



HAL
open science

Delaunay triangulations of closed Euclidean d -orbifolds

Manuel Caroli, Monique Teillaud

► **To cite this version:**

Manuel Caroli, Monique Teillaud. Delaunay triangulations of closed Euclidean d -orbifolds. *Discrete and Computational Geometry*, 2016, 55 (4), pp.827–853. 10.1007/s00454-016-9782-6 . hal-01294409

HAL Id: hal-01294409

<https://inria.hal.science/hal-01294409>

Submitted on 29 Mar 2016

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 closed Euclidean d -orbifolds *

Manuel Caroli[†]

Monique Teillaud[‡]

March 21, 2016

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 incremental algorithm to compute the Delaunay triangulation of the manifold defined by a given set of input points, if it exists. Otherwise, the algorithm returns the Delaunay triangulation of a finite-sheeted covering space of the manifold. The algorithm has optimal randomized worst-case time and space complexity. It extends to closed Euclidean orbifolds. An implementation for the special case of the 3D flat torus has been released in CGAL 3.5. To the best of our knowledge, this is the first general result on this topic.

Keywords Delaunay triangulation, orbit space, crystallographic groups, closed Euclidean manifolds, orbifolds, covering space, incremental algorithm, implementation.

*Preliminary versions of this paper were presented in conferences: the first one gave some results in the case of the 3D flat torus [CT09], and the second generalized the results to closed Euclidean d -manifolds [CT11]. Most of the work was done while the authors were working at INRIA Sophia Antipolis – Méditerranée. 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-sop.inria.fr/geometrica/collaborations/triangles/>.

[†]manuel.caroli@gmail.com.

[‡]**Corresponding author.** Monique.Teillaud@inria.fr. INRIA Nancy – Grand Est, LORIA, 615 rue du Jardin Botanique, 54600 Villers-lès-Nancy, France

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 [BY98, dBvKOS00] and there exist several implementations. We rephrase the definition of the Delaunay triangulation to extend it to the case of closed Euclidean d -manifolds. We define a triangulation as a simplicial complex, which allows us to propose an extension of the well-known incremental algorithm [Bow81], which was designed for \mathbb{E}^d , to these manifolds. Closed Euclidean d -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 [Thu97]. For a given closed Euclidean d -manifold, there are sets of points that do not admit a Delaunay triangulation, as we define it; 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 finite-sheeted covering space of the manifold. This scheme extends to general closed Euclidean d -orbifolds.

So far we are not aware of any robust and efficient algorithm for computing Delaunay triangulations in such spaces. In the literature, proved algorithms are restricted to special cases, and they require the computation of the Delaunay triangulation of a certain number of copies of the input points in the Euclidean space: it was shown that there exists a finite number of copies that allows for computing the correct triangulation in a fundamental domain in 2D [MR97, GM01] and in 3D [DH97]. This obviously leads to an important slow-down. Indeed, in practice, data sets are likely to define a Delaunay triangulation; our geometric test allows the algorithm to avoid copying points whenever possible.

In the engineering community, Thompson proposed an implementation for computing what he calls a periodic Delaunay “*tessellation*” in the 3D flat torus, avoiding duplications of points [Tho02]. All cells in the tessellation are simplices but they do not necessarily form a simplicial complex, and not all tetrahedra satisfy the Delaunay property. More importantly, the algorithm heavily relies on the assumption that input points are well distributed.

Our research was originally motivated by the needs of astronomers who study the evolution of the large scale mass distribution in our universe by running dynamical simulations on periodic 3D data, i.e., on points in the three-dimensional flat torus. In fact there are numerous application fields that need robust software for geometric problems in periodic spaces, for instance biomedical computing [Wei08], solid-state chemistry [Rob06], physics of condensed matter [Duq], fluid dynamics [dFC03], astronomy [vdWPV⁺10], this list being far from exhaustive. For simulations, the input data is often not inherently periodic, but the use of periodic spaces is useful to avoid handling boundary conditions and to simulate an infinite working space. Nevertheless, some of these applications are working on truly periodic data, like material engineering for bone scaffolding [Moe08], which requires computations of meshes of periodic surfaces. A selection of applications

was presented at the “CGAL Prospective Workshop on Geometric Computing in Periodic Spaces”¹ and at the workshop “Subdivide and Tile: Triangulating spaces for understanding the world”.² While the case of the flat torus with cubic domain fulfills the needs of many application fields, some of them, like computational biology [Ber09] [LvdSH10, Section 3.2] or nano-structures, require more general manifolds that are quotient spaces of \mathbb{E}^3 by other crystallographic groups.

In the introductory Section 2, we first recall basic notions about Euclidean triangulations (Section 2.1), then we introduce closed Euclidean d -manifolds and orbifolds and their properties (Section 2.2). Section 3 proposes a definition of a triangulation of a Euclidean manifold and presents sufficient conditions ensuring that a given point set defines a Delaunay triangulation. In Section 4 we prove, using the Bieberbach theorem, that for a closed Euclidean d -orbifold, there is always a finite-sheeted covering orbifold in which the Delaunay triangulation is defined for any set of points. Section 5 proposes an algorithm and analyzes it. Section 6 elaborates on the case of the 3D flat torus, and presents experimental observations performed with our CGAL software package [cga, CT]. The software was demonstrated by a video [CT08] and has already been used by researchers in other research fields, eg. by cosmologists [Sou11, SPK11, vdWVE⁺11, vdWPJ⁺11, HvdWV⁺12].

2 Preliminaries

2.1 Euclidean Simplicial Complexes and Triangulations

Let us briefly recall a few elementary definitions. A k -*simplex* σ in the Euclidean space \mathbb{E}^d ($k \leq d$) is the convex hull $\text{Ch}(P_\sigma)$ 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 called a *face* of σ and has σ as a *coface*. This is denoted by $\sigma \geq \tau$ and $\tau \leq \sigma$. A *Euclidean simplicial complex* is a collection K of simplices in \mathbb{E}^d such that :

- (a). $\sigma \in K, \tau \leq \sigma \Rightarrow \tau \in K$
- (b). $\sigma, \sigma' \in K \Rightarrow \sigma \cap \sigma' = \emptyset$ or $\sigma \cap \sigma' \leq \sigma, \sigma'$
- (c). (local finiteness) Every point in a simplex of K has a neighborhood that intersects at most finitely many simplices in K .

We refer to [Lee00, Chapter 5] for more details. We denote as $\text{St}(v)$ the *star* of a vertex v of a simplicial complex K , i.e., the smallest simplicial subcomplex of K that contains all simplices incident to v [BY98, Section 11.1.2] (this would rather be referred to as the *closed star* in some other references).

¹<http://www.cgal.org/Events/PeriodicSpacesWorkshop/>

²<http://www.lorentzcenter.nl/lc/web/2009/357/info.php3?wsid=357>

The simplicial complex K forms a *triangulation* of the union $|K| = \bigcup_{\sigma \in K} \sigma$. A triangulation defined by 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 [BY98, Chapter 17] [dBvKOS00, Chapter 9].

If a point set $\mathcal{P} \subset \mathbb{E}^d$ shows degeneracies, then the Delaunay triangulation $DT(\mathcal{P})$ actually is not uniquely defined. More precisely, let Q_{cos} be a polytope defined by more than $d + 1$ points of \mathcal{P} , all lying on the same sphere S_{cos} , and such that S_{cos} does not enclose any other point of \mathcal{P} ; then each partition of Q_{cos} into simplices is a valid Delaunay triangulation of Q_{cos} . This sometimes leads to rather define the Delaunay *complex*, whose cells can be polytopes with any number of vertices. Triangulating such polytopes in any way yields a triangulation of \mathbb{E}^d in which each simplex satisfies the Delaunay property. Up to a misuse of language, unless specified otherwise, the term *the* Delaunay triangulation actually refers to any triangulation satisfying the Delaunay property.

Let us state a general relation between the Delaunay property and local finiteness for a simplicial complex in \mathbb{E}^d . Surprisingly, we have not found any similar result in the literature.

Proposition 2.1. *Let K be a set of simplices whose set of vertices \mathcal{P}_K is a discrete point set in \mathbb{E}^d that is not contained in any half-space of \mathbb{E}^d . If K fulfills conditions (a) and (b), and the Delaunay property with respect to \mathcal{P}_K , then K satisfies the local finiteness property (c) as well, i.e., K is a simplicial complex.*

Proof. The hypothesis that \mathcal{P}_K is not contained in any half-space, together with the Delaunay property, implies that all balls circumscribing d -simplices in K have a finite radius.

The half-space hypothesis means that no point of \mathcal{P}_K is extremal. So, for each point $p \in \mathcal{P}_K$, there exist at least $d + 1$ points $p_0, p_1, \dots, p_d \in \mathcal{P}_K$ forming a d -simplex that contains p . For each $i \in \{0, \dots, d\}$, let B_i be the closed ball circumscribing the d -simplex $(p_0, \dots, p_{i-1}, p, p_{i+1}, \dots, p_d)$ (taking $p_{d+1} = p_0$). Let now $q \in \mathcal{P}_K$ be a point lying outside the union of balls B_i . Elementary geometry shows that a ball whose boundary contains p and q must contain at least one of the points p_i , so pq cannot be an edge of K . This shows that $\text{St}(p)$ is contained in the union $B_0 \cup \dots \cup B_d$. This union is compact, and \mathcal{P}_K is discrete, so, only finitely many points of \mathcal{P}_K can be incident to p .

A point p in \mathbb{E}^d that is not an element of \mathcal{P}_K can only be in the interior of one d -simplex, or lie on a $(d - k)$ -face. In any case, it has a neighborhood that has non-empty intersection with a finite number of simplices in K . □

As a counter-example for a set of points contained in a half-space, take the discrete set $\{(n, n^2), n \in \mathbb{N}\}$ of points on a half-parabola in the plane \mathbb{E}^2 . The origin $(0, 0)$ is a vertex incident to infinitely many simplices in the Delaunay triangulation.

2.2 Closed Euclidean Manifolds and Orbifolds

This section is dedicated to introducing closed Euclidean manifolds and orbifolds, their properties, and how to construct them. Most of the material of this section is taken from [Thu97, Section 4.2].

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} .

If \mathcal{G} is a group of isometries of \mathbb{E}^d , then \mathcal{G} is said to *act* on \mathbb{E}^d by the binary operator $\mathcal{G} \times \mathbb{E}^d \rightarrow \mathbb{E}^d$, $(g, x) \mapsto g \cdot x$, which satisfies associativity $((g.h) \cdot x = g \cdot (h \cdot x), \forall g, h \in \mathcal{G}, x \in \mathbb{E}^d)$ and identity $(1_{\mathcal{G}} \cdot x = x, \forall x \in \mathbb{E}^d)$, where $1_{\mathcal{G}}$ denotes the unit element of \mathcal{G} . The *quotient space* \mathbb{E}^d/\mathcal{G} is then the set of all *orbits* of the group action, i.e., the set of equivalence classes for the equivalence relation $x \equiv y \Leftrightarrow \exists g \in \mathcal{G}, y = g \cdot x$.

We quickly survey some background on Bieberbach groups. A d -dimensional *Bieberbach group* \mathcal{G} is a discrete group of isometries of \mathbb{E}^d with compact quotient space \mathbb{E}^d/\mathcal{G} . Such groups are also called *crystallographic groups* or *space groups*. From the Bieberbach theorem, such a group contains a subgroup of finite index isomorphic to \mathbb{Z}^d , the so-called *translation subgroup*.

Note that the quotient space \mathbb{E}^d/\mathcal{G} is not necessarily a manifold: If elements of \mathcal{G} fix points, these points do not have a neighborhood in \mathbb{E}^d/\mathcal{G} that is homeomorphic to a neighborhood in \mathbb{E}^d . The quotient space \mathbb{E}^d/\mathcal{G} can always be described by the more general concept of an *orbifold* [BMP03][Thu02, Chapter 13]. For the quotient space to be a manifold, the group must act freely on \mathbb{E}^d , i.e., elements of \mathcal{G} must not fix points. From the Bieberbach theorem, classes of closed Euclidean d -manifolds up to diffeomorphism correspond to groups acting freely on \mathbb{E}^d and containing a translation subgroup; this gives a classification of closed Euclidean manifolds.

The Bieberbach theorem states that there are only finitely many d -dimensional Bieberbach groups, up to isomorphism [Bie10]. In dimension 2 there are 17, in dimension 3 there are 230.³ In two dimensions, there are only two Bieberbach groups acting freely and thus two closed Euclidean manifolds, up to isomorphism: the torus (quotient of \mathbb{E}^2 by a group generated by two independent translations) and the Klein bottle (quotient of \mathbb{E}^2 by a group generated by a translation and a glide-reflection whose direction is perpendicular to the translation). In three dimensions, there are 10 closed Euclidean manifolds, four of which are non-orientable.

From now on, \mathcal{G} denotes a Bieberbach group in \mathbb{E}^d with quotient map $\pi : \mathbb{E}^d \rightarrow \mathbb{E}^d/\mathcal{G}$.

³The number of Bieberbach groups by dimension is assigned the id A006227 in the On-Line Encyclopedia of Integer Sequences [Slo]. The number of *torsion-free* Bieberbach groups is assigned the id A059104.

A *fundamental domain* of \mathcal{G} is a closed and convex subset $D_{\mathcal{G}}$ of \mathbb{E}^d such that (see [Arm82, Section 4.4] for a formal definition):

- $D_{\mathcal{G}}$ 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 $D_{\mathcal{G}}$.

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} . We denote an original domain by \mathcal{D} .

As an illustration, let us consider the 2D torus $\mathbb{E}^2/(\mathbb{Z}^2, +)$ (resp. the 3D torus $\mathbb{E}^3/(\mathbb{Z}^3, +)$). Here, we identify the group generated by the d translations by the unit orthogonal basis vectors of \mathbb{E}^d to the additive group $(\mathbb{Z}^d, +)$. The closed unit square (resp. cube) is a fundamental domain for this torus, whereas the original domain is a half-open unit square (resp. cube).

For the discussions below we need the following two values:

- (1). The minimum distance $\delta(\mathcal{G})$ by which the action of a group \mathcal{G} moves a point of \mathbb{E}^d :

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

Note that if \mathcal{G} acts freely on \mathbb{E}^d , then $\delta(\mathcal{G}) > 0$.

- (2). The diameter $\Delta(\mathcal{S})$ of the largest d -ball in \mathbb{E}^d that does not contain any point of a set \mathcal{S} in its interior. The existence of the translation subgroup shows that $\Delta(\mathcal{G}p)$ is finite for each point $p \in \mathbb{E}^d$.

3 Triangulations of Closed Euclidean Manifolds

Let \mathcal{G}_F be a d -dimensional Bieberbach group acting freely on \mathbb{E}^d (the subscript in notation \mathcal{G}_F is to keep in mind that the action is free), \mathcal{P} a non-empty finite point set in \mathbb{E}^d , $\mathbb{X} = \mathbb{E}^d/\mathcal{G}_F$ a closed Euclidean manifold with quotient map $\pi : \mathbb{E}^d \rightarrow \mathbb{X}$.

3.1 Definitions

The notion of convex hull, on which the usual definition of a simplex in \mathbb{E}^d relies, is not defined in the manifold \mathbb{X} . Therefore, 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. If σ is a simplex with set of vertices \mathcal{P}_{σ} , we denote as $\dot{\sigma}$ the set $\sigma \setminus \mathcal{P}_{\sigma}$.

Definition 3.1. Let σ be a k -simplex ($k \leq d$) in \mathbb{E}^d . If the restriction $\pi|_{\dot{\sigma}}$ of π to $\dot{\sigma}$ is injective, the image of σ 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, except possibly at vertices. See Figure 1 for an illustration in the torus $\mathbb{E}^2/(\mathbb{Z}^2, +)$. Our definition is less restrictive than Wilson’s definition, which does not even allow the projection of vertices to coincide [Wil08, Section 3.2].

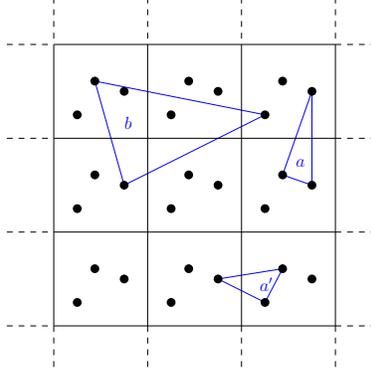


Figure 1: Case of a 2D flat torus. a and a' are simplices, while b is not a simplex.

With this definition of a simplex, we can adopt the same definition of a triangulation in a topological space as in a Euclidean space (Section 2.1):

Definition 3.2. 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). $\sigma, \sigma' \in \mathcal{K} \Rightarrow \sigma \cap \sigma' = \emptyset$ or $\sigma \cap \sigma' \leq \sigma, \sigma'$
- (iii). (local finiteness) Every point in a simplex of \mathcal{K} has a neighborhood that intersects at most finitely many simplices in \mathcal{K} .

A *triangulation* of \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 of \mathbb{X} is *defined by a point set* \mathcal{P} if its set of vertices (0-simplices) is identical to $\pi(\mathcal{P})$.

We refer to [Lee00, Chapter 5] for a complete formal presentation, using the notions of an abstract simplicial complex and its geometric realization.

The main reason why we choose to define a triangulation this way is that we are aiming at extending the use of the standard incremental algorithm [Bow81], which was proved to be very efficient in practice, to the case of Euclidean manifolds. As will be seen in Section 5 (see in particular the last paragraph of Section 5.2), that algorithm strongly relies on the fact that the structure is a simplicial complex after each insertion. So, though the definition has sometimes been relaxed, see e.g., [BS97, FSSB07], we stick to the definition of a triangulation as a simplicial complex, as often done in topology and combinatorics (see for example [BCdVN13], [EH10, Section III.1], [Zom05, Section 2.3.4], [Lee00, Chapter 5], [GGL95, Chapter 34],

[Arm82, Section 6.1], [Hen79, Section 20], [Spa66, Section 3.1]). Moreover this definition allows the computed triangulation to be used by other algorithms that assume that their input is a simplicial complex; this is the case for instance for meshing algorithms [RY07, ART⁺] as well as algorithms to compute α -shapes, for which there are applications in the torus $\mathbb{E}^3/(\mathbb{Z}^3, +)$ [Moe08, vdWPV⁺10].

In the sequel we will also need structures that are more general than triangulations. Let us reuse Zomorodian's term "simplicial set" [Zom10]:

Definition 3.3. A *simplicial set* is a set \mathcal{K} of simplices satisfying conditions (i) and (iii) of Definition 3.2.

The next sections will study the projection $\pi(DT(\mathcal{G}_F\mathcal{P}))$ onto \mathbb{X} of the Delaunay triangulation $DT(\mathcal{G}_F\mathcal{P})$ of the infinite point set $\mathcal{G}_F\mathcal{P}$ in \mathbb{E}^d .

3.2 Properties of the projection $\pi(DT(\mathcal{G}_F\mathcal{P}))$

Assume that the restriction of π to $\dot{\sigma}$ is injective for each σ of $DT(\mathcal{G}_F\mathcal{P})$, i.e., the simplices of $DT(\mathcal{G}_F\mathcal{P})$ project onto simplices on \mathbb{X} . From the fact that incidence relations are maintained by π follows that if τ is a simplex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$ and $\tau' \leq \tau$, then τ' is a simplex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$.

Let us first prove that the projection makes sense, i.e., if the simplices of $DT(\mathcal{G}_F\mathcal{P})$ project onto simplices, then they also "match" under π : all copies of a simplex in $DT(\mathcal{G}_F\mathcal{P})$ are mapped onto the same simplex in \mathbb{X} under π .

To achieve this even in degenerate cases, we choose a Delaunay triangulation in the following way. When more than $d+1$ points of $\mathcal{G}_F\mathcal{P}$ forming a polytope Q_{cos} are all lying on a sphere S_{cos} satisfying the Delaunay property, then all images $g \cdot Q_{cos}, g \in \mathcal{G}_F$, are also polytopes with cospherical vertices since \mathcal{G}_F is a group of isometries. Then we choose any triangulation $T(Q_{cos})$ of Q_{cos} , and triangulate each $g \cdot Q_{cos}$ by the image $g \cdot T(Q_{cos})$. We obtain simplices satisfying the Delaunay property (again, because \mathcal{G}_F is a group of isometries), and which are projecting onto a well-defined triangulation of $\pi(Q_{cos})$ in $\pi(DT(\mathcal{G}_F\mathcal{P}))$.

Lemma 3.4. *If the restriction of π to $\dot{\sigma}$ is injective for each simplex σ of $DT(\mathcal{G}_F\mathcal{P})$, 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.*

The term *internally disjoint* for two simplices means that their respective relative interiors are disjoint.

Proof. Consider a d -simplex σ of $DT(\mathcal{G}_F\mathcal{P})$, whose vertices are a $(d+1)$ -tuple of points $\mathcal{P}_\sigma \subset \mathcal{G}_F\mathcal{P}$. Simplex 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 $DT(\mathcal{G}_F\mathcal{P})$. With the choice made above, it is also true in degenerate cases.

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

Lemma 3.5. *The number of simplices in $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is finite.*

Proof. Since \mathcal{G}_F is a Bieberbach group, it is discrete and $\mathcal{G}_F\mathcal{P}$ is discrete for a finite set \mathcal{P} . In addition, by the Bieberbach theorem \mathcal{G}_F contains a subgroup of translations isomorphic to \mathbb{Z}^d , so no point of $\mathcal{G}_F\mathcal{P}$ is extremal and $DT(\mathcal{G}_F\mathcal{P})$ is locally finite by Proposition 2.1. Followingly, the star of any vertex is finite. All d -simplices have a certain volume larger than some constant, so there are only finitely many d -simplices necessary to fill the original domain \mathcal{D} (see end of Section 2, Page 6, for the definition), which is bounded by definition of a Bieberbach group. The projections of these d -simplices fill \mathbb{X} , and finitely many d -simplices have only finitely many faces. So, overall the size of $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is also finite. \square

Lemma 3.6. *$|\pi(DT(\mathcal{G}_F\mathcal{P}))|$ 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 3.4).
- (2) There are only a finite number of elements in $\pi(DT(\mathcal{G}_F\mathcal{P}))$. \square

The projection $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a simplicial set (Definition 3.3). Note that, even when all simplices of $DT(\mathcal{G}_F\mathcal{P})$ project onto simplices in \mathbb{X} , there are actually cases in which $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is not a simplicial complex: Figure 2 shows an example in the torus $\mathbb{E}^2/\mathbb{Z}^2$, where condition (ii) is not satisfied.

We can now propose a (partial) definition of the Delaunay triangulation of $\pi(\mathcal{P})$ in \mathbb{X} .

Definition 3.7. 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}* .

3.3 Combinatorial Condition

This section is devoted to proving Theorem 3.8, which gives the combinatorial condition ensuring that the Delaunay triangulation is defined, i.e., that the projection $\pi(DT(\mathcal{G}_F\mathcal{P}))$ actually is a simplicial complex.

The *1-skeleton* of a simplicial set is the graph that consists of all its edges and vertices.

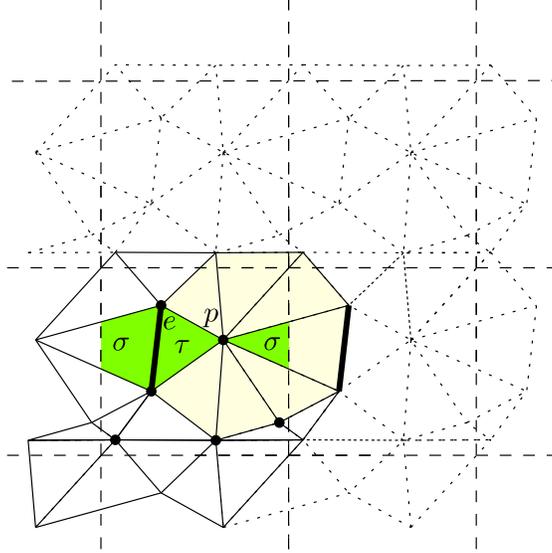


Figure 2: Case of a 2D flat torus. The intersection of σ and τ is $\{p\} \cup e$, which is not a simplex. Note that the projection of the periodic triangulation of \mathbb{E}^2 onto the quotient torus is not injective on the star of p , which contains two copies of edge e .

Theorem 3.8. *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} .*

The next two lemmas will allow us to finally prove this theorem at the end of the section.

Lemma 3.9. *Assume that the restriction of π to each simplex in $DT(\mathcal{G}_F\mathcal{P})$ is injective (i.e. $DT(\mathcal{G}_F\mathcal{P})$ does not contain cycles of length 1). 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 3.4, and we set $\chi_p = \{p\}$. If p is not a vertex in $\pi(DT(\mathcal{G}_F\mathcal{P}))$, then p lies in the relative interiors of σ' and τ' for some proper faces $\sigma' \leq \sigma$ and $\tau' \leq \tau$ because σ and τ are internally disjoint (Lemma 3.4). 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

We can now formulate the following sufficient condition:

Lemma 3.10. *If for each vertex v of $DT(\mathcal{G}_F\mathcal{P})$, the restriction $\pi|_{|\text{St}(v)|}$ of the quotient map 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 $\sigma \subset |\text{St}(v)|$, thus the restriction of $\pi|_{|\text{St}(v)|}$ to σ is injective as well, and \mathcal{K} is a set of simplices (Lemma 3.4).

Conditions (i) and (iii) follow from the above discussion. It remains to show condition (ii): Consider two simplices $\sigma, \tau \in \mathcal{K}$ with $\sigma \cap \tau \neq \emptyset$. From Lemma 3.9, we know that $\sigma \cap \tau$ is a set of simplices in \mathcal{K} . So there exists a vertex $p \in \sigma \cap \tau$. By definition of a simplex, there exist sets $\mathcal{P}_\sigma, \mathcal{P}_\tau$ in $\mathcal{G}_F\mathcal{P}$ such that $\sigma = \pi(\text{Ch}(\mathcal{P}_\sigma))$ and $\tau = \pi(\text{Ch}(\mathcal{P}_\tau))$; we choose \mathcal{P}_σ and \mathcal{P}_τ in such a way that they contain a common point v in the fiber of p . By assumption $\pi|_{|\text{St}(v)|}$ is injective, so π is injective on $\text{Ch}(\mathcal{P}_\sigma)$ and $\text{Ch}(\mathcal{P}_\tau)$, and $\sigma \cap \tau = \pi(\text{Ch}(\mathcal{P}_\sigma) \cap \text{Ch}(\mathcal{P}_\tau)) = \pi(\text{Ch}(\mathcal{P}_\sigma \cap \mathcal{P}_\tau))$. Also, the restriction of $\pi|_{|\text{St}(v)|}$ to $\text{Ch}(\mathcal{P}_\sigma \cap \mathcal{P}_\tau)$ is injective. So from Definition 3.1, 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

Proof of Theorem 3.8. We set $\mathcal{K} = \pi(DT(\mathcal{G}_F\mathcal{P}))$. From Lemma 3.4, 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 3.10 there is a vertex $v \in DT(\mathcal{G}_F\mathcal{P})$ for which $\pi|_{|\text{St}(v)|}$ is not injective. 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 $\text{St}(v)$. 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 $\text{St}(v)$ 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 3.6, which means that $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a triangulation of \mathbb{X} . \square

3.4 Geometric Criterion

The next result follows from Theorem 3.8 by simple geometric reasoning.

Criterion 3.11. *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 $\mathcal{P}' \supseteq \mathcal{P}$.*

The conclusion is stronger than that of Theorem 3.8. Indeed, adding a point in a Delaunay triangulation that is a simplicial complex can create a cycle of length two, as illustrated in the case of the flat torus $\mathbb{E}^2/(\mathbb{Z}^2, +)$ in Figure 3. This is the reason why this criterion is actually used in the incremental algorithm presented later. In addition, it is easy to check in practice.

Proof. The edges of a simplex are completely contained in its circumscribing ball and are thus bounded by the ball diameter. If the diameter of each circumscribing ball is smaller than $\frac{\delta(\mathcal{G}_F)}{2}$, then all edges in the triangulation are also shorter than this quantity. In order to create a cycle of length two, the sum of the lengths of the two edges needs to be at least $\delta(\mathcal{G}_F)$, which is not possible if both edges are shorter than half this length. From Theorem 3.8, it follows that $\pi(DT(\mathcal{G}_F\mathcal{P}))$ is a simplicial complex. If we add more points, the diameter of the largest empty ball cannot become larger. The claim follows. \square

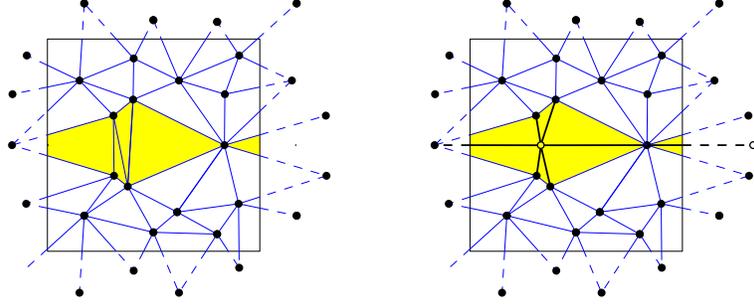


Figure 3: Adding a point in a simplicial complex can create a cycle of length two.

For any Bieberbach group acting freely on \mathbb{E}^d , there are point sets such that Criterion 3.11 is fulfilled, because δ is strictly positive and Δ can be made arbitrarily small by the choice of the point set. So, for any closed Euclidean manifold, there are point sets that define a Delaunay triangulation of the manifold. However, this is not true for general closed Euclidean orbifolds, which is illustrated by Figure 4; in the case when the center does not belong to \mathcal{P} (left), the input point that is closest to the singularity will always create a cycle of length 1 (even if ignoring the dashed diagonal); the right picture shows that when the center belongs to \mathcal{P} , then edges between the same closest point (and maybe others) and the center will form a cycle of length 2. So Definition 3.7 does not extend to orbifolds.

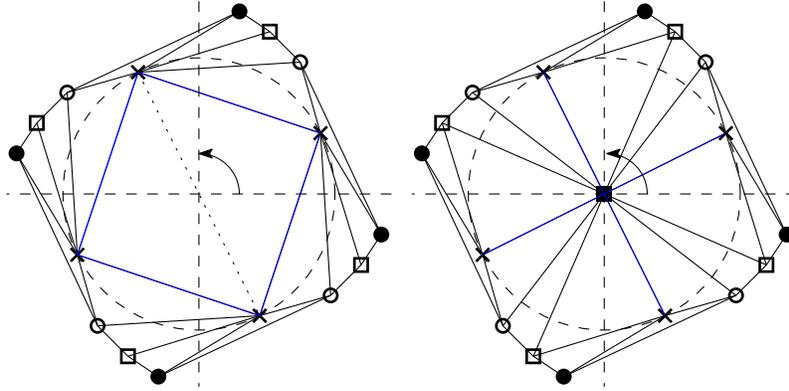


Figure 4: Delaunay triangulation of \mathcal{GP} shown locally around the center of a rotation in \mathcal{G} ; (left) the center does not belong to \mathcal{P} ; (right) the center belongs to \mathcal{P} .

In a given closed Euclidean manifold, there are point sets that do not define a Delaunay triangulation, as mentioned above. In such a case, the following section shows that it is still possible to define a Delaunay triangulation of a finite-sheeted covering space of the manifold; this is also true for a general closed Euclidean orbifolds with a covering orbifold.

4 What if the Delaunay Triangulation is not defined?

This section focuses on cases when the Delaunay triangulation as defined by Definition 3.7 does not exist, either because of a singularity (like in Figure 4), or because the set of input points does not satisfy good properties. We first recall a Bieberbach theorem already mentioned in Section 2.2; we will use it to prove Theorem 4.2.

Theorem 4.1 (Bieberbach [Bie10]). *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. \mathcal{G}_T is the translation subgroup of \mathcal{G}_B .*

The theorem is not restricted to groups acting freely, so all results in this section actually hold for any closed Euclidean orbifold $\mathbb{X} = \mathbb{E}^d/\mathcal{G}_B$, where \mathcal{G}_B is a Bieberbach group. \mathcal{G}_T is a subgroup of d independent translations of \mathcal{G}_B , so $\mathbb{E}^d/\mathcal{G}_T$ is a d -torus.

Theorem 4.2. *There is a normal subgroup \mathcal{G}_C of \mathcal{G}_B of finite index such that the projection of the Delaunay triangulation of $\mathcal{G}_B\mathcal{P} \cup \mathcal{G}_C Q$ in \mathbb{E}^d onto $\mathbb{X}_C = \mathbb{E}^d/\mathcal{G}_C$ is a triangulation for any non-empty finite point set \mathcal{P} in \mathbb{E}^d and any $Q \subseteq \mathcal{G}_B q$ with any $q \in \mathbb{E}^d$.*

As will be detailed after the proof, this theorem shows that there is always a space \mathbb{X}_C in which the Delaunay triangulation is defined. In fact it shows a stronger result: more points can be added to the point set in the initial orbifold, still keeping the Delaunay triangulation of \mathbb{X}_C well defined. This stronger property ensures the correctness of the incremental algorithm presented in Section 5.

In both Criterion 3.11 and Theorem 4.2 we deal with a condition of the form $\Delta(\text{some point set}) < \frac{\delta(\text{some group})}{2}$. Criterion 3.11 is a condition on the point set for Δ to be small enough, whereas Theorem 4.2 gives a condition on the group for δ to be large enough.

Proof. According to Theorem 4.1, there is a group \mathcal{G}_T of d linearly independent translations that is a normal subgroup of \mathcal{G}_B 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 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}_B we have $gg_T g^{-1} \in \mathcal{G}_T$ for each $g \in \mathcal{G}_B, g_T \in \mathcal{G}_T$. 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}_B with index $h = h' \cdot c^d$. Note that $\Delta(\mathcal{G}_C \mathcal{G}_B \mathcal{P}) = \Delta(\mathcal{G}_B \mathcal{P}) \leq \Delta(\mathcal{G}_T p)$ for any $p \in \mathbb{E}^d$. Thus $\Delta(\mathcal{G}_C \mathcal{G}_B \mathcal{P}) < \frac{\delta(\mathcal{G}_C)}{2}$ holds and according to Criterion 3.11 the projection of the Delaunay triangulation of $\mathcal{G}_C \mathcal{G}_B \mathcal{P} = \mathcal{G}_B \mathcal{P}$ 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}_B , using for instance the Reidemeister-Schreier algorithm [Sim94] implemented in GAP [GAP08].

Theorem 4.2 implies that there exists a space \mathbb{X}_C , in which the Delaunay triangulation of the point set $\pi(\mathcal{P})$ is defined. If \mathbb{X} is a manifold, then the space \mathbb{X}_C is a covering space of \mathbb{X} with a finite number of sheets [Arm82, Section 10.4]; If \mathbb{X} is a non-manifold orbifold, then \mathbb{X}_C is a covering orbifold [Thu02, Definition 13.2.2].

The construction of Theorem 4.2 can also be understood as a direct construction of \mathbb{X}_C from \mathbb{X} , as follows.

Each closed Euclidean d -orbifold has a d -torus as covering space with a finite number of sheets: this follows from Theorem 4.1. 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 5). The group generated by g_t and g_g^2 is a translation subgroup of \mathcal{G}_K of index 2.

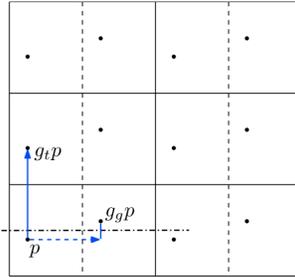


Figure 5: A part of the infinite point grid $\mathcal{G}_K p$.

When the d -torus has been found, i.e., when the translation subgroup with finite index has been found, a subgroup can be chosen so that it fulfills the condition of Theorem 4.2: 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 Criterion 3.11, i.e., to fulfill the inequality $\Delta(\mathcal{G}_C \mathcal{G}_B \mathcal{P}) = \Delta(\mathcal{G}_B \mathcal{P}) < \frac{\delta(\mathcal{G}_C)}{2}$. See Figure 6 for an illustration in two dimensions.

Let us remark that all the criteria that we are using give sufficient conditions. There may be covering spaces of \mathbb{X} with fewer sheets than the number given by this constructions, and still allowing us to construct a triangulation; the study of the flat torus gives an example (see Section 6.1).

As will be seen in Section 5, we actually avoid as much as possible to compute with copies of input points.

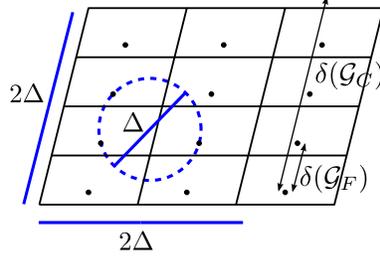


Figure 6: Sufficient number of copies of the fundamental domain.

5 Algorithm and Complexity

Let $\mathbb{X} = \mathbb{E}^d/\mathcal{G}_B$, where \mathcal{G}_B is a Bieberbach group, be a general closed Euclidean orbifold. The h -sheeted covering space $\mathbb{X}_C = \mathbb{E}^d/\mathcal{G}_C$ can be pre-computed as shown in Section 4.

We propose an algorithm derived from Bowyer’s [Bow81] that, for a given set \mathcal{P} of input points, returns the Delaunay triangulation of \mathbb{X} defined by \mathcal{P} if possible, and the Delaunay triangulation of \mathbb{X}_C otherwise.

We first give a general overview of the incremental algorithm in Section 5.1, then we present point insertion in more detail in Section 5.2. Finally, Section 5.3 shows that the algorithm has optimal randomized complexity, expressed as a function of the number of input points.

5.1 Overview

Let the input set of points \mathcal{P} be given. A Delaunay triangulation, initially empty, is computed incrementally by adding the points of \mathcal{P} one by one. The triangulation can be stored as a graph: each d -simplex stores a table of pointers to its vertices and a table of pointers to its 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 Delaunay triangulation whose circumscribing ball diameter is larger than the threshold $\delta(\mathcal{G}_B)/2$. Each such d -simplex of the Delaunay triangulation stores a pointer to its corresponding element in \mathcal{L} .

- (1) The algorithm starts computing the triangulation in \mathbb{X}_C , inserting for each input point the h points forming its fiber in the covering space. These h points are inserted one by one in the triangulation.⁴

The fiber can be determined as follows. Let $\mathcal{G}_Q = \mathcal{G}_B/\mathcal{G}_C$ denote the quotient group of \mathcal{G}_B and \mathcal{G}_C . As \mathcal{G}_C is a subgroup of \mathcal{G}_B of index h , the group \mathcal{G}_Q has h elements. For a given input point p , the h copies to be inserted in the triangulation are the h points of its orbit under the action of \mathcal{G}_Q .

The triangulation of \mathbb{X}_C and the list \mathcal{L} are updated at each insertion.

- (2) Once Criterion 3.11 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} : it iterates over all d -simplices and all vertices to delete

⁴We skip the technicalities of the initialization of the triangulation with the first point of \mathcal{P} .

all copies, keeping only one; it also updates the incidence relations of the d -simplices whose adjacent d -simplices have been deleted.

Then the algorithm inserts each of the remaining input points only once into the triangulation of \mathbb{X} .

If \mathcal{P} is such that Criterion 3.11 is never fulfilled (which will always be the case if \mathbb{X} is not a manifold, i.e., if \mathcal{G}_B does not act freely), then the algorithm never enters (2) and returns the triangulation of the covering space \mathbb{X}_C .

Note however that when \mathbb{X} is a manifold, even when Criterion 3.11 is not fulfilled at the end of the algorithm, it is still possible to check at the end of the computation whether the 1-skeleton of $\pi(DT(\mathcal{GP}))$ 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, there is no guarantee that it will be possible to insert further points in it.

5.2 Point Insertion

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 [Dev02]. 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 to a triangulation of \mathbb{X} when possible. 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 Criterion 3.11. 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 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 guarantees that the structure after the point insertion is a simplicial complex (see also Section 3.1), then the hole is always homeomorphic to a ball (Lemma 3.10 and Theorem 3.8).

5.3 Complexity

The randomized analysis of [Dev02] assumes that the insertion of points of \mathcal{P} is 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 h of sheets), 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 5.1. *The algorithm described in Section 5.1 has optimal randomized worst-case time complexity $O\left(n^{\lceil \frac{d}{2} \rceil} + \log n\right)$ and space complexity $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$, where n is the number of points in \mathcal{P} .*

Note that the complexity is only expressed as a function of the number of input points. The covering space \mathbb{X}_C is independent of the set of points (Section 4); it is precomputed beforehand.

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 5.1, 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.

The randomized worst-case analysis in [Dev02] 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 $O(\alpha)$. This property

extends to the case of computing in \mathbb{X}_C . The vertex set of the Delaunay triangulation in level i is the set \mathcal{P}'_i that contains all copies of the points of \mathcal{P}_i showing up in \mathbb{X}_C . 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 larger, 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}_B , so the asymptotic size of the hierarchy does not change with respect to [Dev02].

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 impact 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 $O(n \log n)$ [Dev02].

6 Further Results and Experiments in the Case of the 3D Cubic Flat Torus

In this section, we focus on the 3D flat torus $\mathbb{X} = \mathbb{T}_{\mathbf{c}}^3 = \mathbb{E}^3/\mathcal{G}_T$ with $\mathcal{G}_T = ((c\mathbb{Z})^3, +)$, where $\mathbf{c} = (c, c, c) \in (\mathbb{E} \setminus \{0\})^3$. The group is generated by three orthogonal translations of length c and the original domain \mathcal{D} is a half-open cube with the same edge length. Other notation are the same as in previous sections.

Implementation. The algorithm presented in Section 5 has been implemented for $\mathbb{T}_{\mathbf{c}}^3$. The package has been reviewed by the CGAL Editorial Board and has been included in the library since version 3.5 [CT]; a package for the 2D case was integrated in CGAL 4.3 [Kru].

We only present the implementation very briefly here; more details can be found in [Car10, Chapter 3]. There are several differences with CGAL 3D triangulations [PT], which are representing triangulations in the Euclidean space. In particular, covering spaces are handled in the following way:

- Each cell of the triangulation stores an *offset* for each of its vertices. A vertex stores a point in the original domain. The offset associated to a vertex is a vector giving the translation that specifies which

element of the orbit of the point is considered.

- We maintain the *covering map* that associates to each vertex the corresponding vertex in the original domain. Conversely, we also maintain the *reverse covering map* that allows to find all vertices associated to a given vertex in the original domain.

Degeneracies are handled using a symbolic perturbation that only relies on the lexicographic comparison between input points [DT03, DT11]. This symbolic perturbation allows us to define the Delaunay triangulation of a degenerate polytope Q_{cos} in a unique way. Since the lexicographic order is preserved by translations in the group $(\mathbb{Z}^3, +)$ each image $g \cdot Q_{cos}, g \in \mathbb{Z}^3$, will automatically be triangulated as the image by g of the triangulation of Q_{cos} , so, all these images project to the same triangulation of $\pi(Q_{cos})$ in $\mathbb{E}^3/\mathbb{Z}^3$.

In the sequel, we report a few observations that are specific to the case of \mathbb{T}_c^3 , and which are used in our CGAL implementation.

6.1 Number of sheets

Let us introduce new notation. The torus $\mathbb{T}_{hc}^3 = \mathbb{E}^3/(hc\mathbb{Z})^3$ is the h^3 -sheeted covering space of \mathbb{T}_c^3 . Its corresponding original domain \mathcal{D}_{hc} is a half-open cube with edge length hc . Let π_{h^3} denote the corresponding projection map.

Criterion 3.11 used in the construction of Theorem 4.2 and in the algorithm can be rewritten here as:

If $\Delta(\mathcal{G}_T\mathcal{P}) < \frac{c}{2}$, then $\pi(DT(\mathcal{G}_T\mathcal{P}'))$ is a triangulation of \mathbb{T}_c^3 for any finite $\mathcal{P}' \supseteq \mathcal{P}$.

A direct application of this criterion shows that a Delaunay triangulation can always be defined in a 64-sheeted covering space of the initial torus:

Corollary 6.1. *Let $\mathcal{P} \subset \mathcal{D}$, $\mathcal{P} \neq \emptyset$ be a finite point set. Then $\pi_{64}(DT(\mathcal{G}_T\mathcal{P}'))$ is a triangulation of \mathbb{T}_{4c}^3 for any finite $\mathcal{P}' \supseteq \mathcal{P}, \mathcal{P}' \subset \mathcal{D}$.*

Proof. For $|\mathcal{P}| = 1$, the largest empty ball in $\mathcal{G}_T\mathcal{P}$ has diameter $\sqrt{3}c$; empty balls cannot become larger when inserting more points. The domain \mathcal{D}_{4c} is a cube of edge length $4c$, which is more than twice $\sqrt{3}c$. Thus the above criterion applies. □

However, using a totally different proof, it is possible to reduce the number sheets from 64 to 27.⁵ This constant 27 is the same as the number of copies previously used in the literature [DH97]. Using the framework of normal subgroups and covering spaces is more powerful and general than using copies of the points in \mathbb{E}^2 (resp. \mathbb{E}^3). Also, there is no boundary in the data structure since whole manifolds are triangulated, and adjacency relations are valid for all simplices.

⁵The incremental algorithm actually requires a slightly stronger result than Proposition 6.2. All details can be found in [Car10, Section 2.3.1], in particular Lemma 2.3.2. We omit them in this paper, since the precise setting and the proof are long and technical.

Proposition 6.2. *For any finite point set $\mathcal{P} \subset \mathcal{D}$, $\mathcal{P} \neq \emptyset$, $\pi_{2\tau}(DT(\mathcal{G}_T\mathcal{P}))$ is a simplicial complex.*

Proof. The proof uses the proof of Dolbilin and Huson [DH97]. They consider the infinite Delaunay triangulation $DT(\mathcal{G}_T\mathcal{P})$ in the universal covering space \mathbb{E}^3 of \mathbb{T}_c^3 . They prove that only the points of $\mathcal{G}_T\mathcal{P}$ contained in \mathcal{D} and the 26 copies that surround it can have an influence on the simplices that are completely contained in \mathcal{D} .

Let $F^{(1)}$ denote the Voronoi cell of the origin O in the Voronoi diagram of the orbit \mathcal{G}_TO of O . Then $F^{(1)}$ is a fundamental domain of \mathcal{G}_T . Let furthermore \mathcal{Q} be a point set in $F^{(1)}$. $A \oplus B$ denotes the Minkowski sum of A and B , and we define $F^{(i)} = F^{(1)} \oplus F^{(i-1)}$ for $i \geq 2$.

Lemma 6.3 ([DH97]). *Let σ and τ be simplices in the infinite periodic Delaunay triangulation $DT(\mathcal{G}_T\mathcal{Q})$ such that one of the vertices of σ lies at O and $\tau \cap F^{(1)} \neq \emptyset$. Then*

- (1). *the center of the circumscribing ball of σ lies in $F^{(1)}$,*
- (2). *σ is completely contained inside $F^{(2)}$.*
- (3). *τ is completely contained inside $F^{(3)}$.*

See Figure 7 for an illustration of Lemma 6.3.

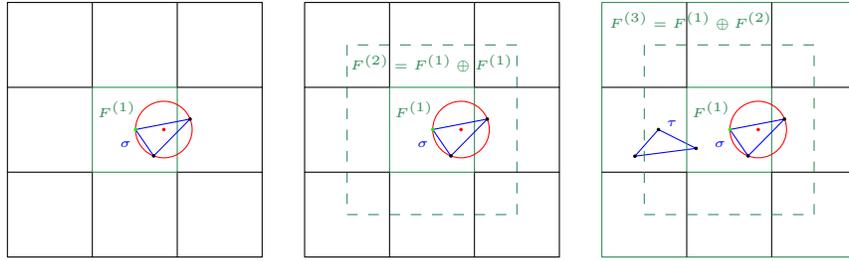


Figure 7: (2D illustration) It is sufficient to consider $F^{(3)}$ in order to compute a Delaunay triangulation of the flat torus.

Let us briefly sketch the proofs of the three properties.

- (1). This follows directly from the fact that $F^{(1)}$ is a Voronoi cell of the Voronoi diagram of \mathcal{G}_TO : If the center of the circumscribing ball of σ was outside $F^{(1)}$, then it would be closer to some other point of $\mathcal{G}_T\mathcal{Q}$, contradicting the fact that it is the center of the circumscribing ball of σ .
- (2). This follows directly from (1), and from the fact that $F^{(2)}$ is the Minkowski sum of $F^{(1)}$ with itself.
- (3). This follows directly from the definition of $F^{(3)}$.

From Lemma 6.3 follows that all triangles with vertices in \mathcal{Q} in the Delaunay triangulation of $\mathcal{G}_T\mathcal{Q}$ are contained inside $F^{(3)}$ and so it is sufficient to compute the finite Delaunay triangulation $DT(\mathcal{G}_T\mathcal{Q} \cap F^{(3)})$ in \mathbb{E}^3 . According to Lemma 6.3, all simplices of $DT(\mathcal{G}_T\mathcal{Q} \cap F^{(3)})$ that have at least one vertex in $F^{(1)}$ are simplices of the infinite periodic Delaunay triangulation $DT(\mathcal{G}_T\mathcal{Q})$. Since the set of these simplices covers $F^{(1)}$, applying the action of \mathcal{G} on it yields an infinite periodic partition of \mathbb{E}^3 .

Now, let us show that there are no cycles of length 2 in $\pi_{27}(DT(\mathcal{G}_T\mathcal{P}))$. Let $\mathcal{D}(i, j, k)$ denote the translated copy of \mathcal{D} by $(i \cdot c, j \cdot c, k \cdot c)$, i.e. $\mathcal{D}(i, j, k) = [i \cdot c, (i + 1) \cdot c) \times [j \cdot c, (j + 1) \cdot c) \times [k \cdot c, (k + 1) \cdot c)$. Assume that there is a cycle of length two in $\pi_{27}(DT(\mathcal{G}_T\mathcal{P}))$. Then there are vertices $v, v', v'' \in DT(\mathcal{G}_T\mathcal{P})$ such that the edges (v, v') and (v', v'') , $(v, v') \neq (v', v'')$, are contained in $DT(\mathcal{G}_T\mathcal{P})$ and that $\pi_{27}(v) = \pi_{27}(v'')$. Let $\mathcal{D}(i, j, k), \mathcal{D}(i', j', k')$, and $\mathcal{D}(i'', j'', k'')$ denote the translations of \mathcal{D} that contain v, v' , and v'' , respectively. According to Lemma 6.3, if a simplex intersects both $\mathcal{D}(i, j, k)$ and $\mathcal{D}(i', j', k')$, then $|i - i'| \leq 1, |j - j'| \leq 1$, and $|k - k'| \leq 1$. Without loss of generality, we can choose v and v' such that $(i, j, k), (i', j', k') \in \{0, 1\}^3$. Note that $\pi_{27}(\mathcal{D}(i, j, k)) = \pi_{27}(\mathcal{D}(i \bmod 3, j \bmod 3, k \bmod 3))$, so at least one of the $i'', j'',$ and k'' must be in $\{0 - 3, 1 - 3, 0 + 3, 1 + 3\}$ for $\pi_{27}(v) = \pi_{27}(v'')$ to hold. This is not possible according to Lemma 6.3.

This concludes the proof of Proposition 6.2. \square

6.2 Geometric criteria

Criterion 3.11 is a condition on the diameter of the largest circumscribing ball. However, computing the circumradius of a tetrahedron is numerically less robust than computing the distance of two points. Therefore, the implementation actually uses Criterion 6.4, which involves edge lengths, which are simpler to compute.

Criterion 6.4. *If $DT(\mathcal{G}_T\mathcal{P})$ contains only edges shorter than $\frac{1}{\sqrt{6}}c$, then $\pi(DT(\mathcal{G}_T\mathcal{P}'))$ is a triangulation of \mathbb{T}_c^3 for any finite $\mathcal{P}' \supseteq \mathcal{P}$.*

Proof. Let B be a ball of diameter d that does not contain any point of $\mathcal{G}_T\mathcal{P}$ in its interior. Consider the tetrahedron t in $DT(\mathcal{G}_T\mathcal{P})$ that contains the center of B . The length of the largest edge of t is bounded from below by the edge length of the regular tetrahedron with circumscribing ball B , which is $\frac{2d}{\sqrt{6}}$. So if all edges in $DT(\mathcal{G}_T\mathcal{P})$ are shorter than $\frac{1}{\sqrt{6}}c$, then the diameter of any empty ball is smaller than $\frac{1}{2}c$. The claim follows from Criterion 3.11 and Lemma 3.6. \square

Both criteria are only sufficient criteria, and Criterion 3.11 is weaker than Criterion 6.4: There are point sets with maximum empty ball diameter shorter than $\frac{1}{2}c$ but edges longer than $\frac{1}{\sqrt{6}}c$.

Section 6.3 presents experimental observations and a theoretical analysis of the number of points needed for a uniformly distributed point set to satisfy the criteria. Section 6.4 shows that computing in covering spaces can be avoided in practice. Section 6.5 shows the running time of the code.

6.3 Satisfying the Criteria

6.3.1 Experimental Study

We tested the implementation on real data from research in cosmology. The input sets consist of up to several hundreds of thousands of points, and their distributions are roughly uniform. This property holds for

most of the applications mentioned in the introduction. With these real data, usually less than 400 points are needed to reach the threshold on the length of edges (Criterion 6.4) and switch to computing in \mathbb{T}_c^3 .

We further tested on randomly generated point sets whether they define a Delaunay triangulation of \mathbb{T}_c^3 . For each number of points between 0 and 600 we computed periodic Delaunay triangulations of 100 different point sets⁶ and calculated the percentage of

- (a). point sets that define a Delaunay triangulation of \mathbb{T}_c^3 (Theorem 3.8),
- (b). point sets whose largest empty circumscribing ball in the Delaunay triangulation has diameter smaller than $\frac{1}{2} c$ (Criterion 3.11),
- (c). point sets for which the longest edge in the Delaunay triangulation is smaller than $\frac{1}{\sqrt{6}} c$ (Criterion 6.4).

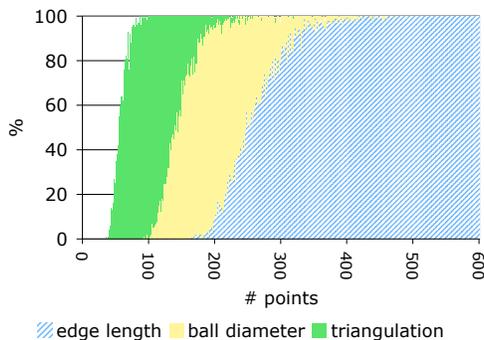


Figure 8: Percentage of random point sets that define a triangulation of \mathbb{T}_c^3 , that have largest ball diameter $< \frac{1}{2} c$, that have longest edge length $< \frac{1}{\sqrt{6}} c$.

Figure 8 shows that almost all sets of 100 points already define a Delaunay triangulation of \mathbb{T}_c^3 (Theorem 3.8); however it can happen that inserting a point in such a triangulation leads to a point set that does not define a Delaunay triangulation of \mathbb{T}_c^3 . In general 100 to 200 points are necessary for the largest empty ball diameter to be smaller than $\frac{1}{2} c$. According to Criterion 3.11 such point sets define a Delaunay triangulation of \mathbb{T}_c^3 even if we add further points. As Criterion 6.4 is a stronger criterion, about 200 to 400 points are required for the algorithm to switch to computing in the original manifold. When the data set becomes large, this number can be considered as a small constant.

6.3.2 Analysis

In this section, we consider the torus \mathbb{T}_1^3 whose fundamental domain is a unit cube. We give an upper bound on the number of points required to fulfill Criteria 3.11 and 6.4 if the input point set is uniformly distributed.

⁶Using the random point generator from CGAL to generate uniformly distributed points in a cube.

Let \mathcal{A} be a finite set of balls of diameter $\frac{1}{4}$ whose union covers \mathbb{T}_1^3 . Lemma 6.5 below describes such a cover with 180 balls. We first observe that any ball B of diameter $\frac{1}{2}$ must contain at least one ball of \mathcal{A} , namely the ball that contains B 's center. So, if \mathcal{P} is a set of n points evenly distributed in \mathbb{T}_1^3 , we get:

$$\begin{aligned}
\text{Prob}(\exists B, \text{diameter}(B) \geq \frac{1}{2}, B \cap \mathcal{P} = \emptyset) &\leq \text{Prob}(\exists A \in \mathcal{A}, A \cap \mathcal{P} = \emptyset) \\
&\leq \sum_{A \in \mathcal{A}} \text{Prob}(A \cap \mathcal{P} = \emptyset) \\
&= 180 \text{Prob}(A_0 \cap \mathcal{P} = \emptyset) \text{ for some } A_0 \in \mathcal{A} \\
&= 180 (1 - \text{volume}(A_0))^n \\
&= 180 \left(1 - \frac{4\pi}{3 \cdot 8^3}\right)^n.
\end{aligned}$$

So Criterion 3.11 will be fulfilled with probability $> 1 - \rho$ if

$$n \geq \frac{\log_2 \rho - \log_2 180}{\log_2 \left(1 - \frac{4\pi}{3 \cdot 8^3}\right)} \simeq 632 + 58 \log_2 \frac{1}{\rho}.$$

The condition will be obtained with probability higher than $\frac{1}{2}$ as soon as about 700 points are inserted and with probability higher than $1 - \frac{1}{1000}$ as soon as about 1500 points are inserted.

It just remains to prove the following:

Lemma 6.5. \mathbb{T}_1^3 can be covered by a set of 180 balls of diameter $\frac{1}{4}$.

Proof. We consider a dense sphere packing in \mathbb{E}^3 consisting of spheres of radius r , built in a classical way (sometimes named a cannonball packing): the x -axis contains centers of spheres equally spaced at distance $2r$ next from another, and the horizontal plane contains centers of spheres forming equilateral triangles of side length $2r$; in the next plane, spheres are placed on top of each hollow between three spheres centered in the first plane; and so on. Spheres pairwise touching form a regular tetrahedra of side length $2r$, and whose circumradius is $\sqrt{3/2}r$. The set of balls having the same centers but with radius $\sqrt{3/2}r$ covers \mathbb{E}^3 (see Figure 9). Let us center one of the spheres at the origin $(0, 0, 0)$. Then the spheres touching it are centered at points $(2r, 0, 0)$, $(r, \sqrt{3}r, 0)$, and $(\frac{1}{2}r, \frac{1}{\sqrt{3}}r, \frac{2\sqrt{2}}{\sqrt{3}}r)$.

We will set $r = \frac{1}{8}\sqrt{\frac{2}{3}} = \frac{1}{4\sqrt{6}}$ in order to get spheres of diameter $\frac{1}{4}$. The projections of spheres whose center lies in the unit cube are enough to cover \mathbb{T}_1^3 . The number of these spheres is:

$$\left\lceil \frac{1}{2r} \right\rceil \times \left\lceil \frac{1}{\sqrt{3}r} \right\rceil \times \left\lceil \frac{\sqrt{3}}{2\sqrt{2}r} \right\rceil = \lceil 2\sqrt{6} \rceil \times \lceil 4\sqrt{2} \rceil \times \lceil 6 \rceil = \lceil 4.9 \rceil \times \lceil 5.7 \rceil \times 6 = 180. \quad \square$$

The longest edge in a Delaunay triangulation is bounded by the diameter of the largest empty ball. Thus the same analysis works for estimating the number of points required such that the longest edge in the triangulation is shorter than $\frac{1}{\sqrt{6}}$, as required by Criterion 6.4. We just have to replace the constant $\frac{1}{2}$ for the ball diameter by $\frac{1}{\sqrt{6}}$. We get a cover of \mathbb{T}_1^3 by 384 spheres of diameter $\frac{1}{2\sqrt{6}}$, and we can conclude that Criterion 6.4 is satisfied with with probability higher than $\frac{1}{2}$ when about 800 points are inserted.

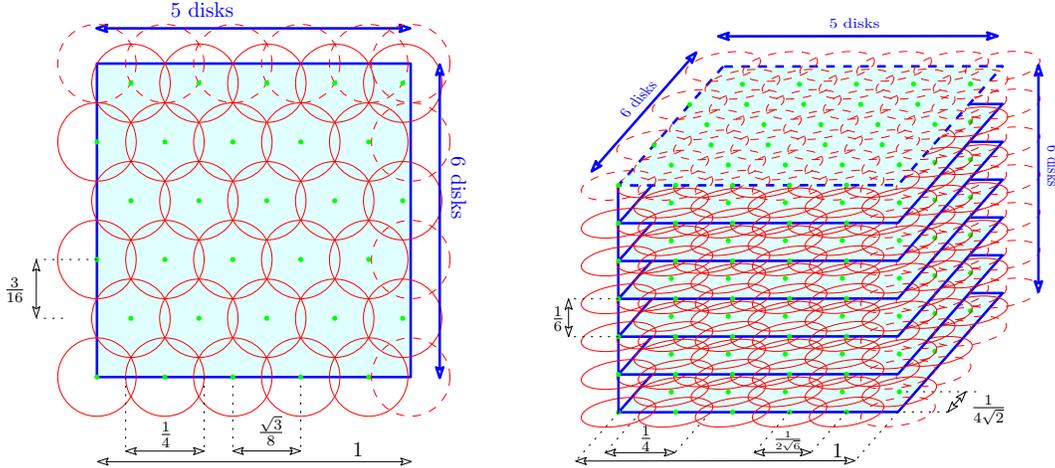


Figure 9: Covering (left) \mathbb{T}_1^2 with 30 disks and (right) \mathbb{T}_1^3 with 180 balls of diameter $\frac{1}{4}$.

The computations of probabilities above use overevaluations; as seen in the previous section, experiments give better results.

6.4 Avoiding Covering Spaces in Practice

We propose an optimization that avoids computing in the 27-sheeted covering space \mathbb{T}_{3c}^3 at all. We start with a set \mathcal{S} of 36 dummy points, chosen such that the diameter of the largest empty sphere is smaller than $\frac{1}{2}c$: Let $\mathcal{S}_{(a,b,c)}^0 = \{(x+a, y+b, z+c) * \mathbf{c} \mid x, y \in \{0, \frac{1}{3}, \frac{2}{3}\}, z \in \{0, \frac{1}{2}\}\}$, then \mathcal{S} is defined as $\mathcal{S} = \mathcal{S}_{(0,0,0)}^0 \cup \mathcal{S}_{(\frac{1}{6}, \frac{1}{6}, \frac{1}{4})}^0$ (see Figure 10). From Criterion 3.11 follows that any point set $\mathcal{S}' \in \mathcal{D}$ with $\mathcal{S} \subset \mathcal{S}'$ defines a Delaunay triangulation of \mathbb{T}_c^3 . So, we can insert all points of the input data set \mathcal{P} and compute the Delaunay triangulation of $\mathcal{S}' = \mathcal{S} \cup \mathcal{P}$ in \mathbb{T}_c^3 . In the end we remove the dummy points, which requires to convert the final result to \mathbb{T}_{3c}^3 in case the data set \mathcal{P} alone does not define a Delaunay triangulation of \mathbb{T}_c^3 .

It would be more intuitive to take 27 dummy points on a regular grid: $\{\frac{1}{3}(x, y, z) \mid x, y, z \in \{0, 1, 2\}\}$. However, as shown on Figure 11, this point set does not prevent cycles of length 2.

6.5 Timings

We compared the running time of our implementation for computing Delaunay triangulations of \mathbb{T}_c^3 with the running time of computing the Delaunay triangulation of \mathbb{E}^3 with the CGAL package [PT]. Our implementation benefits from some of the optimizations that are already available in the CGAL Delaunay triangulations of \mathbb{E}^3 , such as spatial sorting [DD].

For large random point sets, experiments show a factor of about 1.6 between the running time of our current implementation with the above optimization using dummy points, and the CGAL implementation

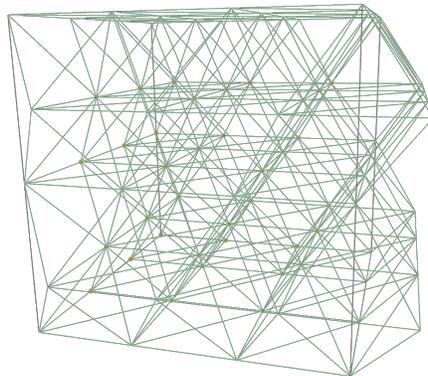


Figure 10: The Delaunay triangulation of the dummy point set.

for \mathbb{E}^3 : the Delaunay triangulation defined by one million points is computed in about 23 seconds for $\mathbb{T}_{\mathcal{C}}^3$, instead of about 14 seconds for \mathbb{E}^3 (on a 2.33 GHz Intel Core 2 Duo processor). The timings have been measured for the case when the original domain is a unit cube, using specialized predicates; if we allow any cube, we lose about 12%. More experiments can be found in [Car10].

7 Future work

An interesting 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. Also, this case would find applications in fields as diverse as computer graphics [RJSG11] and neuromathematics [CF09, CFF11]. The problem is quite challenging even in the simplest case of a the Bolza surface (homeomorphic to a torus having two handles), which corresponds to a group generated by four hyperbolic translations in \mathbb{H}^2 ; it was recently shown that a 128-sheeted covering space is always sufficient [BTV16] but this is probably not tight.

Acknowledgments

The authors wish to thank Olivier Devillers for contributions to Sections 5.3 and 6.3.2, Ramsay Dyer for discussions about the hypothesis in Proposition 2.1, Nico Kruithof for initial work on 3D periodic triangulations, Günter Rote for his comments on a preliminary version of [CT09], Jean-Marc Schlenker for helpful discussions, and Rien van de Weijgaert for providing us with data sets from cosmology research projects. We also acknowledge reviewers of a first version of this paper for their useful comments.

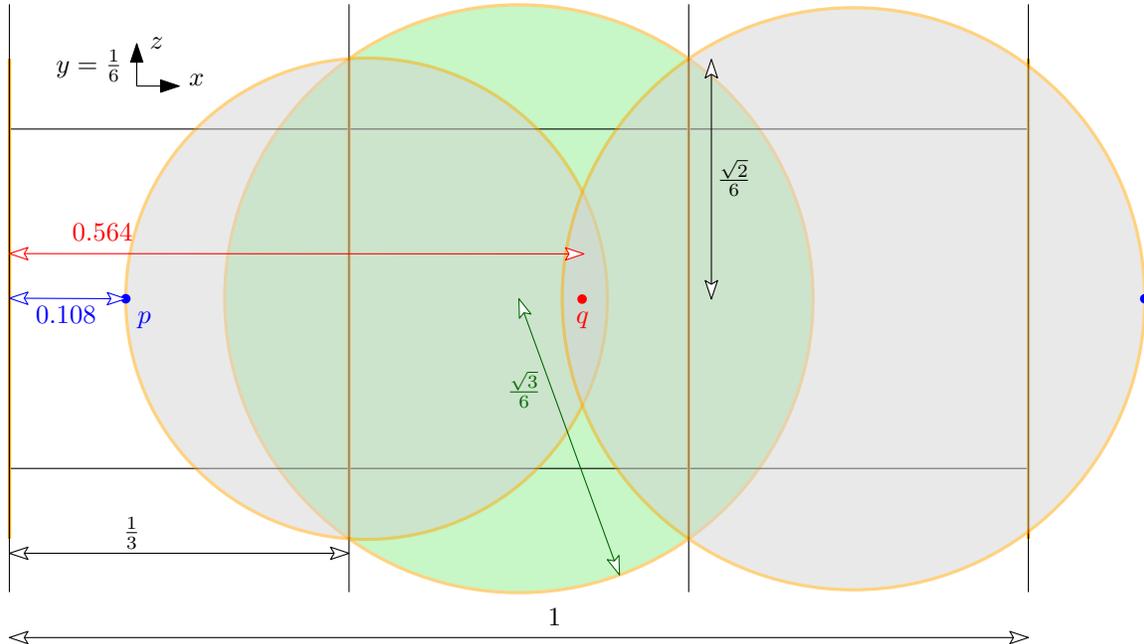


Figure 11: Section in the plane $y = \frac{1}{6}$. 27 points are placed on a regular grid of size $\frac{1}{3}$. The two points $p = (0.108, \frac{1}{6}, \frac{1}{6})$ and $q = (0.564, \frac{1}{6}, \frac{1}{6})$ form a cycle of length 2.

References

- [Arm82] M. A. Armstrong. *Basic Topology*. Springer-verlag, 1982.
- [ART⁺] Pierre Alliez, Laurent Rineau, Stéphane Tayeb, Jane Tournois, and Mariette Yvinec. 3D mesh generation. In *CGAL User and Reference Manual*. CGAL Editorial Board. http://doc.cgal.org/latest/Manual/packages.html#PkgMesh_3Summary.
- [BCdVN13] Alexandre Boulch, Éric Colin de Verdière, and Atsuhiko Nakamoto. Irreducible triangulations of surfaces with boundary. *Graphs and Combinatorics*, 29(6):1675–1688, 2013. Also in arXiv:1103.5364 [math.CO].
- [Ber09] Julie Bernauer. Computational structural biology: Periodic triangulations for molecular dynamics, 2009. Talk at the workshop ‘Subdivide and tile’, <http://www.lorentzcenter.nl/lc/web/2009/357/info.php3?wsid=357>.
- [Bie10] Ludwig Bieberbach. Über die Bewegungsgruppen des n -dimensionalen euklidischen Raumes mit einem endlichen Fundamentalbereich. *Göttinger Nachrichten*, 75, 1910.
- [BMP03] Michel Boileau, Sylvain Maillot, and Joan Porti. *Three-dimensional orbifolds and their geometric structures*. Société mathématique de France, Paris, 2003.
- [Bow81] Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24:162–166, 1981.

- [BS97] Ken A. Brakke and John M. Sullivan. Using symmetry features of the surface evolver to study foams. In Hans-Christian Hege and Konrad Polthier, editors, *Visualization and mathematics: experiments, simulations and environments*, pages 95–118. Springer-Verlag, 1997. <http://portal.acm.org/citation.cfm?id=270122.270129>.
- [BTV16] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations on orientable surfaces of low genus. In *Proceedings of the Thirty-second International Symposium on Computational Geometry*, 2016. <https://hal.inria.fr/hal-01276386>.
- [BY98] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann, <http://www.cup.cam.ac.uk/Scripts/webbook.asp?isbn=0521563224>.
- [Car10] 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/>.
- [CF09] Pascal Chossat and Olivier Faugeras. Hyperbolic planforms in relation to visual edges and textures perception. *PLoS Computational Biology*, 5(12):e1000625, 2009. <https://hal.inria.fr/hal-00807344>.
- [CFF11] Pascal Chossat, Grégory Faye, and Olivier Faugeras. Bifurcation of hyperbolic planforms. *Journal of Nonlinear Science*, 21(4):465–498, 2011. <https://hal.inria.fr/hal-00807355>.
- [cgal] CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
- [CT] Manuel Caroli and Monique Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board. <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3Summary>.
- [CT08] Manuel Caroli and Monique Teillaud. Video: On the computation of 3D periodic triangulations. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, pages 222–223, 2008. <http://www.computational-geometry.org/SoCG-videos/socg08video/>.
- [CT09] 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/>.
- [CT11] Manuel Caroli and Monique Teillaud. Delaunay triangulations of point sets in closed Euclidean d -manifolds. In *Proceedings of the Twenty-seventh Annual Symposium on Computational Geometry*, pages 274–282, 2011. <https://hal.inria.fr/hal-01101094>.
- [dBvKOS00] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.
- [DD] Christophe Delage and Olivier Devillers. Spatial sorting. In *CGAL User and Reference Manual*. CGAL Editorial Board. <http://doc.cgal.org/latest/Manual/packages.html#PkgSpatialSortingSummary>.
- [Dev02] Olivier Devillers. The Delaunay hierarchy. *International Journal of Foundations of Computer Science*, 13:163–180, 2002.

- [dFC03] Gianni de Fabritiis and Peter V. Coveney. Dynamical geometry for multiscale dissipative particle dynamics. <http://xxx.lanl.gov/abs/cond-mat/0301378v1>, 2003.
- [DH97] Nikolai P. Dolbilin and Daniel H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.
- [DT03] 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>.
- [DT11] Olivier Devillers and Monique Teillaud. Perturbations for Delaunay and weighted Delaunay 3D triangulations. *Computational Geometry: Theory and Applications*, 44:160–168, 2011. <http://hal.inria.fr/inria-00560388/>.
- [Duq] Daniel Duque Campayo. Sklogwiki - Boundary conditions. http://www.sklogwiki.org/SklogWiki/index.php/Boundary_conditions.
- [EH10] Herbert Edelsbrunner and John L. Harer. *Computational Topology, An Introduction*. American Mathematical Society, USA, 2010.
- [FSSB07] Matthew Fisher, Boris Springborn, Peter Schröder, and Alexander I. Bobenko. An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing. *Computing*, 81(2-3):199–213, 2007. Special Issue on Industrial Geometry.
- [GAP08] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.4.12*, 2008. <http://www.gap-system.org>.
- [GGL95] Ronald L. Graham, Martin Grötschel, and László Lovász, editors. *Handbook of Combinatorics*. Elsevier, Amsterdam, Lausanne, New York, 1995.
- [GM01] Clara I. Grima and Alberto Márquez. *Computational Geometry on Surfaces*. Kluwer Academic Publishers, 2001.
- [Hen79] Michael Henle. *A Combinatorial Introduction to Topology*. Dover publication, New York, 1979.
- [HvdWV⁺12] Johan Hidding, Rien van de Weygaert, Gert Vegter, Bernard J.T. Jones, and Monique Teillaud. Video: The sticky geometry of the cosmic web. In *Proceedings of the Twenty-eighth Annual Symposium on Computational Geometry*, pages 421–422, 2012. <http://www.computational-geometry.org/SoCG-videos/socg12video/>.
- [Kru] Nico Kruithof. 2D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board. <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic2Triangulation2Summary>.
- [Lee00] John M. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, New York, 2000.
- [LvdSH10] Erik Lindahl, David van der Spoel, and Berk Hess *et al.* Gromacs, a versatile package to perform molecular dynamics, 2010. version 4.5.3, <http://www.gromacs.org/>.

- [Moe08] 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/>.
- [MR97] Marisa Mazón and Tomás Recio. Voronoi diagrams on orbifolds. *Computational Geometry: Theory and Applications*, 8:219–230, 1997.
- [PT] Sylvain Pion and Monique Teillaud. 3D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board. <http://doc.cgal.org/latest/Manual/packages.html#PkgTriangulation3Summary>.
- [RJSJ11] Guodong Rong, Miao Jin, Liang Shuai, and Xiaohu Guo. Centroidal Voronoi tessellation in universal covering space of manifold surfaces. *Computer Aided Geometric Design*, 28(8):475–496, 2011.
- [Rob06] Vanessa Robins. Betti number signatures of homogeneous Poisson point processes. *Physical Review E*, 74(061107), 2006.
- [RY07] Laurent Rineau and Mariette Yvinec. Meshing 3D domains bounded by piecewise smooth surfaces. In *Proceedings of the sixteenth International Meshing Roundtable*, pages 443–460, 2007.
- [Sim94] Charles C. Sims. *Computing with Finitely Presented Groups*. Cambridge University Press, Cambridge, UK, 1994.
- [Slo] N. J. A. Sloane. The on-line encyclopedia of integer sequences. <http://oeis.org/>.
- [Sou11] Thierry Sousbie. The persistent cosmic web and its filament structure I: Theory and implementation. *Monthly Notices of the Royal Astronomical Society*, 414:350383, 2011. Also in arXiv:1009.4015.
- [Spa66] Edwin H. Spanier. *Algebraic Topology*. Springer-Verlag, New York, 1966.
- [SPK11] Thierry Sousbie, Christophe Pichon, and Hajime Kawahara. The persistent cosmic web and its filament structure II: illustrations, 2011. Also in arXiv:1009.4014.
- [Tho02] Karsten E. Thompson. Fast and robust Delaunay tessellation in periodic domains. *International Journal for Numerical Methods in Engineering*, 55:1345–1366, 2002.
- [Thu97] William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, Princeton, New Jersey, 1997.
- [Thu02] William P. Thurston. *The Geometry and Topology of Three-Manifolds*. 2002. <http://www.msri.org/publications/books/gt3m/>.
- [vdWPJ⁺11] Rien van de Weygaert, Pratyush Pranav, Bernard J.T. Jones, E.G. Patrick Bos, Gert Vegter, Herbert Edelsbrunner, Monique Teillaud, Wojciech A. Hellwing, Changbom Park, Johan Hidding, and Mathijs Wintraecken. Probing dark energy with alpha shapes and Betti numbers. Research report, 2011. <http://arxiv.org/abs/1110.5528> [astro-ph.CO].
- [vdWPV⁺10] Rien van de Weygaert, Erwin Platen, Gert Vegter, Bob Eldering, and Nico Kruithof. Alpha shape topology of the cosmic web. In *Proceedings of the 2010 International Symposium on Voronoi Diagrams in Science and Engineering*, ISVD '10, pages 224–234, 2010. Also in arXiv:1006.2765.

- [vdWVE⁺11] Rien van de Weygaert, Gert Vegter, Herbert Edelsbrunner, Bernard J.T. Jones, Pratyush Pranav, Changbom Park, Wojciech A. Hellwing, Bob Eldering, Nico Kruithof, E.G.P. (Patrick) Bos, Johan Hidding, Job Feldbrugge, Eline ten Have, Matti van Engelen, Manuel Caroli, and Monique Teillaud. Alpha, Betti and the megaparsec universe: on the homology and topology of the cosmic web. In *Transactions on Computational Science XIV*, volume 6970 of *Lecture Notes in Computer Science*, pages 60–101. Springer-Verlag, 2011. <http://www.springerlink.com/content/334357373166n902/>.
- [Wei08] Dahlia Weiss. How hydrophobic Buckminsterfullerene affects surrounding water structure. INRIA Geometrica Seminar, <http://www-sop.inria.fr/geometrica>, March 2008.
- [Wil08] P. M. H. Wilson. *Curved Spaces*. Cambridge University Press, Cambridge, 2008.
- [Zom05] Afra Zomorodian. *Topology for Computing*. Cambridge University Press, Cambridge, 2005.
- [Zom10] Afra Zomorodian. The tidy set: A minimal simplicial set for computing homology of clique complexes. In *Proceedings of the Twenty-sixth Annual Symposium on Computational Geometry*, pages 257–266, 2010.