Algorithms for Length Spectra of Combinatorial Tori

Vincent Delecroix
Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400 Talence, France

Matthijs Ebbens
Institut Fourier, CNRS, Université Grenoble Alpes, France

Francis Lazarus
G-SCOP/Institut Fourier, CNRS, Université Grenoble Alpes, France

Ivan Yakovlev
Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400 Talence, France

Abstract

Consider a weighted, undirected graph cellularly embedded on a topological surface. The function assigning to each free homotopy class of closed curves the length of a shortest cycle within this homotopy class is called the marked length spectrum. The (unmarked) length spectrum is obtained by just listing the length values of the marked length spectrum in increasing order.

In this paper, we describe algorithms for computing the (un)marked length spectra of graphs embedded on the torus. More specifically, we preprocess a weighted graph of complexity n in time $O(n^2 \log \log n)$ so that, given a cycle with ℓ edges representing a free homotopy class, the length of a shortest homotopic cycle can be computed in $O(\ell + \log n)$ time. Moreover, given any positive integer k, the first k values of its unmarked length spectrum can be computed in time $O(k \log n)$.

Our algorithms are based on a correspondence between weighted graphs on the torus and polyhedral norms. In particular, we give a weight independent bound on the complexity of the unit ball of such norms. As an immediate consequence we can decide if two embedded weighted graphs have the same marked spectrum in polynomial time. We also consider the problem of comparing the unmarked spectra and provide a polynomial time algorithm in the unweighted case and a randomized polynomial time algorithm otherwise.

2012 ACM Subject Classification Mathematics of computing → Geometric topology; Mathematics of computing → Graphs and surfaces; Mathematics of computing → Enumeration

Keywords and phrases graphs on surfaces, length spectrum, polyhedral norm

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.26

Related Version *Full Version*: `https://arxiv.org/abs/2303.08036`

Funding *Francis Lazarus*: This work is supported by the project TOFU of the LabEx PERSYVAL-Lab (ANR-11-LABX-0025-01) funded by the French program Investissement d'avenir.

Acknowledgements We would like to thank Marcos Cossarini who suggested to look at the norm point of view and Bruno Grenet who enlightened us about the Polynomial Identity Testing problem. We would also like to thank the referees for their helpful comments, in particular for suggesting the complexity improvement in the case of unweighted graphs.

1 Introduction

Combinatorial surfaces are well-studied in computational topology and are usually represented as graphs cellularly embedded on a topological surface. Given a combinatorial surface S with underlying graph G, many algorithms exist for computing the length of its shortest homotopically non-trivial closed walk [23, 10, 13, 4, 3, 9, 2]. Here, the length of a walk is the sum of the weights of its edges if the edges are weighted, or the number of edges if not.

However, relatively little is known about how to compute the *second shortest* non-trivial closed walk, the *third shortest*, etc. More precisely, for every closed walk c in G, we can compute the length of the shortest closed walk freely homotopic to c on S. Obviously, this length only depends on the free homotopy class of c. The ordered sequence of lengths of all free homotopy classes of closed walks is called the *length spectrum* of S with respect to its (weighted) graph G, while the mapping between free homotopy classes of curves and their lengths is called the *marked length spectrum*. The marked length spectrum thus records for every length in the sequence from which free homotopy class it comes from. These notions are well studied in the realm of hyperbolic or Riemannian surfaces [15, 1, 16]. A striking result in that respect is that the marked length spectrum of a non-positively curved surface entirely determines the geometry of the surface [15]. In other words, one may learn the geometry of a surface by just looking at the length of its curves. However, the unmarked length spectrum does not determine the surface even in constant curvature [25].

Analogously, Schrijver [19, Th. 1] proved that embedded graphs that are minor-minimal among graphs with the same marked length spectrum, which he calls *kernels*, are determined by their marked length spectrum up to simple transformations. In [20] Schrijver restricts to unweighted graphs on the torus and notices that the marked length spectrum extends to an integer norm in $\mathbb{R}^2$, i.e. a norm taking integer values at integer vectors. See figure 1. Moreover, its dual unit ball is a finite polygon with integer vertices (see [17, 18]). This allows

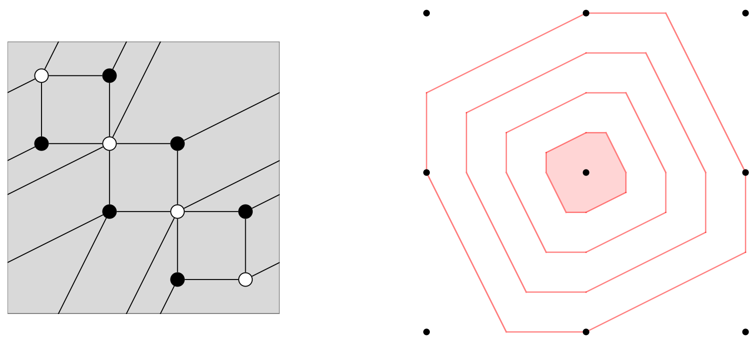

Figure 1 Left, an unweighted graph on the torus. Right, four dilates of the unit ball of the associated norm. Note that the vertical and horizontal generators of the torus have length four as can be seen from the right diagram.

him to reconstruct for every integer norm a graph whose marked length spectrum is given by this norm.

The aim of our paper is threefold. We first extend the results of Schrijver [20] to weighted graphs on the torus. There are good reasons to focus on the torus. For instance, the marked length spectrum of a graph embedded on the torus being a norm is due to the equivalence between homotopy and homology, which is not true for higher genus surfaces. For weighted graphs, the marked length spectrum function still extends to a norm on $\mathbb{R}^2$, that we denote by $N_{G,w}$, but not necessarily to an integer norm. However, we show that it is a polyhedral norm for any choice of weights w, i.e. that the unit ball $B_{G,w} := \{\alpha \in \mathbb{R}^2 \mid N_{G,w}(\alpha) \leq 1\}$ is always a polygon. We also prove that the number of extremal points of this polygon is bounded by a linear function of the number of vertices of G, independent of the weights w.

▶ **Theorem 1.** *Let (G, w) be a weighted graph with $|V|$ vertices cellularly embedded on the torus. Then $N_{G,w}$ is a polyhedral norm. Moreover, its unit ball $B_{G,w}$ is a polygon with no more than $4|V| + 5$ extremal points, and the ratio of the coordinates of each extremal point is rational.*

We also extend Schrijver's reconstruction of a toroidal graph from an integer norm [20] to the weighted case for non-integer polyhedral norms. See Theorem 7 in Section 3.

Our second goal is to provide algorithms to compute the unit ball $B_{G,w}$ and to compute the length spectrum. Our complexity estimates assume the standard RAM model of computation or the standard real-RAM model for non-integer weights supporting constant time arithmetic operations. We denote by n the complexity of G, i.e. its total number of edges and vertices.

▶ **Theorem 2.** *The unit ball $B_{G,w}$ can be computed in $O(n^2 \log\log n)$ time.*

This allows us to compute the length of the shortest closed walk freely homotopic to an input closed walk of ℓ edges in $O(\ell + \log n)$ time. It is a priori not obvious to sort the values of the length spectrum from their homotopy classes. However, by decomposing the unit ball into unimodular sectors, i.e., sectors generated by the columns of unimodular matrices, we are able to compute efficiently the first k values of the length spectrum.

▶ **Theorem 3.** *Let (G, w) be a weighted graph of complexity n cellularly embedded on the torus and let k be a positive integer. After $O(n^2 \log\log n)$ preprocessing time, the first k values of the length spectrum of (G, w) can be computed in $O(k \log n)$ time.*

Recently, Ebbens and Lazarus [7] used shortest path computations in the universal cover of the torus to determine the length spectrum. They compute the first k values of the length spectrum in time $O(kn^2 \log(kn))$. This is to be compared to $O(k \log n)$ in our submission.

Finally, we provide algorithms to check whether two weighted graphs have the same (un)marked length spectrum. In the unweighted case it takes the following simple form.

▶ **Theorem 4.** *The equality of marked and unmarked spectra of two unweighted graphs G and G' embedded on tori can be tested in time $O(n^2)$ and $O\left(n^3\right)$, respectively.*

Our algorithm for the marked length spectrum is also polynomial in the weighted case (Theorem 18). However, to compute the unmarked length spectrum we reduce the equality of length spectra to polynomial identity testing (PIT). See Theorem 19. It becomes deterministic polynomial in the unweighted case as stated in Theorem 4 above.

In contrast with [19], we provide in the full version an example of isospectral toroidal graphs whose associated unit balls are not related by any linear transformation. Hence, they cannot have the same marked spectrum even after applying a homeomorphism on the torus. Also, in the full version on arXiv, we show that in Theorems 2 and 3 the $\log\log n$ factor can be omitted if G is unweighted.

Organization of the paper

We start by discussing some preliminaries in Section 2. We next prove Theorem 1 in Section 3. Theorem 2 is the object of Sections 4 and 5, while Theorem 3 is proved in Section 6. The equality of length spectra is finally discussed in Section 7.

2 Preliminaries

Let $G = (V, E)$ be an undirected graph with vertex set V and edge set E. We allow G to have loop and multiple edges. We denote by $n := |V| + |E|$ the complexity of G. A *weight function* for G is a map $w : E \to \mathbb{R}_+$. The positive value $w(e)$ is the *weight* (or *length*) of the edge $e \in E$. We write (G, w) for a graph G with weight function w. A *walk* is a finite alternating sequence of vertices and edges, starting and ending with a vertex, such that any two successive elements in the sequence are incident. We also use *path* as a synonym for walk.

The *length* $w(c)$ of a walk c is the sum of the weights of its edges, counted with multiplicity. A walk is *closed* when its first and last vertices coincide. This vertex is the *basepoint* of the closed walk. A closed walk without repeated vertices is also called a *simple cycle*.

Throughout this paper, we will use S to denote a topological surface and T to denote the topological orientable surface of genus 1, i.e., a torus. In this paper we assume that G is *cellularly embedded* on S, which means that the complement $S \setminus G$ is a collection of open disks. This embedding can be represented using one of the standard representations, e.g., the incidence graph of flags [8] or rotation systems [14]. A surface together with a cellular embedding of a weighted graph is called a *combinatorial surface*.

Homotopy

Two walks of G are said *homotopic* if they are homotopic as curves in S, i.e., one can be continuously deformed into the other on S while keeping the endpoints fixed. Similarly, two closed walks are *freely homotopic* if they are so as curves in S. Here, we do not require the basepoint to stay fixed during the homotopy. Closed walks (freely) homotopic to a walk reduced to a vertex are said *trivial*. Homotopy is an equivalence relation between walks. The set of homotopy classes of closed walks with fixed basepoint v defines a group under concatenation. It is called the *fundamental group* of S, and denoted by $\pi_1(S, v)$. The fundamental group of the torus is Abelian and isomorphic to $\mathbb{Z}^2$. See e.g. [22]. $\pi_1(T, v)$ is thus in bijection with its set of conjugacy classes, hence with the set of free homotopy classes.

A closed walk is *tight* if it is shortest in its free homotopy class. Note that a homotopy class may contain more than one tight closed walk. Let $\mathcal{C}$ denotes the set of free homotopy classes of S. The map $\mathcal{C} \to \mathbb{R}_+$ that associates to every free homotopy class the length of a tight closed walk in the class is the *marked length spectrum* of S with respect to (G, w). The *unmarked length spectrum* is the list containing in increasing order the lengths of the non-trivial free homotopy classes of G, counted with multiplicity: if two homotopy classes have the same length, then this length will appear twice in the list.

Homology

Let F be the set of faces of the cellular embedding of G in S. We also call a face, an edge or a vertex, a k-cell for $k = 2, 1, 0$, respectively. The group of *2-chains*, C_2, is the group of formal linear combinations of faces with integer coefficients with the obvious addition as group operation. A typical element of C_2 has the form $\Sigma_{f \in F} n_f f$ with $n_f \in \mathbb{Z}$. Likewise, the group C_1 of 1-chains and the group C_0 of 0-chains are the groups of formal linear combinations of edges and vertices, respectively. Cells are assumed to be oriented and a cell multiplied by -1 represents the same cell with opposite orientation.

For $k = 1, 2$, the boundary operator $\partial_k : C_k \to C_{k-1}$ is the linear extension of the map that sends a k-cell to the formal sum of its boundary facets, where the coefficient of a facet in the sum is 1 if its orientation is induced by the orientation of the k-cell and -1 otherwise. The kernel of ∂_k is denoted by Z_k. Its elements are called *k-cycles*, not to be confused with cycles in the graph theoretical sense. The image of ∂_k is denoted by B_{k-1}. The *first homology group* of S with respect to the coefficients $\mathbb{Z}$ is the group $H_1(S; \mathbb{Z}) := \ker \partial_1 / \operatorname{Im} \partial_2$. From homology theory, $H_1(S; \mathbb{Z})$ does not depend on the specific cell decomposition induced by the cellular embedding of G. We can similarly define the first homology group with real coefficients $H_1(S; \mathbb{R})$. Since the 1-chains only depend on the graph G, we also write $Z_1(G; \mathbb{Z})$ for the group of 1-cycles. The Hurewicz theorem states that the map $\pi_1(S, v) \to H_1(S; \mathbb{Z})$ that sends (the homotopy class of) a closed walk to the (homology class of the) formal sum of its oriented edges is onto with kernel the commutator subgroup of $\pi_1(S, v)$. In the case of the torus, $\pi_1(T, v)$ is commutative, so that the above map is an isomorphism. From now on we will identify homotopy and first homology classes on the torus. We will denote by the same letter a closed walk on G and the corresponding 1-cycle in $Z_1(G; \mathbb{Z})$. The homotopy or homology class of a closed walk or 1-cycle c will be indifferently denoted by $[c]$.

Intersection numbers

Given two closed oriented curves c, d on S (endowed with an orientation) with transverse intersections, we may assign a sign to each intersection according to whether the tangents of c and d at the intersection form a positively oriented basis. The sum of the signs over all intersections is called the *algebraic intersection number.* It is a classical result that this number only depends on the homology classes $[c]$ and $[d]$ and that it defines an antisymmetric, nondegenerate bilinear form on $H_1(S; \mathbb{Z})$, denoted by the *pairing* $\langle [c], [d] \rangle$. Of course the total number of intersections of c and d is at least $|\langle [c], [d] \rangle|$.

The universal cover of the torus

We can form a torus by identifying the opposite sides of a square. Equivalently, we can see a torus as the quotient of the plane $\mathbb{R}^2$ by the action of the group of translations generated by $(1, 0)$ and $(0, 1)$, which we identify with the lattice $\mathbb{Z}^2$. Hence, T is identified with $\mathbb{R}^2/\mathbb{Z}^2$ and we have a quotient map $q : \mathbb{R}^2 \to T$. The plane $\mathbb{R}^2$, with the map q, is called the *universal cover* of T. Given a curve c with source point v on T, and a point $\tilde{v} \in q^{-1}(v)$, there is a unique curve $\tilde{c}$ in the plane with source $\tilde{v}$ that *projects* to c, i.e., such that $q(\tilde{c}) = c$. The curve $\tilde{c}$ is a *lift* of c. If c is a closed curve, then the vector from the source to the target of $\tilde{c}$ has integer coordinates and only depends on $[c]$. Hence, each homotopy class can be identified with a lattice translation. Such translations are called *covering transformations* (or translations). A curve is freely homotopic to a simple curve if and only if the coordinates of the corresponding covering translation are coprime [22, Sec. 6.2.2]. By the identification between $\mathbb{Z}^2$, $\pi_1(T, v)$ and $H_1(T; \mathbb{Z})$, any pair (α, β) of homology classes that generates $H_1(T; \mathbb{Z})$ must correspond to an invertible integer transformation, hence to a unimodular matrix. Equivalently, $\langle \alpha, \beta \rangle = \pm 1$. (α, β) is a *positively oriented basis* when $\langle \alpha, \beta \rangle = 1$.

Integer and intersection norms

Let $N : \mathbb{Z}^d \to \mathbb{R}_{\geq 0}$ satisfy the norm axioms:
- $N(\alpha + \beta) \leq N(\alpha) + N(\beta)$ (subadditivity)
- $N(k\alpha) = |k| N(\alpha)$ (absolute homogeneity)
- $N(\alpha) = 0 \implies \alpha = 0$ (separation)

Then N extends to $\mathbb{Q}^d$ using homogeneity, and can be extended to $\mathbb{R}^d$ so that it is continuous. It can be shown that this indeed provides a well-defined norm over $\mathbb{R}^d$ [24]. Such a function N, or its real extension, is called an *integer norm* if $N(\mathbb{Z}^d) \subseteq \mathbb{Z}_{\geq 0}$. Integer norms are *polyhedral*, i.e. their unit ball is a centrally symmetric polytope, and their dual unit ball is a centrally symmetric polytope with integer vertices [24, 20, 17]. See also [5, Sec. 6.0.4]. Integer norms naturally arise as length functions defined over homology classes of curves on surfaces. There are several ways to define curves and their lengths with respect to a graph G embedded on a surface S. One can consider continuous curves on S and define their length as the number of crossings with G. Schrijver [20] applies this framework when S is a torus and shows that this indeed defines a norm. He also considers a framework where the curves are in general position with respect to G, thus avoiding its vertices, and G is required to be 4-regular.

In [19] Schrijver shows that the first framework reduces to the second by considering the medial graph of G. In turn, the second framework reduces to our framework by duality, in the special case where the faces are quadrilaterals and the edges are unweighted.

3 Length spectrum and polyhedral norms on homology

In this section, given a weighted graph (G, w) embedded on a torus T, we introduce a norm on the first homology group of the torus that will be used throughout the article. A correspondence between graphs on the torus and polyhedral norms has been known for some time [20]. But, as far as we know, it has been studied only in the unweighted case and furthermore never analyzed from a computational point of view. For $\alpha \in H_1(T;\mathbb{Z})$ let

$$N_{G,w}(\alpha) := \inf\left\{\sum_{e\in E(G)} |x_e|w(e) : \sum_{e\in E(G)} x_e e \in Z_1(G;\mathbb{Z}) \text{ and } [\sum_{e\in E(G)} x_e e] = \alpha\right\}. \tag{1}$$

In the full version we show that $N_{G,w}$ satisfies the norm axioms. Hence, as explained in the subsection "Integer and intersection norms" of Section 2, $N_{G,w}$ extends to a norm on $H_1(T;\mathbb{R})$ (because $H_1(T;\mathbb{Z})$ is naturally a lattice in $H_1(T;\mathbb{R})$). The next lemma asserts that $N_{G,w}$ is indeed the marked length spectrum of T with respect to (G, w).

▶ **Lemma 5.** *For every $\alpha \in H_1(T;\mathbb{Z})$ we have*

$$N_{G,w}(\alpha) = \inf\left\{\sum_{i\in I} x_i \cdot w(c_i) : [\sum_{i\in I} x_i \cdot c_i] = \alpha \text{ and } x_i \in \mathbb{Z}_{\geq 0} \text{ for } i \in I\right\}, \tag{2}$$

where $\{c_i\}_{i\in I}$ is the (finite) set of all simple cycles in G. The infimum in (2) is attained. Furthermore, for every $\alpha \in H_1(T;\mathbb{Z})$, $N_{G,w}(\alpha)$ is the length of a shortest closed walk c in G with $[c] = \alpha$. In other words, $N_{G,w}$ is the marked length spectrum of T with respect to (G, w).

In the proof of Theorem 1 below we show that the extremal points of the unit ball $B_{G,w} = \{\alpha \in H_1(T;\mathbb{R}) \mid N_{G,w}(\alpha) \leq 1\}$ of $N_{G,w}$ correspond to homology classes that can be represented by simple cycles in G. The next lemma gives a bound on their number.

For a subset X of a real vector space let $\operatorname{conv}(X)$ denote the convex hull of X. Note that $H_1(T;\mathbb{Z})$ is naturally a subset of the real vector space $H_1(T;\mathbb{R})$.

▶ **Lemma 6.** *Let G be a graph with $|V|$ vertices cellularly embedded on the torus T, and let $\mathcal{SC}_G \subset H_1(T;\mathbb{Z})$ be the set of homology classes of curves that can be represented as simple cycles in G. Then, in $H_1(T;\mathbb{R})$ we have $|\operatorname{conv}(\mathcal{SC}_G) \cap H_1(T;\mathbb{Z})| \leq 4|V| + 5$.*

Proof. Identify $H_1(T;\mathbb{Z})$ with $\mathbb{Z}^2$ via an arbitrary positively oriented basis. $H_1(T;\mathbb{R})$ is then identified with $\mathbb{R}^2$, and the algebraic intersection pairing is given by $\langle (x,y),(x',y')\rangle = xy' - x'y$, whose absolute value is the Euclidean area of the parallelogram generated by these vectors.

Let $\alpha, \beta \in \mathcal{SC}_G$ be represented by simple cycles c_α, c_β in G. On the one hand, the number of intersections between c_α and c_β is bounded by $|V|$, since each intersection corresponds to at least one vertex of G, and all these vertices must be different. On the other hand, it is bounded below by the algebraic intersection number $|\langle\alpha,\beta\rangle|$. Hence, $|\langle\alpha,\beta\rangle| \leq |V|$.

Denote by $\|\cdot\|$ the Euclidean norm on $\mathbb{R}^2$ and by $\operatorname{dist}(\cdot,\cdot)$ the Euclidean distance, and consider $\alpha, \beta \in \mathcal{SC}_G$ such that $|\langle\alpha,\beta\rangle|$ is maximal. Then for any $\gamma \in \mathcal{SC}_G$, we have $|\langle\gamma,\alpha\rangle| \leq |\langle\alpha,\beta\rangle|$ and $|\langle\gamma,\beta\rangle| \leq |\langle\alpha,\beta\rangle|$. Note that, since these numbers are the areas of the corresponding parallelograms, $|\langle\gamma,\alpha\rangle| = \|\alpha\| \cdot \operatorname{dist}(\gamma,\mathbb{R}\alpha)$, $|\langle\gamma,\beta\rangle| = \|\beta\| \cdot \operatorname{dist}(\gamma,\mathbb{R}\beta)$ and

$|\langle\alpha,\beta\rangle| = \|\alpha\| \cdot \mathrm{dist}(\beta, \mathbb{R}\alpha) = \|\beta\| \cdot \mathrm{dist}(\alpha, \mathbb{R}\beta)$, where $\mathbb{R}\alpha, \mathbb{R}\beta$ denote the one-dimensional $\mathbb{R}$-subspaces generated by α, β respectively. It follows that $\mathrm{dist}(\gamma, \mathbb{R}\alpha) \leq \mathrm{dist}(\beta, \mathbb{R}\alpha)$ and $\mathrm{dist}(\gamma, \mathbb{R}\beta) \leq \mathrm{dist}(\alpha, \mathbb{R}\beta)$, and so $\mathcal{SC}_G$ is contained in the parallelogram P with vertices $\pm\alpha \pm \beta$, see Figure 2. Clearly, the area $A(P)$ of P is $4|\langle\alpha,\beta\rangle|$. At the same time, by Pick's

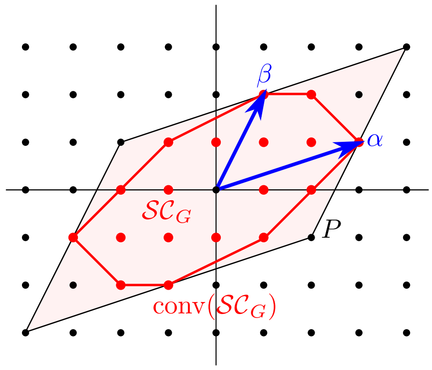

Figure 2 The elements of $\mathcal{SC}_G$ are contained in a parallelogram P of area at most $4|V|$.

theorem $A(P) = I + B/2 - 1$, where I is the number of integer points strictly inside P and B is the number of integer points on its boundary. Since α and β are homology classes represented by simple cycles, their corresponding vectors in $\mathbb{Z}^2$ have coprime coordinates, i.e. the only integer points on the vectors α and β are their endpoints. It follows that the only integer points on the boundary of P are $\pm\alpha, \pm\beta, \pm\alpha \pm \beta$ and so $B = 8$.

Finally, since $\mathcal{SC}_G \subset P$, we have $\mathrm{conv}(\mathcal{SC}_G) \subset P$ as well, and so

$$|\mathrm{conv}(\mathcal{SC}_G) \cap \mathbb{Z}^2| \leq I + B = A(P) + B/2 + 1 = 4|\langle\alpha,\beta\rangle| + 5 \leq 4|V| + 5. \qquad ◀$$

In the full version, it is shown that the order of the bound in Lemma 6 is optimal. We now pass to the proof of Theorem 1. In the unweighted case, the polyhedrality of the norm follows from its integrality [20]. However, this argument does not apply in the weighted case.

Proof of Theorem 1. We refer to the full version for a proof that $N_{G,w}$ satisfies the norm axioms. Hence, as explained in Section 2, $N_{G,w}$ extends to a norm on $H_1(T;\mathbb{R})$. To prove the polyhedrality of this norm, we show that

$$B_{G,w} = \mathrm{conv}\left(\left\{\frac{[c_i]}{w(c_i)} \mid i \in I\right\}\right), \tag{3}$$

where $\{c_i\}_{i\in I}$ is the (finite) set of all oriented simple cycles in G, as in Lemma 5.

Denote the right-hand side of (3) by $B'_{G,w}$. Clearly, for every $i \in I$ we have $N_{G,w}([c_i]) \leq w(c_i)$, so $B'_{G,w} \subset B_{G,w}$. Conversely, take any homology class $\alpha \in H_1(T;\mathbb{Z})$. By Lemma 5, we have $N_{G,w}(\alpha) = \sum_{i\in I} x_i \cdot w(c_i)$ for some $x_i \in \mathbb{Z}_{\geq 0}$ such that $\alpha = [\sum_{i\in I} x_i \cdot c_i]$. Then

$$\frac{\alpha}{N_{G,w}(\alpha)} = \frac{\sum_{i\in I} x_i \cdot [c_i]}{\sum_{i\in I} x_i \cdot w(c_i)} = \sum_{i\in I} \frac{x_i w(c_i)}{\sum_{j\in I} x_j \cdot w(c_j)} \cdot \frac{[c_i]}{w(c_i)}$$

is a representation of $\frac{\alpha}{N_{G,w}(\alpha)}$ as a convex combination of $\frac{[c_i]}{w(c_i)}, i \in I$. Hence $B_{G,w} \subset B'_{G,w}$ and we get (3). By definition, the $[c_i]$ can be represented by simple cycles in G. By Lemma 6 their number is at most $4|V| + 5$, and so the number of extremal points of $B_{G,w}$ is also at most $4|V| + 5$. The slopes of $[c_i]/w(c_i)$ are rational since the $[c_i]$ belong to $H_1(T;\mathbb{Z})$. ◀

Finally, we show how to reconstruct a weighted graph (G, w) embedded on the torus T from a polyhedral norm on $\mathbb{R}^2$.

▶ **Theorem 7.** *Let $N : \mathbb{R}^2 \to \mathbb{R}$ be a polyhedral norm all of whose extremal points have rational slopes. Let $\{\pm(p_i, q_i)\}_{i=1,\dots,n}$ be the set of non-zero integral vectors closest to the origin on the rays issued from the origin in the direction of the extremal points of the unit ball $\{v \in \mathbb{R}^2 : N(v) \leq 1\}$. Then there exists a weighted 4-valent graph (G, w) embedded on the torus T with $\sum_{1\leq i<j\leq n} |p_i q_j - p_j q_i|$ vertices so that $N_{(G,w)} = N$.*

4 Good short basis

Our computation of the length spectrum and of its unit ball relies on the initial computation of a *good short basis*. By a *short basis* we mean a pair of tight simple cycles (a, b) in G such that a is a shortest non-trivial closed walk and b is a shortest non-trivial closed walk satisfying $\langle [a], [b] \rangle = 1$. We say that (a, b) is a *good basis* if $([a], [b])$ is a positively oriented basis of $H_1(T; \mathbb{Z})$ and a and b intersect along a connected path, possibly reduced to a vertex.

▶ **Lemma 8.** *Let (G, w) be a weighted graph of complexity n cellularly embedded on the torus. A good short basis can be computed in $O(n \log n)$ time.*

Sketch of proof. We first compute a shortest non-trivial closed walk a in $O(n \log n)$ time following Kutz [13, Th. 1]. This closed walk must be a tight simple cycle as otherwise it could be decomposed into shorter non-trivial closed walks. We claim that among all shortest non-trivial closed walks b satisfying $\langle [a], [b] \rangle = 1$ there is one that intersects a along a connected path. See full version. Cutting T along a yields an annulus A with two copies a' and a'' of a as boundary components. By the above claim, b intersects A in a shortest path connecting two copies of the same vertex respectively on a' and a''. We find this shortest path using the multiple-source shortest path algorithm of Klein; see [12] and [2, Th. 3.8]. This algorithm builds a data structure in $O(n \log n)$ time that allows to query for the distance between any vertex on a' and any other vertex in A in $O(\log n)$ time. We need to query for the $O(n)$ pairs of copies of vertices of a and retain a pair (u', u'') that minimizes the distance. In order to find an explicit representative of b, we can in a second step run Dijkstra's algorithm with source u' in A. Finally, to ensure that b intersects a along a connected path, we can replace the subpath between u' and the last occurrence of a vertex on a' by a subpath of a' with the same length and do similarly on a''. The total running time is $O(n \log n)$. We obtain b by gluing back the two copies of a. ◀

We shall always express a homology class in the basis (a, b) and identify the class with a vector in $\mathbb{Z}^2$. Hence, $[a]$ and $[b]$ are identified with $(1, 0)$ and $(0, 1)$, respectively.

5 Computing the unit ball

Here, we provide an algorithm for computing the unit ball $B_{G,w}$ of $N_{G,w}$ corresponding to the weighted graph (G, w). Let $\mathcal{TSC}_{G,w} \subset H_1(T; \mathbb{Z})$ be the set of homology classes that admit a tight and simple representative in G. Of course, $\mathcal{TSC}_{G,w} \subseteq \mathcal{SC}_G$, and the homology classes of a and b computed in Section 4 are in $\mathcal{TSC}_{G,w}$ by construction. In Section 3, we proved that $B_{G,w}$ is the convex hull of a set $\{\alpha/N_{G,w}(\alpha)\}_{\alpha \in \mathcal{SC}_G}$ containing $O(|V|)$ classes. We shall compute a subset H of $\mathcal{TSC}_{G,w}$ whose normalized vectors, $\{\alpha/N_{G,w}(\alpha)\}_{\alpha \in H}$, include all the extremal points of $B_{G,w}$. Since the coordinates of each element of $\mathcal{TSC}_{G,w}$ must be coprime, the set of directions defined by the elements of $\mathcal{TSC}_{G,w}$ are pairwise distinct and naturally ordered angularly. We search for H by exploring the whole set of directions using dichotomy together with a simple pruning strategy. Suppose we need to explore the angular sector $\angle(\alpha, \beta)$, where (α, β) forms a basis of $H_1(T; \mathbb{Z})$. The dichotomy consists in splitting the sector into the sectors $\angle(\alpha, \gamma)$ and $\angle(\gamma, \beta)$ with $\gamma := \alpha + \beta$. Note that (α, γ) and (γ, β) are again bases of $H_1(T; \mathbb{Z})$. In particular, the coordinates of γ are coprime. Since for any nonzero $\eta \in H_1(T; \mathbb{Z})$, the vector $\eta/N_{G,w}(\eta)$ lies on the boundary of the unit ball, it follows by convexity of $B_{G,w}$ that the segment $[\frac{\alpha}{N_{G,w}(\alpha)}, \frac{\beta}{N_{G,w}(\beta)}]$ is a subset of a supporting line of $B_{G,w}$ whenever $\gamma/N_{G,w}(\gamma)$ lies on this segment. This last condition has a simple certificate.

▷ Claim 9. $\frac{\gamma}{N_{G,w}(\gamma)}$ lies on the segment $[\frac{\alpha}{N_{G,w}(\alpha)}, \frac{\beta}{N_{G,w}(\beta)}]$ if and only if $N_{G,w}(\alpha + \beta) = N_{G,w}(\alpha) + N_{G,w}(\beta)$.

It follows from the previous discussion that if $N_{G,w}(\alpha + \beta) = N_{G,w}(\alpha) + N_{G,w}(\beta)$, then the interior of the sector $\angle(\alpha, \beta)$ cannot contain any extremal point and we can prune this whole sector in our search. This leads to the pseudo-code of Algorithm 1 for computing H. In the sequel, we say that a pair of closed walks in G is **good** if they are simple and tight cycles, if their homology classes form a basis of $H_1(T; \mathbb{Z})$, and if they moreover intersect along a connected path, possibly reduced to a vertex.

Algorithm 1 Compute H.

Require: A weighted graph (G, w) cellularly embedded on the torus

Ensure: A short basis (a, b) and a sorted list $H = [((x_i, y_i), c_i, w(c_i))]$ where c_i is a simple tight cycle in G, $(x_i, y_i) \in \mathbb{Z}^2$ represents its homology class $[c_i] = x_i[a] + y_i[b]$, and $w(c_i) = N_{G,w}([c_i])$. Also, the extremal points of $B_{G,w}$ are contained in the set of vectors $\{[c_i]/w(c_i) : i \in \{0, \dots, size(H) - 1\}\}$.

1: Compute a good short basis (a, b) as explained in Section 4
2: $h_1 := ((1, 0), a, w(a))$ {Note that $N_{G,w}([a]) = w(a)$}
3: $h_2 := ((0, 1), b, w(b))$ {and that $N_{G,w}([b]) = w(b)$.}
4: $\overline{h_1} := ((-1, 0), \overline{a}, w(a))$
5: $H := \{h_1, h_2\}$ {Initialise H.}
6: $S := \{\angle(h_1, h_2), \angle(h_2, \overline{h_1})\}$ {Initialise a set of sectors to explore with the upper quadrants.}
7: **while** $S \neq \emptyset$ **do**
8: Extract and remove from S a sector $\angle(h, h')$
9: $(x, y), c, \ell := h$ {Note that $N_{G,w}([c]) = \ell$.}
10: $(x', y'), c', \ell' := h'$ {Similarly $N_{G,w}([c']) = \ell'$.}
Require: (c, c') is a good pair
11: Compute a tight representative c'' of $\gamma'' := [c] + [c']$ with its norm $\ell'' := N_{G,w}(\gamma'') = w(c'')$
12: **if** $\ell'' < \ell + \ell'$ **then**
13: $h'' := ((x + x', y + y'), c'', \ell'')$
14: Insert h'' in H between h and h'
15: $S := S \cup \{\angle(h, h''), \angle(h'', h')\}$
16: **end if**
17: **end while**
18: $H := H \cup \overline{H}$ {Add the symmetric of H w.r.t. the origin.}

By subadditivity of the norm, the test in Line 12 of Algorithm 1 may only fail when $N_{G,w}([c'']) = N_{G,w}([c]) + N_{G,w}([c'])$. It then follows from Claim 9 and the preceding discussion on our pruning strategy that we are not missing any direction of extremal points in the upper plane when adding homology classes in Line 14. Moreover, Line 18 and the central symmetry of the unit ball ensure that the above algorithm indeed computes a set of homology classes whose normalized vectors contains the extremal points of $B_{G,w}$. It remains to explain how to perform the computation in Line 11 and to analyse the complexity of Algorithm 1.

▶ **Lemma 10.** *Let (α, β) be a homology basis such that α and β have representatives, respectively c_α and c_β, forming a good pair. We can compute a tight representative $c_{\alpha+\beta}$ for $\alpha + \beta$ in $O(n \log\log n)$ time.*

Sketch of proof. By hypothesis, c_α and c_β intersect along a connected path $p_{\alpha\beta}$. The path $p_{\alpha\beta}$ may be oriented the same way or not in c_α and c_β. We consider the case where it is oriented the same way (see the full version for the other case). Then $c_\alpha = p_\alpha \cdot p_{\alpha\beta}$ and $c_\beta = p_\beta \cdot p_{\alpha\beta}$ for some paths p_α, p_β in G. We cut T along $c_\alpha \cup c_\beta$, viewed as a subgraph of G. We obtain a hexagonal plane domain $\mathcal{D}$ with sides p_β, $p_{\alpha\beta}$, p_α, $\overline{p_\beta}$, $\overline{p_{\alpha\beta}}$, $\overline{p_\alpha}$ in the clockwise order around the boundary of $\mathcal{D}$. See Figure 3. The universal cover of T is tessellated by

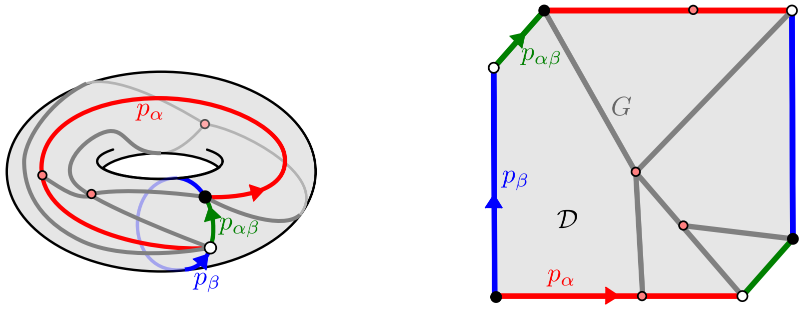

Figure 3 Cutting T along $c_\alpha \cup c_\beta$.

translated copies of $\mathcal{D}$ glued along their sides so that the side p of a domain is glued to the side $\overline{p}$ of the adjacent domain. See Figure 4. As before, let $\gamma = \alpha + \beta$. Since (α, γ) is a positively oriented basis, we know that $\langle \alpha, \gamma \rangle = 1$. Hence, any representative of γ must cross c_α. Let c_γ be a tight representative of γ with a lift $\tilde{c}_\gamma$ in the universal cover starting from a vertex $\tilde{v}$ on the side p_α or $p_{\alpha\beta}$ of a domain $\mathcal{D}_0$ and ending at the vertex $\tilde{w} := \tilde{v} + \tau_\alpha + \tau_\beta$, where τ_α and τ_β are the covering translations corresponding to α and β respectively. There are two situations according to whether $\tilde{v}$ lies on the side p_α or $p_{\alpha\beta}$ of $\mathcal{D}_0$. See Figure 5.

- If $\tilde{v}$ lies on the side p_α of $\mathcal{D}_0$ then $\tilde{w}$ belongs to the side $\overline{p_\alpha}$ of $\mathcal{D}_1 := \mathcal{D}_0 + \tau_\alpha$. We claim that $\mathcal{D}_0 \cup \mathcal{D}_1$ is convex, i.e., any two vertices in $\mathcal{D}_0 \cup \mathcal{D}_1$ can be joined by a shortest path contained in $\mathcal{D}_0 \cup \mathcal{D}_1$. Indeed, since c_α is tight, any bi-infinite concatenation of its lifts is a geodesic line in the weighted lift of G in the universal cover of T. Similarly, any bi-infinite concatenation of lifts of c_β is a geodesic line and thus delimits two convex half-planes. See the dotted and broken lines in Figure 5. It follows that $\mathcal{D}_0 \cup \mathcal{D}_1$ is the intersection of four half-planes, hence is convex. We can thus assume that c_γ has a lift in $\mathcal{D}_0 \cup \mathcal{D}_1$ with endpoints $\tilde{v}$ and $\tilde{w}$ on the boundary of $\mathcal{D}_0 \cup \mathcal{D}_1$. We can glue the p_α side of $\mathcal{D}_0$ with the $\overline{p_\alpha}$ side of $\mathcal{D}_1$ and search for the shortest generating cycle of the resulting annulus in $O(n \log n)$ time as in [6, Prop. 2.7(e)], or, more efficiently, in time $O(n \log\log n)$ as in [11, Theorem 7].
- If $\tilde{v}$ lies on the side $p_{\alpha\beta}$ of $\mathcal{D}_0$ we can compute $\tilde{c}_\gamma$ in linear time. See the full version.

In all cases we may compute a tight representative of γ in $O(n \log\log n)$ time. ◀

▶ **Lemma 11.** *The tight representative $c_{\alpha+\beta}$ computed in Lemma 10 can be modified in $O(n)$ time in order to satisfy the additional following property (P): The intersection of any lift of $c_{\alpha+\beta}$ in the universal cover with any line is either empty or a common connected subpath.*

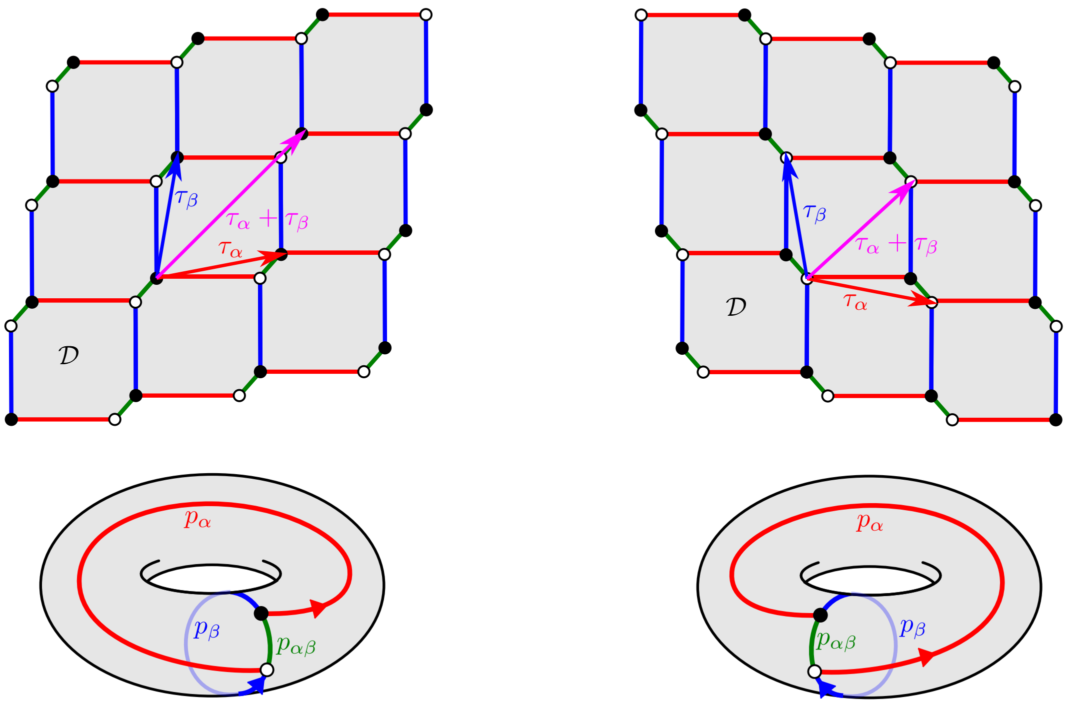

Figure 4 The universal cover of T. Left: $p_{\alpha\beta}$ is oriented consistently with both c_α and c_β. Right: $p_{\alpha\beta}$ has opposite orientation in c_α and c_β.

▶ **Lemma 12.** *Let (α, β) be a homology basis such that α and β have representatives, respectively c_α and c_β, forming a good pair. If $N_{G,w}(\alpha + \beta) < N_{G,w}(\alpha) + N_{G,w}(\beta)$, then the tight representative $c_{\alpha+\beta}$ computed in Lemma 11 is such that $(c_\alpha, c_{\alpha+\beta})$ and $(c_{\alpha+\beta}, c_\beta)$ are good pairs.*

Since in Algorithm 1 we only add $\alpha + \beta$ to H at line 14 when α, β are already in H with $N_{G,w}(\alpha + \beta) < N_{G,w}(\alpha) + N_{G,w}(\beta)$, Lemma 12 immediately implies

▶ **Corollary 13.** $H \subset \mathcal{TSC}_{G,w} \subset \mathcal{SC}_G$

▶ **Corollary 14.** *The number of iterations in the while loop of Algorithm 1, from Line 7 to 17, is bounded by twice the size of $\mathcal{TSC}_{G,w}$.*

▶ **Proposition 15.** *Algorithm 1 runs in $O(n^2 \log\log n)$ time.*

Proof. From Corollary 14, Algorithm 1 enters at most $2|\mathcal{TSC}_{G,w}|$ times the while loop between Lines 7 and 17. This is $O(n)$ iterations by Lemma 6. Each iteration takes $O(n \log\log n)$ time for executing Line 11 by Lemmas 10 and 11. Lemma 12 ensures that only good pairs are stored at Line 15, so that the requirement for executing Line 11 is always satisfied. Line 18 moreover takes time $O(|H|) = O(n)$. Since every other line takes constant time to execute, the total time for running Algorithm 1 is $O(n \cdot n \log\log n + n) = O(n^2 \log\log n)$. ◀

We are now ready to prove that $B_{G,w}$ can be computed in $O(n^2 \log\log n)$ time.

Proof of Theorem 2. Proposition 15 states that we can compute in $O(n^2 \log\log n)$ time a list H of $O(n)$ vectors, with their norms, that contains the directions of the extremal points of the unit ball $B_{G,w}$. After normalising the vectors we compute their convex hull in $O(n \log n)$ time with any classical convex hull algorithm. Overall, this leads to an $O(n^2 \log\log n)$ time algorithm for computing $B_{G,w}$. ◀

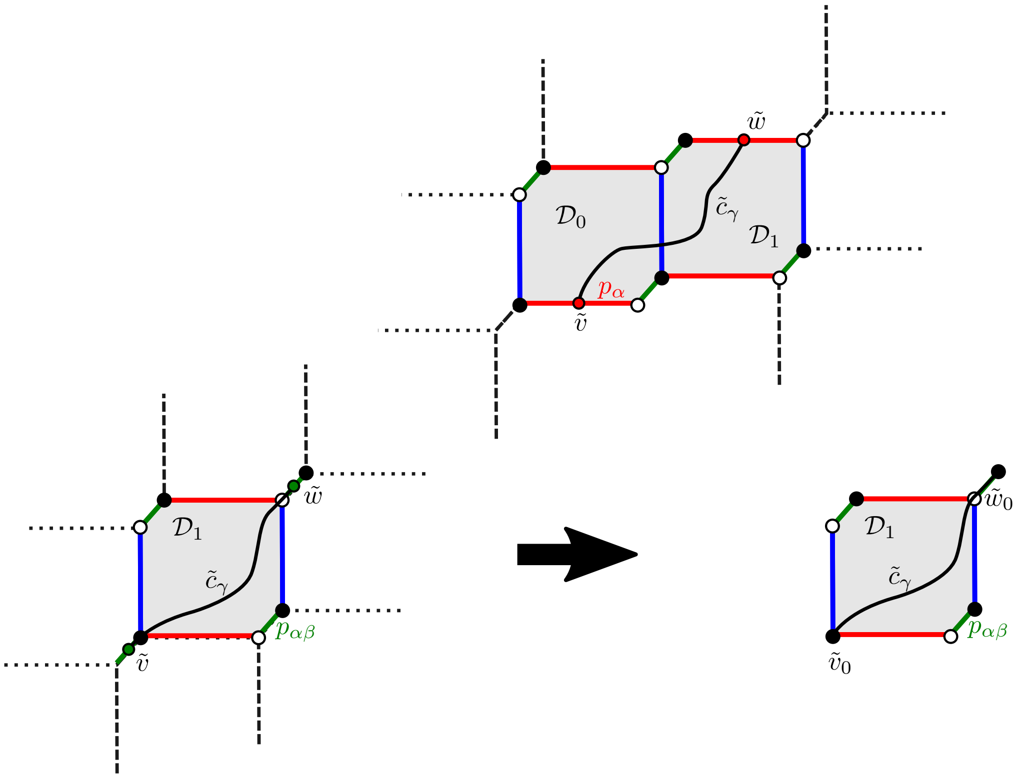

Figure 5 Top: if c_γ crosses p_α, then c_γ has a lift in $\mathcal{D}_0 \cup \mathcal{D}_1$. The dotted and broken lines are supporting geodesics for the considered regions. Bottom: when c_γ crosses $p_{\alpha\beta}$ it has a lift contained in the union of $\mathcal{D}_1$ with two lifts of $p_{\alpha\beta}$. We can shift the origin of the lift to $\tilde{v}_0$.

For further reference, we establish a useful property of the ordered set of elements in H. Namely, two consecutive cycles in H define a *unimodular* cone.

▶ **Lemma 16.** *The sorted list $H = [((x_i, y_i), c_i, w(c_i))]$ computed by Algorithm 1 is such that the rays $\mathbb{R}_{\geq 0}c_i$ are ordered cyclically by angle, and $\langle [c_i], [c_{i+1}] \rangle = 1$ for all i. In particular, the half-open cones $C_i = \mathbb{Z}_{\geq 0}[c_i] + \mathbb{Z}_{>0}[c_{i+1}]$ constitute a partition of $H_1(T; \mathbb{Z})$.*

6 Computing the length spectrum

We now give a proof of Theorem 3.

Proof. First, as described in Sections 4 and 5, we compute in time $O(n^2 \log\log n)$ a short basis (a, b) and a list H of triples $((x_i, y_i), c_i, w(c_i))$, where c_i is a simple tight cycle in G, $(x_i, y_i) \in \mathbb{Z}^2$ represents its homology class $[c_i] = x_i[a] + y_i[b]$ and $w(c_i) = N_{G,w}([c_i])$. From Lemma 16, this provides a partition of $H_1(T; \mathbb{Z})$ into half-open cones $C_i = \mathbb{Z}_{\geq 0}[c_i] + \mathbb{Z}_{>0}[c_{i+1}]$. Note that by construction $N_{G,w}$ is linear over each C_i.

We then compute the values of the length spectrum iteratively, in increasing order, storing them into a list Λ, initially empty. Intuitively, the algorithm consists in sweeping $H_1(T; \mathbb{Z})$ by increasing the radius of the λ-ball $\lambda B_{G,w}$ from $\lambda = 0$. Each time a lattice point is swept, its norm λ is added to Λ. We actually sweep the cones C_i in parallel.

By Lemma 16, the half-open cones C_i decompose the ball $\lambda B_{G,w}$ into sectors $C_i^\lambda := \lambda B_{G,w} \cap C_i$. For each sweeping value λ and each i we store two ordered subsets of C_i^λ into dequeues (double-ended queues) F_i^h and F_i^v corresponding to the *horizontal* and *vertical* sweeping front, respectively. Formally $F_i^h = ((x_1^h, y_1^h), \ldots, (x_{i_h}^h, y_{i_h}^h))$, where each $(x_j^h, y_j^h) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{>0}$ is such that $x_j^h[c_i] + y_j^h[c_{i+1}] \in C_i^\lambda$ and $(x_j^h + 1)[c_i] + y_j^h[c_{i+1}] \notin \lambda B_{G,w}$. Moreover, the homology classes in F_i^h are ordered by their norms in increasing order. Similarly, $F_i^v = ((x_1^v, y_1^v), \ldots, (x_{i_v}^v, y_{i_v}^v))$ contains the list (ordered by norm in increasing order) of coordinates of the homology classes contained in $\lambda B_{G,w}$ but whose translates by $[c_{i+1}]$ have norms larger than λ. Initially F_i^h is the empty dequeue and F_i^v contains the coordinates $(0,0)$ of the zero class. See Figure 6

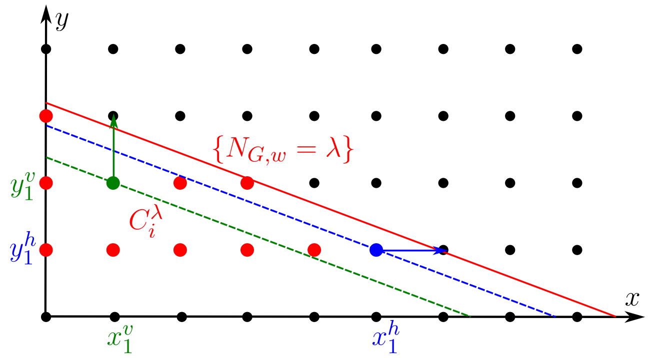

Figure 6 Here, the point (x, y) represents the class $x[c_i] + y[c_{i+1}]$. The solid red line represents the points whose norm is λ. The dashed lines correspond to the norms of (x_1^h, y_1^h) and (x_1^v, y_1^v).

▷ Claim 17. The coordinates in the $([c_i], [c_{i+1}])$ basis of the homology class in $C_i \setminus C_i^\lambda$ with the smallest norm is either $(x_1^h + 1, y_1^h)$ or $(x_1^v, y_1^v + 1)$.

We are now ready to describe the sweeping algorithm. We store the indices i of the sectors C_i^λ in a (balanced) binary search tree $\mathcal{S}$ allowing minimum extraction, deletion and insertion in logarithmic time. The key of sector C_i^λ used for comparisons in the tree is the minimum norm of a homology class in $C_i \setminus C_i^\lambda$. From the previous claim it can be computed in constant time from F_i^h and F_i^v as

$$\begin{aligned} \text{key}[i] &= \min(N_{G,w}([(x_1^h + 1)c_i + y_1^h c_{i+1}]), N_{G,w}([x_1^v c_i + (y_1^v + 1)c_{i+1}])) \\ &= \min((x_1^h + 1)w(c_i) + y_1^h w(c_{i+1}), x_1^v w(c_i) + (y_1^v + 1)w(c_{i+1})) \end{aligned}$$

Suppose we have computed the m first values of the length spectrum, i.e., $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$, and we want to compute λ_{m+1}. We extract and remove from $\mathcal{S}$ the sector $C_i^{\lambda_m}$, with $i = \mathcal{S}.\min()$, i.e., with a non-swept homology class α of minimal norm. Hence, we have $\lambda_{m+1} = \text{key}[i]$. We update F_i^h and F_i^v as follows. If $\alpha = (x_1^h + 1)[c_i] + y_1^h[c_{i+1}]$, then we remove (x_1^h, y_1^h) from the bottom of F_i^h and push $(x_1^h + 1, y_1^h)$ on its top. Likewise, if $\alpha = x_1^v[c_i] + (y_1^v + 1)[c_{i+1}]$, we remove (x_1^v, y_1^v) from the bottom of F_i^v and push $(x_1^v, y_1^v + 1)$ on its top. We do both if $\alpha = (x_1^h + 1)[c_i] + y_1^h[c_{i+1}] = x_1^v[c_i] + (y_1^v + 1)[c_{i+1}]$. Clearly, the updated dequeues contain the required lattice points with respect to the sweeping value $\lambda = \lambda_{m+1}$. We then update $\mathcal{S}$ by inserting i with its new key resulting from the updates of F_i^h and F_i^v. We finally add λ_{m+1} to Λ. By Corollary 13 and Theorem 1, $\mathcal{S}$ contains $O(n)$ items. The running time for a sweeping step is thus $O(\log n)$ time for interacting with $\mathcal{S}$ plus constant time for updating F_i^h, F_i^v and Λ. We can thus compute the first k values of the length spectrum in $O(n^2 \log\log n + k \log n)$ time. ◀

7 Deciding equality of length spectra

We now present an application of Algorithm 1 to the following decision problem: given two weighted graphs (G, w) and (G', w') embedded on tori do they have the same length spectra? This question actually covers two problems: the equality of the marked and of the unmarked length spectrum. As we show, the former reduces to the linear equivalence of polyhedral norms which has a straightforward quadratic time solution. In contrast the latter reduces to the problem of polynomial identity testing which is only known to be in the co-RP complexity class [21]. In particular, this problem is in co-NP. For unweighted graphs however, the complexity is deterministic polynomial.

We first aim to compare the length spectrum as maps from $H_1(T;\mathbb{Z}) \to \mathbb{R}$. We are given two weighted graphs (G, w) and (G', w') embedded on tori T and T' respectively. We say that (G, w) and (G', w') have the same *marked spectrum* if there exists a homeomorphism $\phi : T \to T'$ such that for all $\gamma \in H_1(T;\mathbb{Z})$ we have $N_{G',w'}(\phi_*(\gamma)) = N_{G,w}$ where $\phi_*(\gamma)$ denotes the class $[\phi(c)]$ where c is a curve representative of γ.

▶ **Theorem 18.** *Let (G, w) and (G', w') be two weighted graphs cellularly embedded on tori T and T', each with complexity bounded by n. Then there is an algorithm that answers whether (G, w) and (G', w') have the same marked spectrum in time $O(n^2 \log\log n)$.*

Now we consider the more delicate question of comparing unmarked spectra. That is, we want to decide whether the list of values $\{N_{G,w}(\alpha) : \alpha \in H_1(T;\mathbb{Z})\}$ and $\{N_{G',w'}(\alpha') : \alpha' \in H_1(T';\mathbb{Z})\}$ coincide where each value comes with multiplicity according to the number of homology classes that realize this length. This equality of unmarked length spectra is always decidable and we show that it belongs to the co-RP complexity class, i.e. our algorithm can detect if the unmarked spectra of (G, w) and (G', w') are different in random polynomial time. For this specific test, we need to have access to all integral linear relations between the weights at once. That is, our algorithm needs to have access to the $\mathbb{Q}$-vector space $\{(x_e)_{e\in E(G)} \in \mathbb{Q}^{E(G)} : \sum_e x_e w_e = 0\}$. We assume that the weights are given in the following form : we are given r real numbers $o = (o_1, o_2, \ldots, o_r)$ that do not satisfy any integral linear relations, and for each edge $e \in E(G)$ its weight is given as a linear combination of these real numbers with integral coefficients $w_e = w_{e,1}o_1 + \ldots + w_{e,r}o_r$. We call *complexity* of these weights the sum $\|w\|_o := \sum_{e\in E(G)} \sum_{i=1}^{r} |w_{e,i}|$. Note that this complexity depends on the choice of the numbers $o_1, \ldots, o_r$ and not only on the values w_e as real numbers.

▶ **Theorem 19.** *Let (G, w) and (G', w') be two weighted graphs cellularly embedded on tori T and T', each with complexity bounded by n, where each weight is specified as $w_e = w_{e,1}o_1 + \ldots + w_{e,r}o_r$ with $w_{e,i} \in \mathbb{Z}$ and $o_1, \ldots, o_r$ are r given real numbers. There is an algorithm to decide whether (G, w) and (G', w') have different (unmarked) spectra that runs in random polynomial time in the total input size $n + \log(\|w\|_o + \|w'\|_o)$. Moreover, for fixed r, there is a deterministic algorithm that runs in time $O(n^2 \cdot (\|w\|_o + \|w'\|_o)^r)$.*

Let us explain how to deduce Theorem 4 from Theorems 18 and 19. We emphasize that Theorem 19 allows us to deduce deterministic polynomial time only in the unweighted case. Even with rational weights we are not aware of a deterministic polynomial time algorithm.

Proof of Theorem 4. The case of marked spectrum is simply a particular case of Theorem 18, where we use the fact that the unit ball can be computed in quadratic time in the unweighted case. For the equality of unmarked spectrum, we have $r = 1$ and $o_1 = 1$. The second part of Theorem 19 hence gives deterministic polynomial time in $O(n^2 \cdot (\sum_{e\in E(G)\cup E(G')} 1)) = O(n^3)$. ◀

References

1 Peter Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, 1992.

2 Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM Journal on Computing*, 42(4):1542–1571, 2013.

3 Sergio Cabello, Éric Colin de Verdière, and Francis Lazarus. Algorithms for the edge-width of an embedded graph. *Computational Geometry*, 45(5-6):215–224, 2012.

4 Sergio Cabello and Bojan Mohar. Finding shortest non-separating and non-contractible cycles for topologically embedded graphs. *Discrete & Computational Geometry*, 37(2):213–235, 2007.

5 Marcos Cossarini. *Discrete surfaces with length and area and minimal fillings of the circle*. PhD thesis, Instituto de Matemática Pura e Aplicada (IMPA), 2018.

6 Éric Colin de Verdière and Jeff Erickson. Tightening nonsimple paths and cycles on surfaces. *SIAM Journal on Computing*, 39(8):3784–3813, 2010.

7 Matthijs Ebbens and Francis Lazarus. Computing the length spectrum of combinatorial graphs on the torus. 38th International Symposium on Computational Geometry: Young Researchers Forum, 2022.

8 David Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 599–608, 2003.

9 Jeff Erickson. Combinatorial optimization of cycles and bases. In Afra Zomorodian, editor, *Advances in Applied and Computational Topology*, Proceedings of Symposia in Applied Mathematics, 2012.

10 Jeff Erickson and Sariel Har-Peled. Optimally cutting a surface into a disk. *Discrete & Computational Geometry*, 31(1):37–59, 2004.

11 Giuseppe F. Italiano, Yahav Nussbaum, Piotr Sankowski, and Christian Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 313–322, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/1993636.1993679.

12 Philip N Klein. Multiple-source shortest paths in planar graphs. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 146–155, 2005.

13 Martin Kutz. Computing Shortest Non-Trivial Cycles on Orientable Surfaces of Bounded Genus in Almost Linear Time. In *22nd Annual ACM Symposium on Computational Geometry*, pages 430–437, 2006.

14 Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*, volume 10. JHU press, 2001.

15 Jean-Pierre Otal. Le spectre marqué des longueurs des surfaces à courbure négative. *Annals of Mathematics*, 131(1):151–162, 1990.

16 Hugo Parlier. Interrogating surface length spectra and quantifying isospectrality. *Mathematische Annalen*, 370(3):1759–1787, 2018. URL: https://math.uni.lu/parlier/Papers/Isospectral2016-11-07.pdf.

17 Mikael de la Salle. On norms taking integer values on the integer lattice. *Comptes Rendus Mathématique*, 354(6):611–613, 2016.

18 Abdoul Karim Sane. Intersection norms and one-faced collections. *Comptes Rendus Mathématique*, 358(8):941–956, 2020.

19 Alexander Schrijver. On the uniqueness of kernels. *Journal of Combinatorial Theory, Series B*, 55(1):146–160, 1992.

20 Alexander Schrijver. Graphs on the torus and geometry of numbers. *Journal of Combinatorial Theory, Series B*, 58(1):147–158, 1993.

21 Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. Assoc. Comput. Mach.*, 27:701–717, 1980. doi:10.1145/322217.322225.

22 John Stillwell. *Classical topology and combinatorial group theory*, volume 72. Springer Science & Business Media, 1993.

23 Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *Journal of Combinatorial Theory, Series B*, 48(2):155–177, 1990.

24 William P. Thurston. A norm for the homology of 3-manifolds. In *Mem. Amer. Math. Soc.*, volume 339, pages 99–130. AMS, 1986.

25 Marie-France Vignéras. Isospectral and non-isometric Riemannian manifolds. *Ann. Math. (2)*, 112:21–32, 1980. doi:10.2307/1971319.