Symmetries of unlabelled planar triangulations[1]

Mihyun Kang[2] Philipp Sprüssel[2]

Institute of Discrete Mathematics
Graz University of Technology
Graz, Austria

{kang,spruessel}@math.tugraz.at

Submitted: Jun 2, 2016; Accepted: Jan 26, 2018; Published: Feb 16, 2018
Mathematics Subject Classifications: 05C10, 05C30, 57M15

Abstract

We derive a decomposition scheme of unlabelled triangulations rooted at a single cell, where the decomposition depends on whether the automorphism group of the triangulation contains reflections, rotations, or both. Furthermore, the decomposition scheme is *constructive* in the sense that for each of the three cases, there is a $k \in \mathbb{N}$ such that the scheme defines a one-to-k correspondence between the respective triangulations and their decompositions.

1 Introduction

Graphs embedded on a surface, called *maps*, have been among the most studied objects in graph theory and combinatorics [23, 24, 30, 31, 32], geometry [2, 4, 28, 29], discrete probability theory [3, 22], and theoretical physics [8, 15, 25].

Symmetries of maps play an important role in discrete geometry, enumerative combinatorics, and random sampling [2, 9, 19, 21, 29]. In his seminal work [31], Tutte conjectured that almost all planar *maps* (i.e. graphs *embedded* on the sphere) are asymmetric—a conjecture that was later proved by Richmond and Wormald [26].

Classically, when maps are considered in terms of enumerative combinatorics, they are given with a *rooting*, that is, a vertex, an edge, and a face that are mutually incident are fixed. This kind of rooting is also known as *Tutte rooting* [17, 30, 31]. Maps with Tutte rooting are intrinsically asymmetric, which simplifies the theory necessary to enumerate them. In contrast, enumerating *unrooted* maps [16, 18] requires a better understanding of their symmetries and of how they decompose into isomorphic parts.

---

[1]An extended abstract of this paper has been published in the Proceedings of the European Conference on Combinatorics, Graph Theory and Applications (EuroComb15), Electronic Notes in Discrete Mathematics (2015), 587–594.

[2]Supported by Austrian Science Fund (FWF): P27290 and W1230 II.

In the enumeration of graphs and maps, rooted or unrooted, decompositions play a central role. The most basic example is that every graph is the disjoint union of its components, which can in turn be decomposed into 2-connected blocks. For 2-connected graphs, there are two well-known decompositions. On one hand, there is the *ear-decomposition* that decomposes a 2-connected graph into a cycle and edge-disjoint paths. On the other hand, Tutte [33] proved that every 2-connected graph belongs to one of three disjoint graph classes. For each of these three classes, there is a unique decomposition of its graphs into smaller building blocks, where 3-connected graphs are one of the base cases.

All of the decompositions above can be reversed to construct graphs from their components, 2-connected blocks etc. However, there is a significant difference between the ear-decomposition and the others: while the other decompositions are *unique*, the number of ear-decompositions of a 2-connected graph is not easy to determine. In an enumerative sense, Tutte's decomposition is therefore much better suited to construct 2-connected graphs. For if all base cases are known, then the number of 2-connected graphs can be derived from that information. Observe that it would also be possible to derive the number of 2-connected graphs from the base cases if the decomposition were not unique, but every 2-connected graph had *the same number of decompositions*.

Motivated by this observation, we introduce the following notation. Suppose that $\mathcal{A}, \mathcal{B}$ are classes of graphs or maps and let $\mathcal{D}$ consist of finite collections of elements of $\mathcal{B}$. We say that a *decomposition scheme* of the elements of $\mathcal{A}$ is a function that assigns to each $A \in \mathcal{A}$ a non-empty set $D_A \subset \mathcal{D}$. Each element of D_A is called a *decomposition* of A; the elements of $\mathcal{B}$ are the *building blocks*. A decomposition scheme is called *constructive decomposition* if $\mathcal{D} = \bigcup_A D_A$ and furthermore, for all $A_1, A_2 \in \mathcal{A}$, we have $D_{A_1} \cap D_{A_2} = \emptyset$ and $|D_{A_1}| = |D_{A_2}| < \infty$. In other words, a constructive decomposition defines a one-to-k correspondence (for some $k \in \mathbb{N}$) from $\mathcal{A}$ to $\mathcal{D}$, that is, every $D \in \mathcal{D}$ is the decomposition of a unique element of $\mathcal{A}$, and each element of $\mathcal{A}$ has precisely k decompositions.

In this paper, we shall study *planar triangulations*, that is, planar maps in which each face is bounded by a triangle. Symmetries of triangulations can be related to isometries of the sphere. Finite subgroups of the isometry group $O(3)$ of the sphere have been well studied (see e.g.. [10] for an overview). In particular, isometries can be classified to be either rotations, reflections, or glide reflections. A planar triangulation can then be constructed by triangulating the orbifold (the quotient space of the sphere obtained by factoring out the symmetry group). However, this is *not* a constructive decomposition, because a given triangulation might be obtained by several different ways of triangulating the orbifold.

A different approach was taken by Tutte [32], who studied the symmetries of unrooted planar triangulations and proved that almost all of them are asymmetric. He also derived (non-constructive) decompositions of triangulations with reflective symmetries and those of triangulations with rotative symmetries with the additional property that the order of the automorphism is prime.

In this paper we consider triangulations rooted *at a single cell*, which might be a vertex, an edge, or a face. This is motivated by the problem of enumerating unlabelled cubic planar graphs. Based on a general strategy by Chapuy, Fusy, Kang, and Shoilekova [11], this problem can be reduced to enumerating triangulations rooted at a single cell. See Section 6 for more details.

Observe that by rooting triangulations at a single cell, we exclude glide reflections from the automorphism group, since glide reflections do not have fixed-points. For these rooted triangulations, we complete and strengthen the decompositions of Tutte [32] mentioned above in the sense that (i) we develop decomposition schemes for *all* possible symmetries (reflective, rotative, *or both types of symmetries*) and (ii) for each of the three cases, the scheme is a constructive decomposition, which was not the case for Tutte's decomposition.

Our decomposition of triangulations with reflective symmetries will be similar to Tutte's decomposition, but furthermore, we prove it to be constructive—a proof in which the decomposition of triangulations with both reflective and rotative symmetries will play an essential role. For triangulations with rotative symmetries, Tutte's decomposition is *not* a constructive decomposition and moreover, Tutte only studied the case when the order of the rotations is prime. To cover all cases, we introduce a new class of auxiliary graphs (called fyke nets), with the help of which we derive a constructive decomposition for rotative symmetries. Triangulations with both reflective and rotative symmetries have not been considered before.

The constructive decomposition of triangulations will be composed of two parts: (i) the characterisation of the basic building blocks (consisting of building frames and flagstones) and (ii) the description of how flagstones will be inserted into building frames in order to construct a triangulation with a specified symmetry.

The building frames capture the essential characteristics of each type of symmetries that a triangulation T has. There are three classes of building frames, called girdles, fyke nets, and skeletons, each one corresponding to reflective symmetries, rotative symmetries, or both reflective and rotative symmetries of T. The flagstones consist of special classes of planar maps (called near-triangulations) which encompass each type of symmetries of T. There are three classes of flagstones, each corresponding to one of the three types of symmetries. In each type of symmetries that T has, T will contain a unique subgraph G from the respective class of building frames. Vice versa, we will show that T can be constructed from G by inserting flagstones into some of the faces of G (see Section 2 for a precise definition of this construction). The process of inserting flagstones into faces is similar to that used to obtain stack triangulations, objects that proved to have various applications in geometry [1, 5, 13].

This paper is organised as follows. We start by stating the necessary notation and basic facts in Section 2. In Sections 3 to 5, we then derive the constructive decompositions into building frames and flagstones of triangulations with reflective symmetries, with rotative symmetries, and with both types of symmetries.

2 Preliminaries

A *planar embedding* of a graph G is a drawing of G on the 2-dimensional sphere $\mathcal{S}$ without crossing edges. Given an embedding of G, the image of G decomposes $\mathcal{S}$ into connected components, the *faces* of the embedding. If every face is homeomorphic to an open disc, the embedding is called *(planar) map.* A *triangulation* is a map all of whose faces are bounded by triangles. We refer to the vertices, edges, and faces of a map as its *cells* of dimension 0, 1, and 2, respectively. Two cells of different dimension are called *incident* if one is contained in the (topological) boundary of the other. Two cells of the same dimension are *adjacent* if there is a third cell incident with both.

All graphs and maps considered in this paper are *unlabelled* (i.e. are isomorphism classes of labelled graphs) and *simple* (i.e. the end vertices of an edge are always distinct and no two edges have the same two end vertices). Classically, maps are *rooted* by fixing a vertex, an edge, and a face that are mutually incident. We call such a rooting a *Tutte rooting.* In this paper, triangulations will mostly be given not with a Tutte rooting, but with a rooting at *either* a vertex, an edge, or a face.

Call a triangulation *trivial* if it has at most four vertices, so its underlying graph is a triangle or the complete graph K_4 on four vertices. In view of the results of this paper, these trivial triangulations represent degenerate cases of the structures considered. We will thus consider only non-trivial triangulations for the rest of this paper. Note that in a non-trivial triangulation, no two faces have the same set of vertices.

An *isomorphism* between planar maps G, H is a bijection $\varphi\colon G \to H$ that maps each cell to a cell of the same dimension and preserves incidencies. If G, H are given with the same type of rooting (i.e. they either both have a Tutte rooting or are both rooted at a single cell), then an isomorphism is additionally supposed to map the root(s) of G to the root(s) of H. If $G = H$, then we call φ an *automorphism.* It is well known that for a planar map with Tutte rooting, the identity is the only automorphism.

A cell c of a planar map G is *invariant* under a given automorphism φ if $\varphi(c) = c$. We also say that φ *fixes* c.[1] Analogously, we call a set A of cells invariant if $\varphi(A) = A$; note that the elements of A do not have to be invariant themselves. If G is rooted at c, then c is invariant under all automorphisms by definition. In this case, we denote the group of these automorphisms by $\mathrm{Aut}(c, T)$.

Throughout this paper, we let T be a non-trivial triangulation and choose a cell c_0 as the root of T. The fact that T is non-trivial implies that T is 3-connected and thus T is the unique embedding (up to automorphisms of the sphere) of its underlying graph G by Whitney's theorem [37]. The theorem of Mani [20] states that every 3-connected planar graph H is the skeleton of a convex polyhedron P in $\mathbb{R}^3$ and furthermore, the automorphisms of H are precisely the isometries of P. Therefore, automorphisms of T correspond to isometries of a polyhedron, which immediately implies the following.

Lemma 1. *Every automorphism in* $\mathrm{Aut}(c_0, T)$ *is uniquely determined by its action on the cells incident with* c_0.

The only non-trivial isometries of a polyhedron that have fixed-points are *rotations* (around an axis through c_0) and *reflections* (across a plane through c_0). We will distinguish

[1] Note that φ fixes an edge uv as soon as $\varphi(\{u, v\}) = \{u, v\}$; the vertices u, v themselves do not have to be invariant under φ. Similarly, a face can be invariant although none of its incident vertices or edges is invariant.

whether $\mathrm{Aut}(c_0, T)$ contains reflections, rotations, or both. Correspondingly, we say that T has *reflective symmetry*, *rotative symmetry*, or both.

A planar map N with a Tutte rooting consisting of a face f_N, an edge e_N, and a vertex v_N is called a *near-triangulation* if f_N is bounded by a cycle of any length $\geqslant 3$ while all other faces are bounded by triangles (see Figure 1). The root face f_N is called the *outer face* of N, all vertices and edges on its boundary—in particular the root vertex v_N and the root edge e_N—are called *outer vertices* or *outer edges* of N, respectively. All other vertices, edges, and faces of N are its *inner vertices*, *inner edges*, or *inner faces*, respectively. Deleting the outer face of a near-triangulation yields a bijection between the class of near-triangulations and the class of triangulations of polygons with a root edge on the boundary of the polygon and a root vertex incident with the root edge.

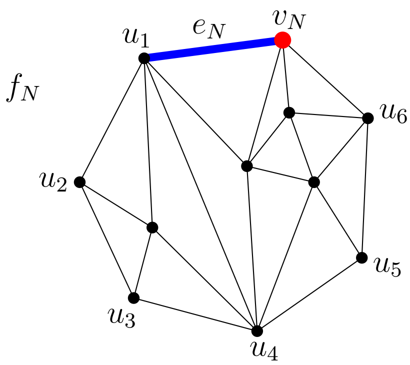

Figure 1: A near-triangulation N with root face f_N, root edge e_N (blue), and root vertex v_N (red). The outer vertices of N are $v_N, u_1, \dots, u_6$; the outer edges are $e_N = v_N u_1, u_1 u_2, \dots, u_5 u_6, u_6 v_N$.

If a map G contains a cycle C and an edge e that does not belong to C but connects two vertices of C, then we call e a *chord of* C. An inner edge of a near-triangulation N is a *chord of* N if it is a chord of the cycle bounding the outer face (e.g. the edge $u_1 u_4$ in Figure 1 is a chord of N).

In order to describe a constructive decomposition of triangulations, we shall use the operation of *inserting near-triangulations* into faces of a given planar map. To make this operation precise, let N be a near-triangulation with m outer vertices and let G be a planar map rooted at a cell c; denote by $\mathcal{S}_N$ and $\mathcal{S}_G$ the spheres on which N and G are embedded, respectively. Suppose that f is a face of G that is bounded by a cycle of length m; let e be an edge on the boundary of f and let v be one of the end vertices of e. We obtain a new planar map H as follows: deleting the outer face of N from the sphere $\mathcal{S}_N$ results in a space D_N homeomorphic to the unit disc; similarly, deleting f from the sphere $\mathcal{S}_G$ results in a space D_G homeomorphic to the unit disc. Note that by construction the boundary C_N of D_N (respectively the boundary C_G of D_G) is the boundary of the outer face of N (respectively that of f) and thus the point set of a cycle of length m. Let $\sigma \colon C_N \to C_G$ be a homeomorphism that

- maps vertices to vertices;

- maps the root vertex v_N of N to v; and
- maps the (point set of the) root edge e_N of N to (the point set of) e.

The quotient space $(D_N \cup D_G)/\sigma$ obtained from the union $D_N \cup D_G$ by identifying every point $x \in C_N$ with $\sigma(x) \in C_G$ is a sphere on which a graph H is embedded. We say that H is obtained from G by *inserting N into f at v and e* (see Figure 2). If the root c of G is *not* the face f, then c is also a cell of H and we consider H to be rooted at c. On the other hand, if $c = f$, then we choose the unique inner face of N incident with e_N as the root of H.

Definition 2. If T is a triangulation and G is a 2-connected subgraph of T, then T can be obtained from G by inserting near-triangulations into several of its faces in the following manner. Let $\mathcal{F}$ be the set of faces of G. Suppose that for each face $f \in \mathcal{F}$, we choose on its boundary C_f an edge e_f and one of its end vertices v_f. On $C_f \cup f$, T induces a triangulation of a polygon; we denote its corresponding near-triangulation by N_f and say that N_f is the near-triangulation *induced by T on f, rooted at v_f and e_f*. Then T is obtained from G by inserting N_f into f at v_f and e_f for each $f \in \mathcal{F}$ (in an arbitrary order). Vice versa, this gives us a decomposition of T into $(G; N_f, f \in \mathcal{F})$.

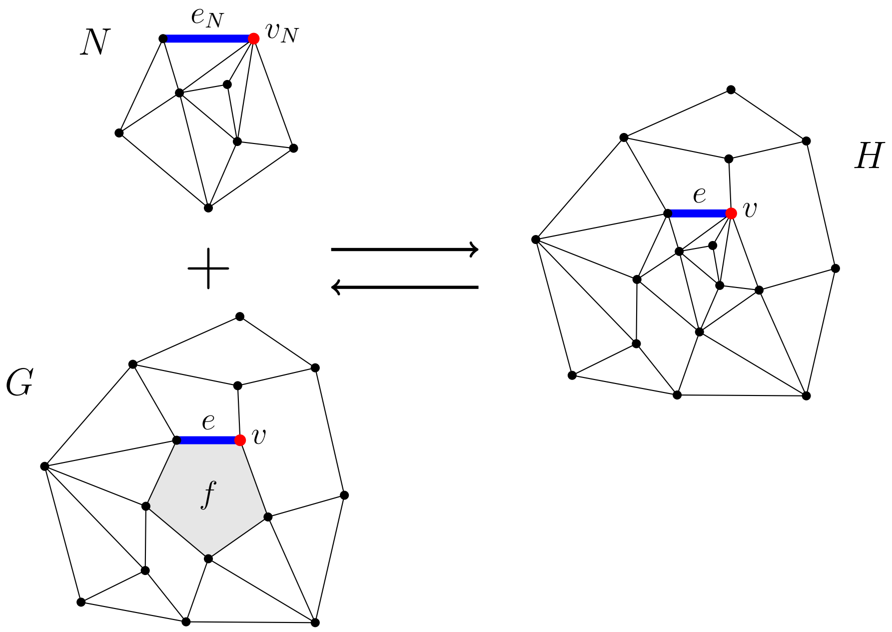

Figure 2: Inserting a near-triangulation N (e.g. 'flagstone') into the face f of a map G (e.g. 'building frame') resulting in a map H; vice versa, this is a decomposition of H into $(G; N_f = N)$.

For every cell c of T of a given dimension k, the numbers of incident cells of dimensions $k+1 \pmod 3$ and $k+2 \pmod 3$ are the same. We call this number the *degree* of c and denote it by $\deg(c)$. Clearly, for a vertex this notion of degree equals the graph theoretical definition; every edge has degree 2; every face of T has degree 3.

Let c, c' be cells of T. A path P in (the underlying graph of) T is called a path *from c to c'* if

- for the first vertex u of P we either have $c = u$ or c is incident with u, but with no other vertex of P; and
- for the last vertex v of P we either have $c' = v$ or c' is incident with v, but with no other vertex of P.

The *length* of P is the number of edges in P.

Given a cell c of T, the set of cells incident with c has a cyclic order $(c_1, c_2, \ldots, c_{2\deg(c)})$ in which two cells are consecutive if and only if they are incident in the triangulation (see Figure 3). We shall think of $c_1, c_2, \ldots$ to lie around c in counterclockwise order; with this convention, said order is unique (up to cyclic permutation). Two cells c_α, c_β with $\alpha, \beta \in \{1, 2, \ldots, 2\deg(c)\}$ are said to *lie opposite* at c if $|\alpha - \beta| = \deg(c)$. We observe that if c is a face, then its boundary is a triangle and every vertex v of this triangle is opposite at c to the edge of the triangle that is not incident with v. If c is an edge, then its two incident faces lie opposite at c and so do its end vertices. If c is a vertex, the situation depends on the parity of $\deg(c)$: for $\deg(c)$ even, every incident edge lies opposite to another incident edge while every face lies opposite to a face. For $\deg(c)$ odd, every edge lies opposite to a face.

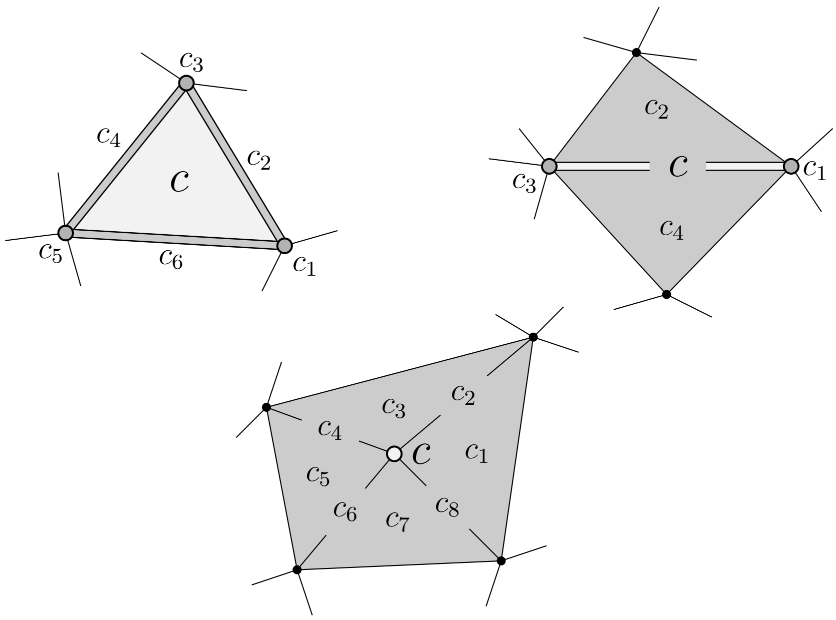

Figure 3: A cyclic order $(c_1, c_2, \ldots, c_{2\deg(c)})$ of the cells incident with a cell c.

The cyclic order of the cells incident with c_0 provides a purely combinatorial way of characterising reflections and rotations. If $\varphi \in \mathrm{Aut}(c_0, T)$ is not the identity, then

(i) φ is a reflection if and only if it changes the orientation of the cyclic order of the cells incident with c_0;

(ii) φ is a rotation if and only if it does not change the orientation of the cyclic order.

Reflections and rotations can also be defined via invariant cells incident with c_0. Indeed, $\varphi \in \mathrm{Aut}(c_0, T)$ is a reflection if and only if it fixes precisely two cells incident with c_0; these cells lie opposite at c_0. On the other hand, φ is a rotation if and only if it fixes no cell incident with c_0.

The characterisation of isometries of the sphere immediately yields the following characterisation of subgroups of $\mathrm{Aut}(c_0, T)$.

Theorem 3. *For every subgroup H of $\mathrm{Aut}(c_0, T)$ that contains at least one non-trivial automorphism, the following holds.*

(i) If H contains a reflection but no rotation, then it is isomorphic to the 2-element group $\mathbb{Z}_2$.

(ii) If H contains $k \geqslant 1$ rotations but no reflection, then it is isomorphic to the cyclic group $\mathbb{Z}_{k+1}$ where $k+1$ is a divisor of $\deg(c_0)$.

(iii) If H contains both reflections and rotations, then it is isomorphic to a dihedral group D_n where $n \geqslant 2$ is a divisor of $\deg(c_0)$.

3 Reflective symmetries

3.1 Building blocks

In this section, we suppose that $\mathrm{Aut}(c_0, T)$ contains a reflection φ.

Our first lemma is a structural result that was first obtained by Tutte [32]. We present (a modified version of) its proof in the appendix for the sake of completeness.

Lemma 4. *There is a cyclic sequence $(c_0, \ldots, c_\ell)$ of pairwise distinct cells such that for each cell c in the sequence the following holds.*

(i) c is invariant under φ;

(ii) the predecessor and the successor of c in the sequence are incident with c and lie opposite at c; and

(iii) no other cell in the sequence is incident with c.

For every edge in the sequence from Lemma 4, its predecessor and its successor are either its two end vertices or its two incident faces. Every face f in the sequence is preceded and followed by a vertex and its opposite edge on the boundary of f.

The invariant cells from Lemma 4 play a central role in the constructive decomposition of T in the case of a reflective symmetry: we will shortly see that these cells are the only cells invariant under φ and thus, they provide a way to define a unique subgraph of T that will be the basic building frame in our constructive decomposition.

Definition 5 (Girdle)**.** Let G be the planar map obtained by taking the union of all vertices and edges that either lie in the sequence $(c_0, \ldots, c_\ell)$ from Lemma 4 or on the boundary of a face in this sequence. We call G the *girdle* with respect to φ. Its cells from the cyclic sequence are called *central cells*, the other ones (which are only part of G because they lie on the boundary of a face from the sequence) are called *outer cells*. By construction, every face in the sequence is also a face of G (and hence a central cell); the other faces of G are called its *hemispheres*. For every face in the sequence, precisely one of the edges on its boundary is a central cell and so is the other face incident with this edge. The union of such two faces and their boundaries is called a *diamond*. Note that every girdle has at least two central vertices; let $j(G)$ be the smallest index for which $c_{j(G)}$ is a vertex (e.g. $j(G) = 3$ in Figure 4).

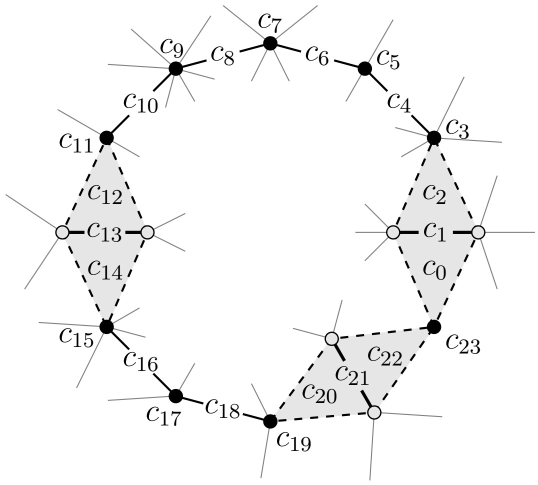

Figure 4: The sequence of cells from Lemma 4. The vertices in this picture, together with all black and all dashed edges, form the girdle G of T (see Definition 5). The central cells of the girdle are the black vertices, the black edges, and the gray faces. The outer cells are the gray vertices and the dashed edges. This girdle has three diamonds.

It is easy to see that there are precisely two hemispheres.

Lemma 6 ([32])**.** *The girdle G has exactly two hemispheres f_1, f_2.*

For each hemisphere f_i, if we fix a vertex v_i and an incident edge e_i on the boundary of f_i, then T induces on f_i a near-triangulation rooted at v_i and e_i (see Definition 2). We choose the root vertices and edges as follows.

Definition 7. Set $v_1 = v_2 := c_{j(G)}$. Then $c_{j(G)+1}$ is either an edge or a face.

(i) If $c_{j(G)+1}$ is an edge, we set $e_1 = e_2 := c_{j(G)+1}$;

(ii) if $c_{j(G)+1}$ is a face, then for $i = 1, 2$, we let e_i be the unique edge on the boundary of f_i that is incident with $c_{j(G)+1}$.

We denote the near-triangulation that T induces on f_i, rooted at v_i and e_i, by N_i (see Figure 5).

Lemma 8 ([32])**.** *The reflection* φ

(i) induces an isomorphism between N_1 *and* N_2 *and*

(ii) fixes precisely the central cells of G.

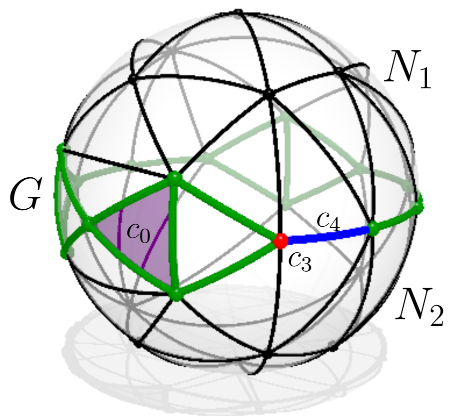

Figure 5: The girdle G (coloured vertices and edges) of a triangulation T rooted at a face c_0 (purple) with $j(G) = 3$. The near-triangulations N_1 and N_2 that T induces on the hemispheres of G, rooted at c_3 (red) and c_4 (blue), are isomorphic by Lemma 8(i).

Like for Lemma 4, we provide proofs of Lemmas 6 and 8 in the appendix for the sake of completeness. By Lemma 8, we can decompose T into its girdle G and two isomorphic near-triangulations N_1, N_2. What other properties do G, N_1, and N_2 have to satisfy? Clearly, each hemisphere of G is bounded by a cycle whose length matches the number of outer vertices of N_1 and N_2. We call this number the *length* of the girdle. The following lemma gives a characterisation of the near-triangulations that can occur.

Lemma 9. *Let* G *be a graph that occurs as the girdle of some triangulation and let* N *be a near-triangulation. There exists a triangulation* T *with a reflection* φ*,* G *as its girdle with respect to* φ*, and* N *as the near-triangulation from Lemma 8 if and only if*

(i) the number of outer vertices of N *is the same as the length of* G *and*

(ii) every chord of N *has at least one end vertex that is an outer vertex of* G.

Proof. First assume that the triangulation T exists. Property (i) is immediate. In order to prove (ii), let $e = uv$ be a chord of N. If u and v are central vertices of G, then Lemma 8(i) would imply that φ maps e to an edge $\varphi(e)$ with the same end vertices. Since e is not contained in G, Lemma 8(ii) shows that $\varphi(e) \neq e$, contradicting the fact that there are no double edges.

For the reverse implication, assume that N and G satisfy (i) and (ii). Let $\tilde{T}$ be any triangulation of which G is a girdle. Then $\tilde{T}$ defines vertices v_1, v_2 and edges e_1, e_2 on the boundaries of the hemispheres f_1 and f_2 of G, respectively. By (i) we can insert N into each hemisphere f_i at v_i and e_i. The result of this operation does not have any double edges by (ii); since all its faces are triangular, it is the desired triangulation T. □

3.2 Constructive decomposition

In this section, we formalise the decomposition scheme by showing how the graphs that can serve as girdles can be constructed and how triangulations with reflective symmetry arise from their girdle and the near-triangulations characterised by Lemma 9.

Constructing all possible girdles is rather easy. Once the length ℓ of the girdle and the number d of diamonds are fixed, all that is left is to consider all arrangements of d diamonds on a girdle of length ℓ. Note that $d \leqslant \frac{\ell}{2}$ is necessary; in the case of c_0 being a face, we furthermore have $d \geqslant 1$.

Let a girdle G be given. The near-triangulations that can be inserted into the hemispheres of G in order to give rise to a triangulation with reflective symmetry have to satisfy the conditions of Lemma 9. In particular, the distribution of chords is restricted. This is formalised in the following definition.

Definition 10. Let N be a near-triangulation and let D be a subset of its set of outer vertices. We call N *chordless outside* D if every chord of N has at least one end vertex in D.

More generally, let a cycle C with a root vertex v_C and a root edge e_C incident with v_C be given and let D_C be a set of vertices in C. Suppose that the length of C is the same as the number of outer vertices of N and let α be the unique isomorphism from C to the boundary C_N of the outer face of N that maps v_C to the root vertex v_N of N and e_C to the root edge e_N of N. We call N *chordless outside* D_C if it is chordless outside $\alpha(D_C)$.

Recall that $j(G)$ is the smallest index for which $c_{j(G)}$ is a vertex and let v_1, v_2, e_1, e_2 be given as in Definition 7. Denote by C_G the cycle in G bounding f_1 and let D_G be the set of outer vertices of G in C_G. With this notation, Lemmas 8 and 9 give rise to the following.

Theorem 11. *The triangulations T with a reflective symmetry in* $\mathrm{Aut}(c_0, T)$ *are precisely the ones that can be constructed by choosing*

- *a girdle G that contains c_0 as a central cell and*
- *a near-triangulation N that is chordless outside D_G*

and inserting a copy of N into each hemisphere f_1 (respectively f_2) of G at v_1 and e_1 (respectively at v_2 and e_2).

Remark 12. The construction in Theorem 11 is a one-to-two correspondence. To see this, suppose first that T has precisely one reflection φ. Then the two possible orientations of the girdle G_φ yield two different constructions. On the other hand, if T has more than one reflection, then it has both reflective and rotative symmetries. In order to prove that we also have a one-to-two correspondence in this case, we shall need some of the notation which we develop in Section 4. We defer the proof of this case to Remark 23.

4 Reflective and rotative symmetries

4.1 Building blocks

In this section, we assume that $\mathrm{Aut}(c_0, T)$ has a subgroup H that contains both reflections and rotations. By Theorem 3, H is isomorphic to D_n where $n \geqslant 2$ is a divisor of $\deg(c_0)$, i.e. there are n reflections and $n-1$ rotations (and the identity).

Since the rotations and the identity form a cyclic group, there is a unique cell $c_1 \neq c_0$ that is invariant under all rotations; we call c_0 the *north pole* and c_1 the *south pole* of T. For each reflection φ, there is a girdle G_φ by the results of Section 3 (see Definition 5 for the definition of a girdle and associated notation).

Clearly, no two girdles are the same by Lemma 1 and every girdle contains the north pole c_0 by definition. Thus, there are $2n$ cells incident with c_0 that are invariant under some reflection; denote them by $a_0, \dots, a_{2n-1}$, enumerated in the same order they lie around c_0 in counterclockwise direction. Then for every reflection, there is an $i \in \{0, \dots, n-1\}$ such that the invariant cells incident with c_0 are a_i and a_{n+i}; denote this automorphism by φ_i and its girdle by G_i.

Lemma 13. *The girdles $G_0, \dots, G_{n-1}$ have the following properties.*

(i) North and south pole are central cells of every girdle.

(ii) The two poles are the only cells that are central cells of more than one girdle.

Proof. The north pole c_0 is a central cell of every girdle by definition. Let G_i, G_j be two distinct girdles. We first show that there is a cell $c \neq c_0$ that is central in both of them and then prove that c is the south pole c_1. This will prove both (i) and (ii).

Since for each of φ_i, φ_j, the invariant cells incident with c_0 lie opposite, the central cells of G_i incident with the north pole lie in different hemispheres (or on the boundaries of different hemispheres) of G_j. Since the central cells of a girdle separate its hemispheres, G_i and G_j meet in at least one central cell apart from the north pole. Let c be such a cell.

Consider the automorphism $\varphi_i \circ \varphi_j$. Since c is invariant both under φ_i and under φ_j, it is also invariant under $\varphi_i \circ \varphi_j$. But the composition of two distinct reflections is always a rotation and thus, the only cells invariant under $\varphi_i \circ \varphi_j$ are the north and south pole, implying that c is the south pole. □

Since the cells $a_0, \dots, a_{2n-1}$ form a cyclic sequence around c_0, we will consider their indices modulo $2n$. In other words, with a slight abuse of notation, we shall write a_i instead of $a_{i \pmod{2n}}$. The same kind of notation will be used for the girdles $G_0, \dots, G_{n-1}$ (modulo n instead of modulo $2n$).

The rotations can be enumerated as $\rho_1, \dots, \rho_{n-1}$ so that every ρ_i satisfies

$$\rho_i(a_j) = a_{j+2i}$$

for all j. Thus, we have $\rho_1^i = \rho_i$ for all $i = 1, \dots, n-1$ (and $\rho_1^n = \mathrm{id}$).

Lemma 14. *For every $i = 1, \dots, n-1$, the following holds (see also Figure 6).*

(i) For every j, the rotation ρ_i induces an isomorphism between the girdles G_j and G_{j+2i}.

(ii) If n is odd, all girdles are isomorphic.

(iii) If n is even, $\rho_{\frac{n}{2}}$ induces a symmetry of each girdle and every two girdles G_i, G_j with $i-j$ even are isomorphic.

Proof. First observe that (i) immediately implies (ii) and (iii). On the other hand, for every $j = 0, \dots, n-1$, we have $\rho_i(a_j) = a_{j+2i}$ by definition and thus (i) follows directly from the fact that the central cells of a girdle are characterised by Lemma 4(ii). $\square$

Recall that Lemma 13 tells us that any two girdles cross precisely twice: once at each of the poles. However, while a central cell of a girdle cannot be a central cell of another girdle (unless it is one of the poles), it might well be an *outer* cell of another girdle.

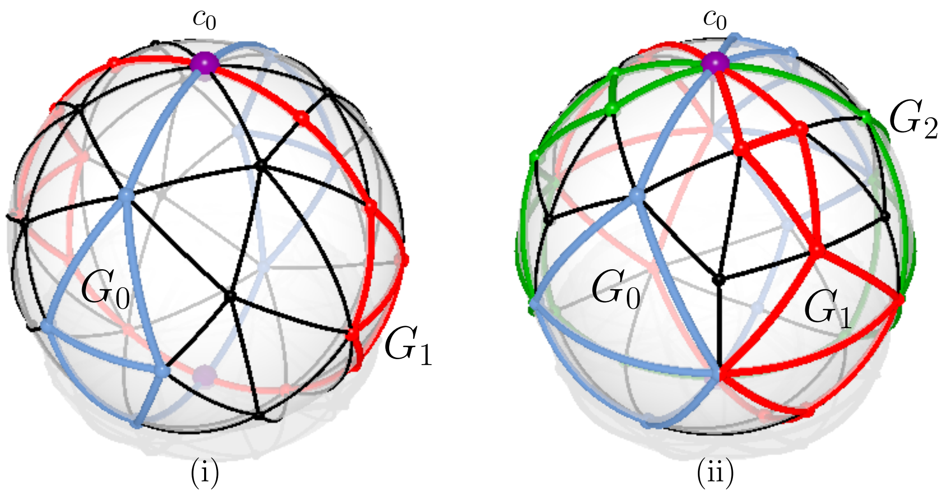

Figure 6: Two triangulations and their girdles with respect to a group $H \simeq D_n$ of automorphisms. In (i), we have $n = 2$ and the only rotation ρ_1 induces a symmetry of each girdle. In (ii), we have $n = 3$. The rotation ρ_1 maps G_0 to G_2, G_1 to G_0, and G_2 to G_1; the second rotation is $\rho_2 = \rho_1^{-1}$. All three girdles are isomorphic.

The poles divide every girdle G_i into two parts in a natural way: if $(x_j)_{j \in \mathbb{Z}_m}$ is the cyclic sequence from Lemma 4 with $x_0 = c_0$ (observe that by Lemmas 4 and 8(ii) this sequence is unique up to orientation), then $x_k = c_1$ for some k and we can consider the sequences $x_0, x_1, \dots, x_k$ and $x_k, x_{k+1}, \dots, x_{m-1}, x_0$. One of the sequences contains a_i, so we denote the union of its elements and their boundaries by M_i. The other sequence contains a_{n+i}, we denote the union of its elements and their boundaries by M_{n+i}. We call M_i and M_{n+i} *meridians*, the cells from the respective sequence of x_j's are the *central cells* of M_i and M_{n+i}, respectively. The other cells are *outer cells*, as before. Note that a central cell of G_i that lies on the boundary of one of the poles will be contained in both M_i and M_{n+i}. However, it will only be a central cell in one of them. Clearly, $G_i = M_i \cup M_{n+i}$ and thus $\bigcup_{i=0}^{n-1} G_i = \bigcup_{i=0}^{2n-1} M_i$.

Like the girdles, the meridians form a cyclic sequence; for simplicity, we will write M_i instead of $M_{i \pmod{2n}}$.

Definition 15 (Skeleton)**.** The union $S := \bigcup_{i=0}^{2n-1} M_i$ is called the *skeleton* of T with respect to the group $H \subseteq \mathrm{Aut}(c_0, T)$ (see Figure 7). For every $i = 0, \dots, 2n-1$, we say that the meridians M_i and M_{i+1} are *adjacent*. Every face of S that is not a central cell of at least one of the meridians is called a *segment* of S.

Observe that the skeleton of T is unique since all the girdles are.

Lemma 16. *The skeleton S has the following properties (see also Figure 7).*

(i) Every reflection φ_i, $0 \leqslant i \leqslant n-1$, induces an isomorphism between M_{i-j} and M_{i+j} for every $j = 1, \dots, n-1$.

(ii) Every rotation ρ_i, $1 \leqslant i \leqslant n-1$, induces an isomorphism between M_j and M_{j+2i} for every $j = 0, \dots, 2n-1$.

(iii) There is an isomorphism in H that maps M_i to M_j if and only if $i-j$ is even.

(iv) For every central cell c of a meridian M_i, $0 \leqslant i \leqslant 2n-1$, exactly one of the following holds.

- (C1) *c is a pole;*
- (C2) *c lies on the boundary of a pole;*
- (C3) *c is not contained in any other meridian;*
- (C4) *c is an outer cell of both meridians adjacent to M_i and not contained in any other meridian.*

(v) Every segment of S is bounded by a cycle that is contained in the union of two adjacent meridians.

(vi) There is a non-negative integer s such that for every pair (M_i, M_{i+1}) of adjacent meridians there are precisely s such segments.

Proof. Claims (i) and (ii) follow from Lemma 1 and the way φ_i and ρ_i act on $a_1, \dots, a_{2n}$. Claim (iii) is an immediate corollary of (i) or (ii).

To prove (iv), let c be a central cell of M_i. Note first that only one of the cases (C1)–(C4) can hold. Now assume that (C1)–(C3) do not hold, i.e., c is neither a pole nor lies on the boundary of a pole and there is at least one meridian M_j with $j \neq i$ that contains c. By Lemma 13(ii), c is an outer cell of every such meridian M_j.

The central cells of M_{i-1} and M_{i+1} separate the sphere into two parts, one of which contains the central cells of M_i (apart from the poles) while the other contains the central cells (apart from the poles) of all other meridians. This implies that c is an outer cell of at least one of M_{i-1}, M_{i+1} and not contained in any other meridian; it remains to show that c is an outer cell of both M_{i-1} and M_{i+1}. By (i), φ_i (or φ_{i-n} if $i > n$) induces an isomorphism between M_{i-1} and M_{i+1} and since c is invariant under φ_i, it is an outer cell of both meridians adjacent to M_i. This proves (iv).

For (v), observe first that every segment of S is bounded by a cycle, since S is 2-connected. To prove the other half of the statement, choose, for each $i = 0, \dots, 2n-1$, an arc A_i (an injective topological path) from c_0 to c_1 in the union of the central cells of M_i. By (iv), the A_i meet only in the poles and thus divide the sphere into $2n$ discs, each having a boundary that is contained in the union of two arcs M_i, M_{i+1}. Since every segment of S is contained in such a disc and no other meridian contains a point in this disc, (v) follows.

Finally, (vi) follows directly from (i) and (ii). □

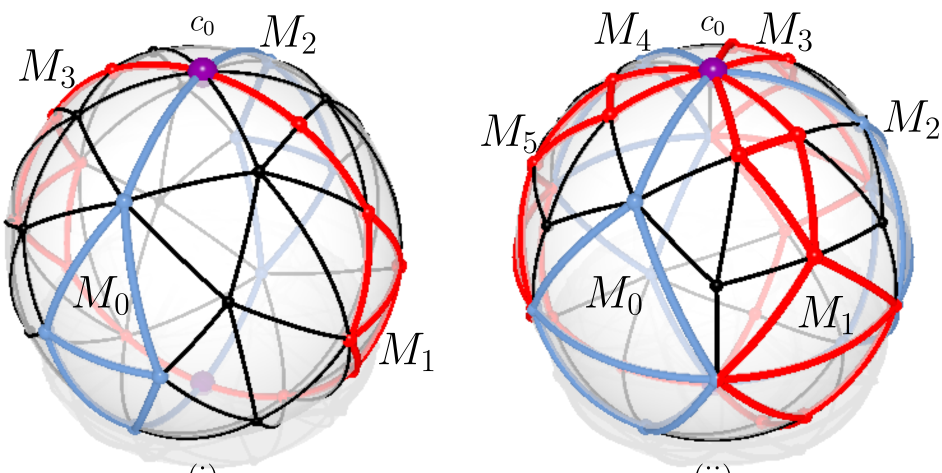

Figure 7: Two triangulations and their skeletons. Meridians with the same colour are isomorphic; each pair (M_i, M_{n+i}) of meridians forms a girdle. In (i), there is a unique segment between any two adjacent meridians, i.e. $s = 1$ in Lemma 16(vi). In (ii), we have $s = 2$.

By Lemma 16(vi), we can denote the segments of S whose boundaries are contained in the union of M_i and M_{i+1} by $f_1^i, \dots, f_s^i$. Note that the cycle from Lemma 16(v) bounding f_j^i is the union of a subpath of M_i and a subpath of M_{i+1}. These paths meet in their end vertices; denote by v_j^i their end vertex closer (in S) to the north pole and by w_j^i the one closer to the south pole. Without loss of generality, we assume that the enumeration of $f_1^i, \dots, f_s^i$ is chosen so that v_j^i is closer to the north pole than $v_{j'}^i$ whenever $j < j'$. Finally, let e_j^i be the edge on the boundary of f_j^i that is incident with v_j^i and

(i) contained in M_i if i is even, or

(ii) contained in M_{i+1} if i is odd.

With this notation and Lemma 16(i) and (ii), we obtain the following.

Lemma 17. *Let $j \in \{1, \dots, s\}$.*

(i) T induces a near-triangulation N_j^i on f_j^i (in terms of Definition 2), rooted at v_j^i and e_j^i for every i.

(ii) The near-triangulations $N_j^0, \dots, N_j^{2n-1}$ are isomorphic.

For a complete description of all possible skeletons, we need to characterise their structure at the poles and at other points where two adjacent meridians meet.

Lemma 18. *Let S be a skeleton and c be one of the poles of T. Then the structure of S at c is the following.*

(i) If c is a vertex, then either

(a) no two meridians meet in a cell incident with c or

(b) there is a number $k \geqslant 1$ such that every two adjacent meridians meet in their first k edges starting from c.

(ii) If $c = uv$ is an edge, then two non-adjacent meridians, say M_0 and M_2, have u respectively v as a central cell and the other two have its incident faces as central cells. No two meridians meet in an edge $e \neq c$ incident with u or v.

(iii) If c is a face f with vertices u, v, w on its boundary, then three mutually non-adjacent meridians, say M_0, M_2, M_4, have u, v, respectively w as a central cell and the other three have the edges incident with f as central cells. Either

(a) no two meridians meet in a cell incident with exactly one of u, v, w or

(b) there is a number $k \geqslant 1$ such that every two adjacent meridians meet in their first k edges starting from c.

Proof. Statements (i) and (iii) follow from Lemma 16(i) and (ii). The first claim in (ii) is immediate, since each of the four meridians contains a different cell incident with c as a central cell. For $i \in \{1, 3\}$, denote by f_i the face incident with c that is a central cell of M_i (see Figure 8(ii)). Since u, v are incident with f_1 and f_3, there are unique vertices w_1, w_3 different from u and v that are incident with f_1 and f_3, respectively. Note that w_1 and w_3 are distinct, since otherwise f_1 and f_3 would have the same set of incident vertices,

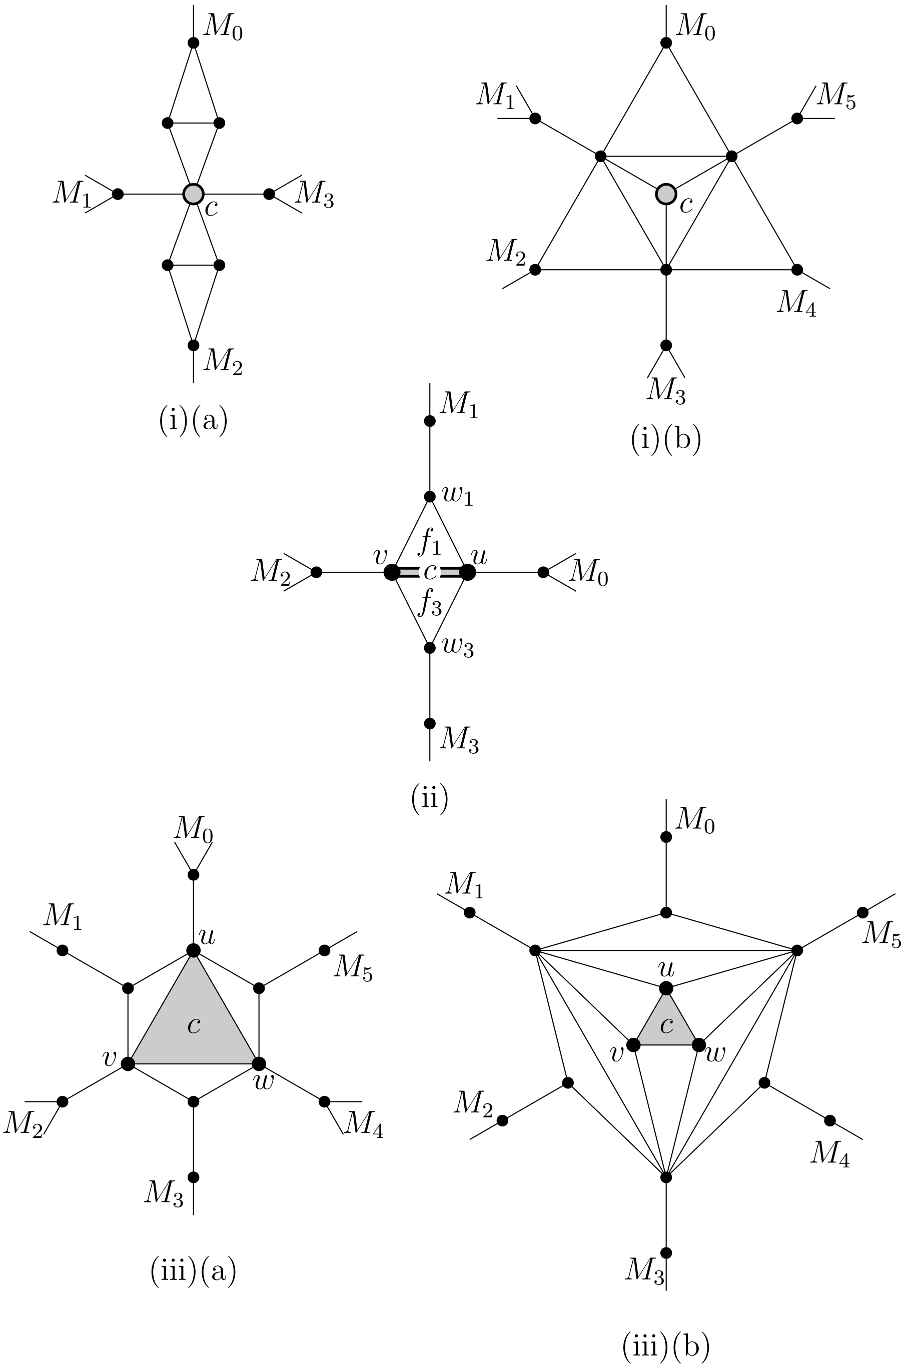

Figure 8: The possible structures of a skeleton at a pole as stated in Lemma 18, with $k = 1$ in the cases (i)(b) and (iii)(b).

which is not possible as our triangulations are simple and non-trivial. Suppose that M_0 meets M_1 in an edge $e \neq c$ incident with u; this has to be the edge uw_1. By applying φ_0, we see that M_0 meets M_3 in the edge uw_3.

As $w_1 \neq w_3$, both uw_1 and uw_3 are outer cells of M_0. Thus, there is a face that is central in M_0 and incident with u, w_1, and w_3. Hence w_1 and w_3 are connected by an edge e_0 that is central in M_0. Applying φ_1 shows that M_2 also has a central edge e_2 that connects w_1 and w_3. Since our triangulations are simple, the edges e_0 and e_2 are identical and T is a K_4. Since we assume all triangulations to be non-trivial, this is a contradiction. We have thus shown (ii). □

Lemma 18 describes the structure of the skeleton at the poles. The following lemma deals with the intersections of adjacent meridians between two segments.

Lemma 19. *Let $j \in \{1, \dots, s-1\}$ be fixed. Then there is a number $k_j \geqslant 0$, such that for every i, the intersection of M_i and M_{i+1} has a component that is a path of length k_j from w^i_j to v^i_{j+1}.*

Proof. This follows immediately from Lemma 16(i). □

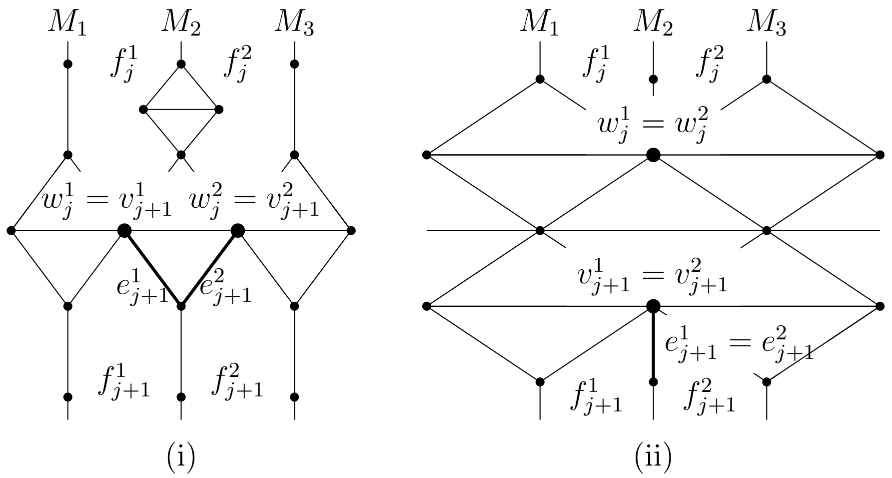

Figure 9: The structure of a skeleton between two segments as described in Lemma 19. In Case (i), we have $k_j = 0$, in Case (ii) $k_j = 2$.

All possible skeletons can thus be constructed by first choosing the numbers $n \geqslant 2$ and s and the dimensions of the poles. Note that a pole can only be an edge if $n = 2$ and it can only be a face if $n = 3$. Then choose the structure at the poles according to Lemma 18 and between the segments according to Lemma 19.

All triangulations with both reflective and rotative symmetry can be obtained by first taking a skeleton and then inserting the same near-triangulation in each type of segment according to Lemma 17. Similarly to the case of reflective symmetries, the near-triangulation inserted into a segment is only allowed to have chords that do not produce double edges by reflecting. In this case, this means that for every chord of the near-triangulation and each meridian bounding the corresponding segment, not both end vertices of the chord are contained in the meridian and central in it.

4.2 Constructive decomposition

In this section, we describe how to construct all possible skeletons and give a definition of the class of near-triangulations that can be inserted into the segments of the skeleton. With this information, we complete the decomposition scheme of triangulations with both reflective and rotative symmetries by showing how each such triangulation can be constructed from its skeleton S by inserting near-triangulations into the segments of S. Finally, we prove that this decomposition scheme is a one-to-two correspondence and thus in particular a constructive decomposition.

The maps that can serve as a skeleton of a triangulation can be constructed as follows. Suppose that the number n of reflections is given. Then we can choose

- the number s of isomorphism classes of segments of the skeleton;
- the structure of the skeleton at the poles according to Lemma 18;
- the numbers $k_1, \ldots, k_{s-1}$ from Lemma 19; and
- the distances of v_j^1 and w_j^1 on M_1 and on M_2 for every $j = 1, \ldots, s$ as well as the number and distribution of diamonds on these meridians between this two vertices. For arbitrary i, the structure of M_i at the boundaries of the segments is identical to that of M_1 or M_2, depending on the parity of i, by Lemma 16(iii).

The near-triangulations that can be inserted into a segment are similar to those that can be inserted into a hemisphere of a girdle: if such a near-triangulation had a chord both of whose end vertices are central cells of the same meridian, then applying the reflection that corresponds to that meridian shows that there is a double edge, a contradiction. In other words, a chord is only allowed if its end vertices are not in the same meridian or if they are in the same meridian, but at least one of them is an outer vertex of that meridian. This is formalised in the following definition.

Definition 20. Let N be a near-triangulation with root vertex v_N and root edge e_N. Suppose that a vertex $w_N \neq v_N$ is fixed, then the boundary of N is the union of two paths from v_N to w_N; denote the path that contains e_N by R and the other path by L. We call L and R the *sides* of the boundary. If vertex sets D_L and D_R on L and R are given, we say that N is 2-*sided chordless outside* D_L *and* D_R if every chord of N whose end vertices both lie on L or both lie on R has least one end vertex in D_L or in D_R, respectively.

More generally, let a cycle C with a root vertex v_C and a root edge e_C incident with v_C be given. If $w_C \neq v_C$ is given, let us define subpaths L_C and R_C as before. Suppose that D_{L_C} and D_{R_C} are sets of vertices on L_C and R_C, respectively. If there is an isomorphism α from C to the boundary of N that respects the rooting and maps w_C to w_N, then we say that N is 2-*sided chordless outside* D_{L_C} *and* D_{R_C} if it is 2-sided chordless outside $\alpha(D_{L_C})$ and $\alpha(D_{R_C})$.

Consider a segment f_j^0 of the skeleton; let C_j be its boundary. The two sides of C_j are its intersections R_j with M_0 and L_j with M_1, the set D_{R_j} (respectively D_{L_j}) is the set of outer vertices of M_0 on R_j (respectively of M_1 on L_j). The near-triangulations that can be inserted into f_j^0 are precisely those that are 2-sided chordless outside D_{L_j} and D_{R_j}. We thus have the following characterisation of triangulations with both reflective and rotative symmetries.

Theorem 21. *The triangulations T for which $\mathrm{Aut}(c_0, T)$ has a subgroup H isomorphic to D_n are precisely those that can be constructed by choosing*

- *a skeleton S_H and*
- *for every $j = 1, \ldots, s$, a near-triangulation N_j that is 2-sided chordless outside D_{L_j} and D_{R_j},*

and inserting a copy of N_j into f_j^i at v_j^i and e_j^i for every $j = 1, \ldots, s$ and $i = 0, \ldots, 2n-1$.

Remark 22. The construction from Theorem 21 is a one-to-two correspondence. To see this, let $S_{\max}$ be the skeleton of T with respect to the entire automorphism group $\mathrm{Aut}(c_0, T)$. Once we choose which meridian of $S_{\max}$ to take for M_1 in S_H, the skeleton S_H is completely defined. There are $|\mathrm{Aut}(c_0, T)|$ many ways to choose a meridian from $S_{\max}$. Two choices M, M' of a meridian yield the same construction (in terms of Theorem 21) if and only if there is an automorphism of T that maps M to M'. Lemma 16(iii) thus implies that there are precisely two different constructions.

Remark 23. An analogous argument to the one used in Remark 22 also shows that the construction from Theorem 11 is a one-to-two correspondence in the case of T having at least two reflective symmetries (and thus at least one rotative symmetry), which finishes the proof from Remark 12.

5 Rotative symmetries

5.1 Building blocks

In this section, suppose that $\mathrm{Aut}(c_0, T)$ contains a rotation φ. Note that the subgroup H of $\mathrm{Aut}(c_0, T)$ generated by φ contains no reflections by definition and hence is isomorphic to a cyclic group by Theorem 3. We fix the group H for the rest of this section; let m be its order. For every cell c incident with c_0, the cells $c, \varphi(c), \ldots, \varphi^{m-1}(c)$ are distinct, since $\varphi, \ldots, \varphi^{m-1}$ are rotations and thus have no invariant cells incident with c_0. Without loss of generality, we can choose φ in such a way that $c, \varphi(c), \ldots, \varphi^{m-1}(c)$ are arranged around c_0 in that order (in counterclockwise direction, say) for every cell c incident with c_0 (see Figure 10).

As in Section 4, there is a unique cell $c_1 \neq c_0$ that is invariant under φ. Again, we call c_0 the north pole and c_1 the south pole of T. Consider a shortest path P in T from c_0 to c_1. Denote its first and last vertex by v_0 and w_0, respectively, and write $v_i := \varphi^i(v_0)$ and $w_i := \varphi^i(w_0)$ for every $i \in \mathbb{N}$.

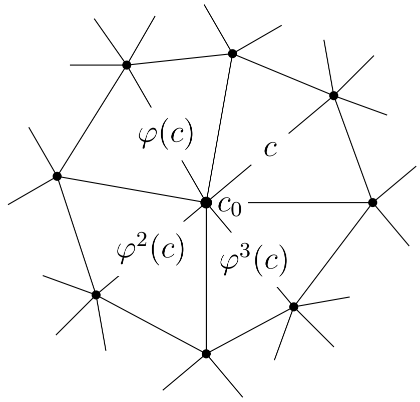

Figure 10: The images of a cell c incident with c_0 under a rotation φ of order 4.

Lemma 24. *The paths $P, \varphi(P), \ldots, \varphi^{m-1}(P)$ do not share any internal vertices. If c_0 is an edge or a face, the first vertices $v_0, \ldots, v_{m-1}$ of the paths are distinct. The same is true for c_1 and the last vertices $w_0, \ldots, w_{m-1}$.*

Proof. First note that the paths $P, \varphi(P), \ldots, \varphi^{m-1}(P)$ are distinct, because $\varphi^i(P) = \varphi^j(P)$ for $i \neq j$ would imply that $\varphi^i(c) = \varphi^j(c)$ for some cell c incident with c_0, which would contradict the fact that φ is a rotation. The same argument shows that two paths can only share an end vertex if it is c_0 or c_1.

Suppose two paths, without loss of generality P and $\varphi^i(P)$ with $i \neq 0$, share an internal vertex. Its distance from the first vertex has to be the same in both paths, since otherwise the union of the two paths would contain a path from c_0 to c_1 shorter than P, a contradiction to the choice of P. But then $\varphi^i(c) = c$ for some internal vertex c of P, meaning that c would be a third fixed cell of the rotation φ^i, a contradiction. □

Our proof of Lemma 24 is inspired by Tutte's proof [32, Statement 5.6] of the special case when m is prime. In said proof, the uniqueness of c_0, c_1 as cells invariant under rotation is not applied; instead, m being prime is used in order to prove that no two paths can share an internal vertex.

Lemma 24 implies that the paths $P, \varphi(P), \ldots, \varphi^{m-1}(P)$ together with c_0 and c_1 divide the triangulation into m parts. The union of these paths and cells *might* thus serve as a building frame in our constructive decomposition.

Definition 25. Let S be the union of the poles c_0, c_1, their boundaries, and paths $P, \varphi(P), \ldots, \varphi^{m-1}(P)$ satisfying the statement of Lemma 24 (see Figure 11). We call S a *spindle* of T with respect to the group $H \subseteq \mathrm{Aut}(c_0, T)$. A face of a spindle which is neither c_0 nor c_1 is called a *segment* of the spindle. For every $i \in \mathbb{N}$, denote by v_i and e_i the first vertex and edge of $\varphi^i(P)$, respectively.

Similarly to reflections, we immediately get the following result.

Lemma 26. *A spindle S has the following properties.*

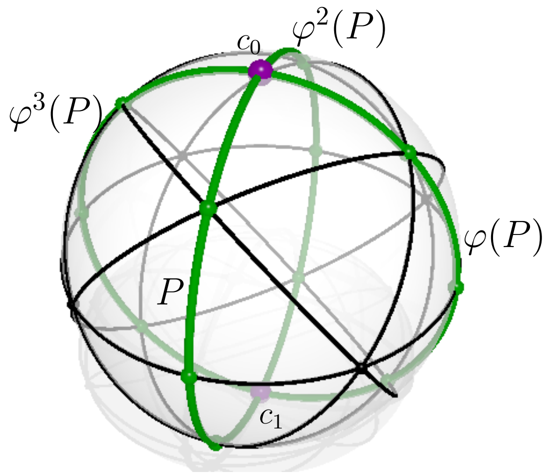

Figure 11: A triangulation and a spindle (coloured vertices and edges) with respect to the group $H = \{\mathrm{id}, \varphi, \varphi^2, \varphi^3\}$ of automorphisms, in which the poles c_0 and c_1 (purple) both are vertices.

(i) *S has exactly m segments $f_1, \dots, f_m$ with each f_i being bounded by a cycle containing $\varphi^{i-1}(P)$ and $\varphi^i(P)$;*

(ii) *for each $i \in \{1, \dots, m\}$, T induces a near-triangulation N_i on f_i, rooted at v_i and e_i; and*

(iii) *for every i, φ is an isomorphism from N_i to N_{i+1}.*

By Lemma 26, we can obtain all triangulations with rotative symmetry by first constructing all possible spindles and then inserting the same near-triangulation in each segment. However, unlike the girdle, a spindle is *not unique* since there might be different choices for the path P. Some triangulations, e.g. the one shown in Figure 12, have a unique spindle. In contrast, Figure 13 shows two different spindles of the same triangulation. Since the near-triangulation inserted in the segments in the first case is not isomorphic to the one used in the second case, there are at least two non-equivalent ways to construct this triangulation by taking a spindle and inserting the same near-triangulation in each segment. Thus, decomposing triangulations with rotative symmetry into their spindle and segments is *not* a constructive decomposition.

In order to find a constructive decomposition, let us refine the definition of a spindle so that it will be a unique substructure of T. To this end, we first aim to find a nested sequence $F_0, \dots, F_k$ of disjoint subgraphs of T that are invariant under φ. Each graph F_i will be *outerplanar*, i.e. it is embedded on the sphere such that all its vertices on the boundary of a common face, the *outer face*.

Definition 27. A graph is called a *cactus* if it is connected and every two cycles in it have at most one vertex in common. Every *block* of a cactus—a maximal subgraph that cannot be disconnected by deleting a single vertex—is a cycle or an edge. If a cactus

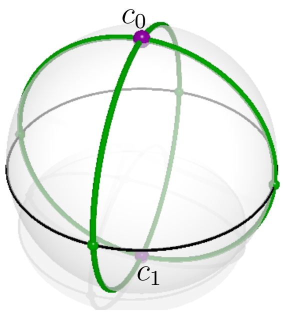

Figure 12: A triangulation with a unique spindle (green).

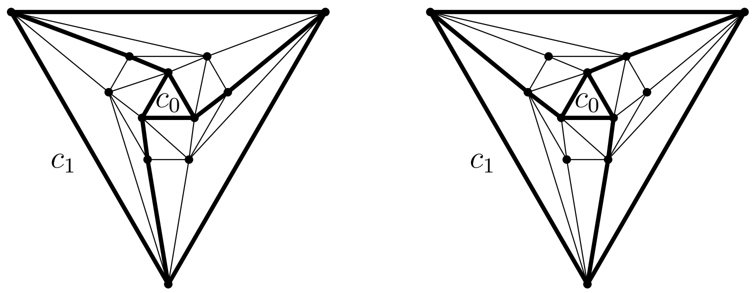

Figure 13: Two spindles (bold) of the same triangulation.

G has a root vertex, this induces a natural order on its set of blocks, similar to a tree order: consider the *block graph* of G—the graph whose vertices are the blocks of G and the vertices separating G and in which a block is adjacent to all separating vertices it contains (see Figure 14). This block graph is always a tree and if we choose its root to be

- the root of G if it is a separating vertex, or otherwise
- the unique block of G containing the root,

then this induces a tree order on the block graph and hence in particular a partial order on the set of blocks of G. In this order the blocks that contain the root are the minimal elements.

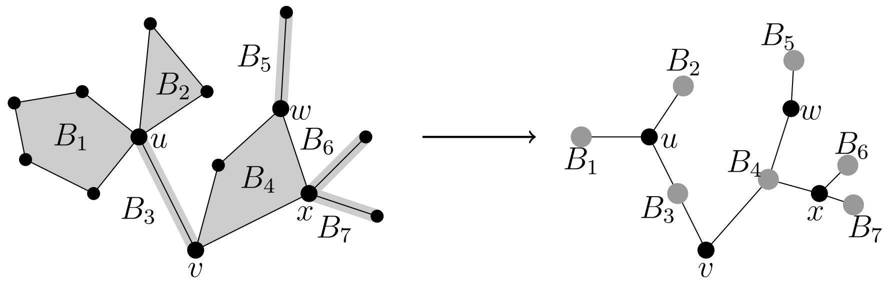

Figure 14: A cactus and its block graph.

Let G be an outerplanar subgraph of T for which the north pole c_0 lies in its outer face. We call G a *plane H-invariant cactus* if it is a cactus and invariant under all elements of the group $H \subseteq \mathrm{Aut}(c_0, T)$. Every cell of G that is invariant under φ contains a fixed-point by Brouwer's fixed-point theorem. Thus, only the outer face of G (because it contains c_0) and the cell that contains the south pole c_1 of T are invariant. If c_1 is a cell of G, then we call G *antarctic* (see Figure 15). If G is antarctic, we define the *centre* of G to be c_1 (if c_1 is a vertex), c_1 plus its end vertices (if c_1 is an edge), or the boundary of c_1 (if c_1 is a face).

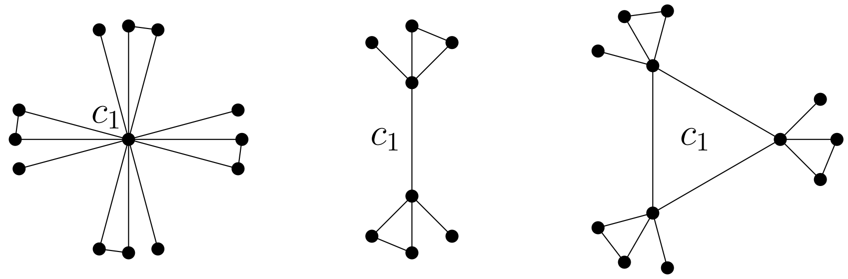

Figure 15: The three types of antarctic plane H-invariant cacti.

If G is not antarctic, we define the centre of G as follows. Choose a vertex v of G and an index $i \in \{1, \dots, m-1\}$ so that the distance d between v and $\varphi^i(v)$ is minimal. Let P be a shortest path from v to $\varphi^i(v)$ in G. The paths $P, \varphi(P), \dots, \varphi^{m-1}(P)$ do not share internal vertices; for otherwise there would be an internal vertex w of P and an index j such that the distance between w and $\varphi^j(w)$ is smaller than d, contradicting the choice of v and i. Thus, the union $P \cup \varphi(P) \cup \cdots \cup \varphi^{m-1}(P)$ is a cycle and invariant under φ. This cycle is called the *centre* of the non-antarctic H-invariant cactus G.

If G is any H-invariant cactus (antarctic or not), then the maximal connected subgraphs of G that share precisely one vertex with the centre are called *branches* of G (see Figure 16); the vertex of a branch B that lies in the centre of G is called the *base* of B. Observe that if G is antarctic and c_1 is a vertex, then the only branch of G is G itself.

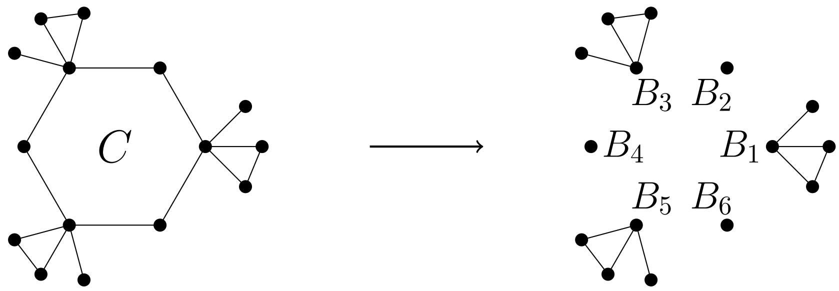

Figure 16: A plane H-invariant cactus for $|H| = 3$ with centre C and branches $B_1, \dots, B_6$.

If G is not antarctic and in addition the boundary of the south pole c_1 of T meets the boundary of the centre of G, then the south pole has to be a face or an edge and by symmetry all vertices on its boundary lie in the centre of G. In this case, we call G *pseudo-antarctic* (see Figure 17).

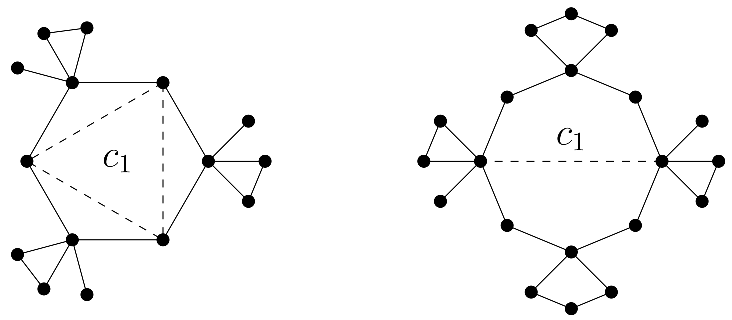

Figure 17: The two possibilities for a pseudo-antarctic plane H-invariant cactus.

Note that the above definition allows the case that the branches of a plane H-invariant cactus are just the vertices of its centre, in particular every invariant cycle is a plane H-invariant cactus. Furthermore, a plane H-invariant cactus is also a plane H'-invariant cactus for every subgroup H' of H.

The following lemma shows that plane H-invariant cacti appear in a natural way when we move from the north pole towards the south pole of the triangulation.

Lemma 28. *Let C be a cycle in T that is invariant under φ (and thus a plane H-invariant cactus). Suppose that C is neither antarctic nor pseudo-antarctic and let f be the face of C that contains the south pole. Denote by $\mathcal{F}$ the set of all faces of T that are contained in f and whose boundaries meet C. Let F be the subgraph of T consisting of all vertices and edges that lie on the boundary of a face $f' \in \mathcal{F}$ but do neither lie in C nor have an incident vertex in C. Then F has a unique component that is a plane H-invariant cactus.*

Proof. By construction, F is outerplanar and all its edges lie on the boundary of its outer face. Thus, no two of its cycles can meet in more than one vertex, showing that all components of F are cacti. The south pole c_1 is not contained in the outer face of F by construction, therefore there is a unique component F_1 of F such that either

- c_1 is contained in F_1 or
- c_1 is contained in a face of F_1 that is not its outer face.

In either case, F_1 is invariant under φ (and hence under all elements of H) and thus a plane H-invariant cactus. □

Repeated application of Lemma 28 gives rise to a finite sequence $F_0, \dots, F_k$ of plane H-invariant cacti in T as follows. We start by letting F_0 be the invariant cycle closest to c_0 like in Figure 18: if c_0 is a face, let F_0 be its boundary. If c_0 is an edge, let F_0 consist of all vertices and edges, apart from c_0 itself, that lie on the boundary of a face incident with c_0. Finally, if c_0 is a vertex, let F_0 consist of all vertices adjacent to c_0 and all edges that lie opposite to c_0 at some face incident with c_0. Note that in either case, F_0 is a cycle whose length is a multiple of $m = |H|$.

If F_0 is antarctic or pseudo-antarctic, then the sequence ends with $k = 0$; otherwise, by applying Lemma 28 with $C = F_0$, we obtain a plane H-invariant cactus F_1. If F_1 is

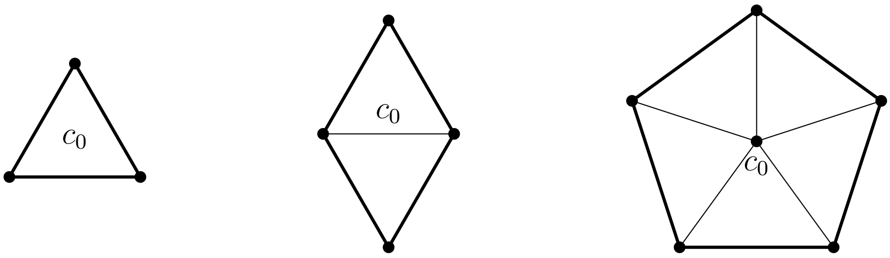

Figure 18: Finding an invariant cycle.

antarctic or pseudo-antarctic, we stop; otherwise, we apply Lemma 28 with C being the centre of F_1 to obtain another plane H-invariant cactus F_2. We continue this way until we obtain an antarctic or pseudo-antarctic plane H-invariant cactus F_k. We call the graphs $F_0, \ldots, F_k$ the *levels* of T (see Figure 19) and denote their centres by $C_0, \ldots, C_k$.

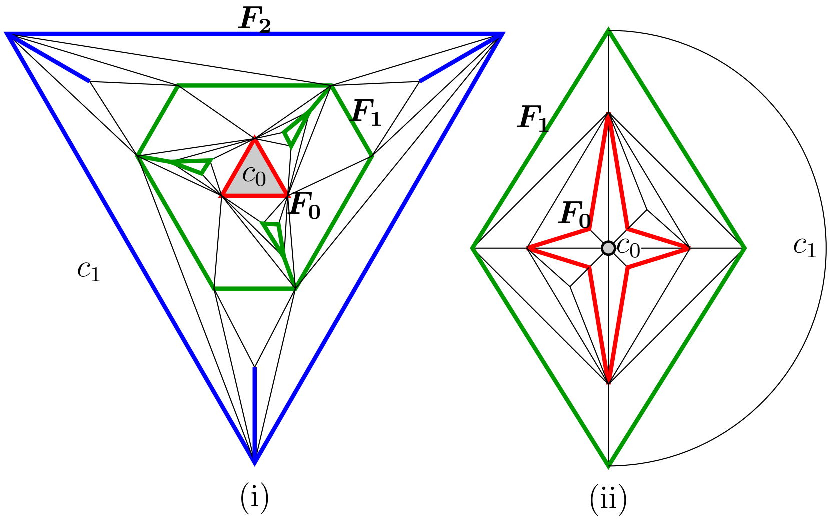

Figure 19: Two triangulations and their levels.

(i) A triangulation with levels F_0 (red), F_1 (green), and F_2 (blue), each of which a plane H-invariant cactus for $|H| = 3$. The last level F_2 is antarctic.

(ii) A triangulation with levels F_0 (red) and F_1 (green), both plane H-invariant cacti for $|H| = 2$. The last level F_1 is pseudo-antarctic.

The idea behind our refined version of a spindle is as follows. For a constructive decomposition, we shall need a *unique* substructure of T; something that the spindle was not able to provide, since the path P was chosen arbitrarily. Instead of connecting the north pole and the south pole by paths, we will base our construction on the levels of T and connect them by edges. Those edges have to be chosen in a unique way, which we will guarantee by always picking the 'leftmost' edge from a given vertex to the next level—a construction that will be made precise shortly. Moreover, it will not always be enough to have m edges from each level to the next. Indeed, if the north pole c_0 is a vertex, then its degree might be a multiple of m and there is no criterion which of the $\deg(c_0)$ edges we should choose. We thus have to start with all these edges.

The starting point of our construction will be vertices $u_0, \ldots, u_{am-1}$ on $F_0 = C_0$ (precise construction follows in Construction 29). We would then like to choose an edge from each u_j to the level F_1. However, not every vertex u_j necessarily has a neighbour in F_1. We will thus walk along the centre C_0 in clockwise direction from each u_j until we find a vertex v_j that has a neighbour in F_1. In order to decide which edge from v_j to F_1 we will pick, let e be one of the two edges of C_0 at v_j and let $e_j = v_j w_j$ be the first edge in clockwise direction around v_j, starting at e, with $w_j \in F_1$. Note that this definition does not depend on which edge of C_0 we choose as e. We call e_j the *leftmost edge* from v_j to F_1 and w_j the *leftmost neighbour* of v_j in F_1. We then continue the construction in F_1 by first going to the base of the branch that contains w_j, then walk along the cycle C_1 until we find a vertex that has a neighbour in F_2 and so on. We will now make this construction precise.

Construction 29 (liaison edges, sources, targets)**.** We begin our construction by choosing vertices $u_0, \ldots, u_{am-1}$ on $F_0 = C_0$ as follows (see also Figure 20): if c_0 is a vertex, let $a := \deg(c_0)/m$ and let $u_0, \ldots, u_{am-1}$ be all vertices of F_0, where the enumeration is in counterclockwise direction around the north pole. If c_0 is an edge, let $a := 1$ and let u_0 and u_1 be the vertices of F_0 that are not end vertices of c_0. Finally, if c_0 is a face, let $a := 1$ and let u_0, u_1, u_2 be the vertices on its boundary in counterclockwise order. Note that by the choice of $u_0, \ldots, u_{am-1}$, we have $\varphi(u_j) = u_{j+a \pmod{am}}$ for every j. With a slight abuse of notation, we will omit the modulo term in the index and simply write u_i instead of $u_{i \pmod{am}}$. We will use this notation also for all other cyclic sequences of vertices throughout this section.

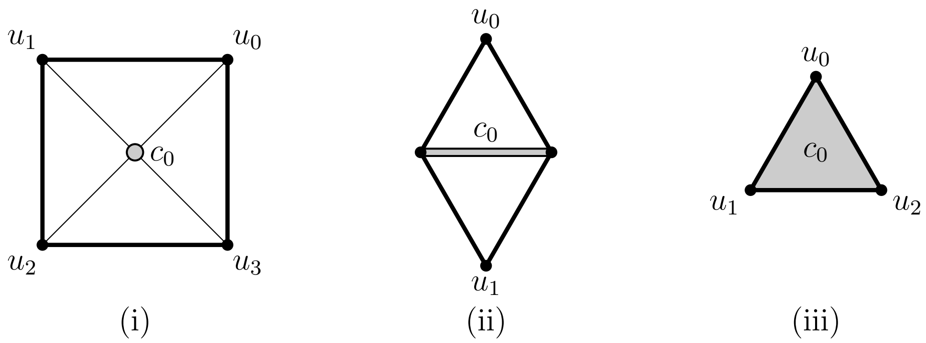

Figure 20: The vertices $u_0, \ldots, u_{am-1}$ for the north pole c_0 being (i) a vertex, (ii) an edge, (iii) a face. Note that in Case (i), we can either have $m = 4$, $a = 1$ or $m = a = 2$.

For each $j = 0, \ldots, am-1$, we define the vertices $u_j^0 := u_j, u_j^1, \ldots, u_j^k$, $v_j^0, \ldots, v_j^{k-1}$, and $w_j^1, \ldots, w_j^k$ as follows (see Figure 21): recursively for $0 \leqslant i \leqslant k-1$

(i) let v_j^i be the first vertex starting from u_j^i along the centre C_i in clockwise direction around the north pole that has an edge to F_{i+1};

(ii) let w_j^{i+1} be the leftmost neighbour of v_j^i in F_{i+1}; and

(iii) let u_j^{i+1} be the base of the branch of F_{i+1} that contains w_j^{i+1}.

The vertices $u_j^0, \dots, u_j^k$, $v_j^0, \dots, v_j^{k-1}$, and $w_j^1, \dots, w_j^k$ are uniquely defined by (i)-(iii). We have $\varphi(u_j^i) = u_{j+a}^i$, $\varphi(v_j^i) = v_{j+a}^i$, and $\varphi(w_j^i) = w_{j+a}^i$ for all i, j by the symmetry of T and the fact that $\varphi(u_j^0) = u_{j+a}^0$. Note that in (i), we encounter v_j^i before we reach u_{j-a}^i. Indeed, if the subpath of C_i from u_j^i to u_{j-a}^i contains no vertex that has a neighbour in F_{i+1}, then by the fact that $\varphi(u_{j-a}^i) = u_j^i$, no vertex of C_i has a neighbour in F_{i+1}, a contradiction to the definition of F_{i+1}.

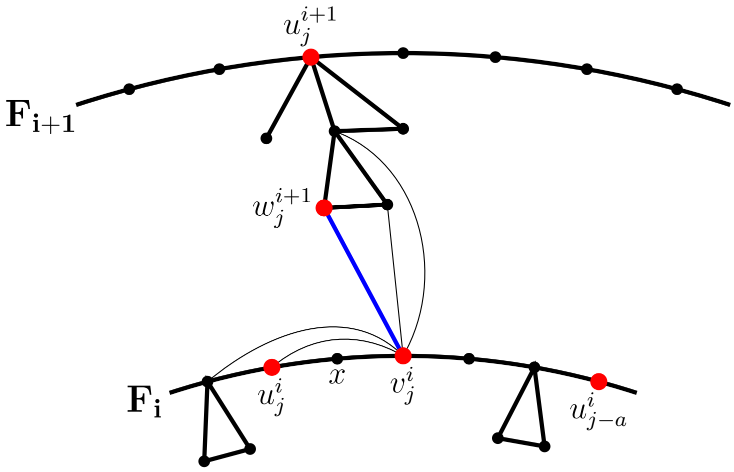

Figure 21: Constructing the sources v_j^i, targets w_j^{i+1}, and bases u_j^i. Note that if x is a base, say $x = u_{j'}^i$, then the construction yields $v_j^i = v_{j'}^i$.

The edges $v_j^i w_j^{i+1}$ are called *liaison edges*. For every liaison edge, we call v_j^i its *source* and w_j^{i+1} its *target*. Note that sources, targets, and bases do not have to be distinct. Clearly, two targets that lie in the same branch will always result in the same base, but also two bases will result in the same source if there is no eligible choice for a source between them on the cycle, and two sources may result in the same target if their leftmost edges lead to the same vertex. It is important to note that the sources $v_j^i, v_{j+a}^i, \dots, v_{j+a(m-1)}^i$ are always distinct, since they form an orbit under φ by the symmetry of the construction. The same holds for targets and bases up to the $(k-1)$-st level.

With the levels $F_0, \dots, F_k$ and the liaison edges $v_j^i w_j^{i+1}$, we are now able to define our refined spindles, called *fyke nets*.

Definition 30 (Fyke net)**.** Let $\tilde{F}$ be the union of

- the levels $F_0, \dots, F_k$ of T;

- all liaison edges $v_j^i w_j^{i+1}$;
- the north pole c_0 of T;
- all edges from c_0 to F_0 if c_0 is a vertex; and
- the south pole c_1 and its boundary if the last level F_k is pseudo-antarctic.

The *fyke net* of T with respect to the group $H \subseteq \mathrm{Aut}(c_0, T)$ is the maximal 2-connected subgraph F of $\tilde{F}$ that contains both poles c_0, c_1 (respectively their boundaries, if the poles are faces).

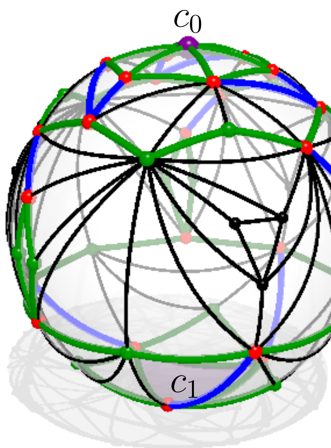

Figure 22: A triangulation and its fyke net (coloured vertices and edges) with respect to the group $H = \mathrm{Aut}(c_0, T)$, in which the north pole c_0 (purple) is a vertex and the south pole c_1 (purple) is a face. The liaison edges are drawn in blue; the sources, targets, and bases are drawn in red.

The intersection of the fyke net with the level F_i is called its *ith layer* and denoted by L_i. Note that every layer is a plane H-invariant cactus of order m by the symmetry of the construction. The fyke net has up to five different types of faces:

(i) faces at the north pole c_0 (either c_0 itself if it is a face, or all faces of T that are incident with c_0);

(ii) faces that are bounded by cycles in a branch of a layer; we call such faces *ears*;

(iii) faces bounded by two consecutive liaison edges and two subpaths of the two layers connecting their sources and their targets; we call such faces *segments*;

(iv) if the last layer L_k is pseudo-antarctic, m faces that are bounded by a subpath of the centre of L_k and the south pole c_1 if it is an edge, or one of its incident edges if it is a face; we call such faces *pseudo-antarctic*;

(v) the south pole c_1 if it is a face.

The following properties of the fyke net are easy to show, using the 2-connectedness of the fyke net and the structure of the levels of T.

Proposition 31. *Let F be the fyke net of a triangulation T. Then every segment of F is bounded by a cycle. An edge e of T lies in F if and only if*

(i) e is the north pole c_0 or an edge from c_0 to F_0,

(ii) e is the south pole c_1,

(iii) e is a liaison edge,

(iv) e lies in the centre of a level F_i, or

(v) e and a target w^i_j are contained in the same branch B of a level F_i and e lies on a path from w^i_j to u^i_j in B. Equivalently, the block $B(e)$ of B containing e and the smallest block (in the tree order on the block graph induced by choosing the base u^i_j of B as its root) $B(w^i_j)$ containing w^i_j satisfy $B(w^i_j) \geqslant B(e)$ (in said tree order).

Proof. The fyke net is 2-connected by definition and thus every face of F is bounded by a cycle. Denote by $\hat{F}$ the subgraph of $\tilde{F}$ consisting of all edges (and their end vertices) listed in (i)–(v). We first show that $\hat{F}$ is 2-connected and thus a subgraph of F. By (iv), $\hat{F}$ contains all centres of levels. The edges listed in (iii) and (v) all lie on paths between the centres of consecutive levels. For any $1 \leqslant i \leqslant k$, the targets $w^i_0, w^i_a, \dots, w^i_{(m-1)a}$ form an orbit under φ and therefore lie in distinct branches of F_i (unless $i = k$ and the centre of F_k is just the vertex c_1). Thus, there are at least m disjoint paths between the centres of F_{i-1} and F_i. Moreover, if c_0 is a vertex, then $\hat{F}$ contains $am \geqslant 2$ edges from c_0 to the centre of F_0. This proves that $\hat{F}$ is 2-connected.

In order to prove that $F = \hat{F}$, let e be an edge in $\tilde{F} \setminus \hat{F}$. Then e lies in a branch B of some level F_i. Let $B(e)$ be the block of B that contains e and let v be the unique vertex of $B(e)$ that separates $B(e)$ from the base of B. Then v also separates $B(e)$ from the centre of F_i and all targets w^i_j, because otherwise e would lie in $\hat{F}$ by (v). Thus, $e \notin F$, which shows $F = \hat{F}$, as desired. □

Note that the orbits of φ partition the set of faces of the fyke net F into sets of size m. In particular, the near-triangulations that T induces at two faces of F are isomorphic if these faces lie in the same orbit.

Unlike spindles, the fyke net is *unique* and thus, we can obtain all triangulations with rotative symmetry by first choosing a fyke net and then the near-triangulations that are to be inserted into the ears, segments, and (possibly) pseudo-antarctic faces.

As in the cases of reflective symmetries or both symmetries, we have to be more specific on which near-triangulations we are allowed to insert into segments.

Lemma 32. *Let f be a segment of the fyke net of T. Then there exist* unique *indices i, j satisfying the following properties (see Figure 23).*

(i) The boundary C of f consists of two liaison edges $v_j^i w_j^{i+1}$, $v_{j+1}^i w_{j+1}^{i+1}$ and paths P_i, P_{i+1}, where P_i is a path in the centre of the layer L_i from v_j^i to v_{j+1}^i and P_{i+1} is a path in the layer L_{i+1} from w_j^{i+1} to w_{j+1}^{i+1}, both in counterclockwise direction around c_0.

(ii) The path P_i has at least one edge.

(iii) The base u_j^i is a vertex on $P_i \setminus \{v_{j+1}^i\}$.

Proof. Property (i) is part of the definition of a segment and (ii) is immediate by the definition of the sources and targets. Property (iii) is clear by the way the source v_j^i has been chosen. □

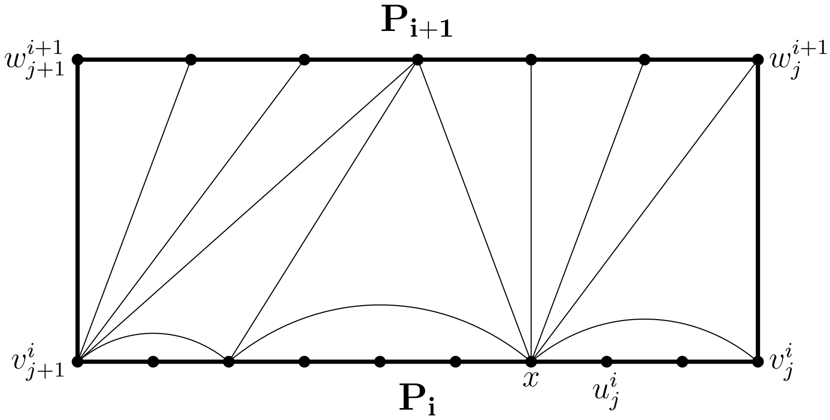

Figure 23: The structure of the near-triangulation in a segment of the fyke net.

Definition 33. Given a segment f, let i, j be the indices from Lemma 32. We write $f_j^i = f$ (observe that no pair (i, j) can occur for more than one segment, because no two segments have the same boundary) and denote by N_j^i the near-triangulation that T induces on f_j^i (see Definition 2), rooted at the vertex v_j^i and the edge $v_j^i w_j^{i+1}$.

Lemma 34. *The near-triangulations N_j^i have the following properties.*

(i) The edge $v_j^i w_j^{i+1}$ is part of the boundary of a non-root face of N_j^i whose third vertex x lies in the subpath of P_i from v_{j+1}^i to the predecessor of u_j^i.

(ii) Every edge of P_{i+1} is part of the boundary of a non-root face of N_j^i whose third vertex is in P_i.

(iii) If $m = 2$ and $\varphi(v_j^i) = v_{j+1}^i$, then there is no edge in N_j^i from v_j^i to v_{j+1}^i.

Proof. Property (i) follows from the existence of a non-root face having the edge $v_j^i w_j^{i+1}$ on its boundary and the fact that its third vertex x cannot be

- a vertex in P_{i+1}, since this would contradict the choice of w_j^{i+1} as the leftmost neighbour of v_j^i;
- an internal vertex of N_j^i, since then x would have been in the $(i+1)$-st level of T, again contradicting the choice of w_j^{i+1};
- a vertex on the subpath of P_i from u_j^i to v_j^i, since by the choice of v_j^i no vertex on this path has a neighbour in the $(i+1)$-st level of T.

In order to prove (ii), let e be an edge of P_{i+1}. It is part of the boundary of a unique non-root face of N_j^i and by the definition of F_{i+1} it is also part of the boundary of a face of T whose third vertex is in F_i. We will show that this latter face is also a face of N_j^i, thus showing (ii).

We will prove this for the edges in P_{i+1} one by one, starting from the edge at w_{j+1}^{i+1}. Let x_1 be the last neighbour of w_{j+1}^{i+1} on P_i (starting from v_{j+1}^i). Together with the root face of N_j^i, the edge $x_1 w_{j+1}^{i+1}$ divides N_j^i into two parts, let N_1 be the part which contains all of P_{i+1}. The edge $x_1 w_{j+1}^{i+1}$ is part of the boundary of a unique non-root face of N_1, denote the third vertex of this face by y_1 (see Figure 24). If y_1 is the neighbour of w_{j+1}^{i+1} on P_{i+1}, then we have found the desired face. Otherwise, it cannot be a vertex of P_{i+1} since the edge $w_{j+1}^{i+1} y_1$ is in F_{i+1} and would thus also have been in L_{i+1}. Since x_1 was the last neighbour of w_{j+1}^{i+1} on P_i, y_1 has to be an internal vertex of N_j^i. Now repeat the construction with y_1 instead of w_{j+1}^{i+1} to obtain a vertex x_2 on P_i (possibly $x_2 = x_1$), a near-triangulation $N_2 \subseteq N_1$ and a vertex y_2. As before, y_2 cannot lie on P_i by the definition of x_2 and not in $P_{i+1} \setminus \{w_{j+1}^{i+1}\}$ by the definition of L_{i+1}. It also cannot be w_{j+1}^{i+1}, since then the edge $x_2 w_{j+1}^{i+1}$ would either contradict the choice of x_1 as the last neighbour of w_{j+1}^{i+1} on P_i (if $x_1 \neq x_2$) or it would yield a double edge (if $x_1 = x_2$), also a contradiction. We can thus continue the construction and will always obtain internal vertices $y_1, y_2, \dots$ of N_j^i. Since these vertices are distinct and N_j^i is finite, this is a contradiction, implying that y_1 must have been the neighbour of w_{j+1}^{i+1} on P_{i+1}.

The same construction for every later edge of P_{i+1} proves (ii). Finally, note that (iii) is immediate since otherwise there would be a double edge in T between v_j^i and v_{j+1}^i. □

Observe that 34(ii) immediately implies that no two vertices in P_{i+1} are connected by a chord in N_j^i.

The near-triangulations inserted into ears or pseudo-antarctic faces, however, do not have any restrictions. Indeed, chords do neither contradict the construction of the layers by Lemma 28 nor can they result in double edges.

5.2 Constructive decomposition

By the results of Section 5.1, a triangulation with rotative symmetry can thus be constructed by first choosing a fyke net, then choosing, for every isomorphism class of ears or pseudo-antarctic faces, any near-triangulation to be inserted into each of these ears,

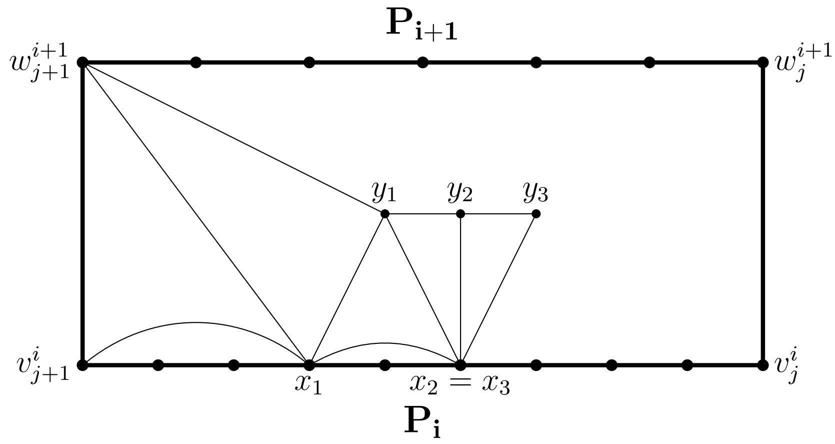

Figure 24: The construction proving Lemma 34(ii). Note that the vertices $x_1, x_2, \dots$ are not necessarily distinct. The vertices $y_1, y_2, \dots$, however, are mutually distinct.

and finally, for every isomorphism class of segments, choosing a near-triangulation that satisfies Lemma 34.

The construction of a map F that can serve as a fyke net requires several steps. Suppose that the desired order m of the automorphism group is already given, as well as the dimensions of the poles and the number $k+1$ of layers $L_0, \dots, L_k$. We construct the fyke net F in the following steps.

- Choose the layer L_0 to be a cycle depending on the dimension of the north pole c_0 like in Figure 18. If c_0 is a vertex, let am be the length of L_0, otherwise we set $a = 1$.

- For $i = 1, \dots, k-1$, let C_i be a cycle whose length is a multiple of m. These cycles will serve as the centres of the layers. The choice of C_k depends on the dimension of the south pole c_1: if c_1 is a vertex, then C_k only consists of c_1 and the layer L_k will be antarctic. If c_1 is an edge, C_k can either be a cycle of even length, in which case L_k will be pseudo-antarctic, or the edge c_1 itself, in which case L_k will be antarctic. Finally, if c_1 is a face, then C_k has to be a cycle whose length is divisible by 3. In that case, L_k will be antarctic if L_k is a triangle and pseudo-antarctic otherwise.

- Choose the bases $u_0^0, \dots, u_{am-1}^0$ in L_0 in a counterclockwise order like in Definition 30: if c_0 is an edge, then choose two opposite vertices as u_0^0, u_1^0 (the other two will then be the end vertices of c_0), otherwise choose all vertices of L_0.

- Choose the sources $v_0^0, \dots, v_{am-1}^0$ as follows: For $j = 0, \dots, a-1$, choose v_j^0 to be a vertex on the path starting at u_j^0 and running along C_0 in clockwise direction around the north pole to the predecessor of u_{j-a}^0 so that $v_1^0, \dots, v_a^0$ appear in counterclockwise order on C_0, but are not necessarily distinct. The remaining sources $v_a^0, \dots, v_{am-1}^0$ are obtained by recursively applying the rotation φ. The set of sources in C_0 is denoted by S_0.

- Recursively for $i = 1, \dots, k$, choose the bases $u_0^i, \dots, u_{am-1}^i$ and sources $v_0^i, \dots, v_{am-1}^i$ on C_i as follows: if C_i has length $a_i m$, pick a subpath consisting of a_i vertices and choose $u_0^i, \dots, u_{a-1}^i$ from this subpath so that they appear in counterclockwise order on C_i. Again, the bases do not have to be distinct. Furthermore, if the sources v_l^{i-1} and v_{l+1}^{i-1} were identical, the corresponding bases u_l^i and u_{l+1}^i should also be identical. The other bases follow again by applying symmetry. After choosing the bases, we can pick the sources like in the previous step (but note that we do not need any sources for $i = k$). We denote the sets of bases and sources in C_i by B_i and S_i, respectively.

- Having fixed the bases and sources for every i, we can now extend the C_i to layers L_i. To that end, we need to add a suitable plane cactus at every base in C_i. For every i, denote by $\pi_i \colon S_{i-1} \to B_i$ the function that maps v to u if there is a j with $v = v_j^{i-1}$ and $u = u_j^i$. For every base u in C_i, we now attach a plane cactus at u that has at most as many blocks that are maximal in its natural order as u has preimages under π_i. Again, it is sufficient to choose the cacti for the first a bases, the others are isomorphic by symmetry.

- Finally, we choose the targets and the liaison edges. For every $u \in B_i$ and every $v_j^{i-1} \in (\pi_i)^{-1}(u)$, we choose a vertex in the cactus at u to be the target w_j^i for the liaison edge $v_j^{i-1} w_j^i$ according to the following rules:
  - The targets are arranged in counterclockwise order for increasing index j and
  - every block that is maximal in the natural order of its branch has at least one vertex that does not belong to any other block and is chosen as a target.

The near-triangulations induced at a segment of the fyke net are characterised by Lemma 34. The following definition formalises this characterisation.

Definition 35. Let C be a cycle that consists of a path P_v from v_1 to v_2, a path P_w from w_1 to w_2, and two edges v_1w_1, v_2w_2. Denote the length of C by ℓ and let u be a vertex on P_v. We choose v_1 as the root vertex of C and v_1w_1 as the root edge. If N is a near-triangulation with ℓ outer edges, root vertex v_N, and root edge e_N, let us denote by α the isomorphism from C to the boundary of the outer face of N that respects the rooting. We say that N is 2*-layered with respect to* v_1*,* v_2*,* w_1*,* w_2*, and* u if it satisfies Lemma 34 with $v_j^i := \alpha(v_1)$, $v_{j+1}^i := \alpha(v_2)$, $w_j^{i+1} := \alpha(w_1)$, $w_{j+1}^{i+1} := \alpha(w_2)$, and $u_j^i := \alpha(u)$.

If T is a triangulation with rotative symmetry and f is a segment of its fyke net F, then by Lemma 32 and Definition 33 we have $f = f_j^i$ for some i, j and N_j^i is 2-layered with respect to v_j^i, v_{j+1}^i, w_j^{i+1}, w_{j+1}^{i+1}, and u_j^i. Here the boundary of f_j^i is rooted at v_j^i and $v_j^i w_j^{i+1}$. For every ear f of the fyke net we choose the root vertex v_f of its boundary to be its vertex closest to the base of the branch it is contained in. As root edge e_f we choose the left of the two edges at v_f on this boundary. Finally, if F has pseudo-antarctic faces, we choose for each such face f the root edge e_f to be either the south pole c_1 (if it is an edge) or the unique edge incident to both f and c_1 (if c_1 is a face). As the root vertex v_f we choose the end vertex of e_f that lies in clockwise direction from e_f around f.

Theorem 36. *The triangulations T with rotative symmetries in* $\mathrm{Aut}(c_0, T)$ *are precisely the ones that can be constructed by choosing*

- *a fyke net F;*
- *for every isomorphism class $[f]$ (under rotation) of ears of F a near-triangulation whose boundary is isomorphic to the boundaries of those ears with respect to the rooting;*
- *for every isomorphism class $[f]$ of pseudo-antarctic faces of F a near-triangulation whose boundary is isomorphic to the boundaries of those faces with respect to the rooting; and*
- *for every isomorphism class $\{f_j^i, f_{j+a}^i, \ldots, f_{j+(m-1)a}^i\}$ of segments of F a 2-layered near-triangulation with respect to v_j^i, v_{j+1}^i, w_j^{i+1}, w_{j+1}^{i+1}, and u_j^i*

and inserting a copy of each near-triangulation into the corresponding faces of F at their root vertices and edges. This construction is a one-to-one correspondence.

6 Discussion

The constructive decomposition presented in Sections 3 to 5 decomposes triangulations into 'building frames' and 'flagstones'. The building frames are girdles (see Definition 5) of triangulations with reflective symmetry, fyke nets (Definition 30) of triangulations with rotative symmetry, and skeletons (Definition 15) in the case of both types of symmetries. Into the faces of the building frames, we insert near-triangulations, or 'flagstones', in order to construct a triangulation with the desired symmetries.

This constructive decomposition is the key to enumerate triangulations with specific symmetries and to sample them uniformly at random, e.g. based on a recursive method [7, 14] or on Boltzmann sampler [6, 12]. To this end, it will be necessary to translate the decomposition into functional equations for the cycle index sums [27, 34, 35, 36] that enumerate these triangulations and the basic structures (i.e. building frames and flagstones) arising in their decomposition. To this end, the cycle index sums will have to distinguish between different types of vertices, e.g. between central vertices and outer vertices (see Definition 5) of a girdle G, because these vertices determine which chords are allowed in the near-triangulation that is to be inserted into the hemispheres of G. The structure of these near-triangulations has to be encoded in a functional equation for their cycle index sum.

Deriving functional equations for cycle index sums that enumerate triangulations with specific symmetries will path the way to an enumeration of unlabelled planar cubic graphs. Chapuy, Fusy, Kang, and Shoilekova [11] developed a general decomposition strategy for a wide variety of graph classes that can in particular be used to relate unlabelled cubic planar graphs to unlabelled 3-connected cubic planar graphs in terms of functional equations for the respective cycle index sums. The base classes of this decomposition include unlabelled 3-connected cubic planar graphs with a distinguished vertex or edge.

By Whitney's theorem [37], we can relate these graphs to cubic *maps* rooted at a single cell, which in turn are the duals of triangulations rooted at a single cell. As Whitney's theorem provides uniqueness of the embedding only up to orientation of the sphere, it follows that this relation between graphs and triangulations can be a one-to-two correspondence or a one-to-one correspondence, depending on the symmetries of the respective triangulation. Thus, encoding the symmetries of triangulations rooted at a single cell in their cycle index sums provides the fundament for enumerating unlabelled cubic planar graphs.

Acknowledgement

The authors thank the anonymous referee for many helpful remarks to improve the presentation of this paper.

References

[1] M. Albenque and J.-F. Marckert. Some families of increasing planar maps. *Electron. J. Probab.*, 13:1624–1671, 2008.

[2] G. Amendola. Decomposition and enumeration of triangulated surfaces *Exper. Math.*, 17:153–166, 2008.

[3] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Commun. Math. Phys.*, 241:191–213, 2003.

[4] B. Benedetti and F. Lutz. Random discrete Morse theory and a new library of triangulations *Exp. Math.*, 23:66–94, 2014.

[5] O. Bernardi and N. Bonichon. Intervals in Catalan lattices and realizers of triangulations. *J. Combin. Theory Ser. A*, 116(1):55–75, 2009.

[6] M. Bodirsky, É. Fusy, M. Kang, and S. Vigerske. Boltzmann samplers, Pólya theory and cycle pointing. *SIAM Journal on Computing*, 40:721–769, 2011.

[7] M. Bodirsky, C. Gröpl, and M. Kang. Generating labeled planar graphs uniformly at random. *Theoret. Comput. Sci.*, 379(3):377–386, 2007.

[8] J. Bouttier, P. Di Francesco, and E. Guitter. Counting colored random triangulations. *Nuclear Phys. B*, 641:519–532, 2002.

[9] G. Brinkmann and B. D. McKay. *plantri: A program for generating planar triangulations and planar cubic graphs.* http://cs.anu.edu.au/people/bdm/plantri/

[10] H. S. M. Coxeter and W. O. J. Moser. *Generators and Relations for Discrete Groups.* 4th edition, Springer Berlin Heidelberg, 1980.

[11] G. Chapuy, É. Fusy, M. Kang, and B. Shoilekova. A complete grammar for decomposing a family of graphs into 3-connected components. *Electron. J. Combin.*, 15:#R148, 2008.

[12] P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer. Boltzmann samplers for the random generation of combinatorial structures. *Combin. Probab. Comput.*, 13:577–625, 2004.

[13] S. Felsner and F. Zickfeld. On the Number of Planar Orientations with Prescribed Degrees. *Electron. J. Combin.*, 15:#R77, 2008.

[14] P. Flajolet, P. Zimmerman, and B. van Cutsem. A calculus for the random generation of labelled combinatorial structures. *Theoret. Comput. Sci.*, 132:1–35, 1994.

[15] P. Di Francesco. 2D quantum gravity, matrix models and graph combinatorics. *Applications of random matrices in physics*, in *NATO Sci. Ser. II Math. Phys. Chem.*, 221:33–88, 2006.

[16] É. Fusy. Counting unrooted maps using tree-decomposition. *Sém. Lothar. Combin.*, 54A:B54Al, 2007.

[17] I. P. Goulden and D. M. Jackson. *Combinatorial Enumeration*. Dover Publications, 2004.

[18] V. A. Liskovets. Enumerative Formulae for Unrooted Planar Maps: a Pattern. *Electron. J. Combin.*, 11(1):#R88, 2004.

[19] J. De Loera, J. Rambau, and F. Santos. *Triangulations*. Structures for Algorithms and Applications Series, Springer-Verlag, Berlin, 2010.

[20] P. Mani. Automorphismen von polyedrischen Graphen. *Math. Ann.*, 192:279–303, 1971.

[21] B. D. McKay. Isomorph-free exhaustive generation. *J. Algorithms*, 26:306–324, 1998.

[22] G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013.

[23] J. W. Moon and L. Moser. Triangular dissections of n-gons. *Canad. Math. Bull.*, 6:175–178, 1963.

[24] D. Osthus, H. J. Prömel, and A. Taraz. On random planar graphs, the number of planar graphs and their triangulations. *J. Combin. Theory Ser. B*, 88:119–134, 2003.

[25] N. Pippenger and K. Schleich. Topological characteristics of random triangulated surfaces. *Random Struct. Alg.*, 28(3):247–288, 2006.

[26] B. Richmond and N. Wormald. Almost all maps are asymmetric. *J. Combin. Theory Ser. B*, 63:1–7, 1995.

[27] R. W. Robinson. Enumeration of non-separable graphs. *J. Combin. Theory*, 9:327–356, 1970.

[28] D. Sleator, R. Tarjan, and W. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–681, 1988.

[29] T. Sulanke and F. Lutz. Isomorphism-free lexicographic enumeration of triangulated surfaces and 3-manifolds. *European J. Combin.*, 30:1965–1979, 2009.

[30] W. T. Tutte. A census of planar triangulations. *Canad. J. Math.*, 14:21–38, 1962.

[31] W. T. Tutte. A census of planar maps. *Canad. J. Math.*, 15:249–271, 1963.

[32] W. T. Tutte. On the enumeration of convex polyhedra. *J. Combin. Theory Ser. B*, 28:105–126, 1980.

[33] W. T. Tutte. *Graph Theory.* Addison-Wesley, 1984.

[34] T. R. S. Walsh. Counting non-isomorphic three-connected planar maps. *J. Combin. Theory Ser. B*, 32:33–44, 1982.

[35] T. R. S. Walsh. Counting unlabelled three-connected and homeomorphically irreducible two-connected graphs. *J. Combin. Theory Ser. B*, 32:12–32, 1982.

[36] T. R. S. Walsh and V. A. Liskovets. The enumeration of non-isomorphic 2-connected planar maps. *Canad. J. Math.*, 35:417–435, 1983.

[37] H. Whitney. Congruent graphs and the connectivity of graphs. *Amer. J. Math.*, 54:150–168, 1932.

A Additional proofs

Proof of Lemma 4

Let I be the set of cells that are invariant under φ. Define an auxiliary graph F with vertex set I by joining two elements of I by an edge whenever they are incident. Note that φ, although chosen as an element of $\mathrm{Aut}(c_0, T)$, is also an element of $\mathrm{Aut}(c, T)$ for every $c \in I$. Since φ is a reflection, every vertex in F has degree 2 and thus, every component is a cycle. Let C be the cycle of F that contains c_0. The vertices of C—arranged in the order they appear on C—form the desired cyclic sequence. Indeed, (i) and (iii) hold by the definition of F and the fact that C is a component of F, while (ii) is an immediate consequence of the fact that φ is a reflection.

Proof of Lemma 6

By Lemma 4, two central cells of G are incident if and only if they are consecutive in the cyclic sequence. We claim that every outer cell is contained in a *unique* diamond, which implies that the subspace of the sphere consisting of G and the faces in its diamonds is contractible to a Jordan curve, which in turn implies Lemma 6 by the Jordan curve theorem. Indeed, every outer cell of G is a vertex or an edge that is contained in a diamond. If two diamonds share an outer edge, they also share an outer vertex v. Now φ maps v to the other outer vertex of each of the two diamonds, hence they also share their second outer vertex. But then the central edges of the two diamonds are distinct and have the same end vertices, contradicting the fact that T has no double edges.

Proof of Lemma 8

First observe that (ii) follows directly from (i) and the fact that all central cells of G are invariant under φ by Lemma 4 and Definition 5. In order to prove (i), let c be any cell of T that is incident with c_0, but is *not* a central cell of G. Then c is contained in one of the hemispheres of G or lies on the boundary of precisely one hemisphere. The reflection φ maps c to a cell that is contained in (or lies on the boundary of) the other hemisphere of G, which yields (i).