Flag Bicolorings, Pseudo-Orientations, and Double Covers of Maps

Hiroki Koike*

Facultad de Ciencias
Universidad Nacional Autónoma de México
Distrito Federal, México

koike@ciencias.unam.mx

Daniel Pellicer†

Centro de Ciencias Matemáticas
Universidad Nacional Autónoma de México
Morelia, Mexico

pellicer@matmor.unam.mx

Miguel Raggi‡

Escuela Nacional de Estudios Superiores
Universidad Nacional Autónoma de México
Morelia, Mexico

mraggi@gmail.com

Steve Wilson

Department of Mathematics and Statistics
Northern Arizona University
Flagstaff, USA

stephen.wilson@nau.edu

Abstract

This paper discusses consistent flag bicolorings of maps, in their own right and as generalizations of orientations and pseudo-orientations. Furthermore, a related doubling concept is introduced, and relationships between these ideas are explored.

Keywords: Map, Orientation, Pseudo-Orientation, Consistent Bicoloring, Flag

1   Introduction

The main goal of this paper is to develop a general theory of flag bicolorings and the related concepts of coverings and pseudo-orientations.

The idea of consistent colorings of the flags of a map with two colors has appeared previously in literature in different contexts. For instance, when the automorphism group of a map which is a polytope has two orbits on the flags, we may color the flags with two colors in such a way that flags in different orbits have different colors. These are

called 2-orbit polytopes in [8], but that investigation requires the extra property that the maps be abstract polytopes. Specific k-colorings of flags are equivalent to the concept of $\mathcal{T}$-*compatible* maps introduced in [13], where $\mathcal{T}$ is a class of k-orbit maps. This concept, generalized to hypermaps, is called $\mathcal{T}$-*conservative* in [1]. In all these instances, the colorings are used as a tool to work with automorphisms of maps.

Although some of the motivations behind the bicoloring of flags come from maps admitting symmetries, the ideas can be applied to general maps. In this work we shall not make assumptions on the automorphism groups of the maps in consideration. Moreover, the definitions of consistent bicolorings of flags and pseudo-orientations can be easily extended to maniplexes [17] (higher dimensional analogues of maps) or to more general incidence geometries [4]. Some results of this paper may even become obvious in the context of operations of hypermaps (see [10]). However, to keep the introductory sections brief, we restrict ourselves to maps. Furthermore, one of the main results of this paper, namely Theorem 8.1 is specifically about maps.

Maps admitting a bicoloring of flags in which adjacent flags have different colors are precisely maps on orientable surfaces (see Proposition 5.1). Although the other types of bicolorings of flags that we consider in this paper do not have a known topological equivalence, some results on orientability do translate directly to analogous results on bicolorability. In this sense we may think of bicolorings of flags as generalizations of orientability in maps.

The paper is organized as follows. Sections 2, 3, 4 introduce the concepts of maps, pseudo-orientations and bicolorings, respectively. We also give a natural group structure to the set of bicolorings of any map. A relationship between pseudo-orientability and bicolorings is provided in Section 5. Section 6 explores the impact on bicolorability of operations on maps. In Section 7, we discuss a natural double cover of every map which is not bicolorable. In Section 8 we determine, for each possible group, which surfaces admit maps with the given group of bicolorings. Finally, in Section 9 we pose some open questions related to our results.

2   Maps

A map $\mathcal{M}$ is, first and foremost, an embedding of a graph (or pseudograph) on a (compact, connected) surface so that the components of the complement of the embedding (called *faces*) are topologically open disks. We can, for example, regard the cube as an embedding of the graph Q_3 on the sphere.

The graph Q_3 is an example of a graph which is *bipartite*, i.e., its vertices can be colored with two colors so that every edge joins vertices of opposite colors. When speaking about maps, we will use the term *vertex-bipartite* to describe a map $\mathcal{M}$ whose underlying graph is bipartite. Similarly, we will call $\mathcal{M}$ *face-bipartite* provided its faces can be colored with two colors so that each edge separates faces of opposite colors. We will call $\mathcal{M}$ *edge-bipartite* provided that the edges can be colored with two colors so that edges which are consecutive around a face (and hence around a vertex) have different colors.

To look more closely at the structure of a map, we find the following subdivision useful:

choose a point in the interior of each face to call its *center* and a point in the relative interior of each edge to be its *midpoint*. Draw dotted lines to connect each face-center with each incidence with the surrounding vertices and edge-midpoints. The original edges and these dotted lines divide the surface into triangles called *flags*. Figure 1 shows the subdivision of the cube into flags.
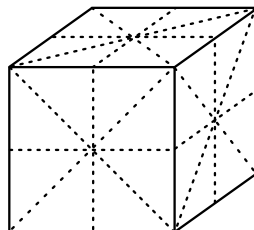

Figure 1: The cube divided into flags

If a face meets a vertex (or an edge) more than once, we emphasize that the face center is connected to each incidence—each appearance—of the vertex or midpoint. For instance, consider the map M_4 shown in the left of Figure 2. This map has only one face, an octagon, four edges and only one vertex. Nonetheless, the dissection into flags draws 16 dotted lines, dividing the octagon, and the map, into 16 flags, as shown on the right.

Figure 2: The map M_4

Each flag corresponds to a mutual incidence of face, edge, and vertex, though different flags may correspond to the same triple.

Let Ω be the set of flags. Then let r_0, r_1, r_2 be the permutations on Ω which match each flag f with its three immediate neighbors, as in Figure 3.

In that figure, we see that f and fr_0 are adjacent along a face-center-to-edge-midpoint line. Thus f and fr_0 differ only in their incidences to a vertex, a 0-dimensional face of $\mathcal{M}$. Similarly, f and fr_1 differ only in their incidences to an edge, a 1-dimensional face, while f and fr_2 differ only in their incidences to a 2-dimensional face. Notice from Figure 3 that the r_2-adjacent flag of fr_0 is also r_0-adjacent to fr_2. In other words, as permutations on Ω, r_0 and r_2 commute.

We can take a slightly more abstract point of view by defining a map to be a pair $(\Omega, [r_0, r_1, r_2])$ where Ω is a set of things called *flags*, the r_i's and $r_0 r_2$ are fixed point free
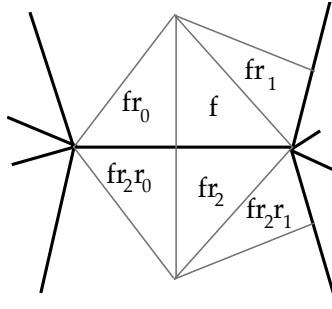
Figure 3: Flags in a map

permutations of order 2 on Ω, the *connection group* $C(\mathcal{M}) = \langle r_0, r_1, r_2 \rangle$ is transitive on Ω, and r_0 and r_2 commute. This $C(\mathcal{M})$ is often called the *monodromy group* of the map. We can then think of vertices in $\mathcal{M}$ as orbits of $\langle r_1, r_2 \rangle$ in $C(\mathcal{M})$. Similarly, edges correspond to orbits of $\langle r_0, r_2 \rangle$ and faces to orbits of $\langle r_0, r_1 \rangle$.

We will use the word *kite* or *corner* for the area within a face where two consecutive edges meet. More formally, a kite is the union of two flags which are r_1-adjacent.

If $\mathcal{M}$ and $\mathcal{N}$ are maps on surfaces $S_{\mathcal{M}}$ and $S_{\mathcal{N}}$, a *projection* from $\mathcal{N}$ to $\mathcal{M}$ is a function ϕ mapping $S_{\mathcal{N}}$ to $S_{\mathcal{M}}$ which is locally a homeomorphism at all points of $\mathcal{N}$ except perhaps at vertices and/or face-centers, and which sends faces to faces, edges to edges, and vertices to vertices.

In combinatorial terms, if $\mathcal{M} = (\Omega, [r_0, r_1, r_2])$ and $\mathcal{N} = (\Omega', [s_0, s_1, s_2])$, a projection of $\mathcal{N}$ to $\mathcal{M}$ is a function ϕ mapping Ω' to Ω such that $s_i \phi = \phi r_i$ for all $i \in R := \{0, 1, 2\}$.

We call such an $\mathcal{N}$ a *cover* of $\mathcal{M}$. Notice that if $\mathcal{N}$ is a cover of $\mathcal{M}$, and the pre-image of some flag in $\mathcal{M}$ has size k, then the projection ϕ is k-to-1 onto every flag of $\mathcal{M}$. We say then that $\mathcal{N}$ is a k-*fold* cover of $\mathcal{M}$.

3   Orientations and Pseudo-Orientations

We call a map *orientable* provided that the surface on which it is embedded is itself orientable. We check that a map is orientable by giving it a *face orientation*; this is an assignment of a circular arrow to each face of $\mathcal{M}$ such that at every edge, the arrows on the faces joined by the edge point along the edge in *opposite* directions, as in Figure 4. We can define a *vertex orientation* similarly, and it is clear that $\mathcal{M}$ has a vertex-orientation if and only if it has a face-orientation, and this happens if and only if $\mathcal{M}$ is orientable.

The word *pseudo-orientation* has been used in two different ways. We will use the term *vertex pseudo-orientation* (VPSO for short) for what is called in [16] simply a pseudo-orientation. Here, we mean an assignment of one circular arrow to each vertex so that at each edge the two arrows cross in the *same* direction, as in Figure 5. If we visualize a gear wheel at each vertex so that cogs on the wheels of adjacent vertices mesh, the map is VPSO if we can turn one wheel, causing all wheels to turn at the same time.

In [16], the idea is used to make an important distinction about k-fold rotary covers
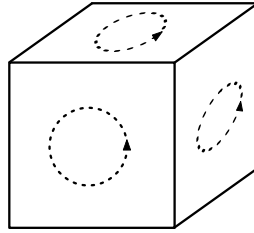
Figure 4: Face orientation of the cube

of a rotary map $\mathcal{M}$ for which the branching is totally ramified at vertices. If $\mathcal{M}$ is non-orientable, k can be larger than 2 only if $\mathcal{M}$ is vertex pseudo-orientable. Vertex pseudo-orientable maps on surfaces of low genus were recently classified in [3].

Similarly, a *face pseudo-orientation* (or FPSO) is an assignment of one circular arrow to each face, as in Figure 5, so that at each edge, the arrows in adjacent faces flow along the edge in the *same* direction. We can simplify this by orienting each edge so that the cycle of edges around each face is consistently oriented.

In [6] this is called a pseudo-orientation, and the result in that paper is that the *Dart Graph* of $\mathcal{M}$ is connected if and only if $\mathcal{M}$ is not face pseudo-orientable.

Finally, we may define an *edge pseudo-orientation* (or EPSO) to be an orientation of the edges such that in every face the direction of the arrows in adjacent edges flows in the *same* direction (either into or out of the face) as in Figure 5.
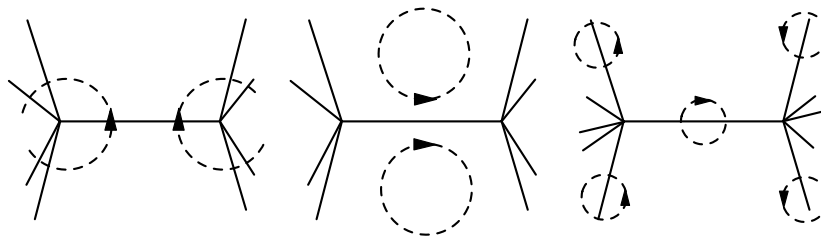

Figure 5: Vertex, Face and Edge Pseudo-orientations

4   Colorings

In this paper, the word 'coloring' will be used to describe what might be more fully notated as a 'consistent flag 2-coloring'. If I is any subset of $R = \{0, 1, 2\}$, then an I-coloring of a map $\mathcal{M}$ is a function $a : \Omega \to \mathbb{Z}_2$ such that for every flag f, if $j \in I$, then $a(fr_j) \neq a(f)$, while if $j \notin I$, then $a(fr_j) = a(f)$. To say that in another way,

$$a(fr_j) + a(f) = \begin{cases} 1 & \text{if } j \in I \\ 0 & \text{if } j \notin I. \end{cases}$$

Figure 6 shows a $\{0\}$-coloring, a $\{1, 2\}$-coloring, and an R-coloring of the cube. The

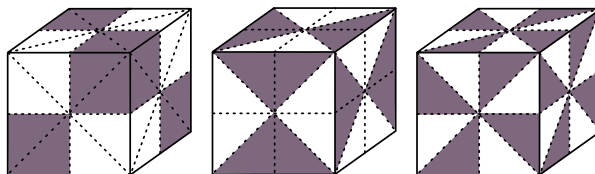reader should check that the colorings extend unambiguously to the unseen faces of the cube as well.

Figure 6: Colorings of the cube

Let us first notice that if a is an I-coloring of $\mathcal{M}$, then so is $1 - a$, and these are, in fact, the only I-colorings of the map.

Give the name *sub-edge* to the segments forming sides of the flags. These are of three types:

Type 0 : midpoint-of-edge to face-center,

Type 1 : face-center to vertex,

Type 2 : vertex to midpoint-of-edge.

Each flag f is adjacent to flag fr_i along their mutual sub-edges of type i. If $\mathcal{M}$ is I-colorable, therefore, the sub-edges bounding each region of a single color are the sub-edges of types in I.

Note that when considering maps in class 2_I as defined in [8], the subset I is used to indicate that i-adjacent flags belong to the *same* orbit. This means that maps in class 2_I admit an $(R \setminus I)$-coloring. In the context of colorings it turns out to be more practical to define I-coloring requiring i-adjacent flags to be of a different color rather than requiring i-adjacent flags to be of the same color if and only if $i \in I$.

4.1 The group $T(\mathcal{M})$

Let Δ stand for the symmetric difference operation: $I \Delta J = (I \cup J) \setminus (I \cap J)$. Let $\mathcal{P}$ be the power set of R. It is well-known that $\mathcal{P}$ is a group under Δ and it is an elementary Abelian group of order eight whose identity element is $\emptyset$.

Proposition 4.1. If a map $\mathcal{M}$ admits an I-coloring and a J-coloring, then it admits an $I \Delta J$-coloring.

Proof: It is easy to verify that the sum of an I-coloring and a J-coloring is an $I \Delta J$-coloring. ∎

From this proposition we see that, for each map $\mathcal{M}$, the sets I for which $\mathcal{M}$ has an I-coloring form a set $T(\mathcal{M})$ which is a subgroup of $\mathcal{P}$.

The reader is encouraged to look at [2, Section 5] for another approach to understanding this group.

4.2 Flag-colorings and map-colorings

Whenever I is a subset of size 1, I-colorable maps have a nice characterization.

Proposition 4.2. A map $\mathcal{M}$ is vertex-bipartite if and only if $\mathcal{M}$ has a $\{0\}$-coloring.

Proof: Given such a coloring, all the flags incident with one vertex will be the same color, while those at an adjacent vertex will be the other color. This is a bipartite coloring of the vertices. Conversely, given a bipartition of the vertices, color all flags incident with black vertices with color 1 and all those incident with white vertices with color 0. This is then a $\{0\}$-coloring of $\mathcal{M}$. ∎

Similar proofs lead us to these:

Proposition 4.3. A map $\mathcal{M}$ is face-bipartite if and only if $\mathcal{M}$ has a $\{2\}$-coloring. (*i.e.* a $\{2\}$-coloring corresponds to a bipartite coloring of the faces of $\mathcal{M}$.)

Proposition 4.4. A map $\mathcal{M}$ is edge-bipartite if and only if $\mathcal{M}$ has a $\{1\}$-coloring. (*i.e.* a $\{1\}$-coloring corresponds to a bipartite coloring of the edges of $\mathcal{M}$.)

4.3 Colorings and words

Consider a *cycle* $f, fr_{i_1}, fr_{i_1}r_{i_2}, fr_{i_1}r_{i_2}r_{i_3}, \ldots, fr_{i_1}r_{i_2}r_{i_3}\ldots r_{i_k} = f$ of flags in a map $\mathcal{M}$. We abbreviate that (f, W), where W is the string $i_1 i_2 i_3 \ldots i_k$. Write r_W for $r_{i_1}r_{i_2}r_{i_3}\ldots r_{i_k}$. We call a cycle (f, W) *I-consistent* if the number of occurrences in W of indices which are in I is even. If (f, W) is I-consistent, this says that, at least along the cycle, we can color flags from two colors so that flags which are i-connected for $i \notin I$ are the same color and those for which $i \in I$ are not. A cycle which is not I-consistent is *I-inconsistent*. It follows, then that $\mathcal{M}$ is I-colorable if and only if every cycle in $\mathcal{M}$ is I-consistent.

Consider this quite general fact: for any function F from a set X to $\mathbb{Z}_2$, extend F to the power set of X by $F(A) = \Sigma_{x \in A} F(x)$ for each subset A of X. Then for any subsets A, B of X, it is clear that $F(A \Delta B) = F(A) + F(B)$.

Lemma 4.5. For any subsets I, J of R, if $K = I \Delta J$, then every cycle in M is consistent either for exactly one of I, J, K or for all three.

Proof: Fix a cycle (f, W) of $\mathcal{M}$, and define F on R with $F(i)$ being the parity of the number of occurrences in W of r_i; furthermore, consider F extended to $\mathcal{P}$. The cycle is I-consistent, then, provided that $F(I) = 0$. Now consider $F(I) + F(J) = F(I \Delta J) = F(K)$. These are elements of $\mathbb{Z}_2$, and so we can rephrase that in this form: $F(I) + F(J) + F(K) = 0$. The number of zeros in $\{F(I), F(J), F(K)\}$ must be 1 or 3; the conclusion follows directly. ∎

We will use this lemma in Section 7.

4.4 Connected sum of maps

A natural operator in topology is the *connected sum* of two surfaces, which is defined to be the surface formed by cutting a small disk from each of the two surfaces and attaching them along the newly created border. We extend this definition to maps by causing the attachment sets to be the boundaries of two suitably chosen faces.

Formally, let $\mathcal{M} = (\Omega, [r_0, r_1, r_2])$ and $\mathcal{N} = (\Omega', [s_0, s_1, s_2])$ be maps and suppose $f_{\mathcal{M}}$ and $f_{\mathcal{N}}$ are flags of these maps for which their respective faces $F_{\mathcal{M}}$ and $F_{\mathcal{N}}$ have the same number k of sides. Furthermore, suppose that no flag of $F_{\mathcal{M}}$ is r_2-connected to any of the flags of $F_{\mathcal{M}}$, and similarly for $F_{\mathcal{N}}$. Then we define the *connected sum* $\mathcal{M} \oplus \mathcal{N}$ with respect to $(f_{\mathcal{M}}, f_{\mathcal{N}})$ to be the map whose flag set is $(\Omega \backslash F_{\mathcal{M}}) \cup (\Omega' \backslash F_{\mathcal{N}})$, with connections t_0, t_1, t_2, where t_0 and t_1 are the restrictions of $r_0 \cup s_0, r_1 \cup s_1$ to this set, and t_2 is the same except that, for each j, $f_{\mathcal{M}}(r_0 r_1)^j r_2$ is t_2-connected to $f_{\mathcal{N}}(s_0 s_1)^j s_2$ and $f_{\mathcal{M}} r_1 (r_0 r_1)^j r_2$ is t_2-connected to $f_{\mathcal{N}} s_1 (s_0 s_1)^j s_2$. Because of the restriction on the faces, these flags are both in the new flag set, and it is easy to check that $t_0 t_2 = t_2 t_0$.
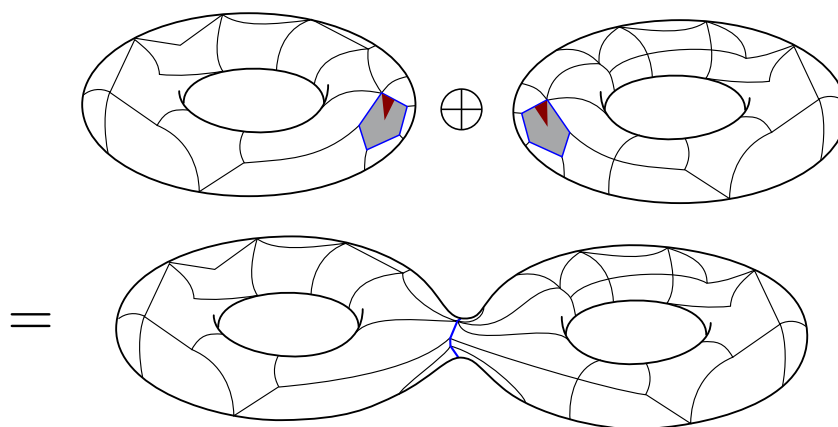
The following Figure illustrates this construction:

Figure 7: A connected sum of two maps

Proposition 4.6. Suppose $\mathcal{M}$ and $\mathcal{N}$ are I-colorable maps for some $I \subseteq R$ and suppose $\mathcal{M} \oplus \mathcal{N}$ exists for some flags $f_{\mathcal{M}}$ and $f_{\mathcal{N}}$. Then $\mathcal{M} \oplus \mathcal{N}$ is I-colorable.

Proof: Choose an I-coloring for $\mathcal{M}$ and one for $\mathcal{N}$ in such a way that $f_{\mathcal{M}}$ and $f_{\mathcal{N}}$ have opposite colors if $2 \in I$ and the same color if $2 \notin I$. Then the induced coloring provides the desired coloring on $\mathcal{M} \oplus \mathcal{N}$. ∎

As we shall see in the following theorem, it is almost always the case that the group $T(\mathcal{M} \oplus \mathcal{N})$ is equal to $T(\mathcal{M}) \cap T(\mathcal{N})$. This can fail, roughly, when for some $I \subseteq R$, $\mathcal{M} \backslash F_{\mathcal{M}}$ is I-colorable but $\mathcal{M}$ is not, or similarly for $\mathcal{N}$.

Theorem 4.7. Suppose $\mathcal{M}$ and $\mathcal{N}$ are maps and suppose $\mathcal{M} \oplus \mathcal{N}$ exists for some flags $f_{\mathcal{M}}$ and $f_{\mathcal{N}}$. Furthermore, suppose that for each $I \notin T(\mathcal{M})$ there exists an I-inconsistent cycle

that doesn't include any of the flags of face $F_{\mathcal{M}}$, and similarly for $\mathcal{N}$. Then $T(\mathcal{M} \oplus \mathcal{N}) = T(\mathcal{M}) \cap T(\mathcal{N})$.

Proof: For $I \subseteq R$, if $I \in T(\mathcal{M}) \cap T(\mathcal{N})$, then $I \in T(\mathcal{M} \oplus \mathcal{N})$ by Proposition 4.6. If $I \notin T(\mathcal{M}) \cap T(\mathcal{N})$, then without loss of generality, assume that $\mathcal{M}$ is not I-colorable. Then $\mathcal{M}$ has cycles which are I-inconsistent, and by the hypothesis, at least one of them includes none of the flags in $F_{\mathcal{M}}$. This cycle, then, occurs in $\mathcal{M} \oplus \mathcal{N}$, and is still I-inconsistent. Thus $I \notin T(\mathcal{M} \oplus \mathcal{N})$. ∎

5   Colorings and Orientations

We now offer and prove a series of propositions connecting the ideas of orientations and colorings. First, we show a collection of easy facts to embody our belief that colorings generalize orientations. We then generalize these facts in Theorem 5.4.

Proposition 5.1. A map is orientable if and only if it has an R-coloring.

Proof: If $\mathcal{M}$ has an R-coloring, orient each face so that, along each edge, the orientation points from the flag with color 0 to the flag with color 1. This is consistent within the face and faces that meet along an edge have orientations which meet correctly. Conversely, given the orientation, assign colors so that along each edge within a face, the orientation points from the flag with color 0 to the flag with color 1. This causes colors to alternate within each face. Because the orientation opposes the orientation in each adjacent face, the colors alternate at each edge, giving an R-coloring. ∎

Similar arguments give us the following two propositions:

Proposition 5.2. A map $\mathcal{M}$ has a face pseudo-orientation if and only if $\mathcal{M}$ admits a $\{0, 1\}$-coloring.

Proposition 5.3. A map $\mathcal{M}$ has a vertex pseudo-orientation if and only if $\mathcal{M}$ admits a $\{1, 2\}$-coloring.

5.1 I-Pseudo-Orientability

Our aim in this subsection is to generalize the definition of pseudo-orientability and the previous results. For any subset I of R (except $\emptyset$), we form the map $X = X(\mathcal{M}, I)$ whose faces are the regions formed by conjoining flags which are connected by each r_i for which $i \notin I$. For example, if $\mathcal{M}$ is the cube, then $X(\mathcal{M}, \{1, 2\})$ and $X(\mathcal{M}, \{0, 2\})$ are shown in Figure 8.

Notice that the boundaries of the faces in $X(\mathcal{M}, I)$ are the sub-edges of types that are *in I*. Thus $\mathcal{M}$ is I-colorable if and only if $X(\mathcal{M}, I)$ is face-bipartite. However, the map $X(\mathcal{M}, I)$ can be constructed even if $\mathcal{M}$ is not I-colorable.

We call $\mathcal{M}$ *I-pseudo-orientable* (or I-PSO for short) when $X(\mathcal{M}, I)$ is face pseudo-orientable. If this happens, we can orient the edges of X so that around each face of
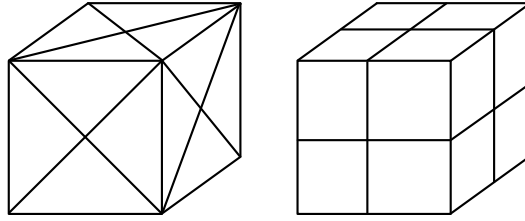
Figure 8: $X(\mathcal{M}, \{1, 2\})$ and $X(\mathcal{M}, \{0, 2\})$

X, the arrows all point the same way. We say a map is $\emptyset$-pseudo-orientable when it is orientable.

It is clear that $X(\mathcal{M}, \{2\}) = \mathcal{M}$ and so $\{2\}$-PSO is equivalent to FPSO. Observation 5.2 shows that this happens exactly when $\mathcal{M}$ is $\{0, 1\}$-colorable. Similarly, $\mathcal{M}$ is $\{0\}$-PSO when it is VPSO, and this happens exactly when $\mathcal{M}$ is $\{1, 2\}$-colorable. Surprisingly, this generalizes completely:

Theorem 5.4. Let $\mathcal{M}$ be a map. Then $\mathcal{M}$ is I-pseudo-orientable if and only if $\mathcal{M}$ is $(R \backslash I)$-colorable.

Proof: First, recall that the sub-edges of $\mathcal{M}$ are of three different types:

Type 0 : midpoint-of-edge to face-center,

Type 1 : face-center to vertex,

Type 2 : vertex to midpoint-of-edge.

Furthermore, let us think of the sub-edges as directed. Call the directions above the 'standard' orientations for the sub-edges, and the reverse of these are the non-standard orientations.

In order to form $X(\mathcal{M}, I)$, we delete or ignore sub-edges of each type $i \notin I$, so we only consider sub-edges of type in I. Suppose $\mathcal{M}$ is already $(R \backslash I)$-colored. Then in $X(\mathcal{M}, I)$, both flags around each sub-edge have the same color. Then we may assign this as the color of the sub-edge. Now give every sub-edge with color 0 the standard orientation, and each sub-edge with color 1 the non-standard direction.

A case by case analysis easily shows that for any choice of I the resulting orientation of the edges induces the desired pseudo-orientation. Note that because of our definition of $\emptyset$-pseudo-orientable as orientable, this theorem works even in this case.

For the converse we may use the orientation to color in the way prescribed to match the previous list. ∎

5.2 The Cheat Sheet

We summarize here the results of previous sections relating to the question of which maps have I-colorings for a given I.

I	$\mathcal{M}$	Requires:
$\emptyset$	All flags have the same color	-
$\{0\}$	Vertex-bipartite	Faces have even degree
$\{1\}$	Edge-bipartite	Faces and vertices have even degree
$\{2\}$	Face-bipartite	Vertices have even degree
$\{0,1\}$	Face Pseudo-Orientable	Vertices have even degree
$\{0,2\}$	Edge Pseudo-Orientable	Faces and vertices have even degree
$\{1,2\}$	Vertex Pseudo-Orientable	Faces have even degree
R	Orientable	-

6 Operators

In this section we shall see how colorings and orientations interact with common map operators such as dual, Petrie, opposite and medial.

Let $\mathcal{M} = (\Omega, [r_0, r_1, r_2])$ be any map. Then $D(\mathcal{M})$ (or the *dual* of $\mathcal{M}$) is the map $(\Omega, [d_0, d_1, d_2])$, where $d_j = r_{2-j}$. This corresponds exactly to the classical geometric dual of a polyhedron. The maps $\mathcal{M}$ and $D(\mathcal{M})$ lie on the same surface.

The *Petrie* of $\mathcal{M}$, $P(\mathcal{M})$, is the map $(\Omega, [p_0, p_1, p_2])$, where $p_0 = r_0 r_2, p_1 = r_1$, and $p_2 = r_2$. This is a less familiar operator on maps. The faces of $P(\mathcal{M})$ are the Petrie paths (see [15]) of $\mathcal{M}$ and vice-versa. Note that the vertices and the edges are preserved by the operation.

We call a map formed from $\mathcal{M}$ by any composition of the operation D and P a *direct derivate* of $\mathcal{M}$. Of special interest among the direct derivates of $\mathcal{M}$ is $opp(\mathcal{M}) = PDP(\mathcal{M}) = DPD(\mathcal{M})$. Formally, this is $(\Omega, [s_0, s_1, s_2])$ where $s_0 = r_0, s_1 = r_1$ and $s_2 = r_0 r_2$. More intuitively, $opp(\mathcal{M})$ is formed from $\mathcal{M}$ by cutting along the edges and then re-attaching along the same matching edges but with the reverse local orientation. See [15] for more information about these operators.

The *medial* of a map $\mathcal{M}$ is drawn on the same surface as $\mathcal{M}$. The vertices of the medial are the edge-midpoints of $\mathcal{M}$ and two are connected by an edge diagonally across each kite to which both belong.
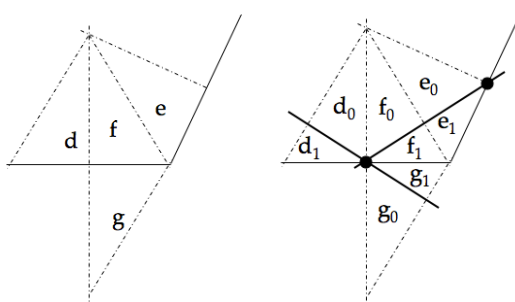

Figure 9: Flags in a map and its medial

Figure 9 shows, on the left, the flags adjacent to a flag f in a map. On the right we see the edges of the medial drawn in heavy lines. These divide each flag into two. This suggests the combinatorial form: the medial of $\mathcal{M} = (\Omega, [r_0, r_1, r_2])$ is the map $(\Omega \times \mathbb{Z}_2, [s_0, s_1, s_2])$ where s_0, s_1, s_2 are defined by

1. $(f, i)s_0 = (fr_1, i)$

2. $(f, 0)s_1 = (fr_0, 0)$

3. $(f, 1)s_1 = (fr_2, 1)$

4. $(f, i)s_2 = (f, 1 - i)$.

6.1 Operators and colorings

In this subsection, we describe the ways in which the colorings of a map come from the colorings of maps related to it by operators. We omit the proofs whenever they follow directly from the definitions.

Lemma 6.1. If $\mathcal{M}$ admits an I-coloring, then $D(\mathcal{M})$ admits a J-coloring, where $J = \{2 - i : i \in I\}$

Proposition 6.2. If $\mathcal{M}$ admits an I-coloring, then $opp(\mathcal{M})$ admits a J-coloring, where:

$$ J = \begin{cases} I & \text{if } 0 \notin I \\ I\Delta\{2\} & \text{if } 0 \in I. \end{cases} $$

Proof: Suppose $\mathcal{M}$ is I-colored. We shall take opposites but keep the coloring of each flag. Let f be a flag. Then the r_0 and r_1-adjacent flags in $opp(\mathcal{M})$ are the same as in $\mathcal{M}$, and the r_2-adjacent flag is $fr_2 r_0$. If $0 \notin I$, then fr_2 is the same color as $fr_2 r_0$, which means the coloring doesn't change when taking opposites. Analogously, if $0 \in I$, fr_2 is the opposite color as $fr_2 r_0$, so taking opposite changes the color of the flag r_2-adjacent to f. ∎

Corollary 6.3. If $\mathcal{M}$ admits an I-coloring, then $P(\mathcal{M})$ admits a J-coloring, where:

$$ J = \begin{cases} I & \text{if } 2 \notin I \\ I\Delta\{0\} & \text{if } 2 \in I. \end{cases} $$

Corollary 6.4. A map $\mathcal{M}$ has all colorings, *i.e.* $T(\mathcal{M}) = \mathcal{P}$, if and only if the surfaces of $\mathcal{M}$, $opp(M)$ and $P(M)$ are orientable.

Proof: Indeed, if $R \in T(\mathcal{M}) \cap T(P(\mathcal{M})) \cap T(opp(\mathcal{M}))$, by the preceding theorems we conclude that $\{0, 1\}$ and $\{1, 2\}$ are both in $T(\mathcal{M})$, which means $R\Delta\{0, 1\} = \{2\}$ and $R\Delta\{1, 2\} = \{0\}$ are also in $T(\mathcal{M})$, and so $T(\mathcal{M}) = \mathcal{P}$, since $\{\{0\}, \{2\}, R\}$ is a generating set of $\mathcal{P}$ under Δ.

Now, suppose $T(\mathcal{M}) = \mathcal{P}$. By previous results, if $\mathcal{M}$ admits every coloring, then $opp(\mathcal{M})$ and $P(\mathcal{M})$ admit every coloring. Thus all three admit an R-coloring and so are orientable. ∎

Theorem 6.5. The medial map of any map can always be {2}-colored. Furthermore, for each I in the first column of the table below, a map $\mathcal{M}$ has an I-coloring if and only if the medial of $\mathcal{M}$ has a J-coloring and a $J\Delta\{2\}$-coloring shown in columns 2 and 3.

I	J	$J\Delta\{2\}$
$\emptyset$	$\emptyset$	$\{2\}$
$\{1\}$	$\{0\}$	$\{0,2\}$
$\{0,2\}$	$\{1\}$	$\{1,2\}$
R	R	$\{0,1\}$

Proof: The medial of any map always admits a {2}-coloring, since we have two types of faces: those that surround vertices of $\mathcal{M}$, and those that surround face-centers of $\mathcal{M}$. For the second row, if the map is edge-bipartite, then we may color each vertex of the medial with the color corresponding to the color of the edge it was the midpoint of and vice-versa. For the third, note that an edge of the medial corresponds to a kite of the original. Finally, the medial map lies on the same surface as the original. Coloring $J\Delta\{2\}$ is a consequence of the fact that the medial always admits a {2}-coloring and Proposition 4.1. ∎

7 Doubles

Given a map $\mathcal{M} = (\Omega, [r_0, r_1, r_2])$ and a subset I of $R = \{0, 1, 2\}$, we define the I-*double* of $\mathcal{M}$, denoted $I \succ \mathcal{M}$, to be the map $\mathcal{N} = (\Omega \times \mathbb{Z}_2, [s_0, s_1, s_2])$, where

$$f_i s_j = \begin{cases} (fr_j)_i & \text{if } j \notin I \\ (fr_j)_{1-i} & \text{if } j \in I. \end{cases}$$

Here, we write "f_i" for (f, i).

Note that this definition is the particular case for maps of the double coverings described in [2, Section 6], where some properties of the double coverings dealt with in this section have been discussed. Here we reproduce some of these in our context.

Notice that the function sending f_i to f is a projection of $\mathcal{N}$ onto $\mathcal{M}$, and so $I \succ \mathcal{M}$ is a covering of $\mathcal{M}$. Also notice that the function which sends f_i to $i \in \mathbb{Z}_2$, is an I-coloring of $I \succ \mathcal{M}$.

Now, if $\mathcal{M}$ has an I-coloring, then $I \succ \mathcal{M}$ consists of two disjoint copies of $\mathcal{M}$; in this case we discard one copy and say that $I \succ \mathcal{M}$ is isomorphic to $\mathcal{M}$. If $\mathcal{M}$ does not have an I-coloring then $I \succ \mathcal{M}$ truly is a double cover (i.e., a 2-fold topological covering which may or may not be branched at vertices and/or face-centers) of $\mathcal{M}$.

We have remarked above that the map $I \succ \mathcal{M}$ always has an I-coloring. Moreover, we claim, it is universal minimal in the sense of the following proposition ([2, Section 6]):

Proposition 7.1. If $\mathcal{N}$ is any map that is a covering of $\mathcal{M}$ and which is I-colorable, then $\mathcal{N}$ must itself cover $I \succ \mathcal{M}$.

An example of a double covering is $\mathcal{N} = R \succ \mathcal{M}$ when $\mathcal{M}$ is a non-orientable map. Because $\mathcal{M}$ does not have an R-coloring, $\mathcal{N}$ is truly a double cover of it, and since $\mathcal{N}$ then *does* have an R-coloring, it must be orientable and so $\mathcal{N}$ is the orientable twofold cover (some say the *canonical* cover) of $\mathcal{M}$.

In Figure 10 we show the $\{2\}$-double cover of the tetrahedron. Every face is a triangle and there are six triangles around each vertex. This map lies on the torus.
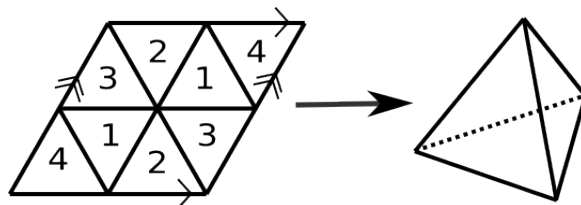

Figure 10: The $\{2\}$-double cover of the tetrahedron

Another example of the double construction in the literature comes from Sherk's 1962 paper [14]. At that time few non-toroidal chiral maps were known. Sherk constructed an infinite family of such maps on an infinite number of different surfaces, using a simple and clever technique. Starting with a chiral map $\mathcal{M}$ on the torus, with triangles as faces and six of them incident to each vertex, he made a new map $\mathcal{N}$ in the following way. For each vertex v of $\mathcal{M}$, he created vertices v_0, v_1; for each edge $\{u, v\}$, he made edges $\{u_0, v_1\}$ and $\{u_1, v_0\}$; and each triangular face $[u, v, w]$ became a hexagonal face $[u_0, v_1, w_0, u_1, v_0, w_1]$. This map has hexagonal faces, still six meeting at a vertex. Thus if $\mathcal{M}$ has D vertices, $2D$ faces and $3D$ edges, then $\mathcal{N}$ has $2D$ vertices, $2D$ faces and $6D$ edges. It follows that $\mathcal{N}$ must lie on the orientable surface of characteristic $-2D$, i.e., of genus $D - 1$. From our point of view, we can construct $\mathcal{N}$ as $\{0\} \succ \mathcal{M}$.

7.1 Properties of double covers

We present here some of the facts about I-doubles. When the proofs are straightforward, they are omitted.

Proposition 7.2. If $\mathcal{M}$ has a I-coloring then so does any $J \succ \mathcal{M}$.

Proposition 7.3. If $\mathcal{M}$ has a J-coloring then any $I \succ \mathcal{M}$ is isomorphic to $(I \Delta J) \succ \mathcal{M}$.

Proof: For any map $\mathcal{M}$, $I \succ \mathcal{M}$ is the smallest covering of $\mathcal{M}$ which admits an I-coloring. If $\mathcal{M}$ has a J-coloring, then $I \succ \mathcal{M}$ has a J-coloring as well. Thus, by Proposition 4.1, $I \succ \mathcal{M}$ admits an $(I \Delta J)$-coloring, and whether $\mathcal{M}$ has an I-coloring or not, this covering of $\mathcal{M}$ must be minimal of those admitting an $(I \Delta J)$-coloring. Thus $I \succ \mathcal{M} \cong (I \Delta J) \succ \mathcal{M}$. ∎

The following two results summarize the examples given above.

Proposition 7.4. If $\mathcal{M}$ is non-orientable, then $R \succ \mathcal{M}$ is isomorphic to the orientable double covering of $\mathcal{M}$.

We say that a map $\mathcal{N}$ is the *Sherk double cover* of a non-vertex-bipartite map $\mathcal{M}$ whenever $\mathcal{N}$ can be obtained from $\mathcal{M}$ by the above procedure used by Sherk to obtain the chiral maps with hexagonal faces.

Proposition 7.5. If $\mathcal{M}$ is not vertex-bipartite, then $\{0\} \succ \mathcal{M}$ is isomorphic to its Sherk double cover. Moreover, the underlying graph of $\{0\} \succ \mathcal{M}$ is the canonical bipartite double cover of the underlying graph of $\mathcal{M}$.

Remark 7.6. Note that $\{0\}, \{1\}, \{2\}$ generate $\mathcal{P}$ under Δ, and so if $\mathcal{M}$ is any map, then $\mathcal{N} = \{0\} \succ (\{1\} \succ (\{2\} \succ \mathcal{M}))$ has $T(\mathcal{N}) = \mathcal{P}$, and so $\mathcal{N}$ is an un-double-able map in our context.

Proposition 7.7. For any map $\mathcal{M}$ and any subset I of R, $T(I \succ \mathcal{M}) = \langle T(\mathcal{M}), I \rangle$

Proof: This is obvious if $I \in T(\mathcal{M})$. If not, it is clear that $\langle T(\mathcal{M}), I \rangle$ is contained in $T(I \succ \mathcal{M})$, and so we need only to show that the opposite containment holds. This is equivalent to showing that if $J \in T(I \succ \mathcal{M})$ and $J \notin T(\mathcal{M})$, then $I \Delta J \in T(\mathcal{M})$ (see Proposition 7.3).

Let $\mathcal{N} = I \succ \mathcal{M}$ and note that, if $J \in T(\mathcal{N})$, then $\mathcal{N}$ is a J-colorable double cover of $\mathcal{M}$; by Proposition 7.1, $\mathcal{N}$ must be isomorphic to $\mathcal{N}' = J \succ \mathcal{M}$. Because $J \notin T(\mathcal{M})$, there is some cycle (f, W) which is J-inconsistent. If any such cycle is I-consistent, then in $\mathcal{N}$, the cycle (f_0, W) has the same length as (f, W), and so is also J-inconsistent, contradicting $J \in T(\mathcal{N})$. Thus every cycle which is J-inconsistent is also I-inconsistent. A similar argument in $\mathcal{N}'$ shows that every cycle in $\mathcal{M}$ which is I-inconsistent is also J-inconsistent. Therefore every cycle is either consistent for both, or inconsistent for both. By Lemma 4.5, every cycle in $\mathcal{M}$ is $I \Delta J$-consistent, and then $(I \Delta J) \in T(\mathcal{M})$, as required. ∎

7.2 Which maps are I-double covers?

We want to generalize the result from [5] which says, loosely speaking, that a reflexible map is a 2-fold orientable cover of a non-orientable reflexible map if and only if its rotation group contains an involutory element which conjugates each of the generators to its inverse.

To re-phrase that in our context, a reflexible map $\mathcal{N}$ which has an R-coloring is an R-double covering of some reflexible map $\mathcal{M}$ which does not have such a coloring if and only if the subgroup of $Aut(\mathcal{N})$ which preserves colors in the coloring contains an involutory element which conjugates each of the generators (of the color-preserving group) to its inverse.

If $\mathcal{N}$ is I-colorable, let a be either of the I-colorings of $\mathcal{N}$ and define $C^+(\mathcal{N}, I)$ to be the subgroup of $C(\mathcal{N})$ consisting of all w such that $a(f) = a(fw)$ for all $f \in \Omega$. Because $\mathcal{N}$ is I-colorable, this group has index 2, and so is normal, in $C(\mathcal{N})$.

Note that Theorem 7.8 generalizes the result from [5] not only to all I-colorings, but also to maps with no assumption on their automorphism group.

Theorem 7.8. A map $\mathcal{N} = (\Omega, [r_0, r_1, r_2])$ is $I \succ \mathcal{M}$ for some non-I-colorable map $\mathcal{M} = (\Theta, [s_0, s_1, s_2])$ if and only if $\mathcal{N}$ is I-colorable and there exists $i \in I$ and $w \in C^+(\mathcal{N}, I)$ of order 2 such that for every element $c \in C^+(\mathcal{N}, I)$, $wcw = r_i c r_i$.

Proof: First, assume $\mathcal{N} = I \succ \mathcal{M}$ with $\mathcal{M}$ a non-I-colorable map. We will prove the existence of w. The group $C^+(\mathcal{N}, I)$ has index 2 in $C(\mathcal{N})$, and there is an isomorphism $f : C^+(\mathcal{N}, I) \to C(\mathcal{M})$ which is the letter-to-letter projection, $f(r_{i_0} r_{i_1} \cdots r_{i_k}) = s_{i_0} s_{i_1} \cdots s_{i_k}$ whenever the word $r_{i_0} r_{i_1} \cdots r_{i_k}$ is in $C^+(\mathcal{N}, I)$ and so has an even number of indices in I.

Because $\mathcal{M}$ is not I-colorable, I is not empty. So consider any $i \in I$ and let $w = f^{-1}(s_i)$. Note that w is an involution since f is an isomorphism. This w is not r_i, since $r_i \notin C^+$; it is, instead, expressible as some longer product of r_j's. Then for any $c \in C^+(\mathcal{M})$ we have $f(wcw) = s_i f(c) s_i$, while $f(r_i c r_i) = s_i f(c) s_i$, since $r_i c r_i$ is in C^+, and f acts there as letter-to-letter projection. Since f is an isomorphism and hence one-to-one, we conclude that $wcw = r_i c r_i$.

Now assume we have $w \in C^+(\mathcal{N}, I)$ of order 2 and $i \in I$ such that for every element $c \in C^+(\mathcal{N}, I)$, $wcw = r_i c r_i$. In particular, with $c = w$, we have $w = www = r_i w r_i$. From this we see that $wr_i = r_i w$. We then construct $\mathcal{M}$ by identifying flags in the following consistent way: let $u = wr_i$, so that u is an involution. We shall identify flag $f \in \Omega$ with fu. Because u is an involution, this identification is unambiguous, and each flag is identified with exactly one other. Since w preserves coloring and r_i changes it, we are identifying a white flag with a black flag. These equivalence classes are the flags of $\mathcal{M}$.

Then we have to define the connectors s_0, s_1 and s_2 of $\mathcal{M}$. Notice that for every $c \in C^+(\mathcal{N}, I)$ we have $u^{-1} c u = c$, so u commutes with all of $C^+(\mathcal{N}, I)$. But $C^+(\mathcal{N}, I)$ has index 2, and $u \notin C^+(\mathcal{N}, I)$ so u commutes with all of $C(\mathcal{N})$. Then we define $([f, fu])s_j = [fr_j, fur_j = fr_j u]$. The s_i's are well defined and we can easily see s_0 and s_2 commute, since r_0 commutes with r_2.

To see that $\mathcal{N}$ is the I-double cover of $\mathcal{M}$, identify each black flag f of $\mathcal{N}$ with the flag $[f, fu]_1$ of $I \succ \mathcal{M}$, and each white flag f with the flag $[f, fu]_0$. ■

In order to use this theorem we would, in theory, need to prove the existence of w and observe that under conjugation, w acts in the same way as some r_i on *every* element of $C^+(\mathcal{N})$. But in practice, we only need to find w that acts in such a way for the *generators* of $C^+(\mathcal{N})$. The element r_i must not be in C^+ and so we will refer to it as an *external generator*. Table I summarizes information about each subset I of R. The second column (adapted from [8]) gives generators for $C^+(\mathcal{N})$ as a subgroup of $C(\mathcal{N})$. The third column gives the $r_i's$ which are not in C^+, and the last column shows the actions on the generators of C^+ under conjugation by this external generator.

8   Which subgroups appear on which surfaces?

The purpose of this section is to prove that every subgroup H of $\mathcal{P}$ occurs on every surface on which it *can* appear. To clarify that statement, we call a group which includes R an

Table I			
I	C^+ generators	external generator	action
$\{0\}$	$a = r_1, b = r_2, c = r_0 r_1 r_0$	r_0	$a \leftrightarrow c, b \leftrightarrow b$
$\{1\}$	$a = r_0, b = r_2, c = r_1 r_0 r_1, d = r_1 r_2 r_1$	r_1	$b \leftrightarrow d, a \leftrightarrow c$
$\{2\}$	$a = r_0, b = r_1, c = r_2 r_1 r_2$	r_2	$a \leftrightarrow a, b \leftrightarrow c$
$\{0,1\}$	$a = r_2, b = r_0 r_1$	r_0	$a \leftrightarrow a, b \leftrightarrow b^{-1}$
		r_1	$a \leftrightarrow b^{-1}ab, b \leftrightarrow b^{-1}$
$\{0,2\}$	$a = r_1, b = r_0 r_2, c = r_0 r_1 r_0$	r_0	$b \leftrightarrow b, a \leftrightarrow c$
		r_2	$b \leftrightarrow b, a \leftrightarrow bcb, c \leftrightarrow bab$
$\{1,2\}$	$a = r_0, b = r_2 r_1$	r_2	$a \leftrightarrow a, b \leftrightarrow b^{-1}$
		r_1	$a \leftrightarrow b^{-1}ab, b \leftrightarrow b^{-1}$
R	$a = r_0 r_1, b = r_1 r_2$	r_0	$a \leftrightarrow a^{-1}, b \leftrightarrow ab^{-1}a^{-1}$
		r_1	$a \leftrightarrow a^{-1}, b \leftrightarrow b^{-1}$
		r_2	$a \leftrightarrow b^{-1}a^{-1}b, b \leftrightarrow b^{-1}$

orientable subgroup of $\mathcal{P}$, and subgroups which do not contain R are *non-orientable*. So we want H to appear on S when they are both orientable or both non-orientable. This is almost true, as explained by the following Theorem:

Theorem 8.1. For every subgroup H of $\mathcal{P}$ and every surface S, there is a map $\mathcal{M}$ on S such that $T(\mathcal{M}) = H$ if and only if S and H are both orientable or both non-orientable, with the following exceptions:

1. The group $\{\emptyset, \{1\}, \{0,2\}, R\}$ does not appear on the sphere.

2. The group $\{\emptyset, \{1\}\}$ does not appear on the projective plane.

3. The group $\{\emptyset, \{0,2\}\}$ does not appear on the projective plane.

The proof of this claim is contained in and scattered through the remaining parts of this section. We will consider the 16 possible subgroups of $\mathcal{P}$ individually and in dual pairs. Such a proof would appear to require 16 constructions, but we can use Lemma 6.1 and Proposition 7.2 for some simplification. Moreover, the insertion techniques that we will introduce and the connected sum construction of Section 4.4 allow even more simplification. We consider the non-orientable cases first; then a simple argument, using the results of Section 7, takes care of the orientable ones.

8.1 $H = \{\emptyset\}$

Consider any triangulation of a non-orientable surface S. Refine the triangulation slightly by placing one new vertex inside some triangular face and connect it to the three vertices on its boundary so as to have one vertex of degree 3. Call this map $\mathcal{M}$. Because of the triangular faces and the degree-3 vertex, the Cheat Sheet of Section 5 shows us that none of the sets of size 1 or 2 are in the group, and because S is non-orientable, it cannot contain R either. Thus $T(\mathcal{M})$ must be the trivial group.

8.2 Insertion

In this subsection, we introduce the technique of *insertion*. We shall prove that any map $\mathcal{M}$ can be easily modified to produce a map which is face or vertex bipartite and also face or vertex pseudo-orientable.

Let us assume, for example, that we wish to modify $\mathcal{M}$ to construct a vertex-bipartite map. Color the vertices of $\mathcal{M}$ at random. For the edges for which both vertices are the same color, add a vertex (of degree 2) to split the edge in two.

The same idea can be applied if we wish vertex pseudo-orientability. We may assign orientations to the vertices at random and then split the edges whose endpoints do not match into two edges. Call this process *vertex-insertion*. This can also be used to turn a map which *is* vertex-bipartite or VPSO into one that is *not*. Note that vertex-insertion does not modify the status of either face-bipartiteness or FPSO of $\mathcal{M}$.

Dually, we may replace an edge by a face bounded by two edges to get either face-bipartition or FPSO. The following picture illustrates this for face pseudo-orientation:
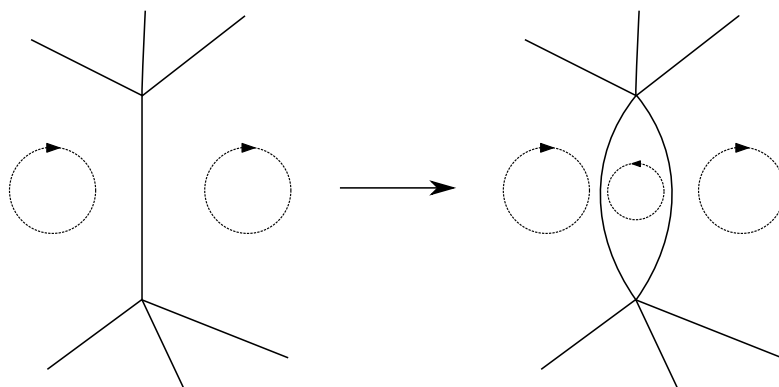

Figure 11: Adjusting one edge

We begin with an arbitrary map on a surface S, and an arbitrary orientation of each face. If there is an edge, as in Figure 11a, where the assigned arrows flow in opposite directions along the edge, replace the edge with a face bounded by a pair of edges as shown in 11b. Orient the resulting lune with an arrow as shown. This gives a map on the same surface with one fewer edge for which the given arrows fail to give a face pseudo-orientation. Continuing this process eventually gives us a map on S admitting a face pseudo-orientation.

Let us call this process *face-insertion*. Here, we used it to produce a FPSO map on any surface. We can use it equally well to turn any map into a face-bipartite map on the same surface. As with vertex insertion, this same process can also be used to convert a map which *is* face bipartite or FPSO to one which is *not*.

Fortunately, when doing vertex insertion to make a map vertex-bipartite there is no danger of accidentally making it vertex pseudoorientable ($\{1, 2\}$-colorable), or vice-versa. If a map is one of vertex bipartite, vertex pseudoorientable and orientable, then the

other two are equivalent. The same is true for maps which are one of face biparte, face pseudoorientable and orientable. Orientability cannot be modified without modifying the surface.

For many groups we can use these two operations to produce maps with these groups on any non-orientable surface. We will show one of these in detail and summarize the rest.

Consider the non-orientable group $H = \{\emptyset, \{0, 1\}\}$. For any non-orientable surface S, we choose any map on S, and, using face-insertion as above, construct a map which is FPSO. Now we must be certain that the group of the map is no bigger than H. The other cosets of H in $\mathcal{P}$ are $H_2 = \{\{0\}, \{1\}\}, H_3 = \{\{0, 2\}, \{1, 2\}\}, H_4 = \{\{2\}, \{0, 1, 2\}\}$. By inserting one vertex on one edge, if necessary, we can cause the map to have a face with odd degree, eliminating $\{0\}$ and $\{1, 2\}$ and hence all of $H_2 \cup H_3$ from the group. The non-orientability of the surface eliminates $R = \{0, 1, 2\}$ and hence all of H_4. Thus the final form $\mathcal{M}$ of the map has $T(\mathcal{M}) = \{\emptyset, \{0, 1\}\}$, as required.

Notice that the group $T(D(\mathcal{M}))$ is then $\{\emptyset, \{1, 2\}\}$.

We produce maps having most of the other non-orientable groups in a similar way. Starting with an arbitrary map on a non-orientable surface we perform insertions in the following ways:

- Use face-insertion to force FPSO and vertex-insertion to force VPSO. The result is a map whose group is $\{\emptyset, \{0, 1\}, \{0, 2\}, \{1, 2\}\}$.

- Use face-insertion to force FPSO and then vertex-insertion to force it to be vertex-bipartite. The resulting map has group $\{\emptyset, \{0\}, \{1\}, \{0, 1\}\}$. Its dual has group $\{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$.

- Use face-insertion to force the map to be face-bipartite and vertex-insertion to force it to be vertex-bipartite. The resulting map has group $\{\emptyset, \{0\}, \{2\}, \{0, 2\}\}$.

- Use face-insertion to force face-bipartiteness and vertex-insertion in order to force some face to have an odd number of sides. The group must be just $\{\emptyset, \{2\}\}$. Its dual then has group $\{\emptyset, \{0\}\}$.

These constructions take care of all non-orientable groups except for two: $\{\emptyset, \{1\}\}$ and $\{\emptyset, \{0, 2\}\}$.

8.3 The groups $H = \{\emptyset, \{1\}\}$ and $H = \{\emptyset, \{0, 2\}\}$

In order to cope with the two non-orientable exceptional groups, we apply the technique of connected sums of maps, developed in Subsection 4.4.

Lemma 8.2. For each group $H \leqslant \mathcal{P}$ with $R \notin H$ and $|H| = 4$, there exists a map $\mathcal{M}$ on every non-orientable surface such that all the following hold:

- $T(\mathcal{M}) = H$.

- There is a face F with exactly 2 sides and exactly 2 vertices.

- $T(\mathcal{M} \setminus F) = H$. That is, the flags of $\mathcal{M}$ that are not in F are *not* I-colorable for every $I \notin H$.

Proof: Let S be an arbitrary non-orientable surface. To construct a map on S with the desired properties, first consider any map for which no face shares a vertex with itself. Because the surface is non-orientable, we may choose an arbitrarily fine subdivision for which there is a Möbius band of faces which does not use all the faces or edges. Then apply the operations described in subsection 8.2 to make the map h-colorable for each h in H. This is possible since any group H with 4 elements with $R \notin H$ must contain two of the following: vertex-bipartite, face-bipartite, FPSO or VPSO, with the other non trivial element being the symmetric difference of the other two.

Now consider an edge that is not in any of the faces of the chosen Möbius band and convert it into three edges, as with face-insertion but with three edges instead of two. This operation does *not* change the status of vertex-bipartition, face-bipartition, VPSO or FPSO as required. Now pick either of the two newly created faces and call it F. Since we still have a Möbius strip, $R \notin T(\mathcal{M} \setminus F)$, but every $I \in T(\mathcal{M})$ is in $T(\mathcal{M} \setminus F)$. Therefore, $T(\mathcal{M} \setminus F) = H$. ∎

Proposition 8.3. For $H = \{\emptyset, \{0, 2\}\}$ and $H = \{\emptyset, \{1\}\}$ we have that for every surface (except the projective plane) there exists a map $\mathcal{M}$ on that surface with $T(\mathcal{M}) = H$.

Proof: Let S be a surface with non-orientable genus $k \geqslant 2$. Consider maps as in the previous lemma $\mathcal{M}_1$ and $\mathcal{M}_2$ such that $\mathcal{M}_1$ lies on a surface with non-orientable genus $k - 1$ and $\mathcal{M}_2$ lies on the projective plane and such that:

- If $H = \{\emptyset, \{1\}\}$, then

$$T(\mathcal{M}_1) = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\} \text{ and } T(\mathcal{M}_2) = \{\emptyset, \{0\}, \{1\}, \{0, 1\}\}.$$

- If $H = \{\emptyset, \{0, 2\}\}$, then

$$T(\mathcal{M}_1) = \{\emptyset, \{0, 1\}, \{0, 2\}, \{1, 2\}\} \text{ and } T(\mathcal{M}_2) = \{\emptyset, \{0\}, \{2\}, \{0, 2\}\}.$$

Then $\mathcal{M}_1 \oplus \mathcal{M}_2$ (along the 2-sided faces as in the previous lemma) lives on the non-orientable surface with genus k and $T(\mathcal{M}_1 \oplus \mathcal{M}_2) = H$ (see Theorem 4.7).

Explicit constructions can be easily constructed and are avoided for brevity. ∎

This concludes the proof for the non-orientable case.

8.4 Orientable Groups

If the subgroup H contains R, then it is generated by R and some subgroup H' not containing R. Each orientable surface S is the orientable twofold cover of some non-orientable surface S'. By the preceding subsections, S' has a map $\mathcal{M}'$ whose group is

H' (unless S is the sphere, S' is the projective plane and H is $\{\emptyset, \{1\}, \{0,2\}, \{0,1,2\}\}$). Then $R \succ \mathcal{M}'$ is a map on S whose group is H. By Proposition 7.7, $T(\mathcal{M})$ is exactly H. Thus, with the one exception, for each subgroup H containing R, every orientable surface admits a map whose group is H. ∎

All that is left to prove now is that the three exceptions mentioned in the theorem are in fact exceptions – namely that there is no map on the specified surfaces with the specified group. We shall prove this first for the group $\{\emptyset, \{1\}, \{0,2\}, R\}$ on the sphere.

Lemma 8.4. There is no map on the sphere with group $\{\emptyset, \{1\}, \{0,2\}, R\}$.

Proof: Suppose $\mathcal{M}$ is one such map. Since $\{0\} \notin T(\mathcal{M})$, the underlying graph of $\mathcal{M}$ is not bipartite, which means it has a cycle C of odd length. By Jordan's Closed Curve Theorem, this divides the sphere into two connected components. Choose one of these two components and erase every edge and vertex exterior to it in order to form the deleted part into a single face whose boundary is C. This gives a new map $\mathcal{M}'$, also on the sphere. Since $\{1\} \in T(\mathcal{M})$, every face of $\mathcal{M}$ must have an even number of sides. Therefore $\mathcal{M}'$ has exactly one face with an odd number of sides. Then the underlying graph of $D(\mathcal{M}')$ has exactly one vertex of odd degree, a well-known impossibility. ∎

Corollary 8.5. No map $\mathcal{M}$ on the projective plane has $T(\mathcal{M}) = \{\emptyset, \{1\}\}$ or $T(\mathcal{M}) = \{\emptyset, \{0,2\}\}$.

Proof: If such a map were to exist, its orientable double cover would violate Lemma 8.4. ∎

The proof of Theorem 8.1 is now complete.

Corollary 8.6. No map on a non-orientable surface has an R-coloring. In all other cases, for any subset I of R, and any surface S, there is a map $\mathcal{M}$ on S which is I-colorable.

9   Open Questions

The constructions in Section 8 lead to quite general maps, although some of them may be somewhat degenerate in the viewpoint of other previous work. It may be the case that more exceptions in Theorem 8.1 are needed if we ask the extra requirement that the maps are polyhedral maps (that is, the intersection of the closure of two distinct faces is either empty, a single vertex, or a single edge. See [9]). However, this is beyond the scope of this paper.

In previous sections we considered colorings of the flags of maps insisting that we use only two colors. In doing so, we ensure that if a given flag is i-adjacent to another flag with the same color then all flags are the same color as their i-adjacent flags.

Consistent colorings with k colors can be defined following the idea of [1] and [13], that is, each flag is assigned a color in $\{1, \dots, k\}$ with the restriction that for every $i \in \{0, 1, 2\}$ and $j \in \{1, \dots, k\}$, the color of the i-adjacent flag to a flag f colored j depends on j but not on f.

When k-coloring a map consistently for $k \geqslant 3$ then it may not be true that a flag is i-adjacent to another flag with the same color if and only if each flag is i-adjacent to another flag with the same color. For example, we can color the flags of the pentagons of a pentagonal prism with color 1, the flags of the edges shared by two squares with color 2, and the remaining flags with color 3, as in Figure 12. The reader can easily verify that this is a consistent 3-coloring, however, flags colored 1 are 1-adjacent to flags colored 1, and flags colored 2 are 1-adjacent to flags colored 3.
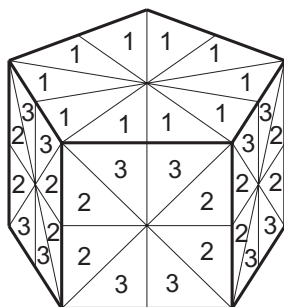

Figure 12: A 3-orbit map

It is unlikely that the concept of pseudo-orientability admits a generalization that preserves the connections between 2-colorings and pseudo-orientations of maps. However, k-colorings may still keep an interesting relation with k-fold covers of maps and maniplexes as well.

Some results in [8], [13], [1] relate symmetry with I-colorings. We suspect that many more results can be obtained linking symmetry with I-colorings.

Acknowledgements

References

[1] A. Breda d'Azevedo, A theory of restricted regularity of hypermaps. *J. Korean Math. Soc.* **43** (2006), no. 5, 991–1018.

[2] A. Breda d'Azevedo, G. Jones, Double coverings and reflexive abelian hypermaps, *Beiträge Algebra Geom.* **41** (2000), no. 2, 371–389.

[3] A. Breda d'Azevedo, D. Catalano; D. Rui, Regular pseudo-oriented maps and hypermaps of low genus. *Discrete Math.* **338** (2015), no. 6, 895–921

[4] F. Buekenhout. Handbook of Incidence Geometry: Buildings and Foundations. *Elsevier Science*, North-Holland, 1995.

[5] M. Conder, S. Wilson, Inner Reflectors and non-orientable regular maps, *Discrete Math.*, **307** (2007), 367–372.

[6] A. Hill, S. Wilson, Four constructions of highly symmetric graphs. *J. Combin. Theory Ser. B*, to appear.

[7] I. Hubard, From geometry to groups and back: The study of highly symmetric polytopes, *Doctoral dissertation*, York University, Toronto, Canada, (2007).

[8] I. Hubard, Two-orbit polyhedra from groups. *European J. Combin.* **31** (2010), no 3, 943–960.

[9] S. Jendrol, H. J. Voss, A local property of polyhedral maps on compact two-dimensional manifolds. *Graph theory (Dörnfeld, 1997). Discrete Math.* **212** (2000), no. 1-2, 111–120.

[10] G. Jones, J. S. Thornton, Operations on Maps, and Outer Automorphisms. *J. Combin. Theory Ser. B* **35** (1983), no. 2, 93–103.

[11] P. McMullen, E. Schulte, Abstract regular polytopes. *Encyclopedia of Mathematics and its Applications*, **92**. Cambridge University Press, Cambridge, (2002). xiv+551 pp.

[12] A. Orbanić, D.Pellicer, T. Pisanski, T. Tucker, Edge-transitive maps of low genus. *Ars Math. Contemp.* **4** (2011), no. 2, 385–402.

[13] A. Orbanić, D. Pellicer, A. Weiss, Map operations and k-orbit maps. *J. Combin. Theory Ser. A* **117** (2010), no. 4, 411–429.

[14] F. A. Sherk, A family of regular maps of type 6,6. *Canad. Math. Bull.* **5** (1962) 13–20.

[15] S. Wilson, Operators over Regular Maps, *Pacific J. Math.* **81** (1979), no. 2, 559–568.

[16] S. Wilson, Riemann Surfaces over Regular Maps, *Canad. J. Math.* **30** (1978), no. 4, 763–782.

[17] S. Wilson, Maniplexes: Part 1: Maps, Polytopes, Symmetry and Operators, *Symmetry* (2012), **4**(2), 265–275.