Tetrahedral triangulations of the sphere

Jeroen Huijben
Supervisor: Frits Beukers

Bachelor thesis Mathematics

Utrecht University

June 2016

Preface

The idea for this thesis comes from an article by William Thurston [1] about triangulations of the sphere and more general shapes of polyhedra. In the first two sections of his article he briefly outlines some methods to construct certain triangulations of the sphere. However, this doesn't seem to be the main focus of his article, as he goes on to work on the more general cone manifolds.

However, his idea is a good start to find some enumerative results. The goal of this thesis is therefore to count a certain class of triangulations, namely *tetrahedral triangulations*, of which the smallest example is indeed the tetrahedron.

The same subject has been addressed in a previous thesis under supervision of Frits Beukers [2], but it resulted in very cumbersome formulas. In this thesis we try to do the counting more efficient to find nicer formulas.

I also chose another approach to the problem. Whereas the approach in [2] is based on polytopes and graphs, I will use the more topological notion of *Delta sets*. These structures are somewhat more general and I think they are more natural in this context. Working with Delta sets is a bit more abstract and a lot of work went into figuring out how the definitions work exactly. The main challenge was working with the triangle maps, because my definition is a bit broader than the definition which usually occurs in literature.

After this rather technical work with triangle maps, we will translate the problem into a problem on symmetries of the regular triangulation of the plane. Once we have achieved this, we can count the number of triangulations by using some number theory. But we start now with some definitions.

Acknowledgements

Foremost I would like to thank Frits Beukers for proposing this enjoyable subject and guiding me during my thesis with useful suggestions. Further I want to thank my friends for various discussions on the subject.

Contents

Chapter 1

Definitions

Triangles are a special case of simplices. In triangulations we will not only encounter triangles, but also segments and points. They are all special cases of the following definition. For every non-negative integer n, the *standard n–simplex* Δ^n is defined by

$$\{(t_0, t_1, \ldots, t_n) \in \mathbb{R}^{n+1} \mid t_0 + t_1 + \ldots + t_n = 1 \text{ and } t_i \geq 0 \text{ for all } i \geq 0\}.$$

We know $\Delta^0, \Delta^1, \Delta^2$ and Δ^3 respectively as a point, segment, triangle and tetrahedron.

A triangulated space can be seen as a collection of such simplices, with certain faces glued together. This is made precise in the following definition. This is actually a special case of a *Delta set* according to [3].

Definition 1.1. A *triangulation* $\mathbf{X}$ consists of a triple of (finite) sets X_0, X_1, X_2, together with *face maps* $d_i : X_n \to X_{n-1}$ for each $n \in \{1, 2\}$ and $0 \leq i \leq n$, such that $d_i d_j = d_{j-1} d_i$ for all $i < j$.

Define for each $0 \leq i \leq n$ the map $d^i : \Delta^{n-1} \to \Delta^n$ by $d^i : (t_0, t_1, \ldots, t_{n-1}) \mapsto (t_0, \ldots, t_{i-1}, 0, t_i, \ldots, t_{n-1})$. Then the *geometric realization* of this triangulation is given by the topological space

$$|\mathbf{X}| = \left(\coprod_{n=0}^{2} X_n \times \Delta^n \right) / \sim,$$

where the sets X_n have the discrete topology, and the equivalence relation $\sim$ is generated by the relation $(\sigma, d^i t) \sim (d_i \sigma, t)$ for all $\sigma \in X_n, t \in \Delta^{n-1}$.

Let's try to make some sense of this definition. The set X_n contains an element for every n–simplex in the triangulation. Intuitively we say X_0 is the set of vertices, X_1 is the set of edges and X_2 is the set of triangles. Each of these n–simplices has $n+1$ faces of dimension $n-1$, and the face maps $d_0, \ldots, d_n$ serve to pick out these faces in X_{n-1}. We can see this going on in the geometric realization. The space $\coprod_{n=0}^{2} X_n \times \Delta^n$ is the disjoint union of a number of simplices, one for every element of X_0, X_1 and X_2. Because $d_i \sigma$ is a face of σ, the simpex corresponding to $d_i \sigma$ should be identified to a face of the simplex corresponding to σ. This is exactly what happens in the relation $(\sigma, d^i t) \sim (d_i \sigma, t)$.

Finally, the condition $d_i d_j = d_{j-1} d_i$ for $i < j$ is also logical when we consider the geometric realization. It is not hard to see that $d^j d^i = d^i d^{j-1}$, hence we find that

$$(d_i d_j \sigma, t) \sim (\sigma, d^j d^i t) = (\sigma, d^i d^{j-1} t) \sim (d_{j-1} d_i \sigma, t).$$

So we find that in the geometric realization, the simplices $d_i d_j \sigma$ and $d_{j-1} d_i \sigma$ are identified. Demanding they be the same in the triangulation ensures the geometric realization is 'faithful'.

We will mainly consider triangulations whose geometric realizations are 2-dimensional manifolds. If this is the case, any point in any edge has a neighbourhood homeomorphic to $\mathbb{R}^2$. This is clearly only possible when there is a triangle on both sides of the edge. More formally, for every edge $e \in X_1$ there are exactly two pairs $(t,i) \in X_2 \times \{0,1,2\}$ such that $e = d_i t$.

Whenever we introduce a new object, we also want to know what the 'natural' maps between those objects look like. In this case, these 'natural' maps are *triangle maps*.

Definition 1.2. A *triangle map* $f : \mathbf{X} \to \mathbf{Y}$ between two triangulations $\mathbf{X}$ and $\mathbf{Y}$ consists of maps $f_n : X_n \to Y_n$ and $f_n^* : X_n \to S_{n+1}$ (where S_{n+1} is the symmetric group on $\{0,1,\ldots,n\}$), such that the following conditions are satisfied:

- $f_{n-1} d_i \sigma = d_{(f_n^*\sigma)(i)} f_n \sigma$ for all $\sigma \in X_n, n \in \{1,2\}$;
- $(f_n^*\sigma) d^i = d^{(f_n^*\sigma)(i)} (f_{n-1}^* d_i \sigma)$ for all $\sigma \in X_n, n \in \{1,2\}$, where an element of S_{n+1} acts on Δ^n by permuting the coordinates.

A triangle map is an *isomorphism* if the maps f_n are bijective.

A triangle map is called *unfolding* if for every pair of different triangles $\sigma, \sigma' \in X_2$ with a common edge, i.e. $d_i\sigma = d_{i'}\sigma'$ for some i, i', we have $(f_2\sigma, (f_2^*\sigma)(i)) \neq (f_2\sigma', (f_2^*\sigma')(i'))$.

The composition $g \circ f : \mathbf{X} \to \mathbf{Z}$ of two triangle maps $f : \mathbf{X} \to \mathbf{Y}$ and $g : \mathbf{Y} \to \mathbf{Z}$ is given by $(g \circ f)_n = g_n \circ f_n$ and $(g \circ f)_n^*(\sigma) = g_n^*(f_n\sigma) f_n^*(\sigma)$. A straightforward calculation verifies this is indeed a triangle map.

NB: From now on we will drop the subscripts in the maps and write $f\sigma$ for $f_n\sigma$ (and $f^*\sigma$ for $f_n^*\sigma$) if $\sigma \in X_n$.

The first condition of the triangle map makes sure that the faces of a simplex $\sigma \in X_n$ are mapped surjectively to the faces of $f\sigma \in Y_n$. If the faces of $f\sigma$ are all different, this uniquely determines the permutation $f^*\sigma$ (given that all faces of σ are different). It may seem that the permutations f^* are therefore superfluous in the definition, but they are important in constructing the geometric triangle map (see next definition). The second condition is needed to make sure the geometric triangle map can be defined.

Theorem 1.1. *A triangle map $f : \mathbf{X} \to \mathbf{Y}$ induces a continuous map $|f| : |\mathbf{X}| \to |\mathbf{Y}|$ by*

$$(\sigma, t) \mapsto (f\sigma, (f^*\sigma)t).$$

This map is called the geometric triangle map. *If f is an isomorphism, $|f|$ is a homeomorphism.*

Proof. We first prove this map is well-defined. To do this, we have to show that $(\sigma, d^i t)$ and $(d_i\sigma, t)$ (which are identified in $|\mathbf{X}|$) are sent to the same point. We see that $(\sigma, d^i t) \mapsto (f\sigma, (f^*\sigma) d^i t)$ and

$$(d_i\sigma, t) \mapsto (f d_i\sigma, (f^* d_i\sigma)t) = (d_{(f^*\sigma)(i)} f\sigma, (f^* d_i\sigma)t) \sim (f\sigma, d^{(f^*\sigma)(i)}(f^* d_i\sigma)t),$$

where we used the first condition of a triangle map. Now the second condition yields $d^{(f^*\sigma)(i)}(f^* d_i\sigma) = (f^*\sigma)d^i$, hence $(\sigma, d^i t)$ and $(d_i\sigma, t)$ are sent to the same point and $|f|$ is well-defined.

To see that $|f|$ is continuous, consider the following diagram:

$$\begin{tikzcd}
\coprod_{n=0}^{2} X_n \times \Delta^n \arrow[r, "\tilde{f}"] \arrow[d, "p_X"'] & \coprod_{n=0}^{2} Y_n \times \Delta^n \arrow[d, "p_Y"] \\
|\mathbf{X}| \arrow[r, dashed, "|f|"'] & |\mathbf{Y}|
\end{tikzcd}$$

where p_X and p_Y are the quotient maps. The map $\tilde{f}$ is the map specified in the theorem, and we have shown that this induces the map $|f|$ such that the diagram commutes. We easily see that $\tilde{f}$ is continuous; the components of $\coprod_{n=0}^{2} X_n \times \Delta^n$ are single simplices and $\tilde{f}$ is an embedding on each of this simplices. Then $|f| \circ p_X = p_Y \circ \tilde{f}$ is continuous as composition of two continuous maps, and by definition of the quotient space, $|f|$ is also continuous.

The second part of the theorem is easy. If f is an isomorphism, is has in inverse triangle map $g : \mathbf{Y} \to \mathbf{X}$ given by $g_n = (f_n)^{-1}$ and $g_n^* : \sigma \mapsto (f_n^* g_n \sigma)^{-1}$. The induced map $|g| : |\mathbf{Y}| \to |\mathbf{X}|$ is also an inverse of $|f|$. Then $|f|$ and its inverse $|g|$ are both continuous, hence $|f|$ is a homeomorphism. $\square$

1.1 Eisenstein triangulation

Now we have seen triangulations and triangle maps, let us consider an explicit triangulation. This specific triangulation will be very important in the rest of the thesis.

Definition 1.3. The *Eisenstein lattice* (also the *Eisenstein integers*) is the ring $\mathbb{Z}[\omega]$, where $\omega = e^{2\pi i/3} = -\frac{1}{2} + \frac{1}{2}i\sqrt{3}$, and is a subring of $\mathbb{C}$. (We will also use $\mathbb{Z}[\omega]$ to denote the underlying additive group.)

If we join the points of $\mathbb{Z}[\omega]$ at distance 1 by segments, we form triangles. The *Eisenstein triangulation*, denoted by **Eis**, has $\mathbb{Z}[\omega]$ as vertices, the segments of length 1 between them as edges, and the triangles are the triangles formed by the edges. If we denote the edge between a en b by $[a, b]$, and the triangle with vertices a, b, c as $[a, b, c]$, the face maps are given by

$$\begin{aligned}
&d_0([z, z+1, z+1+\omega]) = [z, z+1], && d_0([z, z+\omega, z+1+\omega]) = [z, z+\omega],\\
&d_1([z, z+1, z+1+\omega]) = [z, z+1+\omega], && d_1([z, z+\omega, z+1+\omega]) = [z, z+1+\omega],\\
&d_2([z, z+1, z+1+\omega]) = [z+1, z+1+\omega], && d_2([z, z+\omega, z+1+\omega]) = [z+\omega, z+1+\omega],
\end{aligned}$$

$$\begin{aligned}
&d_0([z, z+1]) = z, && d_0([z, z+\omega]) = z, && d_0([z, z+1+\omega]) = z,\\
&d_1([z, z+1]) = z+1, && d_1([z, z+\omega]) = z+\omega, && d_1([z, z+1+\omega]) = z+1+\omega,
\end{aligned}$$

for any $z \in \mathbb{Z}[\omega]$. The geometric realization is canonically homeomorphic to $\mathbb{C}$, and we will simply identify these two spaces.

Here we see the face maps in action: the maps d_0, d_1, d_2 give the three sides of the input triangle as output. To check that this is indeed a triangulation, we have to verify that $d_0 d_1 = d_0 d_0$, $d_0 d_2 = d_1 d_0$ and $d_1 d_2 = d_1 d_1$. Because the maps are explicitly given, this is a straightforward process.

1.2 Statement of the problem

Now we have some definitions, we can rigorously define the problem statement. We consider a triangulation $\mathbf{X}$ with S^2 (the sphere) as geometric realization. Because the sphere is a compact space, all sets of the triangulation must be finite. The vertices and edges naturally form a graph, which is embedded in the sphere. Because there are three face maps from each triangle to an edge, and each edge is reached exactly twice by the face maps, we find that $3|X_2| = 2|X_1|$. Hence there is an integer m such that $|X_2| = 2m$ and $|X_1| = 3m$. Because the sphere has Euler characteristic 2, we also find that $|X_0| - 3m + 2m = |X_0| - |X_1| + |X_2| = 2$, hence $|X_0| = m + 2$. We denote the degrees of these $m + 2$ vertices by $d_1, d_2, \ldots, d_{m+2}$. Then we find that

$$\sum_{i=1}^{m+2}(6 - d_i) = 6(m + 2) - \sum_{i=1}^{m+2} d_i = 6m + 12 - 2|X_1| = 12.$$

We are now interested in those triangulation for which $d_i \leq 6$ for all i. We will call these triangulations *non-negatively curved*. This implies there are at most 12 vertices having degree smaller than 6, we call these vertices *singular*. All non-singular vertices have degree 6.

An interesting question would be to count the number of different non-negatively curved triangulations of S^2, up to the isomorphism in definition 1.2. A rough asymptotic bound on this number is given by $O(n^{10})$, where n is the number of triangles [1], but the exact formula is still unknown. This general problem is outside of the scope of this thesis, but we will solve a special case. Based on ideas from Thurston [1], we will count the number of triangulations with four singular vertices of degree 3. This class will be called *tetrahedral triangulations*.

Chapter 2

Constructing a triangle map

To count the number of triangulations on the sphere, we will translate the problem to a counting problem in the Eisenstein lattice. The first step for this is the construction of an unfolding triangle map from the Eisenstein triangulation **Eis** to a tetrahedral triangulation **T**. The idea for this construction is to start with mapping one triangle in **Eis** to a triangle in **T**, and extending this map one by one to neighbouring triangles. To show that this construction is well-defined, we first have to do some work. We start with a lemma that gives us a tool to extend a triangle map.

Lemma 2.1. *Let* $\mathbf{X}$ *and* $\mathbf{X}'$ *be two triangulations that both contain exactly one triangle and only edges and vertices of this triangle. If* $\mathbf{X}$ *contains three different edges and three different vertices and* f *is a triangle map from one side of* $\mathbf{X}$ *to one side of* $\mathbf{X}'$*, then* f *can be extended to a triangle map from* $\mathbf{X}$ *to* $\mathbf{X}'$*.*

If $\mathbf{X}'$ *has three different edges this extension is unique.*

Proof. Write $X_0 = \{v_0, v_1, v_2\}, X_1 = \{e_0, e_1, e_2\}, X_2 = \{t\}$ such that the face maps are given by $d_i t = e_i$, $d_0 e_j = v_{\max_{i \neq j}(i)}$ and $d_1 e_j = v_{\min_{i \neq j}(i)}$. In the same way write $X'_0 = \{v'_0, v'_1, v'_2\}, X'_1 = \{e'_0, e'_1, e'_2\}, X'_2 = \{t'\}$ (where X'_0 and X'_1 may contains less than three elements), again with face maps $d_i t' = e'_i$, $d_0 e'_j = v'_{\max_{i \neq j}(i)}$ and $d_1 e'_j = v'_{\min_{i \neq j}(i)}$.

Assume f is given on the side e_0 and hence on the vertices v_1 and v_2 as well. For ease of notation we will assume $fe_0 = e'_1, fv_1 = v'_2, fv_2 = v'_0$, in other cases we can find similar constructions. (Note that this is a valid assumption. Because f is a triangle map, f_0 maps the endpoints of e_0, which are v_1 and v_2, surjectively on the endpoints of e'_1, which are v'_0 and v'_2.) Now we consider two cases: f^*e_0 can be either the transposition (01) or the identity.

If $f^*e_0 = (01)$ we define $ft = t', fe_1 = e'_2, fe_2 = e'_0, fv_0 = v'_1$ and $f^*e_1 = (01), f^*e_2 = \mathrm{Id}, f^*t = (012)$. To check that this indeed satisfies the conditions of a triangle map, we need to check the 18 given equations. These calculations are straightforward and show indeed that f is a triangle map.

Now we assume $f^*e_0 = \mathrm{Id}$. Because f is a triangle map on e_0, it follows that $v'_0 = fv_2 = fd_0e_0 = d_0fe_0 = d_0e'_1 = v'_2$. Now we choose $ft = t', fe_1 = e'_0, fe_2 = e'_2, fv_0 = v'_1$ and $f^*e_1 = \mathrm{Id}, f^*e_2 = (01), f^*t = (01)$. Using $v'_0 = v'_2$ it is again straightforward to check the necessary equalities, which shows that f is a triangle map. This concludes the proof of the existence. Now we will prove the statement about uniqueness.

We will again assume that $fe_0 = e'_1, fv_1 = v'_2, fv_2 = v'_0$. We will show that the triangle map f must always be on of the above, depending on f^*e_0. First of all it is clear that $ft = t'$. Because the sides of t must be mapped surjectively on the sides of t', we also have

$fe_1 = e'_2$ and $fe_2 = e'_0$, or $fe_1 = e'_0$ and $fe_2 = e'_2$. The endpoints of e_1, which are v_0 and v_2 must then be mapped surjectively to the endpoints of either e'_2 or e'_0. We now easily see that this forces $fv_0 = v'_1$. Finally, because $e'_1 = fe_0 = fd_0t = d_{(f^*t)(0)}ft = e'_{(f^*t)(0)}$ and because $\mathbf{X}'$ has three different sides, we also have $(f^*t)(0) = 1$.

Now we first assume $f^*e_0 = (01)$. We know that $(f^*t)d^0 = d^{(f*t)(0)}f^*e_0 = d^1(01)$ and if we apply this to a point $(x, y) \in \Delta^1$, we find that $(f^*t)(0, x, y) = (y, 0, x)$, hence $f^*t = (012)$. This implies that $fe_1 = fd_1t = d_2ft = e'_2$ and $fe_2 = fd_2t = d_0ft = e'_0$. Lastly we have to find f^*e_1 and f^*e_2. They follow from $(012)d^i = (f^*t)d^i = d^{(f^*t)(i)}f^*e_i$ and we easily find that $f^*e_1 = (01)$ and $f^*e_2 = \mathrm{Id}$.

If on the other hand $f^*e_0 = \mathrm{Id}$, we find that $(f^*t)d^0 = d^1$, hence $f^*t = (01)$. In the same way we find that $fe_1 = e'_0, fe_2 = e'_2$ and $f^*e_1 = \mathrm{Id}, f^*e_2 = (01)$. This finishes the proof. $\square$

2.1 Triangle paths

Now that we have a tool to extend triangle paths, we define *triangle paths*, along which we will extend the triangle map.

Definition 2.1. A *triangle path* of length n in a triangulation $\mathbf{X}$ is a sequence $(t_1, t_2, \ldots, t_n)$ of triangles such that for each $i \in \{1, \ldots, n-1\}$ the triangles t_i and t_{i+1} satisfy $d_jt_i = d_kt_{i+1}$ (i.e. they have a common edge) for some j, k with $(t_i, j) \neq (t_{i+1}, k)$.

A triangle path $(u_1, \ldots, u_n)$ depends continuously on $(t_1, \ldots, t_n)$ if for every i there is a triangle map $f_i : t_i \to u_i$ such that f_i and f_{i+1} coincide on the common edge of of t_i and t_{i+1}.

Because all degrees in **Eis** and $\mathbf{T}$ are greater than 1, every triangle has three different sides. If $\mathbf{X}$ is one of these triangulations, this means that the condition $(t_i, j) \neq (t_{i+1}, k)$ is equivalent to $t_i \neq t_{i+1}$. We give this slightly more complicated definition to show that similar arguments can be applied if $\mathbf{X}$ would contain vertices of degree 1. Because all triangles in **Eis** and $\mathbf{T}$ have three different sides, we can also apply the second part of 2.1 which gives unique extensions.

There is one more subtlety if there are vertices of degree 2. In that case the common edge of two neighbouring triangles is not unique and this would complicate the definition of the continuously depending triangle paths. To solve this, we have to include in the triangle path which is the common edge of two neighbouring triangles. As this is notationally cumbersome and not necessary in our case, we will not do this.

Note that the subscript in the notation f_n is different from the subscript used in the formal definition of the triangle map. In this case, the subscripts denote different triangle maps.

The next theorem is very important. It shows that it does not matter along which triangle path we extend the triangle map and therefore we can make a valid construction. It is also in this theorem that we rely on the fact that $\mathbf{T}$ has only vertices of degree 3 and 6.

Theorem 2.1. *(a) Let $(u_1, \ldots, u_{n+1})$ be a triangle path in $\mathbf{T}$ which depends continuously on the triangle path $(t_1, \ldots, t_{n+1})$ in* **Eis** *through the maps $f_1, \ldots, f_{n+1}$. If $t_{n+1} = t_1$, then $u_1 = u_{n+1}$ and $f_1 = f_{n+1}$.*

(b) Let $(u_1, \ldots, u_n)$ be a triangle path in $\mathbf{T}$ which depends continuously on the triangle path $(t_1, \ldots, t_n)$ in **Eis** *through the maps $f_1, \ldots, f_n$, and let $(u'_1, \ldots, u'_m)$ be a triangle*

path in $\mathbf{T}$ which depends continuously on the triangle path $(t'_1, \ldots, t'_m)$ in **Eis** *through the maps $f'_1, \ldots, f'_m$. If $t_1 = t'_1$, $f_1 = f'_1$ and $t_n = t'_m$, then $u_n = u'_m$ and $f_n = f'_m$.*

Proof. (a) Let p_i be the midpoint of t_i, then $p_1 p_2 \ldots p_n p_1$ forms a closed curve. Moreover, if we connect the midpoints for every pair of adjacent triangles in **Eis**, we obtain a tessellation of $\mathbb{C}$ with regular hexagons. The loop $p_1 p_2 \ldots p_n p_1$ then contains a finite number of these hexagons. We will now prove the statement by strong induction on $K = n + 5 \cdot$ (number of enclosed hexagons).

Base. The base case is $K = 2$, with $n = 2$ and no hexagons enclosed. In this case f_1 and f_2, as well as f_2 and f_3, coincide on the common edge of t_1 and t_2. This implies that the common edge of u_1 and u_2 is also the common edge of u_2 and u_3. Because there cannot be three triangles with the same common edge, we see that $u_1 = u_3$. Now f_1 and f_3 are both triangle maps from t_1 to u_1 that coincide on one edge. Now lemma 2.1 implies that f_1 and f_3 are the same.

Step. Now take a triangle path $(t_1, \ldots, t_n, t_1)$ with a certain value of K, and assume the result holds for all circular paths with smaller value of K.

Assume $t_i = t_j$ for some $i < j$. Then the path $(u_i, \ldots, u_j)$ depends continuously on $(t_i, \ldots, t_j = t_i)$. This path has length $j - i < n$ and encloses no more hexagons than the original path. Therefore the value of K is smaller and by induction hypothesis we know that $u_i = u_j$ and $f_i = f_j$. Then we see that $(u_1, \ldots, u_{i-1}, u_i = u_j, u_{j+1}, \ldots, u_n, u_{n+1})$ depends continuously on $(t_1, \ldots, t_{i-1}, t_i = t_j, t_{j+1}, \ldots, t_n, t_1)$. This path has length $n - j + i < n$ and also encloses no more hexagons than the original path, so again by the induction hypothesis we have $u_{n+1} = u_1$ and $f_{n+1} = f_1$. This concludes the step in this case.

We now consider the case where $t_i \neq t_j$ for all $i \neq j$. The loop $p_1 p_2 \ldots p_n p_1$ is then a *simple* closed curve if $n > 2$. If the loop encloses no hexagons, we will show that we are in the base case. Namely, if $n > 2$, the two hexagons on both sides of $p_1 p_2$ are separated by the loop $p_1 p_2 \ldots p_n p_1$, and one of them must be contained in the loop. This is a contradiction, hence $n = 2$ and we are in the base case.

Now assume that the loop encloses a positive number of hexagons. The side $p_n p_1$ borders two hexagons, one of which is contained in the loop. If we let $p_n, p_{-3}, p_{-2}, p_{-1}, p_0, p_1$ be the vertices of this hexagon, in order, then $p_1 \ldots p_n p_{-3} p_{-2} p_{-1} p_0 p_1$ is a closed loop which contains one fewer hexagon. Let $t = (t_1, \ldots, t_n, t_{-3}, t_{-2}, t_{-1}, t_0, t_1)$ be the corresponding triangle path. Because it contains one less hexagon and is four triangles longer, the value of K is one smaller and we can apply the induction hypothesis on t.

Let x be the common vertex of the triangles $t_n, t_{-3}, t_{-2}, t_{-1}, t_0, t_1$. Then $f_n(x)$ is a vertex of u_n and u_{n+1}. Because all vertices in $\mathbf{T}$ have degree 3 or 6, we may name the triangles around this vertex as $u_n, u_{-3}, u_{-2}, u_{-1}, u_0, u_{n+1}$, where $(u_n, u_{-3}, u_{-2}) = (u_{-1}, u_0, u_{n+1})$ if the degree of $f_n(x)$ is 3. We will now show that the path $u = (u_1, \ldots, u_n, u_{-3}, u_{-2}, u_{-1}, u_0, u_{n+1})$ depends continuously on t and the corresponding map $t_1 \to u_{n+1}$ is still f_{n+1}.

To do this, we construct maps $f_i : t_i \to u_i$ for $i = -3, -2, -1, 0$. We do this with lemma 2.1: f_{-3} must coincide with f_n on the common edge of t_n and t_{-3}, hence there is a unique extension to a map $f_{-3} : t_{-3} \to u_{-3}$. In a similar fashion we define f_{-2} to f_0. Now f_0 and f_{n+1} map t_0 and t_1 to the adjacent triangles u_0 and u_{n+1}, and they also map the common edge of t_0 and t_1 to the common edge of u_0 and u_{n+1}. Of course they also coincide on x, which implies they also coincide on the other endpoint of the common edge of t_0 and t_1. Now we still have to show that they give the same permutation on this common edge. For this it is enough to show that a point on the common edge near x is sent to the same point by $|f_0|$ and $|f_{n+1}|$. Consider a small circle around x; we map this circle to $|\mathbf{T}|$ with the maps $|f_{n+1}|, |f_n|, |f_{-3}|, \ldots, |f_0|$. If we traverse this circle starting at the common edge of t_0 and t_1 we encounter six edges (including the last edge). If we look where this path is mapped in $|\mathbf{T}|$, we see that we walk in a small loop around $f_n(x)$ and also encounter six edges. Because the degree of $f_n(x)$ is 3 or 6, we must end in the same point as where we started. Hence f_0 and f_{n+1} coincide on the common edge of t_0 and t_1. We now see that u depends continuously on t with the maps $f_1 \ldots, f_n, f_{-3}, f_{-2}, f_{-1}, f_0, f_{n+1}$ and by the induction hypothesis we conclude that $u_{n+1} = u_1$ and $f_{n+1} = f_1$.

This concludes the induction and proves part (a).

(b) From the given we conclude that $(u_n, \ldots, u_1, u_2', \ldots, u_m')$ depends continuously on $(t_n, \ldots, t_1, t_2', \ldots, t_{m-1}', t_n)$. Now we can apply part (a) and we find directly $u_n = u_m'$ and $f_n = f_m'$. □

2.2 Unfolding triangle maps from Eis to T

Now we will go on to construct the triangle map we wanted.

Theorem 2.2. *There exists an unfolding triangle map from* **Eis** *to* **T***.*

Proof. First we will choose for every triangle t in **Eis** a triangle u in **T**, together with a triangle map from t to u. Then we show that these maps are compatible and hence can be put together to form a map $f : \mathbf{Eis} \to \mathbf{T}$.

Choose a vertex A in **T** which has degree 3 and choose a triangle u_0 which has A as vertex. Let $t_0 = \triangle(0, 1, \omega + 1)$, then we can form a triangle map $f_0 : t_0 \to u_0$ such that $f_0(0) = A$. For any other triangle $t \in \mathbf{Eis}$ we can find a triangle path $(t_0, t_1, \ldots, t_n = t)$. We will construct a triangle path $(u_0, u_1, \ldots, u_n)$ that depends continuously on this path. This can easily be done inductively: the common edge of t_0 and t_1 is mapped to an edge of u_0, which is the edge of exactly one other triangle, let this be u_1. Now f_1 has to coincide with f_0 on the common edge of t_0 and t_1, so f_1 has a unique extension to t_1. This argument can be repeated until the whole path $(u_0, u_1, \ldots, u_n)$ and a triangle map $f_n : t \to u_n$ has been constructed. Theorem 2.1 now tells us that u_n and the map $f_n : t \to u_n$ do not depend on the chosen path $(t_0, \ldots, t_n)$.

We can apply the same method for every triangle $t \in \mathbf{Eis}$ to find a triangle $u_t \in \mathbf{T}$ and a triangle map $f_t : t \to u_t$. We want to put all maps together to form a triangle map $f : \mathbf{Eis} \to \mathbf{T}$, but this is only possible if f_t and $f_{t'}$ coincide on $t \cap t'$. We will show that this is indeed the case.

First consider the case that t and t' have one common edge. Then we can choose a triangle path from t_0 to t, and extend it with (t', t). The triangle path in **T** which depends continuously on this path ends in $(u_t, u_{t'}, u_t)$, and the triangle maps are just $f_t, f_{t'}$ and f_t. Then by definition, f_t and $f_{t'}$ coincide on the common edge of t and t'.

If t and t' only coincide in a vertex x, we can either find one triangle t_a that has x as vertex such that (t, t_a, t') is a triangle path, or two triangles t_a, t_b that have x as vertex such that (t, t_a, t_b, t') is a triangle path. Because the triangle maps coincide on the common edges, they coincide pairwise on x, but then f_t and $f_{t'}$ also coincide in x.

This concludes the construction of the triangle map $f : \mathbf{Eis} \to \mathbf{T}$. The only thing left to show is that it is an unfolding map. Consider two triangles $t, t' \in \mathbf{Eis}$ that share a side. The corresponding triangles in **T** are $u_t, u_{t'}$ and the path $(u_t, u_{t'})$ depends continuously on (t, t'). If the common edge of u_t and $u_{t'}$ is $d_j u_t = d_k u_{t'}$, then it follows from the definition of a triangle path that $(u_t, j) \neq (u_{t'}, k)$, which we needed to show. □

When we consider this construction, we see that there is not much choice once we have picked f_0. This might be caused by this specific construction, but there actually is no choice, as the following theorem shows.

Theorem 2.3. *Let two unfolding triangle maps $f, f' : \mathbf{Eis} \to \mathbf{T}$ be given, that coincide on one triangle of $\mathbf{Eis}$. Then $f = f'$.*

Proof. Let t_0 be the triangle on which f and f' coincide, and let t be another triangle. Choose a triangle path $(t_0, t_1, \ldots, t_n = t)$, and define $u_i = f(t_i)$ and $u'_i = f'(t_i)$. We first show that $(u_0, u_1, \ldots, u_n)$ is a triangle path. Consider triangles t_i, t_{i+1} with a common edge, so $d_j t_i = d_k t_{i+1}$. Because f is unfolding, we have $(u_i, (f^* t_i)(j)) \neq (u_{i+1}, (f^* t_{i+1})(k))$. And because $d_{(f^* t_i)(j)} u_i = f d_j t_i = f d_k t_{i+1} = d_{(f^* t_{i+1})(k)} u_{i+1}$ is the common edge of u_t and u_{t+1}, this implies exactly that $(u_0, u_1, \ldots, u_n)$ is a triangle path. It clearly also depends continuously on $(t_0, t_1, \ldots, t_n)$. But the same holds for $(u'_0, u'_1, \ldots, u'_n)$. Applying theorem 2.1, we see that $u_n = u'_n$ and also that f and f' coincide on t. This holds for every triangle $t \in \mathbf{T}$, hence $f = f'$. $\square$

Chapter 3

Symmetry of the Eisenstein triangulation

In the previous section we constructed an unfolding triangle map $f : \mathbf{Eis} \to \mathbf{T}$ and we proved this map to be unique if it was given on one triangle. In this section we will see that f can be regarded as a quotient map from a group action and therefore f can be characterised by this group.

Let G be the group of isometries of $\mathbb{C}$ that map $\mathbb{Z}[\omega]$ to itself. Each isometry induces an isomorphism from $\mathbf{Eis}$ to itself. Note that each of these isomorphisms is uniquely determined by its restriction to one triangle.

Theorem 3.1. *Given an unfolding triangle map $f : \mathbf{Eis} \to \mathbf{T}$, let $G_f = \{g \in G \mid f = f \circ g\}$ be a subgroup of G. The group G_f acts on $\mathbf{Eis}$ and defines orbits of $\mathrm{Eis}_0, \mathrm{Eis}_1$ and Eis_2. Two simplices σ, σ' are in the same orbit if and only if $f\sigma = f\sigma'$.*

Proof. If two simplices σ, σ' are in the same orbit, there is a $g \in G_f$ such that $g\sigma = \sigma'$ and by definition of G_f we see that $f\sigma' = fg\sigma = f\sigma$.

Now assume that $f\sigma = f\sigma'$ for different σ, σ'. First assume both are triangles. There is an isomorphism $g \in G$ that sends σ to σ'. We can choose g such that $f(e) = f \circ g(e)$ for every side e of σ. Now we can easily show that $f^*t = (f \circ g)^*t$ (because $f\sigma$ has three different edges), from which we can deduce that $f^*e = (f \circ g)^*e$ for all the edges and therefore they also coincide on the vertices. Hence $f = f \circ g$ on the triangle σ and because both f and $f \circ g$ are unfolding, we find $f = f \circ g$. Hence $g \in G_f$ and σ and σ' are in the same G_f–orbit.

Now assume σ and σ' are edges. Choose a triangle t which contains σ, then we can find a triangle t' that contains σ' such that $ft = ft'$. Now we repeat the above argument to find an isomorphism $g \in G_f$ such that $gt = t'$. Then $g\sigma$ is an edge of t and it must be σ', because it is the only edge of t' which is sent to $f\sigma$ (because ft has three different sides). Hence σ and σ' are in the same G_f–orbit.

Last we assume σ and σ' are vertices. Then there is an edge e containing σ such that the other endpoint of e is not mapped to $f\sigma$. We also find an edge e' which contains σ' such that $fe = fe'$. Now we use the above argument to see that there is an isomorphism $g \in G_f$ such that $ge = e'$. Then $g\sigma$ is an endpoint of e' and because fe has two different endpoints, we must have $g\sigma = \sigma'$. Hence σ and σ are in the same G_f–orbit. This concludes the proof. $\square$

3.1 From triangulation to lattice

Now we want to do two things. Firstly, given this group, can we reconstruct the space $\mathbf{T}$? In other words, can we form a quotient of a triangulation? We will see that this is indeed possible, but we first address another issue: we want to know how the group looks like. We introduce a bit of notation for this. Denote the *half-turn about* x by $H_x : z \mapsto 2x - z$, and the *translation over* x by $T_x : z \mapsto z + x$.

Theorem 3.2. *Given an unfolding map $f : \mathbf{Eis} \to \mathbf{T}$, let S the set of singular vertices in $\mathbf{T}$, and let $L = f^{-1}(S)$. Then G_f is generated by $\{H_x \mid x \in L\}$.*

Proof. We first prove that $H_x \in G$ if and only if $x \in L$. If $x \in L$, we know that fx is a vertex of degree 3. Consider a triangle path of four triangles with common vertex x, such that the first is the mirror image of the last. This triangle path maps to a triangle path with common vertex fx, and because fx has degree 3, the first and the last triangle are the same. If we also consider where the sides are sent to, we see that f and $f \circ H_x$ coincide on the first triangle of the path. Because both maps are unfolding, we see that $f = f \circ H_x$ and therefore $H_x \in G_f$.
On the other hand, if $H_x \in G_f$ and x is a vertex of $\mathbf{Eis}$, then the edges ending in x are sent pairwise to the same edge in $\mathbf{T}$. This means that fx has degree 3, hence $x \in L$. If x is not a vertex, it must be the midpoint of an edge. In this case the half-turn H_x maps this edge to itself, but mirrored. Then it is impossible that $f = f \circ H_x$ on this edge, which yields a contradiction.

By definition of f, we see that $0 \in L$. If we take $x, y \in L$, we see that $H_x \circ H_y \circ H_0 = H_{x-y}$ and this is an element of G_f. Then $x - y \in L$, so L is an additive group.

We also know that $H_x \circ H_y$ is in G_f for $x, y \in L$, but this composition is a translation over $2(x - y)$. If we denote by T_a the translation over a, this implies $\{T_{2x} \mid x \in L\} \subset G_f$. We want to prove that these half-turns and translations are the only elements of G_f. We do this by considering the different types of possible isometries of $\mathbb{C}$.

Assume $T_y \in G_f$ for a y not in $2L$, then $T_y \circ H_0 = H_{y/2} \in G_f$, so $y/2 \in L$. This contradicts $y \notin 2L$.

If there is a rotation in G_f other than a half-turn, it must be a 3-fold or 6-fold rotation. Its center must be one of the vertices of $\mathbf{Eis}$, and by an argument similar to the one for the half-turns, we see that this vertex maps to a vertex of degree 2 or 1 in $\mathbf{T}$. Because $\mathbf{T}$ contains only vertices of degree 3 or 6, this is a contradiction, hence there are no other rotations in G_f.

Assume there is a reflection in G_f. If the mirror axis is the perpendicular bisector of an edge, this edge is sent to itself but mirrored. We have already seen that this is impossible. Now assume the mirror axis runs along edges and consider one of these edges e. The triangles on both side of e are then in the same orbit. For any other edge e' in the same orbit as e, there is an element of G_f sending e' to e, which sends both triangles bordering e' to the triangles bordering e, hence all those triangles are also in the same orbit. This means that there is only one triangle bordering fe, which contradicts the fact that $\mathbf{T}$ is the triangulation of a 2-dimensional manifold.

For the last case, assume that G_f contains a glide reflection. To handle this case, we first need to find out more about L. First note that $2L$ is in one orbit of G_f, because $\{T_{2x} \mid x \in L\} \subset G_f$. Each of its cosets is therefore also in one orbit, but $\mathbf{T}$ has four singular points. This means there must be at least four cosets of $2L$ in L. Because L is a subgroup of $\mathbb{Z}[\omega]$ which is isomorphic to $\mathbb{Z}^2$, the rank of L is at most 2. If L has rank smaller than 2, the index of $2L$ is at most 2. Hence the rank of $2L$ is exactly 2 and each of its cosets is exactly one G_f–orbit.

Now consider a glide reflection which is a reflection in line ℓ and a translation along v. If ℓ contains a point x from L, we can compose the glide reflection with H_x. This is a reflection in the line perpendicular to ℓ through $x - v/2$, but we already proved that G_f contains no reflections. Hence ℓ contains no points from L. Consider a point $x \in L$ and let x' be its image under the glide reflection. These two points are in the same orbit, hence $x' - x \in 2L$. This means that the midpoint of x' and x is in L, but the midpoint is on ℓ. This is a contradiction, hence G_f contains no glide reflections. This completes the proof. □

3.2 From lattice to triangulation

This theorem tells us what the group G_f looks like and also that the group is completely determined by the *lattice* L. A lattice in $\mathbb{C}$ is a subgroup isomorphic to $\mathbb{Z}^2$ which spans the whole of $\mathbb{C}$. These lattices are well-known, hence the counting problem for these lattices will be easier. So for know we concentrate on the reverse relation: showing that every lattice gives a tetrahedral triangulation.

Theorem 3.3. *Given a sublattice L of $\mathbb{Z}[\omega]$, there is a tetrahedral triangulation $\mathbf{T}$ and an unfolding triangle map $f : \mathbf{Eis} \to \mathbf{T}$ such that $L = f^{-1}(S)$, where S are the singular points of $\mathbf{T}$.*

Proof. Construct the group $G_L = \{H_x \mid x \in L\} \cup \{T_{2x} \mid x \in L\}$. This group acts on $\mathbf{Eis}$ and defines orbits in $\mathrm{Eis}_0, \mathrm{Eis}_1$ and Eis_2. We now define $\mathbf{T}$ by letting T_0, T_1, T_2 be the orbits in respectively $\mathrm{Eis}_0, \mathrm{Eis}_1, \mathrm{Eis}_2$. Note that by definition of G_L it is impossible for the two endpoints of an edge to be in the same orbit. The same goes for the three edges of a triangle.

To finish the definition of $\mathbf{T}$, we need to define the face maps. We start by noting that T_0, T_1 and T_2 are finite, hence we can write $T_0 = \{x_1, \dots, x_n\}$ (choosing a different order will yield a different, but isomorphic triangulation). For an orbit in T_1, we can choose a representative edge $e \in \mathrm{Eis}_1$ and look at its endpoints. The orbit of these endpoints does not depend on the chosen edge, because an isometry in G_L sends the endpoints of an edge to the endpoints of the image of the edge. If the endpoints of e are in the orbits x_i and x_j, with $i < j$, we let $d_0([e]) = x_j$ and $d_1([e]) = x_i$. Similarly for an orbit in T_2 we can choose a representative triangle $t \in \mathrm{Eis}_2$ and look at its edges and vertices. If the vertices are in the orbits x_i, x_j, x_k with $i < j < k$, we let $d_0([t])$ be the orbit of the edge of t between x_j and x_k, $d_1([t])$ is the orbit of the edge of t between x_i and x_k, and $d_2([t])$ is the orbit of the edge of t between x_i and x_j. We can easily check that this is indeed a triangulation according to the definition.

We also must construct an unfolding triangle map $f : \mathbf{Eis} \to \mathbf{T}$ for this new triangulation. The maps $f_n : \mathrm{Eis}_n \to T_n$ are simply the quotient maps. The maps $f_n^* : \mathrm{Eis}_n \to S_{n+1}$ then follow uniquely from the maps f_n, because every triangle has three different edges and every edge has two different vertices. Again we can easily check that this is a triangle map according to the definition. By the explicit form of G_L we also see that two triangles that share an edge are never in the same orbit. This implies that f is also unfolding.

Last we have to show that $\mathbf{T}$ is indeed a tetrahedral triangulation. The degrees of the vertices are good: every coset of $2L$ maps to a vertex of degree 3, which gives four in total. All other vertices have degree 6. We also need to show that $|\mathbf{T}|$ is a sphere. We do this by considering the map $|f| : \mathbb{C} \to |\mathbf{T}|$. It is easy to see that this map is the pointwise quotient map by the action of G_L. To see that it is also the topological quotient map, we need to show that $A \subset |\mathbf{T}|$ is open if and only if $|f|^{-1}(A) \subset \mathbb{C}$ is open. But A is open if and only if it is open in every simplex (without the identification of faces). The map $\tilde{f}$ (introduced in theorem 1.1) is an embedding on every simplex and because it is surjective, it is also a projection, and hence a set is open if and only if its inverse image is open. Prjoecting down again to $\mathbb{C}$ shows the final equivalence. Hence $|\mathbf{T}|$ is homeomorphic to the topological quotient $\mathbb{C}/G_L$, which is homeomorphic to the sphere. □

We now have two constructions: from a triangulation we can construct a lattice and from a lattice we can construct a triangulation. If a triangulation generates a lattice, then we construct from this lattice an isomorphic triangulation. However, there may be multiple lattices belonging to the same triangulation. This is made precise in the following theorem.

Theorem 3.4. *Let D_6 be the group of isometries of* **Eis** *mapping 0 to itself. Two sublattices of $\mathbb{Z}[\omega]$ generate isomorphic triangulations if and only if they are in the same D_6–orbit.*

Proof. If two lattices L, L' are in the same D_6–orbit, there is an isometry which maps L to L'. This isometry also maps the orbits of G_L to the orbits of $G_{L'}$ hence induces an isomorphism between the two triangulations.

Now assume two lattices L and L' have isomorphic triangulations $\mathbf{T}$ and $\mathbf{T}'$ with respective unfolding triangle maps f, f'. We will show that there is an isometry I of **Eis** such that the following diagram commutes:

$$\begin{tikzcd}
\mathbf{Eis} \arrow[r, dashed, "I"'] \arrow[d, "f"'] & \mathbf{Eis} \arrow[d, "f'"] \\
\mathbf{T} \arrow[r, "\sim"'] & \mathbf{T}'
\end{tikzcd}$$

To do this, pick a triangle t in **Eis** and map it via f and $\sim$ to $\mathbf{T}'$. Now choose a triangle t' in its inverse image from f'. Then we can choose an isometry I which maps t to t' such that $\sim \circ f$ and $f' \circ I$ coincide on t. Because both maps are unfolding, they are equal by theorem 2.3. Now we decompose I as composition of an element $g \in D_6$ and a translation. Because 0 is in both L and L', and g fixes 0, the translation maps L' to itself. But then g maps L to L' which is what we wanted to prove. □

The last thing we note is that we can find the number of triangles in a triangulation from the lattice. Consider a sublattice L of $\mathbb{Z}[\omega]$, for each of these two lattices we can choose two generators and they span a parallelogram. The ratio of the areas of these parallelograms is exactly the index $|\mathbb{Z}[\omega] : L|$. The group action of G_L on $\mathbb{C}$ has a fundamental domain formed by joining two of the parallelograms of L along a side, so this fundamental domain has area $2|\mathbb{Z}[\omega] : L| \cdot 2\text{Area}(\triangle(0, 1, \omega + 1))$ and it contains therefore exactly $4|\mathbb{Z}[\omega] : L|$ triangles. In particular we see that the number of triangles in a tetrahedral triangulation is always a multiple of 4.

Now all the information we need about the triangulations is stated in terms of lattices and it only remains to count these lattices.

Chapter 4

Counting lattices

In the previous chapter we concluded that tetrahedral triangulations with $4m$ triangles correspond bijectively to sublattices of $\mathbb{Z}[\omega]$ of index m, modulo the action of a dihedral group of order 12. So let's count these lattices.

Theorem 4.1. *There are exactly $\sum_{d|m} d$ sublattices L of $\mathbb{Z}[\omega]$ with index m.*

Proof. Let a lattice L of index m be given. Pick two generators $n_1 + n_2\omega$ and $n_3 + n_4\omega$ of L (n_i integers), then $n_4(n_1 + n_2\omega) - n_2(n_3 + n_4\omega) = n_1n_4 - n_2n_3$ is a non-zero element of L (if $n_1n_4 - n_2n_3$ were zero, the generators would not be linearly independent). This means that $L \cap \mathbb{Z}$ is a non-trivial subgroup of $\mathbb{Z}$, and therefore it is of the form $d\mathbb{Z}$ for some $d > 0$. Now pick an element $a + b\omega$ of L with $b > 0$ and b minimal. Because $d \in L$, we may also assume $0 \leq a < d$. Then d and $a + b\omega$ also generate L. To see this, let $k + \ell\omega$ be an element of L. Then

$$(k - a\lfloor \tfrac{\ell}{b} \rfloor) + (\ell - b\lfloor \tfrac{\ell}{b} \rfloor)\omega$$

is also in L and $0 \leq \ell - b\lfloor \frac{\ell}{b} \rfloor < b$. Because $a + b\omega$ is an element with b minimal, this means that $\ell - b\lfloor \frac{\ell}{b} \rfloor = 0$. Then $k - a\lfloor \frac{\ell}{b} \rfloor \in \mathbb{Z} \cap L$ is a multiple of d, hence $k + \ell\omega$ is indeed generated by d and $a + b\omega$.

Now the parallelogram spanned by d and $a + b\omega$ has area $bd \cdot \frac{1}{2}\sqrt{3}$, so the index m of L is equal to bd. This means that $d \mid m$ and if we fix d, we also fix b. Then there are exactly d choices for a, so the number of lattices with index m is equal to $\sum_{d|m} d$. □

This is the total number of lattices without accounting for symmetry. So that is what we will do now, with use of Burnside's lemma.

Theorem 4.2. *Let $g(m)$ be the number of sublattices of $\mathbb{Z}[\omega]$ of index m that are invariant under multiplication with ω. Let $t(m)$ be the number of sublattices of $\mathbb{Z}[\omega]$ of index m that are invariant under the conjugation automorphism of $\mathbb{C}$. Then there are*

$$\frac{1}{6}\sum_{d|m} d + \frac{1}{3}g(m) + \frac{1}{2}t(m)$$

tetrahedral triangulations with $4m$ triangles.

Proof. Burnside's lemma states that the number of orbits given by a group action of G on X, is equal to $\frac{1}{|G|}\sum_{g\in G} |X^g|$, where X^g are the fixed points of g. In case of the action of D_6 on the set of sublattices of $\mathbb{Z}[\omega]$ with index m, the group order is 12 and we will count the number of fixed elements of every element. The number of orbits is then exactly the number we want.

For any lattice L we have $x \in L \Leftrightarrow -x \in L$, hence the identity and the half-turn about 0 both leave all lattices fixed. This gives two terms $\sum_{d|m} d$. The other rotations in D_6 correspond with multiplication by $\omega, -\omega, \omega^{-1}, -\omega^{-1}$, and they all leave $g(m)$ lattices invariant. This gives four terms $g(m)$.

Last we have the reflections which can be divided into two conjugacy classes: one contains the reflection in the real axis, the other reflection in the imaginary axis. However, both have the same number of fixpoints: composing one reflection with the half-turn about 0 gives the other reflection. As conjugate elements also have the same number of fixpoints, we find that each reflection has $t(m)$ fixpoints.

Hence the total number of orbits is

$$\frac{1}{12}\left(2\sum_{d|m} d + 4g(m) + 6t(m)\right) = \tfrac{1}{6}\sum_{d|m} d + \tfrac{1}{3}g(m) + \tfrac{1}{2}t(m). \qquad \square$$

4.1 Number theory

We will now derive forms for the functions $g(m)$ and $t(m)$. To determine $g(m)$ we do some number theory in $\mathbb{Z}[\omega]$, hence we need some knowledge of $\mathbb{Z}[\omega]$. It has a norm given by

$$N(a + b\omega) = (a + b\omega)(a + b\overline{\omega}) = a^2 - ab + b^2.$$

Using this norm we can show that $\mathbb{Z}[\omega]$ is a Euclidean domain. Because of this it is also a principal ideal domain (PID) and a unique factorization domain (UFD). Because it is a UFD, the prime elements and the irreducible elements are the same. We will often make implicit use of this. [4]

The number $t(m)$ is easier to derive, it uses only elementary number theory.

Theorem 4.3. *The function $g(m)$ is multiplicative, that is, $g(m)g(n) = g(mn)$ for all coprime integers m, n. For prime powers, g is given by*

$$g(p^k) = \begin{cases} 1 & \text{if } p = 3, \\ k+1 & \text{if } p \equiv 1 \mod 3, \\ 0 & \text{if } p \equiv 2 \mod 3 \text{ and } k \equiv 1 \mod 2, \\ 1 & \text{if } p \equiv 2 \mod 3 \text{ and } k \equiv 0 \mod 2. \end{cases}$$

Proof. This proof is based on a similar result from [5].

We first show that $g(m)$ is equal to the number of numbers $\alpha \in \mathbb{Z}[\omega]$ with norm m, modulo multiplication with units. Note that the norm is defined as $N(\alpha) = \alpha\overline{\alpha}$ and the units are $1, \omega, \omega^2, -1, -\omega, -\omega^2$.

Given a sublattice L that is invariant under multiplication with ω, then L is an ideal in $\mathbb{Z}[\omega]$. Because $\mathbb{Z}[\omega]$ is a PID, L is generated by a single element α. The choice of α is unique up to multiplication with units. We also see that the index of L is equal to the area of the parallelogram spanned by α and $\omega\alpha$, divided by the area of the parallelogram spanned by 1 and ω. This ratio is exactly $|\alpha|^2 = N(\alpha)$. Hence every lattice with index m corresponds to a number α with norm m, modulo multiplication with units.

We first will prove that g is multiplicative. We write $\alpha \equiv \beta$ to indicate the equality is modulo multiplication with units. For all α with norm m and β with norm n, we see that $\alpha\beta$ has norm mn. This yields $g(m)g(n)$ numbers of norm mn, but we need to show that they are all different. So assume $\alpha\beta \equiv \alpha'\beta'$ and let π be an Eisenstein prime that divides α. If it also divides β', then $N(\pi)$ would divide m and n, hence it would divide 1. This is a contradiction, hence α and β' have no common prime factors. As $\alpha \mid \alpha'\beta'$, this implies $\alpha \mid \alpha'$. Similarly we find that $\alpha' \mid \alpha$, hence $\alpha \equiv \alpha'$ and therefore $\beta \equiv \beta'$ as well. This means that $g(mn) \geq g(m)g(n)$.

Now assume γ has norm mn, and write $\alpha \equiv (\gamma, m), \beta \equiv (\gamma, n)$. We want to prove that $\alpha\beta \equiv \gamma$ and that α and β have norms m and n. Let π^k be an Eisenstein prime power. If π^k divides γ, we have $\pi^k \mid \gamma\overline{\gamma} = mn$. As m and n are coprime, this means that either $\pi^k \mid m$ or $\pi^k \mid n$. In the first case we see that $\pi^k \mid \alpha$, because it divides both γ and m, hence $\pi^k \mid \alpha\beta$. In the second case a similar argument shows $\pi^k \mid \alpha\beta$. This holds for all Eisenstein prime powers, hence $\gamma \mid \alpha\beta$. Now if $\pi^k \mid \alpha\beta$, we have either $\pi^k \mid \alpha$ or $\pi^k \mid \beta$, because $(\alpha, \beta) \mid (m, n) = 1$. But $\pi^k \mid \alpha$ gives $\pi^k \mid \gamma$ because $\alpha \mid \gamma$, and in the same way $\pi^k \mid \beta$ gives $\pi^k \mid \gamma$. Hence $\alpha\beta \mid \gamma$, so $\alpha\beta \equiv \gamma$. Comparing the norms gives $N(\alpha)N(\beta) = mn$, from which it follows that $N(\alpha) = m$ and $N(\beta) = n$. This means that $g(mn) \leq g(m)g(n)$. Together with the bound we already proved this gives $g(mn) = g(m)g(n)$.

Now g is completely determined by its value on prime powers. To calculate these values, we will first find out which prime numbers are Eisenstein primes as well. We will prove that a prime number p is an Eisenstein prime if and only if $p \equiv 2 \mod 3$.

First we consider the case $p \equiv 2 \mod 3$ and we assume $p = \alpha\beta$ for two non-unit Eisenstein integers α, β. Taking the norm gives $p^2 = N(\alpha)N(\beta)$, so $N(\alpha) = N(\beta) = p$. But if we write $\alpha = a + b\omega$, this gives $p = a^2 - ab + b^2$. If we consider this equation modulo 3, we see that $2 \equiv p = a^2 - ab + b^2 \equiv a^2 + 2ab + b^2 = (a+b)^2 \mod 3$, but this is a contradiction as 2 is not a quadratic residue modulo 3.

The case $p \equiv 0 \mod 3$ is easily handled, as this implies $p = 3$. We see that we can factor $3 = (2+\omega)(1-\omega)$. Note that $-\omega^2(1-\omega) = \omega^3 - \omega^2 = 2 + \omega$, hence both factors are the same up to multiplication with a unit.

Last, consider the case $p \equiv 1 \mod 3$. We will show that p is not prime by contradiction. Because $p \equiv 1 \mod 3$ we can show with quadratic reciprocity that -3 is a quadratic residue modulo p. Hence there is an integer s with $p \mid s^2 + 3$. But $s^2 + 3 = (s + 1 + 2\omega)(s - 1 - 2\omega)$, so if p is an Eisenstein prime, it must divide either $s + 1 + 2\omega$ or $s - 1 - 2\omega$. But if $p \mid s + 1 + 2\omega$, we must have $p \mid s + 1$ and $p \mid 2$, but $p \mid 2$ is impossible. In the same way $p \mid s - 1 - 2\omega$ gives a contradiction, hence p is not an Eisenstein prime. Hence we can write $p = \pi\pi'$ where π and π' are non-units. Because $p^2 = N(p) = N(\pi)N(\pi')$, we see that $N(\pi) = N(\pi') = p$. This means that $p = \pi\overline{\pi}$ and because $N(\pi)$ is prime, π must be an Eisenstein prime.

We will also need that $\pi \not\equiv \overline{\pi}$, so assume $\pi \equiv \overline{\pi}$. The argument of every unit is a multiple of $\pi/3$, so $\pi \equiv \overline{\pi}$ implies that the argument of π is a multiple of $\pi/6$. If it is even a multiple of $\pi/3$, we see that $\pi \equiv n$ for some integer n, but this is impossible. Else we have $\pi \equiv (2+\omega)n$ for some integer n, but comparing the norms gives $p = 3n^2$, which is impossible for $p \equiv 2 \mod 3$. Hence $\pi \not\equiv \overline{\pi}$.

Now we will calculate the function g on prime powers. We start with $p \equiv 2 \mod 3$, and we want the number of solutions to $p^k = \alpha\overline{\alpha}$. Because p is an Eisenstein prime, α can only contain prime factors p. Hence $\alpha \equiv p^{k/2} \equiv \overline{\alpha}$, this gives exactly one possibility if k is even, and otherwise there are no solutions. Hence $g(p^k) = 1$ for even k and $g(p^k) = 0$ for odd k.

If $p = 3$, we see that $\alpha\overline{\alpha} = 3^k \equiv (2+\omega)^{2k}$ with $2 + \omega$ prime. Hence the only solution is $\alpha \equiv (2+\omega)^k \equiv \overline{\alpha}$, which gives exactly one solution. Hence $g(p^k) = 1$.

If $p \equiv 1 \mod 3$, we have seen that $p = \pi\overline{\pi}$ for a prime π and that π and $\overline{\pi}$ are different primes. If we want to write $\alpha\overline{\alpha} = p^k = \pi^k\overline{\pi}^k$, then α must contain exactly k of the prime factors on the right. Hence $\alpha = \pi^j\overline{\pi}^{k-j}$, and this suffices as $\overline{\alpha} = \overline{\pi}^j\pi^{k-j}$. The number j must be one of the integers $0, 1, \ldots, k$, which gives $k+1$ options. Note that they are all different, because $\mathbb{Z}[\omega]$ has unique factorization. Hence $g(p^k) = k+1$. This concludes the proof. □

Theorem 4.4. *Let $\tau(n)$ be the number of positive divisors of n, and let it be 0 if n is not an integer. Then $t(m) = \tau(m) - \tau(m/2) + 2\tau(m/4)$ for all m.*

Proof. Suppose a lattice L of index m in $\mathbb{Z}[\omega]$ is generated by d and $a + \frac{m}{d}\omega$, for d a positive divisor of m, and a an integer from $\{0, 1, \ldots, d-1\}$. If it is invariant under conjugation, we must have $a + \frac{m}{d}\overline{\omega} \in L$, hence $2a + \frac{m}{d}(\omega + \overline{\omega}) = 2a - \frac{m}{d}$. As this is an integer, we have $d \mid 2a - \frac{m}{d}$. And it is easy to see that this relation implies that L is invariant under conjugation.

Hence we need to find the number of solutions to $d \mid 2a - \frac{m}{d}$. We will split this based on the parity of d and $\frac{m}{d}$. If d is odd, there is an inverse for 2, hence we can find a unique value of a. If d is even and $\frac{m}{d}$ is odd there are no a that satisfy the relation and if both d and $\frac{m}{d}$ are even, there are two values of a that satisfy the relation. Hence we find

$$\begin{aligned}
t(m) &= \sum_{d|m} \begin{cases} 1 & \text{if } d \text{ odd}, \\ 0 & \text{if } d \text{ even}, \frac{m}{d} \text{ odd}, \\ 2 & \text{if } d \text{ even}, \frac{m}{d} \text{ even}, \end{cases} \\
&= \sum_{d|m} 1 - \sum_{d|m} \begin{cases} 0 & \text{if } d \text{ odd}, \\ 1 & \text{if } d \text{ even}, \end{cases} + 2\sum_{d|m} \begin{cases} 1 & \text{if } d \text{ even}, \frac{m}{d} \text{ even}, \\ 0 & \text{else}. \end{cases} \\
&= \tau(m) - \tau(m/2) + 2\tau(m/4).
\end{aligned}$$ □

If m has at least two factors 2, we can even simplify this expression to $t(m) = \tau(m) + \tau(m/8)$. This follows from writing $m = 2^k n$ with n odd. We see that $\tau(2^\ell n) = (\ell+1)\tau(n)$ for all $\ell \geq -1$, hence if $k \geq 2$ we have

$$\begin{aligned}
t(m) &= \tau(m) - \tau(\tfrac{m}{2}) + 2\tau(\tfrac{m}{4}) \\
&= \tau(m) - k\tau(n) + 2(k-1)\tau(n) \\
&= \tau(m) + (k-2)\tau(n) \\
&= \tau(m) + \tau(\tfrac{m}{8}).
\end{aligned}$$

Now that we have derived the forms for $g(m)$ and $t(m)$, we have found an explicit formula for the number of tetrahedral triangulations. In this calculation we need the prime factorization of the number m, which can be obtained very efficiently if m is not too large.

Chapter 5

Generalizations

We have enumerated the tetrahedral triangulations, but this is only one special case of the non-negatively curved triangulations. You might wonder if the same ideas can be applied to other cases, but this is only the case for a few other cases. In our construction we used that all degrees in a tetrahedral are divisors of 6. So in any triangulation where this is the case as well, we could carry out a similar argument. This could be applied to the case with 3 singular vertices of degrees 1, 2 and 3, and the case with 3 singular vertices of degree 2.

A possible approach for other triangulations could be to replace the Eisenstein triangulation with a triangulation of the hyperbolic plane. In the hyperbolic plane there is a triangulation with equilateral triangles such that the degree is 60 at all points. In quite a similar way we could create a triangle map from this hyperbolic triangulation to a non-negatively cured triangulation, and this will always work, because all numbers from 1 to 6 divide 60. Again we can form a group of isometries of the hyperbolic plane such that the triangulation is a quotient of the hyperbolic triangulation by this group. However, this group will be much larger and finding the exact structure will be more difficult.

5.1 Projective triangulations

We will now deal with a different generalization, where we consider different spaces to triangulate. Ideally we would consider spaces that are 2-dimensional manifolds with positive Euler characteristic. Apart from the sphere there is exactly one such manifold that is compact, namely the projective plane with Euler characteristic 1. We will count the number of 'tetrahedral' triangulations for the projective plane.

Let $\mathbf{P}$ be a triangulation of the projective plane where two singular vertices have degree 3, and the other vertices have degree 6. We will call such triangulation a *projective tetrahedral triangulation*. Note that there is a 2-sheeted covering map from the sphere to the projective plane. This allows us to lift the triangulation $\mathbf{P}$ to a triangulation of the sphere. Because the covering is two-sheeted, we obtain 4 points of degree 3, hence a tetrahedral triangulation. We already have constructed an unfolding triangle map from the Eisenstein triangulation to this tetrahedral triangulation, and we can compose it with the covering map to obtain an unfolding triangle map $f : \mathbf{Eis} \to \mathbf{P}$. We will find out how the group G_f looks like in this case.

Theorem 5.1. *Let S bet the set of singular vertices in $\mathbf{P}$, and let $L = f^{-1}(S)$. Then G_f*

is generated by $\{H_x \mid x \in L\}$ *and a glide reflection.*

Proof. We have constructed f as composition of a map from **Eis** to a tetrahedral triangulation (with the covering map). We already know that this map is unchanged upon composition with an H_x for $x \in L$, because L is exactly the inverse image of the singular points. Hence $\{H_x \mid x \in L\}$ and $\{T_{2x} \mid x \in L\}$ are contained in G_f. Furthermore, we can use similar argument as in theorem 3.2 to show that G_f contains no other rotations, translations or reflections. However, if G_f contains no other elements than the given translation and half-turns, the triangulation $\mathbf{P}$ must be tetrahedral. Hence G_f will also contain a glide reflection, say a reflection in ℓ followed by a translation along v. Applying the glide reflection twice gives a translation over $2v$, hence $v \in L$. This means that T_v maps L to itself, then the reflection in ℓ also maps L to itself. In the same way as in theorem 3.2 we find that ℓ contains no point from L.

Given these facts we will show that L is a rectangular lattice. Take a point $x \in L$ and let x' be its reflection in ℓ. We may assume there are no points on the segment xx', otherwise we take as x the point on this segment closest to ℓ. This point cannot lie on ℓ, so $x \neq x'$. Then we can take $a_1 = x' - x$ as one of the generators of L. Now consider $L/(a_1\mathbb{Z})$ which is isomorphic to $\mathbb{Z}$, hence it is generated by an element $a_2 + a_1\mathbb{Z}$. We choose the representative a_2 such that $x + a_2$ has minimal distance to ℓ. Consider the reflection y of $x + a_2$ in ℓ. Then $y - (x + a_2)$ is parallel to xx', hence $y - (x + a_2)$ is a multiple of a_1. Then the distance from $x + a_2$ to ℓ is at least $|a_1|/2$. But if the distance is more than $|a_1|/2$, one of $x + a_2 + a_1$ and $x + a_2 - a_1$ has distance less than $|a_1|/2$ to ℓ, which contradicts that $x + a_2$ has minimal distance to ℓ. So $x + a_2$ has distance $|a_1|/2$ to ℓ. If $x + a_2$ is on the same side of ℓ as x, we see that $a_1 \perp a_2$. If $x + a_2$ is on the other side of ℓ, we see that $x + a_2 - a_1$ is on the same side of ℓ and also has distance $|a_1|/2$ to ℓ. Then $a_1 \perp a_2 - a_1$ and because a_1 and $a_2 - a_1$ also generate L, we still have a perpendicular set of generators. From now on we will assume a_2 is chosen such that it is parallel to ℓ and that a_1 and a_2 are the perpendicular generators.

We see as well that ℓ maps $\mathbb{Z}[\omega]$ to itself, so it must be parallel to one of the mirror axes in D_6. Then a_1 and a_2 are also parallel to mirror axes of D_6, hence L is fixed by two of the reflections in D_6.

Now we still have to find out how G_f looks like. For the given glide reflection, we see that $v \notin 2L$. If this were the case, composing the glide reflection with the translation over v (which are both in G_f) gives a reflection, which cannot be in G_f. So v is an odd multiple of a_2. It is now a simple exercise to see that the following glide reflections are also in G_f: reflect in a line through the point $(n + \frac{1}{2})a_i$ for an integer n, which is perpendicular to a_i, and compose it with a translation over $(2m + 1)a_{3-i}$ for an integer m. These are all the glide reflections which map L to itself such that the mirror axis doesn't pass through a point of L, hence they are the only glide reflections in G_f. □

We see again that G_f is completely determined by the lattice L, but now the lattice must be rectangular. Hence we need to count the number of rectangular lattices modulo the action of D_6. In this case it is somewhat simpler to count. Note that the number of triangles in the triangulation corresponding to L is now $2|\mathbb{Z}[\omega] : L|$.

Theorem 5.2. *The number of projective tetrahedral triangulations with* $2m$ *triangles is* $t(m) = \tau(m) - \tau(m/2) + 2\tau(m/4)$.

Proof. For each rectangular lattice, there are two reflections with perpendicular axes in D_6 which fix the lattice. If there is another reflection which fixes the lattice, we can compose the reflections to give a rotation (not a half-turn) that fixes the lattice. But a lattice is fixed by a rotation only if it is an ideal. However, the ideals are not rectangular, so there are only two reflections which fix the lattice. For the same reason there cannot be rotations that fix the lattice. So every lattice is fixed by exactly four elements of D_6: the identity, the half-turn and two perpendicular reflections.

Now for each of the three pairs of perpendicular reflections, we can find $t(m)$ different rectangular lattices with index m, which gives a total of $3t(m)$. But every lattice is in an orbit of size 3 (because the stabilizer has order 4), hence there are $t(m)$ projective tetrahedral triangulations with $2m$ triangles. □

We have seen that with not too much extra effort we can find the number of tetrahedral triangulations of the projective plane. Even for the torus or the Klein bottle we can count the number of triangulations having degree 6 in all vertices. Other connected, compact, 2-dimensional manifolds all have negative Euler characteristic, hence do not admit non-negatively curved triangulations.

Bibliography

[1] William P Thurston. “Shapes of polyhedra and triangulations of the sphere”. In: *Geometry & Topology Monographs* 1 (1998), pp. 511–549.

[2] Marte Koning. *Triangulating tetrahedra*. 2008.

[3] Greg Friedman. “Survey article: An elementary illustrated introduction to simplicial sets”. In: *Rocky Mountain Journal of Mathematics* 42.2 (2012), pp. 353–423. DOI: 10.1216/RMJ-2012-42-2-353.

[4] Margherita Barile; John Renze and Eric W Weisstein. *Principal Ideal Domain*. From MathWorld – A Wolfram Web Resource. URL: http://mathworld.wolfram.com/PrincipalIdealDomain.html.

[5] Frits Beukers. *Elementary Number Theory*. Department of Mathematics, UU, 2012.