Computing 3D Periodic Triangulations

Manuel Caroli, Monique Teillaud

Computing 3D Periodic Triangulations*

Manuel Caroli and Monique Teillaud

INRIA Sophia Antipolis – Méditerranée
{Manuel.Caroli,Monique.Teillaud}@sophia.inria.fr

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 into tetrahedra 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; our algorithmic test detects when such a duplication can be avoided, which is usually possible in practical situations. Even in cases where point duplication is necessary, our algorithm always computes a triangulation that is homeomorpic to the flat torus. To the best of our knowledge, this is the first algorithm of this kind whose output is provably correct. The implementation will be released in CGAL [7].

1 Introduction

Computing Delaunay triangulations of point sets is a well-studied problem in Computational Geometry. Several algorithms as well as implementations [31, 26, 19, 38, 25, 21] 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* [35].

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 that need 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 [36], solid-state chemistry [29], physics of condensed matter [15], fluid dynamics [10], this list being far from exhaustive.

* 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/.

[1] http://www.cgal.org/Events/PeriodicSpacesWorkshop/

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 [23, 17], or with 27 copies in 3D [14], which obviously leads to a huge slow-down. Additionally, their output is a triangulation in $\mathbb{R}^d, d = 2, 3$, of the copies of the points, whereas our approach always outputs triangulations of the flat torus.

In the engineering community, an implementation for computing a periodic Delaunay *"tessellation"* was proposed, avoiding duplications of points [34]. 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. Figure 1 shows a simple partition of the 2D torus that is not a triangulation. However, in practice, input data sets are likely to admit a Delaunay triangulation.

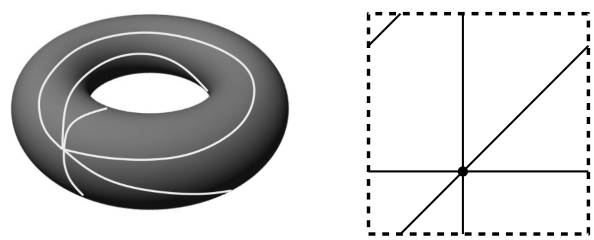

Fig. 1. The partition of the torus (left) and the flat torus (right) is not a triangulation: All simplices have a unique vertex.

Let us insist here on the fact that computing a "true" triangulation, i.e. a *simplicial complex*, is important for several reasons. First, a triangulation is defined as a simplicial complex in the literature [2, 9, 16, 20, 22, 33, 39]. Moreover, designing a data structure to efficiently store tetrahedral tessellations that are non-simplicial complexes (e.g. Δ-complexes [18]) would be quite involved. The CGAL 3D triangulation data structure, that we reuse in our implementation, assumes the structure to be a simplicial complex [24]. Even more importantly, algorithms using a triangulation as input are heavily relying on the fact that the triangulation is a simplicial complex; this is the case for instance for meshing algorithms [27, 28], 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 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.

We present an adaptation of the well-known incremental algorithm in $\mathbb{R}^3$ [3] 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 [25]; 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 [32, 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, where h depends on some parameters of the flat torus, i.e. a triangulation that is still homeomorphic to the flat torus and containing $h > 1$ explicit copies of the input point set. 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.

Our implementation of the algorithm has been accepted for version 3.5 of the CGAL library [7]. We presented a video demonstration of the software [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 in 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 show the correctness of the algorithm and finish with its complexity analysis and experimental observations. Proofs are omitted in this paper due to lack of space. They can be found in [6].

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 $\mathcal{P}_\sigma = \{p_0, p_1, \ldots, p_k\}$. 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 [39, 30]. In the sequel, we deal with infinite simplicial complexes, so, we use the definition given in [22]:

Definition 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 *star* of a subcomplex $\mathcal{L}$ of $\mathcal{K}$ consists of the cofaces of simplices in $\mathcal{L}$: $\mathrm{St}(\mathcal{L}) = \{\sigma \in \mathcal{K} | \sigma \geq \tau \in \mathcal{L}\}$. 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 $|\mathrm{St}(\mathcal{L})|$.

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

At first we give a precise definition of the space of study $\mathbb{T}^3_{\boldsymbol{c}}$. 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 $\mathbb{T}^3_{\boldsymbol{c}}$.

Definition 2 ($\mathbb{T}^3$). *Let $\boldsymbol{c} := (c_x, c_y, c_z) \in (\mathbb{R} \setminus \{0\})^3$ and G be the group $(\boldsymbol{c} * \mathbb{Z}^3, +)$, where $*$ denotes coordinate-wise multiplication*[2]*. The quotient space $\mathbb{T}^3_{\boldsymbol{c}} = \mathbb{R}^3/G$ is called flat torus [35]. We denote the quotient map by* $\boxed{\pi : \mathbb{R}^3 \to \mathbb{T}^3_{\boldsymbol{c}}}$.

The elements of $\mathbb{T}^3_{\boldsymbol{c}}$ are the equivalence classes under the equivalence relation $p_1 \sim p_2 \Leftrightarrow p_1 - p_2 \in \boldsymbol{c} * \mathbb{Z}^3$, for $p_1, p_2 \in \mathbb{R}^3$. Hence, these equivalence classes are isomorphic to $\mathbb{Z}^3$ and $\mathbb{T}^3_{\boldsymbol{c}} \times \mathbb{Z}^3$ is isomorphic to $\mathbb{R}^3$. We also call the points of $\mathbb{T}^3_{\boldsymbol{c}}$ *orbits* and refer to their elements as *representatives*. $\mathbb{T}^3_{\boldsymbol{c}}$ 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 $\mathbb{T}^3_{\boldsymbol{c}}$ is homeomorphic to the hypersurface of a 4-dimensional torus. Consider the closed cuboid $[u, u + c_x] \times [v, v + c_y] \times [w, w + c_z]$. Identifying the pairs of opposite sides results in a space homeomorphic to $\mathbb{T}^3_{\boldsymbol{c}}$. Such a cuboid is usually called a *fundamental domain* or a *fundamental region*. A fundamental domain contains at least one representative of each orbit. The half-open cuboid $\boxed{\mathcal{D}_{\boldsymbol{c}} = [0, c_x) \times [0, c_y) \times [0, c_z)}$ contains exactly one representative for each element of $\mathbb{T}^3_{\boldsymbol{c}}$. We call it the *original domain*. The map

$$\begin{aligned} \varphi_{\boldsymbol{c}} : \mathcal{D}_{\boldsymbol{c}} \times \mathbb{Z}^3 &\to \mathbb{R}^3 \\ (p, \zeta) &\mapsto p + \boldsymbol{c} * \zeta \end{aligned}$$

is bijective. The longest diagonal of $\mathcal{D}_{\boldsymbol{c}}$ has length $\|\boldsymbol{c}\|$, which denotes the L_2-norm of $\boldsymbol{c}$. 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 $p \in \mathcal{D}_{\boldsymbol{c}}$ such that $p_1, p_2 \in \varphi_{\boldsymbol{c}}(\{p\} \times \mathbb{Z}^3)$.

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

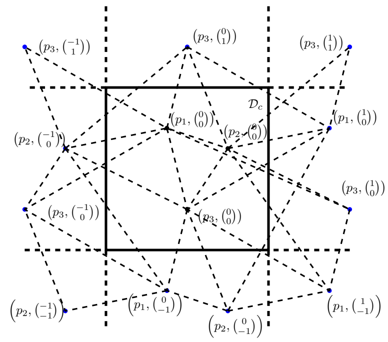

Fig. 2. (2D case) The three points p_1, p_2, and p_3 do not uniquely define a triangle. Intuitively, the offset allows to know which way the triangle "wraps around" the torus.

Now we turn towards the definition of simplices in $\mathbb{T}^3_{\boldsymbol{c}}$. There is no meaningful definition of a convex hull in $\mathbb{T}^3_{\boldsymbol{c}}$ 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 2). In the above definition of $\varphi_{\boldsymbol{c}}$, the offsets are the numbers $\zeta \in \mathbb{Z}^3$. We can adapt the definition of a simplex in $\mathbb{R}^3$ in the following way to $\mathbb{T}^3_{\boldsymbol{c}}$ [37]:

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

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

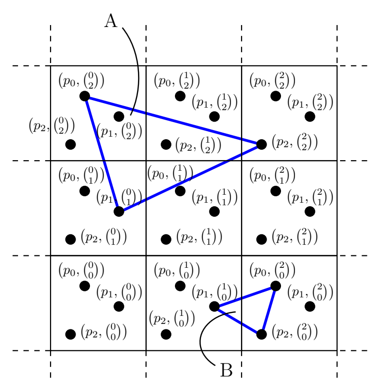

Fig. 3. (2D case) $\pi(A)$ is not a simplex; however, $\pi(B)$ is a simplex.

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 $\mathcal{P}_\sigma \subseteq \mathcal{D}_{\boldsymbol{c}} \times \mathbb{Z}^3$. A simplex τ defined by a set $\mathcal{P}_\tau \subseteq \mathcal{D}_{\boldsymbol{c}} \times \mathbb{Z}^3$ 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 triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$

This section is organized as follows: At first we give a definition of the Delaunay triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$. We observe that there are point sets in $\mathbb{T}^3_{\boldsymbol{c}}$ 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 $\mathbb{T}^3_{\boldsymbol{c}}$. In the last part we discuss how to deal with point sets that do not have a Delaunay triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$.

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

Lemma 1. *For any finite point set $\mathcal{S} \subset \mathcal{D}_{\boldsymbol{c}}$, a set of simplices $\mathcal{K}$ in $\mathbb{R}^3$ that fulfills (i) and (ii) in Definition 1 as well as the Delaunay property with respect to $\mathcal{S}^{\boldsymbol{c}}$ is a simplicial complex in $\mathbb{R}^3$.*

Since $\mathcal{S}^{\boldsymbol{c}}$ contains points on an infinite grid, any point $p \in \mathbb{R}^3$ is contained in some simplex defined by points in $\mathcal{S}^{\boldsymbol{c}}$. Together with Lemma 1, this implies that the set of all simplices with points of $\mathcal{S}^{\boldsymbol{c}}$ as vertices and respecting the Delaunay property is a Delaunay triangulation of $\mathbb{R}^3$ and we denote it by $DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})$. Since $|DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})|$ is homeomorphic to $\mathbb{R}^3$ and π is surjective, then $\pi(|DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})|)$ is homeomorphic to $\mathbb{T}^3_{\boldsymbol{c}}$. So, if $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ is a simplicial complex, it is also a triangulation of $\mathbb{T}^3_{\boldsymbol{c}}$. We can now define a Delaunay triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$:

Definition 4. *Let $DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})$ be the Delaunay triangulation of the point set $\mathcal{S}^{\boldsymbol{c}}$ in $\mathbb{R}^3$. If $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ is a simplicial complex in $\mathbb{T}^3_{\boldsymbol{c}}$, then we call it the* Delaunay triangulation *of $\mathcal{S}$ in $\mathbb{T}^3_{\boldsymbol{c}}$ and denote it by $DT_{\mathbb{T}}(\mathcal{S})$.*

We show now that Definition 4 actually makes sense: Lemma 2 is used to prove Theorem 1, which gives a sufficient condition for $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ to be a simplicial complex.

Lemma 2. *If the restriction of π to any simplex in $DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})$ is a homeomorphism, then conditions (i) and (iii) in Definition 1 are fulfilled.*

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

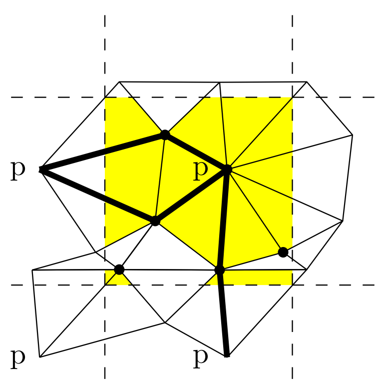

Fig. 4. (2D case) The shaded region is $\varphi_{\boldsymbol{c}}(\mathrm{St}(p) \times \mathbb{Z}^3) \cap \mathcal{D}_{\boldsymbol{c}}$. There are several cycles of length two originating from p.

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* $DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})$ *is a homeomorphism. If the 1-skeleton of* $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ *does not contain any cycle of length less than or equal to two, then* $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ *forms a simplicial complex.*

See Figure 4 for an illustration of Theorems 1 and 2.

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

Definition 5. *[2] Let* $\mathbb{X}$ *be a topological space. A map* $\rho : \widetilde{\mathbb{X}} \to \mathbb{X}$ *is called a* covering map *and* $\widetilde{\mathbb{X}}$ *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* $\widetilde{\mathbb{X}}$*, in such a way that* $\rho|_{U_\alpha}$ *is 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* $\widetilde{\mathbb{X}}$ *is called an* h-sheeted covering space.

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

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

Dolbilin and Huson [14] showed that only the points of $\mathcal{S}^{\boldsymbol{c}}$ contained in $\mathcal{D}_{\boldsymbol{c}}$ and the 26 copies that surround it can have an influence on the Delaunay property for simplices that are completely contained in $\mathcal{D}_{\boldsymbol{c}}$. The ideas of their proof can be used to show the following:

Theorem 3. $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ *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 [14], 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 $\mathbb{T}^3_{\boldsymbol{c}}$. 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 $\mathbb{T}^3_{\boldsymbol{c}}$ 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 [3] to compute a periodic Delaunay triangulation.

We propose the following algorithm:

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

In this way, unlike [14], 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 $\pi_{\boldsymbol{h}}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$. Note that, before switching to computing in $\mathbb{T}^3_{\boldsymbol{c}}$, it is not sufficient to test whether $\pi(DT_{\mathbb{R}}(\mathcal{S}'^{\boldsymbol{c}}))$ is a simplicial complex. Indeed, adding a point could create a cycle of length two (see Figure 5). So, a stronger condition is needed before the switch.

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

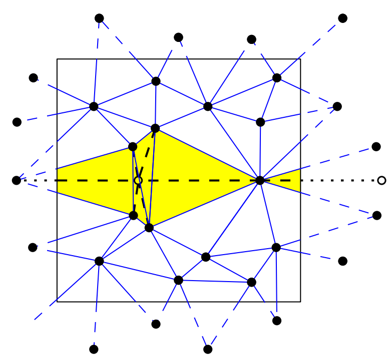

Fig. 5. (2D case) Adding a point in a simplicial complex can create a cycle of length two.

Algorithm 1 Compute Delaunay triangulation of a point set in $\mathbb{T}^3_{\boldsymbol{c}}$

Input: Set $\mathcal{S}$ of points in $\mathcal{D}_{\boldsymbol{c}}$, $\boldsymbol{c}$ such that $\mathcal{D}_{\boldsymbol{c}}$ is a cube with edge length $c \in \mathbb{R}^3 \setminus \{0\}$.
Output: $DT_{\mathbb{T}}(\mathcal{S})$ if possible, otherwise $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$

$\mathcal{S}' \Leftarrow \mathcal{S}$
Pop p from $\mathcal{S}'$
$\mathcal{S} \Leftarrow \{p\}$
$TR_{27} \Leftarrow \pi_{27}(DT_{\mathbb{R}}(\varphi_{\boldsymbol{c}}(\{p\} \times \mathbb{Z}^3)))$ `// can be precomputed`
while the longest edge in TR_{27} is longer than $\frac{1}{\sqrt{6}}c$ **do**
  Pop p from $\mathcal{S}'$; $\mathcal{S} \Leftarrow \mathcal{S} \cup \{p\}$
  for all $p' \in \{p + \boldsymbol{c} * \zeta \mid \zeta \in \{0,1,2\}^3\}$ **do**
    Insert p' into TR_{27}
  end for `//` $TR_{27} = \pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$
  if $\mathcal{S}' = \emptyset$ **then return** $TR_{27} = \pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ `// non-triangulable point set`
end while
Compute $DT_{\mathbb{T}}(\mathcal{S})$ from TR_{27} `// switch to` $\mathbb{T}^3_{\boldsymbol{c}}$
Insert all points remaining in $\mathcal{S}'$ into $DT_{\mathbb{T}}(\mathcal{S})$ one by one
return $DT_{\mathbb{T}}(\mathcal{S})$

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

1. After each insertion, TR_{27} is a Delaunay triangulation in $\mathbb{T}^3_{3\boldsymbol{c}}$. 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 $\mathcal{D}_{\boldsymbol{c}}$ and $\mathcal{T}_p \subseteq \varphi_{\boldsymbol{c}}(\{p\} \times \mathbb{Z}^3) \cap \mathcal{D}_{3\boldsymbol{c}}$, i.e. $\mathcal{T}_p$ is a subset of the grid of 27 copies of p that lie within $\mathcal{D}_{3\boldsymbol{c}}$. Then TR_{27} is always of the form $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}} \cup \mathcal{T}_p^{3\boldsymbol{c}}))$ with $\mathcal{T}_p^{3\boldsymbol{c}} = \varphi_{3\boldsymbol{c}}(\mathcal{T}_p \times \mathbb{Z}^3)$. Lemma 3 shows that this is a triangulation.
2. If all edges in $\pi_{27}(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ are shorter than $\frac{1}{\sqrt{6}}\,c$, then we can switch to computing in $\mathbb{T}^3_{\boldsymbol{c}}$.

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

Lemma 4 gives a criterion to decide whether $\pi(DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}}))$ is a simplicial complex and thus a triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$.

Lemma 4. *If the 1-skeleton of $DT_{\mathbb{R}}(\mathcal{S}^{\boldsymbol{c}})$ contains only edges shorter than $\frac{1}{\sqrt{6}}\,c$, where c is the edge length of $\mathcal{D}_{\boldsymbol{c}}$, then $\pi(DT_{\mathbb{R}}(\mathcal{T}^{\boldsymbol{c}}))$ is a simplicial complex for any finite $\mathcal{T} \subset \mathcal{D}_{\boldsymbol{c}}$ with $\mathcal{S} \subseteq \mathcal{T}$.*

Note that the criterion in Lemma 4 is only sufficient: There are triangulations without cycles of length two that have edges longer than $\frac{1}{\sqrt{6}}\,c$.

Lemmas 3 and 4 prove the correctness of Algorithm 1 in the case of a cubic domain. The above discussion still remains valid if the original domain $\mathcal{D}_{\boldsymbol{c}}$ is a general cuboid, i.e. $\boldsymbol{c} = (c_x, c_y, c_z)$. Only the constants, like the number of sheets of the covering space to start with and the edge length threshold need to be adapted. Analogously, the algorithm can be extended to weighted Delaunay triangulations. For a more detailed discussion see [6].

6 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 $\mathbb{T}^3_{\boldsymbol{h}*\boldsymbol{c}}$ to the triangulation in $\mathbb{T}^3_{\boldsymbol{c}}$.

(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 proportional to 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 $\mathbb{T}^3_{\boldsymbol{h}*\boldsymbol{c}}$ to $DT_{\mathbb{T}}(\mathcal{S})$ when we switch to $\mathbb{T}^3_{\boldsymbol{c}}$, 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 using the Delaunay hierarchy [12] the following result can be shown:

The randomized complexity of Algorithm 1 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 $\mathbb{R}^3$ [25], such as the spatial sorting [11].

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 they are sufficiently well distributed to have triangulations in $\mathbb{T}^3_{\boldsymbol{c}}$. This property holds for most of the applications mentioned in the introduction. With these real data, usually less than 400 points are needed for Algorithm 1 to reach the threshold on the edge length and switch to computing in $\mathbb{T}^3_{\boldsymbol{c}}$.

We compared the running time of our implementation for computing Delaunay triangulations in $\mathbb{T}^3_{\boldsymbol{c}}$ with the running time of computing the Delaunay triangulation in $\mathbb{R}^3$ with the CGAL package [25]. Table 1 shows for large random point sets a factor of about 1.6 between the running time of our current implementation, using the above optimization, and the CGAL implementation for $\mathbb{R}^3$. The timings have been measured for the unit cube $\mathcal{D}_{\boldsymbol{c}} = [0, 1)^3$ using specialized predicates; if we allow $\mathcal{D}_{\boldsymbol{c}}$ to be any cube, we currently lose about 12%. More experiments can be found in [6].

No. of points	$\mathbb{T}^3$	$\mathbb{R}^3$	factor
1000	0.032	0.012	2.65
10000	0.230	0.128	1.79
100000	2.24	1.36	1.65
1000000	23.0	14.2	1.62

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

7 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 $\mathbb{T}^3_{\boldsymbol{c}}$, we output a triangulation in a covering space that is homeomorphic to $\mathbb{T}^3_{\boldsymbol{c}}$. The algorithm has optimal randomized worst case complexity. Note that the main parts of the discussion are not bound to three-dimensional space and will still hold for higher dimensions. The constants in the geometric criteria and the complexity of the underlying algorithm for computing the Delaunay triangulation will have to be adapted.

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 [4]. We thank Mridul Aanjaneya for fruitful discussions on Section 4 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.
2. Mark A. Armstrong. *Basic Topology*. Springer-Verlag, 1982.
3. Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24:162–166, 1981.
4. 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.
5. 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.
6. Manuel Caroli and Monique Teillaud. Computing 3D periodic triangulations. Research Report 6823, INRIA, 2009. `http://hal.inria.fr/inria-00356871`.
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. Robert J. Daverman and Richard B. Sher, editors. *Handbook of Geometric Topology*. Elsevier, Amsterdam, London, Paris, 2002.
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*. 3.4 edition, 2008.
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. Nikolai P. Dolbilin and Daniel H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.
15. Daniel Duque Campayo. Sklogwiki - Boundary conditions. `http://www.sklogwiki.org/SklogWiki/index.php/Boundary_conditions`.
16. Ronald L. Graham, Martin Grötschel, and László Lovász, editors. *Handbook of Combinatorics*. Elsevier, Amsterdam, Lausanne, New York, 1995.

17. Clara I. Grima and Alberto Márquez. *Computational Geometry on Surfaces*. Kluwer Academic Publishers, 2001.
18. Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
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. Michael Henle. *A Combinatorial Introduction to Topology*. Dover publication, New York, 1979.
21. Susan Hert and Michael Seel. dD convex hulls and Delaunay triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.4 edition, 2008.
22. John M. Lee. *Introduction to Topological Manifolds*. Springer-Verlag, New York, 2000.
23. Marisa Mazón and Tomás Recio. Voronoi diagrams on orbifolds. *Computational Geometry: Theory and Applications*, 8:219–230, 1997.
24. Sylvain Pion and Monique Teillaud. 3D triangulation data structure. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.4 edition, 2008.
25. Sylvain Pion and Monique Teillaud. 3D triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.4 edition, 2008.
26. Qhull. `http://www.qhull.org`.
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. Laurent Rineau and Mariette Yvinec. 3D surface mesh generation. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.4 edition, 2008.
29. Vanessa Robins. Betti number signatures of homogeneous Poisson point processes. *Physical Review E*, 74(061107), 2006.
30. 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.
31. 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.
32. 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.
33. Edwin H. Spanier. *Algebraic Topology*. Springer-Verlag, New York, 1966.
34. Karsten E. Thompson. Fast and robust Delaunay tessellation in periodic domains. *International Journal for Numerical Methods in Engineering*, 55:1345–1366, 2002.
35. William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, Princeton, New Jersey, 1997.
36. Dahlia Weiss. How hydrophobic Buckminsterfullerene affects surrounding water structure. INRIA Geometrica Seminar, `http://www-sop.inria.fr/geometrica`, March 2008.
37. P M. H. Wilson. *Curved Spaces*. Cambridge University Press, Cambridge, 2008.
38. Mariette Yvinec. 2D triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual*. 3.4 edition, 2008.
39. Afra Zomorodian. *Topology for Computing*. Cambridge University Press, Cambridge, 2005.