Covering spaces and Delaunay triangulations of the 2D flat torus

Mikhail Bogdanov, Monique Teillaud, Gert Vegter

Covering spaces and Delaunay triangulations of the 2D flat torus

Mikhail Bogdanov* Monique Teillaud† Gert Vegter‡

EuroCG 2012

Abstract

A previous algorithm was computing the Delaunay triangulation of the flat torus, by using a 9-sheeted covering space [5]. We propose a modification of the algorithm using only a 8-sheeted covering space, which allows to work with 8 periodic copies of the input points instead of 9. The main interest of our contribution is not only this result, but most of all the method itself: this new construction of covering spaces generalizes to Delaunay triangulations of surfaces of higher genus.

1 Introduction

The Delaunay triangulation is a widely used structure in Computational Geometry. Delaunay triangulations of the *flat torus* [13] are used in different domains of science, like astronomy [10, 11, 14]. More references can be found in [5, 6, 3].

The Delaunay triangulation of points on a 2D flat torus was previously obtained from a triangulation of the convex hull of nine periodic copies of the points in the plane $\mathbb{E}^2$, laid in a 3x3 pattern [9, 8]. A more recent algorithm directly computes the Delaunay triangulation in a flat torus [5, 4], thus providing all adjacency relations between triangles. The algorithm is incremental, it resorts to a 9-sheeted covering space of the torus, and switches to computing in the initial torus as soon as possible.

After recalling some background on orbit spaces in Section 2, we propose in Section 3 a construction of a family of 2^k sheeted-covering spaces of the flat torus, $k > 0$. In Section 4 we show that the Delaunay triangulation of any finite set of $n > 1$ points in the 8-sheeted covering space of this family is well defined, we give conditions ensuring that the Delaunay triangulation can be defined in the 2- and 4-sheeted covering spaces of the family, and we propose a modified incremental algorithm computing with covering spaces of no more than eight sheets. Section 5 finally gives hints about the extension of the construction to surfaces of higher genus.

2 Background

The flat torus. The unit sphere $\mathcal{S}^1$ can be mapped into $\mathbb{E}^2$ by the map $c : \mathbb{E} \to \mathbb{E}^2$, defined by $c(s) = (\cos s, \sin s)$. Similarly, the torus $\mathbb{T}^2$, being homeomorphic to the Cartesian product $\mathcal{S}^1 \times \mathcal{S}^1$, can be mapped into $\mathbb{E}^4$ by the map $(s,t) \mapsto (\cos s, \sin s, \cos t, \sin t)$. The corestriction of this map to its image yields a map $\pi : \mathbb{E}^2 \to \mathbb{T}^2$. Since this torus, endowed with the ambient metric of $\mathbb{E}^4$, has zero Gaussian curvature, it is called the *flat torus* (as opposed to the 'standard' torus of revolution in $\mathbb{E}^3$, which has regions of positive and negative Gaussian curvature). The periodicity of the maps c and π corresponds to an action of the groups $\mathbb{Z}$ and $\mathbb{Z}^2$, respectively. Such group actions are crucial in our approach, so we first provide the necessary background.

Covering Spaces. In this paper, all topological spaces are assumed to be connected. A *covering map* is a continuous surjective map $\rho : \mathbb{Y} \to \mathbb{X}$ from a topological space $\mathbb{Y}$ to a topological space $\mathbb{X}$, such that each point $x \in \mathbb{X}$ is evenly covered, i.e., there is a an open neighborhood V of x such that $\rho^{-1}(V)$ is the disjoint union of a family $\{U_\alpha\}$ of open subsets of $\mathbb{Y}$ such that $\rho|_{U_\alpha}$ is a homeomorphism for each α. One of our key examples is the map $\pi : \mathbb{E}^2 \to \mathbb{T}^2$ introduced in the previous paragraph. If a point in $\mathbb{X}$ has a finite

*INRIA Sophia Antipolis - Méditerranée, France, Mikhail.Bogdanov@inria.fr
†INRIA Sophia Antipolis - Méditerranée, France, Monique.Teillaud@inria.fr
‡Department of Mathematics and Computing Science, University of Groningen, The Netherlands, G.Vegter@rug.nl

number k of pre-images under the covering map ρ, then the connectedness of $\mathbb{X}$ implies that each point of $\mathbb{X}$ has k pre-images. In this case $\mathbb{Y}$ is called a k-sheeted cover of $\mathbb{X}$. Taking $\mathbb{X} = \mathbb{T}^2$ and $\mathbb{Y} = \mathbb{T}^2$, the map $\sigma : \mathbb{Y} \to \mathbb{X}$, mapping the point $(\cos s, \sin s, \cos t, \sin t) \in \mathbb{Y}$ to the point $(\cos 3s, \sin 3s, \cos 3t, \sin 3t) \in \mathbb{X}$, is an example of a nine-sheeted covering map of the torus (by another torus). If, moreover, $\mathbb{Y}$ is simply connected, i.e., if every closed curve in $\mathbb{Y}$ can be contracted to a point, then $\mathbb{Y}$ is called the *universal cover* of $\mathbb{X}$. It covers all (connected) covers of $\mathbb{X}$ in the sense that for every covering map $\sigma : \mathbb{M} \to \mathbb{X}$ there is a map $f : \mathbb{Y} \to \mathbb{M}$ such that $\pi = \sigma \circ f$. In our earlier example of the nine-sheeted covering map $\sigma : \mathbb{T}^2 \to \mathbb{T}^2$ the map $f : \mathbb{E}^2 \to \mathbb{T}^2$ is given by $f(s,t) = (\cos \frac{1}{3}s, \sin \frac{1}{3}s, \cos \frac{1}{3}t, \sin \frac{1}{3}t)$.

Group Actions and Orbit Spaces. If u and v are orthogonal vectors in $\mathbb{E}^2$ of equal length, then the map $\mathbb{Z}^2 \times \mathbb{E}^2 \to \mathbb{E}^2$, given by $(m,n) * p \mapsto p + mu + nv$, defines an action of the group $\mathbb{Z}^2$ on $\mathbb{E}^2$. In general, an *action* of a group $\mathcal{G}$ on a topological space $\mathbb{X}$ is a map $\mathcal{G} \times \mathbb{X} \to \mathbb{X} : (g,x) \mapsto g * x$ such that
(i) for $g \in \mathcal{G}$, the map $\mathbb{X} \to \mathbb{X} : x \mapsto g * x$, is a homeomorphism on $\mathbb{X}$;
(ii) the identity element $e \in \mathcal{G}$ corresponds to the identity on $\mathbb{X}$, i.e., $e * x = x$, for $x \in \mathbb{X}$;
(iii) $(hg) * x = h * (g * x)$, for $g, h \in \mathcal{G}$ and $x \in \mathbb{X}$.
Each element $g \in \mathcal{G}$ correponds to a homeomorphism $x \mapsto g{*}x$, and this correspondence is an isomorphism according to (ii) and (iii). In our case, $(m,n) \in \mathbb{Z}^2$ corresponds to the translation $a_u^m b_v^n$, where a_u and b_v are the translations over u and v, respectively.
The *orbit* of a point $x \in \mathbb{X}$ under the action of a group $\mathcal{G}$ is the set $\mathcal{G}x = \{g * x \mid g \in \mathcal{G}\}$. Properties (ii) and (iii) imply that the orbits form a partition of $\mathbb{X}$. In other words, the group action induces an equivalence relation on $\mathbb{X}$, given by $x \sim y$ if $x = g * y$ for some $g \in \mathcal{G}$. The *orbit space* $\mathbb{X}/\mathcal{G}$ is the set of all orbits of $\mathbb{X}$ under the action of $\mathcal{G}$. A *fundamental region* is a subset of $\mathbb{X}$ which contains at least one point of each $\mathcal{G}$-orbit with at most one point of each orbit in its interior.
The orbit space of the $\mathbb{Z}^2$-action introduced at the beginning of this paragraph is the torus. Each orbit forms a lattice in $\mathbb{E}^2$. The minimal distance by which $\mathcal{G} = \mathbb{Z}^2$ moves a point of $\mathbb{E}^2$ is $\delta(\mathcal{G}) = \| u \| = \| v \|$, where $\| \cdot \|$ denotes the Euclidean norm. Any closed square of edge length $\delta(\mathcal{G})$ whose edges are parallel to u and v is a fundamental domain of $\mathcal{G}$. For a point $p \in \mathbb{E}^2$, the half-open square $\mathcal{D}(p) = \left\{ p + t_u u + t_v v \mid t_u, t_v \in [-\frac{1}{2}, \frac{1}{2}) \right\}$ contains exactly one representative of each element of $\mathbb{T}^2$. We call $\mathcal{D}(O)$ the *original domain*, where O denotes the origin.

Group Actions and Covering Spaces. In all our examples, the group actions have the property that each $x \in \mathbb{X}$ has a neighborhood in which each point belongs to a different $\mathcal{G}$-orbit. (Technically, this follows from the fact that the action is *discontinuous* and *fixed point free*; See [12] for details.) If, moreover, $\mathbb{X}$ is a *metric space* with metric d, as in our examples, then the orbit space $\mathbb{M} := \mathbb{X}/\mathcal{G}$ can also be endowed with a metric $d_{\mathbb{M}}$ given by $d_{\mathbb{M}}(\mathcal{G}x, \mathcal{G}y) = \min\{d(x', y') \mid x' \in \mathcal{G}x, y' \in \mathcal{G}y\}$. The orbit space is *locally isometric* to the space $\mathbb{X}$. In this case, the quotient map $\varrho : \mathbb{X} \to \mathbb{X}/\mathcal{G}$ is a covering map, which is a local isometry. In particular, if $\mathbb{X} = \mathbb{E}^2$, as in our examples, orbit spaces are *flat*.
A normal subgroup $\mathcal{H}$ of $\mathcal{G}$ has index k if it has k distinct right cosets in $\mathcal{G}$, i.e., there are k elements $e = g_1, g_2, \ldots, g_k$ of $\mathcal{G}$ such that $\mathcal{G}$ is the disjoint union of $g_1\mathcal{H}, g_2\mathcal{H}, \ldots, g_k\mathcal{H}$. The $\mathcal{H}$-action on $\mathbb{X}$ induced by the $\mathcal{G}$-action gives rise to an orbit space $\mathbb{X}/\mathcal{H}$. This space is a k-sheeted cover of the orbit space $\mathbb{X}/\mathcal{G}$, since the natural map $\rho : \mathbb{X}/\mathcal{H} \to \mathbb{X}/\mathcal{G} : \mathcal{H}x \mapsto \mathcal{G}x$ is k-to-1. Indeed, each $\mathcal{G}$-orbit is partitioned into k orbits of $\mathcal{H}$: $\mathcal{G}x = g_1\mathcal{H}x \cup g_2\mathcal{H}x \cup \cdots \cup g_k\mathcal{H}x$. This is the key construction of covering spaces with finite covering degree.

3 Some finitely-sheeted covering spaces of the flat torus

Let a, b denote translations of the same length along the x- and y-axis in $\mathbb{E}^2$, respectively. We denote as $\mathcal{G}_1$ the group $< a, b >$ generated by a and b. Since $\mathcal{G}_1$ is Abelian, any $x \in \mathcal{G}_1$ can be uniquely written as $x = a^\alpha b^\beta$, where $\alpha, \beta \in \mathbb{Z}$. The *length* $\lambda_1(x)$ of x can thus be defined as the sum $|\alpha| + |\beta|$. We are going to study some covering spaces of the flat torus $\mathbb{T}_1^2 := \mathbb{E}^2/\mathcal{G}_1$. The original domain of $\mathcal{G}_1$ is denoted by $\mathcal{D}_1$. The edge length of $\mathcal{D}_1$ is denoted by $l = \delta(\mathcal{G}_1) = \|a\| = \|b\|$.

In the sequel, the inverse x^{-1} of an element x of $\mathcal{G}_1$ will alternatively be denoted by $\overline{x}$.

All subgroups are normal, since $\mathcal{G}_1$ is Abelian. Let us consider the subgroup $\mathcal{G}_2$ consisting of elements of $\mathcal{G}_1$ of even length. It is easy to see that $\mathcal{G}_2 = \langle \overline{a}b, ab \rangle$.

Lemma 1 *$\mathcal{G}_2$ is a subgroup of index 2 in $\mathcal{G}_1$.*

Proof. Let $\varphi : \mathcal{G}_1 \to \mathbb{Z}_2$ be the group homomorphism defined by $\varphi(x) = \lambda_1(x) \mod 2$. The subgroup $\mathcal{G}_2$ is the kernel of φ. According to the First Isomorphism Theorem (see, e.g., [1]) $\ker\varphi$ is a normal subgroup of $\mathcal{G}_1$, and $G/\ker\varphi \cong \phi(G)$. Therefore, $\mathcal{G}_1/\mathcal{G}_2 \cong \mathbb{Z}_2$. $\square$

In an inductive way, for $k > 1$, if $\mathcal{G}_{2^{k-1}}$ is a subgroup of index 2^{k-1} of $\mathcal{G}_1$ generated by two elements, i.e.,
$\mathcal{G}_{2^{k-1}} = \langle g, h\rangle$, $g, h \in \mathcal{G}_1$, we construct the subgroup $\mathcal{G}_{2^k}$ of $\mathcal{G}_{2^{k-1}}$ as follows: $\mathcal{G}_{2^k} = \langle \bar{g}h, gh\rangle$. The proof of the following lemma is similar to that of Lemma 1.

Lemma 2 *For any $k > 0$, $\mathcal{G}_{2^k}$ is a subgroup of index 2^k in $\mathcal{G}_1$.*

By construction, for any $k \geq 0$, the generators of $\mathcal{G}_{2^k}$ are two translations whose vectors are orthogonal and of equal length $\delta(\mathcal{G}_{2^k}) = \sqrt{2}^k l$. Following Section 2, the orbits of $\mathcal{G}_{2^k}$ are isomorphic to $\mathbb{Z}^2$ and the original domain $\mathcal{D}_{2^k}$ of $\mathcal{G}_{2^k}$ is a half-open square of edge length $\delta(\mathcal{G}_{2^k})$ (See Figure 1). For simplicity, x both denotes an element x of $\mathcal{G}_1$ and the image xO of the origin O by x.

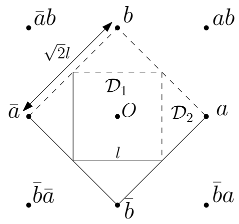

Figure 1: The original domains $\mathcal{D}_1$ and $\mathcal{D}_2$ of $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively.

For any $k > 0$, $\mathbb{T}^2_{2^k} = \mathbb{E}^2/\mathcal{G}_{2^k}$ is a flat torus, with the corresponding projection map $\pi_{2^k} : \mathbb{E}^2 \to \mathbb{T}^2_{2^k}$. $\mathbb{T}^2_{2^k}$ is a covering space of $\mathbb{T}^2_1$ together with the covering map $\rho_k := \pi_1 \circ \pi^{-1}_{2^k}$. Since the index of $\mathcal{G}_{2^k}$ in $\mathcal{G}_1$ is 2^k, we get the following result:

Proposition 3 *For any $k > 0$, $\mathbb{T}^2_{2^k}$ is a 2^k-sheeted covering space of $\mathbb{T}^2_1$.*

4 Delaunay triangulations via 2^k-sheeted covering spaces, $k > 0$

As in [5, 6], we stick to the definition of a *triangulation* of a topological space $\mathbb{Y}$ as a geometric simplicial complex K such that $K = \cup_{\sigma\in K}\sigma$ is homeomorphic to $\mathbb{Y}$. A triangulation defined by a point set $\mathcal{P}$ is a triangulation whose set of vertices is identical to $\mathcal{P}$.

The Delaunay triangulation of a point set $\mathcal{S}$ in $\mathbb{E}^2$ is a triangulation of the convex hull of $\mathcal{S}$ such that the circumscribing disk of any triangle does not contain any point of $\mathcal{P}$ in its interior. It can be defined uniquely even in denegerate cases [7].

We use the following definition of a Delaunay triangulation of a flat torus $\mathbb{T}^2 = \mathbb{E}^2/\mathcal{G}$ defined by a set of points $\mathcal{S} \in \mathbb{E}^2$.

Definition 1 ([5]) (Delaunay triangulation of $\mathbb{T}^2$). *Let $DT(\mathcal{GS})$ be the Delaunay triangulation of $\mathcal{GS}$ in $\mathbb{E}^2$. If $\pi(DT(\mathcal{GS}))$ is a simplicial complex in $\mathbb{T}^2$, then we call it the* Delaunay triangulation *of $\mathbb{T}^2$ defined by $\mathcal{S}$ and we denote it by $DT_{\mathbb{T}}(\mathcal{S})$.*

As shown in [5] this Delaunay triangulation is not always defined. However, there are always some covering spaces of the torus in which the Delaunay triangulation can be defined and computed.

Let $\Delta(\mathcal{S})$ denote the diameter of the largest disk in $\mathbb{E}^2$ that does not contain any point of a set $\mathcal{S}$ in its interior. Note that for any $p \in \mathbb{E}^2, \Delta(\mathcal{G}_{2^k}p) = \sqrt{2}^{k+1} l$.

Proposition 4 ([5]) *If $\Delta(\mathcal{GS}) < \frac{\delta(\mathcal{G})}{2}$, then $\pi(DT(\mathcal{GS}'))$ is a triangulation of $\mathbb{T}^2$ for any finite $\mathcal{S}' \supseteq \mathcal{S}$.*

Proposition 5 *If $\Delta(\mathcal{G}_1\mathcal{S}) < \frac{1}{2}\delta(\mathcal{G}_{2^k})$, then $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{S} \cup \mathcal{G}_{2^k}\mathcal{T}))$ is a triangulation of $\mathbb{T}^2_{2^k}$ for any finite point set $\mathcal{T}$ in $\mathbb{E}^2$. In particular, then $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{S}'))$ is a triangulation of $\mathbb{T}^2_{2^k}$ for any finite $\mathcal{S}' \supseteq \mathcal{S}$, $\mathcal{S}' \in \mathbb{E}^2$.*

Proof. Let $\mathcal{S}_{2^k}$ denote the set $\mathcal{G}_1\mathcal{S} \cap \mathcal{D}_{2^k}$. By Proposition 4, if $\Delta(\mathcal{G}_{2^k}\mathcal{S}_{2^k}) < \frac{1}{2}\delta(\mathcal{G}_{2^k})$, then $\pi_{2^k}(DT(\mathcal{G}_{2^k}\mathcal{S}_{2^k} \cup \mathcal{G}_{2^k}\mathcal{T}))$ is a triangulation. Let us note that $\mathcal{G}_{2^k}\mathcal{S}_{2^k} = \mathcal{G}_1\mathcal{S}$.

For $\mathcal{S}' \supseteq \mathcal{S}$, $\mathcal{G}_1\mathcal{S}' = \mathcal{G}_1\mathcal{S} \cup \mathcal{G}_1(\mathcal{S}' \setminus \mathcal{S}) = \mathcal{G}_1\mathcal{S} \cup \mathcal{G}_{2^k}(\mathcal{G}_1(\mathcal{S}' \setminus \mathcal{S}) \cap \mathcal{D}_{2^k})$. It remains to take $\mathcal{T} := \mathcal{G}((\mathcal{S}' \setminus \mathcal{S}) \cap \mathcal{D}_{2^k})$. $\square$

By Proposition 5, if the maximum empty disk diameter $\Delta(\mathcal{G}_1\mathcal{S})$ is smaller than $\frac{1}{2}\delta(\mathcal{G}_4) = l$, then $\pi_4(DT(\mathcal{G}_1\mathcal{S} \cup \mathcal{G}_4\mathcal{T})))$ is a simplicial complex for any finite $\mathcal{T}$ in $\mathbb{E}^2$. If it is smaller than $\frac{1}{2}\delta(\mathcal{G}_2) = \frac{\sqrt{2}}{2}l$, then $\pi_2(DT(\mathcal{G}_1\mathcal{S} \cup \mathcal{G}_2\mathcal{T})))$ is a simplicial complex.

Note that, for $k > 3$, $\frac{1}{2}\delta(\mathcal{G}_{2^k}) = \frac{1}{2}\sqrt{2}^k l > \sqrt{2}l = \Delta(\mathcal{G}_1 p)$, for any $p \in \mathcal{S}$. Therefore, by Proposition 5, $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{S}))$ is a simplicial complex. Let us now consider the case $k = 3$.

Corollary 6 *For any finite set $\mathcal{P} \subset \mathcal{D}_1$ of $n > 1$ points, $\pi_8(DT(\mathcal{G}_1\mathcal{P}))$ is a simplicial complex.*

Proof. The maximum empty disk diameter $\Delta(\mathcal{G}_1 p)$, for any $p \in \mathcal{P}$, is $\sqrt{2}l$, that is equal to $\frac{1}{2}\delta(\mathcal{G}_8)$ (See Figure 2). We are going to show that $\Delta(\mathcal{G}_1\mathcal{P})$ is strictly less than $\sqrt{2}l$.

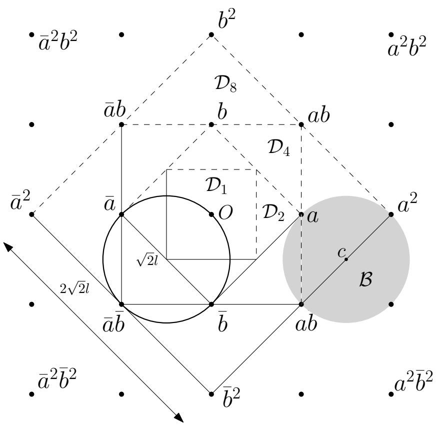

Figure 2: The original domain $\mathcal{D}_8$. Maximum empty disk with diameter $\Delta(\mathcal{G}_1 O)$.

Let us suppose that $\mathcal{B}$ is a disk with diameter $\sqrt{2}l$ centered at some point $c \in \mathbb{E}^2$ and containing no point of $\mathcal{G}_1\mathcal{P}$ in its interior (See Figure 2). $\mathcal{B}$ contains $\mathcal{D}_1(c)$. Since $\mathcal{D}_1(c)$ is a half-open square, there is only one point of $\mathcal{D}_1(c)$ on the boundary of $\mathcal{B}$. For p and $q \neq p, q \in \mathcal{P}$, there is a representative of $\mathcal{G}_1 p$ and a representative of $\mathcal{G}_1 q$ in $\mathcal{D}_1(c)$. At least one of these representatives lies in the interior of $\mathcal{B}$. This contradicts that the interior of $\mathcal{B}$ is empty.

If we add more points, the diameter of the largest empty disk cannot become larger. By Proposition 5, $\pi_8(DT(\mathcal{G}_1\mathcal{P}))$ is a simplicial complex. $\square$

Algorithm. For a finite set of points $\mathcal{P} \subset \mathcal{D}_1$, the algorithm of [5], inspired by the standard incremental algorithm [2], computes $DT_{\mathbb{T}}(\mathcal{G}_1\mathcal{P})$ via a 9-sheeted covering space. As soon as the condition of Proposition 4 is fulfilled upon insertion of a point, it switches to computing in $\mathbb{T}^2$.

Assuming that $\mathcal{P}$ contains at least two points, we modify this algorithm and use the covering spaces $\mathbb{T}^2_8, \mathbb{T}^2_4$, and $\mathbb{T}^2_2$ using Proposition 5 and Corollary 6.

• The algorithm starts by precomputing $\pi_8(DT(\mathcal{G}_1\{p, q\}))$, for any $\{p, q\} \subset \mathcal{P}$. Then it adds points one by one in the Delaunay triangulation of the 8-sheeted covering space $\mathbb{T}^2_8$.

• If after insertion of some set of points $\mathcal{S} \subseteq \mathcal{P}$, the maximum empty disk diameter becomes less than l, then $\pi_4(DT(\mathcal{G}_1\mathcal{S}'))$ is guaranteed to be a triangulation for any finite $\mathcal{S}' \supseteq \mathcal{S}$. So, we can discard all redundant periodic copies of simplices of $\pi_8(DT(\mathcal{G}_1\mathcal{S}))$ and switch to incrementally computing $\pi_4(DT(\mathcal{G}_1\mathcal{S}'))$ in the 4-sheeted covering space $\mathbb{T}^2_4$ for $\mathcal{S} \subseteq \mathcal{S}' \subseteq \mathcal{P}$.

• Similarly, if after some more insertions the maximum empty disk diameter becomes less than $\frac{\sqrt{2}}{2}l$, then we switch to incrementally computing triangulations in the 2-sheeted covering space $\mathbb{T}_2^2$.

• If it becomes less than $\frac{1}{2}l$, then we switch to computing $\pi(DT(\mathcal{G}_1\mathcal{P}))$ in $\mathbb{T}^2$.

For any $\mathcal{S} \subset \mathcal{P}$, $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{S}))$ contains 2^k periodic copies of each point of $\mathcal{S}$. Hence, using covering spaces whose number of sheets is as small as possible improves the efficiency of the algorithm, even though the asymptotic complexity is unchanged: it stays randomized worst-case time and space optimal [5].

5 Future work. Covering spaces and Delaunay triangulations of the double torus.

Whereas flat tori studied above, which have one handle, are a Euclidean surfaces, 2-tori, i.e., tori with two handles, are hyperbolic surfaces. A 2-torus can be constructed as the orbit space under action on the hyperbolic plane $\mathbb{H}^2$ of a discrete group generated by four hyperbolic translations. Due to lack of space, we only give here a few hints on how to extend the approach above to this case. Details and proofs will be given in a forthcoming paper.

A major difference with the Euclidean case is that hyperbolic translations do not commute.

Let $a, b, c,$ and d denote four hyperbolic translations and $\overline{a}, \overline{b}, \overline{c},$ and $\overline{d}$ their respective inverse translations. The group denoted as $\mathcal{G}_1^{\mathbb{H}} := \left\langle a, b, c, d \mid ab\overline{a}\overline{b}cd\overline{c}\overline{d} \right\rangle$ is defined as the quotient group of $\langle a, b, c, d\rangle$ by its smallest normal subgroup containing the element $ab\overline{a}\overline{b}cd\overline{c}\overline{d}$. For any $p \in \mathbb{H}^2$, $\mathcal{G}_1^{\mathbb{H}}p$ is a lattice. The fundamental domain of $\mathcal{G}_1^{\mathbb{H}}$ is a regular octagon in $\mathbb{H}^2$.

As in the Euclidean case, we can define the subgroup $\mathcal{G}_2^{\mathbb{H}}$ of $\mathcal{G}_1^{\mathbb{H}}$ consisting of the elements of even length. $\mathcal{G}_2^{\mathbb{H}}$ is a normal subgroup of index 2 of $\mathcal{G}_1^{\mathbb{H}}$, and $\mathbb{H}^2/\mathcal{G}_2^{\mathbb{H}}$ is a two-sheeted covering space of the 2-torus $\mathbb{H}^2/\mathcal{G}_1^{\mathbb{H}}$. The surface $\mathbb{H}^2/\mathcal{G}_2^{\mathbb{H}}$ does not have the same genus: it is a 3-torus.

In an inductive way, we define the normal subgroup $\mathcal{G}_{2^k}^{\mathbb{H}}, k > 1$ of index 2^k in $\mathcal{G}_1^{\mathbb{H}}$ as the subgroup of elements of even length in $\mathcal{G}_{2^{k-1}}^{\mathbb{H}}$ (the length is defined in terms of the generators of $\mathcal{G}_{2^{k-1}}^{\mathbb{H}}$). The orbit space $\mathbb{H}^2/\mathcal{G}_{2^k}^{\mathbb{H}}, k > 0$ is a 2^k-sheeted covering space of the 2-torus.

We expect that there exists $k > 1$, such that the Delaunay triangulation of $\mathbb{H}^2/\mathcal{G}_{2^k}^{\mathbb{H}}$ be well defined for any set of points in $\mathbb{H}^2$.

References

[1] M. Artin. *Algebra.* Prentice Hall, 1991.

[2] A. Bowyer. Computing Dirichlet tesselations. *Comput. J.*, 24:162–166, 1981.

[3] M. Caroli. *Triangulating Point Sets in Orbit Spaces.* Thèse de doctorat en sciences, Université de Nice-Sophia Antipolis, France, 2010. http://tel.archives-ouvertes.fr/tel-00552215/.

[4] M. Caroli and M. Teillaud. 3D periodic triangulations. In CGAL Editorial Board, editor, *CGAL User and Reference Manual.* 3.5 – 3.9 edition, 2009 – 2011. http://www.cgal.org/.

[5] M. Caroli and M. Teillaud. Computing 3d periodic triangulations. In *European Symposium on Algorithms*, volume 5757 of *LNCS*, pages 37–48, 2009.

[6] M. Caroli and M. Teillaud. Delaunay triangulations of point sets in closed Euclidean d-manifolds. In *Proc. 27th Annual Symposium on Computational Geometry*, pages 274–282, 2011.

[7] O. Devillers and M. Teillaud. Perturbations for Delaunay and weighted Delaunay 3D triangulations. *Computational Geometry: Theory and Applications*, 44:160–168, 2011.

[8] C. I. Grima and A. Màrquez. *Computational Geometry on Surfaces.* Kluwer Academic Publishers, 2001.

[9] M. Mazón and T. Recio. The computation of Voronoi diagrams on the Euclidean and spherical 2-orbifolds. In *Abstracts 10th EuroCG*, pages 12–15, 1994.

[10] T. Sousbie. The persistent cosmic web and its filamentary structure: I - theory and implementation. *Monthly Notices of the Royal Astronomical Society*, 414:350–383, 2011. Preliminary version available as http://arxiv.org/abs/arXiv:1009.4015.

[11] T. Sousbie, C. Pichon, and H. Kawahara. The persistent cosmic web and its filamentary structure: II - illustrations. *Monthly Notices of the Royal Astronomical Society*, 414:384–403, 2011. Preliminary version available as http://arxiv.org/abs/arXiv:1009.4014.

[12] J. Stillwell. *Geometry of Surfaces.* Universitext. Springer Verlag, Berlin, Heidelberg, New York, 1992.

[13] W. Thurston. *Three-Dimensional Geometry and Topology, Volume 1.* Princeton University Press, New Jersey, 1997.

[14] R. van de Weygaert, G. Vegter, H. Edelsbrunner, B. Jones, P. Pranav, C. Park, W. Hellwing, B. Eldering, N. Kruithof, E. Bos, J. Hidding, J. Feldbrugge, E. ten Have, M. van Engelen, M. Caroli, and M. Teillaud. Alpha, Betti and the Megaparsec universe: On the topology of the cosmic web. In *Transactions on Computational Science XIV*, volume 6970 of *LNCS*, pages 60–101. Springer-Verlag, 2011. http://www.springerlink.com/content/334357373166n902.