EMBEDDING DIGRAPHS ON ORIENTABLE SURFACES

C. PAUL BONNINGTON, MARSTON CONDER, PATRICIA MCKENNA, AND MARGARET MORTON

DECEMBER 3, 1998

ABSTRACT. We consider a notion of embedding digraphs on orientable surfaces, applicable to digraphs in which the indegree equals the outdegree for every vertex, i.e., Eulerian digraphs. This idea has been considered before in the context of "compatible Euler tours" or "orthogonal A-trails" by Andsersen at al [1] and by Bouchet [4]. This prior work has mostly been limited to embeddings of Eulerian digraphs on predetermined surfaces, and to digraphs with underlying graphs of maximum degree at most 4. In this paper, a foundation is laid for the study of all Eulerian digraph embeddings. Results are proved which are analogous to those fundamental to the theory of undirected graph embeddings, such as Duke's Theorem [5], and an infinite family of digraphs which demonstrates that the genus range for an embeddable digraph can be any nonnegative integer is given. We show that it is possible to have genus range equal to one, with arbitrarily large minimum genus, unlike in the undirected case. The difference between the minimum genera of a digraph and its underlying graph is considered, as is the difference between the maximum genera. We say that a digraph is upper-embeddable if it can be embedded with 2 or 3 regions, and prove that every regular tournament is upper-embeddable.

1. INTRODUCTION

Graph embeddings and their generalisations have been studied by many authors over the years. For a survey of results in topological graph theory, the reader is referred to an article by Archdeacon [2]. Fundamental to the subject has been the study of the maximum and minimum orientable genus of a graph, problems which have been proved polynomial (Furst et al [8]) and NP-complete (Thomassen [16]), respectively.

In this paper, we consider 2-cell embeddings of loopless digraphs on compact connected orientable two-manifolds or *surfaces*, as we'll call them. Research in this area has until now been restricted to embeddings with exactly two regions; this has been explored by Kotzig [13], Las Vergnas [14], Bouchet [4], and Andersen et al [1]. Much of this literature focuses on digraphs with indeg(x)=outdeg(x) = 2 for each vertex x, and the 2-region embeddings are on surfaces as we've defined them. For larger degrees, the 2-region embeddings are sometimes on pseudosurfaces, a topic which we do not address in this work. More comparisons between our results and previous findings are made in the section on upper-embeddable digraphs.

We consider only Eulerian digraphs with indeg(x) = outdeg(x) ≥ 2 for each vertex x. It is not necessary to consider digraphs which contain a vertex x with indeg(x) = outdeg(x) = 1, since such a vertex and its arcs can simply be replaced by one arc from the in-neighbour of x to the out-neighbour of x. Hence a digraph $D = (V,A)$ is an *embeddable digraph* if for every $x \in V$, indeg(x) = outdeg(x) ≥ 2, and the graph which underlies D is connected. Note that our decision to exclude digraphs which contain at least one vertex x with indeg(x) =

The authors acknowldege support from the Auckland University Research Committee and the NZ Marsden Fund.

outdeg(x) = 1 results in the discussion of some digraphs in Section 2.4 which have multigraphs as their underlying graph. However, the theory (other than Proposition 4) does not change if such vertices are allowed, since troublesome arcs can be split into two in order to avoid a multigraph as the underlying graph, without changing the embedding properties of the digraph.

By an *embedding* of a digraph on a surface, we mean that the arcs and vertices of the digraph are placed on the surface, with arcs meeting only at mutually incident vertices in such a way that the orientation of a region is consistent with the orientation of the arcs which make up its boundary, explaining the restriction to orientable surfaces. As with graph embeddings, the *regions* of an embedding are the components of the complement of the digraph on the surface; with 2-cell embeddings each such component is homeomorphic to an open disc. The term "faces" is saved for a specific type of region, as explained below.

Vertex rotation schemes are employed to represent embeddings. The condition that arc directions on region boundaries be consistent forces in-neighbours and out-neighbours to alternate in the rotation scheme for each vertex, hence the requirement that indeg(x) = outdeg(x) for each vertex x. Rotation schemes which do not have in-neighbours and out-neighbours alternating in this way at every vertex will not be discussed, as they serve no purpose in this setting. It is useful to note that a rotation scheme, together with the orientation of the surface, yields the set of regions of an embedding. In the context of embeddable digraphs, each arc is on the boundary of exactly two regions, one we call a *face* (this uses the arc in the forward direction) and the other we call an *antiface* (each arc is traversed against its given orientation).

Let the genus of a surface S be denoted $\gamma(S)$. The genus, $\gamma(D)$, and maximum genus, $\gamma_M(D)$, of an embeddable digraph D, are the smallest and largest respectively of the numbers $\gamma(S)$ for surfaces S on which D can embed. The difference between these two numbers is called the *genus range* of the digraph. This notation is the same as that used in the undirected case. For a rigorous development of these underpinnings in the context of undirected graphs, the reader is referred to the book by Gross and Tucker [11]. The particulars which pertain to embeddable digraphs are perhaps best illustrated with an example.

Example 1. Figure 1 shows an embeddable digraph (in fact, a regular tournament) and a rotation scheme. Notice that the rotation at any vertex is an alternating list of in-neighbours and out-neighbours of the vertex. The seven faces and seven antifaces listed beside the figure are dictated by the given rotation scheme. Euler's formula ($|V|-|A|+|R| = 2-2g$) shows that the given embedding is on the surface of genus 1, the torus.

We make use of some informal language in order to make the paper more readable. For example, we say that a region "visits" or "touches" a vertex x if x is on the boundary of the region; we might also say that x is "on" the region. The "regions about a vertex x" are those regions which have x on their boundaries. A "corner" of a region consists of two consecutive arcs of the region's boundary. If a region visits a particular vertex more than once, we are sometimes interested in the "distance between two particular consecutive occurrences of the vertex on the boundary of the region"; by this we mean the number of consecutive arcs on the boundary of the region between the two occurrences of the vertex. If the vertex is on the region exactly twice, then we use the shorter of the two distances. In the tournament section, we are concerned with the number of antifaces which have a particular vertex on their boundary, and call the vertex a two-antiface vertex (for example) if this number is two. Any other informal language used is self-explanatory or explained in context.

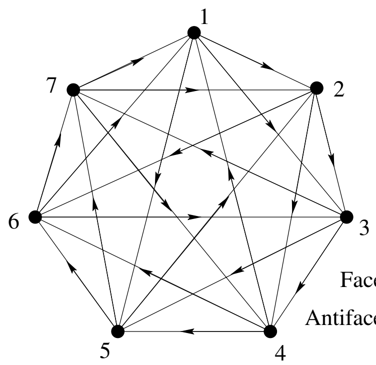

1: 5 6 2 4 3 7
2: 6 7 3 5 4 1
3: 5 2 7 1 4 6
4: 1 2 5 7 6 3
5: 2 3 6 1 7 4
6: 3 4 7 2 1 5
7: 2 6 4 5 1 3

Faces: {1,2,6}, {1,3,4}, {1,5,7}, {2,3,7}, {2,4,5}, {3,5,6}, {4,6,7}
Antifaces: {2,1,4}, {3,1,7}, {5,1,6}, {3,2,5}, {6,2,7}, {4,3,6}, {5,4,7}

FIGURE 1. *The given rotation scheme yields an embedding on the torus with the listed faces and antifaces.*

The paper is organised as follows. In Section 2, we prove several basic results on digraph embeddings which are analogous to results from the undirected case. The focus of Section 3 is a special class of embeddable digraphs, namely regular tournaments, and we prove that all regular tournaments are upper-embeddable. Finally, in Section 4 we mention some of the directions which further research on this topic might follow.

2. EMBEDDING DIGRAPHS

In this section we present some fundamental results on digraph embeddings which parallel work done in the undirected case and whose justifications use similar proof techniques. In addition, some natural questions are raised which are particular to the directed case. We offer them here for the sake of completeness.

2.1. **Parity.** The following proposition is justified by Euler's formula for graph embeddings ($|V|-|E|+|R|=2-2g$), since an embedding of an embeddable digraph is an embedding of the underlying graph.

Proposition 1. *If $D=(V,A)$ is an embeddable digraph, then for any rotation scheme σ of D*

$$|V|-|A|+|R|=2-2g$$

where $|R|$ is the number of regions of the embedding and g is the genus of the embedding surface.

It follows from Proposition 1 that the numbers of regions in two distinct embeddings of an embeddable digraph have the same parity.

2.2. **Adjacent Embeddings: New rotation schemes via minimum change.** Given the rotation scheme σ_1 for an embeddable digraph D, we can find another embedding of D by creating the rotation scheme σ_2 as follows. Choose one vertex v of D, and switch the position of exactly two in-neighbours or exactly two out-neighbours of v in the row corresponding to v of σ_1; the resulting rotation scheme is σ_2. If one rotation scheme can be obtained from another in this manner, then the two corresponding embeddings are said to be *adjacent*.

Proposition 2. *Suppose the rotation scheme σ_1 for an embeddable digraph D yields an embedding with f regions. If the rotation scheme σ_2 yields an adjacent embedding, then the latter embedding has f, $f+2$, or $f-2$ regions.*

Proof: Suppose the two rotation schemes differ only at vertex x, and assume without loss of generality that two outarcs at x have been switched. Assume that the number of regions at x in the original embedding σ_1 is m, and recall that indeg(x) = outdeg(x) ≥ 2. Figure 2 shows the (without loss of generality) three possibilities for the number of regions which involve arcs (x,b) and (x,e). The top row of the figure shows the arrangements of faces and antifaces around vertex x before switching these two outarcs, while the bottom row of the figure shows the corresponding arrangements of faces and antifaces about x after the switch.

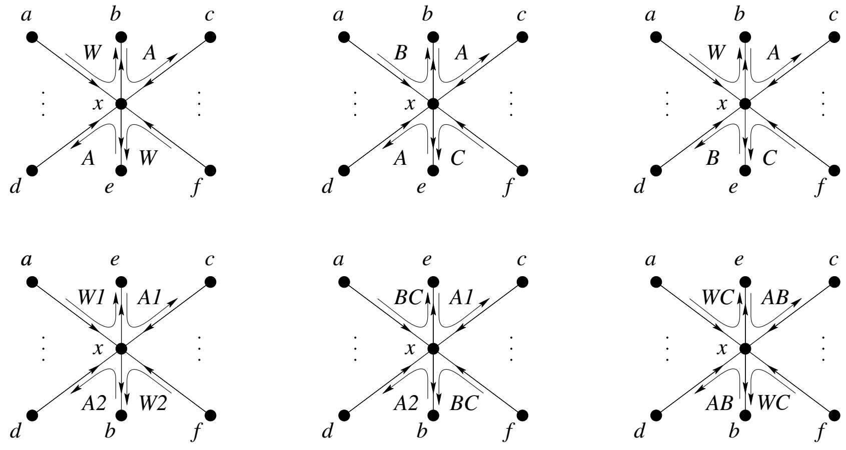

FIGURE 2. *The three possible arrangements of faces and antifaces about x before and after switching arcs (x,b) and (x,e).*

In the first scenario, arcs (x,b) and (x,e) are on exactly one face (W) and one antiface (A). After switching these two outarcs, arcs (x,b) and (x,e) are on exactly two distinct faces ($W1$ and $W2$) and two distinct antifaces ($A1$ and $A2$), and it follows that the number of regions which visit vertex x in the new embedding is $m+2$.

Next we consider the possibility that arcs (x,b) and (x,e) are on exactly two faces (B and C) and one antiface (A). In this case, faces labeled B and C before the switch become one face (BC) after the switch, while the antiface labeled A splits into two antifaces ($A1$ and $A2$). This gives a new embedding in which the number of regions at vertex x is m.

Lastly, if the two arcs are on two distinct faces (A and B) and two distinct antifaces (W and C) before the switch, then they are each on just one face (AB) and one antiface (WC) after the switch, giving $m-2$ regions at vertex x after the switch.

The three cases yield adjacent embeddings of D with $f+2$, f, and $f-2$ as the number of regions, respectively. ■

2.3. **An analogue to Duke's Theorem** [5]. Let $\{i, i+1, \ldots, i+j\}$ be the set of genera of the orientable surfaces on which a particular embeddable digraph D embeds; we call this the *genus list* of D, and denote it $GL(D)$.

Proposition 3. *Let D be an embeddable digraph. Then the genus list $GL(D)$ is an unbroken interval of integers.*

Proof: Let σ_1 and σ_2 be two rotation schemes for D. We know from Proposition 2 and basic knowledge of permutations that σ_1 can be obtained from σ_2 by a sequence of switches of pairs of outarcs and switches of pairs of inarcs at the vertices of D. Hence there exists a sequence of adjacent embeddings of D beginning with one on the surface of genus $\gamma(D)$ and ending with one on the surface of genus $\gamma_M(D)$. By Proposition 2, the numbers of regions in two adjacent embeddings of D differ by two or zero, so adjacent embedding surfaces differ by at most one in genus. The result follows. ■

2.4. **Upper-Embeddable Digraphs.** It is impossible to embed with 1-region any embeddable digraph, since each arc must be on one face and one antiface. If an embeddable digraph D has an embedding with exactly two regions on the surface of genus p, then $\gamma_M(D) = p$. Two Euler circuits in an embeddable digraph are said to be *compatible* if they have no pair of consecutive arcs in common. We give a necessary and sufficient condition for a 2-region embedding to exist if indeg(x) = outdeg(x) = 2 for all vertices x. The condition is closely related to work done by Andersen et al. [1] on orthogonal A-trails (which are defined in relation to a particular embedding), and perhaps explains their interest in the underlying graph having valency 4. See Figure 3 for an example of two compatible Euler circuits in a digraph.

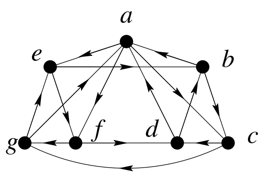

Two compatible Euler circuits:
a,c,d,b,c,g,e,f,g,a,e,b,a,f,d,a
a,e,f,d,b,a,c,g,a,f,g,e,b,c,d,a

A rotation scheme which has the two compatible Euler circuits as boundaries of regions of the embedding:

a: d c b f g e
b: d c e a
c: a d b g
d: c b f a
e: g f a b
f: e g a d
g: c e f a

FIGURE 3. *A two-region embedding means the digraph has two compatible Euler circuits, but the converse is not necessarily true.*

Proposition 4. *Let D be an embeddable digraph D with indeg(x) = outdeg(x) =2 for every vertex x. D is embeddable with exactly two regions if and only if D has a pair of compatible Euler circuits.*

Proof: If D has a two-region embedding, then the two regions are an Euler antiface and an Euler face. Since every vertex of D has indegree 2, these two regions do not have any corners made up of the same two arcs. Hence the two Euler circuits are compatible.

For sufficiency, let one of the Euler circuits determine the rotation scheme for the digraph. In the resulting embedding, the two Euler circuits enclose the two regions. ■

Note that irrespective of the degrees of the vertices of an embeddable digraph D, the existence of a pair of compatible Euler circuits is a necessary condition for the existence of a 2-region embedding in a surface, but it is not a sufficient condition. A counterexample is the rotational tournament on nine vertices in which each vertex v_i has outset $\{v_{i+1}, v_{i+2}, v_{i+3}, v_{i+4}\}$ (with additon modulo 9). The number of regions in an embedding of this tournament has to be odd, yet the tournament does have a pair of compatible Euler circuits.

As a result, in the general case (with larger degrees), finding compatible Euler circuits does not help with the embedding problem. This is unfortunate, since a substantial body of research exists on the problem of finding sets of pairwise compatible Euler circuits (see for example Fleischner et al. [6], Fleischner and Jackson [7] and Jackson [12]). Instead, it is necessary to find pairs of Euler circuits which which have the added property that the two circuits can arise as face and antiface of a 2-region embedding of the digraph.

We say an embeddable digraph is *upper-embeddable* if it can be embedded with two regions or with three regions. It is not true that every embeddable digraph is upper-embeddable. A counterexample is the "bracelet" digraph with an even number of "beads", or the even bracelet digraph. A bracelet digraph is a directed cycle with the reversal of each arc added; we are interested in those bracelet digraphs which have an even number of vertices (and, hence an even number of 2-cycles or "beads"). Since there are just two distinct rotation schemes at each vertex, colouring the vertices black and white is an easy way to indicate which rotation scheme is in use at each vertex for a given embedding. If a vertex is coloured black we use an anticlockwise rule, while at white vertices we use a clockwise rule. Figure 4 shows a drawing of the bracelet digraph on six vertices, with a listing of the arcs of the regions dictated by the given rotation scheme.

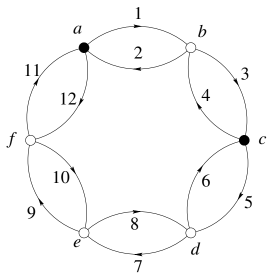

Faces:
{1,3,4,2}, {5,7,9,11,12,10,8,6}

Antifaces:
{1, 11, 12, 2}, {3, 4, 6, 5}, {7, 8}, {9, 10}

FIGURE 4. *The bracelet digraph on six vertices.*

Proposition 5. *The bracelet digraph on* $2k$ *vertices has maximum genus* 1.

Proof: Consider the bracelet digraph on $2k$ vertices drawn in the standard way (see Figure 4), with black vertices indicating anticlockwise schemes and white indicating clockwise schemes. If all vertices are white, then the embedding is in the plane, as given. The same is true if all vertices are black. Assume there is at least one vertex of each colour. Then the number of faces is equal to the number of black vertices and the number of antifaces is equal to the number of white vertices. Hence the total number of regions equals $2k$, and by Euler's formula, the genus of such an embedding is 1. ■

It is interesting to note that the argument in the proof above indicates that the embedding distribution of the bracelet digraph on $2k$ vertices consists of two embeddings of genus 0, and $2^{2k}-2$ embeddings of genus 1.

2.5. **Genus range.** The following family of embeddable digraphs shows that the *genus range*, or difference between maximum genus and minimum genus, for an embeddable digraph can be arbitrarily large. We call the family of digraphs "pinwheel digraphs", and denote the pinwheel digraph on $2k$ vertices ($k \geq 3$) as PW_k. The digraph PW_k on vertices labeled 0, 1, ..., $2k-1$ consists of two directed k-cycles, $0,2,4,\ldots,2k-2,0$ and $1,3,5,\ldots,2k-1,1$, with the additional arcs $(i,i+1)$ for all even i and $(i,i-3)$ for all odd i, with subtraction modulo $2k$. Consequently, PW_k has $4k$ arcs. A planar drawing of the pinwheel digraph PW_6 is shown in Figure 5 (the rotation scheme at each vertex is anticlockwise).

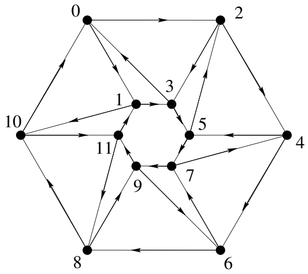

FIGURE 5. *The pinwheel digraph PW_6.*

Proposition 6. *The pinwheel digraph PW_k has minimum genus* 0 *and maximum genus k.*

Proof: It is obvious that PW_k embeds on the plane. In order to prove that it embeds on the surface of genus k and no surface with greater genus, we give a 2-region embedding of the digraph. The two regions of such an embedding are enclosed by two compatible Euler circuits. We achieve this with the following two Euler circuits of PW_k:

$$0,2,4,\dots,2k-2,2k-1,2k-4,2k-3,2k-6,\dots,3,0,1,3,5,\dots,2k-1,1,2k-2,0,2$$

and

$$0,2,3,5,2,4,5,7,4,6,7,\dots,2k-4,2k-2,2k-1,1,2k-2,0,1,3,0,2$$

Using this construction and Proposition 4, we see that the pinwheel graph PW_k has a 2-region embedding, since we can always find two compatible Euler circuits. It remains to calculate the genus g of such an embedding. For this we use the Euler formula $|V|-|A|+|R|=2-2g$, and find that $g=k$. ∎

Payan and Xuong [15] proved that the genus range of a graph exceeds one whenever the genus exceeds one. It is natural to wonder if the same result (or something similar) is true for embeddable digraphs. The answer is no. Each member of the next family of digraphs we discuss has genus range one, while the genus is at least two. Further, this family shows that there is no analogous "special" minimum genus in the directed case, since the minimum genus of one of these digraphs can be arbitrarily large, with embedding range remaining one.

We direct the edges of a particular family of circulant graphs to obtain the aforementioned embeddable digraphs, which we'll call *spoke digraphs*. For each digraph, the number $n \geq 5$ of vertices is odd, and the vertices are labeled with the integers 0 through $n-1=2k$. For each vertex i, there is an arc from i to $i+1$ and from i to $i+\frac{n+1}{2}$, with addition modulo n. Consequently there are $2n$ arcs and the shortest directed circuits have length $\frac{n+1}{2}$. The spoke digraph on 11 vertices is shown in Figure 6.

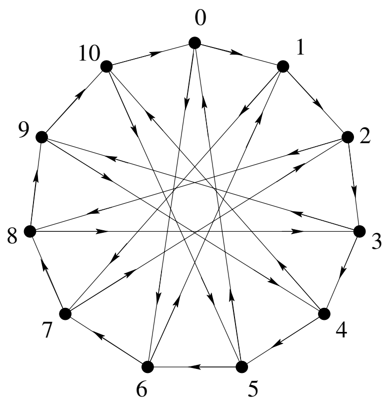

FIGURE 6. *The spoke digraph on eleven vertices.*

Proposition 7. *The spoke digraph on $n = 2k+1$ vertices has maximum genus k and minimum genus $k-1$.*

Proof: Euler's formula, together with the fact that the genus list is an unbroken interval of integers (Proposition 3), convinces us that it is sufficient to show that each spoke digraph has a 3-region embedding, a 5-region embedding, and no 7-region embedding.

We begin by showing that the spoke digraph on n vertices does not have a 7-region embedding. To this end, suppose that it does have a 7-region embedding; this embedding has at least four faces or four antifaces. Suppose without loss of generality that the embedding has at least four faces. The digraph has $2n$ arcs, each of which is counted exactly once in the sum of the face lengths. Hence the average face length is at most $\frac{2n}{4} = \frac{n}{2}$. Since n is odd, there aren't any faces of this length, so there must be at least one of length less than or equal to $\frac{n-1}{2}$. This is a contradiction, since the shortest directed circuit in a spoke digraph on n vertices is of length $\frac{n+1}{2}$.

If the digraph is drawn in standard form, as is the one on 11 vertices in the figure, we have a 5-region embedding of the digraph if vertices 0, $\frac{n-1}{2}$, and $\frac{n+1}{2}$ are coloured black (anticlockwise), and all others are coloured white (clockwise). For a 3-region embedding, colour vertices 0 through $\frac{n-1}{2}$ black and all remaining vertices white. ■

2.6. **Embeddings of the underlying graph.** It is an interesting pursuit to compare the genus and maximum genus of an embeddable digraph D with the genus and maximum genus respectively of the graph G which underlies D. Certainly $\gamma(G) \leq \gamma(D) \leq \gamma_M(D) \leq \gamma_M(G)$. It is natural to wonder how big the difference can be between $\gamma(G)$ and $\gamma(D)$ or between $\gamma_M(D)$ and $\gamma_M(G)$. In this section we give families of embeddable digraphs which demonstrate that these two differences can be arbitrarily large.

The cartesian product of the directed cycle C_n with itself (see Figure 7 for $C_4 \times C_4$) gives a family of embeddable digraphs which demonstrates that the difference between $\gamma(G)$ and $\gamma(D)$ can be arbitrarily large. In this case, the undirected graph has genus 1, while the genus of the directed graph grows without bound as n grows.

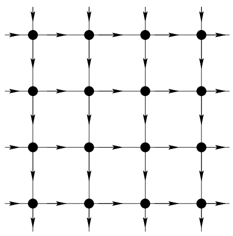

FIGURE 7. *The cartesian product of the directed cycle C_4 with itself, $C_4 \times C_4$. (Note: partial arcs "connect in the back" as expected; also, inarcs do not alternate with outarcs in this drawing, so it does not represent an embedding of the digraph.)*

Proposition 8. *The cartesian product $D = C_n \times C_n$ of the directed cycle C_n with itself has genus equal to $\frac{n^2-3n+2}{2}$.*

Proof: If we use the rotation scheme which is illustrated in Figure 8 (for $C_4 \times C_4$) on $C_n \times C_n$, then the faces are $2n$ directed cycles of length n, while the antifaces are n directed cycles of length $2n$. To see that this embedding is a minimum-genus embedding, note that it makes optimal use of the shortest directed circuits in the digraph. The digraph has exactly $2n$ directed circuits of length n, and these are the shortest in the digraph. The second shortest directed circuit length is $2n$, confirming that the embedding is a minimum genus embedding. Since the embedding has $n + 2n = 3n$ regions total, Euler's formula ($|V| - |A| + |R| = 2 - 2g$) convinces us that the genus of $C_n \times C_n$ is $\frac{n^2-3n+2}{2}$. ■

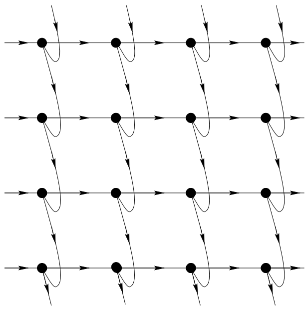

FIGURE 8. *The rotation scheme depicted yields a minimum genus embedding of $C_4 \times C_4$, if a clockwise rule is used at every vertex.*

The even bracelet digraphs of Section 2.4 provide the needed evidence that the difference between $\gamma_M(D)$ and $\gamma_M(G)$ can be arbitrarily large. We proved in Proposition 5 that the maximum genus of the bracelet digraph on $2k$ vertices is 1. However, the maximum genus of the underlying multigraph is k. To see this, use the generalisation of the embedding scheme shown in Figure 9 for the bracelet graph on six vertices. This scheme yields a 2-face embedding of the graph, and Euler's formula gives the desired result.

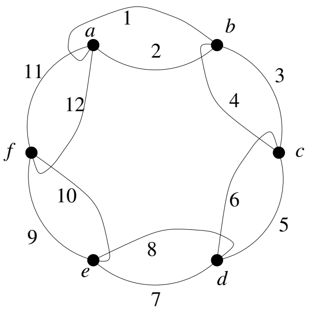

Regions:
{1, 3, 5, 7, 9, 11, 2, 4, 6, 8, 10, 12}
{1, 11, 10, 8, 5, 4, 2, 12, 9, 8, 5, 4}

FIGURE 9. *The undirected bracelet graph on six vertices has a two-face embedding.*

3. EMBEDDING TOURNAMENTS

We have found that every regular tournament is upper-embeddable and present these results here. Note that a regular tournament has an odd number of vertices. The question of minimum genus of a regular tournament on n vertices is one worthy of further study, and is closely related (for relevant congruence classes of n (mod 12)) to face-2-colourable triangular embeddings of the complete graph on n vertices, K_n, which have been studied by Grannell et al. [9] and Bonnington et al. [3]). The latter work yields an exponential family of regular tournaments T with $\gamma(T) = \gamma(K_{|V(T)|})$. The question of the maximum genus of a regular tournament is answered in Theorem 1.

Theorem 1. *Every regular tournament is upper-embeddable.*

A computer check shows that regular tournaments on five and seven vertices are upper-embeddable, so we restrict our discussion to the regular tournaments on nine or more vertices. We prove four technical lemmas before proving Theorem 1.

Lemma 1. *If T is a regular tournament on $n \geq 9$ vertices, then there exists an embedding of T which has each vertex on one face and at most two antifaces.*

Proof: Let E be an Euler circuit of T, and σ a rotation scheme induced by E. The following describes how to perform a series of switches to σ, to create a rotation scheme σ' which preserves E as one face of the embedding, and has every vertex on at most three regions total. Note that E is incident with each vertex $\frac{n-1}{2}$ times, alternating with antifaces.

Given the rotation scheme (T, σ), consider a vertex v which is on face E and m antifaces, where $m \geq 3$. If corners of three antifaces A, B, and C occur consecutively around vertex v in the embedding (alternating with corners of face E), as shown in Figure 10, then simultaneously switching b with d and c with e in the rotation scheme at v yields an embedding with two fewer antifaces, since A, B and C become one antiface.

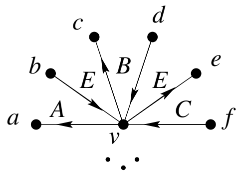

FIGURE 10. *The corners of three antifaces A, B, and C occur consecutively around vertex v in the embedding, alternating with corners of face E.*

Repeat the above process until there are no three distinct antifaces which occur consecutively at v.

Let A, B, and C be three distinct antifaces, each of which has vertex v on its boundary. We choose a pair of corners of E at vertex v which have all three antifaces A, B, and C as neighbouring antifaces. There are three possibilities (see Figure 11), all of which can be rearranged so that the three distinct antifaces become one, while the Euler face E is preserved.

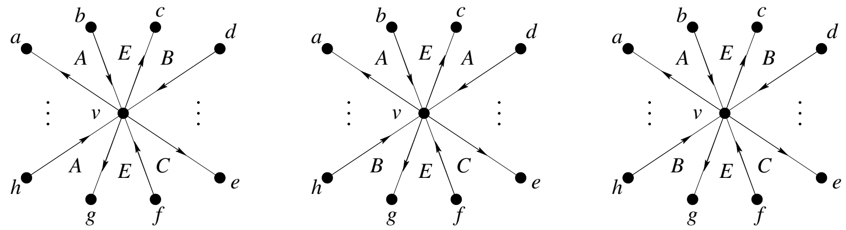

FIGURE 11. *The three possibilities for three distinct antifaces and two corners of face E at vertex v.*

In the first two cases shown above, switching the vertex b with f and c with g in the rotation at v yields an embedding with two fewer antifaces. In the third case, the aforementioned switches would yield an embedding with the same number of antifaces as the given rotation, so a different switch is required to reduce the total number of antifaces at vertex v. In this case note the following: the vertices y and z which immediately precede a in the clockwise rotation scheme for vertex v are distinct from g and h respectively (otherwise three distinct antifaces occur consecutively at v). The clockwise rotation scheme for v contains the triple $y\ z\ a$. Together the arcs (v,y) and (z,v) form a corner of either antiface A or B (not C, since this gives three distinct consecutive antifaces at vertex v), and in either case we have one of the first two scenarios of the figure above. Hence by switching y with g and z with g in the rotation scheme for vertex v, we reduce the number of antifaces at vertex v by two.

By employing the above techniques at every vertex, we construct an embedding of T in which each vertex appears only on face E and on at most two antifaces. ■

We must reduce the total number of regions in the embedding. To this end, Lemma 2 convinces us of the existence of a "special" antiface, and Lemma 3 shows that there are two "special" vertices on that antiface. Finally, Lemma 4 gives a method for making use of these two vertices to reduce the total number of regions in the embedding.

Lemma 2. *Let T be a regular tournament on $n \geq 9$ vertices with rotation scheme σ. If every vertex is on at most three regions (one Euler face and at most two antifaces) and there are four or more antifaces in all, then there exists an antiface which visits every vertex at least three times.*

Proof: Suppose we have distinct antifaces A, B, C, and D. Assume there exists a vertex which is on both antiface A and antiface B; we call such a vertex an AB vertex and say that AB is its "type". Existence of an AB vertex implies there are no CD vertices, since the digraph is a tournament. Since antiface C exists, we must have either a vertex of type AC or BC.

<u>Case i</u> Suppose there exists a vertex of type AC. Then there are no vertices of type BD. Since D is an antiface, there must be vertices of type AD, and therefore no vertices of type

BC. Hence the only vertex types are *AB*, *AC*, and *AD* (moreover, there are at least three vertices of each type), plus possibly type *A* vertices (which are only on one antiface, *A*). Antiface *A* visits every vertex at least once.

<u>Case ii</u> Suppose there exists a vertex of type *BC*. Then arguments similar to those in Case i lead us to conclude that antiface *B* visits every vertex at least once.

Without loss of generality, suppose antiface *A* visits every vertex at least once, so that the only possible vertex types for vertices on two antifaces are *AB*, *AC*, and *AD*. Let x be a vertex of type *AB*. Since T is a tournament, the neighbourhood of x (consisting of both in-neighbours and out-neighbours) contains at least three vertices of type *AC* and at least three vertices of type *AD*. We consider the number of times antiface *A* visits vertex x. The worst case scenario is that the three guaranteed neighbours of x which are of type *AC* and the three guaranteed neighbours of x which are of type *AD* occur consecutively (in any order) in the rotation scheme for x as shown in Figure 12, in which case we see that antiface *A* visits vertex x at least three times. ■

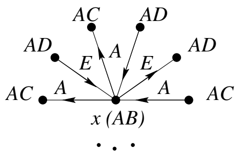

FIGURE 12. *The worst-case scenario: the six vertices (three of type AC and three of type AD) occur consecutively in the rotation scheme at x. We see that antiface A visits vertex x at least three times in this case.*

Lemma 3. *Let T be a regular tournament on $n \geq 9$ vertices with rotation scheme σ. If every vertex is on at most three regions (one Euler face and at most two antifaces) and there exist four or more antifaces in total, then there exists a pair of vertices α and β which are on distinct antifaces, and are interlaced on the boundary of a third antiface which visits each vertex of the tournament at least three times.*

Proof: Choose a vertex α which is on two distinct antifaces, such that between some two occurrences of α on the boundary of antiface *A* lies a vertex β of a type different from α. Hence β is on the boundary of an antiface which does not touch vertex α. Further conditions for the choice of α are that the "distance" between the two occurrences of α on the portion of the boundary of antiface *A* which contains β is minimised.

Suppose α is of type *AB* and β is of type *AC*. There are at least two more occurrences of vertex β on the boundary of antiface *A*, since *A* visits each vertex at least three times. If one of them occurs outside of the two previously identified occurrences of α, then we have an"interlacing" $\ldots\alpha\ldots\beta\ldots\alpha\ldots\beta\ldots$ as desired. If not, then all copies of β on the boundary of antiface *A* occur between the two described occurrences of α. Consider the two "extreme" β's in the listing, i.e., the two which are closest to the two α's. Because of our choice of this pair of α's we know that all vertices listed between the two "extreme" β's are of the same type (*AC*) as β or are on just one antiface *A*. Hence, there are at most two arcs of the boundary of antiface *A* which have β on one end and a two-antiface vertex of a different type on the other end. This contradicts the fact that β has at least six two-antiface neighbours whose types are not *AC*. We conclude that α and β must be "interlaced" on the boundary of antiface *A*. ■

Lemma 4. *Every regular tournament T can be embedded with four or fewer regions.*

Proof: We have confirmed that the one regular tournament on five vertices and the three regular tournaments on seven vertices are upper-embeddable, so we restrict our discussion to regular tournaments on $n \geq 9$ vertices.

We begin with an embedding of T which satisfies the hypotheses of Lemmas 2 and 3. That is, the embedding has one Euler face, at least four antifaces, and at most two antifaces at each vertex. Let α and β be a pair of vertices of distinct type interlaced on the boundary of antiface A, which visits each vertex at least three times. The existence of antiface A is guaranteed by Lemma 2, and the existence of α and β is guaranteed by Lemma 3. Suppose that α is a type AB vertex and β is a type AC vertex. Then the listings for the boundaries of antifaces A, B, and C look as follows:

$A: \ \ldots a\alpha b \ldots c\beta d \ldots e\alpha f \ldots g\beta h \ldots$
$B: \ \ldots p\alpha q \ldots$
$C: \ \ldots r\beta s \ldots$

We find that in the rotation scheme for α, vertex a is followed by b, vertex e is followed by f and vertex p is followed by q. Similarly, in the rotation scheme for vertex β, vertex c is followed by d, vertex g is followed by h, and vertex r is followed by s.

Change the rotation scheme at vertex α so that vertex a is followed by f, vertex e is followed by q and vertex p is followed by b, while keeping each of these six vertices with its Euler "partner" at α from the original rotation scheme dictated by E. Similarly, arrange the rotation scheme at β so that vertex c is followed by h, vertex g is followed by s, and vertex r is followed by d, again keeping each of the six with its original Euler "partner" at β.

The result is a new antiface ABC, consisting of all of the pieces of antifaces A, B, and C. The listing for the boundary of this antiface is:

$ABC: \ \ldots a\alpha f \ldots g\beta s \ldots r\beta d \ldots e\alpha q \ldots p\alpha b \ldots c\beta h \ldots$

Notice that this technique works even when, for example, the vertex b listed in the rotation scheme at vertex α is the Euler "partner" at α of the vertex p. Consideration of the possible cases is left to the reader.

Given an embedding of T with one Euler face and at least four antifaces, having each vertex on at most three regions total, we can use the techniques presented here to find an embedding with two fewer antifaces. The resulting embedding has one Euler face (the same Euler face as before), and each vertex is still on at most three regions in total. Therefore it is possible to repeat the process as long as the resulting embedding has at least four antifaces. The process must stop when an embedding with four or fewer regions is achieved. ■

We make use of the rotation switch just described in the proof of Theorem 1. Note that if such a pair of interlaced vertices exists on the boundary of some antiface, then the described switch can be applied, reducing the total number of antifaces by 2. For convenience, we restate Theorem 1 here, and give its proof next. The maximum genus for K_n is $\lfloor \frac{1-n+\frac{n(n-1)}{2}}{2} \rfloor$ (Xuong [17]), a fact which is useful in the proof.

Theorem 1. *Every regular tournament is upper-embeddable.*

Proof: If $n \equiv 1 \pmod 4$, then the number of regions in any embedding of the tournament is odd, and the process described in the proof of Lemma 4 ends with an embedding of the tournament with exactly one face and two antifaces. Hence, if the number of vertices is equivalent to 1 (mod 4), then the tournament is upper-embeddable.

The case with $n \equiv 3 \pmod 4$ requires more work, since the process described in the proof of Lemma 4 results in an embedding of the tournament with one Euler face and at most three antifaces; also, there may be times when the process cannot be implemented at all, due to an early shortage of antifaces (i.e., three antifaces from the beginning).

Suppose we have an embedding of T with exactly one Euler face E and three antifaces A, B, and C. We can assume that each vertex is on at most two antifaces by Lemma 1. Then the six possible vertex types are: AB, AC, BC, A, B, and C.

Suppose one of AB, AC, BC is not a vertex type, say BC. Then there are vertices of type AB (at least three), AC (at least three) and possibly A. Since an AB vertex x has at least three AC neighbours, antiface A visits x at least twice. Similarly, antiface A visits each AC vertex and each A vertex at least twice. Using the techniques of Lemma 3, we find an interlaced pair of 2-antiface vertices of distinct type on the boundary of antiface A. If we switch their rotation schemes as described in the proof of Lemma 4, then the resulting embedding has exactly two regions.

If two of AB, AC, BC are not vertex types, say AC and BC, then there are vertices of type AB and C only, a contradiction.

Suppose we have at least one vertex of each type, AB, AC, BC. Then we have no vertices of type A, B, or C. We claim that there exists a pair P, Q from the set $\{AB, AC, BC\}$ such that $|P| \geq 2$ and $|Q| \geq 3$. To see this, suppose not. Assume for all pairs P, Q from the set $\{AB, AC, BC\}$ either $|P| < 2$ or $|Q| < 3$. Consider the pair AB, AC; one of these is of size less than 3. Suppose $|AB| < 3$. Consider the pair BC, AC; one of these is of size less than 3. Suppose $|AC| < 3$.

- If $|AB| = |AC| = 1$ then antiface A has only two vertices, a contradiction.
- If $|AB| = 1$ and $|AC| = 2$, then since $n \geq 9$, $|BC| \geq 6$ and the sets AC and BC are the P and Q we seek respectively.
- If $|AB| = |AC| = 2$ then $|BC| \geq 5$, and the desired P and Q do exist.

It remains to show that there exists an interlaced pair of vertices of distinct types on the boundary of some antiface. Suppose $|AB| \geq 2$ and $|AC| \geq 3$. Let $\{x,z\} \subseteq AB$, $\{y,a,b\} \subseteq AC$, and $\{x'\} \subseteq BC$.

Since T is a tournament, there is an arc between x and y. This arc must be on the boundary of antiface A. We consider it to be the "first" arc in the boundary listing of antiface A.

$$A: \; x\, y\, \ldots$$

Suppose the boundary of antiface A has no interlaced pairs of vertices of distinct types. Then all appearances of y in the listing for the boundary of antiface A must be before the second appearance of x.

$$A: \; x\, y\, \ldots \text{ (all other copies of } y) \ldots\, x$$

Since y is adjacent to every vertex of AB, and each resulting arc is on the boundary of antiface A, all vertices of AB except possibly one, say z, must be listed between the first and last y in the boundary listing for antiface A.

$$A: \; x\, y\, \ldots \text{ (all other copies of } y \text{ and elements of } AB - \{x,z\}) \ldots\, y\, z \ldots\, x$$

Suppose $AB - \{x,z\}$ is nonempty; let $v \in AB - \{x,z\}$. Since there are no interlaced pairs on the boundary of antiface A, all copies of v must be between two consecutive y's (otherwise we interlace v and y). Hence vertices a and b appear between these two consecutive y's, since v is adjacent to both a and b on the boundary of antiface A. Also, a

and b appear to the right of the rightmost y, since they are both adjacent to x. So we have an interlacing of x and a or of x and b on the boundary of antiface A, a contradiction.

If $AB - \{x,z\}$ is empty, then $|AB| = 2$. There are three cases to consider.

- If both a and b are listed between the first and last y's on the boundary of antiface A, then vertex x is interlaced with one of them. This is because x is listed with each of them to the right of the rightmost y.
- If neither a nor b is listed between the first and last y's, then both arcs between y and $\{a,b\}$ are on the boundary of antiface C. In this case, we can argue as above, but with sets AC and BC (recall $x' \in BC$), and find an interlacing on the boundary of antiface C of vertices x' and a or of vertices x' and b.
- Suppose just one of a and b, say a, appears on the boundary of face A between the first and last y's.

  $A: \ x\, y \dots y\, a \dots y\, z \dots x$

  Since vertex a is adjacent to x and to z on the boundary of antiface A, vertices x and a are listed together to the right of the rightmost y. Also, z and a are listed together to the right of the rightmost y (otherwise y and z are interlaced on the boundary of antiface A). Every possibility for the arrangements of a with x and a with z to the right of the rightmost y yields an interlacing either of x with a, a with z, or of z with b.

We conclude that there exists an interlaced pair of vertices of distinct types on the boundary of some antiface. Using the rotation switch described in the proof of Lemma 4, we produce an embedding of T with exactly two regions. ∎

4. Future Work

Some of the open problems which have arisen from this research are listed below.

- Which tournaments on n vertices have genus $\lceil \frac{(n-3)(n-4)}{12} \rceil$, the genus of K_n?
- Characterise those embeddable digraphs which are upper-embeddable. Is there an analogue to the splitting tree result used to classify upper-embeddable graphs?
- Is the embedding distribution of an embeddable digraph always (strongly) unimodal, as is conjectured to be the case in the study of undirected graphs by Gross et al. [10]?

References

[1] L. D. Andersen, A. Bouchet, and B. Jackson. Orthogonal A-trails of 4-regular graphs embedded in surfaces of low genus. *J. Combin. Thy. (B)*, 66:232–246, 1996.

[2] D. Archdeacon. Topological graph theory: A survey. *Congressus Numerantium*, 115:5–54, 1996.

[3] C.P. Bonnington, M.J. Grannell, T.S. Griggs, and J. Širáň. Exponential families of non-isomorphic triangulations of complete graphs. Submitted.

[4] A. Bouchet. Maps and Δ-matroids. *Discrete Mathematics*, 78:59–71, 1989.

[5] R.A. Duke. The genus, regional number, and Betti number of a graph. *Canad. J. Math.*, 18:817–822, 1966.

[6] H. Fleischner, A. J. W. Hilton, and B. Jackson. On the maximum number of pairwise compatible Euler cycles. *Journal of Graph Theory*, 14(1):51–63, 1990.

[7] H. Fleischner and B. Jackson. Compatible Euler tours in Eulerian digraphs. *Cycles and Rays*, G. Hahn et al. (eds.):95–100, 1990.

[8] M. Furst, G. Gross, and L.A. McGeoch. Finding a maximum-genus graph imbedding. *J. Assoc. Comput. Mach.*, 35:523–534, 1988.

[9] M.J. Grannell, T.S. Griggs, and J. Širáň. Face 2-colourable triangular embeddings of complete graphs. submitted.

[10] J.L. Gross, D. Robbins, and T.W. Tucker. Genus distributions for bouquets of circles. *J. Combin. Thy. (B)*, 47:292–306, 1989.

[11] J.L. Gross and T.W. Tucker. *Topological Graph Theory*. Wiley-Interscience, New York, 1987.

[12] B. Jackson. Compatible Euler tours for transition system s in Eulerian graphs. *Discrete Mathematics*, 66:127–131, 1987.

[13] A. Kotzig. Eulerian lines in finite 4-valent graphs and their transformations. In Theory of Graphs (Erdös and Katona Eds.), Proceedings colloquium held at Tihany (Hungary, 1966). North-Holland, Amsterdam, 219–230, 1986.

[14] M. Las Vergnas. Eulerian circuits of 4-valent graphs imbedded in surfaces. Algebraic Methods in Graph Theory (L. Lovasz and V. Sos Eds), Proceedings colloquium held at Szeged (Hungary, 1978). North Holland, Amsterdam., 1981.

[15] C. Payan and N. H. Xuong. Graph theory newsletter. 8:1, 1978.

[16] C. Thomassen. The graph genus problem is NP-complete. *J. Algorithms*, 10:568–576, 1989.

[17] N. H. Xuong. How to determine the maximum genus of a graph. *J. Combin. Thy. (B)*, 26:217–225, 1979.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OFAUCKLAND, PRIVATE BAG 92019, AUCKLAND, NEW ZEALAND

E-mail address: `bonning@scitec.auckland.ac.nz`

E-mail address: `conder@math.auckland.ac.nz`

E-mail address: `mckenna@math.auckland.ac.nz`

E-mail address: `morton@math.auckland.ac.nz`