



HAL
open science

Computing 3D Periodic Triangulations

Manuel Caroli, Monique Teillaud

► **To cite this version:**

Manuel Caroli, Monique Teillaud. Computing 3D Periodic Triangulations. [Research Report] RR-6823, 2009. inria-00356871v3

HAL Id: inria-00356871

<https://inria.hal.science/inria-00356871v3>

Submitted on 2 Feb 2009 (v3), last revised 12 Jun 2009 (v4)

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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Computing 3D Periodic Triangulations

Manuel Caroli — Monique Teillaud

N° 6823

Janvier 2009

Thème SYM



R
rapport
de recherche

Computing 3D Periodic Triangulations

Manuel Caroli* , Monique Teillaud*

Thème SYM — Systèmes symboliques
Équipes-Projets Géométrica

Rapport de recherche n° 6823 — Janvier 2009 — 18 pages

Abstract: This work is motivated by the need for software computing 3D periodic triangulations in numerous domains including astronomy, material engineering, biomedical computing, fluid dynamics etc.

We design an algorithmic test to check whether a partition of the 3D flat torus actually forms a triangulation (which subsumes that it is a simplicial complex). We propose an incremental algorithm that computes the Delaunay triangulation of a set of points in the 3D flat torus without duplicating any point, whenever possible. As far as we know, this is the first algorithm of this kind whose output is provably correct. Proved algorithms found in the literature are in fact always computing with 27 copies of the input points, while our algorithmic test detects when such a duplication can be avoided, which is usually possible in practical situations.

The algorithm was implemented and is going to be submitted to the CGAL Editorial Board[†]. A video of this work was presented at SoCG'08[‡].

Key-words: triangulation, periodic, flat torus, Delaunay

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/>.

* INRIA Sophia Antipolis – Méditerranée [Email: {Manuel.Caroli,Monique.Teillaud}@sophia.inria.fr]

[†] www.cgal.org

[‡] <http://www.computational-geometry.org/SoCG-videos/socg08video>

Calculer des triangulations périodiques en 3D

Résumé : Ce travail est motivé par le besoin d'un logiciel pour calculer des triangulations périodiques 3D dans de nombreux domaines dont l'astronomie, l'ingénierie des matériaux, l'informatique bio-médicale, la dynamique des fluides, etc.

Nous concevons un test algorithmique pour vérifier qu'une partition du tore plat 3D forme une triangulation (ce qui comprend la propriété d'être un complexe simplicial). Nous proposons un algorithme incrémental qui calcule la triangulation de Delaunay d'un ensemble de points dans le tore plat 3D, sans duplication de points si possible. C'est à notre connaissance le premier algorithme de ce genre pour lequel la correction du résultat soit prouvée. Les algorithmes prouvés que l'on trouve dans la littérature calculent avec 27 copies des points, alors que notre test détecte les cas où cette duplication peut être évitée, ce qui est en général possible dans les situations pratiques.

L'algorithme a été programmé et sera soumis au comité éditorial de la bibliothèque CGAL[§]. Une vidéo de ce travail a été présentée à SoCG'08[¶].

Mots-clés : triangulation, périodique, tore plat, Delaunay

[§] www.cgal.org

[¶] www.computational-geometry.org/SoCG-videos/socg08video

1 Introduction

Computing Delaunay triangulations of point sets is a well-studied problem in Computational Geometry. Several algorithms [3, 9] as well as implementations [30, 25, 19, 37, 24, 20] are available. However, these algorithms are mainly restricted to triangulations in \mathbb{R}^d . In this paper, we take interest in triangulations of a *periodic* space, represented as the so-called *flat torus* [33].

This 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. In fact there are numerous application fields needing robust software for geometric problems in periodic spaces. A small sample of these needs, in fields like astronomy, material engineering for prostheses, mechanics of granular materials, was presented at the CGAL Prospective Workshop on Geometric Computing in Periodic Spaces.¹ Many other diverse application fields could be mentioned, for instance biomedical computing [35], solid-state chemistry [28], physics of condensed matter [16], fluid dynamics [10], this list being far from exhaustive.

So far we are not aware of any robust and efficient algorithm for computing Delaunay triangulations from a given point set \mathcal{S} in a periodic space. In the literature, proved algorithms usually need to compute with 9 copies of each input point in the planar case [22, 18], or with 27 copies in 3D [15], which obviously leads to a huge slow-down.

In the engineering community, an implementation for computing a periodic Delaunay “*tessellation*” was proposed, avoiding duplications of points [32]. However, the tessellation is not necessarily a simplicial complex. Moreover, the algorithm heavily relies on the assumption that input points are well distributed.

In fact, as shown in Section 4, using copies of the input points may actually be necessary: in some cases, the flat torus may be partitioned into tetrahedra having the points as vertices and satisfying the Delaunay property, but such a partition does not always form a simplicial complex. However, in practice, data sets are likely to admit a Delaunay triangulation.

Let us insist here on the fact that computing a “true” triangulation, i.e. a *simplicial complex*, is important for several reasons. First, of course, a triangulation is always defined as a simplicial complex in the literature. Moreover, designing a data structure to efficiently store tetrahedral tessellations that are non-simplicial complexes would be quite involved and the problem is not even well-defined; The CGAL 3D triangulation data structure, that we reuse in our implementation, actually assumes the structure to be a simplicial complex [23]. Even more importantly, algorithms using a triangulation as input are actually heavily relying on the fact that the triangulation is a simplicial complex; this is the case for instance for meshing algorithms [27, 26], as well as algorithms to compute α -shapes, which are actually needed in the periodic case by several applications mentioned at the beginning of this introduction. We are planning to use the 3D periodic triangulation as the fundamental ingredient for computing these structures in the future.

Contributions of the paper

We prove conditions ensuring that the Delaunay triangulation can be actually computed without duplicating the input points. To this aim, we design an algorithmic test for checking whether a set \mathcal{K} of simplices in the flat torus forms a simplicial complex.

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

We present an adaptation of the well-known incremental algorithm in \mathbb{R}^3 [4] that allows to compute three-dimensional Delaunay triangulations in the flat torus. We focus on the incremental algorithm for several reasons: Its practical efficiency has been proved in particular by the fully dynamic implementation in CGAL [24]; moreover, a dynamic algorithm, allowing to freely insert (and remove) points, is a necessary ingredient for all meshing algorithms and software based on Delaunay refinement methods (see for instance [31, 27, 8]).

For sets of points that cannot be triangulated in the flat torus, our algorithm outputs a triangulation of an h -sheeted covering space, i.e. a triangulation of $h > 1$ explicit copies of the point set on the flat torus, where h depends on some parameters of the flat torus. However, as soon as the above mentioned conditions are fulfilled, the algorithm switches to a 1-sheeted covering and so does not duplicate points. In this way, the algorithm always computes a triangulation and is provably correct. It has optimal randomized worst case complexity.

The algorithm was implemented. The software is going to be submitted to the CGAL Editorial Board for inclusion in a future version of the CGAL library [7]. This software was demonstrated by a video [6]. A potential re-design of the CGAL 3D triangulation package, allowing a future extension to other spaces was proposed [5].

The paper is organized as follows. In Section 2 we review some general notions about triangulations and simplicial complexes. In the next section, we adapt the definition of simplicial complexes to the flat torus. In Section 4 we give a criterion to decide whether a point set has a triangulation on the flat torus. We give a second criterion that is based on the same idea but can be verified easily by the algorithm that is presented in Section 5. We prove the correctness of the algorithm and discuss its extension to computing weighted Delaunay triangulations. We finish with its complexity analysis and experimental observations.

2 Triangulations

Before talking about triangulations we need to recapitulate the well-known notions of simplices and simplicial complexes. A k -simplex σ in \mathbb{R}^3 ($k \leq 3$) 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 $\tau \leq \sigma$. Note that $\sigma \geq \sigma$ and $\sigma \leq \sigma$.

The following definitions are completely combinatorial. With an appropriate definition of a simplex, they will remain valid in any topological space \mathbb{X} .

There exist several definitions of simplicial complexes in the literature. Often they restrict to a finite number of simplices [38, 29]. In the sequel, we deal with infinite simplicial complexes, so, we use the definition given in [21]:

Definition 2.1 (Simplicial complex). 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 if \mathcal{K} is finite, then the third condition is always fulfilled.

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 of a point set* \mathcal{S} is a triangulation such that the set of vertices of the triangulation is identical to \mathcal{S} .

Some more definitions are needed for the following discussion: Let \mathcal{K} be a simplicial complex. If a subset of \mathcal{K} is a simplicial complex as well, we call it *subcomplex* of \mathcal{K} . The *closure* $\overline{\mathcal{L}}$ of a subset $\mathcal{L} \subseteq \mathcal{K}$ is the smallest subcomplex containing \mathcal{L} : $\overline{\mathcal{L}} = \{\tau \in \mathcal{K} | \tau \leq \sigma \in \mathcal{L}\}$. The *star* of \mathcal{L} consists of the cofaces of simplices in \mathcal{L} : $\text{St}(\mathcal{L}) = \{\sigma \in \mathcal{K} | \sigma \geq \tau \in \mathcal{L}\}$. In the following sections, we will be interested in the union of simplices in the star of a set \mathcal{L} of simplices, denoted as $|\text{St}(\mathcal{L})|$. We denote the interior of a simplex σ by $\overset{\circ}{\sigma}$. Two simplices σ and τ are said to be *internally disjoint* if their interiors are disjoint, i.e. $\overset{\circ}{\sigma} \cap \overset{\circ}{\tau} = \emptyset$.

3 The flat torus $\mathbb{T}_{\mathbf{c}}^3$

At first we give a precise definition of the space of study $\mathbb{T}_{\mathbf{c}}^3$. Then we review some of its well-known properties and establish the notations used in the following discussion. Finally, we give a definition of simplices in $\mathbb{T}_{\mathbf{c}}^3$.

Definition 3.1 (\mathbb{T}^3). Let $\mathbf{c} := (c_x, c_y, c_z) \in (\mathbb{R} \setminus \{0\})^3$ and G be the group $(\mathbf{c} * \mathbb{Z}^3, +)$, where $*$ denotes coordinate-wise multiplication². The quotient space $\mathbb{T}_{\mathbf{c}}^3 = \mathbb{R}^3 / G$ is called *flat torus* [33]. We denote the quotient map by $\pi : \mathbb{R}^3 \rightarrow \mathbb{T}_{\mathbf{c}}^3$.

Another commonly used denomination is *locally euclidean torus* [36, 33]. The space $\mathbb{T}_{\mathbf{c}}^3$ is a special case of an *orbifold* [34]. The elements of $\mathbb{T}_{\mathbf{c}}^3$ are the equivalence classes under the equivalence relation $p_1 \sim p_2 \Leftrightarrow p_1 - p_2 \in \mathbf{c} * \mathbb{Z}^3$, for $p_1, p_2 \in \mathbb{R}^3$. Hence, these equivalence classes are isomorphic to \mathbb{Z}^3 and $\mathbb{T}_{\mathbf{c}}^3 \times \mathbb{Z}^3$ is isomorphic to \mathbb{R}^3 . We also call the points of $\mathbb{T}_{\mathbf{c}}^3$ *orbits* and refer to their elements as *representatives*. $\mathbb{T}_{\mathbf{c}}^3$ is a metric space with $d_{\mathbb{T}}(\pi(p), \pi(q)) := \min d_{\mathbb{R}}(p', q')$ for $p' \sim p, q' \sim q$. Note that π is continuous.

The space $\mathbb{T}_{\mathbf{c}}^3$ is homeomorphic to the hypersurface of a 4-dimensional torus. Consider the closed cube $[u, u + c_x] \times [v, v + c_y] \times [w, w + c_z]$. Identifying the pairs of opposite sides results in a space homeomorphic to $\mathbb{T}_{\mathbf{c}}^3$. Such a cube is usually called *fundamental domain* or *fundamental region*; A fundamental domain contains at least one representative of each orbit. The half-open cube $\mathcal{D}_{\mathbf{c}} = [0, c_x) \times [0, c_y) \times [0, c_z)$ contains exactly one representative for each element of $\mathbb{T}_{\mathbf{c}}^3$. We call it *original domain*.

The map $\varphi_{\mathbf{c}} : \mathcal{D}_{\mathbf{c}} \times \mathbb{Z}^3 \rightarrow \mathbb{R}^3$
 $(p, \zeta) \mapsto p + \mathbf{c} * \zeta$

is bijective. The longest diagonal of $\mathcal{D}_{\mathbf{c}}$ has length $\|\mathbf{c}\|$, which denotes the L_2 -norm of \mathbf{c} . We say that two points $p_1, p_2 \in \mathbb{R}^3$ are *periodic copies* of each other if they both lie in the same orbit, or equivalently if there is a point $p \in \mathcal{D}_{\mathbf{c}}$ such that $p_1, p_2 \in \varphi_{\mathbf{c}}(\{p\} \times \mathbb{Z}^3)$.

Now we turn towards the definition of simplices in $\mathbb{T}_{\mathbf{c}}^3$. There is no meaningful definition of a convex hull in $\mathbb{T}_{\mathbf{c}}^3$ and a tetrahedron is not uniquely defined by four points. We attach with each vertex an integer vector, named *offset*, that specifies one representative out of an orbit (see Figure 1). In the above definition of $\varphi_{\mathbf{c}}$, the offsets are the numbers $\zeta \in \mathbb{Z}^3$.




Figure 1: (2D case) The three points p_1, p_2 , and p_3 do not uniquely define a triangle. Intuitively, the offset allows to know which way the triangle “wraps around” the torus.

²coordinate-wise multiplication: $(a_x, a_y, a_z) * (b_x, b_y, b_z) := (a_x b_x, a_y b_y, a_z b_z)$

We can adapt the definition of a simplex in \mathbb{R}^3 in the following way to $\mathbb{T}_{\mathbf{c}}^3$ [36]:

Definition 3.2 (simplex). Let \mathcal{P} be a set of $k+1$ ($k \leq 3$) point offset pairs (p_i, ζ_i) in $\mathcal{D}_{\mathbf{c}} \times \mathbb{Z}^3$, $0 \leq i \leq k$. Let $\text{Ch}(\mathcal{P})$ denote the convex hull of $\varphi_{\mathbf{c}}(\mathcal{P}) = \{p_i + \mathbf{c} * \zeta_i \mid 0 \leq i \leq k\}$ in \mathbb{R}^3 . If the restriction $\pi|_{\text{Ch}(\mathcal{P})}$ of π to the convex hull of \mathcal{P} is a homeomorphism, the image of $\text{Ch}(\mathcal{P})$ by π is called a k -simplex in $\mathbb{T}_{\mathbf{c}}^3$.

In other words, the image under π of a simplex in \mathbb{R}^3 is a simplex in $\mathbb{T}_{\mathbf{c}}^3$ only if it does not self-intersect or touch. Figure 2 shows the convex hulls A and B of three point-offset pairs in $[0, 1]^2 \times \mathbb{Z}^2$; $(p_1, \binom{0}{2})$ is a representative of the equivalence class of a vertex of A that lies *inside* A .

There are infinitely many sets of point-offset pairs specifying the same simplex. The definition of face and coface is adapted accordingly: Let σ be a k -simplex defined by a set $\mathcal{P}_{\sigma} \subseteq \mathcal{D}_{\mathbf{c}} \times \mathbb{Z}^3$. A simplex τ defined by a set $\mathcal{P}_{\tau} \subseteq \mathcal{D}_{\mathbf{c}} \times \mathbb{Z}^3$ is a *face* of σ and has σ as a *coface* if and only if there is some $\zeta \in \mathbb{Z}^3$ such that $\{(p_i, \zeta_i + \zeta) \mid (p_i, \zeta_i) \in \mathcal{P}_{\tau}\} \subseteq \mathcal{P}_{\sigma}$.

4 Delaunay triangulations in $\mathbb{T}_{\mathbf{c}}^3$

This section is organized as follows: At first we give a definition of the Delaunay triangulation in $\mathbb{T}_{\mathbf{c}}^3$. We observe that there are point sets in $\mathbb{T}_{\mathbf{c}}^3$ whose Delaunay triangulation is in fact not defined. The second part elaborates on this question, finally giving a criterion to decide whether or not a point set has a Delaunay triangulation in $\mathbb{T}_{\mathbf{c}}^3$. In the last part we discuss how to deal with point sets that do not have a Delaunay triangulation in $\mathbb{T}_{\mathbf{c}}^3$.

4.1 Definition

Let us recall that a triangulation of a point set \mathcal{S} in \mathbb{R}^3 is a *Delaunay* triangulation *iff* each tetrahedron satisfies the Delaunay property, i.e. its circumscribing ball does not contain any point of \mathcal{S} in its interior [3, 9].

If the point set is not degenerate, i.e. if no five points of \mathcal{S} are cospherical, then its Delaunay triangulation is uniquely defined. Still, even for degenerate point sets, it is possible to specify a unique Delaunay triangulation, using a symbolic perturbation [13]. In the sequel we always assume Delaunay triangulations in \mathbb{R}^3 to be uniquely defined in that way (see Lemma 4.3).

Let \mathcal{S} now denote a finite point set in $\mathcal{D}_{\mathbf{c}}$. We want to define the Delaunay triangulation of $\pi(\mathcal{S})$ in $\mathbb{T}_{\mathbf{c}}^3$. The idea is to use the projection under π of a Delaunay triangulation of the infinite periodic point set $\boxed{\mathcal{S}^{\mathbf{c}} := \varphi_{\mathbf{c}}(\mathcal{S} \times \mathbb{Z}^3)}$ in \mathbb{R}^3 .

From now on, (i), (ii) and (iii) always denote the three conditions of Definition 2.1 (Simplicial complex).

Lemma 4.1. *For any finite point set $\mathcal{S} \subset \mathcal{D}_{\mathbf{c}}$, a set of simplices \mathcal{K} in \mathbb{R}^3 that fulfills (i), (ii) and the Delaunay property with respect to $\mathcal{S}^{\mathbf{c}}$ is a simplicial complex in \mathbb{R}^3 .*




Figure 2: (2D case) $\pi(A)$ is not a simplex; however, $\pi(B)$ is a simplex.

Proof. We need to show that \mathcal{K} has the local finiteness property (iii).

Assume there is a vertex v with an infinite number of incident simplices and thus an infinite number of incident edges. Since \mathcal{S} contains only a finite number of points, there must be at least one point q in \mathcal{S} of which infinitely many periodic copies are adjacent to v . The periodic copies of q form a grid, in which the diameter of the largest empty ball is bounded by $\|c\|$. So circumscribing balls of tetrahedra that are cofaces of edges between v and periodic copies of q that are further away than $\|c\|$ cannot be empty. This is a contradiction and hence all vertices in \mathcal{K} are incident to only a finite number of simplices.

Let now consider a point p in \mathbb{R}^3 that is not a vertex in \mathcal{K} . If it lies in the interior of a tetrahedron, then it has a neighborhood that intersects only one simplex. If it lies in the interior of a triangle, then it has a neighborhood that intersects three simplices: the triangle and the two incident tetrahedra. Assume now p lies in the interior of an edge, then it has a neighborhood intersecting only the triangles and tetrahedra that are incident to the two endpoints of the edge. According to the above discussion, these are only finitely many. \square

Since \mathcal{S}^c contains points on an infinite grid, any point $p \in \mathbb{R}^3$ is contained in some simplex defined by points in \mathcal{S}^c . Together with Lemma 4.1, this implies that the set of all simplices with points of \mathcal{S}^c as vertices and respecting the Delaunay property is a Delaunay triangulation of \mathbb{R}^3 and we denote it by $DT_{\mathbb{R}}(\mathcal{S}^c)$. Since $|DT_{\mathbb{R}}(\mathcal{S}^c)|$ is homeomorphic to \mathbb{R}^3 and π is surjective, then $\pi(|DT_{\mathbb{R}}(\mathcal{S}^c)|)$ is homeomorphic to $\mathbb{T}_{\mathcal{C}}^3$. So, if $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a simplicial complex, it is also a triangulation of $\mathbb{T}_{\mathcal{C}}^3$. We can now define a Delaunay triangulation in $\mathbb{T}_{\mathcal{C}}^3$:

Definition 4.2. Let $DT_{\mathbb{R}}(\mathcal{S}^c)$ be the Delaunay triangulation of \mathcal{S}^c in \mathbb{R}^3 . If $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a simplicial complex in $\mathbb{T}_{\mathcal{C}}^3$, then we call it the *Delaunay triangulation* of \mathcal{S} in $\mathbb{T}_{\mathcal{C}}^3$ and denote it by $DT_{\mathbb{T}}(\mathcal{S})$.

4.2 Existence criterion

In this section, we show that Definition 4.2 actually makes sense: At first, we verify that the simplices “match” under π , i.e. that all periodic copies of a simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ are mapped onto the same simplex in $\mathbb{T}_{\mathcal{C}}^3$ under π . We prove as well that if $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a set of simplices, then it fulfills conditions (i) and (iii).

Let us start with the first lemma:

Lemma 4.3. *If the restriction of π to any simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ is a homeomorphism, then $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a set of internally disjoint simplices in $\mathbb{T}_{\mathcal{C}}^3$ that do not contain any point of $\pi(\mathcal{S})$ in their interior.*

Proof. Consider a tetrahedron σ of $DT_{\mathbb{R}}(\mathcal{S}^c)$, whose vertices are a four-tuple of points $\mathcal{P}_{\sigma} \subset \mathcal{S}^c$. σ satisfies the Delaunay property, so all periodic copies $\varphi_c(\mathcal{P}_{\sigma} \times \mathbb{Z}^3)$ also have an empty circumscribing ball. This shows that all these periodic copies form tetrahedra of $DT_{\mathbb{R}}(\mathcal{S}^c)$.

Note that this is even true in degenerate cases: If we handle degeneracies as in [13], then the Delaunay triangulation of a set of cospherical points only depends on their lexicographic order. As translating the set of points does not change their lexicographic order, all periodic copies of that point set are triangulated in the same way.

Followingly, π collapses precisely all the periodic copies of σ onto its equivalence class in $\mathbb{T}_{\mathcal{C}}^3$. As any lower-dimensional simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ is incident to some tetrahedron, 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 $DT_{\mathbb{R}}(\mathcal{S}^c)$ 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 $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is finite: $DT_{\mathbb{R}}(\mathcal{S}^c)$ is locally finite (Lemma 4.1), i.e. the star of any vertex is finite. As \mathcal{S} is discrete also \mathcal{S}^c is discrete and all tetrahedra have a certain volume larger than some constant. Followingly there are only finitely many tetrahedra necessary to fill the original domain \mathcal{D}_c and thus $\mathbb{T}_{\mathcal{C}}^3$. Finitely many tetrahedra have only finitely many faces so the overall number of simplices in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is finite as well.

So far we know that if all simplices in $DT_{\mathbb{R}}(\mathcal{S}^c)$ are mapped as simplices onto $\mathbb{T}_{\mathcal{C}}^3$, then the whole triangulation is mapped onto a set of simplices in $\mathbb{T}_{\mathcal{C}}^3$. We now consider the incidence relation.

Observation 4.4. *Assume the restriction of π to any simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ is a homeomorphism. If τ is a simplex in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ and $\tau' \leq \tau$, then τ' is a simplex in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. This follows immediately from the fact that incidence relations are maintained by π and from Lemma 4.3.*

It only remains to show condition (ii), i.e. the intersection of two simplices σ and τ in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is another simplex χ that is incident to both σ and τ .

Lemma 4.5. *Assume the restriction of π to any simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ is a homeomorphism. Let $\sigma, \tau \in \pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ be any two simplices in $\mathbb{T}_{\mathcal{C}}^3$, then $\sigma \cap \tau$ is a set of simplices in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$.*

Proof. W.l.o.g. 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 $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. The union is finite because there are only finitely many simplices in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. Consider a point $p \in \sigma \cap \tau$. If p is a vertex of $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$, then it is not contained in the interior of any other simplex, according to Lemma 4.3, and we set $\chi_p = \{p\}$. If p is not a vertex in $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$, then $p \in \sigma'$ and $p \in \tau'$ for some proper faces $\sigma' \leq \sigma$ and $\tau' \leq \tau$ because σ and τ are internally disjoint (Lemma 4.3). 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 $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. \square

We can now formulate the following sufficient condition for $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ to be a simplicial complex:

Theorem 1. *If for all vertices v of $DT_{\mathbb{R}}(\mathcal{S}^c)$ the restriction of the quotient map $\pi|_{|\text{St}(v)|}$ is a homeomorphism, then $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ forms a simplicial complex.*

Proof. We set $\mathcal{K} = \pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. Let σ be a simplex of $DT_{\mathbb{R}}(\mathcal{S}^c)$ and v an incident vertex. Then $\sigma \subseteq |\text{St}(v)|$, thus the restriction of $\pi|_{|\text{St}(v)|}$ to σ is a homeomorphism as well, and \mathcal{K} is a set of simplices (Lemma 4.3).

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 $\mathcal{D}_c \times \mathbb{Z}^3$ such that $\sigma = \pi(\text{Ch}(\mathcal{P}_\sigma))$ and $\tau = \pi(\text{Ch}(\mathcal{P}_\tau))$. From Lemma 4.5, 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 a homeomorphism, 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 a homeomorphism. So from Definition 3.2, 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

In the following theorem we give another criterion that is algorithmically easier to check. Let us recall that the *1-skeleton* of a simplicial complex is the subcomplex that consists of all edges and vertices.

Theorem 2. *Assume the restriction of π to any simplex in $DT_{\mathbb{R}}(\mathcal{S}^c)$ is a homeomorphism. If the 1-skeleton of $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ does not contain any cycle of length ≤ 2 , then $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ forms a simplicial complex.*

Proof. We set $\mathcal{K} = \pi(DT_{\mathbb{R}}(\mathcal{S}^c))$. From the above discussion 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 Theorem 1 there is a vertex $v \in \mathcal{K}$ for which $\pi|_{|\text{St}(v)|}$ is not a homeomorphism. 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 $\sigma'_R \in \pi^{-1}(\sigma)$ and $\sigma''_R \in \pi^{-1}(\sigma)$ containing p and q , respectively. Thus σ'_R and σ''_R are both elements of $\overline{\text{St}(v)}$. Let u, w be vertices different from v with $u \leq \sigma'_R$ and $w \leq \sigma''_R$. The vertices u, w are as well elements of $\overline{\text{St}(v)}$ and thus there are edges (u, v) and (v, w) in $DT_{\mathbb{R}}(\mathcal{S}^c)$. Since $\pi(u) = \pi(w)$, the projection of these two edges under π forms a cycle of length two in $\mathbb{T}_{\mathbf{c}}^3$. \square




Figure 3: (2D case) The shaded region is $\varphi_c(\text{St}(p) \times \mathbb{Z}^3) \cap \mathcal{D}_c$. There are several cycles of length two originating from p .

See Figure 3 for an illustration of Theorems 1 and 2.

4.3 Point sets without a Delaunay triangulation in $\mathbb{T}_{\mathbf{c}}^3$

In this section, we explain how we can give a finite representation of the periodic triangulation $DT_{\mathbb{R}}(\mathcal{S}^c)$ that is a simplicial complex, even if $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ is not a simplicial complex.

Definition 4.6. [2] Let \mathbb{X} be a topological space. A map $\rho : \tilde{\mathbb{X}} \rightarrow \mathbb{X}$ is called a *covering map* and $\tilde{\mathbb{X}}$ is said to be a *covering space* of \mathbb{X} if the following condition holds: For each point $x \in \mathbb{X}$ there is an open neighborhood V , and a decomposition of $\rho^{-1}(V)$ as a family $\{U_\alpha\}$ of pairwise disjoint open subsets of $\tilde{\mathbb{X}}$, in such a way that $\rho|_{U_\alpha}$ a homeomorphism for each α . Let h_x denote the cardinality of the family $\{U_\alpha\}$ corresponding to some $x \in \mathbb{X}$. If the maximum $h := \max_{x \in \mathbb{X}} h_x$ is finite, then $\tilde{\mathbb{X}}$ is called an *h -sheeted covering space*.

\mathbb{R}^3 with the quotient map π as covering map is a *universal covering* of $\mathbb{T}_{\mathbf{c}}^3$, which means that it is a covering space for all covering spaces of $\mathbb{T}_{\mathbf{c}}^3$ [2].

Let $\mathbf{h} = (h_x, h_y, h_z) \in \mathbb{N}^3$. $\mathbb{T}_{\mathbf{h}*\mathbf{c}}^3$ is a covering space of $\mathbb{T}_{\mathbf{c}}^3$ together with the covering map $\rho_{\mathbf{h}} := \pi \circ \pi_{\mathbf{h}}^{-1}$, where $\pi_{\mathbf{h}} : \mathbb{R}^3 \rightarrow \mathbb{T}_{\mathbf{h}*\mathbf{c}}^3$ denotes the quotient map of $\mathbb{T}_{\mathbf{h}*\mathbf{c}}^3$. As $\rho_{\mathbf{h}}^{-1}(p)$ for any $p \in \mathbb{T}_{\mathbf{c}}^3$ consists of $h_x \cdot h_y \cdot h_z$ different points, $\mathbb{T}_{\mathbf{h}*\mathbf{c}}^3$ is a $h_x \cdot h_y \cdot h_z$ -sheeted covering space. The original domain is $D_{\mathbf{h}*\mathbf{c}} = [0, h_x c_x) \times [0, h_y c_y) \times [0, h_z c_z)$. If $h_x = h_y = h_z$ we use the notation $\pi_h := \pi_{\mathbf{h}}$ with $h := h_x \cdot h_y \cdot h_z$, like for π_{27} in Theorem 3 below.

Dobbin and Huson [15] showed that only the points of \mathcal{S}^c contained in \mathcal{D}_c and the 26 copies that surround it can have an influence on the Delaunay property for simplices that are completely contained in \mathcal{D}_c . The ideas of their proof could be used to show the following:

Theorem 3. $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$ is always a simplicial complex.

We prefer to use the framework of covering spaces, rather than just talk about copies of the points as in [15], for several reasons: A major part of the code can be reused for any finite covering space. Also, the simplicial complex we compute is actually homeomorphic to \mathbb{T}_c^3 . So we do not have any artificial boundaries in the data structure and we get all adjacency relations between simplices.

The algorithm we use to compute triangulations of \mathbb{T}_c^3 requires a slightly stronger result, which we present in the next section.

5 Algorithm

As mentioned in the introduction, there is a strong motivation for reusing the standard incremental algorithm [4] to compute a periodic Delaunay triangulation. We propose the following algorithm:

- We start computing in some finitely-sheeted covering space $\mathbb{T}_{h \times c}^3$ of \mathbb{T}_c^3 , with h chosen such that $\pi_h(DT_{\mathbb{R}}(\mathcal{S}^c))$ is guaranteed to be a triangulation.
- In practice, if the point set is large and reasonably well distributed, it is quite likely that after having inserted all the points of a subset $\mathcal{S}' \subset \mathcal{S}$, all the subsequent $\pi(DT_{\mathbb{R}}(\mathcal{S}''^c))$ for $\mathcal{S}' \subset \mathcal{S}'' \subset \mathcal{S}$ are actually simplicial complexes in \mathbb{T}_c^3 . In this case, we discard all periodic copies of simplices of $\pi_h(DT_{\mathbb{R}}(\mathcal{S}'^c))$ and switch to computing $\pi(DT_{\mathbb{R}}(\mathcal{S}^c))$ in \mathbb{T}_c^3 by adding all the points left in $\mathcal{S} \setminus \mathcal{S}'$.

In this way, unlike [15], we avoid duplicating points as soon as this is possible. However, if \mathcal{S} is a small and/or badly distributed point set, the algorithm never enters the second phase and returns $\pi_h(DT_{\mathbb{R}}(\mathcal{S}^c))$.

Note that, before switching to computing in \mathbb{T}_c^3 , it is not sufficient to test whether $\pi(DT_{\mathbb{R}}(\mathcal{S}'^c))$ is a simplicial complex. Indeed, adding a point could create a cycle of length two (see Figure 4). So, a stronger condition is needed before the switch.

The following observation will be useful in the subsequent proofs:

Observation 5.1. Let Δ denote a tetrahedron in $DT_{\mathbb{R}}(\mathcal{S}^c)$ and \mathcal{B}_{Δ} its circumscribing ball. If the diameter of \mathcal{B}_{Δ} is smaller than $c_{\min} := \min\{c_x, c_y, c_z\}$, then $\pi|_{\mathcal{B}_{\Delta}}$ is a homeomorphism, and $\pi|_{\Delta}$ is a homeomorphism as well since $\Delta \subset \mathcal{B}_{\Delta}$. Thus Δ is a simplex in \mathbb{T}_c^3 .




Figure 4: (2D case) Adding a point in a simplicial complex can create a cycle of length two.

Cubic domain. If the original domain \mathcal{D}_c is a cube with edge length c , the incremental algorithm uses the 27-sheeted covering space \mathbb{T}_{3c}^3 , where $3c = (3c, 3c, 3c)$. The original domain for this covering space is \mathcal{D}_{3c} .

We prove below that once all edges in the triangulation in \mathbb{T}_{3c}^3 are shorter than $\frac{1}{\sqrt{6}}c$, computing a simplicial complex in \mathbb{T}_c^3 is possible.

See Algorithm 1 for a pseudo-code listing of the algorithm.

Algorithm 1 Compute Delaunay triangulation of a point set in \mathbb{T}_c^3

Input: Set \mathcal{S} of points in \mathcal{D}_c , c such that \mathcal{D}_c is a cube with edge length $c \in \mathbb{R}^3 \setminus \{0\}$.

Output: $DT_{\mathbb{T}}(\mathcal{S})$ if possible, otherwise $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$

```

1:  $\mathcal{S}' \leftarrow \mathcal{S}$ 
2: Pop  $p$  from  $\mathcal{S}'$ 
3:  $\mathcal{S} \leftarrow \{p\}$ 
4:  $TR_{27} \leftarrow \pi_{27}(DT_{\mathbb{R}}(\varphi_c(\{p\} \times \mathbb{Z}^3)))$  // can be precomputed
5: while the longest edge in  $TR_{27}$  is longer than  $\frac{1}{\sqrt{6}}c$  do
6:   Pop  $p$  from  $\mathcal{S}'$ ;  $\mathcal{S} \leftarrow \mathcal{S} \cup \{p\}$ 
7:   for all  $p' \in \{p + c * \zeta \mid \zeta \in \{0, 1, 2\}^3\}$  do
8:     Insert  $p'$  into  $TR_{27}$ 
9:   end for //  $TR_{27} = \pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$ 
10:  if  $\mathcal{S}' = \emptyset$  then return  $TR_{27} = \pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$  // non-triangulable point set
11: end while
12: Compute  $DT_{\mathbb{T}}(\mathcal{S})$  from  $TR_{27}$  // switch to  $\mathbb{T}_c^3$ 
13: Insert all points remaining in  $\mathcal{S}'$  into  $DT_{\mathbb{T}}(\mathcal{S})$  one by one
14: return  $DT_{\mathbb{T}}(\mathcal{S})$ 

```

Two central points must be established to prove the correctness of the algorithm:

1. After each insertion, TR_{27} is a Delaunay triangulation in \mathbb{T}_{3c}^3 . Let us emphasize on the fact that Theorem 3 cannot be used here because in the inner loop (step 8), the set of points present in TR_{27} does not contain all the periodic copies of p .
Let p be a point in \mathcal{D}_c and $\mathcal{T}_p \subseteq \varphi_c(\{p\} \times \mathbb{Z}^3) \cap \mathcal{D}_{3c}$, i.e. \mathcal{T}_p is a subset of the grid of 27 copies of p that lie within \mathcal{D}_{3c} . Then TR_{27} is always of the form $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c \cup \mathcal{T}_p^{3c}))$ with $\mathcal{T}_p^{3c} = \varphi_{3c}(\mathcal{T}_p \times \mathbb{Z}^3)$. Lemma 5.2 shows that this is a triangulation.
2. If all edges in $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$ are shorter than $\frac{1}{\sqrt{6}}c$, then we can switch to computing in \mathbb{T}_c^3 . See Lemma 5.4 for a proof.

Lemma 5.2. *Let $\mathcal{S} \subset \mathcal{D}_c$ be a finite point set and $p \in \mathcal{D}_c$ a point. If \mathcal{D}_c is a cube, then $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c \cup \mathcal{T}_p^{3c}))$ is a triangulation.*

Proof. We first consider a triangulation of only one point, i.e. $\mathcal{S} = \{p\}$ for some $p \in \mathcal{D}_c$. Then $\pi_{27}(\mathcal{S}^c)$ consists of 27 points arranged as a regular grid. W.l.o.g. we can assume $p = (0, 0, 0)$ because triangulations of point sets are invariant with respect to translations. The point set $\pi_{27}(\mathcal{S}^c)$ is highly degenerate: There are 27 different empty balls that all have eight points on their boundary. They are centered at $((\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) + \zeta) * c$ with $\zeta \in \{0, 1, 2\}^3$ and have radius $r = \frac{\sqrt{3}}{2}c \approx 0.866c$. W.l.o.g. we now consider the ball \mathcal{B} centered at $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. It is easy to check that the intersection of the image of any pair of balls under π_{27} consists of one connected set, i.e. that π_{27} restricted to the union of any pair of balls is a homeomorphism.

As the length of an edge of the cube \mathcal{D}_{3c} is $3c$ and the diameter of \mathcal{B} is $\sqrt{3}c$, an edge of length $(3 - \sqrt{3})c$ would be enough to introduce a cycle of length two. One largest empty ball containing such an edge would be defined by the points $(2, 0, 0)$, $(2, 0, 1)$, $(2, 1, 0)$, $(2, 1, 1)$ and $(\frac{1+\sqrt{3}}{2}, \frac{1}{2}, \frac{1}{2})$ and thus its radius would be $0.712c$. There exist other such balls along the y - and z -axis but since \mathcal{D}_c is a cube, their radii are the same. Then the maximum overlap between \mathcal{B} and this ball would be $< (2 \cdot 0.712 + \sqrt{3} - 3)c < 0.155c$. The smallest empty ball would have radius $\frac{3-\sqrt{3}}{2}c > 0.633c$. Let \mathcal{B}_2 denote such a ball. Then $\pi_{27}|_{\mathcal{B} \cup \mathcal{B}_2}$ is not a homeomorphism and $\pi_{27}(\mathcal{B} \cap \mathcal{B}_2)$ consists of two disjoint connected components.

Now we have to exploit a special property of Algorithm 1: $\pi_{27}(\mathcal{S}^c)$ consists of a regular orthogonal grid of 27 points for each point in \mathcal{S} and $\pi_{27}(\mathcal{T}_p^{3c})$ is a subset of the grid of 27 copies of p . From this we know that after adding a point $p_1 \in \mathcal{T}_p^{3c}$, the next point will differ by some vector in $c * \mathbb{Z}^3$. However, to be able to form a cycle of length two crossing the balls \mathcal{B}_1 and \mathcal{B}_2 its length along one axis would have to be between $1.266c$ and $1.424c$. This is not possible because $[1.266, 1.424] \cap \mathbb{Z} = \emptyset$.

There is no further possibility for a cycle of length two to occur because the shortest diagonals through \mathcal{D}_{3c} have length $\sqrt{18}c$ which is larger than $2\sqrt{3}c$.

Using Theorem 2 this proves that $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a triangulation for \mathcal{S} consisting of only one point. If \mathcal{S} contains more points, then the empty balls can only be smaller, which even more avoids cycles of length two. \square

As this previous proof is quite technical and involved we give a much simpler proof for a slightly weaker result here:

Lemma 5.3. *Let $\mathcal{S} \subset \mathcal{D}_c$ and $\mathcal{T} \in \mathcal{D}_{4c}$ be finite point sets. If \mathcal{D}_c is a cube, then $\pi_{64}(DT_{\mathbb{R}}(\mathcal{S}^c \cup \mathcal{T}^{4c}))$ is a triangulation in \mathbb{T}_{4c}^3 .*

Proof. We first consider the case that $\mathcal{S} = \{p\}$ for some $p \in \mathcal{D}_c$. Then $\pi_{64}(\mathcal{S}^c)$ is a set of 64 points arranged in a regular grid. Thus the largest empty ball has diameter $d = \|c\| = \sqrt{3}c$. As edges are contained in empty balls, the edge length is bounded by $d < 2c$. The cube \mathcal{D}_{4c} that forms the original domain of \mathbb{T}_{4c}^3 has edge length $4c$. Since all edges in $\pi_{64}(DT_{\mathbb{R}}(\mathcal{S}^c))$ are shorter than $2c$, no cycle of length two can occur in the 1-skeleton of $\pi_{64}(DT_{\mathbb{R}}(\mathcal{S}^c))$. From Observation 5.1 and Theorem 2, $\pi_{64}(DT_{\mathbb{R}}(\mathcal{S}^c))$ is a simplicial complex. If we add more points, the diameter of the largest empty ball cannot become larger. The claim follows. \square

Note that unlike Lemma 5.2, Lemma 5.3 does not require \mathcal{T} to be a subset of a point grid $\varphi_c(\{p\} \times \mathbb{Z}^3)$ for some point $p \in \mathcal{D}_c$.

We now prove the edge-length criterion.

Lemma 5.4. *If the 1-skeleton of $DT_{\mathbb{R}}(\mathcal{S}^c)$ contains only edges shorter than $\frac{1}{\sqrt{6}}c$, where c is the edge length of \mathcal{D}_c , then $\pi(DT_{\mathbb{R}}(\mathcal{T}^c))$ is a simplicial complex for any finite $\mathcal{T} \subset \mathcal{D}_c$ with $\mathcal{S} \subseteq \mathcal{T}$.*

Proof. Assume there is a ball \mathcal{B} of diameter d that does not contain any point of \mathcal{S}^c in its interior. Consider the tetrahedron Δ in $DT_{\mathbb{R}}(\mathcal{S}^c)$ that contains the center of \mathcal{B} . The length of the largest edge of Δ is bounded from below by the edge length of the regular tetrahedron with circumscribing ball \mathcal{B} , which is $\frac{2d}{\sqrt{6}}$. So if all edges in $DT_{\mathbb{R}}(\mathcal{S}^c)$ 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 with the same argumentation as in the proof of Lemma 5.3. \square

Note that Lemma 5.4 only gives a sufficient criterion. There are point sets with maximum empty ball diameter shorter than $\frac{1}{2}c$ but edges longer than $\frac{1}{\sqrt{6}}c$. However, it is much easier in practice to test for the longest edge length rather than for the largest empty ball diameter.

Lemmas 5.2 and 5.4 prove the correctness of Algorithm 1 in the case of a cubic domain.

Non-cubic domain. The above discussion still remains valid if the original domain \mathcal{D}_c is a general cuboid, i.e. $\mathbf{c} = (c_x, c_y, c_z)$. Only the constants, like the number of sheets of the covering space to start with and the edge length threshold need to be adapted.

In the proof of Lemma 5.3, the necessary condition is that the diameter of an empty sphere is always half the length of the shortest edge of the original domain. The diameter of the largest empty ball in a periodic triangulation of one point is given by $\|\mathbf{c}\|$. Consider $D_{\mathbf{h}*\mathbf{c}}$; The condition boils down to $c_i h_i > 2\|\mathbf{c}\|$ for $i = x, y, z$. So, the covering space to start with is $T_{\mathbf{h}*\mathbf{c}}^3$ with $h_i = \left\lceil 2 \frac{\|\mathbf{c}\|}{c_i} \right\rceil$.

In order to switch to the 1-sheeted covering space \mathbb{T}_c^3 , we need to make sure that the largest empty balls have diameter smaller than $\frac{1}{2}c_{\min}$, where $c_{\min} = \min\{c_x, c_y, c_z\}$. The threshold in Lemma 5.4 simply becomes $\frac{1}{\sqrt{6}}c_{\min}$.

Weighted Delaunay triangulation

The weighted Delaunay triangulation, or regular triangulation, generalizes the Delaunay triangulation when the sites are spheres, also called *weighted points* (see for instance [17]).

Like the Delaunay triangulation, the regular triangulation can be defined uniquely, even in degenerate cases [14]. Here, we show that the whole discussion in Section 4 works in the same way for regular triangulations. We also prove an edge-length criterion similar to Lemma 5.4, which allows to generalize Algorithm 1 to compute a periodic regular triangulation as well.

Let \mathcal{S} be a set of weighted points in \mathbb{R}^3 , i.e. pairs $p^w = (p, w_p) \in \mathbb{R}^3 \times \mathbb{R}$. p^w can also be seen as a sphere of center p and radius $\sqrt{w_p}$. The *power product* between p^w and q^w is defined as $\Pi(p^w, q^w) = \|p - q\|^2 - w_p - w_q$. p^w and q^w are said to be *orthogonal* if $\Pi(p^w, q^w) = 0$, i.e. the corresponding spheres intersect orthogonally. In a non-degenerate case, a four-tuple of weighted points admit exactly one common orthogonal weighted point. Let $o^w = (o, w_o)$ denote such an orthogonal weighted point. The four-tuple of points is said to have the *weighted Delaunay property* if for any other point p^w of \mathcal{S} , $\Pi(o^w, p^w) \geq 0$. The weighted Delaunay triangulation, or regular triangulation, is the set of all tetrahedra formed by four-tuples of points having this property. Note that if we change the weights of all points in \mathcal{S} by the same amount, the weighted Delaunay triangulation does not change.

Let \mathcal{D}_c be a cube of edge length c and \mathcal{S} a set of weighted points in \mathcal{D}_c . Let W and w respectively denote the largest and smallest weight in \mathcal{S} . Let $RT_{\mathbb{R}}(\mathcal{S}^c)$ denote the regular triangulation of \mathcal{S}^c in \mathbb{R}^3 .

Lemma 5.5. *If the longest edge of $RT_{\mathbb{R}}(\mathcal{S}^c)$ is smaller than $\sqrt{\frac{1}{6}c^2 - \frac{8}{3}(W - w)}$, then $\pi(RT_{\mathbb{R}}(\mathcal{T}^c))$ is a triangulation in \mathbb{T}_c^3 for any finite $\mathcal{T} \subset \mathcal{D}_c$ with $\mathcal{S} \subseteq \mathcal{T}$.*

Note that if $W = w$, this bound reduces to $\frac{1}{\sqrt{6}}c$, which is the edge length threshold for Delaunay triangulations.

The proof of this lemma uses the following auxiliary result:

Lemma 5.6. *Consider $RT_{\mathbb{R}}(\mathcal{S}^c)$. If the orthogonal weighted point of a tetrahedron has weight w_o , then the length of the longest edge is not larger than $2\sqrt{w_o + W - w}$.*

Proof. Let $o^w = (o, w_o)$ be an orthogonal weighted point and $(p, w_p), (q, w_q)$ two of the weighted points orthogonal to o^w . Then $\|o - p\| = \sqrt{w_p + w_o}$. The edge length $\|p - q\|$ attains its maximum if the segment $[p, q]$ contains o . Then $\|p - q\| = \sqrt{w_o + w_p} + \sqrt{w_o + w_q}$. This expression is maximized if the weight is maximized. We know that changing the weights of all points in the triangulation by the same amount does not change the triangulation. So we can assume $W - w$ to be the maximum possible weight. The claim follows. See Figure 5 for an illustration in 2D. \square




Figure 5: The longest possible edge of a tetrahedron with orthogonal sphere of radius $\sqrt{w_o}$.

Proof of Lemma 5.5. Let λ be the length of the longest edge in $RT_{\mathbb{R}}(\mathcal{S}^c)$. Then the circumradius of any tetrahedron in $RT_{\mathbb{R}}(\mathcal{S}^c)$ cannot be larger than $\sqrt{\frac{3}{8}}\lambda$. From Lemma 5.6 follows

that the longest possible edge of a tetrahedron is bounded by $\lambda_{max} := 2\sqrt{\left(\sqrt{\frac{3}{8}}\lambda\right)^2 + W - w}$.

If we now choose λ to be smaller than $\sqrt{\frac{1}{6}c^2 - \frac{8}{3}(W - w)}$, then λ_{max} is smaller than $\frac{1}{2}c$. Thus no cycle of length two can occur and $\pi(RT_{\mathbb{R}}(\mathcal{S}^c))$ is a triangulation according to Theorem 2 together with the equivalent of Observation 5.1 for regular triangulations. Adding further points cannot increase the size of the largest circumscribing ball and so $\pi(RT_{\mathbb{R}}(\mathcal{T}^c))$ is a triangulation as well. \square

Corollary 5.7. *Let $\mathcal{T}_p^{\eta c}$ be chosen as in Lemma 5.2 with $\eta \in \mathbb{N}$. If $\eta = \lceil 8 \cdot (3 + \frac{4}{c}(W - w))^{3/2} \rceil$, then for $h = \eta^3$, $\pi_h(RT_{\mathbb{R}}(\mathcal{S}^c \cup \mathcal{T}_p^{\eta c}))$ is a triangulation in $\mathbb{T}_{\eta c}^3$.*

Proof. If \mathcal{S} consists of only one point, then the regular triangulation of \mathcal{S}^c actually is a Delaunay triangulation because all periodic copies of the point have the same weight. So the largest orthogonal sphere coincides with the largest circumsphere and has radius $\frac{\sqrt{3}}{2}c$. With the same argument as in the proof of Lemma 5.5 the diameter of the largest orthogonal sphere after adding any other point is bounded by $2 \cdot \sqrt{\left(\frac{\sqrt{3}}{2}c\right)^2 + W - w}$. To make sure that this expression is smaller than $2 \cdot \eta \cdot c$ we choose $\eta = \lceil 8 \cdot (3 + \frac{4}{c}(W - w))^{3/2} \rceil$. \square

6 The Delaunay Hierarchy in the periodic case

In a nutshell, the structure of [12] 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 some vertices in different levels corresponding to the same input point. The structure allows for fast point location in the complete triangulation. We refer the reader to the original paper for a more complete description of this data structure.

This structure can be adapted to our algorithm. Let \mathbb{T}_{h*c}^3 be the covering space chosen as presented at the beginning of Section 5. Each new level of the hierarchy stores a triangulation

in $\mathbb{T}_{h^*c}^3$ when it is created, and this triangulation is converted when possible to a triangulation in \mathbb{T}_c^3 . Note that, if a given level l stores a triangulation in \mathbb{T}_c^3 , then the next level $l+1$ is also in \mathbb{T}_c^3 , since it contains more points and thus also stores a triangulation in \mathbb{T}_c^3 by Lemma 5.4. However, some level l can store a triangulation in $\mathbb{T}_{h^*c}^3$ while the next level $l+1$ is converted in \mathbb{T}_c^3 . In this case, for all the vertices corresponding to periodic copies of a given input point in that level l , their pointer to the level $l+1$ just all lead to the same vertex corresponding to this non-duplicated point in \mathbb{T}_c^3 .

The randomized analysis of [12] assumes the insertion of points of \mathcal{S} to be performed in a random order. The changes to this analysis, when computing in $\mathbb{T}_{h^*c}^3$, 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 result remains the same.

7 Theoretical and Practical Analysis

Complexity analysis. Let us first discuss the following two points: (1) How to test for the length of the longest edge and (2) how to switch from the triangulation in $\mathbb{T}_{h^*c}^3$ to the triangulation in \mathbb{T}_c^3 .

(1) We maintain an unsorted data structure \mathcal{E} that references all edges that are longer than the threshold $\frac{1}{\sqrt{6}} c_{\min}$. As soon as \mathcal{E} is empty, we know that the longest edge is smaller than the threshold. The total number of edges that are inserted and removed in \mathcal{E} is not larger than the total number of simplices that are created and destroyed during the algorithm. We can have direct access from the simplices to their edges in \mathcal{E} . Hence, the maintenance of \mathcal{E} does not change the algorithm complexity.

(2) To convert the triangulation in $\mathbb{T}_{h^*c}^3$ to $DT_{\mathbb{T}}(\mathcal{S})$ when we switch to \mathbb{T}_c^3 , we need to iterate over all cells and all vertices to delete all periodic copies, keeping only one; furthermore, we need to update the incidence relations of those tetrahedra whose neighbors have been deleted. This is linear in the size of the triangulation and thus dominated by the main loop.

The overall algorithm is incremental and can be implemented with the Delaunay hierarchy [12]. From the discussion in Section 6 and the two points above, we can conclude that:

The randomized complexity of the algorithm of Section 5 is the same as the complexity of [12], and thus it has randomized worst-case optimal complexity $O(n^2)$.

Experimental observations. Algorithm 1 has been implemented in CGAL, so, it benefits from some of the optimizations that are already available in the CGAL Delaunay triangulations in \mathbb{R}^3 [24], such as the spatial sorting [11]. In this paragraph, we only report a few observations that are specific to the periodic case.

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 random. This 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 number of edges and switch to

No. of points	\mathbb{R}^3	\mathbb{T}^3	factor
1000	0.015	0.022	1.47
10000	0.12	0.21	1.75
100000	1.3	2.1	1.61
1000000	15	23	1.53

Table 1: Current running times in seconds on a 2.33 GHz Intel Core 2 Duo processor.

computing in $\mathbb{T}_{\mathcal{C}}^3$. With randomly generated point sets, usually between 400 and 800 points are needed. These numbers can be considered as constant compared to the total number of input points.

The current implementation shows a factor of about 1.5 between the running time of our implementation for computing the Delaunay triangulation in $\mathbb{T}_{\mathcal{C}}^3$ and the running time of computing the Delaunay triangulation in \mathbb{R}^3 with the CGAL package [24] for large random point sets (1 million points). See Table 1.

The package will be submitted to the CGAL Editorial board for inclusion in the next release of the library. We aim at a high interusability with other CGAL components such as the surface mesher and the computation of α -shapes.

8 Conclusion and future work

We proposed an algorithm to compute 3D periodic Delaunay triangulations. The algorithm is guaranteed to produce a correct finite representation of the periodic triangulation for any given point set. We avoid duplications of points whenever possible, and if there is no triangulation for some point set in the flat torus $\mathbb{T}_{\mathcal{C}}^3$, we output a triangulation in a covering space that is homeomorphic to $\mathbb{T}_{\mathcal{C}}^3$. The algorithm has optimal randomized worst case complexity.

Future work will mainly concentrate on two topics: (1) Extend in a similar way some meshing and α -shape algorithms based on Delaunay triangulations so that they can handle periodic data. (2) Extend this work to more general orbifolds. There is ongoing work to unify our results with the results of [1].

Acknowledgments

We are very grateful to Nico Kruithof, who initiated this work and made important contributions when he visited INRIA for a short post-doctoral stay in 2006. We thank Mridul Aanjaneya for fruitful discussions on Section 4.2 and Olivier Devillers for discussions on his Delaunay hierarchy.

Also, we wish to thank Rien van de Weijgaert for providing us with data sets from cosmology research projects to test our implementation.

References

- [1] Mridul Aanjaneya and Monique Teillaud. Triangulating the real projective plane. In *Mathematical Aspects of Computer and Information Sciences*, 2007. Available from: <http://www-spiral.lip6.fr/MACIS2007/Papers/aanjaneya-teillaud-macis07-final.pdf>.
- [2] Mark A. Armstrong. *Basic Topology*. Springer-Verlag, 1982.
- [3] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann. Available from: <http://www.cup.cam.ac.uk/Scripts/webbook.asp?isbn=0521563224>.
- [4] Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24:162–166, 1981.
- [5] Manuel Caroli, Nico Kruithof, and Monique Teillaud. Decoupling the CGAL 3D triangulations from the underlying space. In *Workshop on Algorithm Engineering and Experiments*, pages 101–108, 2008. Available from: <http://www.siam.org/proceedings/alnex/2008/alx08.010carolim.pdf>.
- [6] 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. Available from: <http://www.computational-geometry.org/SoCG-videos/socg08video/>.
- [7] CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.

-
- [8] Siu-Wing Cheng, Tamal K. Dey, and Joshua A. Levine. A practical Delaunay meshing algorithm for a large class of domains. In *Proceedings of the sixteenth International Meshing Roundtable*, pages 477–494, 2007.
- [9] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edition, 2000.
- [10] Gianni de Fabritiis and Peter V. Coveney. Dynamical geometry for multiscale dissipative particle dynamics. <http://xxx.lanl.gov/abs/cond-mat/0301378v1>, 2003.
- [11] Christophe Delage. Spatial sorting. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:SpatialSorting.
- [12] Olivier Devillers. The Delaunay hierarchy. *International Journal of Foundations of Computer Science*, 13:163–180, 2002.
- [13] 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.
- [14] Olivier Devillers and Monique Teillaud. Perturbations and vertex removal in Delaunay and regular 3D triangulations. Research Report 5968, INRIA, 2006. Available from: <https://hal.inria.fr/inria-00090522>.
- [15] Nikolai P. Dolbilin and Daniel H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.
- [16] Daniel Duque Campayo. Sklogwiki - Boundary conditions. http://www.sklogwiki.org/SklogWiki/index.php/Boundary_conditions.
- [17] Herbert Edelsbrunner and Nimish R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15:223–241, 1996.
- [18] Clara I. Grima and Alberto Márquez. *Computational Geometry on Surfaces*. Kluwer Academic Publishers, 2001.
- [19] Martin Held. Vroni: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments. *Computational Geometry: Theory and Applications*, 18:95–123, 2001.
- [20] Susan Hert and Michael Seel. dD convex hulls and Delaunay triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:ConvexHullD.
- [21] John M. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, New York, 2000.
- [22] Marisa Mazón and Tomás Recio. Voronoi diagrams on orbifolds. *Computational Geometry: Theory and Applications*, 8:219–230, 1997.
- [23] Sylvain Pion and Monique Teillaud. 3D triangulation data structure. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:TDS3.
- [24] Sylvain Pion and Monique Teillaud. 3D triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:Triangulation3.
- [25] Qhull. <http://www.qhull.org>.
- [26] Laurent Rineau and Mariette Yvinec. 3D surface mesh generation. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:SurfaceMesher3.
- [27] 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.
- [28] Vanessa Robins. Betti number signatures of homogeneous Poisson point processes. *Physical Review E*, 74(061107), 2006.
- [29] Günter Rote and Gert Vegter. Computational topology: An introduction. In Jean-Daniel Boissonnat and Monique Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, pages 277–312. Springer-Verlag, Mathematics and Visualization, 2006.

-
- [30] Jonathan R. Shewchuk. Triangle: Engineering a 2d quality mesh generator and Delaunay triangulator. In *First Workshop on Applied Computational Geometry*. Association for Computing Machinery, May 1996.
 - [31] Jonathan R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. In *Proceedings of the fourteenth Annual Symposium on Computational Geometry*, pages 86–95. ACM Press New York, NY, USA, 1998.
 - [32] Karsten E. Thompson. Fast and robust Delaunay tessellation in periodic domains. *International Journal for Numerical Methods in Engineering*, 55:1345–1366, 2002.
 - [33] William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, Princeton, New Jersey, 1997.
 - [34] William P. Thurston. *The Geometry and Topology of Three-Manifolds*. www.msri.org, 2002.
 - [35] Dahlia Weiss. How hydrophobic Buckminsterfullerene affects surrounding water structure. INRIA Geometrica Seminar, <http://www-sop.inria.fr/geometrica>, March 2008.
 - [36] P M. H. Wilson. *Curved Spaces*. Cambridge University Press, Cambridge, 2008.
 - [37] Mariette Yvinec. 2D triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.3 edition, 2007. Available from: http://www.cgal.org/Manual/3.3/doc.html/cgal_manual/packages.html#Pkg:Triangulation2.
 - [38] Afra Zomorodian. *Topology for Computing*. Cambridge University Press, Cambridge, 2005.



Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399