

Delaunay triangulations of point sets in closed Euclidean d-manifolds

Manuel Caroli, Monique Teillaud

► **To cite this version:**

Manuel Caroli, Monique Teillaud. Delaunay triangulations of point sets in closed Euclidean d-manifolds. 27th Annual Symposium on Computational Geometry, Jun 2011, Paris, France. pp.274-282, 10.1145/1998196.1998236 . hal-01101094

HAL Id: hal-01101094

<https://hal.inria.fr/hal-01101094>

Submitted on 9 Jan 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Delaunay Triangulations of Point Sets in Closed Euclidean d -Manifolds*

Manuel Caroli[†]

Monique Teillaud[‡]

ABSTRACT

We give a definition of the Delaunay triangulation of a point set in a closed Euclidean d -manifold, i.e. a compact quotient space of the Euclidean space for a discrete group of isometries (a so-called Bieberbach group or crystallographic group). We describe a geometric criterion to check whether a partition of the manifold actually forms a triangulation (which subsumes that it is a simplicial complex). We provide an algorithm to compute the Delaunay triangulation of the manifold for a given set of input points, if it exists. Otherwise, the algorithm returns the Delaunay triangulation of a finitely-sheeted covering space of the manifold. The algorithm has optimal randomized worst-case time and space complexity.

Whereas there was prior work for the special case of the flat torus, as far as we know this is the first result for general closed Euclidean d -manifolds. This research is motivated by application fields, like computational biology for instance, showing a need to perform simulations in quotient spaces of the Euclidean space by more general groups of isometries than the groups generated by d independent translations.

*This work was partially supported by the ANR (*Agence Nationale de la Recherche*) under the “Triangles” project of the *Programme blanc* (No BLAN07-2_194137) <http://www.inria.fr/sophia/geometrica/collaborations/triangles/>.

[†]This work was done while the first author was working at INRIA

[‡]INRIA Sophia Antipolis - Méditerranée

1. INTRODUCTION

The Delaunay triangulation of a point set in \mathbb{E}^d is a well-studied structure in computational geometry. Efficient algorithms are known and there exist various implementations. We extend the well-known incremental algorithm that computes the Delaunay triangulation in \mathbb{E}^d [5] to the case of closed Euclidean d -manifolds. Such manifolds can be represented as quotient spaces of \mathbb{E}^d for a certain class of discrete groups of isometries, the so-called Bieberbach groups or crystallographic groups [29]. For a given closed Euclidean d -manifold, there are sets of points that do not define a Delaunay triangulation; we describe a geometric test that can be used to check this while running the incremental algorithm. In such cases, the algorithm actually computes the Delaunay triangulation of copies of the input points in a finitely-sheeted covering space of the manifold.

This paper is a generalization of [9], which discusses the case of the three-dimensional flat torus \mathbb{T}^3 , and which is accompanied by a CGAL software package [10]. The flat torus is the quotient space of \mathbb{E}^3 by a group of three independent translations. While this case fulfills the needs of many application fields, some of them, like computational biology [2] [21, Section 3.2], require more general manifolds that are quotient spaces of \mathbb{E}^3 by other crystallographic groups.

In the introductory Section 2, we recall basic notions (Section 2.1), summarize the work done on the flat torus (Section 2.2), and introduce closed Euclidean d -manifolds and their properties (Section 2.3). Section 3 studies Delaunay triangulations in closed Euclidean d -manifolds and shows, using the Bieberbach theorem, that there is always a finitely-sheeted covering space of the manifold, in which the Delaunay triangulation is defined for any set of points. Section 4 proposes an algorithm and analyzes it.

2. PRELIMINARIES

2.1 Simplicial Complexes and Triangulations

Let us briefly recall a few elementary definitions.

A k -simplex σ in \mathbb{E}^d ($k \leq d$) is the convex hull of $k + 1$ affinely independent points $\mathcal{P}_\sigma = \{p_0, p_1, \dots, p_k\}$. A simplex τ defined by $\mathcal{P}_\tau \subseteq \mathcal{P}_\sigma$ is a *face* of σ and has σ as a *coface*. This is denoted by $\sigma \geq \tau$ and $\tau \leq \sigma$.

In this paper, we consider triangulations with an infinite number of vertices, so, we use the following definition of a simplicial complex:

DEFINITION 1 ([20]). A simplicial complex is a set \mathcal{K} of simplices such that:

- (i). $\sigma \in \mathcal{K}, \tau \leq \sigma \Rightarrow \tau \in \mathcal{K}$
- (ii). $\sigma, \sigma' \in \mathcal{K} \Rightarrow \sigma \cap \sigma' \leq \sigma, \sigma'$
- (iii). Every point in a simplex of \mathcal{K} has a neighborhood that intersects at most finitely many simplices in \mathcal{K} (local finiteness).

Note that the previous definition is completely combinatorial; with an appropriate definition of a simplex, it remains valid in any topological space \mathbb{X} . We will propose a definition of a simplex adapted to the case of a closed Euclidean manifold in Section 3.

A *triangulation* of a topological space \mathbb{X} is a simplicial complex \mathcal{K} such that $|\mathcal{K}| = \bigcup_{\sigma \in \mathcal{K}} \sigma$ is homeomorphic to \mathbb{X} . A triangulation is defined by a point set \mathcal{P} if its set of vertices (0-simplices) is identical to \mathcal{P} .

A triangulation of a point set \mathcal{P} in \mathbb{E}^d is a *Delaunay triangulation* iff each simplex satisfies the *Delaunay property*, i.e. its circumscribing ball does not contain any point of \mathcal{P} in its interior [6, 11].

There are several reasons why we define a triangulation as a simplicial complex. Firstly, a triangulation is defined this way in the literature [1, 13, 17, 19, 20, 28, 34]. Moreover, designing a data structure to store tessellations that are non-simplicial complexes (e.g. Δ -complexes [18]) would be quite involved and less efficient in terms of both time and space complexity. Even more importantly, algorithms using a triangulation as input rely on the fact that the triangulation is a simplicial complex; this is the case for instance for meshing algorithms [24, 25], as well as algorithms to compute α -shapes, for which there are applications in the periodic case, too [22, 32]. Moreover, the incremental algorithm presented in Section 4 relies on the fact that the structure is a triangulation after each insertion.

2.2 Review of Triangulations in \mathbb{T}^3

Let \mathcal{P} be a finite point set in the three-dimensional Euclidean space \mathbb{E}^3 . The three-dimensional *flat torus* \mathbb{T}^3 is defined as the quotient space \mathbb{E}^3/\mathcal{G} , where $\mathcal{G} = (\mathbb{Z}^3, +)$ or equivalently the group generated by the three orthogonal unit translations of \mathbb{E}^3 . Intuitively, \mathbb{T}^3 is obtained by identifying all three pairs of opposite facets of a cube of \mathbb{E}^3 .

Let $\pi : \mathbb{E}^3 \rightarrow \mathbb{T}^3$ denote the quotient map and $DT(\mathcal{GP})$ denote the Delaunay triangulation of the infinite point set $\mathcal{GP} = \{p + z \mid p \in \mathcal{P}, z \in \mathbb{Z}^3\}$ in \mathbb{E}^3 .

DEFINITION 2 ([9]). *If $\pi(DT(\mathcal{GP}))$ is a simplicial complex in \mathbb{T}^3 , then we call it the Delaunay triangulation of $\pi(\mathcal{P})$ in \mathbb{T}^3 .*

There are point sets in \mathbb{T}^3 that do not define a Delaunay triangulation: For example, if \mathcal{P} consists of one point only, then $\pi(DT(\mathcal{GP}))$ is not a simplicial complex: the condition (ii) of Definition 1 is violated. Let us recall that the *1-skeleton* of a simplicial complex is the graph that consists of all edges and vertices.

THEOREM 1 ([9]). *If the 1-skeleton of $\pi(DT(\mathcal{GP}))$ does not contain cycles of length ≤ 2 , then $\pi(DT(\mathcal{GP}))$ is a triangulation of \mathbb{T}^3 .*

This yields a geometric criterion for $\pi(DT(\mathcal{GP}))$ to be a triangulation. Note that it also holds for supersets of \mathcal{P} , which will be useful for the algorithm later on.

COROLLARY 2 ([9]). *Let \mathbb{B} denote the largest 3-ball in \mathbb{E}^3 that does not contain points of \mathcal{GP} in its interior. If \mathbb{B} has diameter $< \frac{1}{2}$, then $\pi(DT(\mathcal{GP}'))$ is a triangulation in \mathbb{T}^3 for any finite $\mathcal{P}' \supseteq \mathcal{P}$.*

Consider the quotient space $\mathbb{T}_{27}^3 := \mathbb{E}^3/\mathcal{G}_{27}$ with $\mathcal{G}_{27} := ((3\mathbb{Z})^3, +)$. Then \mathbb{T}_{27}^3 is a 27-sheeted covering space of \mathbb{T}^3 (see e.g. [1] for an introduction on covering spaces). The following theorem ensures that, for any set of points, it is always possible to compute a Delaunay triangulation in \mathbb{T}_{27}^3 .

THEOREM 3 ([9]). *For any finite point set \mathcal{P} in \mathbb{E}^3 , the projection of the Delaunay triangulation of \mathcal{GP} in \mathbb{E}^3 onto \mathbb{T}_{27}^3 is a triangulation.*

Here, the infinite point set is still \mathcal{GP} whereas the group defining the quotient is \mathcal{G}_{27} .

Theorem 3 and Corollary 2 lead to a modified version of the incremental algorithm for computing Delaunay triangulations in \mathbb{E}^3 [5]:

- It starts with inserting 27 copies per input point, computing their Delaunay triangulation in \mathbb{T}_{27}^3 .
- Once the largest 3-ball not containing any vertex in its interior has diameter smaller than $\frac{1}{2}$, the algorithm switches to computing in \mathbb{T}^3 and inserts each of the remaining points only once.

While computing in \mathbb{T}_{27}^3 , 27 copies of points of \mathcal{P} are inserted one by one. So, in fact the following extended version of Theorem 3 is needed for the algorithm to work:

THEOREM 4 ([9]). *Theorem 3 still holds if we replace \mathcal{GP} by $\mathcal{GP} \cup \mathcal{G}_{27}Q$ for any $Q \subseteq \mathcal{G}$ with any $p \in \mathbb{E}^3$.*

As shown in [9], depending on the set of points \mathcal{P} , the algorithm computes either a triangulation of \mathbb{T}^3 if possible, or a triangulation of \mathbb{T}_{27}^3 , which is homeomorphic to \mathbb{T}^3 . In practice, data sets are likely to define a Delaunay triangulation of \mathbb{T}^3 . The algorithm has optimal randomized worst-case complexity.

2.3 Closed Euclidean Manifolds

This section is dedicated to introducing closed Euclidean manifolds, their properties, and how to construct them. Most concepts mentioned in this section are taken from [30].

A *closed manifold* is a compact manifold without boundary. A d -manifold is called *Euclidean* or *flat*, if every point has a neighborhood isometric to a neighborhood in \mathbb{E}^d .

We need some more notions: Let \mathcal{G} be a group and \mathcal{H} denote a subgroup of \mathcal{G} . \mathcal{H} is called *normal* in \mathcal{G} if it is invariant under conjugation, i.e., if for all $h \in \mathcal{H}$ and $g \in \mathcal{G}$, $ghg^{-1} \in \mathcal{H}$. For a group element $g \in \mathcal{G}$, the set $\{gh \mid h \in \mathcal{H}\}$ is called a *coset* of \mathcal{H} in \mathcal{G} . The *index* of a subgroup \mathcal{H} in \mathcal{G} is defined as the number of cosets of \mathcal{H} in \mathcal{G} .

A d -dimensional *Bieberbach group* \mathcal{G}_B is a discrete group of isometries of \mathbb{E}^d with compact quotient space $\mathbb{E}^d/\mathcal{G}_B$. Such groups are also called *crystallographic groups* or *space groups*.

THEOREM 5 (BIEBERBACH [3]).

- Let \mathcal{G}_B be a d -dimensional Bieberbach group. There is a group \mathcal{G}_T of d linearly independent translations that is a normal subgroup of \mathcal{G}_B of finite index. We call \mathcal{G}_T the translational subgroup of \mathcal{G}_B .
- For any d , there is only a finite number of d -dimensional Bieberbach groups, up to isomorphism.

Note that the quotient space $\mathbb{E}^d/\mathcal{G}_B$ is not necessarily a manifold: If \mathcal{G}_B leaves points fixed, these points do not have a neighborhood in $\mathbb{E}^d/\mathcal{G}_B$ that is homeomorphic to a neighborhood in \mathbb{E}^d . The quotient space $\mathbb{E}^d/\mathcal{G}_B$ can always be described by the more general concept of an orbifold [4, 31]. For the quotient space to be a manifold, the group must not have fixed points. In other words the group must be torsion-free, i.e., the identity must be the only element of finite order.

If \mathcal{G}_T is a subgroup of d independent translations of \mathcal{G}_B , then $\mathbb{E}^d/\mathcal{G}_T$ is a d -torus.

THEOREM 6 ([30]). Any closed Euclidean d -manifold corresponds up to diffeomorphism to exactly one quotient space $\mathbb{E}^d/\mathcal{G}_B$, where \mathcal{G}_B is a torsion-free d -dimensional Bieberbach group.

This means that it is sufficient to consider torsion-free Bieberbach groups to completely classify closed Euclidean manifolds.

According to Theorem 5, there are only finitely many d -dimensional Bieberbach groups, up to isomorphism. In dimension 2 there are 17, in dimension 3 there are 230.¹ In two dimensions, there are only two torsion-free Bieberbach groups and thus two closed Euclidean manifolds, up to isomorphism: the torus and the Klein bottle. In three dimensions, there are 10 closed Euclidean manifolds, four of which are non-orientable.

3. TRIANGULATIONS IN CLOSED EUCLIDEAN MANIFOLDS

The goal of this section is to generalize the results of [9] recalled in Section 2.2.

Let \mathcal{G}_F be a torsion-free d -dimensional Bieberbach group, \mathcal{P} a finite point set in \mathbb{E}^d , $\mathbb{X} := \mathbb{E}^d/\mathcal{G}_F$ a closed Euclidean manifold with quotient map $\pi : \mathbb{E}^d \rightarrow \mathbb{X}$, and $DT(\mathcal{G}_F\mathcal{P})$ the Delaunay triangulation of the infinite point set $\mathcal{G}_F\mathcal{P}$ in \mathbb{E}^d .

3.1 Definition

To be able to consider triangulations of \mathbb{X} , as defined in Section 2, we first give a definition for a simplex in such a manifold. A similar definition can be found in [33].

DEFINITION 3. Let \mathcal{P}_σ be a set of $k+1$ ($k \leq d$) points in \mathbb{E}^d . If the restriction $\pi|_{\text{Ch}(\mathcal{P}_\sigma)}$ of π to the convex hull $\text{Ch}(\mathcal{P}_\sigma)$ of \mathcal{P}_σ is injective, the image of $\text{Ch}(\mathcal{P}_\sigma)$ by π is called a k -simplex in \mathbb{X} .

In more intuitive terms, this definition requires simplices not to self-intersect in the quotient manifold. We can now adapt Definition 2 to the Delaunay triangulation of $\pi(\mathcal{P})$ in \mathbb{X} .

¹The number of Bieberbach groups by dimension is assigned the id A006227 in the On-Line Encyclopedia of Integer Sequences [27]. The number of torsion-free Bieberbach groups is assigned the id A059104.

DEFINITION 4. If $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a triangulation of \mathbb{X} (which subsumes that it is a simplicial complex in \mathbb{X}), then we call it the Delaunay triangulation of $\pi(\mathcal{P})$ in \mathbb{X} .

Note that there are actually cases in which $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is not a simplicial complex, as shown in Figure 1 for a torus $\mathbb{X} = \mathbb{E}^2/(\mathbb{Z}^2, +)$.




Figure 1: The intersection of σ and τ is $\{p\} \cup e$, which is not a simplex

For the discussions below we need the following two values:

1. The minimum distance $\delta(\mathcal{G})$ by which a group \mathcal{G} moves a point:

$$\delta(\mathcal{G}) = \min_{p \in \mathbb{E}^d, g \in \mathcal{G}, g \neq 1_{\mathcal{G}}} \text{dist}(p, gp),$$

where $1_{\mathcal{G}}$ denotes the unit element of \mathcal{G} . Note that if \mathcal{G} is torsion-free and discrete, then $\delta(\mathcal{G}) > 0$ holds.

2. The diameter $\Delta(\mathcal{S})$ of the largest d -ball \mathbb{B} in \mathbb{E}^d that does not contain any point of a set \mathcal{S} in its interior.

We now generalize Theorem 1:

THEOREM 7. If the 1-skeleton of $\pi(DT(\mathcal{G}_F\mathcal{P}))$ does not contain cycles of length ≤ 2 then $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a triangulation of \mathbb{X} .

Most parts of the proof of Theorem 1 are completely combinatorial and do not depend on the space, so they extend directly to \mathbb{X} and we postpone them to the appendix. We only prove the generalized version of Lemma 4.1 of [9]:

LEMMA 8. Let \mathcal{K} be a set of simplices in \mathbb{E}^d whose vertices are exactly the elements of $\mathcal{G}_F\mathcal{P}$, and that fulfills conditions (i) and (ii) of Definition 1, and the Delaunay property with respect to $\mathcal{G}_F\mathcal{P}$. Then \mathcal{K} satisfies the local finiteness property (iii) as well. Thus, \mathcal{K} is a simplicial complex.

PROOF. Let us assume that there is a vertex $v \in \mathcal{K}$ with an infinite number of incident simplices and thus an infinite number of incident edges in \mathcal{K} . Since \mathcal{P} contains only a finite number of points, there must be at least one point q in \mathcal{P} such that infinitely many points of the discrete point set \mathcal{G}_Fq are adjacent to v . Note that $\delta(\mathcal{G}_F) > 0$ and $\Delta(\mathcal{G}_Fq) < \infty$ hold because \mathcal{G}_F is a torsion-free Bieberbach group. Projecting all the edges from v to points in \mathcal{G}_Fq onto the unit d -sphere S centered in v yields an infinite point set \mathcal{P}_S . As S is bounded, \mathcal{P}_S must have an accumulation point. We choose q_1 and q_2 from \mathcal{G}_Fq such

that the distance between their projections onto S is smaller than ε for some $\varepsilon < \frac{\delta(\mathcal{G}_F)^3}{\Delta(\mathcal{G}_F q)^3}$ and $\varepsilon > 0$. Without loss of generality, we assume that $\text{dist}(v, q_2) \geq \text{dist}(v, q_1)$. We give a lower bound on the diameter D of the circumcircle of the triangle vq_1q_2 (see Fig. 2): D is given by the product of the three edge lengths divided by twice the triangle's area. The three edge lengths are each at least $\delta(\mathcal{G}_F)$. Also $\text{dist}(v, q_2) \leq \Delta(\mathcal{G}_F q)$ and the height of the triangle corresponding to the segment vq_2 is at most $\Delta(\mathcal{G}_F q)\varepsilon$. This yields $D \geq \frac{\delta(\mathcal{G}_F)^3}{\Delta(\mathcal{G}_F q)^2 \varepsilon} > \frac{\delta(\mathcal{G}_F)^3}{\Delta(\mathcal{G}_F q)^2 \frac{\delta(\mathcal{G}_F)^3}{\Delta(\mathcal{G}_F q)^3}} = \Delta(\mathcal{G}_F q)$. So, the d -ball

B_{vq_2} with v and q_2 on its boundary must have diameter larger than $\Delta(\mathcal{G}_F q)$ to not contain q_1 in its interior. As the largest empty d -ball has diameter $\Delta(\mathcal{G}_F q)$, B_{vq_2} cannot be empty, which is a contradiction to \mathcal{K} having the Delaunay property.




Figure 2: The gray area shows the possible positions of q_1

Let us now consider a point p in \mathbb{E}^d that is not a vertex in \mathcal{K} . Let σ denote the simplex that contains p in its interior and let v_σ denote a vertex of σ . Let $St(v_\sigma)$ denote the set of simplices that v_σ is incident to. Above we have shown that $St(v_\sigma)$ contains only finitely many elements. The set $St(\sigma)$ of simplices that σ is incident to is a subset of $St(v_\sigma)$, thus it is finite. There is a neighborhood $U(p)$ that has non-empty intersection with exactly the elements $St(\sigma)$. \square

Corollary 2 mentions the threshold $\frac{1}{2}$, which depends on the group \mathcal{G} . The generalized version of this corollary follows by simple geometric reasoning from Theorem 7.

COROLLARY 9. *If $\Delta(\mathcal{G}_F \mathcal{P}) < \frac{\delta(\mathcal{G}_F)}{2}$, then $\pi(DT(\mathcal{G}_F \mathcal{P}'))$ is a triangulation of \mathbb{X} for any finite $\mathcal{P}' \supseteq \mathcal{P}$.*

For any torsion-free Bieberbach group there are point sets such that the condition of Corollary 9 is fulfilled, because δ is strictly positive and Δ can be made arbitrarily small by the choice of the point set.

However, as mentioned above and illustrated by Figure 1, there are point sets that do not define a Delaunay triangulation. The following section shows that it is still possible to define a Delaunay triangulation in a finitely-sheeted covering space of the manifold.

3.2 Point sets that do not define a Delaunay triangulation of \mathbb{X}

This section focuses on giving a generalized version of Theorem 4.

THEOREM 10. *There is a normal subgroup \mathcal{G}_C of \mathcal{G}_F of finite index such that the projection of the Delaunay triangulation of $\mathcal{G}_F \mathcal{P} \cup \mathcal{G}_C Q$ in \mathbb{E}^d onto $\mathbb{X}_C = \mathbb{E}^d / \mathcal{G}_C$ is a triangulation for any finite point set \mathcal{P} in \mathbb{E}^d and any $Q \subseteq \mathcal{G}_F q$ with any $q \in \mathbb{E}^d$.*

PROOF. According to Theorem 5, there is a group \mathcal{G}_T of d linearly independent translations that is a normal subgroup of \mathcal{G}_F with finite index h' . We choose generators g_1, \dots, g_d of \mathcal{G}_T in the following way: Let g_1 be the shortest translation in \mathcal{G}_T . Let g_{i+1} be the shortest translation in \mathcal{G}_T that is linearly independent of the translations g_1, \dots, g_i . Note that $\Delta(\mathcal{G}_T p)$ does not depend on a specific choice of p and thus can be considered constant. We can find an integer coefficient c such that for each g_i the inequality $\text{dist}(p, g_i^c p) > 2\Delta(\mathcal{G}_T p)$ holds for any $p \in \mathbb{E}^d$. The group \mathcal{G}_C generated by g_1^c, \dots, g_d^c is a subgroup of \mathcal{G}_T of index c^d with the property $\delta(\mathcal{G}_C) > 2\Delta(\mathcal{G}_T p)$ for any $p \in \mathbb{E}^d$. As \mathcal{G}_T is normal in \mathcal{G}_F we have $gg_T g^{-1} \in \mathcal{G}_T$ for each $g \in \mathcal{G}_F, g_T \in \mathcal{G}_T$. By construction of \mathcal{G}_C there is a bijection between the $g_T \in \mathcal{G}_T$ and the $g_C \in \mathcal{G}_C$ given by $g_C = g_T^c$. Now it is easy to see that \mathcal{G}_C is a normal subgroup of \mathcal{G}_F with index $h = h' \cdot c^d$. Note that $\Delta(\mathcal{G}_C \mathcal{G}_F \mathcal{P}) = \Delta(\mathcal{G}_F \mathcal{P}) \leq \Delta(\mathcal{G}_T p)$ for any $p \in \mathbb{E}^d$. Thus $\Delta(\mathcal{G}_C \mathcal{G}_F \mathcal{P}) < \frac{\delta(\mathcal{G}_C)}{2}$ holds and according to Corollary 9 the projection of the Delaunay triangulation of $\mathcal{G}_C \mathcal{G}_F \mathcal{P} = \mathcal{G}_F \mathcal{P}$ onto \mathbb{X}_C forms a triangulation, which remains true even when adding further points. \square

Note that the proof is constructive, i.e. it describes how to construct \mathcal{G}_C from \mathcal{G}_T . The group \mathcal{G}_T can be constructed from \mathcal{G}_F , using for instance the Reidemeister-Schreier algorithm [26] implemented in GAP [16]. Theorem 10 means that there exists a space \mathbb{X}_C , in which the Delaunay triangulation of the point set $\pi(\mathcal{P})$ is defined. The space \mathbb{X}_C is a covering space of \mathbb{X} with a finite number of sheets [1]. Theorem 10 can also be understood by constructing \mathbb{X}_C from \mathbb{X} directly, as follows.

We first recall the notion of a *fundamental domain* (see [1] for a formal definition): Let \mathcal{G} be a discrete groups of isometries in \mathbb{E}^d with quotient map $\pi : \mathbb{E}^d \rightarrow \mathbb{E}^d / \mathcal{G}$. A fundamental domain of \mathcal{G} is a closed and convex subset $D_{\mathcal{G}}$ of \mathbb{E}^d such that

- $D_{\mathcal{G}}$ contains at least one point of the preimage by π of any point in $\mathbb{E}^d / \mathcal{G}$.
- If $D_{\mathcal{G}}$ contains more than one point of the same preimage, then all points of this preimage lie on the boundary of $D_{\mathcal{G}}$.

For instance the unit cube is a fundamental domain of \mathbb{T}^3 as defined in Section 2.2.

Each closed Euclidean d -manifold has a d -torus as covering space with a finite number of sheets. This follows from Theorem 5 as discussed above. A fundamental domain of the d -torus is a d -dimensional hyperparallelepiped. By gluing two of these hyperparallelepipeds together, we get a new covering space that is again a d -torus. We can construct \mathbb{X}_C by gluing as many copies of the fundamental domain as necessary to fulfill the condition in Corollary 9, i.e., $\Delta(\mathcal{G}_C \mathcal{G}_F \mathcal{P}) = \Delta(\mathcal{G}_F \mathcal{P}) < \frac{\delta(\mathcal{G}_C)}{2}$. See Figure 3 for an illustration in two dimensions.

As an example we consider the flat Klein bottle $\mathbb{E}^2 / \mathcal{G}_K$, where \mathcal{G}_K is the group generated by a translation g_t and a glide-reflection g_g , that is a reflection together with a translation parallel to the reflection axis (see Figure 4). The group generated by g_t and g_g^2 is a translational subgroup of \mathcal{G}_K of index 2. Now we can choose a subgroup of this translational subgroup with finite index that fulfills the condition of Theorem 10 as in Figure 3.




Figure 3: Sufficient number of copies of the fundamental domain




Figure 4: A part of the infinite point grid \mathcal{G}_{kp}

Note that in both Corollary 9 and Theorem 10 we deal with the condition of the form $\Delta < \frac{\delta}{2}$. Corollary 9 gives a condition on the point set \mathcal{P} for Δ to be small enough, whereas Theorem 10 gives a condition on the group \mathcal{G}_C for δ to be large enough.

4. ALGORITHM AND COMPLEXITY

The algorithm described at the end of Section 2.2 generalizes to any closed Euclidean manifold $\mathbb{X} = \mathbb{E}^d / \mathcal{G}_F$. We first give a general overview, then we detail a few aspects more precisely.

4.1 Overview

A preprocessing step consists in computing the h -sheeted covering space $\mathbb{X}_C = \mathbb{E}^d / \mathcal{G}_C$ as shown in Section 3.2. This preprocessing is independent from the point set, so, its complexity expressed as a function of the number of points is a constant depending only on \mathbb{X} .

Let now the set \mathcal{P} be given. A triangulation, initially empty, is computed incrementally by adding the points of \mathcal{P} one by one. The triangulation can be stored as a graph: d -simplices store a table of pointers to their vertices and a table of pointers to their adjacent d -simplices. Each vertex contains the coordinates of the point it corresponds to. Additionally, an external list \mathcal{L} stores a pointer for each d -simplex of the triangulation whose circumscribing ball is larger than the threshold $\delta(\mathcal{G}_F)$. Each such d -simplex of the triangulation stores a pointer to its corresponding element in \mathcal{L} .

- (1) The algorithm starts computing in \mathbb{X}_C , inserting h copies per input point.

The copies can be determined as follows. Let $\mathcal{G}_Q := \mathcal{G}_F / \mathcal{G}_C$ denote the quotient group of \mathcal{G}_F and \mathcal{G}_C . As \mathcal{G}_C is a subgroup of \mathcal{G}_F of index h , the group \mathcal{G}_Q has

h elements. For a given input point p , we insert all h points of its orbit under the action of \mathcal{G}_Q . \mathcal{L} is updated at each insertion.

- (2) Once the condition of Corollary 9 is met for the current point set, i.e. once \mathcal{L} is empty, the algorithm converts the triangulation of \mathbb{X}_C to a triangulation of \mathbb{X} , and inserts each of the remaining input points only once into the triangulation of \mathbb{X} .

To perform the conversion from \mathbb{X}_C to \mathbb{X} , the algorithm iterates over all d -simplices and all vertices to delete all periodic copies, keeping only one; it also updates the incidence relations of the d -simplices whose adjacent d -simplices have been deleted.

If \mathcal{P} is such that the condition of Corollary 9 is never fulfilled, then the algorithm never enters (2) and returns the triangulation of the covering space \mathbb{X}_C .

Note however that in this case, it is still possible to check at the end of the computation, whether the 1-skeleton of $\pi(DT(\mathcal{G}_P))$ contains no cycle of length ≤ 2 , and if so, to convert the triangulation of \mathbb{X}_C to a triangulation of \mathbb{X} as above. However, it will be impossible to insert further points in it.

4.2 Point insertion

As usual in an incremental algorithm, the insertion of each point consists of two steps: a geometric step (locating the point, then identifying all simplices that disappear after the point is inserted) and a combinatorial step (creating new simplices).

Geometric step.

For efficiency reasons, this step uses a point location data structure, the Delaunay hierarchy, originally designed for efficient computation of Delaunay triangulations of \mathbb{E}^d . We refer the reader to the original paper for a complete description of this data structure [12]. In a nutshell, it is built incrementally and has several levels: The intermediate levels store the Delaunay triangulations of an increasing sequence of subsets of the set of input points, while the last level stores the complete triangulation. There are pointers between vertices in different levels corresponding to the same input point. The Delaunay hierarchy can be adapted to our algorithm. Each new level of the hierarchy stores a triangulation of \mathbb{X}_C when it is created, and this triangulation is converted when possible to a triangulation of \mathbb{X} . Note that, if a given level l stores a triangulation of \mathbb{X} , then the next level $l+1$ is also in \mathbb{X} , since it contains more points and thus also stores a triangulation of \mathbb{X} by Corollary 9. However, some level l can store a triangulation of \mathbb{X}_C while the next level $l+1$ is converted into \mathbb{X} . In this case, for all the vertices corresponding to periodic copies of a given input point p in that level l , their pointers to the level $l+1$ all lead to the same vertex corresponding to p in \mathbb{X} .

After locating the point using the Delaunay hierarchy, the simplices that will not satisfy the Delaunay property after insertion of the new point must be identified. The algorithm starts at the d -simplex containing the point and does a simple traversal using the adjacency relations to detect these simplices.

Each d -simplex of the triangulation stores the information on how to map it isometrically into \mathbb{E}^d , i.e. the appropriate

element of its preimage under the quotient map π (or the quotient map of \mathbb{X}_C , denoted as π_C). This is used to evaluate the two predicates that are necessary to perform this geometric step:

- testing the orientation of $d+1$ points (to locate the point), and
- computing whether a point lies inside or outside a d -ball circumscribing $d+1$ points (to check the Delaunay property).

Both predicates are evaluated in \mathbb{E}^d : For each d -simplex on which we need to evaluate a predicate, we take its preimage under π (or π_C) from the data structure and evaluate the predicate on this preimage. This works exactly the same way even if \mathbb{X} is non-orientable: The orientation of a preimage under π is computed using the orientation predicate in \mathbb{E}^d .

Combinatorial step.

This step consists in updating the triangulation: the simplices that do not satisfy the Delaunay property after the insertion are removed, and the resulting hole is then triangulated by simplices incident to the new point.

Note that this approach only works if the hole is homeomorphic to a d -ball. The hole is the union of all simplices of the star of the newly inserted point. Since the algorithm guaranties that the structure after the point insertion is a simplicial complex, then the hole is always homeomorphic to a ball (Lemma 16 in appendix and Theorem 7).

4.3 Complexity

The randomized analysis of [12] assumes the insertion of points of \mathcal{P} to be performed in a random order. The changes to this analysis, when computing in \mathbb{X}_C , are minor: The points are inserted in sets of constant size (the number of periodic copies), and these sets are inserted in random order. The randomized worst-case complexity of the algorithm remains equal to the one of the algorithm for computing the Delaunay triangulation of \mathbb{E}^d .

THEOREM 11. *The algorithm described in Section 4.1 has optimal randomized worst-case time and space complexity $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$, where n is the number of points in \mathcal{P} .*

PROOF. The vertex set of the Delaunay triangulation in level i is denoted by \mathcal{P}_i and the levels of the hierarchy are numbered from bottom to top, that is $\mathcal{P}_0 = \mathcal{P}$. Let $1/\alpha$ denote the probability that a point is in \mathcal{P}_{i+1} given that it is in \mathcal{P}_i . In the algorithm of Section 4.1, the points of \mathcal{P} can be inserted in random order but when computing in \mathbb{X}_C , a constant number of copies of each point are inserted consecutively. Let \mathcal{P}'_i be the set containing all copies of the points of \mathcal{P}_i that show up in \mathbb{X}_C .

The randomized worst-case analysis in [12] shows that, if the input points are inserted in a random order, the expected cost of the walk in level i is linear in the product of α and the average vertex degree in a triangulation, which is $O\left(n^{\lceil \frac{d-2}{2} \rceil}\right)$ in the worst case. The proof is based on the fact that the number of points in \mathcal{P}_i that are closer to a query point q than to any other point in \mathcal{P}_{i+1} is in $O(\alpha)$. This property extends to the case of computing in \mathbb{X}_C . The vertex set of the Delaunay triangulation in level i is \mathcal{P}'_i . As for each point of \mathcal{P}_i there is a constant number of copies in \mathcal{P}'_i , the above bound can be at most a constant times

bigger, which is still in $O(\alpha)$. Thus the cost of the walk in level i is still in $O\left(\alpha n^{\lceil \frac{d-2}{2} \rceil}\right)$. Summing up over all levels yields an expected $O\left(n^{\lceil \frac{d-2}{2} \rceil} + \log n\right)$ complexity for one point insertion.

We insert at most hn points into the Delaunay triangulation, where h is a constant depending on the group \mathcal{G}_F , so the asymptotic size of the hierarchy does not change with respect to [12].

Furthermore, the conversion from \mathbb{X}_C to \mathbb{X} in step (2) of the algorithm is linear in the size of the triangulation, and it is applied only once during the algorithm run, so, it does not increase the overall complexity.

The total number of elements that are inserted to and removed from \mathcal{L} is at most proportional to the total number of d -simplices that are created and destroyed during the algorithm. Also, the maximum size of \mathcal{L} is the maximum number of d -simplices in the triangulation. Hence, the maintenance of \mathcal{L} does not change the complexity of the algorithm. \square

The bound for the point insertion comes from the worst-case average vertex degree in a triangulation. Thus for triangulations where the vertex degrees stay small throughout the construction, the complexity is even better, since it is dominated by the point insertion. For instance, for random distributions the expected complexity is $O(n \log n)$ [12].

5. CONCLUSION

We generalized the work of [9] to any closed Euclidean d -manifold.

A natural question is how to extend the results to more general orbit spaces. The results of Section 3 exclude Bieberbach groups with fixed points. However, from the Bieberbach theorem we know that any Euclidean orbit space under a Bieberbach group is an orbifold that has a closed Euclidean manifold as a finitely-sheeted covering space, on which our approach works. So, while the approach cannot compute a triangulation in the orbifold, it can always compute a Delaunay triangulation in a covering space.

Another question is whether Delaunay triangulations can be computed in hyperbolic manifolds that are quotients of hyperbolic spaces \mathbb{H}^d by a Fuchsian group. The question is natural and exciting from a mathematical point of view, since these groups are much richer than the crystallographic groups. Preliminary attempts show the problem to be quite challenging even in the simplest case of a group of four hyperbolic translations in \mathbb{H}^2 : Determining how many copies of the point set must be used is still open [7, Section 4.4]. This case would already find applications in fields as diverse as computer graphics [23] and neuromathematics [8, 15].

Acknowledgments

The authors wish to thank Jean-Marc Schlenker for helpful discussions and Günter Rote for his comments on a preliminary version of this paper.

6. REFERENCES

- [1] M. A. Armstrong. *Basic Topology*. Springer-verlag, 1982.
- [2] Julie Bernauer. Computational structural biology: Periodic triangulations for molecular dynamics, 2009. Talk at the workshop 'Subdivide and tile', <http://www.lorentzcenter.nl/1c/web/2009/357/info.php3?wsid=357>.
- [3] Ludwig Bieberbach. Über die Bewegungsgruppen des n -dimensionalen euklidischen Raumes mit einem endlichen Fundamentalbereich. *Göttinger Nachrichten*, 75, 1910.
- [4] Michel Boileau, Sylvain Maillot, and Joan Porti. *Three-dimensional orbifolds and their geometric structures*. Société mathématique de France, Paris, 2003.
- [5] Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24:162–166, 1981.
- [6] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.
- [7] Manuel Caroli. *Triangulating Point Sets in Orbit Spaces*. Thèse de doctorat en sciences, Université de Nice-Sophia Antipolis, France, 2010. <http://tel.archives-ouvertes.fr/tel-00552215/>.
- [8] Pascal Chossat and Olivier Faugeras. Hyperbolic planforms in relation to visual edges and textures perception, 2009. <http://arxiv.org/abs/0907.0963>.
- [9] Manuel Caroli and Monique Teillaud. Computing 3D periodic triangulations. In *Proceedings 17th European Symposium on Algorithms*, volume 5757 of *Lecture Notes in Computer Science*, pages 37–48, 2009. <http://hal.inria.fr/inria-00356871/>.
- [10] Manuel Caroli and Monique Teillaud. 3D periodic triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.7 edition, 2010. <http://www.cgal.org/>.
- [11] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.
- [12] Olivier Devillers. The Delaunay hierarchy. *Internat. J. Found. Comput. Sci.*, 13:163–180, 2002.
- [13] Robert J. Daverman and Richard B. Sher, editors. *Handbook of Geometric Topology*. Elsevier, Amsterdam, London, Paris, 2002.
- [14] Olivier Devillers and Monique Teillaud. Perturbations and vertex removal in a 3D Delaunay triangulation. In *Proceedings of the fourteenth ACM-SIAM Symposium on Discrete Algorithms*, pages 313–319, 2003. <http://hal.inria.fr/inria-00166710/>.
- [15] Grégory Faye, Pascal Chossat and Olivier Faugeras. Some theoretical results for a class of neural mass equations, 2010. <http://arxiv.org/abs/1005.0510>.
- [16] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.12*, 2008.
- [17] Ronald L. Graham, Martin Grötschel, and László Lovász, editors. *Handbook of Combinatorics*. Elsevier, Amsterdam, Lausanne, New York, 1995.
- [18] Allan Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [19] Michael Henle. *A Combinatorial Introduction to Topology*. Dover publication, New York, 1979.
- [20] John M. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, New York, 2000.
- [21] Erik Lindahl, David van der Spoel, Berk Hess *et al.* Gromacs, a versatile package to perform molecular dynamics, 2010. version 4.5.3, <http://www.gromacs.org/>.
- [22] Maarten Moesen. Periodicity and the design of bone scaffolds, 2008. Talk at the CGAL Prospective Workshop on Geometric Computing in Periodic Spaces, <http://www.cgal.org/Events/PeriodicSpacesWorkshop/>.
- [23] Guodong Ron, Miao Jin, and Xiaohu Guo. Hyperbolic centroidal Voronoi tessellation. In *ACM Symposium on Solid and Physical Modeling*, pages 117–126, 2010.
- [24] Laurent Rineau and Mariette Yvinec. Meshing 3D domains bounded by piecewise smooth surfaces. In *Meshing Roundtable conference proceedings*, pages 443–460, 2007.
- [25] Laurent Rineau and Mariette Yvinec. 3D surface mesh generation. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.7 edition, 2010.
- [26] Charles C. Sims. *Computing with Finitely Presented Groups*. Cambridge University Press, Cambridge, UK, 1994.
- [27] N. J. A. Sloane. The on-line encyclopedia of integer sequences. <http://www.research.att.com/~njas/sequences/>.
- [28] Edwin H. Spanier. *Algebraic Topology*. Springer-Verlag, New York, 1966.
- [29] William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, Princeton, New Jersey, 1997.
- [30] William P. Thurston. *Three Dimensional Geometry and Topology, vol. I*. Princeton University Press, Princeton, New Jersey, 1997.
- [31] William P. Thurston. *The Geometry and Topology of Three-Manifolds*. 2002. <http://www.msri.org/publications/books/gt3m/>.
- [32] Rien van de Weijgaert, Gert Vegter, Erwin Platen, Bob Eldering, and Nico Kruithof. Alpha shape topology of the cosmic web, 2010. <http://arxiv.org/abs/1006.2765>.
- [33] P. M. H. Wilson. *Curved Spaces*. Cambridge University Press, Cambridge, 2008.
- [34] Afra Zomorodian. *Topology for Computing*. Cambridge University Press, Cambridge, 2005.

Appendix: Proof of Theorem 7

As announced in the main part of the paper, the proof follows the proof of Theorem 1 given in [9] very closely.

We call *original domain* a subset of a fundamental domain that contains exactly one element of the preimage by π of any point in \mathbb{E}^d/\mathcal{G} . Let \mathcal{D} be an original domain.

At first, we show that Definition 4 actually makes sense: We verify that the simplices “match” under π , i.e. that all copies of a simplex in $DT(\mathcal{G}_F\mathcal{P})$ are mapped onto the same simplex in \mathbb{X} under π . Then we can show that $|\pi(DT(\mathcal{G}_F\mathcal{P}))|$ is homeomorphic to \mathbb{X} . We also prove that if $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a set of simplices, then it fulfills conditions (i) and (iii). Finally, we discuss under which circumstances condition (ii) is fulfilled, which yields the sufficient condition on $\pi(DT(\mathcal{G}_F\mathcal{P}))$ to be a triangulation. (Here, as well as in the sequel, (i), (ii), and (iii) always refer to the conditions in Definition 1.)

Let us start with the first lemma:

LEMMA 12. *If the restriction of π to any simplex in $DT(\mathcal{G}_F\mathcal{P})$ is injective, then $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a set of internally disjoint simplices in \mathbb{X} that do not contain any point of $\pi(\mathcal{G}_F\mathcal{P})$ in their interior.*

PROOF. Consider a d -simplex σ of $DT(\mathcal{G}_F\mathcal{P})$, whose vertices are a $(d+1)$ -tuple of points $\mathcal{P}_\sigma \subset \mathcal{G}_F\mathcal{P}$. σ satisfies the Delaunay property, so all copies $\mathcal{G}_F\mathcal{P}_\sigma$ also have an empty circumscribing ball. This shows that all these copies form d -simplices of $DT(\mathcal{G}_F\mathcal{P})$. This can be ensured even in degenerate cases: If we handle degeneracies as in [14], then the Delaunay triangulation of a set of cospherical points only depends on some intrinsic ordering between them. By choosing an ordering that is compatible with \mathcal{G}_F , all copies of that point set are triangulated in the same way.

Followingly, π collapses precisely all the copies of σ onto its equivalence class in \mathbb{X} . As any lower-dimensional simplex in $DT(\mathcal{G}_F\mathcal{P})$ is incident to some d -simplex, and thus is defined by a subset of its vertices, the same holds for simplices of any dimension.

Now the projections under π of two internally disjoint k -dimensional simplices σ and τ in $DT(\mathcal{G}_F\mathcal{P})$ are either equal or internally disjoint for $k \geq 1$, due to the bijectivity of π between both simplices and their respective images. The same argument implies that the interior of a simplex cannot contain any vertex. \square

We observe that $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is finite: $DT(\mathcal{G}_F\mathcal{P})$ is locally finite (Lemma 8), i.e. the star of any vertex is finite. As \mathcal{P} is finite, $\mathcal{G}_F\mathcal{P}$ is discrete and all d -simplices have a certain volume larger than some constant. Followingly, there are only finitely many d -simplices necessary to fill the original domain \mathcal{D} and thus \mathbb{X} . Finitely many d -simplices have only finitely many faces so the overall number of simplices in $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is finite as well.

Since $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a triangulation by definition, to prove that it is a triangulation of \mathbb{X} , it only remains to show:

LEMMA 13. *$|\pi(DT(\mathcal{G}_F\mathcal{P}))|$ is homeomorphic to \mathbb{X} .*

PROOF. By its construction $|DT(\mathcal{G}_F\mathcal{P})| = \mathbb{E}^d$ and π is surjective. Followingly, $\pi(|DT(\mathcal{G}_F\mathcal{P})|)$ is equal to \mathbb{X} . Then,

the chain of equalities

$$\begin{aligned} \pi(|DT(\mathcal{G}_F\mathcal{P})|) &= \pi\left(\bigcup_{\sigma \in DT(\mathcal{G}_F\mathcal{P})} \sigma\right) \\ &\stackrel{(1)}{=} \pi\left(\bigcup_{\tau \in \pi(DT(\mathcal{G}_F\mathcal{P}))} \pi^{-1}(\tau)\right) \\ &\stackrel{(2)}{=} \bigcup_{\tau \in \pi(DT(\mathcal{G}_F\mathcal{P}))} \tau \\ &= \bigcup_{\sigma \in DT(\mathcal{G}_F\mathcal{P})} \pi(\sigma) \\ &= |\pi(DT(\mathcal{G}_F\mathcal{P}))| \end{aligned}$$

holds with the following arguments:

- (1) This step just regroups the order of the simplices but does not change the set (cf. Lemma 12).
- (2) There is only a finite number of elements in $\pi(DT(\mathcal{G}_F\mathcal{P}))$.

\square

So far we know that if all simplices in $DT(\mathcal{G}_F\mathcal{P})$ are mapped as simplices onto \mathbb{X} , then the whole triangulation is mapped onto a set of simplices in \mathbb{X} . We now consider the incidence relation.

OBSERVATION 14. *Assume that the restriction of π to any simplex in $DT(\mathcal{G}_F\mathcal{P})$ is injective. If τ is a simplex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$ and $\tau' \leq \tau$, then τ' is a simplex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$. This follows immediately from the fact that incidence relations are maintained by π and from Lemma 12.*

It only remains to show condition (ii), i.e. the intersection of two simplices σ and τ in $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is another simplex χ that is incident to both σ and τ .

LEMMA 15. *Assume that the restriction of π to any simplex in $DT(\mathcal{G}_F\mathcal{P})$ is injective. Let $\sigma, \tau \in \pi(DT(\mathcal{G}_F\mathcal{P}))$ be any two simplices in \mathbb{X} , then $\sigma \cap \tau$ is a set of simplices in $\pi(DT(\mathcal{G}_F\mathcal{P}))$.*

PROOF. Without loss of generality, we assume that $\sigma \cap \tau \neq \emptyset$. We show that $\sigma \cap \tau = \bigcup_{p \in \sigma \cap \tau} \chi_p$, where χ_p is a simplex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$. The union is finite because there are only finitely many simplices in $\pi(DT(\mathcal{G}_F\mathcal{P}))$. Consider a point $p \in \sigma \cap \tau$. If p is a vertex of $\pi(DT(\mathcal{G}_F\mathcal{P}))$, then it is not contained in the interior of any other simplex, according to Lemma 12, and we set $\chi_p = \{p\}$. If p is not a vertex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$, then $p \in \overset{\circ}{\sigma}'$ and $p \in \overset{\circ}{\tau}'$ for some proper faces $\sigma' \leq \sigma$ and $\tau' \leq \tau$ because σ and τ are internally disjoint (Lemma 12). Since σ' and τ' are again either internally disjoint or identical, it follows that they are the same face and we set $\chi_p := \sigma' = \tau'$. By condition (i) the simplex χ_p is contained in $\pi(DT(\mathcal{G}_F\mathcal{P}))$. \square

Remember that $|\text{St}(v)|$ denotes the union of the simplices in the star of v . We can now formulate the following sufficient condition:

LEMMA 16. *If for all vertices v of $DT(\mathcal{G}_F\mathcal{P})$ the restriction of the quotient map $\pi|_{|\text{St}(v)|}$ is injective, then $\pi(DT(\mathcal{G}_F\mathcal{P}))$ forms a simplicial complex.*

PROOF. We set $\mathcal{K} = \pi(DT(\mathcal{G}_F\mathcal{P}))$. Let σ be a simplex of $DT(\mathcal{G}_F\mathcal{P})$ and v an incident vertex. Then $\sigma \subseteq |\text{St}(v)|$, thus the restriction of $\pi|_{|\text{St}(v)|}$ to σ is injective as well, and \mathcal{K} is a set of simplices (Lemma 12).

Conditions (i) and (iii) follow from the above discussion. It remains to show condition (ii): Consider two simplices $\sigma, \tau \in \mathcal{K}$ with $\sigma \cap \tau \neq \emptyset$. By definition of a simplex, there exist sets $\mathcal{P}_\sigma, \mathcal{P}_\tau$ in $\mathcal{G}_F\mathcal{D}$ such that $\sigma = \pi(\text{Ch}(\mathcal{P}_\sigma))$ and $\tau = \pi(\text{Ch}(\mathcal{P}_\tau))$. From Lemma 15, we know that $\sigma \cap \tau$ is a set of simplices in \mathcal{K} . So there exists a vertex $v \in \sigma \cap \tau$ and $\sigma, \tau \in \text{St}(v)$. By assumption $\pi|_{|\text{St}(v)|}$ is injective, so π is injective on σ and τ , and $\sigma \cap \tau = \pi(\text{Ch}(\mathcal{P}_\sigma)) \cap \pi(\text{Ch}(\mathcal{P}_\tau)) = \pi(\text{Ch}(\mathcal{P}_\sigma \cap \mathcal{P}_\tau))$. Also, the restriction of $\pi|_{|\text{St}(v)|}$ to $\text{Ch}(\mathcal{P}_\sigma \cap \mathcal{P}_\tau)$ is injective. So from Definition 3, it follows that $\sigma \cap \tau$ is a simplex. Since $\sigma \cap \tau \subseteq \sigma, \tau$, we have $\sigma \cap \tau \leq \sigma, \tau$. \square

We can now prove Theorem 7. We set $\mathcal{K} = \pi(DT(\mathcal{G}_F\mathcal{P}))$. From Lemma 12 and Observation 14, we know that \mathcal{K} is a finite set of simplices that fulfills conditions (i) and (iii).

Assume that \mathcal{K} is not a simplicial complex. From Lemma 16 there is a vertex $v \in \mathcal{K}$ for which $\pi|_{|\text{St}(v)|}$ is not injective. As π is continuous by definition, this implies the existence of two different points $p, q \in |\text{St}(v)|$ with $\pi(p) = \pi(q)$. Let σ denote the simplex of \mathcal{K} that contains $\pi(p) = \pi(q)$ in its interior. Then there are two different simplices $\sigma'_\mathbb{E} \in \pi^{-1}(\sigma)$ and $\sigma''_\mathbb{E} \in \pi^{-1}(\sigma)$ containing p and q , respectively. Thus $\sigma'_\mathbb{E}$ and $\sigma''_\mathbb{E}$ are both elements of $\overline{\text{St}(v)}$ (the closure $\overline{\mathcal{L}}$ of a subset $\mathcal{L} \subseteq \mathcal{K}$ is the smallest subcomplex of \mathcal{K} containing \mathcal{L}). Let u, w be vertices different from v with $u \leq \sigma'_\mathbb{E}$ and $w \leq \sigma''_\mathbb{E}$. The vertices u, w are also elements of $\overline{\text{St}(v)}$ and thus there are edges (u, v) and (v, w) in $DT(\mathcal{G}_F\mathcal{P})$. From $\pi(\sigma'_\mathbb{E}) = \pi(\sigma''_\mathbb{E})$ follows that $\pi(u) = \pi(w)$, and so the projection of (u, v) and (v, w) under π forms a cycle of length two in \mathbb{X} , which contradicts the assumption that $\pi|_{|\text{St}(v)|}$ is injective. So \mathcal{K} must be a simplicial complex that is homeomorphic to \mathbb{X} by Lemma 13, which means that $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a triangulation of \mathbb{X} .