Minimum Weight Bases

Francis Lazarus and Arnaud de Mesmay

October 26, 2017

Contents

In the second lecture we saw that a graph could be associated with a vector space, called the cycle space. We will see that this cycle space can be extended to surfaces giving birth to the *first homology group*. We also introduced the fundamental group of a graph or of a surface in another lecture. Hence, we now have two group structures that encode the topology of a space X, where X is either a graph or a surface. These structures are both generated by closed walks in the graph of X and we call a **basis** any generating set with the minimum number of closed walks. In order to derive a more informative notion of minimality we assume that the edges of the considered graph have a positive weight. This allows to define the weight of a closed walk as the sum of its edge weights (counted with multiplicity). A **minimum weight basis** is then a basis such that the sum of the weights of its members is minimum. The computation of minimum weight bases has received much attention when X is a graph and was studied more recently for combinatorial surfaces. Good references on the subject include a comprehensive survey on cycle bases in graphs by Kavitha et al. [KLM+09] and another survey on optimization of cycles and bases on surfaces by Erickson [Eri12]. We shall use the qualifiers *minimum* and *shortest* interchangeably to designate a walk, tree or subgraph of minimum weight.

1 Minimum Basis of the Fundamental Group of a Graph

Let G be a connected graph with basepoint v and let $|.|: E \to \mathbb{R}_+$ be a weight function. The fundamental group $\pi_1(G, v)$ is a free group whose rank is the number of chords of any spanning tree of G, which is $1 - n + m$, where n and m are respectively the number of vertices and edges of G. Indeed, as we saw, every chord e of a spanning tree T gives rise to a loop $\gamma_{v,e}^T$ obtained by connecting v to each endpoint of the chord using paths in the tree, and these loops form a basis of $\pi_1(G, v)$. Not all bases arise this way but a minimum one may indeed be obtained by this construction. For this, we take for T a **shortest path tree** with root v: every vertex w of G is linked to v by a path in T whose weight is minimum among all the paths from v to w in G. When all the weights are equal a shortest path tree can be computed by a breadth-first search traversal in time $O(m)$. In the general case, one may use Dijkstra's algorithm [CLRS09] to compute a shortest path tree in $O(m + n\log n)$ time. Remark that $\gamma_{v,e}^T$ is a shortest loop through the chord e.

Theorem 1.1. *The basis of $\pi_1(G, v)$ associated with a shortest path tree with root v is a minimum weight basis.*

The following proof is based on an purely algebraic preliminary lemma. First note that a free group F over a set $(x_1, x_2, \ldots, x_r)$ gives rise to a *free Abelian group* (this is the same a free $\mathbb{Z}$-module) F^{ab} by making all the generators commute. Hence, if we let R be the set of relations $\{x_i x_j = x_j x_i\}_{1\le i<j\le r}$, a presentation for F^{ab} is $<\{x_1, x_2, \ldots, x_r\} \mid R>$. We thus have a quotient $F \twoheadrightarrow F^{ab} = F/<R>$ and we denote by $[x] \in F^{ab}$ the image of any $x \in F$. Note that $[x]$ can be uniquely written as a linear combination of the $[x_i]$'s.

Lemma 1.2. *Let $(x_1, x_2, \ldots, x_r)$ and $(y_1, y_2, \ldots, y_r)$ be two bases of a free group F. Denote by $y_j(x_1, x_2, \ldots, x_r)$ the expression of y_j in terms of the basis $(x_1, x_2, \ldots, x_r)$. Then, there exists a permutation σ of $\{1, \ldots, r\}$ such that for each i the coefficient of $[x_i]$ in $[y_{\sigma(i)}(x_1, x_2, \ldots, x_r)]$ is nonzero.*

PROOF. The automorphism of F defined by $x_i \mapsto y_i$, $1 \le i \le r$, quotients to an automorphism of F^{ab}. Let c_{ij} be the coefficient of $[x_j]$ in $[y_i(x_1, x_2, \ldots, x_r)]$. Viewing F^{ab} as a free $\mathbb{Z}$-module over the $[x_i]$'s, the matrix $(c_{ij})_{1\le i,j\le r}$ of this automorphism has nonzero determinant. It follows that at least one term $\prod_{1\le i\le r} c_{i\sigma(i)}$ of the usual Leibnitz expansion of the determinant must be nonzero. This implies the lemma. □

PROOF OF THEOREM 1.1. Let T be a shortest path tree from v. We denote by $e_1, e_2, \ldots, e_r$ the chords of T in G. Let $(b_1, b_2, \ldots, b_r)$ be a basis for $\pi_1(G, v)$. According to the preliminary lemma, there is a permutation σ of $\{1, \ldots, r\}$ such that the coefficient of $[\gamma_{v,e_i}^T]$ in $[b_{\sigma(i)}]$ is nonzero. It follows that $b_{\sigma(i)}$ goes through e_i, hence is at least as long as γ_{v,e_i}^T by the remark before the theorem. As a direct consequence $\sum_i |b_i| \ge \sum_i |\gamma_{v,e_i}^T|$. □

2 Minimum Basis of the Cycle Space of a Graph

As we saw, the set of Eulerian subgraphs $Z(G)$ of a connected graph G can be given a vector space structure over the coefficient field $\mathbb{Z}/2\mathbb{Z}$. We also observed that a basis could be obtained from any spanning tree T of G by considering for each chord e of T the cycle γ_e^T composed of e and the path in T connecting e's endpoints. Such a basis is called a *fundamental cycle basis*. As opposed to the case of the fundamental group, a minimum weight basis of the cycle space is not always a fundamental cycle basis. The counterexample in Figure 1 was found by Hartvigsen and Mardon [HM93]. In general, looking for the minimum weight *fundamental* basis is NP-hard [DPeK82].

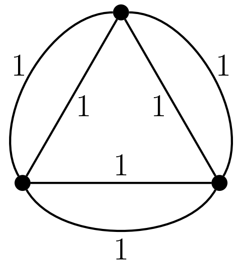

Figure 1: Each spanning tree in this graph is a path of length 2. The corresponding fundamental basis is composed of two cycles of length 2 and two cycles of length 3 leading to a fundamental cycle basis of total weight 10. However, a minimum weight basis of total weight 9 is given by the three outer cycles of length 2 and the central triangle.

However, Horton [Hor87] proved that computing a minimum weight basis with $\mathbb{Z}/2\mathbb{Z}$ coefficients can be done in polynomial time. His algorithm is based on the greedy algorithm over combinatorial structures called *matroids*.

2.1 The Greedy Algorithm

As a vector space, $Z(G)$ inherits a **matroid** structure. A matroid is indeed an abstraction of a vector space that only retains linear dependencies. It is defined by a **ground set** S (intuitively the set of vectors) and a nonempty **family of independent sets** $\mathscr{I} \subset 2^S$ that satisfies

- **the hereditary property:** $J \in \mathscr{I}$ and $I \subset J$ implies $I \in \mathscr{I}$, and
- **the exchange property:** $I, J \in \mathscr{I}$ and $|I| < |J|$ implies that $I \cup \{x\} \in \mathscr{I}$ for some $x \in J \setminus I$.

A **basis** is just a maximally independent set. By the exchange property, all the bases have the same cardinality. Matroid theory was introduced by Hassler Whitney (1935) and has many applications including combinatorial optimization, discrete geometry, etc. When the elements of the ground set are weighted, there is a famous *greedy algorithm* that determines a minimum weight basis. It works as follows: maintain an independent set starting from the empty set, and iteratively add an element x to the current set I if $I \cup \{x\}$ is independent and if x has minimum weight among such elements. The algorithm stops when no x can be found, *i.e.* when I is a basis. In practice, the elements of S are scanned in increasing order of weights, so that each time an x is found such that $I \cup \{x\}$ is independent it can be added to the current I. The whole set S is thus scanned only once during the algorithm.

Theorem 2.1. *The greedy algorithm returns a minimum weight basis.*

PROOF. Let $(x_1, x_2, \dots, x_r)$ be the basis returned by the greedy algorithm, where x_i is the ith inserted element. By the choice of each element we have $|x_1| \le |x_2| \le \cdots \le |x_r|$, where $|x|$ is the weight of x. Consider any other basis $(y_1, y_2, \dots, y_r)$ indexed in non-decreasing order: $|y_1| \le |y_2| \le \cdots \le |y_r|$. Suppose by way of contradiction that there is some index i such that $|y_i| < |x_i|$ and choose such i as small as possible. Then, by the exchange property we can find $y \in \{y_1, \dots, y_i\}$ such that $\{x_1, \dots, x_{i-1}, y\}$ is independent. Since $|y| \le |y_i| < |x_i|$ this would contradict the choice of x_i. It follows that $|y_i| \ge |x_i|$ for all i, implying that $(x_1, x_2, \dots, x_r)$ has minimum weight. □

Since the cycle space contains 2^r cycles, the greedy algorithm per se does not seem very efficient. In order to restrict the search of a new basis element at each step of the algorithm, Horton [Hor87] gave a characterization of the cycles that may belong to a minimum weight basis.

Lemma 2.2. *Suppose $b = c + d$ is a cycle of a basis B of $Z(G)$. Then either $B \setminus \{b\} \cup \{c\}$ or $B \setminus \{b\} \cup \{d\}$ is a basis.*

PROOF. If c and d were both in the linear span of $B \setminus \{b\}$, then so would b. □

Corollary 2.3. *Assuming positive weights, the cycles of a minimum weight basis are simple.*

PROOF. Suppose that b is a non-simple cycle of a minimum weight basis B. Then b can be written as the sum $b = c + d$ of two edge disjoint cycles. In particular, b is longer than c or d. By the preceding lemma, we can replace b by c or d in B to get a shorter basis, contradicting the minimality of B. □

Note: if some of the weights cancel, then basically the same proof shows the existence of a minimum weight basis with simple cycles only.

Lemma 2.4. *Let b be a cycle of a minimum weight basis. Let p and q be two edge disjoint paths such that $b = p \cdot q^{-1}$. Then p or q is a shortest path.*

PROOF. Let t be a shortest path from the common initial vertex of p and q to their common last vertex. With a little abuse of notation, we can write $b = p \cdot t^{-1} + t \cdot q^{-1}$. By Lemma 2.2, b must be no longer than $p \cdot t^{-1}$ or $t \cdot q^{-1}$, implying with Corollary 2.3 that either q or p is a shortest path. □

Corollary 2.5. *Let v be a vertex of a cycle b of a minimum weight basis. Then b decomposes into $p \cdot a \cdot q^{-1}$ where a is an arc and p, q are two shortest paths with v as initial vertex.*

PROOF. Consider the arc sequence $(a_1, a_2, \dots, a_k)$ of b with v the origin vertex of a_1 and the target of a_k. Let i be the maximal index such that $(a_1, a_2, \dots, a_i)$ is a shortest path. Then $b = (a_1, a_2, \dots, a_i) \cdot a_{i+1} \cdot (a_{i+2}, \dots, a_k)$ and the previous lemma implies that $(a_{i+2}, \dots, a_k)$ is a (possibly empty) shortest path $\square$

When there is a unique shortest path between every pair of vertices, this corollary allows us to reduce the number of scanned cycles at each addition step of the greedy algorithm to nm cycles, one for each (vertex, edge) pair. For the rest of this section, we assume uniqueness of shortest paths and discuss the general case in the next section. We denote by $\gamma_{v,e}$ the cycle obtained by connecting the endpoints of e with shortest paths to v. By the uniqueness of shortest paths, $\gamma_{v,e} = \gamma^T_{v,e}$ where T is the shortest path tree rooted at v.

Proposition 2.6. *Let $G = (V, E)$ be a connected graph with n vertices and m edges and let $r = 1 - n + m$ be the rank of its cycle space. A minimum weight basis of $Z(G)$ can be computed in $O(n^2 \log n + r^2 nm) = O(nm^3)$ time.*

PROOF. By Corollary 2.5, we can restrict the scan step of the greedy algorithm to the cycles $\gamma_{v,e}$ with $(v, e) \in V \times E$. For each vertex v, we compute a shortest path tree T in $O(n \log n + m)$ time using Dijkstra's algorithm. There are r nontrivial cycles of the form $\gamma^T_{v,e}$, each of size $O(n)$. Their computation and storage for all the vertices v thus requires $O\big(n(n \log n + m + rn)\big)$ time. They can be sorted according to their length in $O(rn \log(rn))$ time. In order to check if a cycle is independent of the current family of basis elements, we view a cycle as a vector in $(\mathbb{Z}/2\mathbb{Z})^E$. We use Gauss elimination to maintain the current family in row echelon form. This family has at most r vectors and testing a new vector against this family by Gauss elimination needs $O(rm)$ time. The cumulated time for testing independence is thus $O(r^2 nm)$. The whole greedy algorithm finally takes time

$$O\big(n(n \log n + m + rn) + rn \log(rn) + r^2 nm\big) = O(n^2 \log n + r^2 nm).$$

$\square$

Note that the above scan can be further reduced by discarding the cycles $\gamma_{v,e}$ that are not simple. We can also decompose a cycle into a linear combination of a fixed fundamental basis associated to a tree. The decomposition of a cycle is just given by its trace over the chords of that tree. This allows to represent the current family of basis elements by a matrix of size $r \times r$ instead of $r \times m$. Further improvements were proposed [KMMP04, KMMP08, MM09], often based on randomization. In particular, the algorithm by Kavitha et al. [KMMP08] runs in $O(m^2 n + mn^2 \log n)$ time. Using integer coefficients rather than $\mathbb{Z}/2\mathbb{Z}$ gives a more general notion of cycle space. However, this space does not form a matroid in general and the greedy algorithm cannot be applied anymore. The status of the computation of a minimal weight cycle basis with integer coefficients is still unknown.

3 Uniqueness of Shortest Paths

The proof of Proposition 2.6 is based on the uniqueness of shortest paths. In fact, the proof can be adapted to show that the same algorithm works even if we do not assume that there is a unique shortest path between every pair of vertices. See [Hor87] or [Laz14, Lem. 1.6.7]. It may happen for other applications that we strongly need uniqueness to ensure correctness of the algorithms. We usually get the uniqueness by a *perturbation schema*, where the weight of each simple path is replaced by a slightly different one. Let P_{xy} be the set of simple paths with minimal unperturbed weight between vertices x and y. The perturbation should be such that P_{xy} contains a unique path of minimum perturbed weight. In other words, the aim is to get an order on each P_{xy} so that we can choose the smallest path as the unique shortest path. This can be achieved by adding an infinitesimal weight of the form (i) $\varepsilon^{c(e)}$ or (ii) $c(e)\varepsilon$ to every edge e, where $\varepsilon > 0$ is some arbitrarily small number and $c(e)$ is an appropriately chosen coefficient.

Using the exponential form (i) we can simply choose pairwise distinct edge coefficients, for example the edge indices, assuming that they are indexed from 1 to m. This way, distinct paths are perturbed by distinct polynomials in ε and get distinct weights for ε small enough. We can view the polynomials as bit vectors of length m where a 1 coordinate at index i indicates the presence of the monomial ε^i. The ordering in P_{xy} is simply the lexicographic ordering on the bit vectors. This perturbation schema would a priori require an extra $O(m)$ time for comparing path lengths. Cabello et al. [CCE13, Sec. 6.2] propose to reduce the comparison time to $O(\log m)$ using some sophisticated data structure. However, their algorithm assumes that two paths need to be compared only when they intersect along a common prefix.

We can avoid this restriction using the linear form (ii), that is when the weight of an edge e is perturbed by $c(e)\varepsilon$. The perturbation of a path is now ε times the sum of its edge coefficients. Choosing the edge coefficients such that there is a unique minimum weight sum in each P_{xy} is more tricky than for the form (i). Cabello et al. [CCE13, Sec. 6.1] propose the following random perturbation schema based on the Isolating Lemma of Mulmuley et al. [MVV87, Lem. 1].

Lemma 3.1 (Isolating –)**.** *Let $\mathscr{I}$ be an arbitrary family of subsets of $\{1,\dots,m\}$. For a vector $c=(c_1,c_2,\dots,c_m)$ of m integers and for $I\in\mathscr{I}$, we put $c(I)=\sum_{i\in I}c_i$. Choosing c uniformly at random in $\{1,\dots,M\}^m$, the probability that $c(I)$ is minimized by a unique $I\in\mathscr{I}$ is at least $1-m/M$.*

PROOF. We suppose that $\mathscr{I}$ contains at least two subsets, since otherwise the lemma is trivial. For $i\in\{1,\dots,m\}$ we set

$$\mathscr{I}_i^+=\{I\in\mathscr{I}\mid i\in I\}\quad\text{and}\quad\mathscr{I}_i^-=\{I\in\mathscr{I}\mid i\notin I\}$$

Suppose that none of $\mathscr{I}_i^+$ and $\mathscr{I}_i^-$ is empty. Note that the quantities $\min_{I\in\mathscr{I}_i^+}c(I)-c_i$ and $\min_{I\in\mathscr{I}_i^-}c(I)$ do not depend on c_i. Fixing all the coefficients c_j, for $j\neq i$, the constant $\min_{I\in\mathscr{I}_i^-}c(I)-(\min_{I\in\mathscr{I}_i^+}c(I)-c_i)$ equals c_i with probability at most $1/M$. It follows that $\min_{I\in\mathscr{I}_i^+}c(I)=\min_{I\in\mathscr{I}_i^-}c(I)$ holds with unconditional probability at most $1/M$. Hence, with probability at least $1-m/M$, $\min_{I\in\mathscr{I}_i^+}c(I)\neq\min_{I\in\mathscr{I}_i^-}c(I)$ for all i such that $\mathscr{I}_i^+$ and $\mathscr{I}_i^-$ are both nonempty. Consider a vector c such that this occurs and let $I_0 \in \mathscr{I}$ for which $c(I_0)$ is minimum. Then, any other $J \in \mathscr{I}$ must differ from I_0 by some index i. If $i \in I_0$ and $i \notin J$ then $c(I_0) = \min_{I\in\mathscr{I}} c(I) = \min_{I\in\mathscr{I}_i^+} c(I)$ while $c(J) \geq \min_{I\in\mathscr{I}_i^-} c(I)$. Since $\min_{I\in\mathscr{I}_i^+} c(I) \neq \min_{I\in\mathscr{I}_i^-} c(I)$ we deduce $c(J) > c(I_0)$. Likewise, we again obtain $c(J) > c(I_0)$ if $i \notin I_0$ and $i \in J$. □

Lemma 3.2. *Choose for each of the m edges of an edge weighted graph G an integral coefficient in $\{1,\dots,m^4\}$ uniformly and independently at random. Consider the linear perturbation schema (ii) as described above. With probability at least $1-\frac{1}{2m}$, there is a unique shortest path between any pair of vertices.*

PROOF. For each pair $\{x, y\}$ of vertices, Let $\mathscr{I}_{xy}$ be the family of subsets of edge indices corresponding to the paths in P_{xy}. Applying Lemma 3.1 to $\mathscr{I}_{xy}$, we deduce that with probability at least $1-1/m^3$ there is a unique shortest path between x and y for the perturbed weights. There are $n \leq m$ vertices in G (we may assume that G is not a tree). Hence, the $\binom{n}{2}$ pairs of vertices are each connected by a unique shortest path with probability at least $1-\binom{n}{2}/m^3 \geq 1-\frac{1}{2m}$. □

We shall turn to the computation of minimum bases on surfaces. We first extend the notion of cycle space to surfaces.

4 First Homology Group of Surfaces

4.1 Back to Graphs

First recall that the cycle space $Z(G)$ of a graph $G = (V, E)$ is the space of its Eulerian subgraphs. One can define such subgraphs thanks to the **boundary operator**. This operator δ_1 sends any edge to the mod 2 sum of its endpoints. In particular, if e is a loop-edge, $\delta_1 e = 0$. By linear extension, δ_1 defines a linear map from the vector space $(\mathbb{Z}/2\mathbb{Z})^E$ of formal mod 2 sum of edges to the space $(\mathbb{Z}/2\mathbb{Z})^V$ of mod 2 sum of vertices. Viewing a subgraph as a mod 2 sum of its edges, it is easily seen that Eulerian subgraphs correspond to the mod 2 sum of edges with empty boundary. In other words,

$$Z(G) = \ker \delta_1.$$

We define the **mod 2 abelianization** of a group A as the quotient $A/S(A)$ by the subgroup $S(A)$ generated by its squares. Note that $S(A)$ is normal and contains the **derived subgroup** $[A, A]$ generated by the **commutators** $[a, b] = aba^{-1}b^{-1}$. Indeed, one check that

$$[a, b] = (aba^{-1})^2a^2(a^{-1}b^{-1})^2 \quad \text{and} \quad asa^{-1} = s[s^{-1}, a]$$

Hence, if s is a product of squares, so is any conjugate asa^{-1}. We can now relate the cycle space with the fundamental group of a graph thanks to the following mod 2 version of the Hurewicz theorem.

Proposition 4.1. *For any vertex v of a connected graph G the cycle space $Z(G)$ is isomorphic to the mod 2 abelianization of $\pi_1(G,v)$.*

PROOF. Denote by $\mathscr{L}$ the set of loops of G with basepoint v. Consider the map $\varphi : \mathscr{L} \to Z(G)$ defined by $(a_1, a_2, \ldots, a_k) \mapsto \sum_{i=1}^{k} a_i$, where the coefficient in the sum are taken modulo 2. Adding or removing a spur in a loop does not change its image by φ. The map φ thus quotients to a morphism $\overline{\varphi} : \pi_1(G,v) \to Z(G)$. Let T be a spanning tree of G and let $C = E(G) \setminus E(T)$ be its set of chords. We know that $Z(G)$ is generated by the cycles $\{\gamma_e^T\}_{e\in C}$. Since $\varphi(\gamma_{v,e}^T) = \gamma_e^T$, the map $\overline{\varphi}$ is onto. Let $\gamma = \gamma_{v,e_1}^T \cdot \gamma_{v,e_2}^T \cdots \gamma_{v,e_k}^T$ be a representative of some element of $\pi_1(G,v)$ written over the basis $\{\gamma_{v,e}^T\}_{e\in C}$. Then $\varphi(\gamma) = \sum_{e\in C} n_e \gamma_e^T$ where n_e is the cumulated exponent of $\gamma_{v,e}^T$ in γ. Hence, the homotopy class of γ belongs to $\ker \overline{\varphi}$ if and only if all the n_e cancel. This is exactly saying that γ belongs to the subgroup $S(\pi_1(G,v))$ of $\pi_1(G,v)$. We thus have

$$Z(G) \simeq \pi_1(G,v)/\ker \overline{\varphi} = \pi_1(G,v)/S(\pi_1(G,v))$$

□

Exercise 4.2. Given a product w in the generators $\{x_1, x_2, \ldots, x_r\}$ (and their inverses) of a group Γ, show that $w = x_1^{n_1} \cdot x_2^{n_r} \cdots x_1^{n_r} \cdot p$ where each n_i is the cumulated exponent of x_i in w and p is a product of commutators. (It might be useful to notice the relation $ba = ab[b^{-1}, a^{-1}]$.) Deduce that $S(\Gamma)$ is equal to the set of products whose cumulated exponents are all even.

4.2 Homology of Surfaces

The graph G of a combinatorial surface M has its own cycle space $Z(G)$. However, a topological surface may have distinct cellularly embedded graphs with non-isomorphic cycle spaces. In order to get a topologically invariant notion of cycle space, we further quotient $Z(G)$ by identifying cycles that bound together a subset of faces of M. More formally, let $C_2(M) := (\mathbb{Z}/2\mathbb{Z})^F$ be the vector space of subsets of the set F of faces of M. The elements of $C_2(M)$ are called **2-chains**. The boundary $\partial_2 f$ of a face $f \in F$ is the mod 2 sum of the edges of its facial walk. It is clearly a cycle of $Z(G)$, meaning that $\partial_1 \partial_2 f = 0$. This boundary ∂_2 extends linearly to a boundary operator $\delta_2 : C_2(M) \to Z(G)$. Two cycles $c, d \in Z(G)$ are said **homologous**, which we write $[c] = [d]$, if their mod 2 sum is the boundary of some 2-chain $\sigma \in C_2(M)$: $c - d = \partial_2 \sigma$. We can now define the **first homology group** of M as the space of homology classes:

$$H_1(M) := \ker \delta_1 / \mathrm{Im}\, \delta_2.$$

The fact that this homology group is indeed a topological invariant is an immediate consequence of the invariance of the fundamental group and of the following mod 2 version of the Hurewicz theorem for surfaces.

Proposition 4.3. *For any vertex v of a connected map M the first homology group $H_1(M)$ is isomorphic to the mod 2 abelianization of $\pi_1(M,v)$.*

PROOF. Let G be the graph of M. As in the proof of Proposition 4.1, denote by $\mathscr{L}$ the set of loops of G with basepoint v and by $\varphi : \mathscr{L} \to Z(G)$ the mapping defined by $\varphi(a_1, a_2, \ldots, a_k) = \sum_{i=1}^{k} a_i$. The composition $[\varphi] : \mathscr{L} \to Z(G) \to H_1(M)$ is compatible with elementary homotopies in M. This is obvious for the addition or removal of a spur. If $\lambda \cdot p \cdot \mu \mapsto \lambda \cdot q \cdot \mu$ is an elementary homotopy with $p \cdot q^{-1}$ the facial walk of a face f, then $\varphi(\lambda \cdot p \cdot \mu) - \varphi(\lambda \cdot q \cdot \mu) = \partial_2 f \in \operatorname{Im} \partial_2$. Whence $[\varphi(\lambda \cdot p \cdot \mu)] = [\varphi(\lambda \cdot q \cdot \mu)]$ in $H_1(M)$. It follows that $[\varphi]$ descends to the quotient $\overline{\varphi} : \pi_1(M, v) \to H_1(M)$. On the other hand, homotopic loops in G are homotopic in M so that we have an onto morphism $\pi_1(G, v) \twoheadrightarrow \pi_1(M, v)$. We also know from the proof of Proposition 4.1 that the morphism $\pi_1(G, v) \twoheadrightarrow Z(G)$ is onto. We thus have two equal compositions $\pi_1(G, v) \twoheadrightarrow Z(G) \twoheadrightarrow H_1(M)$ and $\pi_1(G, v) \twoheadrightarrow \pi_1(M, v) \xrightarrow{\overline{\varphi}} H_1(M)$ implying that $\overline{\varphi}$ is onto.

It remains to prove that $\ker \overline{\varphi}$ is the subgroup $S(\pi_1(M, v))$ generated by the squares of $\pi_1(M, v)$ to conclude that $H_1(M) \simeq \pi_1(M, v) / \ker \overline{\varphi}$ is the mod 2 abelianization of $\pi_1(M, v)$. Since multiplication by 2 gives zero in $H_1(M)$, we have $S(\pi_1(M, v)) \subset \ker \overline{\varphi}$. For the reverse inclusion we consider a loop γ whose homotopy class is in $\ker \overline{\varphi}$, *i.e.* such that $[\varphi(\gamma)] = 0$. Hence, there must be a 2-chain $\sum_j f_j$ such that $\varphi(\gamma) = \sum_j \partial_2 f_j$. For each j, we choose a vertex v_j incident to f_j and we let p_j be the facial walk of f_j starting at v_j. Using the path γ^T_{v, v_j} from v to v_j in T we form the loop $\gamma_j := \gamma^T_{v, v_j} \cdot p_j \cdot (\gamma^T_{v, v_j})^{-1}$ with basepoint v. On the one hand, since $\varphi(\gamma_j) = \partial_2 f_j$, we have $\varphi(\gamma) = \varphi(\prod_j \gamma_j)$ in $Z(G)$. Equivalently, $\varphi(\gamma) + \varphi(\prod_j \gamma_j) = 0$. By Proposition 4.1, the homotopy class of $\gamma \cdot \prod_j \gamma_j$ is in $S(\pi_1(G, v))$. It is thus in $S(\pi_1(M, v))$ viewed as a loop in M. On the other hand, since each γ_j is contractible in M, the loops $\gamma \cdot \prod_j \gamma_j$ and γ are homotopic in M. It follows that the homotopy class of γ is in $S(\pi_1(M, v))$. □

Corollary 4.4. *Let M be a combinatorial surface of genus g without boundary. We have*

$$H_1(M) \simeq \begin{cases} (\mathbb{Z}/2\mathbb{Z})^{2g} & \text{if } M \text{ is orientable, and} \\ (\mathbb{Z}/2\mathbb{Z})^{g} & \text{otherwise.} \end{cases}$$

PROOF. If M is orientable, we know that its fundamental group as combinatorial presentation $\pi_1 \simeq < a_1, b_1, \ldots, a_g, b_g \mid [a_1, b_1] \cdots [a_g, b_g] >$. By Proposition 4.3, we have $H_1(M) \simeq \pi_1 / S(\pi_1)$. Since $[a_1, b_1] \cdots [a_g, b_g] \in S(\pi_1)$, we also have $\pi_1 / S(\pi_1) \simeq F_{2g} / S(F_{2g})$ where $F_{2g} := < a_1, b_1, \ldots, a_g, b_g \mid - >$ is the free group over $\{a_1, b_1, \ldots, a_g, b_g\}$. Now, it is easily seen that the mod 2 abelianization of a free group of rank r is the $\mathbb{Z}/2\mathbb{Z}$-vector space of dimension r, whence $H_1(M) \simeq F_{2g} / S(F_{2g}) \simeq (\mathbb{Z}/2\mathbb{Z})^{2g}$. A similar proof holds when M is non-orientable. □

5 Minimum Basis of the Fundamental Group of a Surface

Let M be a combinatorial surface with graph G. As in Section 1, we assume that the edges of G are positively weighted. Given a vertex v of M, a *minimum weight basis* of $\pi_1(M, v)$ is a set of loops with basepoint v whose homotopy classes form a basis of

$\pi_1(M,v)$ and whose total weight is minimum. Erickson and Whittlesey [EW05] have proposed a simple algorithm to compute a minimum weight basis. We first describe how to formally cut M along a subgraph of G.

5.1 Dual Maps and Cutting

The **dual map** M^* of M is obtained by inverting the roles of vertices and edges in M. Its graph G^* is the dual graph of G. If H is a subgraph of G, we denote by H^* the subgraph of G^* induced by the edges dual to the edges of H. We also denote by $M\backslash\backslash H$ the map with boundary obtained by cut opening M along H. It boils down to double the edges of H, updating the rotation system of M to include these new edges. Equivalently, if one views M as a polygonal schema, *i.e.* as a gluing of polygons by pairwise identifications of their sides, cutting along H amounts to forbid the identification between the sides that correspond to edges in H. In the dual map, the effect is to delete the corresponding dual edges. Hence,

Lemma 5.1. *The adjacency graph of the faces of $M\backslash\backslash H$ is $G^* - E(H^*)$. In particular, the connected components of $M\backslash\backslash H$ and of $G^* - E(H^*)$ are in 1-1 correspondence.*

As usual, $E(H^*)$ designates the set of edges of H^*.

5.2 Homotopy Basis Associated with a Tree-Cotree Decomposition

Recall that a tree-cotree decomposition (T, D^*, C) of M is given by a spanning tree T of G, a spanning tree D^* of $G^* - E(T^*)$, and the complementary set of edges $C = E(G)\setminus(E(T)\cup E(D))$.

Lemma 5.2. *If (T, D^*, C) is a tree-cotree decomposition of M, then C contains $2-\chi(M)$ edges. In particular the cycle spaces of the graphs $T\cup C$ and $D^*\cup C^*$ have dimension $2-\chi(M)$*

Proof. The trees T and D^* being spanning we have $|E(T)| = |V(M)|-1$ and $|E(D^*)| = |F(M)|-1$. Thanks to Euler formula, we can write

$$|V(M)|+|F(M)|-\chi(M) = |E(M)| = |E(T)|+|E(D^*)|+|C| = |V(M)|+|F(M)|-2+|C|,$$

whence $|C| = 2-\chi(M)$. $\square$

In analogy with the basis of the fundamental group of a graph associated with a spanning tree, we can associate a basis of the fundamental group of M with a tree-cotree decomposition.

Lemma 5.3. *Let v be a vertex of M, and let (T, D^*, C) be a tree-cotree decomposition of M. The set of loops $\{\gamma^T_{v,c}\}_{c\in C}$ is a basis of $\pi_1(M,v)$.*

PROOF. Since D^* is a tree, the gluing of faces of M along the edges of D is a disk. In other words, $M \backslash\backslash (T \cup C)$ is a disk. Hence, every edge $d \in E(D)$ cuts this disk into two disks. Choose one of those disks. Its boundary writes $(d, e_1, \dots, e_k)$ where each e_i is an edge of $T \cup C$. This boundary is obviously contractible. By inserting a round-trip to v in T at each vertex along this boundary, we see that $\gamma^T_{v,d} \cdot \gamma^T_{v,e_1} \cdots \gamma^T_{v,e_k}$ is contractible. This shows that $\gamma^T_{v,d}$ is in the span of $\{\gamma^T_{v,c}\}_{c \in C}$ since γ^T_{v,e_i} is contractible for every e_i in T. Now, since $\{\gamma^T_{v,e}\}_{e \in E(D) \cup C}$ is a (fundamental) basis of $\pi_1(G, v)$, it is also a generating set for $\pi_1(M, v)$. In turn this generating set is generated by $\{\gamma^T_{v,c}\}_{c \in C}$. Finally, by Lemma 5.2 we note that C contains the minimum number of elements required for a basis of $\pi_1(M, v)$. □

5.3 The Greedy Homotopy Basis

For each chord e of T, the loop $\gamma^T_{v,e}$ is a shortest loop through e with basepoint v and we define the weight of the edge e^* dual to e as

$$w(e^*) = |\gamma^T_{v,e}|,$$

where $|.|$ denotes the given weight function in G. We consider a maximum weight spanning tree K^* of $G^* - E(T^*)$ with respect to the weight function w, and we let C be the set of edges primal to the chords of K^* in $G^* - E(T^*)$. We thus have a tree-cotree decomposition (T, K^*, C) and the set of loops

$$\Gamma := \{\gamma^T_{v,e}\}_{e \in C}$$

is the associated basis of $\pi_1(M, v)$. Following [EW05], we call Γ a **greedy homotopy basis**. The name comes from a greedy computation of the maximum spanning tree K^* which makes the loops in Γ appear in a greedy fashion. It results from Proposition 4.3 that the set of homology classes of the loops in Γ is a basis of $H_1(M)$. A **greedy factor** of a loop ℓ with basepoint v is any loop in Γ which appears with a non-zero coefficient in the decomposition of ℓ in this homology basis.

Lemma 5.4. *The weight $w(e^*)$ of any chord e of T in G is larger or equal to the weights (with respect to $|.|$) of the greedy factors of $\gamma^T_{v,e}$.*

PROOF. The set of chords of T is the disjoint union $E(K) \cup C$. If $e \in C$, then $\gamma^T_{v,e}$ is its own and unique greedy factor and the result is trivial. We now assume that $e \in E(K)$. We put $C_1 := \{c \in C \mid w(c^*) \le w(e^*)\}$ and $C_2 := \{c \in C \mid w(c^*) > w(e^*)\}$. We consider the connected graph $K^*_e := G^* - (E(T^*) \cup C^*_1) = K^* + C^*_2$. We claim that $K^*_e - e^*$ is not connected. Otherwise, e^* would belong to a cycle of K^*_e. This cycle would contain an edge c^* in C^*_2 and exchanging e^* with c^* in K^* would produce a spanning tree with strictly larger weight, contradicting the maximality of K^*. It ensues from Lemma 5.1 that $M \backslash\backslash (T \cup C_1)$ is connected while $M \backslash\backslash (T \cup C_1 \cup \{e\})$ is not. Hence, e appears exactly once in the boundary of each component of $M \backslash\backslash (T \cup C_1 \cup \{e\})$. Considering the formal sum of the faces of one component and its image by the boundary operator, we obtain that $e + \kappa$ is 0-homologous for some chain κ with support in $T \cup C_1$. We conclude that the greedy factors of $\gamma^T_{v,e}$ are contained in $\{\gamma^T_{v,c}\}_{c \in C_1}$, as desired. □

Lemma 5.5. *Let ℓ be a loop with basepoint v in G. Any greedy factor of ℓ has weight at most $|\ell|$.*

PROOF. We consider ℓ as a loop of G and express its homotopy class in the free basis of $\pi_1(G, v)$ associated with the chords of T in G: $\ell \sim \gamma^T_{v,e_1} \cdot \gamma^T_{v,e_2} \cdots \gamma^T_{v,e_k}$. We assume this expression reduced, so that each e_i, $1 \le i \le k$, occurs at least once in ℓ. In particular, $|\ell| \ge w(e_i^*)$. Since any greedy factor of ℓ must occur as a greedy factor of some γ^T_{v,e_i}, we can apply Lemma 5.4 to e_i and conclude. □

We denote by $\gamma_1, \ldots \gamma_{|C|}$ the loops in the greedy homology basis Γ. Similarly to Lemma 1.2, we can easily show that

Lemma 5.6. *For any basis $\{\ell_i\}_{1 \le i \le |C|}$ of $\pi_1(M, v)$, there exists a permutation τ of $\{1 \ldots |C|\}$ such that for each $i \in \{1 \ldots |C|\}$, the loop γ_i is a greedy factor of $\ell_{\tau(i)}$.*

It directly follows from the two preceding lemmas that

Proposition 5.7. *Any greedy homotopy basis is a minimum weight basis.*

In order to compute a greedy homotopy basis one needs to compute a shortest path tree and a maximum weight spanning tree. A shortest path tree of a graph with n vertices and m edges can be computed in $O(n \log n + m)$ time using Dijkstra's algorithm. Classic maximum (or minimum) weight spanning tree algorithms run[1] in $O(n \log n + m)$ time [Tar83]. Since a homotopy basis of a map of genus g has $O(g)$ loops, and since each loop of a greedy basis may have size $O(n)$ we obtain

Theorem 5.8 ([EW05])**.** *Let M be a finite connected map of genus g without boundary with n vertices and m non-negatively weighted edges. Given a vertex v of M, a minimum weight basis of $\pi_1(M, v)$ can be computed in $O(n \log n + gn + m)$ time.*

6 Minimum Basis of the First Homology Group of a Surface

6.1 Homology Basis Associated with a Tree-Cotree Decomposition

In analogy with the fundamental cycle basis of a graph associated with a spanning tree, we can associate a basis of $H_1(M)$ with a tree-cotree decomposition.

Lemma 6.1. *Let (T, D^*, C) be a tree-cotree decomposition of M. The set of cycles $\{\gamma^T_c\}_{c \in C}$ is a basis of $H_1(M)$.*

PROOF. We can either reproduce the proof of Lemma 5.3, replacing contractible by 0-homologous, or directly apply Proposition 4.3. □

[1] A faster $O(m)$ algorithm exists for embedded graphs. See for instance Sec. 3.1. of Éric Colin de Verdière course notes.

6.2 The Greedy Homology Basis

We again assume that the edges of the graph G of M are positively weighted. We also assume uniqueness of shortest path between each pair of vertices in G, see Section 3. Analogously to Section 2, we look for a basis of $H_1(M)$ such that the sum of the weights of the cycles in the basis is minimal. Since $H_1(M)$ is a vector space, the greedy matroidal algorithm of Section 2.1 remains valid as well as the characterization in Corollary 2.3 and 2.5 of the cycles in a minimum basis. For each vertex v of M we let T_v be a shortest path tree rooted at v. By Corollary 2.5, we can restrict the cycle scan in the greedy algorithm to simple cycles of the form $\gamma_{v,e} := \gamma_{v,e}^{T_v}$, one for each chord e of T_v. In fact, we can further restrict the scan to a subset of $O(g)$ candidate cycles per vertex.

Lemma 6.2. *The set of loops $\mathscr{L}_v = \{\gamma_{v,e} \mid e \in E(G)\backslash E(T_v)\}$ contains at most $3(1-\chi(M)) = O(g)$ distinct homology classes. Furthermore, we can select in $O(n\log n + m)$ time a subset $\mathscr{S}_v \subset \mathscr{L}_v$ of at most $3(1-\chi(M))$ loops that contains a homologous loop of minimal weight for each homology class in $\mathscr{L}_v$.*

PROOF. Following Section 5.1 we use a * superscript to denote duality. Put $K^* := G^* - E(T_v^*)$. Since T_v is a tree, it can be completed to a tree-cotree decomposition of M and it results from Lemma 5.2 that the cycle space $Z(K^*)$ has dimension $2-\chi(M)$. If $e_1^*, \dots, e_k^*$ are the edges incident to a vertex dual to a face f of M, then $\partial_2 f = \sum_i e_i = \sum_i \gamma_{v,e_i}$, so that $\sum_i \gamma_{v,e_i}$ is null-homologous. This sum can be restricted to $e_i \in E(K)$ because γ_{v,e_i} is null-homologous otherwise. It follows that $\gamma_{v,e}$ is also null-homologous whenever e^* is a pendant edge in K^*. We can delete recursively all the pendant edges in K^* since their corresponding cycle is null-homologous. We are left with a subgraph K_1^* without degree one vertex and with the same cycle space as K^*. If two edges e^* and e'^* share a degree two vertex in K_1^* we also have that $\gamma_{v,e}$ and $\gamma_{v,e'}$ are homologous. It follows that the number of distinct homology classes is at most the number of maximal chains, *i.e.* of maximal paths with degree two internal vertices in K_1^*. This is also the number of edges of the graphs K_2^* obtained by contracting each such maximal chain to a single edge. Because each vertex of K_2^* has degree three or more, we have $2|E(K_2^*)| \geq 3|V(K_2^*)|$ by double counting of the vertex-edge incidences. On the other hand,

$$2-\chi(M) = \dim Z(K^*) = \dim Z(K_1^*) = \dim Z(K_2^*) = 1 - |V(K_2^*)| + |E(K_2^*)|$$

It ensues that $|E(K_2^*)| \leq 3(|E(K_2^*)| - |V(K_2^*)|) = 3(1-\chi(M))$ as desired. In practice, we first compute T_v and the distance of each vertex to the root v in $O(n\log n + m)$ time using Dijkstra's algorithm. For any edge e of M, the length of $\gamma_{v,e}$ can then be computed in constant time. We recursively remove the pendant edges of K^* and traverse each maximal chain of the resulting graph K_1^* in linear time, only keeping in $\mathscr{S}_v$ the loop $\gamma_{v,e}$ corresponding to the traversed edge e^* if the loop has minimum weight in the maximal chain. □

The greedy matroidal algorithm requires to test if a loop is homologically independent of the already selected loops. To this end we consider a fixed homology basis $\mathscr{B} := \{\gamma_c^T\}_{c\in C}$ associated with some tree-cotree decomposition (T, D^*, C).

Lemma 6.3. *We can compute the homology coordinates with respect to $\mathscr{B}$ of each of the loops in $\mathscr{S}_v$ in $O(gm)$ total time.*

Proof. We first compute for each edge e of M, the coordinates of γ_e^T with respect to $\mathscr{B}$. This can be done in $O(gm)$ time for all the edges in D by a simple traversal of the dual tree D^*. We then traverse the shortest path tree T_v from its root v in order to compute for each vertex x the homology coordinates with respect to $\mathscr{B}$ of the loop $\gamma_v(x) := \gamma_{v,x}^{T_v} \cdot \gamma_{x,v}^T$ composed of the two $(x-v)$-paths in T_v and T respectively. The traversal needs $O(gn)$ time, spending $O(g)$ time per vertex to compute the coordinates of $[\gamma_v(x)] = [\gamma_v(y)] + [\gamma_{yx}^T]$ using the predecessor y of x in T_v. The coordinates of any $\gamma_{v,e}$ in $\mathscr{S}_v$ can now be decomposed into the sum of the coordinates of $\gamma_v(x)$, γ_e^T and $\gamma_v(y)$ where x, y are the endpoints of e. It thus takes $O(g)$ time to compute the coordinates of any loop in $\mathscr{S}_v$ and the whole computation needs $O(gm)$ time. □

Theorem 6.4 ([EW05])**.** *Let M be a finite connected map of genus g with n vertices and m weighted edges. A minimum weight basis of $H_1(M)$ can be computed in $O(n^2 \log n + gnm + g^3 n)$ time.*

Proof. We can select $O(gn)$ loop candidates for the minimal weight basis and compute their weights in $O(n^2 \log n + nm)$ time according to Lemma 6.2. Their homology coordinates with respect to $\mathscr{B}$ is computed in $O(gnm)$ time following Lemma 6.3. After sorting the $O(gn)$ candidate loops according to their weight in $O(gn \log n)$ time, the greedy algorithm consists in scanning the candidate loops in increasing order, keeping the scanned loop in the minimal basis if it is homologically independent of the previously selected loops. This last test can be answered in $O(g^2)$ time using Gauss elimination to maintain the $O(g)$ selected loops in row echelon form. The whole scan thus takes $O(g^3 n)$ time. Summing up all the steps we may conclude the theorem. □

When $g = o(n^{1/3})$, a faster $O(g^3 n \log n)$ algorithm was obtained by Borradaile et al. [BCFN16]. It combines the approach of Kavitha et al. [KMMP08] for the minimum cycle basis with the use of a certain cyclic covering to compute each cycle of the minimum weight basis.

References

[BCFN16] Glencora Borradaile, Erin Wolf Chambers, Kyle Fox, and Amir Nayyeri. Minimum cycle and homology bases of surface embedded graphs. In *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik, 2016.

[CCE13] Sergio Cabello, Erin W. Chambers, and Jeff Erickson. Multiple-source shortest paths in embedded graphs. *SIAM Journal on Computing*, 42(4):1542–1571, 2013.

[CLRS09] Thomas Cormen, Charles Leiserson, Ronald Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, third edition, 2009.

[DPeK82] Narsingh Deo, Gurpur M. Prabhu, and Mukkai S. et Krishnamoorty. Algorithms for generating fundamental cycles in a graph. *ACM Trans. Math. Software*, 8:26–42, 1982.

[Eri12] 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.

[EW05] Jeff Erickson and Kim Whittelsey. Greedy optimal homotopy and homology generators. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1038–1046, 2005.

[HM93] David Hartvigsen and Russel Mardon. When Do Short Cycles Generate the Cycle Space. *Journal of Combinatorial Theory, Series B*, 57:88–99, 1993.

[Hor87] Joseph D. Horton. A Polynomial-Time Algorithm to Find the Shortest Cycle Basis of a Graph. *SIAM Journal of Computing*, 16(2):358–366, 1987.

[KLM+09] Telikepalli Kavitha, Christian Liebchen, Kurt Mehlhorn, Dimitrios Michail, Romeo Rizzi, Torsten Ueckerdt, and Katharina A Zweig. Cycle bases in graphs characterization, algorithms, complexity, and applications. *Computer Science Review*, 3(4):199–243, 2009.

[KMMP04] Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna Paluch. A Faster Algorithm for Minimum Cycle Basis of Graphs. In *ICALP*, 2004.

[KMMP08] Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E. Paluch. An Õ(m2 n) algorithm for minimum cycle basis of graphs. *Algorithmica*, 52:333–349, 2008.

[Laz14] Francis Lazarus. *Combinatorial Graphs and Surfaces from the Computational and Topological Viewpoint Followed by some notes on The Isometric Embedding of the square Flat Torus.* Mémoire d'habilitation à diriger des recherches, Université Joseph Fourier, September 2014.

[MM09] Kurt Mehlhorn and Dimitrios Michail. Minimum Cycle Bases: Faster and Simpler. *ACM Transactions on Algorithms*, 6(1):art. 8, 2009.

[MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. In *Proceedings of the nineteenth annual ACM Symposium on Theory of Computing (STOC)*, pages 345–354. ACM, 1987.

[Tar83] Robert E. Tarjan. *Data Structures and Network Algorithms.* Number 44 in CBMS-NFS Regional conference series in applied mathematics. SIAM, 1983.