ROTATION SYSTEMS: THEORY AND APPLICATION

TOMAS J. BOOTHBY

ABSTRACT. This paper provides an introduction to the theory of *rotation systems*, a means of describing and computing with graphs embedded on surfaces. Then, we define a *direct connection* as a pairing of a finite number of points on opposite sides of a simple closed curve, according to some permutation. We then compute the minimal genus of a surface which is bounded by this curve and contains a set of non-intersecting curves which connect these pairs of points.

This paper is devoted to a question of drawing a certain kind of diagram on a surface without edge crossings. Specifically, one can visualize the action of a permutation on n letters by writing two rows of those letters, and drawing a line from each letter in the top row to its image in the bottom row; for example, the permutation $\tau = [254361]$ yields the diagram:

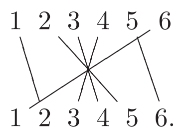

We will show that this diagram can be drawn on a genus 2 surface, and under certain reasonable restrictions, it cannot be drawn on a lower genus surface without crossing lines. In general, we find that one can draw the above diagram for a permutation τ on a surface with genus g, where g is equal to the minimal number of terms in a factorization of τ into *block interchanges*.

1. Introduction to Rotation Systems

Rotation systems are computational objects used to describe graphs embedded in orientable surfaces. There are a number of differing definitions of rotation systems; notably, Mohar (for example, in [7], denotes a rotation system $\pi = \{\pi_v : v \in V(G)\}$, where π_v is a cyclic permutation of the edges incident with v. The author finds the following definition to be easier to deal with from a computational standpoint, and we show equivalence to Mohar's definition in Remark 1.9.

Definition 1.1. Let D be a finite set with even order, $|D| = m = 2n$, and let $\nu, \varepsilon \in S_m$, where ε and $\phi := \nu\varepsilon$ are fixed-point free, ε is an involution, and $\langle \nu, \varepsilon \rangle$ acts transitively on D. We call the ordered pair (ν, ε) a *rotation system*, and the elements $d \in D$ *darts*.

We denote the orbits of the cyclic groups generated by ν, ε, and ϕ

$$V = \{v \subset D : \langle \nu \rangle \text{ acts transitively on } v\}$$
$$F = \{f \subset D : \langle \phi \rangle \text{ acts transitively on } f\}$$
$$E = \{e \subset D : \langle \varepsilon \rangle \text{ acts transitively on } e\}$$

and refer to the elements in V, F, and E as *vertices*, *faces*, and *edges* respectively. If an edge e is contained in a vertex, we call e a *loop*. Finally, we call a rotation

The author wishes to express his appreciation to the National Science Foundation for supporting in part this work through grant DMS-0354131.

system *simple* if it contains no loops, and if, when an edge e intersects vertices u and v, it is the only such edge.

Using the terms edge, vertex, and face in this way may appear to be abuse of notation, but we will promptly dissolve any ambiguity. For any rotation system, we can construct a graph embedded upon a surface whose vertices correspond to vertices in the rotation system; edges to edges, and faces to faces. Further, simple rotation systems correspond to simple graphs.

Non-simple graphs, then, are looped multigraphs. If we relax the requirement that ε is an involution, the correspondence is to hypergraphs. This is covered extensively in [3], and not of interest for the purpose of this paper.

1.1. **Constructing Embedded Graphs.**

Proposition 1.2. *Every edge intersects at most two vertices.*

Proof. Since $\langle \nu, \varepsilon \rangle$ acts transitively on D, each $e \in E$ must intersect at least one vertex. Also, since ε is a fixed-point free involution, $e = \{k, l\} \subset D$. Since the vertices are mutually disjoint, e can intersect at most two of them. Note that if (ν, ε) is a simple rotation system, every edge will intersect exactly two vertices. □

Proposition 1.3. *Every edge intersects at most two faces.*

Proof. Apply the the argument of Proposition 1.2, replacing ν with ϕ. □

Remark 1.4. Note that there exist seemingly degenerate cases in which an edge intersects a single face, but this does not pose any significant problems to us. For example, in Lemma 1.7, we associate a face with a "polygon with $|f|$ sides". The reader is, no doubt, comfortable with this situation when $|f| > 2$, but for the case when $|f| = 2$, we merely take the polygon to be a disc, and the "sides" to be halves of the boundary circle. When this face intersects only one edge, we identify these sides to construct a surface homeomorphic to the sphere.

Lemma 1.5. *For each rotation system (ν, ε), we can construct a labelled graph G, and the construction of G is well-defined.*

Proof. Label $D = \{1, 2, \ldots, m\}$, and order the vertices in V by $v_i < v_j$ if $\min v_i < \min v_j$, and label them in ascending order. Since $V = \{v_1, \ldots, v_n\}$ is a complete, disjoint partition of a finite set, this ordering is well-defined.

Now, we construct an $(n \times n)$ matrix, $X = (x_{ij})$ by

$$x_{ij} = |\{e \in E : e \cap v_i \neq \emptyset, e \cap v_j \neq \emptyset\}|,$$

and let G be the graph corresponding to the adjacency matrix X, and see that this construction is well-defined. □

Corollary 1.6. *In the case that an angle system (ν, ε) is simple, the associated graph is simple.*

Proof. By Proposition 1.2 and Lemma 1.5, it follows from our definition of a simple angle system that G will neither have loops, nor multiple edges. □

The following assumes a certain level of comfort with constructing surfaces by pasting pieces together. A detailed proof that this process works can be found in Chapter 6 of [6]. We differ from Lee in that we are not concerned with any differentiability. Specifically, we describe orientation as a concept of clockwise rotation – if moving from any point on the surface continuously does not change the direction of rotation, then the surface is orientable.

Lemma 1.7. *Every simple rotation system corresponds to a graph embedded upon a compact, orientable Riemann surface.*

FIGURE 1. Details of the construction used in Lemma 1.7 for (ν, ε) where $\nu = (7396)\cdots$, $\varepsilon = (17)(68)(95)(34)\cdots$, and $\phi = (274)\cdots$.

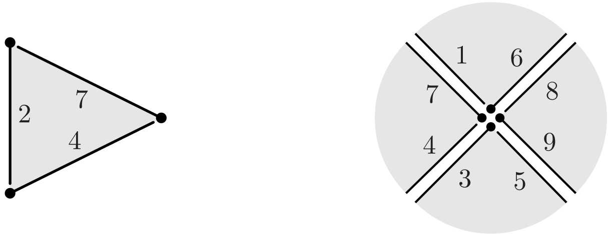

(a) $P_{\{2,7,4\}}$

(b) Polygons at a vertex before identification.

Sketch of proof. Let F be the faces of a simple rotation system, and to each $f \in F$, associate an $|f|$-sided planar polygon P_f, homeomorphic to a closed disk, and choose an orientation (that is, a direction of rotation to call clockwise). Now, separate the boundary of the polygon into intervals bounded at the corners; open at the counterclockwise end and closed at the clockwise end and label each by a dart in f according to the action of ϕ on f. We draw these intervals as "match sticks" as in Figure 2(a).

Then, we glue each side of P_f with label d to the side of the polygon P_g labeled $\varepsilon(d) \in g \in F$, scaling either side to match the other (or equivalently, we can require that each polygon is regular). Further, we require that the corners labelled d and $\varepsilon(d)$ are not identified. By Proposition 1.3, there is only one such g, and this procedure is well-defined.

We claim that the surface constructed by this procedure is a 2 dimensional, orientable compact topological manifold. The curious reader will find a detailed proof of this claim in [6]; we will merely describe our construction at the corners. As our goal is a natural correspondence between rotation systems and embedded graphs, the reader will expect that corners will correspond to vertices.

Note now that ε is an involution, so $\nu = \phi\varepsilon$. So, let v be an orbit of ν, with $d \in v$. Then, $\nu(d) = (\phi \circ \varepsilon)(d)$, the action of ε corresponds to crossing an edge, and the action of ϕ corresponds to walking clockwise around a face. Since the sides of a polygon agree with their clockwise corner, the corner labelled by $\nu(d)$ is on the face one step counterclockwise from that of d. See Figure 2(b) for an illustration of this

Next, we show orientability. Above, we defined the action of ϕ to be clockwise rotation in each face, so we consider what happens on the boundaries between faces. By requiring that the corners labelled by d and $\varepsilon(d)$ are not identified, we have forced consistency. If one looks at the sides of the polygon in Figure 2(a) as matches, the head of each match points clockwise – and this stays consistent over edges and vertices, as seen in Figure 2(b).

We conclude with an observation that our surface is the result of pasting polygons together. Note that we have identified sides to sides, and corners to corners, so that there is no boundary left over. The identified sides correspond to orbits of ε, and the identified corners of the polygons correspond to orbits of ν – that is, edges in the rotation system are realized as edges on our surface, and vertices in the rotation system are vertices on the surface. Brief inspection shows that this graph is precisely the one obtained in Lemma 1.5. □

Remark 1.8. Though the correspondence between rotation systems and embedded graphs is many-to-one, this procedure allows one to transfer between the objects without loss of generality.

Proof. We can run the construction backwards easily – given a graph embedded upon a surface, we slice the surface along edges, labelling darts as we go. We record the label pairs that correspond to edges to find ε, and the labels around vertices (in counterclockwise order!) to find ν. If we keep the original labels, we end up with the same rotation system. Otherwise, we can easily permute the labels to find the original. This leads to a strong notion of an isomorphism of rotation systems which will not be further discussed in this paper. □

Remark 1.9. Finally, we note that we need not start with a surface. As in Remark 1.8, we can construct a rotation system from a graph merely by assigning labels to coincident edge-vertex pairs, and picking an ordering of these labels at each vertex. By Lemma 1.7, any rotation system obtained in this way yields an embedding of our graph. When a graph is simple, this is equivalent to choosing a cyclic permutation of the edges incident to each vertex.

Lemma 1.10. *Let (ν, ε) be a rotation system, and $\mathcal{R}$ the surface constructed in Lemma 1.7. Then, the genus of $\mathcal{R}$ is*

$$g = 1 - \frac{1}{2}(|V| - |E| + |F|).$$

Proof. See Lemma 4.1.4 in [8]. □

2. Application of Rotation Systems

It is trivial to use rotation systems to compute the minimal genus of a graph. That is, one generates all possible dart orderings around vertices, and applies Lemma 1.10, whence the genus of a graph G is

$$g = \min\left\{1 - \frac{1}{2}(|\langle\nu\rangle| - |\langle\varepsilon\rangle| + |\langle\nu\varepsilon\rangle|) : (\nu, \varepsilon) \text{ a rotation system of } G\right\}.$$

A naïve implementation of this will require considering $(\deg(v) - 1)!$ orderings at every vertex, so the expected runtime is a product of factorials.

In [7], Mohar presents an algorithm to determine if a graph G can be embedded on a fixed surface of genus g in $O(|V|)$ time. However, when the genus is unbounded, runtime explodes. This algorithm is a vast improvement over the naïve, but one cannot expect to sit down and implement it in an afternoon.

We present a family of graphs whose genus is sharply bounded between 0 and $|V|/4$, and an algorithm to compute that genus in linear time.

Problem 2.1. *Let D be a surface homeomorhpic to an open disc, and assign the points on ∂D a circular ordering. Then, pick two distinct points $x, y \in \partial D$, and n distinct points from each interval (x, y) and (y, x), and label them*

$$\{p_1, p_2, \ldots, p_n\} \subset (x, y) \text{ and } \{q_1, q_2, \ldots, q_n\} \subset (y, x)$$

such that

$$x > p_1 > \cdots > p_n > y > q_n > \cdots > q_2 > q_1 > x.$$

Then, let $\tau \in S_n$. If $\mathcal{R}$ is a compact, orientable Riemann surface such that $D \subset \mathcal{R}$, with mutually disjoint parametric curves $\gamma_i : [0, 1] \to \mathcal{R} \setminus D$,

$$\gamma_i(0) = p_i \text{ and } \gamma_i(1) = q_{\tau(i)},$$

what is the minimal genus that $\mathcal{R}$ can have?

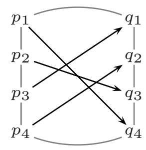

FIGURE 2. A naïve crossing diagram for $\tau = [4312]$

From here on, we will avoid discussion of topology, since the author is fond of the discrete calculation afforded by rotation systems. We call the collection of curves $\{\gamma_1, \ldots, \gamma_n\}$ a *connection*, and these curves together with ∂D will be called the *face graph*. In [4], Curtis and Morrow discuss this situation when a (more general) connection is planar, for the purpose of recovering the conductivities of an electrical network based on the boundary response.

3. Definitions

Definition 3.1. A *block interchange* is a permutation which exchanges two contiguous, non-overlapping sequences. For example,

$$[15674238] : 1\overline{23}4\underline{567}8 \mapsto 1\underline{567}4\overline{23}8,$$

$$[15672348] : 1\overline{234}\underline{567}8 \mapsto 1\underline{567}\overline{234}8,$$

etc., are block interchanges, as are transpositions. Therefore, every permutation can be factored into block interchanges. We call the minimal number of terms in any such factorization of a permutation φ the *block interchange distance*, denoted $d_I(\varphi)$.

Note, we underline/overline blocks for clarity in the above example. We will not continue to use this notation.

Computational biologists are commonly interested in genome rearrangements. Typical edits performed on a genome are *block interchanges, block transpositions*, and *reversals*. So, computational biology has recently emerged as a source of interesting problems concerning symmetric groups S_n for very large n.

Definition 3.2 (bigraph)**.** A *bigraph* is a pair of simple directed graphs which have a common vertex set. We denote a bigraph $G = (V, E_0, E_1)$ where $E_0, E_1 \subset V \times V$. Loosely, we describe a bigraph as a graph whose edges are assigned one of two colors, but in general, we want to allow a nonempty intersection of E_0 and E_1 and avoid discussion of multigraphs (or edges assigned multiple colors).

Definition 3.3. We define a *cycle* in two ways. In general, we will denote $c(\cdot)$ as the number of cycles of an object.

(1) For a bigraph G, an *alternating cycle* is a sequence of edges

$$(v_0, v_1), (v_1, v_2), \ldots, (v_n, v_0)$$

whose edges alternate in color. Here, $c(G)$ denotes the number of alternating cycles in G. Note, this is not generally a good definition. However, for bigraphs discussed in this paper, the inbound and outbound degree is 1 for both grey and black edges at every node. Hence, the edge sets uniquely decompose into disjoint alternating cycles, so for our purposes, $c(G)$ is well-defined.

(2) For a permutation $\pi \in S_n$, we define a cycle as the orbit of an element i under the permutation π. Then,

$$c(\pi) = |\{s \subset \{1, \cdots, n\} : \langle\pi\rangle \text{ acts transitively on } s\}|,$$

the number of disjoint cycles under the action of π.

Definition 3.4. Given a permutation $\tau \in S_n$, we define the *cycle graph*, as the bigraph $G_\tau = (V, E_g, E_b)$ with vertices

$$V = \{q_0, q_1, \dots, q_n\},$$

black edges

$$E_b = \{(q_{\tau(j)}, q_{\tau(i)}) : 0 \leq i \leq n, j = i + 1 \bmod n + 1\}$$

where we define $\tau(0) = 0$, and grey edges

$$E_g = \{(q_i, q_j) : 0 \leq i \leq n, j = i + 1 \bmod n + 1\}.$$

Bafna and Pevzner introduce a slightly different form of the cycle graph in [1], which is also used in [2]. Later, Doignon and Labarre use the the form above in [5]. The difference is that our vertex q_0 is the result of merging vertices 0 and $n+1$ from the original, and doesn't make or break cycles – so doesn't invalidate our use of Christie's result.

Theorem 3.5. *The block interchange distance, of $\tau \in S_n$ is given by*

$$d_I(\tau) = (n + 1 - c(G_\tau))/2.$$

Proof. See Theorem 4 in [2]. □

Definition 3.6. Finally, we define the *face graph* of a permutation $\tau \in S_n$, as the bigraph $F_\tau = (V, E_g, E_b)$ with vertices

$$V = \{p_1, \dots, p_n, q_1, \dots q_n\},$$

black edges

$$E_b = \{(p_i, q_{\tau(i)}) : i = 0, \dots, n\} \cup \{(q_{\tau(i)}, p_i) : i = 0, \dots, n\},$$

and grey edges

$$E_g = \{(p_n, p_{n-1}), \dots, (p_2, p_1), (p_1, q_1), (q_1, q_2), \dots, (q_{n-1}, q_n), (q_n, p_n)\}.$$

In Theorem 4.1, we denote the set-theoretic graph of a function $f : X \to Y$ by

$$\Gamma(f) = \{(x, f(x)) : x \in X\} \subset X \times Y.$$

In [5], Doignon and Labarre use the same notation for the looped graph with vertices X and edges $\Gamma(f)$ as defined here.

4. Tracing Faces

Now, we construct the rotation system of F_τ. We will refer to this as the *preferred combinatorial embedding* of F_τ, since we are interested in the genus of this embedding. We assign the darts labels

$$B = \{a_1, b_1, r_1, a_2, \dots, a_n, b_n, r_n\} \cup \{c_1, d_1, s_1, c_2, \dots, c_n, d_n, s_n\},$$

where a_i, b_i and r_i are associated to p_i, and c_i, d_i, s_i to q_i (see Figures 4 and 4). Then, we write ν and ε in cycle notation,

$$\nu = \prod_{i=1}^{n} (a_i b_i r_i)(c_i d_i s_i)$$

and

$$\varepsilon = (b_1 s_1)(d_n r_n) \prod_{i=1}^{n-1} (r_i b_{i+1})(d_i s_{i+1}) \prod_{i=1}^{n} (a_i c_{\tau(i)}).$$

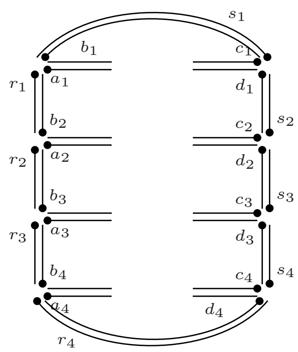

FIGURE 3. The darts of a connection graph. Note, the crossings are not shown.

Then, we compute $\phi = \nu\varepsilon$ to act on $B \setminus \{r_i, s_i : 1 \le i \le n\}$ by

$$\phi : \begin{cases} a_i & \mapsto & d_{\tau(i)}, \\ b_i & \mapsto & a_{i-1}, & i \ne 1, \\ b_1 & \mapsto & c_1, \\ c_{\tau(i)} & \mapsto & b_i, \\ d_i & \mapsto & c_{i+1}, & i \ne n, \\ d_n & \mapsto & a_n, \end{cases}$$

and the action on $\{r_i, s_i : 1 \le i \le n\}$ produces a single cycle, $(s_1 s_2 \dots s_n r_n \dots r_2 r_1)$.

Note that the face graph is merely a jazzed-up version of our diagram in Figure 2. We've colored the edges corresponding to the connection black, those corresponding to the boundary circle grey, and assigned directions for a nice correspondence to G_τ. This correspondence to G_τ is the motivation for the name *face graph* – cycles in the cycle graph correspond to faces in the embedded face graph.

Theorem 4.1. *For $\tau \in S_n$, the number of faces in the preferred combinatorial embedding of F_τ is given by $c(\phi) = 1 + c(F_\tau)$.*

Proof. Denote $I = \{1, 2, \cdots, n\}$, and

$$X = \{a_i, c_i : i \in I\}, Y = \{b_i, d_i : i \in I\}, Z = \{r_i, s_i : i \in I\}$$

so $B = X \cup Y \cup Z$, a disjoint union. Note that $\phi(Z) = Z$, $\phi(X) = Y$, and $\phi(Y) = X$. Then, let $\ell : B \to V(F_\tau)$ assigns darts to vertices; $a_i, b_i, r_i \mapsto p_i$, $c_i, d_i, s_i \mapsto q_i$, and extend the map to ordered pairs accordingly:

$$\ell(x, y) = (\ell(x), \ell(y)).$$

By this construction, we can formally write $\ell(x) \in V(F_\tau)$ for all $x \in B$.

Furthermore, for all $1 \le i \le n$,

$$\ell(a_i, \phi(a_i)) = \ell(a_i, d_{\tau(i)}) = (p_i, q_{\tau(i)}) \in E_b$$

and

$$\ell(c_{\tau(i)}, \phi(c_{\tau(i)})) = \ell(c_{\tau(i)}, b_i) = (q_{\tau(i)}, p_i) \in E_b.$$

Similarly,

$$\ell(b_i, \phi(b_i)) = (p_i, p_{i-1}) \in E_g \text{ if } i \ne 1,$$
$$\ell(b_1, \phi(b_1)) = (p_1, q_1) \in E_g,$$

and

$$\ell(d_i, \phi(d_i)) = (q_i, q_{i+1}) \in E_g, \text{ if } i \ne n,$$

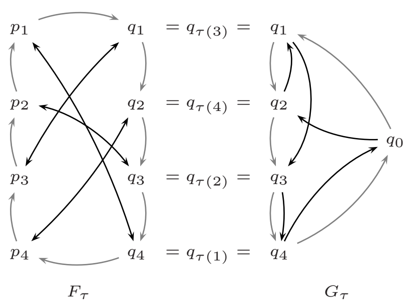

FIGURE 4. F_τ and G_τ for $\tau = [4312]$.

$$\ell(d_n, \phi(d_n) = (q_n, p_n) \in E_g.$$

Therefore, $\ell(\Gamma(\phi|_X)) = E_b$ and $\ell(\Gamma(\phi|_Y)) = E_g$. Thus, since $\phi(X) = Y$ and $\phi(Y) = X$, every cycle of $\phi|_{X\cup Y}$ corresponds to an alternating cycle in F_τ and by a counting argument, all edges are covered by this correspondence. Finally, we have one additional cycle from Z, so $c(\phi) = 1 + c(F_\tau)$. □

Remark 4.2. In Theorem 4.1, we construct ℓ to behave as a functor; taking objects to objects and relations to relations in a cohesive manner. A stronger property is that ℓ covers cycles in F_τ with orbits of ϕ. In the following, we rely on a similar construction. However, there are not enough vertices in the target graph for a bijection. So, we skirt this issue by mapping segments of cycles to segments of cycles, taking care to only consider segments whose endpoints appear in both graphs.

Theorem 4.3. *For all $\tau \in S_n$, $c(F_\tau) = c(G_\tau)$.*

Proof. Let ℓ map $q_i \in V(F_\tau)$ to $q_i \in V(G_\tau)$. We will proceed to extend ℓ to map alternating paths between edges in F_τ to edges in G_τ such that color, direction, and vertex labels match at the path boundaries. Denoting an alternating sequence $[e_1e_2 \dots e_n]$; for $1 \le i < n$ let

$$\ell([(q_i, q_{i+1})]) := [(q_i, q_{i+1})],$$

(here grey $\mapsto$ grey) and,

$$\ell([(q_{\tau(i+1)}, p_{i+1})(p_{i+1}, p_i)(p_i, q_{\tau(i)})]) := [(q_{\tau(i+1)}, q_{\tau(i)})],$$

(so black-grey-black $\mapsto$ black).

There are four edges in each F_τ and G_τ not covered by this mapping, so we map them in the only way that makes sense,

$$\ell([(q_{\tau(1)}, p_1)(p_1, q_1)]) := [(q_{\tau(1)}, q_0)(q_0, q_1)],$$
$$\ell([(q_n, p_n)(p_n, q_{\tau(n)})]) := [(q_n, q_0)(q_0, q_{\tau(n)})].$$

where these sequences are grey-black and black-grey, respectively. Note that we don't explicitly map any vertex in F_τ to q_0 in G_τ. However, by covering every edge by nonoverlapping cycle segments such that the endpoints agree, we have established a one-to-one correspondence of cycles in F_τ and G_τ as desired. □

Theorem 4.4. *The preferred combinatorial embedding of F_τ has genus $g = d_I(\tau)$.*

Proof. By Corollary 1.10 and the above,

$$2g = 2 - (|V(F_\tau)| - |E(F_\tau)| + 1 + c(G_\tau)) .$$

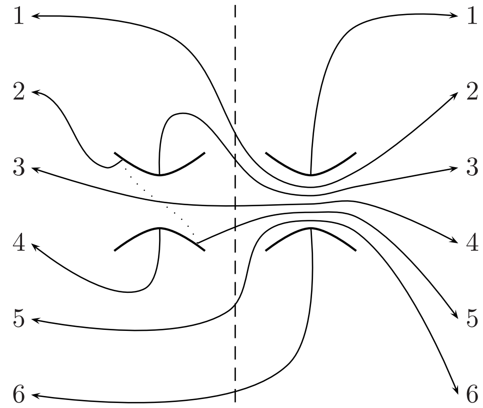

FIGURE 5. An embedding of [254361] on a genus 2 surface.

Since $|V(F_\tau)| = 2n$ and $|E(F_\tau)| = 3n$,

$$2g = 2 - (2n - 3n + 1 + c(G_\tau)) = 2 - (1 - n + c(G_\tau))$$

and finally, by Theorem 3.5, $d_I(\tau) = (n + 1 - c(G_\tau))/2$ so

$$2g = 2 - (2 - 2d_I(\tau)),$$

and $g = d_I(\tau)$ as desired. □

5. An Example

Now, consider the permutation $\tau = [254361] = [143256] \cdot [234561]$. The rotation system associated to this permutation is

$$\nu = (a_1b_1r_1)\cdots(a_6b_6r_6)(c_1d_1s_1)\cdots(c_6d_6s_6),$$

$$\varepsilon = (r_1b_2)\cdots(r_5b_6)(r_6d_6)(s_6d_5)\cdots(s_2d_1)(s_1b_1)(a_1c_2)(a_2c_5)\cdots(a_6c_1).$$

Then,

$$\begin{aligned}\phi = \nu\varepsilon \quad &= \quad (s_1s_2s_3s_4s_5s_6r_6r_5r_4r_3r_2r_1)(a_1d_2c_3b_4a_3d_4c_5b_2)\\ &\quad\quad (a_2d_5c_6b_5a_4d_3c_4b_3)(a_5d_6a_6d_1c_2b_1c_1b_6).\end{aligned}$$

By the Euler characteristic, our diagram has genus

$$\begin{aligned} g &= 1 - \frac{1}{2}(c(\nu) - c(\varepsilon) + c(\phi)) \\ &= 1 - \frac{1}{2}(2 \cdot 6 - 3 \cdot 6 + 4) \\ &= 2. \end{aligned}$$

Indeed, we can embed this connection on a surface with two handles, as in Figure 5. Furthermore, note that we have drawn the figure to "present" the factorization of [254361] into block interchanges. One immediate corollary of Theorem 4.4 is that we can embed a connection in a genus 1 surface if and only if it is a block interchange.

Remark 5.1. By the above, we have reduced the problem of drawing a connection on its minimal surface to drawing each block interchange in a minimal factorization of the permutation.

6. Extensions, Conclusion

In [5], Doignon and Labarre derive an explicit formula for *Hultman numbers*, $\mathcal{S}_H(n,k)$, the number of permutations in S_n whose cycle graph decomposes into k alternating cycles. They apply Christie's result to find that the number of permutations τ in S_n whose block interchange distance is k to be $\mathcal{S}_H(n, n+1-2k)$. For instance the number of permutations which require maximal genus for a fixed n arise where G_τ has exactly one cycle. So, for even n, there are $2\frac{n!}{n+2}$ such permutations. For odd n, there is only one, the reverse permutation.

Also, Christie describes an algorithm to explicitly compute a minimal factorization of a permutation into block interchanges, which we will not reproduce here. A careful implementation of this algorithm takes linear time in n. Hence, we can compute the genus of F_τ in $O(|V|)$, regardless of genus. Further, as in Remark 5.1, we can draw a reasonably nice picture of a minimal embedding of F_τ with ease.

7. Acknowledgements

Joseph Mitchell initially posed this question while searching for an analogue of the *medial graph* in non-planar graphs. Joe and I worked closely together on this problem, and together we achieved a positive resolution in the Summer of 2007. However, the resulting exposition lacked clarity and rigor.

Many thanks go to Jim Morrow, for his encouragement, advice, and relentless diligence even after dozens of readings of the same document. Also, without Sloane's tables[9], Joe and I may never have solved this problem – we made little real progress before [5] was found through his website. Finally, to my family, and Megan, for everything.

References

[1] Vineet Bafna and Pavel A. Pevzner. Sorting by transpositions. *SIAM Journal of Discrete Mathematics*, 11:224–240, 1998.
[2] David A. Christie. Sorting permutations by block-interchanges. *Information Processing Letters*, 60(4):165–169, 1996.
[3] Robert Cori and Antonio Machì. Maps, hypermaps and their automorphisms: a survey. *Expositiones Mathematicæ*, 10, 1992.
[4] Edward B. Curtis and James A. Morrow. *Inverse Problems for Electrical Networks*. World Scientific Publishing Company, first edition, 2000.
[5] Jean-Paul Doignon and Anthony Labarre. On Hultman Numbers. *Journal of Integer Sequences*, 10(6), June 2007.
[6] John M. Lee. *Introduction to Topological Manifolds*. Springer, first edition, 2000.
[7] Bojan Mohar. Embedding graphs in an arbitrary surface in linear time. In *STOC '96: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 392–397, New York, NY, USA, 1996. ACM.
[8] Bojan Mohar and Cartsen Thomassen. *Graphs on Surfaces*. The Johns Hopkins University Press, first edition, 2001.
[9] N. J. A. Sloane. The on-line encyclopedia of integer sequences, 2007.