

Counting geodesics between surface triangulations

BY HUGO PARLIER

*Department of Mathematics, FSTM, University of Luxembourg, Esch-sur-Alzette,
L-4364, Luxembourg.*

e-mail: hugo.parlier@uni.lu

AND LIONEL POURNIN

LIPN, Université Paris 13, Villetaneuse, 93430, France.

e-mail: lionel.pournin@univ-paris13.fr

(Received 14 August 2023; revised 04 December 2024)

Abstract

Given a surface Σ equipped with a set P of marked points, we consider the triangulations of Σ with vertex set P . The flip-graph of Σ is the graph whose vertices are these triangulations, and whose edges correspond to flipping arcs in these triangulations. The flip-graph of a surface appears in the study of moduli spaces and mapping class groups. We consider the number of geodesics in the flip-graph of Σ between two triangulations as a function of their distance. We show that this number grows exponentially provided the surface has enough topology, and that in the remaining cases the growth is polynomial.

2020 Mathematics Subject Classification: 57K20 (Primary)

1. Introduction

The number of geodesics between two points in a metric space appears in many contexts, and is related to understanding the dynamical, geometric and topological properties of this space. In particular, for compact manifolds, the problem of counting closed geodesics has many facets and is related to the volume entropy. These notions enjoy counterparts in combinatorial settings and in particular for graphs [2]. Certain manifolds, such as the standard round sphere or Euclidean spaces, enjoy “blocking” properties, meaning that removing a small number of points (or cutting away a small volume submanifold) can block all the geodesics between two points (see for instance [5, 13]). The unicity, or local unicity, of geodesics between points is a feature of non-positive curvature, but it can fail even in coarse settings. The curve graph is a relevant example: while it is Gromov hyperbolic [7], there are infinitely-many geodesic paths between pairs of vertices at distance 2. In the related pants graph however [1], it is still unclear whether a similar property holds.

Flip-graphs are naturally associated to finite-type surfaces, just like curve graphs or pants graphs, and we explore the problem of counting the geodesics between their vertices. More precisely, we consider an orientable surface Σ of genus g , with b boundary curves, and p punctures. We equip Σ with finitely-many marked points: each puncture of Σ is a marked




Fig. 1. The behaviour of $\Delta_k(\Sigma)$ when $b \leq 3, g \leq 2$ and $p \leq 5$. Symbols indicate the behaviour of $\Delta_k(\Sigma)$ as k goes to infinity: solid disks indicate exponential behaviour, striped disks polynomial behaviour, solid circles that $\Delta_k(\Sigma) = 1$, and dotted circles represent unsolved cases. Triangles mean that $\mathcal{F}(\Sigma)$ is finite. When $b = 0$ and there is no symbol, $\mathcal{F}(\Sigma)$ is empty.

point and all the other marked points of Σ are placed in its boundary curves in such a way that each boundary curve of Σ contains at least one marked point. A triangulation of Σ can be thought of as a set of pairwise non-crossing arcs between the marked points of Σ that decompose Σ into triangles (see Section 2 for a precise definition). The flip-graph of Σ is the graph $\mathcal{F}(\Sigma)$ that admits the triangulations of Σ as its vertices and such that two of these triangulations are connected by an edge of $\mathcal{F}(\Sigma)$ when they differ by exactly one arc. We think of $\mathcal{F}(\Sigma)$ as a metric graph whose edges have length 1.

The vertex degrees of $\mathcal{F}(\Sigma)$ are bounded and with the exception of certain topological types, $\mathcal{F}(\Sigma)$ is infinite. These flip-graphs are related to the study of moduli spaces and mapping class groups [4, 8]. Polygons form a prominent example of surfaces the flip-graphs of which are finite. In this case, the marked points are the vertices of the polygon. The resulting flip-graphs are the 1-skeletons of associahedra [6, 15–17] and their geometric properties are well-studied [11, 12, 14]. The punctured disks with marked points on the boundary are another example of surfaces whose flip-graphs are finite [10]. The only other surface with a finite flip-graph is the 3-punctured sphere (that flip-graph is depicted in Figure 2).

Here, we are primarily interested in infinite flip-graphs, whose lengths of geodesics can get arbitrarily large. The quantity we study is the largest number $\Delta_k(\Sigma)$ of geodesics in $\mathcal{F}(\Sigma)$ between any two vertices at distance k . The vertex degrees of $\mathcal{F}(\Sigma)$ being uniformly bounded implies that $\Delta_k(\Sigma)$ is always finite. More precisely, there is an immediate upper bound on $\Delta_k(\Sigma)$ which grows exponentially as a function of k (with the maximal vertex degree of $\mathcal{F}(\Sigma)$ minus 1 as a basis for the exponential). Our first theorem identifies the surfaces for which there is a matching exponential lower bound.

THEOREM 1. *If b and p are not both equal to 0 and $2g + p + b$ is at least 3, then $\Delta_k(\Sigma)$ grows at least exponentially as a function of k as k tends to infinity, except possibly when:*

- (i) Σ is the 4-punctured sphere (g and b are equal to 0 and p is equal to 4);
- (ii) Σ is a genus one surface with a unique boundary component and no puncture (g and b are equal to 1 and p is equal to 0); or
- (iii) Σ is a 2-punctured disk (g is equal to 0, b to 1, and p to 2).

We also investigate the cases that are not covered by Theorem 1. Figure 1 provides a visual diagram of when the exponential behaviour kicks in for the low complexity cases. As will be discussed in the sequel, some cases are very easy (namely when the graphs are finite)

and in a select few number of cases, we were unable to determine the growth rate with our methods. When Σ is a cylinder without punctures, we show the following.

THEOREM 2. *If Σ is a cylinder without punctures, then the growth of $\Delta_k(\Sigma)$ as a function of k is at most polynomial as k goes to infinity.*

Note that, when Σ is a cylinder without punctures, $b = 2$ and $g = p = 0$ and thus $2g + p + b$ is at most 2. This explains why this case is not listed as an exception in the statement of Theorem 1. We also study the case of a doubly punctured disk.

THEOREM 3. *If Σ is a 2-punctured disk, then the growth of $\Delta_k(\Sigma)$ as a function of k is at most polynomial as k goes to infinity.*

In fact Theorems 2 and 3 will be obtained as a consequence of a more general result. We will show that, as k goes to infinity, $\Delta_k(\Sigma)$ depends only polynomially on two quantities related to the surface Σ^* obtained from Σ by removing all but one marked points on each boundary curve. The first of these quantities, $\tilde{\Delta}_k(\Sigma^*)$ is the maximum value of $\Delta_i(\Sigma^*)$ when i ranges between 1 and k . The second quantity, $\Lambda_k(\Sigma^*)$ is the maximum number of triangulations in any ball of radius k in $\mathcal{F}(\Sigma^*)$.

THEOREM 4. *$\Delta_k(\Sigma)$ is upper bounded by a polynomial function of $\tilde{\Delta}_k(\Sigma^*)$ and $\Lambda_k(\Sigma^*)$.*

The bounds on $\Delta_k(\Sigma)$ stated by Theorems 1, 2 and 4 will be given explicitly (see Theorems 18, 20 and 27). Finally, we point out that there are only two remaining cases, for which we do not know whether $\Delta_k(\Sigma)$ grows exponentially or polynomially as a function of k . These two unsolved cases are the 4-punctured sphere and the genus one surface with a single boundary component and no puncture.

2. Surface triangulations, flip-graphs, and contractions

Throughout the paper, we consider a finite-type orientable surface Σ , possibly with boundary curves and with a non-empty set P of marked points. In order to ensure that P can be the vertex set of a triangulation of Σ , each boundary curve must contain at least one marked point. Marked points that lie in the interior of the surface we call *punctures*.

An arc is a subset of Σ that is obtained as the continuous image of a closed interval, which is an embedding on its interior, and with the endpoints of the interval sent to P . We are interested in arcs up to the homotopies that preserve P and we only consider arcs which are not homotopic to a point (which can only occur if the two endpoints are the same). Two arcs are said to be disjoint if they are disjoint in their interiors and a triangulation of Σ is a set of pairwise disjoint arcs that is maximal with respect to inclusion.

We call the number of arcs in a triangulation (that does not depend on the choice of the triangulation) the *arc complexity* of Σ and denote it by $\kappa(\Sigma)$.

We also consider the unoriented graph $\mathcal{F}(\Sigma)$ whose vertices are the triangulations of Σ and whose edges link two triangulations that differ by a single arc. In other words, two triangulations T and T' of Σ are the endpoints of an edge of $\mathcal{F}(\Sigma)$ when T' can be obtained from T by replacing an arc with a different one. This operation is called a *flip* and the graph $\mathcal{F}(\Sigma)$ the *flip-graph* of Σ . We will think of $\mathcal{F}(\Sigma)$ as a metric space whose elements are the triangulations of Σ where the distance $d(T, T')$ between two triangulations T and T' in $\mathcal{F}(\Sigma)$ is the minimal number of flips needed to transform T into T' .




Fig. 2. The flip-graph of a 3-punctured sphere.

A minimal length path between two triangulations is called geodesic. For any two triangulations T and T' of Σ , we denote by $\#(T, T')$ the number of geodesic paths between T and T' in $\mathcal{F}(\Sigma)$. For any k in \mathbb{N} , we consider the quantity

$$\Delta_k(\Sigma) = \max \left\{ \#(T, T') : (T, T') \in \mathcal{F}(\Sigma)^2, d(T, T') = k \right\}.$$

Note that $\Delta_k(\Sigma)$ is always well-defined and finite. This is because the number of possible flips in a triangulation of Σ is bounded by $\kappa(\Sigma)$ and hence the number of possible sequences of k flips that start at a given triangulation is always at most $\kappa(\Sigma)^k$. Since a flip in a geodesic path cannot be immediately followed by the inverse flip, we get the following.

PROPOSITION 5. $\Delta_k(\Sigma) \leq \kappa(\Sigma)(\kappa(\Sigma) - 1)^{k-1}$.

Throughout the paper, p will denote the number of punctures of Σ , g its genus, and b the number of its boundary components. It will be important to remember that p only counts punctures (marked points in the interior of Σ) but ignores the number of marked points of P that lie on the boundary curves of Σ . If p and b are simultaneously equal to 0, then $\mathcal{F}(\Sigma)$ is empty and for this reason, we will assume that their sum is positive. If b and g are both equal to 0, then Σ is a sphere. If in addition, p is equal to 1 or 2, there are not enough marked points to build a triangulation of Σ (at best, one can decompose Σ into bigons when p is equal to 2). Therefore, we assume that if b and g are both equal to 0, then p is at least 3.

Let us also recall the three cases when $\mathcal{F}(\Sigma)$ is a finite graph. These three cases are marked with a triangle in Figure 1. The first case is the 3-times punctured sphere (when b and g are equal to 0 and p to 3). In that case $\mathcal{F}(\Sigma)$ is finite and the full graph is depicted in Figure 2. The second case is the disk (when g and p are equal to 0 and b is equal to 1). In that case, $\mathcal{F}(\Sigma)$ is empty when there are less than 3 marked points on the boundary of the disk and non-empty otherwise. This particular flip-graph is a graph made up of the vertices and edges of a polytope called the associahedron [6, 15–17], whose number of vertices is counted by the Catalan numbers. The third case where $\mathcal{F}(\Sigma)$ is finite is the once-punctured disk (that is when b and p are both equal to 1 and g to 0) [10]. In this case $\mathcal{F}(\Sigma)$ is never empty and its size depends on the number of marked points on the boundary curve.

We distinguish two kinds of arcs in a triangulation T of Σ : if an arc in T can be homotoped to the boundary of Σ then we call that arc a *boundary* arc of T . Any arc in T that is not a boundary arc will be called an *interior* arc. Observe that all the triangulations of Σ have the




Fig. 3. The contraction of α in a triangulation of Σ .

same set of boundary arcs. For this reason, we will also refer to these arcs as the boundary arcs of Σ . When a boundary component γ of Σ contains a single point p from P , then p splits γ into a single boundary arc α of Σ . In that case, α is twice incident to p and we call this boundary arc a *boundary loop*. In the remainder of the paper, we denote the number of boundary arcs of Σ by n . As each boundary component of Σ contains at least one point from P , it gives rise to at least one boundary arc and therefore n is not less than b . Note that when n is equal to b , all the boundary arcs of Σ are boundary loops.

Let us now describe an operation on Σ and its triangulations that will be useful throughout the paper. Consider a boundary arc α of Σ and assume that α is not a boundary loop. In that case, we can build a new surface $\Sigma \setminus \alpha$ by removing one of the vertices of α from the marked points of Σ . In other words, $\Sigma \setminus \alpha$ is identical to Σ except that its set of marked points is $P \setminus \{a\}$ instead of P , where a is a vertex of α . Note that the two surfaces obtained from this construction using, for a , one or the other vertex of α are homeomorphic. For this reason, we will identify $\Sigma \setminus \alpha$ with any of these surfaces.

Now consider a triangulation T of Σ and denote by t the triangle of T incident to α . Further denote by a and a' the vertices of α as shown on the left of Figure 3. We can obtain a triangulation $T \setminus \alpha$ of $\Sigma \setminus \alpha$ from T as follows. First displace a and a' within α until they merge into a point that we will still denote by a' . During that motion, the arcs in T incident to a and a' should be kept incident to the moving vertices as shown in Figure 3. This makes α disappear and transforms t into the bigon that is striped in the center of Figure 3. In particular, the set of arcs resulting from that operation is not a triangulation. However, this set of arcs can be turned into the announced triangulation $T \setminus \alpha$ of $\Sigma \setminus \alpha$ by removing one of the arcs bounding the bigon. The only thing that remains of the triangle t is the other arc bounding the bigon, that is labelled by β on the right of Figure 3. This operation is referred to as the *contraction* of α within either Σ or T . It will be important to keep in mind that this contraction operation can only be performed when α is not a boundary loop.

The contraction operation was first used in [11] in order to compute the diameter of the flip-graph of a disk without punctures and later in [8–10] to do the same for the modular flip-graphs of more general surfaces up to homeomorphism. It is also used more generally in [12] to compute distances within the flip-graph of a disk without punctures. Let us recall a property of the contraction operation that is instrumental in all of these articles and will also turn up later here. This property is related to the following definition.

Definition 6. Consider two triangulations T and T' of Σ that are related by a flip. We say that this flip is incident to a boundary arc α of Σ when the triangle incident to α in T is not the same as the triangle incident to α in T' .

In other words, a flip is incident to α when it exchanges the diagonals of a quadrilateral which has α as one of its sides. Now recall that all the triangulations of Σ contain α and assume that α is not a boundary loop. In that case, for any two triangulations T and T' that

are related by a flip, $T \setminus \alpha$ and $T' \setminus \alpha$ are two triangulations of $\Sigma \setminus \alpha$ that either coincide or are still related by a flip. More precisely $T \setminus \alpha$ and $T' \setminus \alpha$ coincide precisely when the flip that relates T and T' is incident to α . We therefore obtain the following.

LEMMA 7. *Consider a non-loop boundary arc α of Σ . If v flips are incident to α along some geodesic path in $\mathcal{F}(\Sigma)$ between two triangulations T and T' , then*

$$d(T, T') \geq d(T \setminus \alpha, T' \setminus \alpha) + v.$$

3. The exponential regime

Denote by Γ the surface with two boundary components, genus zero, no puncture, and a single marked point on each boundary component. In this section, we prove Theorem 1: our aim is to show that, except for a few cases, $\Delta_k(\Sigma)$ is at least an exponential function of k when k goes to infinity. In order to do that, we are going to exhibit strongly convex subgraphs of $\mathcal{F}(\Sigma)$ that are isomorphic to the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$. Indeed, we shall see that the largest possible number of geodesic paths between two vertices of that cartesian product is an exponential function of their distance. We recall that a subgraph G of $\mathcal{F}(\Sigma)$ is strongly convex when the geodesic paths in $\mathcal{F}(\Sigma)$ between two vertices of G are all entirely contained in G . Strong convexity in flip-graphs has been studied in different disguises for instance in [3, 4, 12, 14]. Further recall that Σ^* denotes the surface obtained from Σ by removing, on each boundary component, all the marked points except one.

LEMMA 8. *$\mathcal{F}(\Sigma)$ has a strongly convex subgraph isomorphic to $\mathcal{F}(\Sigma^*)$.*

Proof. Denote by α_1 to α_b the boundary components of Σ . As there is at least a marked point in any of them, we can pick one of these marked points, say a_i in α_i . Cutting α_i at a_i results in a boundary loop α_i^* twice incident to a_i that contains in its interior all the marked points in α_i but a_i . Now consider a sequence β_1 to β_b arcs in Σ such that β_i is twice incident to a_i and the portion of Σ bounded by $\alpha_i \cup \beta_i$ is a topological disk Σ_i . By construction, Σ^* is homeomorphic to the surface obtained by cutting the disks Σ_1 to Σ_b out from Σ . We will therefore identify Σ^* with this surface.

Now consider a sequence T_1 to T_b of triangulations of Σ_1 to Σ_b , respectively and observe that, for any triangulation T^* of Σ^* ,

$$\left(\bigcup_{i=1}^b T_i \right) \cup T^*$$

is a triangulation of Σ . It immediately follows that $\mathcal{F}(\Sigma)$ contains a copy of $\mathcal{F}(\Sigma^*)$ as an induced subgraph. According to [4], that subgraph is strongly convex in $\mathcal{F}(\Sigma)$.

Using Lemma 8, we can state sufficient conditions on Σ under which $\mathcal{F}(\Sigma)$ has a strongly convex subgraph isomorphic to $\mathcal{F}(\Gamma)$.

PROPOSITION 9. *If Σ has at least one boundary curve and at least two punctures, then $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ has at least one boundary curve and at least two punctures. Denote by β one of the boundary curves of Σ and by x one of its punctures. Pick an arc α in Σ twice incident to x that separates the surface into a cylinder Γ' without punctures bounded by α




Fig. 4. The constructions in Propositions 9, 10 and 11.

and β and a surface Σ' that contains all the remaining topology of Σ as shown on the left of Figure 4. Note that α is not contractible because Σ' has at least one puncture. Therefore, Γ' is indeed an essential cylinder. Pick a triangulation T' of Σ' and consider the subgraph \mathcal{G} of $\mathcal{F}(\Sigma)$ induced by the triangulations that contain all the arcs in T' .

By construction, \mathcal{G} is isomorphic to $\mathcal{F}(\Gamma')$ and according to [4], \mathcal{G} is a strongly convex subgraph of $\mathcal{F}(\Sigma)$. Therefore, $\mathcal{F}(\Sigma)$ has a strongly convex subgraph isomorphic to $\mathcal{F}(\Gamma')$. This does not immediately provide the desired result because $\mathcal{F}(\Gamma')$ is not necessarily isomorphic to $\mathcal{F}(\Gamma)$. Indeed, β can contain several marked points and, in that case, Γ' is not homeomorphic to Γ . However, $[\Gamma']^*$ is homeomorphic to Γ and, by Lemma 8, $\mathcal{F}(\Gamma')$ has a strongly convex subgraph isomorphic to $\mathcal{F}(\Gamma)$, which proves the proposition.

PROPOSITION 10. *If Σ has at least one boundary curve and genus at least one, then $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ has at least one boundary curve β and denote by x one of the marked points in β . If in addition, Σ has genus at least one, we can pick an arc α twice incident to x in the interior of Σ that separates Σ into two surfaces one of which is a genus one surface without punctures and a single boundary component and the other contains the remaining topology of Σ as shown in the center of Figure 4. Denote by Σ' the latter surface and note that if α is homotopic to β , then Σ' is just a bigon. Now pick an interior arc α' twice incident to x in $\Sigma \setminus (\Sigma' \setminus \alpha)$ and denote by Γ' the cylinder obtained by cutting $\Sigma \setminus (\Sigma' \setminus \alpha)$ along α' . Observe that Γ' has no puncture, a single marked point on one of its boundary components, and two on the other, all three of which are copies of x .

If α and β are homotopic, then denote by \mathcal{G} the subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulations that contain α' . Otherwise, consider a triangulation T' of Σ' and denote by \mathcal{G} the subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulations that contain α' and all the arcs in T' . In both cases, \mathcal{G} is isomorphic to $\mathcal{F}(\Gamma')$. Moreover, according to [4], \mathcal{G} is a strongly convex subgraph of $\mathcal{F}(\Sigma)$. Finally, observe that $[\Gamma']^*$ is homeomorphic to Γ . Hence, by Lemma 8, \mathcal{G} admits a copy of $\mathcal{F}(\Gamma)$ as a strongly convex subgraph.

PROPOSITION 11. *If Σ has at least two boundary curves, then its flip-graph $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ has two boundary curves, denote by β one of these curves, and by x one of the marked points in β . Pick an arc α twice incident to x in the interior of Σ that cuts the surface into a cylinder Γ' without punctures and a surface Σ' that contains the remaining topology of Σ , as shown on the right of Figure 4. Observe that α and β might be homotopic. However, if this happens, Σ is a cylinder without punctures and Σ^* is homeomorphic to Γ . Therefore, the result immediately follows from Lemma 8 in that case.

Now if α and β are not homotopic, consider a triangulation T' of Σ' . The subgraph \mathcal{G} induced in $\mathcal{F}(\Sigma)$ by the triangulations that contain all the arcs in T' is isomorphic to $\mathcal{F}(\Gamma')$ and by [4, theorem 1.1] it is a strongly convex subgraph of $\mathcal{F}(\Sigma)$. Since Σ^* is homeomorphic to Γ , the proposition therefore follows from Lemma 8.

Using the above three propositions, we can exhibit strongly convex subgraphs of $\mathcal{F}(\Sigma)$ isomorphic to $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ except in a select number of cases.

LEMMA 12. *If Σ has at least five punctures, then $\mathcal{F}(\Sigma)$ admits a copy of the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ has at least five punctures. Consider an arc α in the interior of Σ twice incident to one of these punctures and that separates Σ into two sub-surfaces Σ' and Σ'' one of which is a 2-punctured disk. The subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulation that contain α is isomorphic to $\mathcal{F}(\Sigma') \times \mathcal{F}(\Sigma'')$ and according to [4], that subgraph is strongly convex in $\mathcal{F}(\Sigma)$. As both Σ' and Σ'' have at least one boundary and at least two punctures, the lemma therefore follows from Proposition 9.

LEMMA 13. *If Σ has genus at least two and at least one puncture or at least one boundary component, then $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ at least one puncture or one boundary component. In that case, one can pick a marked point x in Σ (which is either a puncture or a marked point in a boundary component). If in addition, Σ has genus at least two, then we can choose an arc α in the interior of Σ that is twice incident to x and that separates Σ into a genus one surface Σ' and a positive genus surface Σ'' . By construction, the subgraph induced in $\mathcal{F}(\Sigma)$ by the triangulations that admit α as an arc is isomorphic to $\mathcal{F}(\Sigma') \times \mathcal{F}(\Sigma'')$. Moreover, by [4] that subgraph is strongly convex in $\mathcal{F}(\Sigma)$. As Σ' and Σ'' each have at least one boundary component and positive genus, the result follows from Proposition 10.

LEMMA 14. *If Σ has at least three boundary components or exactly two boundary components and at least one puncture, then $\mathcal{F}(\Sigma)$ admits a copy of the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ either has at least three boundary components or just two boundary components but at least one puncture. In the latter case, let x denote a puncture and, in the former, a marked point in one of the boundary components. Now pick an arc α in the interior of Σ twice incident to x . We can require that α separates Σ into two surfaces Σ' and Σ'' , both of which have at least two boundary components. Note that α is a boundary component of these two surfaces and the other is one of the boundary components of Σ that does not contain x . The subgraph induced in $\mathcal{F}(\Sigma)$ by the triangulations that contain α is isomorphic to $\mathcal{F}(\Sigma') \times \mathcal{F}(\Sigma'')$ and that subgraph is strongly convex in $\mathcal{F}(\Sigma)$ according to Theorem 1.1 from [4]. By Proposition 11, $\mathcal{F}(\Sigma')$ and $\mathcal{F}(\Sigma'')$ each admit a copy of $\mathcal{F}(\Gamma)$ as a strongly convex subgraph, which proves the lemma.

LEMMA 15. *If Σ has positive genus and either at least two boundary components or exactly one boundary component but at least one puncture, then $\mathcal{F}(\Sigma)$ admits a copy of the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*




Fig. 5. The flip-graph of a cylinder without punctures and a marked point on each boundary.

Proof. Assume that Σ at least two boundary components or that it has a single boundary component but at least one puncture. In the former case, denote by x a marked point in one of the boundary components of Σ and in the latter case denote by x one of its punctures. Under the additional assumption that Σ has positive genus, there exists a non-separating arc α twice incident to x in Σ . Cutting Σ along α results in a surface Σ' with at least three boundaries. By construction, the subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulations that contain α is isomorphic to $\mathcal{F}(\Sigma')$ and according to [4], that subgraph is strongly convex in $\mathcal{F}(\Sigma)$. By Lemma 14, $\mathcal{F}(\Sigma')$ admits a copy of $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.

LEMMA 16. *If Σ is a disk with at least three punctures then $\mathcal{F}(\Sigma)$ admits a copy of the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ is a punctured disk with at least three punctures. Consider an arc α twice incident to one of these punctures in the interior of Σ . We can assume that α cuts Σ into a (possibly punctured) cylinder Σ' and a 2-punctured disk Σ'' . The subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulations that contain α is then isomorphic to $\mathcal{F}(\Sigma') \times \mathcal{F}(\Sigma'')$ and is strongly convex in $\mathcal{F}(\Sigma)$ by [4]. The result then follows from Propositions 9 and 11.

LEMMA 17. *If Σ does not have a boundary, but has at least two punctures and genus exactly one, then $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ as a strongly convex subgraph.*

Proof. Assume that Σ has genus one and at least two punctures but no boundary. Consider a non-separating arc α in Σ that is twice incident to one of the punctures. Cutting Σ along α results in a cylinder Σ' with at least one puncture. As in the above proofs, the subgraph of $\mathcal{F}(\Sigma)$ induced by the triangulations that admit α as an arc is strongly convex and isomorphic to $\mathcal{F}(\Sigma')$. Hence, the result is immediately obtained from Lemma 14.

We are now ready to prove Theorem 1. The proof relies on the above lemmas and on the fact that $\mathcal{F}(\Gamma)$ is a bi-infinite simple path. More precisely, as already observed in [8], $\mathcal{F}(\Gamma)$ is the infinite connected graph whose vertices each are incident to exactly two edges, as shown in Figure 5. In other words, $\mathcal{F}(\Gamma)$ can be modeled as the graph whose vertices are the integers in \mathbb{Z} and whose edges connect any two consecutive integers. In turn, the cartesian product $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ is isomorphic to the graph \mathcal{G} whose vertices are the points from \mathbb{Z}^2 and whose edges connect two points that coincide in one coordinate and differ by 1 in the other. For any two positive integers i and j , the distance in the graph \mathcal{G} between the origin $(0,0)$ of \mathbb{Z}^2 and the point (i, j) is precisely $i + j$ and the geodesics starting at the former point and ending at the latter are made of edges along which one coordinate increases by 1 while the other coordinate does not change. As a consequence, there are

$$\frac{(i + j)!}{i!j!}$$

geodesic paths in \mathcal{G} between these two points of \mathbb{Z}^2 . As a consequence, if k is any integer greater than or equal to 2, then taking

$$\begin{cases} i = \lfloor \frac{k}{2} \rfloor \text{ and} \\ j = \lceil \frac{k}{2} \rceil \end{cases}$$

and using Stirling’s approximation of the factorial, one obtains that there are two vertices of $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ whose distance is equal to k such that the number of geodesics between these vertices in $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$ is greater than

$$\left(1 - \frac{1}{k}\right) \sqrt{\frac{2}{k\pi}} 2^k e^{-\frac{1}{4(k-1)}}.$$

With this estimate, we can prove the following theorem that gives an explicit exponential lower bound on $\Delta_k(\Sigma)$ and therefore implies Theorem 1.

THEOREM 18. *If b and p are not both equal to 0 and $2g + p + b$ is at least 3, then for every integer k greater than or equal to 2,*

$$\Delta_k(\Sigma) > \left(1 - \frac{1}{k}\right) \sqrt{\frac{2}{k\pi}} 2^k e^{-\frac{1}{4(k-1)}}$$

except possibly when:

- (i) Σ is the 4-punctured sphere (g and b are equal to 0 and p is equal to 4);
- (ii) Σ is a genus one surface with a unique boundary component and no puncture (g and b are equal to 1 and p is equal to 0); or
- (iii) Σ is a 2-punctured disk (g is equal to 0, b to 1, and p to 2).

Proof. This is a direct consequence of Lemmas 12 to 17 and of the discussion above on the number of geodesics between the vertices of $\mathcal{F}(\Gamma) \times \mathcal{F}(\Gamma)$.

Remark 19. When g , p , and b grow, $\mathcal{F}(\Sigma)$ will admit as strongly convex subgraphs cartesian products of the form $\mathcal{F}(\Gamma)^r$ where r is greater than 2. In that case, the lower bound on $\Delta_k(\Sigma)$ provided by Theorem 18 can be improved. However, the asymptotic behaviour of $\Delta_k(\Sigma)$ will remain exponential as we pointed out in the introduction.

4. General estimates

For any triangulation T of Σ and any non-negative integer k , we consider the set $B_\Sigma(T, k)$ of all the triangulations in $\mathcal{F}(\Sigma)$ at distance at most k from T . Equivalently, $B_\Sigma(T, k)$ is the ball of radius k centered at T in $\mathcal{F}(\Sigma)$. Recall that the number of triangulations contained in $B_\Sigma(T, k)$ is bounded above by $\kappa(\Sigma)^k$. We will denote by $\Lambda_k(\Sigma)$ the largest possible number of triangulations in any such ball:

$$\Lambda_k(\Sigma) = \max \left\{ |B_\Sigma(T, k)| : T \in \mathcal{F}(\Sigma) \right\}.$$

We also denote by $\tilde{\Delta}_k(\Sigma)$ the largest value of $\Delta_i(\Sigma)$ when $1 \leq i \leq k$:

$$\tilde{\Delta}_k(\Sigma) = \max \left\{ \Delta_i(\Sigma) : 1 \leq i \leq k \right\}.$$

The aim of this section is to prove the following theorem.

THEOREM 20. *For every positive integer k ,*

$$\Delta_k(\Sigma^*) \leq \Delta_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{(n-b)r} \Lambda_k(\Sigma^*)^{2(r-1)} \tilde{\Delta}_k(\Sigma^*)^r \tag{1}$$

where r is equal to $(2\kappa(\Sigma))^{n-b}$.

The first inequality in (1) is easily obtained. Recall that a subgraph G' of a graph G is strongly convex when all the geodesic paths in G between any two vertices of G' remain in G' . Hence, we immediately get the inequality $\Delta_k(\Sigma^*) \leq \Delta_k(\Sigma)$ from Lemma 8.

Hence, we turn our attention to proving the second inequality in (1). As a first step, we show that when a boundary arc of Σ is not a boundary loop, the number of flips incident to this arc along any geodesic path of length k in $\mathcal{F}(\Sigma)$ can be bounded above independently from k . In order to show this, consider a boundary arc α of Σ that is not a boundary loop and a vertex x of α . Let β be the boundary arc of Σ other than α that admits x as a vertex. Consider the only triangle t contained in Σ that is incident to both α and β and observe that t exists precisely because α is not a boundary loop. Now denote by ε the edge of t that is distinct from α and β . We will denote by $\mathcal{F}_x(\Sigma)$ the subgraph induced in $\mathcal{F}(\Sigma)$ by the triangulations that admit ε as an arc. Equivalently, $\mathcal{F}_x(\Sigma)$ is made up of all the triangulations of Σ that do not contain any interior arc incident to x . Observe that the map $\psi_{x,\alpha}: \mathcal{F}_x(\Sigma) \rightarrow \mathcal{F}(\Sigma \setminus \alpha)$ that sends a triangulation T in $\mathcal{F}_x(\Sigma)$ to $T \setminus \alpha$ is an isomorphism (and therefore, since we think of flip-graphs as metric spaces, an isometry). In the sequel, we denote by $\text{deg}_T(x)$ the number of incidences between x and the arcs of T . It will be important to keep in mind that any arc in T that is twice incident to x contributes 2 to $\text{deg}_T(x)$.

LEMMA 21. *Consider a boundary arc α of Σ and a vertex x of α . If α is not a loop then, for any triangulation T of Σ ,*

$$d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) \leq \text{deg}_T(x) - 2.$$

Proof. Consider a triangulation T of Σ and the triangle t of T incident to α . Denote

$$i(T) = \begin{cases} \text{deg}_T(x) & \text{if } t \text{ does not have an edge twice incident to } x \text{ and} \\ \text{deg}_T(x) - 1 & \text{if } t \text{ has an edge twice incident to } x. \end{cases}$$

Under the assumption that α is not a loop, $i(T)$ is necessarily at least 2. Indeed, in that case, x must be incident to two boundary arcs of Σ (none of which is a boundary loop).

We will prove by induction on $i(T)$ that

$$d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) \leq i(T) - 2. \tag{2}$$

As $i(T)$ is at most $\text{deg}_T(x)$, the desired result will follow.

Observe that, if $i(T)$ is equal to 2, then T belongs to $\mathcal{F}_x(\Sigma)$. As a consequence, T coincides with $\psi_{x,\alpha}^{-1}(T \setminus \alpha)$, which establishes the base case for the induction. Now assume that $i(T)$ is at least 3. Denote by y the vertex of α distinct from x and by γ the edge of t that is not incident to y . Note that γ is a flippable interior arc of T . Moreover, γ is the only edge of t that can possibly be twice incident to x because y is distinct from x . Denote by T' the triangulation of Σ obtained by flipping γ in T . Since T and T' are related by a flip incident




Fig. 6. The flip between T and T' .

to α , the triangulations $T' \setminus \alpha$ and $T \setminus \alpha$ coincide. In particular,

$$d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) \leq d(T', \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) + 1. \tag{3}$$

We shall now prove that

$$i(T') = i(T) - 1. \tag{4}$$

Denote by t' the triangle of T' incident to α and by γ' the edge of t' that is not incident to y . The flip between T and T' is shown in Figure 6 depending on whether γ and γ' are twice incident to x or just once. For instance, on the left of the figure, neither γ nor γ' are twice incident to x . In that case, $i(T)$ coincides with $\deg_T(x)$ and $i(T')$ with $\deg_{T'}(x)$. Hence, (4) follows from the observation that $\deg_{T'}(x)$ is less than $\deg_T(x)$ by exactly one. The situation depicted next on Figure 6 is when γ is not twice incident to x but γ' is. Therefore, $i(T)$ is equal to $\deg_T(x)$ and $i(T')$ to $\deg_{T'}(x) - 1$. However, here $\deg_{T'}(x)$ is equal to $\deg_T(x)$ and we recover (4) as well. In the situation depicted third on the figure, γ is twice incident to x and γ' only once. Hence, $i(T)$ is equal to $\deg_T(x) - 1$ and $i(T')$ to $\deg_{T'}(x)$. As in this case, $\deg_{T'}(x)$ is less than $\deg_T(x)$ by two, (4) holds again. In the situation shown last, γ and γ' are both twice incident to x and $i(T') - i(T)$ is equal to $\deg_{T'}(x) - \deg_T(x)$. As $\deg_{T'}(x)$ is less than $\deg_T(x)$ by one, this proves (4), as announced. Therefore, by induction,

$$d(T', \psi_{a,\alpha}^{-1}(T' \setminus \alpha)) \leq i(T') - 2. \tag{5}$$

Combining (3), (4) and (5) proves (2), as desired.

THEOREM 22. *Consider a boundary arc α of Σ . If α is not a loop, then at most $2\kappa(\Sigma) - 4$ flips are incident to α along any geodesic path in $\mathcal{F}(\Sigma)$.*

Proof. Let x and y be the vertices of α and consider two triangulations T and T' of Σ . As α is not a loop, x and y are distinct and

$$\deg_T(x) + \deg_T(y) + \deg_{T'}(x) + \deg_{T'}(y) \leq 4\kappa(\Sigma).$$

We can therefore assume without loss of generality that

$$\deg_T(x) + \deg_{T'}(x) \leq 2\kappa(\Sigma) \tag{6}$$

by, if needed, exchanging the labels of x and y . Now consider a geodesic path between T and T' in $\mathcal{F}(\Sigma)$ and denote by ν the number of flips along it that are incident to α . We shall

prove that $\nu \leq 2\kappa(\Sigma) - 4$. According to Lemma 7,

$$d(T \setminus \alpha, T' \setminus \alpha) + \nu \leq d(T, T').$$

Since $\psi_{x,\alpha}$ is an isometry, this can be rewritten as

$$d(\psi_{x,\alpha}^{-1}(T \setminus \alpha), \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) + \nu \leq d(T, T'). \tag{7}$$

According to Lemma 21,

$$\begin{cases} d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) \leq \deg_T(x) - 2 \text{ and} \\ d(T', \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) \leq \deg_{T'}(x) - 2. \end{cases}$$

Combining these inequalities with (6) yields

$$d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) + d(T', \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) \leq 2\kappa(\Sigma) - 4.$$

Summing this with (7), one obtains that $2\kappa(\Sigma) - 4 - \nu$ is at least

$$d(T, \psi_{x,\alpha}^{-1}(T \setminus \alpha)) + d(\psi_{x,\alpha}^{-1}(T \setminus \alpha), \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) + d(T', \psi_{x,\alpha}^{-1}(T' \setminus \alpha)) - d(T, T')$$

and it follows from the triangle inequality that $\nu \leq 2\kappa(\Sigma) - 4$.

Recall that n is the number of the boundary arcs of Σ and b is the number of its boundary components. The second ingredient for the proof of the second inequality in (1) is the following lemma according to which $\Lambda_k(\Sigma)$ only differs from $\Lambda_k(\Sigma^*)$ by a factor that does not depend on k but only on $\kappa(\Sigma)$, n , and b .

LEMMA 23. $\Lambda_k(\Sigma^*) \leq \Lambda_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{n-b} \Lambda_k(\Sigma^*)$.

Proof. According to Lemma 8, $\mathcal{F}(\Sigma)$ admits a copy of $\mathcal{F}(\Sigma^*)$ as a strongly convex subgraph. We obtain as an immediate consequence that

$$\Lambda_k(\Sigma^*) \leq \Lambda_k(\Sigma).$$

The other inequality is proven by induction on n . The base case is immediate as Σ coincides with Σ^* when n is equal to b . Assume that n is greater than b . In that case, Σ admits a boundary arc α that is not a loop. By Lemma 7, for any two triangulations T and T' of Σ ,

$$d(T \setminus \alpha, T' \setminus \alpha) \leq d(T, T').$$

Therefore, the map $T' \mapsto T' \setminus \alpha$ sends $B_\Sigma(T, k)$ to a subset of $B_{\Sigma \setminus \alpha}(T \setminus \alpha, k)$. Now consider a triangulation T' in $B_\Sigma(T, k)$. Recall that, when α is contracted within T' , then its two vertices are merged into a single vertex a' and only one of the edges of the triangle t of T' incident to α remains in $T' \setminus \alpha$ as shown in Figure 3 where that edge is labelled β . We can recover all the triangulations T'' in $B_\Sigma(T, k)$ such that $T'' \setminus \alpha$ is equal to $T' \setminus \alpha$ by reversing the process shown in Figure 3. It suffices to pick any arc β of $T' \setminus \alpha$ incident to a' , to insert another arc β' homotopic to β , thus creating a bigon bounded by $\beta \cup \beta'$ and then changing this bigon into a triangle by splitting its boundary at vertex a' and connecting the two obtained copies of a' with the boundary arc α . Note that if β is twice incident to a' , then the bigon can be completed into a triangle in two different ways because its boundary can be split at both of its vertices. The number of triangulations in $B_\Sigma(T, k)$ that can be obtained by this reversed procedure is therefore the number of arcs in $T' \setminus \alpha$ incident to a' where the

arcs twice incident to a' are counted twice. This number is precisely $\text{deg}_{T' \setminus \alpha}(a')$ which is, in turn, not greater than $2\kappa(\Sigma \setminus \alpha)$.

This shows that the map $T' \mapsto T' \setminus \alpha$ sends at most $2\kappa(\Sigma \setminus \alpha)$ triangulations from $B_\Sigma(T, k)$ to the same triangulation of $\Sigma \setminus \alpha$ and as a consequence,

$$|B_\Sigma(T, k)| \leq 2\kappa(\Sigma \setminus \alpha) |B_{\Sigma \setminus \alpha}(T \setminus \alpha, k)|.$$

Assume that T is a triangulation of Σ such that $|B_\Sigma(T, k)|$ is equal to $\Lambda_k(\Sigma)$. In that case, the previous inequality can be rewritten as

$$\Lambda_k(\Sigma) \leq 2\kappa(\Sigma \setminus \alpha) |B_{\Sigma \setminus \alpha}(T \setminus \alpha, k)|.$$

Now observe that $\kappa(\Sigma \setminus \alpha)$ is equal to $\kappa(\Sigma) - 2$. As in addition $|B_{\Sigma \setminus \alpha}(T \setminus \alpha, k)|$ is not greater than $\Lambda_k(\Sigma \setminus \alpha)$, we obtain that

$$\Lambda_k(\Sigma) \leq (2\kappa(\Sigma) - 4) \Lambda_k(\Sigma \setminus \alpha). \tag{8}$$

However $\Sigma \setminus \alpha$ has one boundary arc less than Σ . Therefore, by induction,

$$\Lambda_k(\Sigma \setminus \alpha) \leq (2\kappa(\Sigma \setminus \alpha) - 4)^{n-1-b} \Lambda_k(\Sigma^*). \tag{9}$$

As $\kappa(\Sigma \setminus \alpha) \leq \kappa(\Sigma)$, combining (8) and (9) completes the proof.

Remark 24. Note that, according to the proof of Lemma 23, the upper bound on $\Lambda_k(\Sigma)$ provided by this lemma can immediately be improved to

$$2^{n-b} \Lambda_k(\Sigma^*) \prod_{i=1}^{n-b} (\kappa(\Sigma) - 2i)$$

when $n - b$ is greater than 1 but we will not make use of this in the sequel.

Just as Lemma 23, the second inequality in (1) will be proven by induction on n . The third ingredient in this proof is the induction step, provided by the following lemma.

LEMMA 25. *If a boundary arc α of Σ is not a loop, then for any positive k ,*

$$\Delta_k(\Sigma) \leq \tilde{\Delta}_k(\Sigma \setminus \alpha)^{2\kappa(\Sigma)} \Lambda_k(\Sigma)^{2\kappa(\Sigma)}.$$

Proof. Let α be a non-loop boundary arc of Σ . Consider two triangulations T and T' of Σ whose distance in $\mathcal{F}(\Sigma)$ is equal to k and such that $\#(T, T')$ is equal to $\Delta_k(\Sigma)$. Denote

$$S = \bigcup_{m=0}^{2\kappa(\Sigma)-1} B_\Sigma(T, k)^m.$$

We associate an element of S to each geodesic in $\mathcal{F}(\Sigma)$ between T and T' as follows. Denote by T_i the triangulation of Σ at distance i from T along such a geodesic. Denote by $\phi(1)$ to $\phi(m)$ the indices such that the flip between $T_{\phi(i)}$ and $T_{\phi(i)+1}$ is incident to α . We assume without loss of generality that ϕ is an increasing function of i . Observe that m might be equal to 0 if there is no flip incident to α along the considered geodesic. Moreover, it follows from Lemma 22 that m is less than $2\kappa(\Sigma)$. If $m = 0$, then the element s of S that we associate with the considered geodesic is the empty set (which is the only element of

$B_{\Sigma}(T, k)^0$). Otherwise, we associate the geodesic with

$$s = (T_{\phi(1)}, T_{\phi(2)}, \dots, T_{\phi(m)}). \tag{10}$$

Observe that some elements of S may not be associated to any geodesic between T and T' . For instance, if T and T' do not contain the same triangle incident to α , then the empty set cannot be associated to any geodesic between T and T' as all of them must contain a flip incident to α . Now let us bound the number of geodesics between T and T' that are associated to s . If s is the empty set, then the triangle incident to s is the same in T and in T' . Moreover, all the triangulations along all the geodesics associated to s contain that triangle. Therefore, the number of geodesics associated to s is at most the number of geodesics between $T \setminus \alpha$ and $T' \setminus \alpha$ which, in turn is at most $\tilde{\Delta}_k(\Sigma \setminus \alpha)$ as, by Lemma 7, $d(T \setminus \alpha, T' \setminus \alpha) \leq k$.

Now, if s is given by (10), then observe that along any geodesic associated with s , the triangle incident to α in $T_{\phi(i+1)}$ and $T_{\phi(i+1)}$ must be the same. In particular, the flip that relates $T_{\phi(i)}$ and $T_{\phi(i+1)}$ is prescribed by s when $i < m$ and by T' when $i = m$. Therefore, the number of geodesics in $\mathcal{F}(\Sigma)$ between T and T' that are associated to s is precisely

$$\#(T, T_{\phi(1)}) \left[\prod_{i=1}^{m-1} \#(T_{\phi(i+1)}, T_{\phi(i+1)}) \right] \#(T_{\phi(m)}, T').$$

However, since the triangle incident to α is not modified along the portion between $T_{\phi(i+1)}$ and $T_{\phi(i+1)}$ of any geodesic associated to s ,

$$\#(T_{\phi(i+1)}, T_{\phi(i+1)}) \leq \#(T_{\phi(i+1)} \setminus \alpha, T_{\phi(i+1)} \setminus \alpha).$$

By Lemma 7, $T_{\phi(i+1)} \setminus \alpha$ and $T_{\phi(i+1)} \setminus \alpha$ have distance at most k in $\mathcal{F}(\Sigma \setminus \alpha)$. Hence,

$$\#(T_{\phi(i+1)} \setminus \alpha, T_{\phi(i+1)} \setminus \alpha) \leq \tilde{\Delta}_k(\Sigma \setminus \alpha)$$

and in turn,

$$\#(T_{\phi(i+1)}, T_{\phi(i+1)}) \leq \tilde{\Delta}_k(\Sigma \setminus \alpha)$$

By the same argument, we obtain that $\#(T, T_{\phi(1)})$ and $\#(T_{\phi(m)}, T')$ are also both bounded above by $\tilde{\Delta}_k(\Sigma \setminus \alpha)$. Therefore, the number of geodesics associated to s is at most $\tilde{\Delta}_k(\Sigma \setminus \alpha)^{m+1}$. Hence, any element of $B_{\Sigma}(T, k)^m$ is associated with at most this many geodesics, and we can bound the number of geodesics in $\mathcal{F}(\Sigma)$ between T and T' as

$$\Delta_k(\Sigma) \leq \sum_{m=0}^{2\kappa(\Sigma)-1} |B_{\Sigma}(T, k)|^m \tilde{\Delta}_k(\Sigma \setminus \alpha)^{m+1}$$

The right-hand side of this inequality can be rewritten as

$$\tilde{\Delta}_k(\Sigma \setminus \alpha) \frac{|B_{\Sigma}(T, k)|^{2\kappa(\Sigma)} \tilde{\Delta}_k(\Sigma \setminus \alpha)^{2\kappa(\Sigma)} - 1}{|B_{\Sigma}(T, k)| \tilde{\Delta}_k(\Sigma \setminus \alpha) - 1}$$

Observe that, under the assumption that k is positive, $\tilde{\Delta}_k(\Sigma \setminus \alpha)$ is also positive and $|B_{\Sigma}(T, k)|$ is at least 2. In particular, not only is the denominator of the above fraction non-zero, but it is also least $\tilde{\Delta}_k(\Sigma \setminus \alpha)$. As a consequence,

$$\Delta_k(\Sigma) \leq |B_{\Sigma}(T, k)|^{2\kappa(\Sigma)} \tilde{\Delta}_k(\Sigma \setminus \alpha)^{2\kappa(\Sigma)}$$

Since $|B_\Sigma(T, k)|$ is at most $\Lambda_k(\Sigma)$, this completes the proof.

We are now ready to prove Theorem 20.

Proof of Theorem 20. By Lemma 8, we only need to prove the second inequality from (1). As announced, we will prove it by induction on n . Note that the base case, when n is equal to b , is immediate. Assume that n is greater than b . In that case, Σ has a boundary arc α that is not a loop. According to Lemma 25,

$$\Delta_k(\Sigma) \leq \Delta_i(\Sigma \setminus \alpha)^{2\kappa(\Sigma)} \Lambda_k(\Sigma)^{2\kappa(\Sigma)},$$

where i denotes an integer satisfying $1 \leq i \leq k$. In turn, Lemma 23 makes it possible to bound $\Lambda_k(\Sigma)$ above in terms of $\Lambda_k(\Sigma^*)$ and we obtain

$$\Delta_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{2(n-b)\kappa(\Sigma)} \Delta_i(\Sigma \setminus \alpha)^{2\kappa(\Sigma)} \Lambda_k(\Sigma^*)^{2\kappa(\Sigma)}. \tag{11}$$

By induction,

$$\Delta_i(\Sigma \setminus \alpha) \leq (2\kappa(\Sigma \setminus \alpha) - 4)^{(n-b-1)r'} \Lambda_i(\Sigma^*)^{2(r'-1)} \tilde{\Delta}_i(\Sigma^*)^{r'}, \tag{12}$$

where

$$r' = (2\kappa(\Sigma \setminus \alpha))^{n-b-1}.$$

Recall that $\kappa(\Sigma \setminus \alpha)$ is at most $\kappa(\Sigma)$. As in addition $\Lambda_i(\Sigma^*)$ and $\tilde{\Delta}_i(\Sigma^*)$ both are nondecreasing functions of i , (12) yields

$$\Delta_i(\Sigma \setminus \alpha) \leq (2\kappa(\Sigma) - 4)^{(n-b-1)r'} \Lambda_k(\Sigma^*)^{2(r'-1)} \tilde{\Delta}_k(\Sigma^*)^{r'}.$$

Combining this with (11), one obtains

$$\Delta_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{2(n-b+(n-b-1)r')\kappa(\Sigma)} \Lambda_k(\Sigma^*)^{2\kappa(\Sigma)(2r'-1)} \tilde{\Delta}_k(\Sigma^*)^{2\kappa(\Sigma)r'}.$$

Now denote

$$r = (2\kappa(\Sigma))^{n-b}$$

and recall that $\kappa(\Sigma \setminus \alpha) \leq \kappa(\Sigma)$. Hence, $2\kappa(\Sigma)r' \leq r$. Moreover,

$$2\kappa(\Sigma)(2r' - 1) \leq 2r - 2\kappa(\Sigma).$$

As $\kappa(\Sigma)$ is positive, the right-hand side of this inequality is at most $2(r - 1)$. Therefore,

$$\Delta_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{2(n-b+(n-b-1)r')\kappa(\Sigma)} \Lambda_k(\Sigma^*)^{2(r-1)} \tilde{\Delta}_k(\Sigma^*)^r. \tag{13}$$

As $2\kappa(\Sigma)r' \leq r$, the exponent of $2\kappa(\Sigma) - 4$ in (13) can be bounded as

$$2(n - b + (n - b - 1)r')\kappa(\Sigma) \leq 2(n - b)\kappa(\Sigma) + (n - b - 1)r \tag{14}$$

However, as $\kappa(\Sigma)$ and $n - b$ are both at least 1, we have

$$n - b \leq (2\kappa(\Sigma))^{n-b-1}.$$

Therefore, the first term in the right-hand side of (14) is at most r and in turn, that right-hand side is at most $(n - b)r$. Since $2\kappa(\Sigma) - 4$ is at least 2 (otherwise Σ would not admit a triangulation), the desired upper bound on $\Delta_k(\Sigma)$ follows from (13).

5. Polynomial cases

Recall that Γ is a cylinder without punctures and exactly one marked point on each boundary component. As already mentioned at the end of Section 3, $\mathcal{F}(\Sigma)$ is isomorphic to a bi-infinite simple path. We obtain the following statement as an immediate consequence.

PROPOSITION 26. $\Delta_k(\Gamma) = 1$ and $\Lambda_k(\Gamma) = 2k + 1$.

When Σ is a cylinder without punctures, Σ^* is equal to Γ . Hence, Theorem 20 and Proposition 26 provide Theorem 2. More precisely, we have the following statement.

THEOREM 27. *If Σ is a cylinder without punctures, then*

$$\Delta_k(\Sigma) \leq (2\kappa(\Sigma) - 4)^{(n-b)r} (2k + 1)^{2(r-1)},$$

where r is equal to $(2\kappa(\Sigma))^{(n-b)}$.

This case corresponds to the striped disk in Figure 1. Let us turn our attention to the case of the 2-punctured disk. For the remainder of the section, we assume that Σ is a 2-punctured disk. According to Theorem 20, we only need to consider the 2-punctured disk Σ^* with only one marked point on the boundary. Denote by x and y the two punctures of Σ^* and by z the marked point on its boundary. Note that there is a unique arc with vertices x and y up to homotopy. We denote that arc by ε . Denote by \mathcal{G} the subgraph of $\mathcal{F}(\Sigma^*)$ induced by the triangulations that contain ε and by \mathcal{G}' its subgraph induced by the triangulations that do not contain ε . A portion of $\mathcal{F}(\Sigma^*)$ is shown in Figure 7. The seven triangulations in the two top rows of the figure are in \mathcal{G}' and the eight in the two bottom rows in \mathcal{G} . Observe that \mathcal{G}' contains two kinds of triangulations. The first kind of triangulations have two interior arcs twice incident to z and two non-flippable interior arcs, one incident to x and the other to y . Three such triangulations are shown in the top row of Figure 7. Each of them is adjacent in $\mathcal{F}(\Sigma^*)$ to exactly two triangulations, both in \mathcal{G}' and of the second kind. Four of these second kind triangulations are represented in the second row of Figure 7. In these triangulations, there is a single interior arc twice incident to z . One of the punctures is incident to a non-flippable arc and the other two flippable arcs. These triangulations are adjacent in $\mathcal{F}(\Sigma^*)$ to two triangulations of the first kind and to exactly one triangulation in \mathcal{G} since flipping the interior arc twice incident to z introduces ε . With this description, \mathcal{G}' is, like $\mathcal{F}(\Gamma)$, a bi-infinite simple path because it is a connected infinite subgraph with all vertices of degree 2. Moreover one can move from \mathcal{G}' to \mathcal{G} by a single possible flip in every other triangulation in that path, as shown in Figure 7 and if two triangulations are far enough apart in \mathcal{G}' , it is faster to pass through \mathcal{G} . We get the following as a consequence.

PROPOSITION 28. *Consider two triangulations T and T' that both belong to \mathcal{G}' . If all the triangulations along some geodesic path between T and T' within $\mathcal{F}(\Sigma^*)$ belong to \mathcal{G}' , then the distance of T and T' in $\mathcal{F}(\Sigma^*)$ is at most 6.*




Fig. 7. A portion of the flip-graph of a 2-punctured disk with one marked point in the boundary.

Proof. Consider a geodesic path in $\mathcal{F}(\Sigma^*)$ between T and T' and assume that all the triangulations along that path belong to \mathcal{G}' . Recall that \mathcal{G}' is a simple path in which the triangulations alternate between the two types represented in the first and second rows of Figure 7. Assume for contradiction that the distance of T and T' in $\mathcal{F}(\Sigma^*)$ is at least seven. In particular, there are at least eight triangulations along the considered geodesic. Therefore, one can find a sequence of seven consecutive triangulations T_1 to T_7 in that geodesic such that this sequence starts and ends with triangulations from the second row of Figure 7. In fact, we can assume that T_1 is the first triangulation of the second row in that figure and T_7 the last. As can be seen in the figure, T_1 and T_7 are connected by a path of length five via \mathcal{G} contradicting the assumption that the considered path is a geodesic within $\mathcal{F}(\Sigma^*)$.

We can now prove Theorem 3.

Proof of Theorem 3. According to Theorem 20, it suffices to show that both $\Lambda_k(\Sigma^*)$ and $\Delta_k(\Sigma^*)$ are at most polynomial functions in k .

It is easy to see in Figure 7 that the number of triangulations in any ball of radius k within $\mathcal{F}(\Sigma^*)$ is at most a linear function of k : roughly speaking, this figure represents $\mathcal{F}(\Sigma^*)$ as four infinite rows. In this representation, the edges of $\mathcal{F}(\Sigma^*)$ allow to move within a given column or from one column to the next. In the bottom row, the edges allow to skip one column, but no edge allows to skip more than one column. As a consequence, the number of triangulations in a ball of radius k within $\mathcal{F}(\Sigma^*)$ is at most $8(2k + 1)$ and so is $\Lambda_k(\Sigma^*)$.

According to Proposition 28, the length of a geodesic path that remains outside of \mathcal{G} is at most 6. However, according to Theorem 1.1 from [4], \mathcal{G} is a strongly convex subgraph of $\mathcal{F}(\Sigma^*)$. Therefore, only the first 6 and the last 6 steps along any geodesic in $\mathcal{F}(\Sigma^*)$ can possibly be outside of \mathcal{G} . It then suffices to prove the theorem for the geodesic paths that are entirely contained within \mathcal{G} . Observe that \mathcal{G} is isomorphic to the flip-graph of a cylinder without punctures, with a single marked point on one boundary, and two on the other. (That cylinder is the one obtained by cutting Σ along ε .) It then follows from Theorem 2 that the

number of geodesics in $\mathcal{F}(\Sigma^*)$ between any two triangulations in \mathcal{G} is at most a polynomial function of their distance as that distance goes to infinity.

Acknowledgments. Hugo Parlier is supported by the Luxembourg National Research Fund OPEN grant O19/13865598. Most of the results reported here were obtained in June 2022 while the first author was visiting the second at the University Paris 13, thanks to funding from the MathSTIC (CNRS FR3734) research consortium.

REFERENCES

- [1] J. ARAMAYONA, H. PARLIER, and K. J. SHACKLETON. Totally geodesic subgraphs of the pants complex. *Math. Res. Lett.* **15** (2008), 309–320.
- [2] F. BALACHEFF. Volume entropy, weighted girths and stable balls on graphs. *J. Graph Theory* **55** (2007), 291–305.
- [3] C. CEBALLOS and V. PILAUD. The diameter of type D associahedra and the non-leaving-face property. *J. Combin.* **51** (2016), 109–124.
- [4] V. DISARLO and H. PARLIER. The geometry of flip graphs and mapping class groups. *Trans. Amer. Math. Soc.* **372** (2019), 3809–3844.
- [5] E. GUTKIN and V. SCHROEDER. Connecting geodesics and security of configurations in compact locally symmetric spaces. *Geom. Dedicata* **118** (2006), 185–208.
- [6] C. W. LEE. The associahedron and triangulations of the n -gon. *European J. Combin.* **10** (1989), 551–560.
- [7] H. A. MASUR and Y. N. MINSKY. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* **138** (1999), 103–149.
- [8] H. PARLIER and L. POURNIN. Flip-graph moduli spaces of filling surfaces. *J. Eur. Math. Soc.* **19** (2017), 2697–2737.
- [9] H. PARLIER and L. POURNIN. Modular flip-graphs of one holed surfaces. *European J. Combin.* **67a** (2018)a, 158–173.
- [10] H. PARLIER and L. POURNIN. Once punctured disks, non-convex polygons, and pointihedra. *Ann. Comb.* **22b** (2018)b, 619–640.
- [11] L. POURNIN. The diameter of associahedra. *Adv. Math.* **259** (2014), 13–42.
- [12] L. POURNIN and Z. WANG. Strong convexity in flip-graphs. arXiv: [2106.08012](https://arxiv.org/abs/2106.08012) (2021).
- [13] B. SCHMIDT and J. SOUTO. A characterization of round spheres in terms of blocking light. *Comment Math. Helv.* **85** (2010), 259–271.
- [14] D. SLEATOR R. TARJAN and W. THURSTON. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.* **1** (1988), 647–681.
- [15] J. D. STASHEFF. Homotopy associativity of H-spaces. I. *Trans. Amer. Math. Soc.* **108**(a) (1963), 275–292.
- [16] J. D. STASHEFF. Homotopy associativity of H-spaces. II. *Trans. Amer. Math. Soc.* **108**(b) (1963), 293–312.
- [17] D. TAMARI. Monoïdes préordonnés et chaînes de Malcev. Thèse de Mathématiques, Paris (1951).