{d-1} h_{d-1})$. Since the half-arcs alternate in direction at each v , these rotations specify an orientable directed embedding of D . Each proface that enters v on h_i leaves on $g_i = \gamma(h_i)$, so the profaces join the arcs together in the same way that the circuits in \mathcal{C} do. Thus, the profaces are exactly the elements of \mathcal{C} . \square

This lemma was used implicitly by Bonnington et al. [3, p. 13] in the case where \mathcal{C} consists of a directed euler circuit.

3. Previous results

Graph embeddings with euler circuit faces have been a subject of interest for nearly sixty years. In this section we discuss relevant prior results. We focus on orientable embeddings relative to a circuit decomposition, but briefly discuss more general results first.

Edmonds [9, p. 123] showed that every eulerian graph has a bi-eulerian embedding. A result of Kotzig [23, Theorem 3] can be interpreted as saying that every circuit decomposition of a 4-regular graph can be extended to an embedding by adding an euler circuit. In both Edmonds' and Kotzig's results nothing is specified about the orientability of the resulting embedding. We showed [10, Theorem 6.5] that a circuit decomposition of an arbitrary eulerian graph can be extended to an embedding by adding an euler circuit, and the embedding can be guaranteed to be nonorientable unless every block of the graph is a cycle and the circuit decomposition is the collection of cycles. This extends some special cases in the literature, such as the nonorientable parts of Theorems 3.2 and 3.3 below. In particular, every eulerian graph that is not a cycle has a nonorientable bi-eulerian embedding where one of the euler circuit faces can be specified in advance.

We turn now to orientable embeddings. If we have a circuit decomposition \mathcal{C} of an eulerian graph, then we can find an upper embedding relative to \mathcal{C} if we can complete \mathcal{C} to an orientable embedding using just one face (which is necessarily an euler circuit) or two faces. Similarly, if we have a directed circuit decomposition \mathcal{C} of an eulerian digraph, then we can find an upper directed embedding relative to \mathcal{C} if we can complete \mathcal{C} to a directed embedding using just one face (which must be a directed euler circuit) or two faces. By Observation 2.3 the embedding is necessarily orientable, and we can assume that \mathcal{C} forms the profaces and the added face or faces form the antifaces.

General formulas for relative oriented maximum genus were given by Bonnington [2] and Archdeacon, Bonnington, and Širáň [1], which generalize maximum orientable genus formulas due to Xuong [29] and Nebeský [26], respectively. For a relative oriented embedding of a graph, each specified face has a specified direction that should be clockwise in the surface. If we have a directed circuit decomposition \mathcal{C} of an eulerian digraph D , we can consider this as a collection of oriented circuits in the underlying graph G , and the formulas from [1, 2] can be applied to find the maximum genus of a (necessarily orientable) directed embedding relative to \mathcal{C} . However, it seems to be difficult to use these formulas to derive easily applicable conditions for the existence of a relative upper embedding (although we do know of one situation where we can do this, as we discuss at the end of this section).

The first result that we are aware of giving a specific condition for finding a relative upper embedding of some type of circuit decomposition is the following.

Theorem 3.1 (Bonnington, Conder, Morton, and McKenna, [3]). *Let T be a directed euler circuit in a regular (equivalently, eulerian) tournament D . Then there is an orientable directed embedding of D with T as the only proface and with at most two antifaces, i.e., an upper directed embedding of D relative to $\{T\}$.*

There is a series of papers [15, 18, 19, 20] involving Griggs and Širáň as authors, and also Erskine, Grannell, McCourt, and Psomas, which discuss upper embeddings relative to

a triangular decomposition of a graph or digraph, and more specifically completing such a decomposition to an embedding by adding an euler circuit. They are interested in triangular decompositions that arise from structures in design theory such as Steiner triple systems, symmetric 3-configurations, and latin squares. Their first results were for undirected situations.

Theorem 3.2 (Granell, Griggs and Širáň [18]). *Let \mathcal{C} be a Steiner triple system of order $n \geq 7$, i.e., a decomposition of the complete graph $G = K_n$ into triangles. Then there are both orientable and nonorientable embeddings of G with the elements of \mathcal{C} as faces and exactly one additional euler circuit face.*

Theorem 3.3 (Griggs, Psomas, and Širáň [20]). *Let \mathcal{C} be a latin square of order p , i.e., a decomposition of the balanced complete tripartite graph $G = K_{p,p,p}$ into triangles. If $p \geq 2$ then there is a nonorientable embedding of G with the elements of \mathcal{C} as faces and with exactly one additional euler circuit face. If p is odd then there is an orientable embedding of G with the elements of \mathcal{C} as faces and with exactly one additional euler circuit face.*

Griggs, McCourt and Širáň strengthened Theorems 3.2 and 3.3 by adding specified directions to the triangle decompositions. This effectively changes the setting from undirected graphs to digraphs, and their results can be stated as follows.

Theorem 3.4 (Griggs, McCourt, and Širáň [19, Theorem 1.1]). *Let \mathcal{C} be an oriented Steiner triple system, i.e., a decomposition of a regular tournament D into directed triangles. Then there is an orientable directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly one directed euler circuit antiface.*

Theorem 3.5 (Griggs, McCourt, and Širáň [19, Theorem 1.2]). *Let \mathcal{C} be an oriented latin square of odd order, i.e., a decomposition of an eulerian orientation D of a balanced complete tripartite graph $K_{p,p,p}$ into directed triangles. Then there is an orientable directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly one directed euler circuit antiface.*

Erskine, Griggs, and Širáň also investigated the situation for graphs and digraphs obtained from symmetric 3-configurations, obtaining a positive result for small cases but a negative result in general.

Theorem 3.6 (Erskine, Griggs, and Širáň [15, Theorems 2.1 and 3.2]). *Suppose G is the associated graph of a symmetric configuration n_3 , i.e., G is a 6-regular n -vertex simple graph with a decomposition \mathcal{T} into triangles. Let \mathcal{C} be a collection of directed triangles obtained by directing the elements of \mathcal{T} , and let D be the digraph obtained from G by directing the edges as specified by \mathcal{C} .*

If n is odd and $7 \leq n \leq 19$, then every for every such D and \mathcal{C} there is an orientable directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly one directed euler circuit antiface.

If n is odd and $n \geq 21$ then there exists such a graph $G = G_0$ with decomposition $\mathcal{T} = \mathcal{T}_0$ such that there is no orientable embedding of G_0 with the elements of \mathcal{T}_0 as faces and exactly

one additional face. Thus, for every D and \mathcal{C} derived as above from G_0 and \mathcal{T}_0 there is no orientable directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly one antiface.

The proofs of Theorems 3.2 through 3.6 use arguments that are specific to triangle decompositions.

The current authors investigated maximum genus of embeddings relative to an euler circuit, and of directed embeddings relative to a directed euler circuit, in [10]. We showed [10] that in every eulerian graph or digraph where every vertex has degree $2 \bmod 4$, every (directed) euler circuit can be extended to an orientable (directed) embedding by adding a second (directed) euler circuit, giving a relative upper embedding. One way to derive this result is by applying Bonnington's formula for maximum genus of a relative oriented embedding [2], although that is not the approach used in [10]. Theorem 3.6 above from [15] shows that our result unfortunately does not generalize to extending an arbitrary circuit decomposition; see [10, Section 5] for further discussion.

The main results of this paper (Theorem 4.1 and its corollaries) encompass a number of the results mentioned above, specifically Theorem 3.1, the orientable part of Theorem 3.2, and Theorem 3.4. Our main results are not, however, strong enough to imply the orientable part of Theorem 3.3, or Theorem 3.5, because balanced complete tripartite graphs do not satisfy our density condition.

4. Statement of results

In this section we state our results for ‘dense’ eulerian graphs and digraphs. We also provide results for all eulerian graphs and digraphs with only one or two vertices (but arbitrarily many edges or arcs).

4.1. The main result and some consequences

In this subsection we consider ‘dense’ eulerian digraphs, where ‘dense’ means that each vertex is adjacent to at least roughly $4/5$ of the other vertices. We show that given a collection \mathcal{C} of edge-disjoint directed circuits in a dense eulerian digraph, we can find a directed orientable embedding in which all elements of \mathcal{C} are profaces, and there are at most two antifaces. This means that we can determine the maximum genus of an orientable directed embedding in which the elements of \mathcal{C} are required to be profaces. As a special case, and subject to a necessary parity condition we can find an orientable bi-eulerian directed embedding where one of the euler circuit faces may be specified in advance.

Our main result, and many of our arguments, rely only on information about whether two vertices are adjacent in a digraph, and not how many arcs are between them or directions of the arcs. Thus, it is often convenient to work with the undirected simple graph underlying a digraph D , which we denote by D_{us} . The graph D_{us} has the same vertex set as D , and distinct vertices u and v are adjacent in D_{us} if and only if they are adjacent in D . Moreover, if W is a directed walk (such as a facial walk of a directed embedding) in D , then W_{us} is the corresponding walk in D_{us} : we take the sequence of vertices in W , eliminating consecutive

repetitions of a vertex (created by following loops). We also treat W_{us} as a subgraph of D_{us} ; whether we are thinking of it as a walk or a subgraph should be clear from context.

The main result of this paper is the following.

Theorem 4.1. *Let D be an n -vertex eulerian digraph where $\delta(D_{\text{us}})$, the minimum degree of the underlying simple undirected graph of D , is at least $(4n+2)/5$. Let \mathcal{C} be a directed circuit decomposition of D . Then there is a directed embedding of D in an oriented surface (i.e., an orientable surface with a specific orientation) with the elements of \mathcal{C} as the profaces and with exactly one or two antifaces, depending on whether $|V(D)| + |A(D)| + |\mathcal{C}|$ is odd or even, respectively.*

This embedding has maximum genus among all orientable directed embeddings of D in which all elements of \mathcal{C} are faces.

As mentioned in Section 3, Theorem 4.1 generalizes Theorems 3.1 and 3.4, but not Theorem 3.5, because $K_{p,p,p}$ is not dense enough. Because the proof of Theorem 4.1 is rather long and technical, we defer it to Sections 5 and 6, in order to state here some immediate corollaries that illustrate its impact and application.

If \mathcal{C} consists of a single directed euler circuit, then we have the following.

Corollary 4.2. *Let D be an n -vertex eulerian digraph where $\delta(D_{\text{us}}) \geq (4n+2)/5$. Let T be a directed euler circuit in D . Then there is an oriented directed embedding of D with T as the only proface, and with exactly one or two antifaces, depending on whether $|V(D)| + |A(D)|$ is even or odd, respectively.*

This embedding has maximum genus among all orientable directed embeddings of D . When $|V(D)| + |A(D)|$ is even this embedding is an orientable bi-eulerian directed embedding of D with T as a specified face.

When $|V(D)| + |A(D)|$ is even in Corollary 4.2, we can also apply Observation 2.1(d) to conclude that *every* maximum genus orientable directed embedding of D is bi-eulerian.

By extending a set \mathcal{C} of arc-disjoint directed circuits to a circuit decomposition using an euler circuit in each component of the subgraph induced by the unused arcs, we also obtain the following corollary.

Corollary 4.3. *Let D be an n -vertex eulerian digraph where $\delta(D_{\text{us}}) \geq (4n+2)/5$. Let \mathcal{C} be a set of arc-disjoint directed circuits in D , let $A(\mathcal{C})$ be the set of arcs used by elements of \mathcal{C} , and let $\alpha(\mathcal{C})$ be the number of components of $D - A(\mathcal{C})$ with at least one edge. Then there is an oriented directed embedding of D with $|\mathcal{C}| + \alpha(\mathcal{C})$ profaces and with exactly one or two antifaces, depending on whether $|V(D)| + |A(D)| + |\mathcal{C}| + \alpha(\mathcal{C})$ is odd or even, respectively.*

This embedding has maximum genus among all orientable directed embeddings of D in which all elements of \mathcal{C} are faces.

We also have corollaries for undirected graphs. Let G_s denote the underlying simple graph of a graph G . We state the undirected counterparts of Theorem 4.1 and Corollary 4.2. In each case the first paragraph is obtained by applying the digraph result to an eulerian orientation of the graph G that makes each element of \mathcal{C} into a directed circuit, and the

conclusions about maximum genus in the second paragraph are easily verified. There is also an undirected version of Corollary 4.3, whose statement we leave to the reader.

Corollary 4.4. *Let G be an n -vertex eulerian graph where $\delta(G_s)$, the minimum degree of the underlying simple graph of G , is at least $(4n + 2)/5$. Let \mathcal{C} be a circuit decomposition of G . Then there is a 2-face-colorable orientable embedding of G with the elements of \mathcal{C} as one of the color classes and with exactly one or two faces of the other color, depending on whether $|V(G)| + |E(G)| + |\mathcal{C}|$ is odd or even, respectively.*

This embedding has maximum genus among all orientable embeddings of G in which all elements of \mathcal{C} are faces.

Corollary 4.5. *Let G be an n -vertex eulerian graph where $\delta(G_s) \geq (4n + 2)/5$. Let T be an euler circuit in G . Then there is a 2-face-colorable orientable embedding of G with T as the unique face of one color and with exactly one or two faces of the other color, depending on whether $|V(G)| + |E(G)|$ is even or odd, respectively.*

This embedding has maximum genus among all 2-face-colorable orientable embeddings of G . When $|V(G)| + |E(G)|$ is even this embedding is an orientable bi-eulerian embedding of G with T as a specified face, and the embedding has maximum genus among all orientable embeddings of G .

When $|V(D)| + |A(D)|$ is even in Corollary 4.2, all maximum genus directed embeddings are bi-eulerian. However, when $|V(G)| + |E(G)|$ is even in Corollary 4.5, the embeddings provided by the corollary are bi-eulerian maximum genus embeddings, but there may be other maximum genus embeddings that are not bi-eulerian, as the following example shows.

Example 4.6. Suppose we have an eulerian graph G satisfying our density condition, with a decomposition into three circuits C_1, C_2, C_3 , all of which contain the same vertex v . (Such situations are not hard to construct.) Apply Corollary 4.4 to complete $\{C_1, C_2, C_3\}$ to an embedding Φ using an euler circuit T . By applying an undirected version of part (a) of the Three Face Lemma (Lemma 5.1 below) to Φ at v we can merge C_1, C_2 , and C_3 into a single face, giving a bi-eulerian embedding, or alternatively merge C_1, C_2 and T into a single face, giving a two-face orientable embedding that is not bi-eulerian.

Our results can be combined with known results on the existence of triangular decompositions in dense graphs, to give maximum genus embeddings using these decompositions. We say a graph is *triangle-divisible* if its number of edges is divisible by 3 and every vertex has even degree. In [25], Nash-Williams conjectured that there is some constant c such that for sufficiently large n , every n -vertex triangle-divisible simple graph G with $\delta(G) \geq cn$ has a decomposition into edge-disjoint triangles, and Graham pointed out that c must be at least $3/4$. The current best result on Nash-Williams' conjecture is as follows.

Theorem 4.7 (Delcourt and Postle [5]). *For every $\varepsilon > 0$ there is $N(\varepsilon)$ such that every n -vertex triangle-divisible simple graph G with $n \geq N(\varepsilon)$ and $\delta(G) \geq ((7 + \sqrt{21})/14 + \varepsilon)n$ has a decomposition into edge-disjoint triangles.*

The value of $(7 + \sqrt{21})/14$ is about $0.827327 > 4/5$, and triangle-divisible simple graphs G with $\delta(G) \geq n/2$ are eulerian, so applying Corollary 4.4 gives the following.

Corollary 4.8. *For every $\varepsilon > 0$ there is $N'(\varepsilon)$ with the following properties. Every n -vertex triangle-divisible simple graph G with $n \geq N'(\varepsilon)$ and $\delta(G) \geq ((7 + \sqrt{21})/14 + \varepsilon)n$ has a triangular decomposition \mathcal{C} . For every such decomposition \mathcal{C} there is a 2-face-colorable orientable embedding where all faces of one color are bounded by the elements of \mathcal{C} and there are exactly one or two faces of the other color, depending on whether $|V(G)|$ is odd or even, respectively. This embedding has maximum genus among all orientable embeddings in which all elements of \mathcal{C} are faces.*

This may be regarded as a strengthening of the orientable part of Theorem 3.2, since a complete graph has a triangular decomposition, corresponding to a Steiner triple system, if and only if it is triangle-divisible.

Corollary 4.8 also tells us that a sufficiently large sufficiently dense triangle-divisible simple graph G always has an edge-outer embedding that is optimal in three respects: it has an outer face of minimum possible length (the euler circuit face), the maximum possible number of faces other than the outer face (the triangular faces), and hence it attains the minimum genus numerically possible for an edge-outer embedding of a simple graph with $|V(G)|$ vertices and $|E(G)|$ edges.

Future strengthenings of Theorem 4.7 may also provide improvements to Corollary 4.8.

4.2. Graphs and digraphs with one or two vertices

Since D_{us} is simple, the degree condition in Theorem 4.1 gives $n-1 \geq \delta(D_{\text{us}}) \geq (4n+2)/5$, which requires that $n \geq 7$. However, our techniques can also give results covering all eulerian graphs and digraphs with just one or two vertices (but arbitrarily many edges or arcs).

An *edge-cut* S in a digraph D is the set of all arcs with one end U and other end in $\overline{U} = V(D) - U$ for some proper nonempty subset U of $V(D)$, and a *k-edge-cut* is an edge-cut of cardinality k . In an eulerian digraph D every edge-cut has even cardinality. Recall that for an integer a , $a \bmod 2$ is equal to 0 if a is even and 1 if a is odd.

Proposition 4.9. *Let D be an eulerian digraph with $|V(D)| \leq 2$ and let \mathcal{C} be a directed circuit decomposition of D .*

- (a) *If $|V(D)| = 1$, or if $|V(D)| = 2$ and D has no 2-edge-cut, then there is an oriented directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly one or two antifaces, depending on whether $|V(D)| + |A(D)| + |\mathcal{C}|$ is odd or even, respectively.*
- (b) *Suppose that $|V(D)| = 2$ and D has a 2-edge-cut. Let $V(D) = \{v_1, v_2\}$, for each $i \in \{1, 2\}$ let α_i be the number of directed loops incident with v_i in D , and let γ_i be the number of elements of \mathcal{C} that are incident only with v_i , not with v_{3-i} . Then there is an oriented directed embedding of D with the elements of \mathcal{C} as the profaces and with exactly $((\alpha_1 + \gamma_1) \bmod 2) + ((\alpha_2 + \gamma_2) \bmod 2) + 1$ antifaces.*

In both cases the embedding described has maximum genus among all orientable directed embeddings of D in which all elements of \mathcal{C} are faces.

Proposition 4.9 is proved below, after Lemma 5.2. Immediate corollaries of Proposition 4.9 can also be stated for bi-eulerian embeddings, partial circuit decompositions (collections of arc-disjoint circuits), and undirected graphs. We leave the formulation of these to the reader.

5. Supporting lemmas

The proof of our main theorem, Theorem 4.1, depends on the lemmas of this section. These lemmas facilitate combining antifaces in an oriented directed embedding without changing the profaces. While we developed these lemmas to prove Theorem 4.1, they are also of generic utility for manipulating embedded (di)graphs.

Our proof of Theorem 4.1 was inspired by the approach used to prove Theorem 3.1 in Bonnington et al. [3, Section 3]. However, since we treat general dense graphs instead of the special case of tournaments, we necessarily develop a number of new ideas. Our Three Face Lemma (5.1) and Interlaced Faces Lemma (5.2) simplify and generalize some arguments from [3, Lemma 3.1 and proof of Lemma 3.4]. Our Three Neighbor Lemma (5.3) generalizes [3, Lemma 3.3] and our Diamond Lemma (5.4) summarizes and generalizes some reasoning from [3, proof of Theorem 3.1].

Since we address digraphs here that are not orientations of complete graphs or even simple graphs, we need new sufficient conditions for when the results above can be applied. We provide these in the Three Neighbor Corollary (5.6), Big and Moderate Faces Corollary (5.7), and Diamond Corollary (5.9). We also introduce a number of new ideas in the Bipartite Degeneracy Lemma (5.8), Division Lemma (5.10), Blow Up Lemma (5.11), and how we use touch graphs in Section 6 for the proof of the main result. Our more general setting also requires a significantly more complicated analysis. Moreover, since our approach only manipulates antifaces, and does not depend on the profaces, it can be applied to situations where the profaces form an arbitrary directed circuit decomposition, not just a directed euler circuit.

5.1. General lemmas for merging antifaces

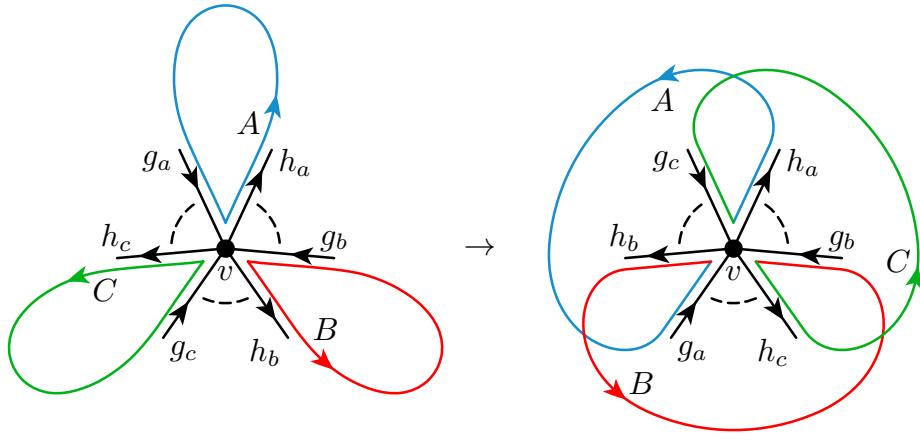
We begin with some general results. The Three Face Lemma, Lemma 5.1 below, is our main tool for merging or rearranging faces. We follow it with a useful consequence, Lemma 5.2.

The Three Face Lemma (5.1) describes what happens if we modify the rotation at a vertex v by breaking it in three places and reassembling it. The places where we break it correspond to faces A , B , and C . The lemma covers the cases where the three faces are distinct, or where B is distinct from $A = C$. (We can also investigate what happens when $A = B = C$, although we do not need this for our results. There are two situations: one is the inverse of part (a), where one antiface splits into three antifaces, and in the other situation the antiface has its boundary rearranged but remains a single antiface.)

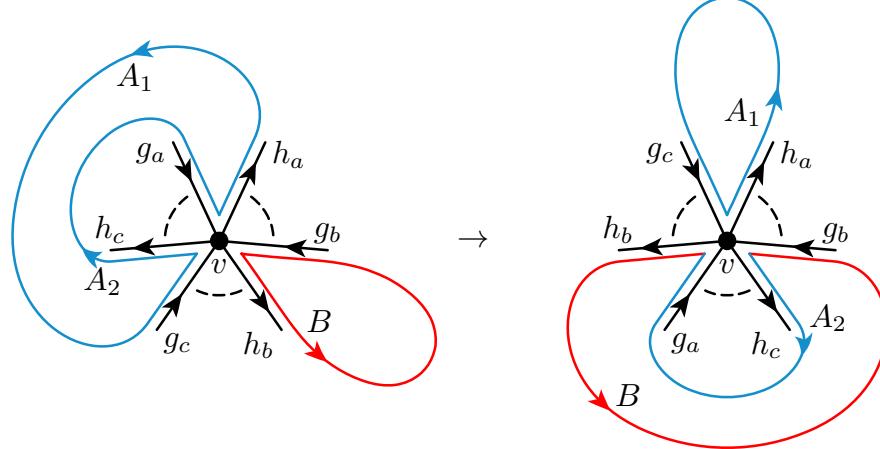
Lemma 5.1 (Three Face Lemma). *Let Φ be an oriented directed embedding of an eulerian digraph D , and $v \in V(D)$.*

(a) Suppose A , B , and C are distinct antifaces that each contain v . Then there is an oriented directed embedding Φ' of D that has the same profaces and antifaces as Φ except that A , B and C are merged into a single antiface.

(b) Suppose A and B are distinct antifaces that each contain v . Suppose further that A can be written as $A = A_1 \cdot A_2$ where A_1 and A_2 are nontrivial vv -walks. Also consider B as a vv -walk. Then there is an oriented directed embedding Φ' of D that has the same profaces and antifaces as Φ , except that A and B are replaced by two new antifaces $A' = A_1 \cdot B$ and $B' = A_2 \cdot B$ for some $i \in \{1, 2\}$.



(a) Three distinct initial antifaces.



(b) Two initial antifaces.

Figure 2: Cases for the Three Face Lemma.

Proof. Suppose the clockwise rotation at v in Φ is $(g_0h_0g_1h_1 \dots g_{d-1}h_{d-1})$ where $A^-(v) = \{g_i \mid i \in \mathbb{Z}_d\}$ and $A^+(v) = \{h_i \mid i \in \mathbb{Z}_d\}$. Given a, b, c in (increasing) cyclic order in

$\mathbb{Z}_d = \{0, 1, \dots, d-1\}$, let Φ' be the oriented directed embedding obtained from Φ by changing the rotation at v from

$$(h_a g_{a+1} h_{a+1} \dots g_b \ h_b g_{b+1} h_{b+1} \dots g_c \ h_c g_{c+1} h_{c+1} \dots g_a)$$

to

$$(h_a g_{a+1} h_{a+1} \dots g_b \ h_c g_{c+1} h_{c+1} \dots g_a \ h_b g_{b+1} h_{b+1} \dots g_c).$$

(a) We may assume that our three distinct faces are represented as vv -walks with the first outgoing half-arc at v of A denoted by h_a , of B by h_b , and of C by h_c . Thus, $A = vh_a \dots g_a v$, $B = vh_b \dots g_b v$, and $C = vh_c \dots g_c v$. Construct Φ' as above. Faces other than A, B, C are unchanged, and by tracing the antiface starting $vh_a \dots$ we see that A, B, C are merged into a single antiface $A \cdot B \cdot C$. See Figure 2(a).

(b) We may assume (by swapping A_1 and A_2 if necessary) that $A_1 = vh_a \dots g_c v$, $A_2 = vh_c \dots g_a v$, and $B = vh_b \dots g_b v$. Construct Φ' as above. Faces other than A and B are unchanged, and by tracing faces we see that A and B are replaced by new antifaces $vh_c \dots g_a v h_b \dots g_b v = A_2 \cdot B$ and $vh_a \dots g_c v = A_1$. See Figure 2(b). \square

In part (b) of the statement of Lemma 5.1, A_i is whichever of A_1 or A_2 the proof relabels as A_2 , which is the one that has its initial and final half-arcs further from g_b and h_b in the rotation around v . In general we cannot choose the value of i .

There is an undirected version of Lemma 5.1(a), which allows us to combine three distinct faces incident with the same vertex into a single face by adding a handle to an orientable embedding of a graph. Ellingham and Weaver [13] explain how this idea can be regarded as a special case of an operation that merges two pairs of faces incident with a given vertex. Similar ideas and special cases involving local modification of an embedding around a single vertex or single face have been used for a long time, for example by Ringel [27, p. 120] in one of the papers proving the Map Color Theorem, by Duke [8, Theorem 3.2] in showing that the orientable genus range of a graph is an interval, and by Xuong [29, Figure 4] in finding maximum genus embeddings.

The operation in Lemma 5.1 does not change the profaces of the embedding at all. In our arguments the profaces will be specified in advance and never modified. However, if desired, by switching the roles of profaces and antifaces, we could also apply Lemma 5.1 to merge or rearrange profaces without changing antifaces.

We say distinct vertices x, y are *interlaced* on a closed walk (or face) W if (possibly after cyclic shifting) W can be written as $x \dots y \dots x \dots y \dots$. Many of the following lemmas give conditions that assure the existence of, or allow us to create, at least one pair of interlaced vertices. We eventually use interlaced pairs of vertices to merge antifaces, thus helping us reach the ultimate goal of only one or two antifaces.

Lemma 5.2 (Interlaced Faces Lemma). *Let Φ be an oriented directed embedding of an eulerian digraph D , and $x, y \in V(D)$. Suppose that there are three distinct antifaces A, B, C such that x and y are interlaced on A , while x occurs on B and y occurs on C . Then there is an oriented directed embedding Φ' of D that has the same profaces and antifaces as Φ except that A, B and C are replaced by a single antiface.*

Proof. Write $A = A_1 \cdot A_2$ where A_1 and A_2 are nontrivial walks of the form $x \dots y \dots x$. By Lemma 5.1(b) we may modify Φ so as to replace A and B by two antifaces $A_i \cdot B$ and A_{3-i} . Now $A_i \cdot B$, A_{3-i} , and C are three distinct antifaces at y , so by Lemma 5.1(a) we may combine them into a single antiface. \square

Now we consider the situation of Theorem 4.1 or Proposition 4.9. We suppose D is an n -vertex eulerian digraph, and \mathcal{C} is a directed circuit decomposition of D . By Lemma 2.4 there exists an oriented directed embedding Φ of D in which the profaces are exactly the elements of \mathcal{C} . Let \mathcal{A} be the set of antifaces of Φ . Our overall goal now is to show that we can decrease $|\mathcal{A}|$ to 1 or 2 without changing the profaces. If there is a vertex that belongs to three antifaces, we will apply Lemma 5.1(a) to reduce the number of antifaces. Otherwise, each vertex belongs to at most two antifaces, and we say the embedding is *locally irreducible*. In our arguments the profaces of the embedding never change.

At this point we have developed the tools necessary to prove Proposition 4.9, which covers graphs with one or two vertices.

Proof of Proposition 4.9. (a) By Lemma 2.4 there is an oriented directed embedding Φ_0 of D with \mathcal{C} as the set of profaces. By repeatedly applying the Three Face Lemma (5.1) beginning with Φ_0 , we can obtain a locally irreducible embedding Φ_1 with \mathcal{C} as the set of profaces. We show that we can obtain Φ_2 with at most two antifaces and with \mathcal{C} as the set of profaces.

If Φ_1 has at most two antifaces, which holds if $|V(D)| = 1$, then we can take $\Phi_2 = \Phi_1$. So we may suppose that $V(D) = \{v_1, v_2\}$ and Φ_1 has at least three antifaces. Some antiface A must appear at both v_1 and v_2 . Let B and C be two other antifaces. If there is a fourth antiface, or if either of B or C is incident with both v_1 and v_2 , or if B and C are incident with the same vertex v_i , then Φ_1 is not locally irreducible. So there are only three antifaces A, B, C , and we may assume without loss of generality that B is only incident with v_1 and C is only incident with v_2 . Thus, A uses all arcs between v_1 and v_2 , and since there are at least four such arcs, v_1 and v_2 are interlaced on A . We can therefore apply the Interlaced Faces Lemma (5.2) to obtain Φ_2 with exactly one antiface.

The precise number of antifaces in Φ_2 follows from Observation 2.3(b). The maximum genus conclusion follows because the number of antifaces is minimum.

(b) For this case we just provide an outline of the proof, as the details are straightforward but tedious to give in full. A 2-vertex eulerian digraph D with a 2-edge-cut can be decomposed into two 1-vertex eulerian digraphs D_1 and D_2 using a *2-edge-cut reduction*. There is a bijection between oriented directed embeddings Φ of D and ordered pairs (Φ_1, Φ_2) of oriented directed embeddings of D_1 and D_2 . (This generalizes [10, Observation 3.3].) Applying this idea and part (a) gives the result. \square

For Theorem 4.1 we try to create situations in locally irreducible embeddings where we have interlacing, so that we can use the Interlaced Faces Lemma (5.2). We wish to concisely describe sets of vertices that belong to specific antifaces. For distinct antifaces A and B we define $AB = V(A) \cap V(B)$ and $A\bar{B} = V(A) - V(B)$. For $i \in \{0, 1\}$ we let Ai denote the set of vertices that belong to A and exactly i other antifaces. If $u \in AB$ or $u \in A0$ we also say u is of *type* AB or $A0$, respectively. If $u \in A\bar{B}$ or $u \in A1$ we say u is of *general*

type $A\overline{B}$ or $A1$, respectively. If the embedding is locally irreducible, then $V(A) = A0 \cup A1$, $A1 = \bigcup_{P \in \mathcal{A}, P \neq A} AP$, and each vertex u has a unique type which specifies exactly which antifaces u belongs to. In that case, for $u \in AB$ and $v \in AC$ where A, B, C are distinct antifaces, every arc of D between u and v must belong to A .

We have two results that allow us to find interlacing in locally irreducible embeddings. For both proofs, note that two vertices are interlaced on W if and only if they are interlaced on the walk W_{us} .

Lemma 5.3 (Three Neighbor Lemma). *Assume we have a locally irreducible embedding with antiface A . Suppose there is nonempty $S \subseteq A1$ such that each $v \in S$ is adjacent to at least three other vertices of S of types different from v . Then there are vertices $x, y \in S$ of different types that are interlaced on A .*

Proof. The hypothesis means that $|S| \geq 4$. Two adjacent vertices of S of different types must be joined by an arc of A , so the hypothesis also means that each vertex of S is incident with at least three edges of A_{us} (the walk in D_{us} corresponding to A), and so occurs at least twice on A_{us} .

Therefore, we may choose an interval (sequence of consecutive vertices) $I = v_0 v_1 \dots v_p$ along A_{us} such that v_0 and v_p are occurrences of some vertex $x \in S$, some v_i with $1 \leq i \leq p-1$ is a vertex $y \in S$ of type different from x , and p is as small as possible over all choices of I , x , and y .

If y appears on A_{us} outside of I , then x and y are interlaced on A_{us} and hence on A as desired. So suppose that all occurrences of y on A_{us} are in the interval I . Let v_m and v_n be the first and last occurrences of y , so that $1 \leq m < n \leq p-1$. No vertex v_i with $m+1 \leq i \leq n-1$ can be a vertex of S with type different from y , or we could replace x by y and reduce p . But then y is adjacent to at most two vertices in S of type different from itself, namely v_{m-1} and v_{n+1} , which contradicts our hypothesis. \square

Lemma 5.4 (Diamond Lemma). *Assume we have a locally irreducible embedding with distinct antifaces A, B, C . Suppose that we have three distinct type AC vertices t, u, v and one type AB vertex x , such that $tu, uv \in E(A_{us})$ and $xt, xu, xv \in E(D_{us})$. Then x is interlaced on A with at least one of t, u or v .*

Proof. Since x is of type AB and t, u , and v are of type AC with $B \neq C$, the arcs of D joining x to t, u , and v must belong to A , so $xt, xu, xv \in E(A_{us})$. Therefore, possibly after reversing A_{us} , we may choose an interval $I = v_0 v_1 \dots v_p$ along A_{us} , where $v_0 = v_p = x$, $v_1 = u$, and no v_i , $1 \leq i \leq p-1$, is equal to x . At most one of xt or xv belongs to I , so we may assume without loss of generality that xt , and hence t , occurs on A_{us} outside of I . If tu is an edge of I then t and x are interlaced on A_{us} and hence on A , and if tu is not an edge of I then u and x are interlaced on A_{us} and hence on A . \square

5.2. Lemmas using density

Now we give some consequences of the above results that depend on various degree conditions. To state these results we define k to be the maximum degree of the complement

of D_{us} , or equivalently $k = n - 1 - \delta(D_{\text{us}})$. In other words, for each $v \in V(D)$ there are at most k vertices different from v that are not adjacent to v in D . The parameter k will have this meaning for the rest of this subsection and in Section 6. Thus, if S is a set of at least $k + 1$ vertices, any vertex v not in S is adjacent to, and hence shares an antiface with each of, at least $|S| - k \geq 1$ elements of S . Moreover, the induced subgraph $D_{\text{us}}[S]$ has minimum degree at least $|S| - k - 1$.

We will use the following observation frequently, often without explicit reference.

Observation 5.5. If $A \in \mathcal{A}$ has $A0 \neq \emptyset$ then $|V(A)| \geq n - k$.

Proof. If $v \in A0$ then v and all its neighbors are in A . Thus, at most k vertices are not in A , i.e., $|V(A)| \geq n - k$. \square

The next four results give conditions under which we can employ the Three Neighbor Lemma (5.3) or the Diamond Lemma (5.4) to find interlacing.

Corollary 5.6 (Three Neighbor Corollary). *Suppose we have a locally irreducible embedding with at least three antifaces. Suppose that for some antiface A we have $|A1| - |AP| \geq k + 3$ for all antifaces $P \neq A$. Then there are vertices x, y both of general type $A1$ but of different types that are interlaced on A .*

Proof. Let $S = A1$. Consider an arbitrary $v \in S$, of type AP . Since $|A1| - |AP| \geq k + 3$ there are at least $k + 3$ vertices of type different from v in S . Since v is nonadjacent to at most k of these, v is adjacent to at least three of them. The result then follows from the Three Neighbor Lemma (5.3). \square

Corollary 5.7 (Big and Moderate Faces Corollary). *Suppose we have a locally irreducible embedding with three distinct antifaces A, B, C where $|V(A)| \geq n - k$ (for example, if $A0 \neq \emptyset$) and $|V(B)|, |V(C)| \geq 2k + 3$. Then there are a vertex of type AB and a vertex of type AC that are interlaced on A .*

Proof. Since $|V(A)| \geq n - k$, at most k vertices of B are not also vertices of A , so $|AB| \geq k + 3$. Similarly, $|AC| \geq k + 3$. Each vertex of AB is therefore adjacent to at least three vertices of AC , and vice versa, so we may apply the Three Neighbor Lemma (5.3) with $S = AB \cup AC$. \square

We will use the following general lemma to apply the Three Neighbor Lemma (5.3) in one particular situation. A graph is d -degenerate if every subgraph has minimum degree at most d .

Lemma 5.8 (Bipartite Degeneracy Lemma). *Let d be a nonnegative integer. If a simple bipartite graph G of order $n \geq 2d$ is d -degenerate then it has at most $d(n - d)$ edges. Thus, a simple bipartite graph of order $n \geq 2d$ with more than $d(n - d)$ edges has a subgraph with minimum degree at least $d + 1$.*

Proof. Construct a sequence of graphs $G_0, G_1, G_2, \dots, G_{n-2d}$, where $G_0 = G$ and for $i \geq 1$, $G_i = G_{i-1} - v_i$ where v_i has degree at most d in G_{i-1} . Since G_{n-2d} is a bipartite simple graph of order $2d$, it follows that $|E(G_{n-2d})| \leq d^2$, and $|E(G)| \leq |E(G_{n-2d})| + d(n - 2d) \leq d^2 + d(n - 2d) = d(n - d)$. \square

This lemma is sharp: if $n \geq 2d$, then $K_{d,n-d}$ is a d -degenerate bipartite graph with exactly $d(n - d)$ edges.

Corollary 5.9 (Diamond Corollary). *Suppose that we have a locally irreducible embedding with distinct antifaces A and B . Suppose that either $k = 0$ and $|AB| \geq 3$, or $|AB| \geq 3k + 4$, and that there exist vertices of types AP and BQ with $P, Q \notin \{A, B\}$ (possibly $P = Q$). Then there is a vertex of type AB that is either interlaced on A with a vertex of type AP or interlaced on B with a vertex of type BQ .*

Proof. Let x be a vertex of type AP , and x' a vertex of type BQ , and let S be the set of vertices of type AB adjacent to both x and x' . Then $|S| \geq |AB| - 2k$. Let H be the subgraph of D_{us} induced by S . Then $\delta(H) \geq |S| - k - 1 \geq |AB| - 3k - 1$. Every edge of H is an edge of A_{us} or B_{us} or both.

If $k = 0$ and $|AB| \geq 3$, then D_{us} is complete, which means that $|S| = |AB| \geq 3$ and H has a triangle. If $|AB| \geq 3k + 4$ then $\delta(H) \geq 3$ so H has a vertex of degree at least 3. In either case H has three edges that pairwise have a vertex in common. At least two of these edges, say tu and uv , belong to the same one of A_{us} or B_{us} .

If $tu, uv \in E(A_{us})$ then we apply the Diamond Lemma (5.4) to t, u, v and x to get x interlaced with one of t, u, v on A . If $tu, uv \in E(B_{us})$ then we apply the Diamond Lemma (5.4) to t, u, v and x' to get x' interlaced with one of t, u, v on B . \square

We could also try to improve the above result by finding a triangle in H when $k \geq 1$. To do this we could use Mantel's Theorem (an n -vertex simple graph with average degree greater than $n/2$ has a triangle) for H . However, this would require $|AB| \geq 4k + 3$, which is at least as strong a condition as $|AB| \geq 3k + 4$ when $k \geq 1$.

In our arguments we need the antifaces to be reasonably large (have many vertices) so that we can force interlacing to occur using the Three Neighbor Lemma (5.3) or the Diamond Lemma (5.4) or their consequences. The next result will be applied to use large antifaces to help 'blow up' smaller antifaces in an embedding, to obtain situations where we can make progress. It involves a cyclically ordered finite set of points with some possible break points (black), some reasonably uniformly distributed points (white) and some arbitrarily distributed points (red). It allows us to split the order at two black points so that we obtain something close to a desired division of the white and red points.

For this result we need to work with cyclic orders, both of a finite set of points and of the points of a circle. In both cases, for distinct points a and b , we denote by (a, b) the set of points encountered moving forward from a until we reach b (excluding a and b); we can then define $[a, b] = (a, b) \cup \{a, b\}$. We define $(a, a) = \emptyset$ and $[a, a] = \{a\}$. For $c \in \mathbb{R}$ we let $c\mathbb{Z} = \{ci \mid i \in \mathbb{Z}\}$. We treat a circle of circumference c as the additive group $\mathbb{R}/c\mathbb{Z}$ of reals modulo c , and identify its points with the interval $[0, c] \subseteq \mathbb{R}$.

Lemma 5.10 (Division Lemma). *Suppose C is a finite cyclic order where each point of C is colored black, white, or red. Let the black points be $b_0, b_1, b_2, \dots, b_{k-1}$ in cyclic order (indexed by elements of \mathbb{Z}_k), and let W and R be the sets of white and red points, respectively. Let m be a positive integer. Suppose that $|W \cap [b_i, b_{i+1}]| \leq m$ for every $i \in \mathbb{Z}_k$, and that $|W| > |R|$. Then for every $p \in \mathbb{R}$ with $m \leq p \leq |W \cup R| - m$ there are b_i and b_j , $i \neq j$, such that $p - m < |(W \cup R) \cap [b_i, b_j]| < p + m$.*

Proof. Let $\omega_i = |W \cap [b_i, b_{i+1}]| \leq m$ and $\rho_i = |R \cap [b_i, b_{i+1}]|$ for $i \in \mathbb{Z}_k$. Let $\ell = |W \cup R| = |W| + |R| = \sum_{i \in \mathbb{Z}_k} (\omega_i + \rho_i)$. Because $m \leq p \leq \ell - m$, we have $\ell \geq 2m$. Since $|W| > |R|$, we have $|W| \geq m + 1$, and so $k \geq 2$. Define (not necessarily distinct) real numbers x_0, x_1, \dots, x_k by

$$x_0 = 0, \quad \text{and} \quad x_{i+1} = x_i + \omega_i + \rho_i \quad \text{for } i \in \{0, 1, 2, \dots, k-1\}.$$

Then $x_k = \ell$, so we may identify $x_0 = 0$ with $x_k = \ell$ in $[0, \ell] \subseteq \mathbb{R}$ and treat the resulting set C' simultaneously as the group $\mathbb{R}/\ell\mathbb{Z}$, as a cyclic order, and as a circle of circumference ℓ . We use λ to denote the length measure on C' .

We now color each point of C' black, white, or red. All points x_i are black. For each $i \in \mathbb{Z}_k$ the interval (x_i, x_{i+1}) is empty if $\omega_i = \rho_i = 0$, but otherwise we color it as follows. If $\omega_i = 0$ and $\rho_i > 0$ we color all of (x_i, x_{i+1}) red. If $\omega_i > 0$ then we color $(x_i, x_i + \frac{1}{2}\omega_i) \cup (x_i + \frac{1}{2}\omega_i, x_{i+1})$ white and $[x_i + \frac{1}{2}\omega_i, x_{i+1} - \frac{1}{2}\omega_i]$ red (even if $\rho_i = 0$, the point halfway between x_i and x_{i+1} is red, which we need for (iii) below). Let B' , W' , and R' be the sets of black, white, and red points in C' , respectively. We note the following properties.

- (i) B' is a set of measure zero.
- (ii) $\lambda([x_i, x_{i+1}]) = \omega_i + \rho_i = |(W \cup R) \cap [b_i, b_{i+1}]|$, with $\lambda(W' \cap [x_i, x_{i+1}]) = \omega_i$, and $\lambda(R' \cap [x_i, x_{i+1}]) = \rho_i$.
- (iii) If $y \in B' \cup W'$ then there exists i with $x_i \in (y - \frac{1}{2}m, y + \frac{1}{2}m)$.

From (i) and (ii) it follows that $\ell = \lambda(C') = |R| + |W|$, with $\lambda(W') = |W|$, and $\lambda(R') = |R| < |W|$. Hence $\lambda(R') < \frac{1}{2}\lambda(C')$. Therefore, letting R'' be R' translated by $-p$ in $C' = \mathbb{R}/\ell\mathbb{Z}$, we have $\lambda(R' \cup R'') \leq \lambda(R') + \lambda(R'') = 2\lambda(R') < \lambda(C')$, so there exists $y \notin R' \cup R''$, which means that $y, y + p \notin R'$. By (iii) there is some $x_i \in (y - \frac{1}{2}m, y + \frac{1}{2}m)$ and some $x_j \in (y + p - \frac{1}{2}m, y + p + \frac{1}{2}m)$; these intervals are disjoint because $m \leq p \leq \ell - m$, so $x_i \neq x_j$. Now $\lambda([x_i, x_j])$ differs by less than m from $\lambda([y, y + p]) = p$. Again from (i) and (ii) we get $\lambda([x_i, x_j]) = |(W \cup R) \cap [b_i, b_j]|$, so our result follows. \square

We will apply the Division Lemma (5.10) with p either an integer or a half-integer. In these cases the bounds on $q = |(W \cup R) \cap [b_i, b_j]|$, which say that $|q - p| < m$, cannot be improved. When p is an integer, the bound becomes $|q - p| \leq m - 1$, and when p is a half-integer it becomes $|q - p| \leq m - \frac{1}{2}$. Consider \mathbb{Z}_{4m+1} , with $m \geq 1$, where $B = \{b_0 = 0, b_1 = m+1\}$, and suppose $[b_0, b_1]$ contains m white points, and $[b_1, b_0]$ contains m white and $2m-1$ red points. For $p = 2m-1$ we must use $|W \cup R \cap [b_0, b_1]| = m = p - (m-1)$ to satisfy the conclusion, and for $p = 2m$ we must use $|W \cup R \cap [b_1, b_0]| = 3m-1 = p+m-1$. These show that the bound $|q - p| \leq m - 1$ cannot be improved when p is an integer. For $p = 2m - \frac{1}{2}$ we

must use either $|(W \cup R) \cap [b_0, b_1]| = m = p - (m - \frac{1}{2})$ or $|(W \cup R) \cap [b_1, b_0]| = 3m - 1 = p + m - \frac{1}{2}$. This shows that the bound $|q - p| \leq m - \frac{1}{2}$ cannot be improved when p is a half-integer.

We also cannot relax the condition $|W| > |R|$ to $|W| \geq |R|$. Consider \mathbb{Z}_{6m+3} , $m \geq 1$, with $B = \{b_0 = 0, b_1 = 2m+1, b_2 = 4m+2\}$, where each interval $[b_i, b_{i+1}]$ contains m white and m red points, so that $|W| = |R| = 3m$. If we let $p = 3m$, then we cannot achieve $2m = p - m < |(W \cup R) \cap [b_i, b_j]| < p + m = 4m$.

We now use the Division Lemma (5.10) to show that we can use a large face A to increase the size of a smaller face B .

Lemma 5.11 (Blow Up Lemma). *Suppose we have a locally irreducible embedding Φ with distinct antifaces A and B , and there exists a vertex x of type AB . Suppose that $|V(A)| \geq 5$ and $n \geq k + 3$. Then there is an embedding Φ' (possibly $\Phi' = \Phi$) that has the same profaces and antifaces as Φ except that A and B are replaced by antifaces A' and B' where $|V(A')|, |V(B')| \geq \min(\frac{1}{2}|V(A)| - 1, \frac{1}{2}(n - k - 1))$ and x has type $A'B'$.*

Proof. Let $a = |V(A)|$ and $b = |V(B)|$. Suppose that the sequence of vertices along A is $(v_0, v_1, v_2, \dots, v_{t-1})$, with indexing by elements of \mathbb{Z}_t . Let $Y = \{y \neq x \mid y = v_{i-1} \text{ or } v_{i+1} \text{ for some } i \in \mathbb{Z}_t \text{ with } v_i = x\}$; in other words, $Y \subseteq V(A)$ is the set of vertices joined to x by at least one arc of A . We know x has at least $n - k - 1$ neighbors, each of which must be joined to x by an arc of A or B (or both). A vertex joined to x by an arc of B must be a vertex of B , and x has at most $b - 1$ neighbors in $V(B) - \{x\}$. Thus, $|Y| \geq (n - k - 1) - (b - 1) = n - k - b$. Let $Z = V(A) - (\{x\} \cup Y)$.

We will color some of the elements of \mathbb{Z}_t as follows. If $v_h = x$ then h is black. For each $y \in Y$ choose one arc e_y of A joining x and y , let v_i be the occurrence of y on A as an end of e_y , and color i white. For each $z \in Z$ choose one occurrence v_j of z on A and color j red. So there are $|Y|$ white points and $|Z|$ red points. There are at most two arcs e_y between consecutive occurrences of x on A , and hence at most two white points of \mathbb{Z}_t between consecutive black points.

Suppose first that $|Y| > |Z|$. Discard all uncolored points of \mathbb{Z}_t , creating a cyclic order C with some black points, $\ell = a - 1 \geq 4$ white or red points, and more white than red. Apply the Division Lemma (5.10) with $m = 2$ and $p = \frac{1}{2}\ell$: since $\ell \geq 4$ we have $2 \leq p \leq \ell - 2$. We obtain distinct black points $h, j \in C$ so that the interval $[h, j]$ of C contains q_1 white or red points where $p - 2 < q_1 < p + 2$. The number of white or red points in $[j, h]$ is $q_2 = \ell - q_1 = 2p - q_1$, which also satisfies $p - 2 < q_2 < p + 2$. Since $p, q_1, q_2 \in \frac{1}{2}\mathbb{Z}$ we get $q_1, q_2 \geq p - \frac{3}{2} = \frac{1}{2}(\ell - 3)$. Write A (after cyclic shifting) as $A_1 \cdot A_2$ where $A_1 = v_h (= x) \dots v_j (= x)$ and $A_2 = v_j \dots v_h$. Each A_i contains q_i distinct points of $V(A) - \{x\}$ and hence $q_i + 1 \geq \frac{1}{2}(\ell - 1)$ distinct points of $V(A)$. Thus, applying part (b) of the Three Face Lemma (5.1) to $A = A_1 \cdot A_2$ and B gives new antifaces A' and B' each with at least $\frac{1}{2}(\ell - 1) = \frac{1}{2}a - 1$ vertices.

Now suppose that $|Y| \leq |Z|$. Let $c = \min(\frac{1}{2}a - 1, \frac{1}{2}(n - k - 1))$. Since $a \geq 5$, $a \geq \frac{1}{2}a - 1 \geq c$. If $b \geq \frac{1}{2}(n - k - 1)$ then $b \geq c$ as well, and we take $A' = A$, $B' = B$. So we may assume that $b \leq \frac{1}{2}(n - k - 2)$. Now $|Y| \geq n - k - b \geq \frac{1}{2}(n - k + 2)$ and therefore $|Y| \geq 3$ since $n \geq k + 3$ and $|Y|$ is an integer. Uncolor some points of Z so that the number of red points is $|Y| - 1$, which is less than the number of white points. Apply the Division Lemma

(5.10) (discarding uncolored points) and part (b) of the Three Face Lemma (5.1) as in the previous paragraph, but with $\ell = 2|Y| - 1 \geq 5 \geq 4$ white or red points. This gives new antifaces A' and B' each with at least $\frac{1}{2}(\ell - 1) = |Y| - 1 \geq \frac{1}{2}(n - k)$ vertices.

Thus, we always have A' and B' each with at least $\min(\frac{1}{2}a - 1, \frac{1}{2}(n - k - 1))$ vertices. In all cases $x \in V(A') \cap V(B')$. \square

6. Proof of the main result

In this section we prove Theorem 4.1. We need one final tool, namely the touch graph of a locally irreducible embedding. The touch graph encodes information about how each antiface in a locally irreducible embedding touches itself and other antifaces at the vertices, and will make our arguments easier to explain and visualize.

Definition 6.1. Let Φ be a locally irreducible oriented directed embedding. The *touch graph* K is an undirected graph which may have loops and multiple edges. The vertices of K are the antifaces of Φ , and the edges of K are the vertices of D , i.e., $V(K) = \mathcal{A}$ and $E(K) = V(D)$. For the incidence relation of K , a vertex of D of type $A0$ in Φ becomes a loop of K at A , and a vertex of type AB becomes a link (non-loop edge) of K joining A and B . Thus in K , $A\bar{B}$ is the set of edges incident with A but not B , $A0$ is the set of loops at A , $A1$ is the set of links at A , and $V(A)$ is $E_K(A)$, the set of all edges incident with A in K . Since Φ is locally irreducible, K is well-defined.

It will sometimes be convenient to represent K in simplified form, where we replace multiple parallel edges (including loops) by a single edge weighted by the number of parallel edges.

Critically, in the touch graph, all sets of vertices of D are interpreted as sets of edges. Thus, symbols such as u, v, w, \dots represent vertices in D but edges in K . See Figures 3, 4, and 5.

Figure 3 shows a directed embedding of a digraph D , where D is an orientation of K_{12} minus a perfect matching. It has two profaces, one dark teal (dark, if color is not visible) and one light teal (light), shown on the left. It has six anti-faces, one light yellow (very light), one orange (light), one dark red (dark), and three black, shown on the right. It is straightforward to check from the rotations at the vertices that this is indeed a directed embedding.

Notice that the three antifaces colored light yellow, orange, and dark red, meet at vertex d in the embedding of D . Thus we can apply Lemma 5.1 to merge these faces into a single orange face, shown in Figure 4. After merging these faces, the embedding of D has only four antifaces and is locally irreducible. We can thus create the touch graph shown in Figure 5.

Note that K is connected when D is connected. This can be seen as follows. Given $A, B \in \mathcal{A}$ choose $v \in V(A)$ and $w \in V(B)$. The sequence of vertices on a vw -path in D_{us} gives a sequence of adjacent edges in K with the first, v , adjacent to A and the last, w , adjacent to B , and so K has a connected subgraph containing A and B .

We now have all the tools we need to prove the main theorem of this paper, which we repeat here for convenience.

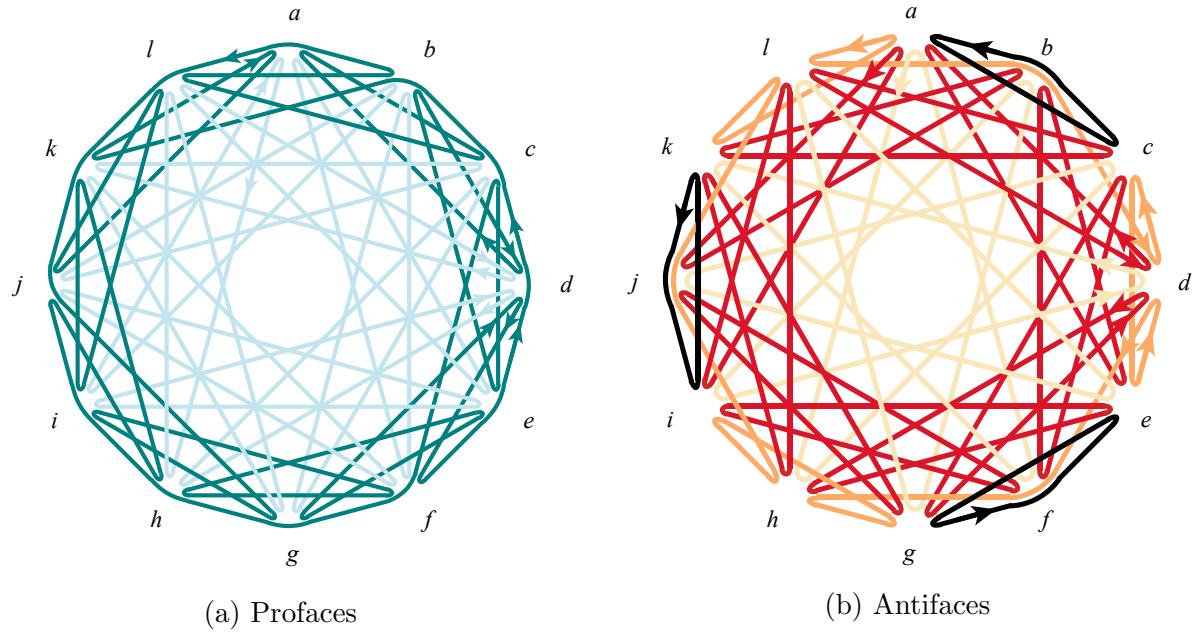


Figure 3: A directed embedding of an orientation of K_{12} minus a perfect matching.

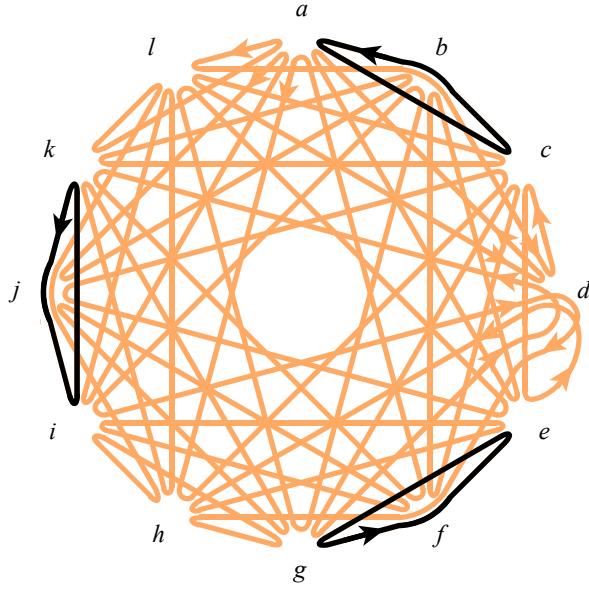


Figure 4: Applying the 3-face lemma to the light yellow, orange, and dark red faces at vertex d to create only a single orange face.

Theorem 4.1. *Let D be an n -vertex eulerian digraph where $\delta(D_{us})$, the minimum degree of the underlying simple undirected graph of D , is at least $(4n + 2)/5$. Let \mathcal{C} be a directed circuit decomposition of D . Then there is a directed embedding of D in an oriented surface*

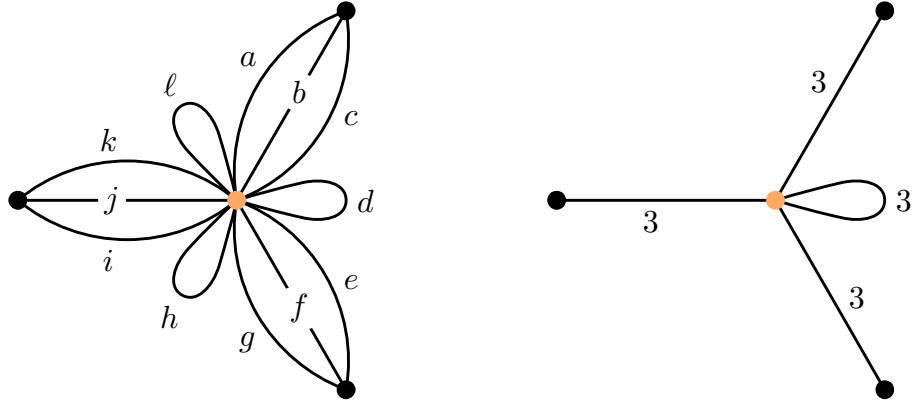


Figure 5: Touch graph for the three black and one orange (light) antifaces of Figure 4, with simplified weighted version at right.

(i.e., an orientable surface with a specific orientation) with the elements of \mathcal{C} as the profaces and with exactly one or two antifaces, depending on whether $|V(D)| + |A(D)| + |\mathcal{C}|$ is odd or even, respectively.

This embedding has maximum genus among all orientable directed embeddings of D in which all elements of \mathcal{C} are faces.

Proof of Theorem 4.1. We have an eulerian digraph D satisfying $\delta(D_{\text{us}}) \geq (4n + 2)/5$. In terms of $k = n - 1 - \delta(D_{\text{us}})$, this means that $n \geq 5k + 7$. By our definition of eulerian, D is connected (this also follows from $\delta(D_{\text{us}}) \geq n/2$).

We are given a directed circuit decomposition \mathcal{C} of D . By Lemma 2.4 there exists an oriented directed embedding Φ of D whose set of profaces is \mathcal{C} . As usual, we let \mathcal{A} denote the set of antifaces. We will show that if $|\mathcal{A}| \geq 3$ then we can reduce the number of antifaces, without changing the profaces. After sufficiently many reductions, we obtain an embedding with one or two antifaces as desired. We use a case analysis, as follows. In Case 1 we dispose of situations where Φ is not locally irreducible. For the remaining cases we consider the structure of the touch graph K . Case 2 is when K has no loops, with subcases 2.1 and 2.2 depending on whether K is not or is a star, respectively. Case 3 is when K has loops, with subcases 3.1 and 3.2 depending on whether there are loops at two or more vertices or at just one vertex, respectively.

Case 1. Suppose some vertex belongs to three distinct antifaces. Then applying part (a) of the Three Face Lemma (5.1) reduces the number of antifaces.

So we may henceforth assume that Case 1 does not happen, i.e., that Φ is locally irreducible. We let K be the touch graph of Φ . We say K is a *star* if K has one vertex A with which every edge is incident; in other words, if $V(A) = E_K(A)$ is equal to $V(D) = E(K)$ for some $A \in \mathcal{A}$.

Case 2. We begin by assuming that K has no loops.

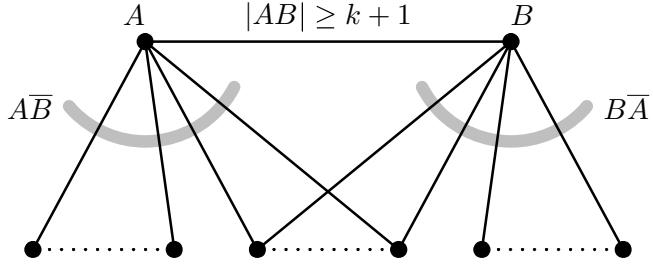


Figure 6: Touch graph for Cases 2.1.2, 2.1.3, and 2.1.4.

Case 2.1. Suppose K has no loops and is not a star, i.e., $P_0 = \emptyset$ and $V(P) \neq V(D)$ for all $P \in \mathcal{A}$. Then $P_1 = V(P)$ for all $P \in \mathcal{A}$. Let $A, B \in \mathcal{A}$ be such that $\max\{|PQ| \mid P, Q \in \mathcal{A}, P \neq Q\} = |AB|$, i.e., A and B have the maximum number of common vertices over all pairs of circuits in \mathcal{A} . Without loss of generality assume that $|V(A)| \geq |V(B)|$, which implies that $|A\bar{B}| \geq |B\bar{A}|$.

Our strategy in Case 2.1 is to use the Three Neighbor Corollary (5.6) (or a similar argument based on the Bipartite Degeneracy Lemma (5.8)) if $|AB|$ is not too large, or the Diamond Corollary (5.9) if it is. Because of the structure of our argument, we deal with small $|AB|$ first, and then the remaining situations in decreasing order of $|AB|$.

We know that every edge of K is adjacent to (shares a vertex with) all but at most k of the other edges. If $S \subseteq E(K)$ with $|S| \geq k + 1$, then every edge of K not in S is adjacent to at least $|S| - k$ edges of S , and hence is incident with an end of some edge of S .

Case 2.1.1. Suppose that $|AB| \leq k$. There is some edge v in K of type AB , and at most k edges are not adjacent to v , i.e., at most k edges of K are incident with neither A nor B . Thus, $|V(A) \cup V(B)| \geq n - k \geq 4k + 7$. Since $|V(A)| \geq |V(B)|$, we have $|A_1| = |V(A)| \geq 2k + 4$. Since $|AP| \leq |AB| \leq k$ for all $P \in \mathcal{A}$ with $P \neq A$, we have $|A_1| - |AP| \geq k + 4$ for all such P . Thus, there is interlacing by the Three Neighbor Corollary (5.6), and by the Interlaced Faces Lemma (5.2) we can reduce the number of antifaces. (This argument also works for $|AB| = k + 1$, although we cover that situation in Case 2.1.4.)

In the remainder of Case 2.1, we may suppose that $|AB| \geq k + 1$. Thus, every vertex not in AB must be adjacent to at least one vertex in AB , and hence every edge of K not in AB is incident with either A or B . Thus, $V(D) = E(K) = V(A) \cup V(B) = A\bar{B} \cup AB \cup B\bar{A}$. Since $A\bar{B} = V(D) - V(B)$ and $V(B) \neq V(D)$ and K is not a star, we have $A\bar{B} \neq \emptyset$. Similarly, $B\bar{A} \neq \emptyset$. Figure 6 shows the (simplified weighted) structure of K in this situation.

Case 2.1.2. Suppose that $|AB| \geq 3k + 4$. Since $A\bar{B}, B\bar{A} \neq \emptyset$, there are $v \in AP, w \in BQ$ with $P, Q \notin \{A, B\}$. So by the Diamond Corollary (5.9) there is interlacing, and we can use the Interlaced Faces Lemma (5.2) to reduce the number of antifaces.

Case 2.1.3. Suppose that $|AB| = 3k + 3$. If $k = 0$, then we can use the Diamond Corollary (5.9) and the Interlaced Faces Lemma (5.2) to reduce the number of antifaces as in Case 2.1.2. Therefore, we may assume that $k \geq 1$.

We have $n = |\overline{AB}| + |\overline{BA}| + |AB| = |\overline{AB}| + |\overline{BA}| + 3k + 3 \geq 5k + 7$, so $|\overline{AB}| + |\overline{BA}| \geq 2k + 4$. Since $|\overline{AB}| \geq |\overline{BA}|$, we have $|\overline{AB}| \geq k + 2$.

Now think of \overline{AB} and AB as vertex sets in D , and consider the bipartite subgraph H of D_{us} consisting of $V(H) = V(A)$ and exactly the edges of D_{us} between \overline{AB} and AB . We claim that H has a subgraph H' with minimum degree at least 3. This will follow from the Bipartite Degeneracy Lemma (5.8) with $d = 2$ provided $|V(H)| \geq 4$ and $|E(H)| > 2(|V(H)| - 2)$. We can satisfy these conditions as follows. We have $|V(H)| = |\overline{AB}| + |AB| = |\overline{AB}| + 3k + 3 \geq 4k + 5 \geq 4$. Every vertex of \overline{AB} is adjacent in D_{us} , and hence in H , to at least $|AB| - k = 2k + 3$ vertices of AB . Since $k \geq 1$, it follows that

$$\begin{aligned} |E(H)| - 2(|V(H)| - 2) &\geq (2k + 3)|\overline{AB}| - 2(|\overline{AB}| + 3k + 1) = (2k + 1)|\overline{AB}| - 6k - 2 \\ &\geq (2k + 1)(k + 2) - 6k - 2 = 2k^2 - k = (2k - 1)k > 0. \end{aligned}$$

Thus, $|E(H)| > 2(|V(H)| - 2)$ and we can apply the Bipartite Degeneracy Lemma (5.8) with $d = 2$. Therefore H' exists, and we can apply the Three Neighbor Lemma (5.3) with $S = V(H')$ to show that there is interlacing, and then use the Interlaced Faces Lemma (5.2) to reduce the number of antifaces.

Case 2.1.4. Suppose that $k + 1 \leq |AB| \leq 3k + 2$. Then $|\overline{AB}| + |\overline{BA}| = n - |AB| \geq (5k + 7) - (3k + 2) = 2k + 5$. Since $|\overline{AB}| \geq |\overline{BA}|$, we have $|\overline{AB}| \geq k + 3$. Now for any $P \neq A$ we have $|A1| - |AP| \geq |A1| - |AB| = |V(A)| - |AB| = |\overline{AB}| \geq k + 3$. Therefore, by the Three Neighbor Corollary (5.6) there is interlacing, and we can use the Interlaced Faces Lemma (5.2) to reduce the number of antifaces.

In the remaining cases, Cases 2.2 and 3 below, we often need to modify the embedding to increase the number of vertices of some faces, using the Blow Up Lemma (5.11). We can always apply the Blow Up Lemma (5.11) to $A, B \in \mathcal{A}$ if $|V(A)| \geq n - k$ and D has a vertex x of type AB , because then $|V(A)| \geq 4k + 7 \geq 5$, and we know that $n \geq 5k + 7 \geq k + 3$. We will say we are *blowing up B using A*.

When we apply the Blow Up Lemma (5.11), the number of antifaces does not change. The two new antifaces A' and B' that replace A and B will have $|V(A')|, |V(B')| \geq \min(\frac{1}{2}(n - k) - 1, \frac{1}{2}(n - k - 1)) = \frac{1}{2}(n - k) - 1 \geq \frac{1}{2}(4k + 7) - 1 = 2k + 2\frac{1}{2}$, but since these are integers, $|V(A')|, |V(B')| \geq 2k + 3$. We therefore create two faces satisfying the ‘moderate face’ condition of the Big and Moderate Faces Corollary (5.7). Also we know that $V(A') \cup V(B') = V(A) \cup V(B)$, and $x \in A'B'$, so that $A'B' \neq \emptyset$.

Case 2.2. Suppose that K has no loops and is a star, i.e., $P0 = \emptyset$ for all $P \in \mathcal{A}$ and $V(A) = V(D) = E(K)$ for some $A \in \mathcal{A}$. Then $AP = P1 = V(P)$ for all $P \in \mathcal{A}$, $P \neq A$. Let $B \in \mathcal{A}$ be such that $\max\{|AP| \mid P \in \mathcal{A}, P \neq A\} = |AB|$.

Our strategy in Case 2.2 is to use the Three Neighbor Corollary (5.6) if $|AB|$ is not too large, but if it is, we need to blow up a third face and use the Big and Moderate Faces Corollary (5.7).

If $|A1| - |AB| \geq k + 3$ then since $|AP| \leq |AB|$ we have $|A1| - |AP| \geq k + 3$ for all $P \neq A$, there is interlacing by the Three Neighbor Corollary (5.6), and by the Interlaced Faces Lemma (5.2) we can reduce the number of antifaces. Therefore, we may assume that

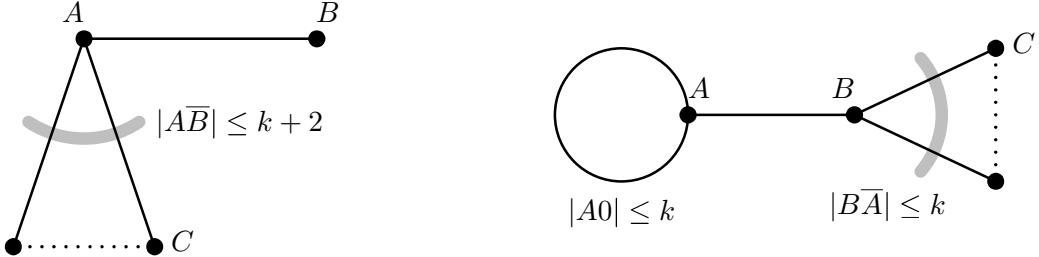


Figure 7: Touch graph for Case 2.2 with $|AB|$ large (left) and for Case 3.2.1 (right).

$|A\bar{B}| = |A1| - |AB| = n - |AB| \leq k + 2$, i.e., that $|V(B)| = |AB| \geq n - k - 2$. Thus, at most $k + 2$ vertices of D do not belong to B . The left side of Figure 7 shows the structure of K in this situation.

There must be some third face C , so $|AC| = |V(C)| > 0$. Since $|V(A)| = n$, we can blow up C using A , obtaining a new embedding Φ' (and touch graph K') with set of antifaces \mathcal{A}' in which A, C are replaced by A', C' with $|V(A')|, |V(C')| \geq 2k + 3$ and $A'C' \neq \emptyset$. If Φ' is not locally irreducible, we may apply Case 1 to reduce the number of antifaces, so we may assume that Φ' is locally irreducible.

Since at most $k + 2$ vertices of D do not belong to B , we have $|BA'| \geq |V(A')| - (k + 2) \geq k + 1$ and similarly $|BC'| \geq k + 1$. Also, $A'C' \neq \emptyset$. Thus, K' contains at least one triangle $(A'BC')$ and hence is not a star.

Assume without loss of generality that $|A'0| \geq |C'0|$. If $|A'0| = 0$ then K' has no loops ($P0$ has not changed for $P \neq A, C$), so we can apply Case 2.1 to reduce the number of antifaces. So we may assume that $|A'0| > 0$. Then $|V(A')| \geq n - k$ by Observation 5.5, $|V(B)| \geq n - k - 2 \geq 4k + 5 \geq 2k + 3$, and $|V(C')| \geq 2k + 3$. Thus, there is interlacing by the Big and Moderate Faces Corollary (5.7), and we can apply the Interlaced Faces Lemma (5.2) to reduce the number of antifaces.

Case 3. Suppose that K has at least one loop, i.e. there is at least one $P \in \mathcal{A}$ with $P0 \neq \emptyset$. By Observation 5.5, $|V(P)| \geq n - k$ for each such P .

Our strategy in Case 3 is to blow up faces until we either obtain a previous case or can apply the Big and Moderate Faces Corollary (5.7).

Case 3.1. Suppose there are loops incident with at least two vertices of K . Thus there are distinct $A, B \in \mathcal{A}$ with both $A0$ and $B0$ nonempty.

Since K is connected and $|\mathcal{A}| \geq 3$, at least one of A and B , say A , is adjacent to a third vertex C of K , i.e., $AC \neq \emptyset$. Since $|V(A)| \geq n - k$ we can blow up C using A , to obtain a new embedding Φ' containing faces B, A', C' with $|V(B)| \geq n - k$ and $|V(A')|, |V(C')| \geq 2k + 3$. If Φ' is not locally irreducible we apply Case 1, and otherwise we have interlacing by the Big and Moderate Faces Corollary (5.7) and use the Interlaced Faces Lemma (5.2), reducing the number of antifaces in either case.

Case 3.2. Suppose there is exactly one vertex of K with one or more incident loops.

Case 3.2.1. Suppose that the one vertex with loops has only one neighbor in K . In other words, $A0 \neq \emptyset$ for a unique $A \in \mathcal{A}$ and $AB \neq \emptyset$ for a unique $B \in \mathcal{A}$, $B \neq A$.

Since $|V(A)| \geq n - k \geq k + 1$, every edge of K not in $V(A)$ must share an end with an edge in $V(A)$, so every edge of K not in $V(A)$ must be incident with B . Thus, $V(D) = E(K) = A0 \cup AB \cup \overline{BA}$. Since K is connected and $|\mathcal{A}| \geq 3$, there must be $C \in \mathcal{A}$ adjacent to B in K . Since $v \in BC$ is nonadjacent to at most k other edges of K , $|A0| \leq k$ and so $|V(B)| \geq n - k$. The right side of Figure 7 shows the structure of K in this situation. We blow up C using B to give Φ' containing antifaces A, B', C' with $|V(A)| \geq n - k$ and $|V(B')|, |V(C')| \geq 2k + 3$. If Φ' is not locally irreducible we apply Case 1, and otherwise we have interlacing by the Big and Moderate Faces Corollary (5.7) and use the Interlaced Faces Lemma (5.2), reducing the number of antifaces in either case.

Case 3.2.2. Suppose there is exactly one vertex of K with one or more incident loops, and this vertex has at least two neighbors in K . In other words, $A0 \neq \emptyset$ for a unique $A \in \mathcal{A}$, and there are at least two $P \neq A$ with $AP \neq \emptyset$. This is our final case, so if we can modify the embedding and obtain an earlier case, we are finished.

Choose $B \in \mathcal{A}$ with $AB \neq \emptyset$. Blow up B using A , obtaining Φ' (and touch graph K') with new antifaces A', B' with $|V(A')|, |V(B')| \geq 2k + 3$. If Φ' is not described by Case 3.2.2, we can apply a previous case to reduce the number of antifaces. Thus, we may assume Φ' is locally irreducible, there is exactly one vertex of K' with one or more incident loops, and this vertex has at least two neighbors in K' . The vertex of K' with incident loops must be one of A' or B' , say A' , because blowing up B using A cannot create loops at other vertices of K . Thus $|V(A')| \geq n - k$. Let C be a neighbor of A' other than B' . Blow up C using A' , obtaining Φ'' (and touch graph K'') with antifaces A'', C'' with $|V(A'')|, |V(C'')| \geq 2k + 3$. Again, we may assume that Φ'' is described by Case 3.2.2, so there is exactly one vertex of K'' with one or more incident loops. This must be one of A'' or C'' , say A'' . Therefore, Φ'' has antifaces A'', B', C'' with $|V(A'')| \geq n - k$ and $|V(B')|, |V(C'')| \geq 2k + 3$. Hence, there is interlacing by the Big and Moderate Faces Corollary (5.7), and we can use the Interlaced Faces Lemma (5.2) to reduce the number of antifaces.

Thus, we can always reduce $|\mathcal{A}|$ if $|\mathcal{A}| \geq 3$, and by repeated reductions we obtain an embedding Φ_0 with $|\mathcal{A}| = 1$ or 2 antifaces and with \mathcal{C} as the set of profaces.

We now verify that the embedding we have obtained satisfies all the conclusions of the theorem. Consider the given digraph D and the given directed circuit decomposition \mathcal{C} , and take any directed embedding of D in which \mathcal{C} is a subset of the faces. By Observation 2.3 this is orientable, we can choose an orientation so that \mathcal{C} is the collection of profaces, and if \mathcal{W} is the collection of other faces then $|\mathcal{W}|$ has the same parity as $|V(D)| + |A(D)| + |\mathcal{C}|$. Hence, $|\mathcal{W}| \geq 1$ or 2 depending on whether $|V(D)| + |A(D)| + |\mathcal{C}|$ is odd or even, respectively.

The embedding Φ_0 is an example of an embedding described in the previous paragraph, with $\mathcal{W} = \mathcal{A}$, and it achieves the lower bound on $|\mathcal{W}|$. It therefore minimizes the number of faces, and hence maximizes the genus, among all directed embeddings of D in which all elements of \mathcal{C} are faces. \square

The proof of Theorem 4.1 describes a procedure for reducing the number of antifaces

in an embedding, until only one or two remain. It is easy to derive an algorithm from this proof. The key manipulation of embeddings that we use is the Three Face Lemma (5.1), which can be implemented very simply using the rotation scheme representation of an embedding. Almost all of the steps in the proof, and the supporting results in Sections 2 and 5, involve explicit constructions that are easily implemented in polynomial time. The one possible exception is our use of the Bipartite Degeneracy Lemma (5.8): we use existence of a subgraph of minimum degree at least $d+1$ in a graph with sufficiently many edges. However, it is not difficult to make this constructive: if a graph is known to be not d -degenerate, then a nonempty subgraph with minimum degree at least $d+1$ can be obtained by repeatedly deleting vertices of degree at most d until it is impossible to continue. Our proof of Theorem 4.1 therefore yields a polynomial time algorithm to find the promised embedding.

7. Conclusion

Naturally, we would like to improve our main result by weakening the required degree condition. This seems difficult to do using our current arguments. The hypothesis $\delta(D_{\text{us}}) \geq (4n+2)/5$, or equivalently $n \geq 5k+7$, is used in a tight way in many places in our proof. Every time we use the Blow Up Lemma (5.11) to create two moderate-sized faces so we can apply the Big and Moderate Faces Corollary (5.7), in Cases 2.2 and 3 of the proof of Theorem 4.1, we require $n \geq 5k+7$ to guarantee that the moderate-sized faces have at least $2k+3$ vertices. We also require $n \geq 5k+7$ for Case 2.1: Case 2.1.4 requires $n \geq 5k+7$ to handle the situation when $|AB| = 3k+2$, and we cannot extend the arguments in the other subcases of Case 2.1 to handle this situation.

However, we have no evidence to suggest that the degree condition in Theorem 4.1 is best possible. There may be a weaker condition that allows us to construct an oriented directed embedding of an eulerian digraph D with a specified directed circuit decomposition \mathcal{C} as the profaces and just one or two antifaces. In particular, this condition might give Theorem 3.5 on latin squares as a corollary. The condition may even be much weaker, involving only a constant lower bound on minimum degree (of D or D_{us}), possibly combined with a connectivity condition such as a constant lower bound on edge-connectivity. We only have very weak restrictions on what these bounds could be. Examples discussed in [10, Section 3] show that there are 4-edge-connected 4-regular (non-simple) graphs with no bi-eulerian embedding, and examples from [15] and [10, Section 5] show that there are 4-edge-connected 6-regular graphs with a triangular decomposition that cannot be completed to an orientable embedding by adding at most two more faces. Digraph examples can be derived from these by suitably directing the edges.

For further discussion of related problems, such as determining maximum genus for orientable directed embeddings of a given eulerian digraph, we refer the reader to [10].

We close by noting that the general problem of finding spanning walks in graphs subject to various restrictions arises in many settings. Our work here was motivated by DNA self-assembly problems, but protein self-assembly problems, which are also related to graph embeddings but require different conditions, have also been studied, for example in [16]. In addition to these applications involving biological structures, and to illustrate the range of

these spanning walk problems, there is also the recent work [6, 7]. These papers give more variations, such as allowing edges to be repeated multiple times in any direction with different turning constraints from the topologically motivated ones given here, with application to robotics. We expect that the interplay among these many and varied settings and problems will prove both interesting and productive.

References

- [1] Dan Archdeacon, C. Paul Bonnington, and Jozef Širáň. “A Nebeský-type characterization for relative maximum genus”. *J. Combin. Theory Ser. B* 73.1 (1998), pp. 77–98.
- [2] C. Paul Bonnington. “The relative maximum genus of a graph”. *J. Combin. Theory Ser. B* 60.2 (1994), pp. 195–206.
- [3] C. Paul Bonnington et al. “Embedding digraphs on orientable surfaces”. *J. Combin. Theory Ser. B* 85.1 (2002), pp. 1–20.
- [4] Yichao Chen and Wenjie Fang. “A character approach to directed genus distribution of graphs: the bipartite single-black-vertex case”. *Discrete Math.* 345.6 (2022), paper 112833 (12 pages).
- [5] Michelle Delcourt and Luke Postle. “Progress towards Nash-Williams’ conjecture on triangle decompositions”. *J. Combin. Theory Ser. B* 146 (2021), pp. 382–416.
- [6] Erik D. Demaine, Yael Kirkpatrick, and Rebecca Lin. “Graph threading”. In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Ed. by Venkatesan Guruswami. Vol. 287. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 38:1–38:18. URL: <https://doi.org/10.4230/lipics.itcs.2024.38>.
- [7] Erik D. Demaine, Yael Kirkpatrick, and Rebecca Lin. *Graph threading with turning costs*. 2024. arXiv: [2405.17953](https://arxiv.org/abs/2405.17953).
- [8] Richard A. Duke. “The genus, regional number, and Betti number of a graph”. *Canadian J. Math.* 18 (1966), pp. 817–822.
- [9] Jack Edmonds. “On the surface duality of linear graphs”. *J. Res. Nat. Bur. Standards Sect. B* 69B (1965), pp. 121–123.
- [10] M. N. Ellingham and Joanna A. Ellis-Monaghan. *Bi-eulerian embeddings of graphs and digraphs*. Submitted. 2024. arXiv: [2404.00325](https://arxiv.org/abs/2404.00325).
- [11] M. N. Ellingham and Joanna A. Ellis-Monaghan. “Edge-outer graph embedding and the complexity of the DNA reporter strand problem”. *Theoret. Comput. Sci.* 785 (2019), pp. 117–127.

[12] M. N. Ellingham and Joanna. A. Ellis-Monaghan. “Maximum genus orientable embeddings from circuit decompositions of dense eulerian graphs and digraphs”. In: *Proceedings of the 12th European Conference on Combinatorics, Graph Theory and Applications (Prague, 2023)*. Masaryk University Press, 2023, pp. 401–406. URL: <https://doi.org/10.5817/CZ.MUNI.EUROCOMB23-056>.

[13] M. N. Ellingham and Adam Weaver. “Constructing all minimum genus embeddings of $K_{3,n}$ ”. In: *The International Conference on Topological and Geometric Graph Theory*. Vol. 31. Electron. Notes Discrete Math. Elsevier Sci. B. V., Amsterdam, 2008, pp. 293–298.

[14] Joanna A. Ellis-Monaghan and Iain Moffatt. *Graphs on Surfaces: Dualities, Polynomials, and Knots*. SpringerBriefs in Mathematics. Springer, New York, 2013, xii+139 pages.

[15] Grahame Erskine, Terry Griggs, and Jozef Širáň. “On the upper embedding of symmetric configurations with block size 3”. *Discrete Math.* 343.4 (2020), paper 111774 (6 pages).

[16] Gašper Fijavž, Tomaž Pisanski, and Jernej Rus. “Strong traces model of self-assembly polypeptide structures”. *MATCH Commun. Math. Comput. Chem.* 71.1 (2014), pp. 199–212.

[17] Herbert Fleischner. *Eulerian Graphs and Related Topics. Part 1. Vol. 1*. Vol. 45. Annals of Discrete Mathematics. North-Holland Publishing Co., Amsterdam, 1990, xiv+387 pages.

[18] Mike J. Grannell, Terry S. Griggs, and Jozef Širáň. “Maximum genus embeddings of Steiner triple systems”. *European J. Combin.* 26.3-4 (2005), pp. 401–416.

[19] Terry S. Griggs, Thomas A. McCourt, and Jozef Širáň. “On the upper embedding of Steiner triple systems and Latin squares”. *Ars Math. Contemp.* 18.1 (2020), pp. 127–135.

[20] Terry S. Griggs, Constantinos Psomas, and Jozef Širáň. “Maximum genus embeddings of Latin squares”. *Glasg. Math. J.* 60.2 (2018), pp. 495–504.

[21] Jonathan L. Gross and Thomas W. Tucker. *Topological Graph Theory*. Reprint of the 1987 original [Wiley, New York; MR0898434 (88h:05034)] with a new preface and supplementary bibliography. Dover Publications, Inc. Mineola, NY, 2001, xvi+361 pages.

[22] Nataša Jonoska, Nadrian C. Seeman, and Gang Wu. “On existence of reporter strands in DNA-based graph structures”. *Theoret. Comput. Sci.* 410.15 (2009), pp. 1448–1460.

[23] A. Kotzig. “Eulerian lines in finite 4-valent graphs and their transformations”. In: *Theory of Graphs (Proc. Colloq., Tihany, 1966)*. Academic Press, New York, 1968, pp. 219–230.

- [24] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001, xii+291 pages.
- [25] C. St.J. A. Nash-Williams. In: *Combinatorial Theory and its Applications, I-III (Proc. Colloq., Balatonfüred, 1969)*. Vol. 4. Colloq. Math. Soc. János Bolyai. Untitled problem concerning decomposition of graphs into triangles, with comment by R. L. Graham. North-Holland, Amsterdam-London, 1970, pp. 1179–1183.
- [26] Ladislav Nebeský. “A new characterization of the maximum genus of a graph”. *Czecho-slovak Math. J.* 31(106).4 (1981), pp. 604–613.
- [27] Gerhard Ringel. “Über das Problem der Nachbargebiete auf orientierbaren Flächen”. *Abh. Math. Sem. Univ. Hamburg* 25 (1961), pp. 105–127.
- [28] Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, Inc., Upper Saddle River, NJ, 1996, xvi+512 pages.
- [29] Nguyen Huy Xuong. “How to determine the maximum genus of a graph”. *J. Combin. Theory Ser. B* 26.2 (1979), pp. 217–225.