J. Eur. Math. Soc. 19, 2697–2737
DOI 10.4171/JEMS/726

Hugo Parlier · Lionel Pournin

Flip-graph moduli spaces of filling surfaces

Received October 15, 2014

Abstract. This paper is about the geometry of the flip-graphs associated to triangulations of surfaces. More precisely, we consider a topological surface with a privileged boundary curve and study the space of its triangulations with n vertices on this curve. The surfaces we consider topologically fill this boundary curve, so we call them *filling surfaces*. The associated flip-graphs are infinite whenever the mapping class group of the surface (the group of self-homeomorphisms up to isotopy) is infinite, and we can obtain moduli spaces of flip-graphs by considering these graphs up to the action of the mapping class group. This always results in finite graphs, which we call modular flip-graphs. Our main focus is on the diameter growth of these graphs as n increases. We obtain general estimates that hold for filling surfaces of any topological type. We find more precise estimates for certain families of filling surfaces and obtain asymptotic growth results for several of them. In particular, we find the exact diameter of modular flip-graphs when the filling surface is a cylinder with a single vertex on the non-privileged boundary curve.

Keywords. Flip-graphs, triangulations of surfaces, combinatorial moduli spaces

1. Introduction

Triangulations of surfaces are very natural objects that appear in the study of topological, geometric, algebraic, probabilistic, and combinatorial aspects of surfaces and related topics. We are interested in a natural structure on spaces of triangulations: *flip-graphs*. Vertices of flip-graphs are triangulations, and two triangulations span an edge if they differ by a single arc (our base surface is a topological object and we consider triangulations up to vertex-preserving isotopy). When edge lengths are all set to one, flip-graphs are geometric objects that provide a measure for how different triangulations can be.

Flip-graphs appear in different contexts and take different forms. As flipping an arc (replacing an arc by another one) does not change either the vertices or the topology of the surface, flip-graphs correspond to triangulations of homeomorphic surfaces with a prescribed set of vertices. Provided the surface has enough topology, flip-graphs are infinite,

H. Parlier: Mathematics Research Unit, University of Luxembourg, Luxembourg;
e-mail: hugo.parlier@uni.lu

L. Pournin: LIPN, Université Paris 13, Villetaneuse, France;
e-mail: lionel.pournin@univ-paris13.fr

Mathematics Subject Classification (2010): Primary 05C12; Secondary 57M50

and self-homeomorphisms of the surface act on this graph as isomorphisms. In fact, modulo some exceptional cases, the group of self-homeomorphisms of the surface up to isotopy (the *mapping class group*) is exactly the automorphism group of the graph [10]. The quotient of a flip-graph via its automorphism group is finite, and thus via the Schwarz–Milnor Lemma (see for example [1]), a flip-graph and the associated mapping class group are quasi-isometric.

Furthermore, if one gives the triangles in a triangulation a given geometry, each triangulation corresponds to a geometric structure on a surface. In this direction, Brooks and Makover [2] defined *random surfaces* to be geometric surfaces coming from a random triangulation where each triangle is an ideal hyperbolic triangle. This notion of a random surface is a way of sampling points in *Teichmüller and moduli spaces*—roughly speaking, the space of hyperbolic metrics on a given topological structure. Although in the above it is only the vertex set of flip-graphs that appears, in the theory of *decorated* Teichmüller spaces, flip-graphs play an actual role [15]. In a similar direction, Fomin, Shapiro, and Thurston [6], and more recently Fomin and Thurston [7], have used flip-graphs and their variants to study cluster algebras that come from the Teichmüller theory of bordered surfaces. For all of these reasons, flip-graphs and their relatives appear frequently and importantly in the study of moduli spaces, surface topology, and mapping class groups.

In a different context, flip-graphs are important objects for the study of triangulations of arbitrary dimension, whose vertices are placed in a Euclidean space and whose simplices are embedded linearly (see [3] and references therein). In this case, flip-graphs are always finite, and they are sometimes isomorphic to the graph of a polytope, or admit subgraphs that have this property. Such flip-graphs emerge for instance from the study of generalized hypergeometric functions and discriminants [9] and from the theory of cluster algebras [8]. The simplest non-trivial case is that of the flip-graph of a polygon, which turns out to be the graph of a celebrated polytope—the associahedron [11]. The study of this graph has an interesting history of its own [19], and one of the reasons it has attracted so much interest is that it pops up in surprisingly different contexts, including theoretical physics and computer science (see for instance [11, 17, 18, 20]).

Associahedra appear, in particular, in the work of Sleator, Tarjan, and Thurston [17] on the dynamic optimality conjecture. They proved the theorem below about the diameter of these polytopes for sufficiently large n, using constructions of polyhedra in hyperbolic 3-space. Their proof, however, does not tell us how large n should be for the theorem to hold. The second author proved this theorem whenever $n > 12$ using combinatorial arguments [16]. Note that for smaller n the diameter behaves differently.

Theorem ([16, 17]). *The flip-graph of a convex polygon with n vertices has diameter $2n - 10$ whenever $n > 12$.*

This theorem is in some sense our starting point. The topology of a polygon is the simplest that one can imagine—it is simply the boundary circle filled by a disk. Our basic question is the following: *what happens when one replaces the disk by a surface with more topology?* These surfaces, which we call *filling* as they fill the boundary circle, give rise to infinite flip-graphs as soon as the mapping class group is infinite. We are interested in precisely these cases here. Up to homeomorphisms that preserve the circle boundary pointwise, we get nice finite combinatorial moduli spaces of triangulations whose geometry, and in particular whose diameter, we study.

We note that the filling surfaces with finite flip-graphs are the disk, the Möbius band, and the disk with a single puncture. In addition to the case of the disk [16, 17], the Möbius band has been discussed in [5], and the once-punctured disk in [14].

Precise definitions and notation can be found in the next section—but in order to state our results, we briefly describe them here. We will consider a filling surface Σ (a topological surface with a privileged boundary curve) and denote by Σ_n the same surface with n marked points on the privileged boundary. The *modular flip-graph* $\mathcal{MF}(\Sigma_n)$ is the flip-graph of Σ_n up to homeomorphism. For example $\mathcal{MF}(\Sigma_n)$ is the graph of the associahedron when Σ is a disk.

Our first result is the following upper bound for the diameter of $\mathcal{MF}(\Sigma_n)$ which does not asymptotically depend on the topology of the filling surface.

Theorem 1.1. *For any filling surface Σ there exists a constant K_Σ such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq 4n + K_\Sigma.$$

A simple consequence of this result and of the monotonicity of $\operatorname{diam}(\mathcal{MF}(\Sigma_n))$, proven in Section 2.3, is that the limit

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathcal{MF}(\Sigma_n))}{n}$$

exists (and is less than or equal to 4). Again, in the case of the associahedron, this limit is 2. It is perhaps not a priori obvious why the limit should not *always* be 2, independently of the topology of Σ, but this turns out not to be the case.

In order to exhibit different behaviors, we study particular examples of filling surfaces. Our examples are surfaces Σ with genus 0 and $k+1$ boundaries, including the privileged one, and each of the non-privileged boundaries contains a single marked or unmarked point. We will refer to these non-privileged boundary curves with a single point as *boundary loops*. Marking or not the point they contain amounts to disallowing or allowing the mapping class group acting on the flip-graph to exchange them. We provide the following upper bounds for the diameters.

Theorem 1.2. *Let Σ be a filling surface with $k \geq 2$ marked boundary loops and no other topology. There exists a constant K_k which only depends on k such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq (4 - 2/k)n + K_k.$$

Similarly:

Theorem 1.3. *Let Σ be a filling surface with $k \geq 1$ unmarked boundary loops and no other topology. There exists a constant K_k which only depends on k such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(3 - \frac{1}{2k}\right)n + K_k.$$

The constants K_k in both these theorems are a priori unrelated. In the case of the associahedron, upper bounds of the correct order (i.e. $2n$) are somewhat immediate, but here the upper bounds, although not particularly mysterious, are more involved.

Now consider the filling surface Γ with a unique boundary loop—being marked or unmarked does not matter in this case. This surface will play a particular role in our paper and we are able to obtain the exact diameter of its modular flip-graphs:

Theorem 1.4. *The diameter of the modular flip-graphs of* Γ *satisfies*

$$\mathrm{diam}(\mathcal{MF}(\Gamma_n)) = \lfloor 5n/2 \rfloor - 2.$$

This shows that Theorem 1.3 is asymptotically sharp when $k = 1$. As for the associahedron, the hard part is the lower bound. We note in the final section that the lower bound from this theorem proves a general lower bound on the diameter of $\mathcal{MF}(\Sigma_n)$, provided Σ has at least one interior marked point and any additional topology (for instance any genus or any additional marked points or boundary loops).

Our final main result is about the filling surface with genus 0 and exactly two marked boundary loops—we call this particular surface Π as we give it special attention.

We prove the following.

Theorem 1.5. *The diameter of* $\mathcal{MF}(\Pi_n)$ *is not less than* $3n$.

This result and the upper bound from Theorem 1.2 when $k = 2$ show that the diameter of $\mathcal{MF}(\Pi_n)$ grows like $3n$ (with constant error term).

Our lower bounds always come from somewhat involved combinatorial arguments, using the methods introduced in [16]. Boundary loops play an important part, since to ensure that two triangulations are far apart, we show that moving these loops necessarily entails a certain number of flips.

The remainder of the article is organized as follows. We begin with a section devoted to preliminaries which include notation and basic or known results we need. As the results may be of interest to people with different mathematical backgrounds, we spend some time talking about the setup in order to keep the article as self-contained as possible. The third section deals with the upper bounds, and the fourth and fifth sections with the lower bounds. In the final section, we discuss some consequences of our results and we conclude the article with several questions and conjectures about what the more general picture might look like.

2. Preliminaries

In this section we describe in some detail the objects we are interested in, introduce notation and some of the tools we use. In particular, the methods from [16] are generalized to arbitrary filling surfaces in Subsection 2.2.

2.1. Filling surfaces and flip-graphs

We consider a topological orientable surface Σ with the following three properties.

Property 1. The surface Σ has at least one boundary curve, and we think of one of the boundary curves as being special. We will refer to it as the *privileged boundary* (it has no marked or unmarked points on it, but will be endowed with them in what follows). The other boundary curves are non-privileged.

Property 2. All non-privileged boundary curves of Σ have at least one marked or unmarked point on it. This is because we need to triangulate Σ and these points are necessary to do so. The distinction between marked and unmarked points will become clearer in the following, but note that if a boundary curve contains one marked point, all the other points on this boundary are naturally marked, as their position relative to the marked point determines a marking. Also note that most of the specific examples we study in more detail have only one point on each non-privileged boundary curve. For this reason, we use the term *boundary loop* for a boundary curve with a single point.

Property 3. The surface Σ is of finite type. It can have genus, marked or unmarked points in its interior or on its non-privileged boundary curves, but only a finite number of each. Another way of saying this is to ask that its group of self-homeomorphisms (up to isotopy) be finitely generated (but generally not finite).

We illustrate Σ in Fig. 1 with its different possible features. Note that if it has no topology, then Σ is simply a disk.

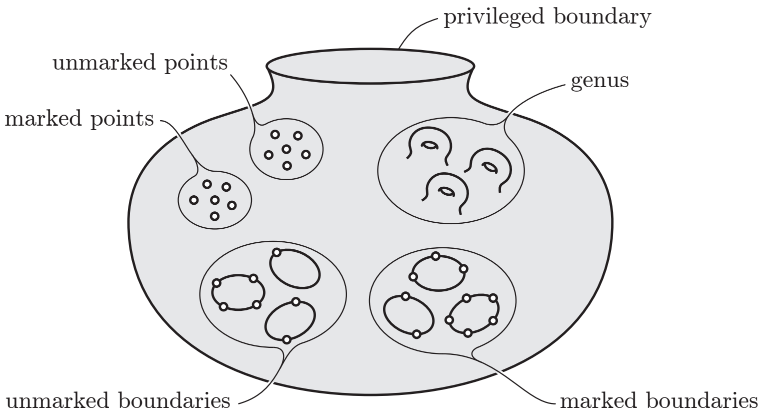

Fig. 1. Σ and its possible features

For any positive integer n, from Σ we obtain a surface Σ_n by placing n marked points on the privileged boundary of Σ. We are interested in triangulating Σ_n and studying the geometry of the resulting flip-graph. We fix Σ_n, and we refer to its set of marked and unmarked points as its *vertices*. An *arc* of Σ_n is an isotopy class of non-oriented simple paths between (not necessarily distinct) vertices. A *multi-arc* is a collection of arcs whose interiors can be realized disjointly (they can share vertices but cannot cross).

From arcs, one can construct a simplicial complex called the arc complex. This complex is well studied in geometric topology; it is built by associating simplices to sets of arcs that can be realized disjointly. A *triangulation* of Σ_n is a maximal collection of arcs that can be realized disjointly. Said differently, triangulations are maximal multi-arcs with respect to inclusion. Although they are not necessarily "proper" triangulations in the usual sense, they do cut the surface into a collection of triangles. Also note that a multi-arc can always be completed to a triangulation.

For fixed Σ_n, the number of interior arcs of a triangulation is a fixed number. Note that, by an Euler characteristic argument, it increases linearly in n.

We now construct the *flip-graph* $\mathcal{F}(\Sigma_n)$. The vertices of $\mathcal{F}(\Sigma_n)$ are the triangulations of Σ_n, and two vertices share an edge if they coincide in all but one arc. Another way of seeing this is that they share an edge if they are related by a single *flip* operation, as shown in Fig. 2. The resulting flip-graph is sometimes finite, sometimes infinite, but it is always locally finite and connected, as any isotopy class of arcs can be introduced into a triangulation by a finite number of flips (see for instance [12]).

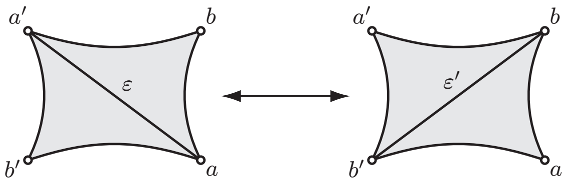

Fig. 2. The flip that exchanges the arcs ε and ε'.

When Σ is a disk, $\mathcal{F}(\Sigma_n)$ is a finite graph (it is the graph of the associahedron). An example of an infinite flip-graph is given by the surface of genus 0 with a unique boundary loop, and no marked or unmarked points in its interior. It is thus a cylinder with one of the boundary curves being the privileged boundary and the other a boundary loop. This surface, which we denote by Γ for future reference, is depicted on the left of Fig. 3. In this figure, the marked point on the boundary loop is denoted by a_0 and the privileged boundary is shown on the outside.

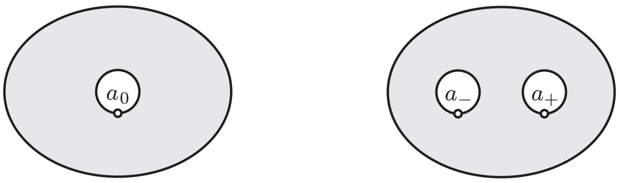

Fig. 3. The filling surfaces Γ and Π.

A triangulation of Γ_1 always contains two interior arcs between a_0 and the other marked point placed on the privileged boundary. Both arcs can be flipped, so $\mathcal{F}(\Gamma_1)$ is everywhere of degree 2. Furthermore, since there are infinitely many isotopy classes of arcs (one can think of arcs winding around the cylinder), there are infinitely many triangulations and $\mathcal{F}(\Gamma_1)$ is infinite. Being connected, infinite, and regular of degree 2, $\mathcal{F}(\Gamma_1)$ is isomorphic to the infinite line graph ($\mathbb{Z}$ with its obvious graph structure).

In general, whenever $\mathcal{F}(\Sigma_n)$ is infinite, there is a non-trivial natural action of the group of self-homeomorphisms of Σ_n on $\mathcal{F}(\Sigma_n)$. This is because homeomorphisms will preserve the property of two triangulations being related by a flip, so they induce a simplicial action on $\mathcal{F}(\Sigma_n)$. It is where the importance of being a marked or an unmarked point plays a part. We allow homeomorphisms to exchange unmarked points (but fix them globally as a set). In contrast, they must fix all marked points individually. We denote by $\mathrm{Mod}(\Sigma_n)$ the group of such homeomorphisms up to isotopy. Note that once $n \geq 3$, by the action on the privileged boundary of Σ, all such homeomorphisms are orientation preserving. As we are primarily interested in large n, we do not need to worry about orientation reversing homeomorphisms.

The combinatorial moduli spaces we are interested in are thus

$$\mathcal{MF}(\Sigma_n) = \mathcal{F}(\Sigma_n)/\mathrm{Mod}(\Sigma_n).$$

Observe that this always gives rise to connected finite graphs. To unify notation, we denote the corresponding flip-graph by $\mathcal{MF}(\Sigma_n)$ even if the homeomorphism group action is trivial. We think of these graphs as discrete metric spaces where points are vertices of the graphs and the distance is the usual graph distance with edge length 1. In particular, some of these graphs have loops (a single edge from a vertex to itself), but adding or removing a loop gives rise to an identical metric space. For this reason, we think of these graphs as not having any loops.

Our main focus is on the diameter of $\mathcal{MF}(\Sigma_n)$, which we denote $\mathrm{diam}(\mathcal{MF}(\Sigma_n))$. For fixed Σ, we will be interested in how this diameter grows as a function of n. In order to exhibit distant triangulations, we will spend some time studying filling surfaces of particular topological types. One of them is Γ, already described above. It has one boundary loop (a non-privileged boundary with a single marked point). Similarly we shall consider the filling surface Π shown on the right of Fig. 3. It has genus 0, exactly two marked boundary loops (we distinguish between them) and no interior marked or unmarked point. This surface is thus a sphere with three holes: one of them is the privileged boundary and the other two are boundary loops, each with a single marked point. These points are respectively denoted by a_- and a_+ in Fig 3.

2.2. *Deleting a vertex on the privileged boundary*

One of the main ingredients used in [16] to obtain lower bounds on flip distances is the operation of deleting a vertex from a triangulation. Here, we will use this operation to the same end. When n is greater than 1, vertices in the privileged boundary will be deleted from triangulations of a given surface Σ_n, resulting in triangulations of Σ_{n-1}.

Consider a filling surface Σ_n. We label the vertices placed on the privileged boundary from a_1 to a_n in such a way that two vertices with consecutive indices are also consecutive on the boundary. Furthermore, the boundary arc with vertices a_p and a_{p+1} will be denoted by α_p, and the boundary arc with vertices a_n and a_1 by α_n.

Now consider a triangulation T of Σ_n. Some triangle t of T, depicted on the left of Fig. 4, is incident to the arc α_p. Assuming that $n > 1$, this triangle necessarily has two other distinct edges. Denote these edges by β_p and γ_p as shown in the figure. Deleting

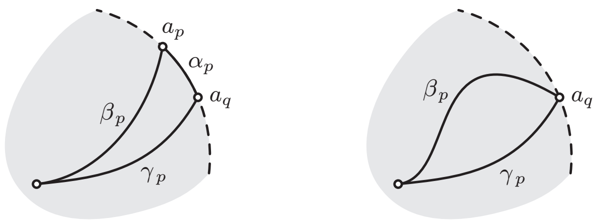

Fig. 4. The triangle incident to α_p in some triangulation of Σ_n (left), and what happens to it when a_p is displaced to the other vertex of α_p (right).

the vertex a_p consists in displacing this vertex along the boundary to the other vertex of α_p, and removing the arc β_p from the resulting set of arcs. Observe in particular that the displacement of the vertex a_p removes a_p from the privileged boundary and the arc α_p from the triangulation, as shown on the right of Fig. 4. Moreover, the arcs β_p and γ_p have then become isotopic, and the removal of β_p results in a triangulation of Σ_{n-1}.

Note that the deletion operation preserves triangulation homeomorphy. Therefore, this operation carries over to moduli of flip-graphs, and transforms any triangulation in $\mathcal{MF}(\Sigma_n)$ into a triangulation in $\mathcal{MF}(\Sigma_{n-1})$. The triangulation obtained by deleting the vertex a_p from T is denoted $T\backslash\backslash p$ in the remainder of the paper, following [16]. This notation will be used for triangulations in both $\mathcal{F}(\Sigma_n)$ and $\mathcal{MF}(\Sigma_n)$.

Consider two triangulations U and V in $\mathcal{MF}(\Sigma_n)$ and assume that they can be obtained from one another by a flip. The following proposition shows that the relation between $U\backslash\backslash p$ and $V\backslash\backslash p$ can be of two kinds.

Proposition 2.1. *Suppose $n \geq 2$. If U and V are triangulations in $\mathcal{MF}(\Sigma_n)$ related by a flip, then $U\backslash\backslash p$ and $V\backslash\backslash p$ are either identical or related by a flip.*

Proof. Consider the quadrilateral whose diagonals are exchanged by the flip relating U and V. The deletion of the vertex a_p either shrinks this quadrilateral to a triangle, deforms it to another quadrilateral, or leaves it unaffected. In the first case, $U\backslash\backslash p$ and $V\backslash\backslash p$ are identical because the deletion then removes the two arcs exchanged by the flip. In the other two cases, $U\backslash\backslash p$ and $V\backslash\backslash p$ can also be identical (while vertex deletion preserves homeomorphy, it does not always preserve non-homeomorphy), but if they are not, they differ exactly on the (possibly deformed) quadrilateral. More precisely, they can be obtained from one another by the flip that exchanges the diagonals of this quadrilateral. □

In the following, a flip between two triangulations U and V in $\mathcal{MF}(\Sigma_n)$ is called *incident to the arc* α_p when $U\backslash\backslash p$ is identical to $V\backslash\backslash p$.

When Σ is a disk, the flips incident to α_p are exactly the ones that affect the triangle incident to this arc within a triangulation [16]. When Σ is not a disk, these flips are still incident to α_p, but they are not necessarily the only ones. For instance, the unique triangulation in $\mathcal{MF}(\Gamma_1)$ and the four triangulations in $\mathcal{MF}(\Gamma_2)$ are depicted in Fig. 5. Since $\mathcal{MF}(\Gamma_1)$ has a single element, we have:

Proposition 2.2. *If T is one of the four triangulations in $\mathcal{MF}(\Gamma_2)$, then any flip performed in T is incident to both α_1 and α_2.*

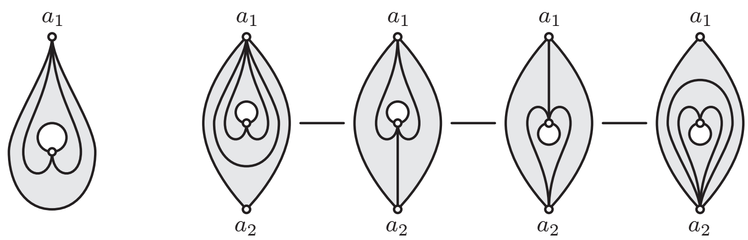

Fig. 5. The unique triangulation in $\mathcal{MF}(\Gamma_1)$ (left) and the four triangulations in $\mathcal{MF}(\Gamma_2)$. The lines between the latter four triangulations depict $\mathcal{MF}(\Gamma_2)$.

Fig. 5 also shows the edges of the modular flip-graph of Γ_2. In this flip-graph, the third triangulation from the left is obtained from the second one by replacing any of the two interior arcs incident to a_1 by an interior arc incident to a_2. Assume that the removed arc is the one on the left. In this case, the triangle incident to α_1 is not affected by the flip. Yet, via Proposition 2.2, this flip is incident to α_1.

Now assume that U and V are two arbitrary triangulations that belong to $\mathcal{MF}(\Sigma_n)$. Consider a sequence $(T_i)_{0\leq i\leq k}$ of triangulations in $\mathcal{MF}(\Sigma_n)$ such that $T_0 = U$, $T_k = V$, and T_{i-1} can be transformed into T_i by a flip whenever $0 < i \leq k$. Such a sequence will be called a *path* of length k from U to V, and can be alternatively thought of as a sequence of k flips that transform U into V. According to Proposition 2.1, removing unnecessary triangulations from the sequence $(T_i\backslash\backslash p)_{0\leq i\leq k}$ results in a path from $U\backslash\backslash p$ to $V\backslash\backslash p$ and the number of triangulations that need be removed from the sequence is equal to the number of flips incident to α_p along $(T_i)_{0\leq i\leq k}$. In other words:

Lemma 2.3. *Let U and V be two triangulations in $\mathcal{MF}(\Sigma_n)$. If f flips are incident to the arc α_p along a path of length k between U and V, then there exists a path of length $k - f$ between $U\backslash\backslash p$ and $V\backslash\backslash p$.*

Note that when Σ is a disk, this lemma is exactly Theorem 3 from [16]. A path between two triangulations U and V in $\mathcal{MF}(\Sigma_n)$ is called a *geodesic* if its length is minimal among all the paths between U and V. The length of any such geodesic is equal to the distance of U and V in $\mathcal{MF}(\Sigma_n)$, denoted by $d(U, V)$.

Invoking Lemma 2.3 with a geodesic between U and V immediately yields:

Theorem 2.4. *Let $n > 1$ and consider two triangulations U and V in $\mathcal{MF}(\Sigma_n)$. If there exists a geodesic between U and V along which at least f flips are incident to the arc α_p, then the following inequality holds:*

$$d(U, V) \geq d(U\backslash\backslash p, V\backslash\backslash p) + f.$$

In well defined situations, at least two flips are incident to a given boundary arc along any geodesic. This may be the case when one of the triangulations at the ends of the geodesic has a well placed *ear*, i.e. a triangle with two edges in the privileged boundary, as shown on the left of Fig. 6. In the figure, these two edges are α_p and α_q, and the vertex they share is a_q. In this case, we will say that the triangulation has an ear *at* a_q.

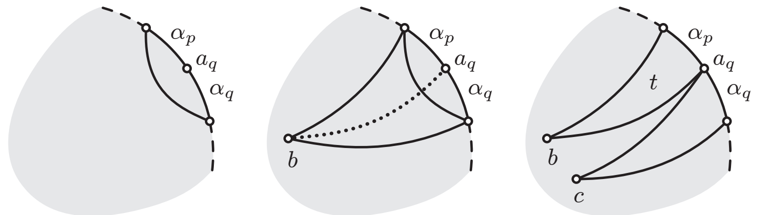

Fig. 6. The triangulations U (left) and V (right) from the statement of Lemma 2.5. The j-th flip along the geodesic used in the proof of this lemma is shown in the middle, where the solid edges belong to T_{j-1}, and the arc introduced is dotted.

The following result, proven in [16] when Σ is a disk, holds in general:

Lemma 2.5. *Consider two triangulations U and V in $\mathcal{MF}(\Sigma_n)$. Further consider two distinct arcs α_p and α_q in the privileged boundary of Σ_n such that a_q is a vertex of α_p. If U has an ear at a_q and if the triangles of V incident to α_p and to α_q do not have a common edge, then for any geodesic between U and V, there exists $r \in \{p, q\}$ such that at least two flips along this geodesic are incident to α_r.*

Proof. Assume that U has an ear at a_q and that the triangles of V incident to α_p and to α_q do not have a common edge. In this case, U and V are as shown on the left and on the right, respectively, of Fig. 6. Note that the vertices b and c represented in this figure can be identical. At least one flip along any path between U and V is incident to the arc α_p because the triangles of U and of V incident to this arc are distinct.

Consider a geodesic $(T_i)_{0\leq i\leq k}$ from U to V and assume that only one of the flips along this geodesic is incident to α_p, say the j-th one. This flip must then be as shown in the middle of Fig. 6. Not only is it incident to α_p but also to α_q. Moreover, the triangle t of V incident to α_p already belongs to T_j.

Now observe that the triangle of T_j incident to α_q shares an edge with t. By assumption, the triangle of V incident to α_q does not have this property. Therefore, at least one of the last $k - j$ flips along $(T_i)_{0\leq i\leq k}$ must affect the triangle incident to α_q. This flip is then the second flip along the geodesic incident to α_q, as desired. □

2.3. A projection lemma

Here we briefly describe a result from [4] in our setting and its implications for our diameter estimates. This lemma is about two triangulations U and V of Σ_n with arcs in common. It says that these arcs must also be arcs of all the triangulations along any geodesic between U and V in the flip-graph $\mathcal{F}(\Sigma_n)$. This generalizes Lemma 3 from [17], originally proven in the case of a disk with marked boundary points. Formally:

Lemma 2.6 (Projection Lemma). *Let U and V be two triangulations of Σ_n. Further consider a geodesic $(T_i)_{0\leq i\leq k}$ from U to V in the graph $\mathcal{F}(\Sigma_n)$. If μ is a multi-arc common to U and V, then μ is also a multi-arc of T_i whenever $0 < i < k$.*

It is essential to note that the above lemma does *not* necessarily hold in $\mathcal{MF}(\Sigma_n)$. However, it clearly does hold if an arc or a multi-arc is invariant under all elements of $\mathrm{Mod}(\Sigma_n)$. Namely, consider an arc α *parallel* to the privileged boundary (by parallel we mean that the portion of Σ_n bounded by this arc and by a part of the privileged boundary is a disk). Then, as any element of $\mathrm{Mod}(\Sigma_n)$ fixes the privileged boundary arcs individually, the arc α is also invariant. In particular, assume that α has vertices a_1 and a_3. By the above, α is never removed along a geodesic between two triangulations containing this arc. So naturally we get a geodesically convex and isometric copy of $\mathcal{MF}(\Sigma_{n-1})$ inside $\mathcal{MF}(\Sigma_n)$. Thus we obtain the following.

Proposition 2.7. $\mathrm{diam}(\mathcal{MF}(\Sigma_{n-1})) \leq \mathrm{diam}(\mathcal{MF}(\Sigma_n))$.

Note that, by observing that there are points of $\mathcal{MF}(\Sigma_n)$ outside the isometric copy of $\mathcal{MF}(\Sigma_{n-1})$, it is not too difficult to see that in fact the above inequality is strict, but we make no use of that in the sequel.

3. Upper bounds

In this section we prove upper bounds on the diameter of modular flip-graphs depending on the topology of the underlying surface.

3.1. A general upper bound

We begin with the following general upper bound.

Theorem 3.1. *For any filling surface Σ there exists a constant K_Σ such that*

$$\mathrm{diam}(\mathcal{MF}(\Sigma_n)) \leq 4n + K_\Sigma.$$

Before proving the theorem, let us give the basic idea of the proof. Consider a triangulation T of Σ_n and a vertex a of this surface. Let us call the number of interior arcs of T incident to a the *interior degree* of a in T. For large enough n the average interior degree of the vertices of T can be arbitrarily close to 2, and thus given any two triangulations U and V the average sum of the interior degrees tends to 4. We can then choose a vertex a (on the privileged boundary) in such a way that its interior degree is at most 4. We perform flips within U to obtain $\tilde{U}$ and flips within V to obtain $\tilde{V}$ so that $\tilde{U}$ and $\tilde{V}$ both have an ear at a. In doing so we can now safely ignore a boundary vertex and repeat the process.

In order to quantify the number of flips each of the steps described above might cost, we first prove the following lemma.

Lemma 3.2. *For $n \geq 2$, consider a vertex a on the privileged boundary of Σ_n and two triangulations U and V of Σ_n. If the interior degrees of a in U and in V sum to at most* 4, *then there exist two triangulations $\tilde{U}$ and $\tilde{V}$ of Σ_n, each with an ear at a, such that*

$$d(U, \tilde{U}) + d(V, \tilde{V}) \leq 4.$$

Proof. We shall prove the lemma by showing that there is always a flip in either U or V that reduces the degree of a, and thus by iteration, one must flip at most four arcs to reach both $\tilde{U}$ and $\tilde{V}$. Let ε be any interior arc incident to a in either U or V.

First suppose that ε is flippable. If flipping ε reduces the degree of a, we flip it. If not, then the flip quadrilateral of ε (shown on the left of Fig. 7) must have a boundary arc, say α, with the vertex a at its two ends. This situation, sketched in the middle of Fig. 7, corresponds to when the vertex labeled b' on the left of the figure is equal to a.

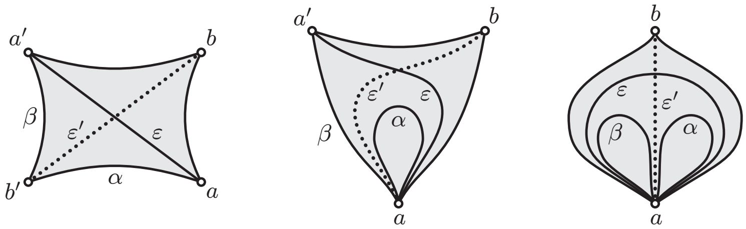

Fig. 7. The flip dealt with in the proof of Lemma 3.2 (left), and a sketch of the surface when this flip does not reduce the degree of a (middle and right).

As n is not less than 2, α must be an interior arc. In addition, α is twice incident to a and is thus flippable. If flipping α reduces the degree of a, we flip α and we can proceed. Therefore, we assume that flipping α does not decrease the degree of a. In this case, the vertex a' shown on the left of Fig. 7 is necessarily the same vertex as a. The arcs α, β and ε (see Fig. 7, right) are now three interior arcs twice incident to a. Thus the interior degree of a is at least 6, which is impossible.

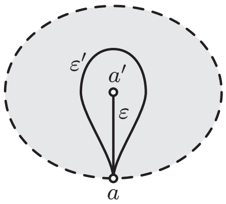

Fig. 8. When ε is not flippable.

Finally, consider the case where ε is not flippable. This arc is then surrounded by another arc ε' twice incident to a, as sketched in Fig. 8. Flipping ε' reduces the degree of a because the flip introduces an arc incident to a'. □

Note that Lemma 3.2 holds *a fortiori* when U and V belong to $\mathcal{MF}(\Sigma_n)$. We can now proceed with the proof of the theorem.

Proof of Theorem 3.1. Consider the surface Σ_1 and insert points in its privileged boundary to obtain Σ_n. The Euler characteristics of these surfaces satisfy

$$\chi(\Sigma_n) = \chi(\Sigma_1).$$

A triangulation T of Σ_n has $n-1$ more vertices and $n-1$ more triangles than a triangulation T' of Σ_1. It also has $n-1$ more boundary arcs. By invariance of the Euler characteristic, this means that T has exactly $n-1$ more *interior* edges than T'. Hence, the number of interior edges of T is exactly

$$n + E_\Sigma,$$

where E_Σ is a precise constant which depends on Σ but not on n. We now focus our attention on the interior degree of the privileged boundary vertices. The total interior degree of all these vertices is at most $2(n+E_\Sigma)$.

The sum of the interior degrees of all vertices in two triangulations U and V in $\mathcal{MF}(\Sigma_n)$ is not greater than $4(n+E_\Sigma)$. Thus the average sum of the interior degrees of the privileged boundary vertices is at most

$$4 + \frac{4}{n}E_\Sigma.$$

Therefore, for $n > 4E_\Sigma$, there exists a privileged boundary vertex a whose interior degrees in U and in V sum to at most 4.

We now apply the previous lemma to flip U and V a total of at most four times into two new triangulations with ears at a. We treat the new triangulations as if they lay in $\mathcal{MF}(\Sigma_{n-1})$ and we repeat the process inductively until $n \le 4E_\Sigma$. We end up with two triangulations $\tilde{U}$ and $\tilde{V}$ that only differ on a subsurface homeomorphic to Σ_{n_0}, where

$$n_0 \le 4E_\Sigma.$$

Hence, there is a path of length at most $\operatorname{diam}(\mathcal{MF}(\Sigma_{n_0}))$ between $\tilde{U}$ and $\tilde{V}$. We therefore obtain the following inequality:

$$d(U,V) \le 4(n-4E_\Sigma) + \operatorname{diam}(\mathcal{MF}(\Sigma_{n_0})) = 4n + K_\Sigma,$$

where $K_\Sigma = \operatorname{diam}(\mathcal{MF}(\Sigma_{n_0})) - 16E_\Sigma$ does not depend on n. □

Before looking at more precise bounds for a given surface topology, we note that, together with the monotonicity from Proposition 2.7, we have the following:

Corollary 3.3. *For any filling surface* Σ *the following limit exists and satisfies*

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathcal{MF}(\Sigma_n))}{n} \le 4.$$

3.2. Upper bounds for Γ

In this section we prove a much stronger and specific upper bound in the case where our surface is Γ, a cylinder with a single boundary loop.

Theorem 3.4. *The diameters of the modular flip-graphs of* Γ *satisfy*

$$\operatorname{diam}(\mathcal{MF}(\Gamma_n)) \le 5n/2 - 2.$$

Proof. Let U and V be triangulations in $\mathcal{MF}(\Gamma_n)$. Denote by a_0 the unique vertex not on the privileged boundary, and α_0 the boundary loop it belongs to. The basic strategy is to perform flips within both triangulations until all interior arcs are incident to a_0 and then find a path between the resulting triangulations.

We begin by observing that a triangulation in $\mathcal{MF}(\Gamma_n)$ has $n+1$ interior arcs. Furthermore, any triangulation T of Γ_n has at least two distinct interior arcs incident to a_0. Indeed, α_0 is incident to a triangle of T whose other two edges must have vertex a_0. These edges are also both incident to the same vertex on the privileged boundary. Hence, they must be interior arcs of the triangulation.

Thus, $n-1$ flips suffice to reach a triangulation with all arcs incident to a_0 from either U or V. Note that such a triangulation is uniquely determined by the privileged boundary vertex of the triangle incident to α_0.

We now perform the above flips within U and V to obtain two triangulations U' and V'. Denote by a_u and a_v the privileged boundary vertices of the triangle incident to α_0 in U' and V', respectively. This necessitates at most $2n-2$ flips.

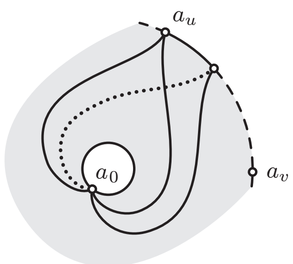

Fig. 9. The flip used in the proof of Theorem 3.4. The arc introduced is dotted.

Now to get from U' to V', we proceed as follows. Note that, thinking of the privileged boundary as a graph, the distance of a_u and a_v along this boundary is at most $n/2$. We can perform a flip in U' to obtain a triangulation similar to U', wherein the privileged boundary vertex of the triangle incident to α_0 is 1 closer to a_v along the privileged boundary (this is illustrated in Fig. 9). Thus, in at most $n/2$ flips the triangulations U' and V' can be transformed into one another. The result follows. □

It turns out that this straightforward upper bound is (somewhat surprisingly) optimal, as will be shown later on. In the next subsection we provide upper bounds for an arbitrary number of boundary loops.

3.3. Upper bounds for surfaces with multiple boundary loops

The first case we treat is that of marked boundary loops.

Theorem 3.5. *Let Σ be a filling surface with $k \geq 2$ marked boundary loops and no other topology. There exists a constant K_k, which only depends on k, such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq (4 - 2/k)\, n + K_k.$$

Proof. We begin by choosing a boundary loop α_0 and its vertex which we denote a_0. Note that as before, any triangulation has at least two interior arcs incident to a_0.

Given two triangulations U and V in $\mathcal{MF}(\Sigma_n)$ we perform flips within both triangulations until all interior arcs are incident to the vertex a_0. This can be done with at most $2n + 8k - 10$ flips for the following reason. A straightforward Euler characteristic argument shows that any triangulation in $\mathcal{MF}(\Sigma_n)$ has exactly $n + 4k - 3$ interior arcs. As observed above, at least two of these arcs are already incident to a_0, so each triangulation is at most $n+4k-5$ flips away from a triangulation with all arcs incident to a_0. We denote the resulting triangulations by U' and V'.

Triangulations with the above property are by no means canonical but they do have a very nice structure. Visually, it is useful to think of the vertex a_0 as the center of the triangulation. Most arcs (at least when n is considerably larger than k) will be arcs going from a privileged boundary vertex a_p to a_0, and will be the unique arcs doing so. However, some of them will have a companion arc (or several) also incident to the same two vertices. For this to happen, as they are necessarily non-isotopic arcs, they must enclose some topology: if they bounded a topological disk they would be isotopic, so there must be at least one loop inside. If we consider two successive arcs like this (by successive we mean belonging to the same triangle), they must be boundary arcs of a triangle with a companion loop incident to a_0. We shall refer to subsurfaces bounded by such two successive arcs as a *pod* and its subsurface bounded by the companion loop as a *pea*. A pod is depicted on the right of Fig. 10, where the pea is the striped region.

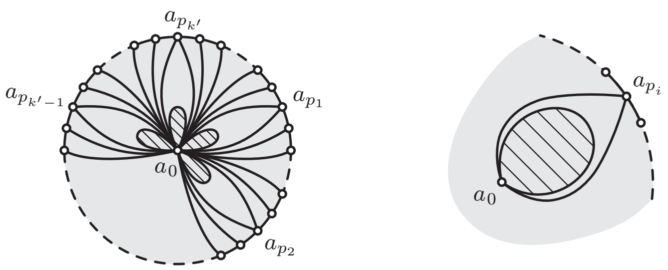

Fig. 10. Peas in pods.

Observe that any pea must contain at least one of the k interior boundary loops but could possibly contain several. Hence there are at most k peas, and as every pod is non-empty, at most k pods. We denote the number of peas and pods by k' and the privileged boundary vertices incident to the pods by a_{p_1} to $a_{p_{k'}}$ clockwise, as shown on the left of Fig. 10. Note that a_{p_j} and $a_{p_{j+1}}$ are possibly the same vertex.

The vertices $a_{p_1}, \ldots, a_{p_{k'}}$ are separated along the privileged boundary by sequences of vertices (possibly none) which have single arcs to a_0 (see Fig. 10, left). We call these sequences *gaps*. As there are n vertices on the boundary separated by at most k' pods, there is always a gap of size at least $n/k' - 1 \geq n/k - 1$, i.e., at least $n/k - 1$ consecutive vertices along the boundary are adjacent to a_0 by a single interior arc. We consider the largest gap in U' and the largest gap in V'. The sets of vertices not found in the gaps are both of cardinality at most $n - n/k + 1$.

We will distinguish two cases. First assume that some vertex, say a_g, on the privileged boundary does not belong to either the gap of U' or the gap of V'.

The strategy here is to flip U' and V' into triangulations with a single pod incident to a_g. They will thus coincide outside of the pod and it will suffice to flip inside the pod a number of times depending only on k to relate the two triangulations. We begin by observing that a pod can be moved to a neighboring vertex by a single flip unless another pod obstructs its passage (see Fig. 11, left). For both triangulations we proceed in the same way. We "condemn" the gap, and move the pods until they reach a_g without passing through the condemned gap as follows. We take one of the pods bounding the gap, say the first one clockwise, and move it clockwise until it reaches another pod or the vertex a_g. In the former case, the two pods are transformed into a single pod by the flip portrayed on the right of Fig. 11. We then continue to move the pod clockwise until reaching another pod (or a_g) etc. Once a_g has been reached, there are no pods left between a_g and the condemned gap on one side. We do the same on the other side, moving the pods counter-clockwise from the other end of the condemned gap to a_g.

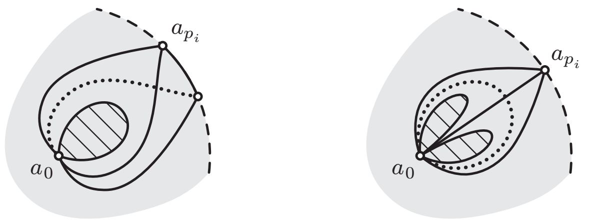

Fig. 11. A flip that moves a pod (left) and joins two pods (right). The arc introduced is dotted.

We now bound above the number of flips that were necessary to perform the transformation. As there were originally at most k pods, at most $k-1$ flips will be necessary to join pods. All of the other flips have reduced by 1 the distance between the pods bounding the condemned gap, thus there were at most $n - n/k$ such flips.

As we performed this in both triangulations, the total number of flips that have been carried out does not exceed

$$(2 - 2/k)\,n + 2k - 2.$$

We now have two triangulations U'' and V'' that differ only on a single pea containing all of the topology and where all arcs are incident to a_0. In particular, we only need to perform flips inside the pea. As a subsurface, it is homeomorphic to Σ_1, thus

$$d(U'', V'') \leq \operatorname{diam}(\mathcal{MF}(\Sigma_1)),$$

and this diameter is equal to some constant K'_k which only depends on k. Using these estimates and our original estimates on the distances to U' and V' we obtain

$$d(U, V) \leq (2 - 2/k)n + 2k - 2 + K'_k + 2n + 8k - 10.$$

If we set $K_k = K'_k + 10k - 12$, this results in the desired inequality:

$$d(U, V) \leq (4 - 2/k)n + K_k.$$

We now review our second case. Assume that each of the vertices on the privileged boundary belongs to the gap of U' or to the gap of V'.

This is the easier case since all the privileged boundary vertices incident to pods of U' lie in a sector disjoint from another sector containing all the privileged boundary vertices incident to pods of V'. As above, we move the pods by flips. Choose a pod in U' at the boundary of the gap, say the first one clockwise, and move it clockwise (i.e. without passing through the gap) towards the other boundary using the flip shown on the left of Fig. 11. This proceeds until U' is transformed into a triangulation with a single pod at the other boundary of the gap. Let a_l be the vertex on the privileged boundary that is incident to the remaining pod. We now move the pods in V' similarly but in the opposite direction: we start from the last pod clockwise, move it counter-clockwise, and merge it along the way with the other pods until we reach a vertex a_l with a single remaining pod. We denote the resulting triangulations by U'' and V''. Note that, as above, there were at most $2k-2$ flips that served to join adjacent pods. All other flips brought the outermost pods 1 closer to a_l. Hence, there were at most $n-1$ such flips. Thus,

$$d(U', U'') + d(V', V'') \leq n - 1 + 2k - 2.$$

Now U'' and V'' differ in a single pea, and so as above

$$d(U'', V'') \leq \mathrm{diam}(\mathcal{MF}(\Sigma_1)).$$

We can conclude as follows, taking the same constant K_k as previously:

$$d(U, V) \leq 3n + K_k \leq (4 - 2/k)n + K_k.$$

Note that the second inequality holds because $k \geq 2$. □

Observe that this implies an upper bound of the order of $3n$ when $k = 2$, that is, when $\Sigma = \Pi$. An adaptation of the above proof for unmarked boundary loops gives stronger upper bounds. In particular the following is true.

Theorem 3.6. *Let Σ be a filling surface with $k \geq 1$ unmarked boundary loops and no other topology. There exists a constant K_k which only depends on k such that*

$$\mathrm{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(3 - \frac{1}{2k}\right)n + K_k.$$

Proof. Let U and V be triangulations in $\mathcal{MF}(\Sigma_n)$. We begin by observing that every boundary loop is *close* to some vertex on the privileged boundary.

More precisely, consider the graph D that is dual to U, whose vertices are the triangles of U and whose edges connect two triangles that share an edge. Observe that D is connected. Let t be the triangle of U incident to some boundary loop. Consider a triangle t' of U incident to the privileged boundary that is closest to t in D, and a geodesic between t and t' in D. The only triangle incident to the privileged boundary along this geodesic is t'. Hence the length of this geodesic cannot depend on n, but only on k. For this reason, the vertex of t' on the privileged boundary is the one we call *close* to the boundary loop. Now observe that flipping the arcs of U dual to the edges of our geodesic from t' to t will introduce a triangle incident to both the boundary loop and the privileged boundary vertex it is close to. We then say that the boundary loop is *hanging off* this vertex.

We carry out the above sequence of flips for every boundary loop. Note that these flips never remove an arc incident to the privileged boundary. Hence, once a boundary loop is hanging off a privileged boundary vertex, it will be left so by the later flips. The number of flips needed to transform both U and V as described above does not depend on n, but only on k. We denote the resulting triangulations by U' and V'.

By construction, all the boundary loops of U' and V' hang off privileged boundary vertices, either alone or in a bunch as depicted in Fig. 12. Observe that two boundary curves hanging off the same vertex are necessarily separated by at least one triangle.

Fig. 12. Boundary loops "hanging off" privileged boundary vertices.

For a moment we forget all of the triangles of U' and V' that are not incident to a boundary loop. We consider the collection of privileged boundary vertices that have boundary loops hanging off them in either U' or V'. There are at most $2k$ such vertices, and as in the previous proof we consider the gaps of successive privileged boundary vertices without a boundary loop hanging off them in either triangulation. We now consider the largest gap, whose size is at least $n/(2k) - 1$.

We choose one of the privileged boundary vertices *contained* in the gap and denote it a_0. We carry out flips within both U' and V' to increase the interior degree of a_0 but (and this is important) without flipping the edges of any triangle incident to a boundary loop. Once this is done, all other arcs are incident to a_0. The vertices in the boundary loops are incident to a unique arc which joins them to a_0, as shown in Fig. 13.

The two triangulations look very similar with the exception of the placement of the boundary loops. They are all found in sectors (which we call *pods*) bounded by two arcs

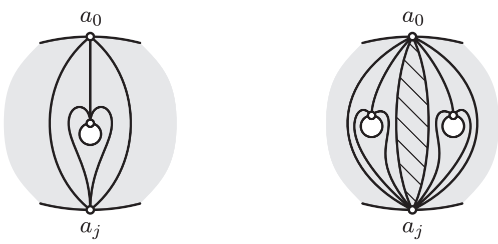

Fig. 13. A pod with a unique boundary loop (left), and one with several (right).

between a_0 and some other privileged boundary vertex a_j, possibly alone, possibly with other boundary loops (see Fig. 13). As in the previous theorem, we want to put these boundary loops in *peas* so that they are easy to move, but this time we use the privileged boundary vertex a_0 as a base for all the peas.

To do this, we perform flips inside each pod so that all the boundary loops inside a given pod become enclosed in a single pea attached to a_0. This may take a certain number of flips, but an upper bound on how many is given by

$$\operatorname{diam}(\mathcal{MF}(\Sigma_2')),$$

where Σ' is the surface inside the pod. As Σ' has at most k interior boundary curves, this is bounded by some function of k. Now, each boundary curve is inside some pea belonging to a pod attached to both a_0 and some other privileged boundary vertex a_j. This a_j is of course the original vertex that the boundary curve was *close* to.

We can now begin to move the pods around. The idea is to move the pods clockwise around a_0 using the flip depicted on the left of Fig. 14.

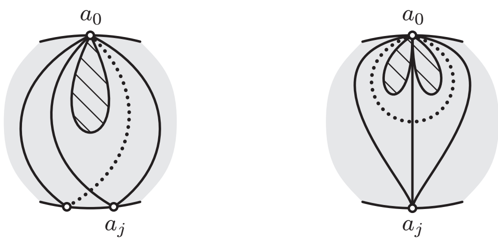

Fig. 14. A flip that moves a pod by one vertex clockwise around a_0 (left) and a flip that joins two pods (right). In each case, the arc introduced is dotted.

We will refer to the number of boundary loops in a pea or in a pod as the pea or the pod's *multiplicity*. We begin as follows: we consider the first pod clockwise around a_0 in either triangulation. If both triangulations have such a pod, we choose the one with the largest multiplicity. If they both have a pod of the same multiplicity, we leave them as they are and look for the next pod clockwise in either triangulation. The selected pod is incident to a_0 and to another privileged boundary vertex a_j.

If one of the triangulations has no pod incident to a_j, we move the pod clockwise within the one that does, until the next vertex incident to a pod on either triangulation is reached. As in the proof of the previous theorem, moving a pod by one vertex requires a single flip. This flip is depicted on the left of Fig. 14.

If however both triangulations have pods with different multiplicities incident to a_j, we first perform flips inside the one with the larger multiplicity to split it into two pods, each containing a pea attached to a_0. We make the first pod (in the direction of our orientation) with the same multiplicity as the pod of the other triangulation, and the second with whatever multiplicity comes from the leftover boundary loops. Again, this splitting operation requires a number of flips, but no more than

$$\operatorname{diam}(\mathcal{MF}(\Sigma_2')),$$

where Σ' is the surface inside the pod, as above. We then move this second pod by flips to the next vertex clockwise incident to a pod on either triangulation. Whenever the moving pod encounters another pod in its own triangulation, we perform a single flip to join them as shown on the right of Fig. 14, and we iterate the process until we reach the last pod clockwise around a_0.

The two resulting triangulations have pods of the same multiplicity incident to the same privileged boundary vertices. More precisely, these triangulations only possibly differ in the way the peas are triangulated. We therefore finally perform flips inside the peas in order to make the two triangulations coincide. Note that the number of these flips does not depend on n but only on k.

Let us now take a look at how many flips we have performed.

We began by tweaking both triangulations so that all boundary loops hung off privileged boundary vertices. This required a number of flips that does not depend on n, but only on k, which we call K'_k. We then increased the interior degree of a_0. By an Euler characteristic argument, this required at most $2n + 4k - 6$ flips. Moving pods from one end of the gap to the other required at most n flips from which the size of the gap must be subtracted, thus at most $n - n/(2k)$ flips.

In several places we had to transform two triangulations in $\mathcal{MF}(\Sigma'_2)$ into one another for some subsurface Σ' of Σ. The number of flips needed to perform every such transformation in any possible subsurface Σ' is bounded above by a number K''_k that does not depend on n. We had to do these transformations at most k times to attach the peas to a_0, and once every time a pod had to be split. The splitting operation was performed at most $2k$ times because the number of pods in the two triangulations is bounded above by $2k$. Hence the total number of flips performed to modify triangulations in $\mathcal{MF}(\Sigma'_2)$ is at most $3kK''_k$. Likewise, we may have had to join pods together, requiring in total at most $2k$ flips. The final flipping inside the peas was bounded above by a number K'''_k that does not depend on n. We therefore obtain an upper bound of

$$\left(3 - \frac{1}{2k}\right)n + K_k$$

on the diameter of $\mathcal{MF}(\Sigma_n)$, where $K_k = K'_k + 3kK''_k + K'''_k + 6k - 6$. □

3.4. A few other cases

The proof of Theorem 3.5 still works when some of the boundary loops are replaced by interior points. The only difference is that some of the peas will enclose interior points instead of boundary loops. Hence:

Theorem 3.7. *Let Σ be a filling surface with l marked boundary loops, k marked interior vertices and no other topology. If $k + l \geq 2$, then there exists a constant K_{k+l} which only depends on $k + l$ such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(4 - \frac{2}{k+l}\right)n + K_{k+l}.$$

Adapting the proof of Theorem 3.6 to surfaces with interior points and boundary loops is not immediate. Indeed, a point and a boundary loop cannot be exchanged. However, if all the boundary loops are replaced by interior vertices, a straightforward adaptation of this proof will work. As above, the only difference is that peas will enclose vertices instead of boundary loops, so we only give the main steps.

Theorem 3.8. *Let Σ be a disk with $k \geq 2$ unmarked interior vertices. There exists a constant K_k which only depends on k such that*

$$\operatorname{diam}(\mathcal{MF}(\Sigma_n)) \leq \left(3 - \frac{1}{2(k-1)}\right)n + K_k.$$

Proof. Given any two triangulations U and V, we begin by choosing any interior vertex and perform flips to increase its incidence in both triangulations. This requires $2n$ flips in total plus a constant that only depends on k. The resulting triangulations now have peas in pods where the peas have the form of a loop surrounding a single arc between an interior vertex and the vertex incident to all interior arcs.

As in the proof of Theorem 3.6, we consider the largest gap between two pods and move them around in an almost identical fashion. The gap is of size at least $n/(2k-2)$ as we have already used one of the interior vertices as the "center" of the triangulation.

The remaining details of the proof are identical to those in the proof of Theorem 3.6 and we leave them to the dedicated reader. ◻

4. Lower bounds for Γ

In this section, we prove the following lower bound on the diameter of $\mathcal{MF}(\Gamma_n)$:

$$\operatorname{diam}(\mathcal{MF}(\Gamma_n)) \geq \lfloor 5n/2 \rfloor - 2. \tag{4.1}$$

This will be done by exhibiting two triangulations A_n^- and A_n^+ in $\mathcal{MF}(\Gamma_n)$ whose distance is equal to the right-hand side of (4.1). These triangulations are built by modifying the triangulation Z_n of Δ_n depicted in Fig. 15, where Δ_n is a disk with n marked vertices on the boundary. The interior arcs of Z_n form a zigzag, i.e., a simple path that alternates between left and right turns. This path starts at the vertex a_n and ends at $a_{n/2}$ when n is even, and at $a_{\lceil n/2 \rceil + 1}$ when n is odd. When $n > 3$, Z_n has an ear at a_1 and another ear

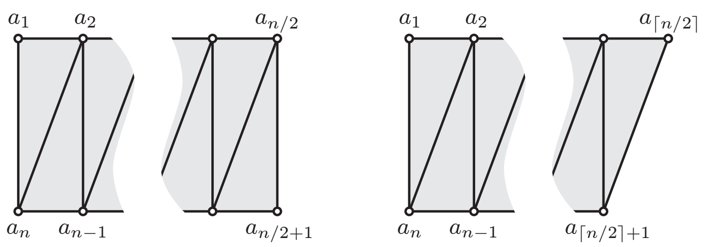

Fig. 15. The triangulation Z_n of Δ_n depicted when n is even (left) and odd (right).

at $a_{\lfloor n/2\rfloor+1}$. When $n=3$, this triangulation is made up of a single triangle which is an ear at all three vertices. Note that Z_n cannot be defined when $n<3$.

Assume that $n\geq 3$. A triangulation A_n^- of Γ_n can be built by considering the ear of Z_n in a_1 and by "piercing" it. Formally, we place a boundary loop α_0 with a vertex a_0 inside the ear and re-triangulate the pierced ear as shown in the top row of Fig. 16. Another triangulation A_n^+ of Γ_n can be built by piercing the ear of Z_n at $a_{\lfloor n/2\rfloor+1}$, by placing the vertex a_0 on the boundary of the resulting hole, thus creating a boundary loop α_0, and by re-triangulating the pierced ear as shown in the bottom row of Fig. 16.

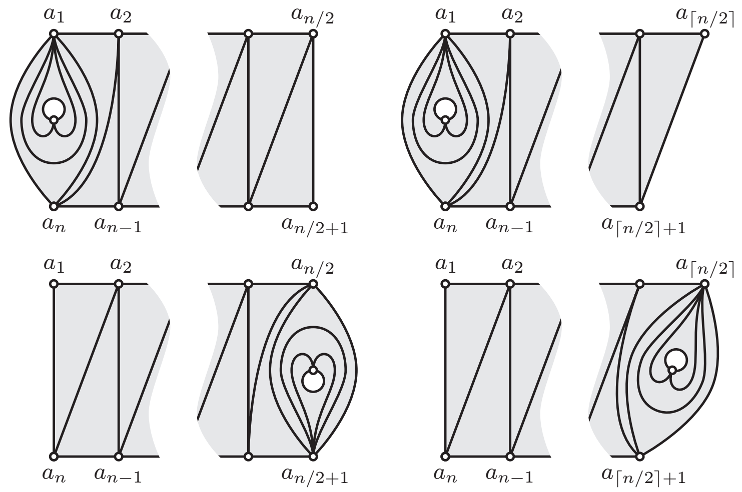

Fig. 16. The triangulations A_n^- (top row) and A_n^+ (bottom row) of Γ_n depicted when n is even (left) and odd (right). For simplicity, the vertex a_0 is unlabeled here.

In the remainder of the section, the triangulations A_n^- and A_n^+ are understood as elements of $\mathcal{MF}(\Gamma_n)$, that is, up to homeomorphism.

We will also define A_n^- and A_n^+ when $1\leq n\leq 2$. The triangulations A_2^- and A_2^+ are the triangulations in $\mathcal{MF}(\Gamma_2)$ that contain a loop arc at a_1 and a_2, respectively, as shown in Fig. 5. The triangulations A_1^- and A_1^+ will both be equal to the unique triangulation in $\mathcal{MF}(\Gamma_1)$, also shown in Fig. 5.

One of the main steps of our proof will be to show that for every $n>2$,

$$d(A_n^-,A_n^+)\geq \min\big(d(A_{n-1}^-,A_{n-1}^+)+3,\,d(A_{n-2}^-,A_{n-2}^+)+5\big). \tag{4.2}$$

This inequality will be obtained using well chosen vertex deletions or sequences of them. For instance, for $n\geq 2$, observe that deleting the vertex a_n from both A_n^- and A_n^+ results in triangulations isomorphic to A_{n-1}^- and A_{n-1}^+. More precisely, once a_n has been deleted, the other vertices need to be relabeled in order to obtain A_{n-1}^- and A_{n-1}^+. If we delete any a_j instead of a_n, then the natural relabeling amounts to relabeling a_i as a_{i-1} whenever $i>j$. This relabeling provides a map onto the triangulations of Γ_{n-1}. For future reference we call any such map a *vertex relabeling*. We can now precisely state the observation we need: the triangulations $A_n^-\backslash\backslash n$, resp. $A_n^+\backslash\backslash n$ are isomorphic to A_{n-1}^-, resp. A_{n-1}^+ via the same vertex relabeling. This can be checked using Fig. 5 when $2 \le n \le 4$ and Fig. 16 when $n \ge 3$. According to Theorem 2.4, it follows from this observation that if there exists a geodesic between A_n^- and A_n^+ with at least three flips incident to α_n, then

$$d(A_n^-, A_n^+) \ge d(A_{n-1}^-, A_{n-1}^+) + 3, \tag{4.3}$$

and inequality (4.2) holds in this case. Now assume that $n \ge 3$ and observe that for any integer i with $1 \le i < n$ and any $j \in \{n-i, n-i+1\}$, deleting the vertices a_i and a_j from A_n^- and A_n^+ results in triangulations of Γ_{n-2} isomorphic to A_{n-2}^- and A_{n-2}^+, respectively. The isomorphism comes from the vertex relabeling described above. Hence, if there exists a geodesic between A_n^- and A_n^+ with at least three flips incident to α_i, and a geodesic between $A_n^-\backslash\backslash i$ and $A_n^+\backslash\backslash i$ with at least two flips incident to α_j, then invoking Theorem 2.4 twice yields

$$d(A_n^-, A_n^+) \ge d(A_{n-2}^-, A_{n-2}^+) + 5, \tag{4.4}$$

and inequality (4.2) also holds in this case. Observe that (4.3) and (4.4) follow from the existence of particular geodesic paths. The rest of the section is devoted to proving the existence of geodesic paths that imply at least one of these inequalities.

Since α_n is not incident to the same triangle in A_n^- and in A_n^+, at least one flip is incident to this arc along any geodesic from A_n^- to A_n^+. We will study the geodesics between A_n^- and A_n^+ depending on which arc is introduced by their first flip incident to α_n. This is the purpose of the next lemmas.

Lemma 4.1. *Let $n > 2$. Consider a geodesic from A_n^- to A_n^+ whose first flip incident to the arc α_n introduces an arc with vertices a_0 and a_n. If α_n is incident to at most two flips along this geodesic, then α_1 is incident to at least three flips along it.*

Proof. Let $(T_i)_{0\le i\le k}$ be a geodesic from A_n^- to A_n^+. Assume that the first flip incident to α_n along $(T_i)_{0\le i\le k}$ is the j-th one, and that it introduces an arc with vertices a_0 and a_n. This flip must then be the one shown on the left of Fig. 17.

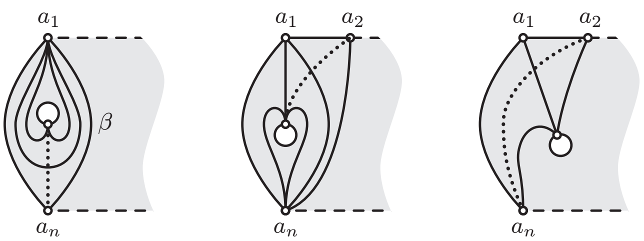

Fig. 17. The j-th (left), l'-th (middle), and l-th (right) flips performed along the path $(T_i)_{0\le i\le k}$ in the proof of Lemma 4.1. In each case, the arc introduced is dotted.

Assume that at most one flip along $(T_i)_{0\le i\le k}$ other than the j-th one is incident to α_n. In this case, there must be exactly one such flip among the last $k-j$ flips of $(T_i)_{0\le i\le k}$, say the l-th one. Moreover, this flip replaces the triangle of T_j incident to α_n by the triangle of A_n^+ incident to α_n. There is only one way to do so, depicted on the right of Fig. 17. As portrayed in the figure, this flip is incident to α_1.

To reach a contradiction, assume that at most one flip along $(T_i)_{0\leq i\leq k}$ other than the l-th one is incident to α_1. In this case, the first flip incident to α_1 along $(T_i)_{0\leq i\leq k}$, say the l'-th one, replaces the triangle of A_n^- incident to the arc α_1 by the triangle of T_{l-1} incident to this arc. There is only one way to do so, depicted in the middle of Fig. 17. One can see that the triangle of $T_{l'-1}$ incident to the boundary loop α_0 cannot be identical to the triangle of A_n^- incident to this arc. Hence one of the first $l'-1$ flips along $(T_i)_{0\leq i\leq k}$, say the j'-th one, removes the triangle of A_n^- incident to α_0.

As j' is less than l', the arc β shown on the left of Fig. 17, belongs to both $T_{j'-1}$ and $T_{j'}$. The portion of each of these triangulations bounded by the arcs α_n and β belongs to $\mathcal{MF}(\Gamma_2)$. According to Proposition 2.2, the j'-th flip along $(T_i)_{0\leq i\leq k}$ is then incident to α_n. As the j-th and l-th flips along this path are also incident to α_n, this contradicts the assumption that α_n is incident to at most two flips along $(T_i)_{0\leq i\leq k}$. Therefore α_1 must be incident to at least three flips along this geodesic. □

Lemma 4.2. *Let $n > 2$. Consider a geodesic from A_n^- to A_n^+ whose first flip incident to α_n introduces an arc with vertices a_1 and a_2. If α_n is incident to at most two flips along this geodesic, then α_1 is incident to at least four flips along it.*

Proof. Let $(T_i)_{0\leq i\leq k}$ be a geodesic from A_n^- to A_n^+ whose first flip incident to α_n, say the j-th one, introduces an arc with vertices a_1 and a_2. This flip must then be the one shown on the left of Fig. 18. Note that it is incident to α_1.

Fig. 18. The j-th (left) and l-th (right) flips performed along the path $(T_i)_{0\leq i\leq k}$ in the proof of Lemma 4.2. In each case, the arc introduced is dotted.

Assume that at most one flip along $(T_i)_{0\leq i\leq k}$ other than the j-th one is incident to α_n. In this case, there must be exactly one such flip among the last $k-j$ flips of $(T_i)_{0\leq i\leq k}$, say the l-th one. This flip replaces the triangle of T_j incident to α_n by the triangle of A_n^+ incident to α_n. There is only one way to do so, depicted on the right of Fig. 18. Note that this flip is also incident to α_1.

Finally, as the arc introduced by the j-th flip along $(T_i)_{0\leq i\leq k}$ is not removed before the l-th flip, there must be two more flips incident to the arc α_1 along this geodesic: the flip that removes the loop arc with vertex a_1 shown on the left of Fig. 18, and the flip that introduces the loop arc with vertex a_2 shown on the right of the figure. This proves that at least four flips are incident to α_1 along $(T_i)_{0\leq i\leq k}$. □

The following lemma provides the existence of a particular ear in a triangulation along some geodesics between A_n^- and A_n^+. This will result in a lower bound on the distance of A_n^- and A_n^+ via Lemma 4.4 below. The existence of such ears will also be instrumental in Section 5 when proving lower bounds on $\mathrm{diam}(\mathcal{MF}(\Pi_n))$.

Lemma 4.3. *For $n \geq 4$, consider a geodesic from A_n^- to A_n^+ whose first flip incident to α_n introduces an arc with vertices a_1 and a_p, where $2 < p < n$. Then some triangulation along this geodesic has an ear at a_q, where $2 \leq q \leq n$ and $q \neq \lfloor n/2 \rfloor + 1$.*

Proof. Let $(T_i)_{0 \leq i \leq k}$ be a geodesic from A_n^- to A_n^+ whose first flip incident to α_n, say the j-th one, introduces an arc with vertices a_1 and a_p, where $2 < p < n$. This flip is depicted in Fig. 19, separately when $p \leq \lceil n/2 \rceil$ and when $p > \lceil n/2 \rceil$.

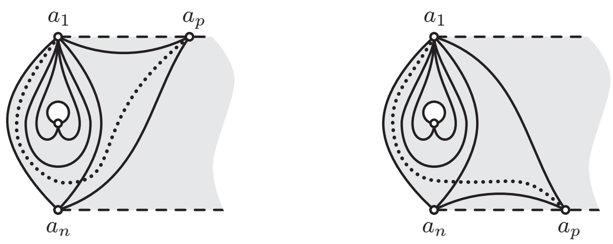

Fig. 19. The j-th flip performed along the path $(T_i)_{0 \leq i \leq k}$ in the proof of Lemma 4.3 when $p \leq \lceil n/2 \rceil$ (left) and when $p > \lceil n/2 \rceil$ (right). The arc introduced is dotted.

First assume that p is not greater than $\lceil n/2 \rceil$. Consider the arc of T_j with vertices a_1 and a_p shown as a solid line on the left of Fig. 19. The portion of T_j bounded by this arc and by the arcs $\alpha_1, \ldots, \alpha_{p-1}$ is a triangulation of the disk Δ_p. If $p > 3$, then this triangulation has at least two ears, and one of them is also an ear of T_j at a_q where $2 \leq q < p$. If $p = 3$ this property still necessarily holds with $q = 2$ since the triangulation of Δ_p induced by T_j is made up of a single triangle.

Now assume that $\lceil n/2 \rceil < p < n$. Consider the arc with vertices a_1 and a_p introduced by the j-th flip along $(T_i)_{0 \leq i \leq k}$ and shown as a dotted line on the right of Fig. 19. The portion of T_j bounded by this arc and by the arcs $\alpha_p, \ldots, \alpha_n$ is a triangulation of Δ_{n-p+2}. Since $n - p + 2 \geq 3$, an argument similar to the one used in the last paragraph shows that T_j has an ear at some vertex a_q where $p < q \leq n$.

Hence, T_j has an ear at a_q where either $2 \leq q < p \leq \lceil n/2 \rceil$ or $\lceil n/2 \rceil < p < q \leq n$. In particular, q is necessarily distinct from $\lfloor n/2 \rfloor + 1$. □

Lemma 4.3 can be combined with the following lemma to obtain inequality (4.2). Note that arguments close to the ones used in the proof of the next lemma will serve to prove Theorem 5.3 in Section 5.

Lemma 4.4. *Let $n \geq 4$. If some triangulation along a geodesic between A_n^- and A_n^+ has an ear at a_q, where $2 \leq q \leq n$ and $q \neq \lfloor n/2 \rfloor + 1$, then*

$$d(A_n^-, A_n^+) \geq d(A_{n-2}^-, A_{n-2}^+) + 5.$$

Proof. Consider a geodesic $(T_i)_{0\leq i\leq k}$ between A_n^- and A_n^+ and assume that, for some j in $\{0,\ldots,k\}$, T_j has an ear at a_q, where $2\leq q\leq n$ and $q\neq\lfloor n/2\rfloor+1$.

Set $r=n-q+1$. The portion of either A_n^- or A_n^+ placed between the arcs α_{q-1}, α_q, and α_r is shown on the left of Fig. 20. As T_j has an ear at a_q, one can split the geodesic $(T_i)_{0\leq i\leq k}$ at the triangulation T_j and invoke Lemma 2.5 for each of the resulting portions. Doing so, we find that either α_{q-1} and α_q are both incident to exactly three flips along this geodesic, or one of these arcs is incident to at least four flips along it.

These two cases will be reviewed separately.

First assume that α_s is incident to at least four flips along $(T_i)_{0\leq i\leq k}$, where s is either $q-1$ or q. In this case, Theorem 2.4 yields

$$d(A_n^-,A_n^+)\geq d(A_n^-\backslash\backslash s,A_n^+\backslash\backslash s)+4. \tag{4.5}$$

Observe that the arc α_r is not incident to the same triangle in $A_n^-\backslash\backslash s$ and in $A_n^+\backslash\backslash s$ (because of the placement of the boundary loop in the two triangulations). Hence, some flip must be incident to this arc along any geodesic between $A_n^-\backslash\backslash s$ and $A_n^+\backslash\backslash s$.

Invoking Theorem 2.4 again, we find

$$d(A_n^-\backslash\backslash s,A_n^+\backslash\backslash s)\geq d(A_n^-\backslash\backslash s\backslash\backslash r,A_n^+\backslash\backslash s\backslash\backslash r)+1. \tag{4.6}$$

As $A_n^-\backslash\backslash s\backslash\backslash r$ and $A_n^+\backslash\backslash s\backslash\backslash r$ are isomorphic to A_{n-2}^- and A_{n-2}^+ by the same vertex relabeling, the desired result is obtained by combining (4.5) and (4.6).

Now assume that α_{q-1} and α_q are both incident to exactly three flips along $(T_i)_{0\leq i\leq k}$. Note that at least one of the first j flips and at least one of the last $k-j$ flips along $(T_i)_{0\leq i\leq k}$ are incident to either α_{q-1} or α_q (because A_n^- and A_n^+ do not have an ear at a_q, while T_j does). We can assume without loss of generality that exactly one of the first j flips and two of the last $k-j$ flips along $(T_i)_{0\leq i\leq k}$ are incident to α_q by, if needed, reversing the geodesic $(T_i)_{0\leq i\leq k}$ (this is possible thanks to the symmetry between A_n^- and A_n^+). In this case, according to Lemma 2.5, exactly two of the first j flips and exactly one of the last $k-j$ flips along this geodesic are incident to α_{q-1}.

It can further be assumed without loss of generality that the j-th flip along $(T_i)_{0\leq i\leq k}$ is the one that introduces the ear at a_q. This flip is therefore incident to α_q, and since there is only one such flip along $(T_i)_{0\leq i\leq j}$, it must be as shown on the right of Fig. 20. Note that it is also incident to α_{q-1}. Now consider the triangle incident to α_{q-1} when this flip is performed, labeled t in the figure. This triangle must be introduced by the first flip incident to α_{q-1}, earlier along the geodesic. Say this flip is the l-th one along the geodesic. It must be as shown in the middle of Fig. 20.

Consider a geodesic $(T_i')_{0\leq i\leq k'}$ from $A_n^-\backslash\backslash q$ to $T_l\backslash\backslash q$, and a geodesic $(T_i'')_{0\leq i\leq k''}$ from $T_l\backslash\backslash q$ to $A_n^+\backslash\backslash q$. Since three flips are incident to α_q along $(T_i)_{0\leq i\leq k}$, splitting $(T_i)_{0\leq i\leq k}$ at the triangulation T_j and invoking Theorem 2.4 for each of the resulting portions yields

$$k'+k''\leq d(A_n^-,A_n^+)-3. \tag{4.7}$$

Observe that the triangles incident to α_r in $A_n^-\backslash\backslash q$ and in $T_l\backslash\backslash q$ are distinct (see Fig. 20). Hence, at least one flip is incident to α_r along $(T_i')_{0\leq i\leq k'}$, and by Theorem 2.4,

$$k'\geq d(A_n^-\backslash\backslash q\backslash\backslash r,T_l\backslash\backslash q\backslash\backslash r)+1. \tag{4.8}$$

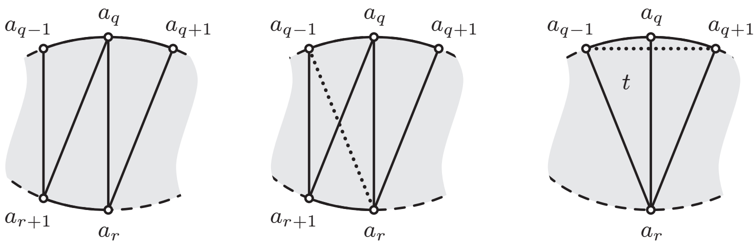

Fig. 20. The portion of either T^- or T^+ next to the vertex a_q (left), the l-th flip along the geodesic used in the proof of Lemma 4.4 (middle), and the j-th flip along this geodesic (right). The arc introduced by each flip is dotted.

Similarly, the triangles incident to α_r in $T_l \backslash\backslash q$ and in $A_n^+ \backslash\backslash q$ are distinct. Hence, at least one flip is incident to α_r along $(T_i'')_{0 \le i \le k''}$, and by Theorem 2.4,

$$k'' \geq d(T_l \backslash\backslash q \backslash\backslash r, A_n^+ \backslash\backslash q \backslash\backslash r) + 1. \tag{4.9}$$

By the triangle inequality, (4.8) and (4.9) yield

$$k' + k'' \geq d(A_n^- \backslash\backslash q \backslash\backslash r, A_n^+ \backslash\backslash q \backslash\backslash r) + 2. \tag{4.10}$$

Since $A_n^- \backslash\backslash q \backslash\backslash r$ and $A_n^+ \backslash\backslash q \backslash\backslash r$ are isomorphic to A_{n-2}^- and A_{n-2}^+ by the same vertex relabeling, the desired inequality is obtained by combining (4.7) and (4.10). □

We are now ready to establish the announced inequality.

Theorem 4.5. *For every $n > 2$,*

$$d(A_n^-, A_n^+) \geq \min\big(d(A_{n-1}^-, A_{n-1}^+) + 3, d(A_{n-2}^-, A_{n-2}^+) + 5\big).$$

Proof. Assume that $n \geq 3$ and consider a geodesic $(T_i)_{0 \le i \le k}$ from A_n^- to A_n^+. If at least three flips are incident to α_n along it, then Theorem 2.4 yields

$$d(A_n^-, A_n^+) \geq d(A_{n-1}^-, A_{n-1}^+) + 3.$$

Indeed, as mentioned above, $A_n^- \backslash\backslash n$ and $A_n^+ \backslash\backslash n$ are isomorphic to A_{n-1}^- and A_{n-1}^+, respectively, via the same vertex relabeling. Therefore in this case, the desired result holds. So we can assume in the remainder of the proof that at most two flips are incident to α_n along $(T_i)_{0 \le i \le k}$. Further assume that the first flip incident to α_n along this geodesic is the j-th one. We review three cases, depending on which arc is introduced by this flip.

First assume that the j-th flip introduces an arc with vertices a_0 and a_n. This flip must be the one depicted on the left of Fig. 17. Consider a geodesic $(T_i')_{0 \le i \le k'}$ from $A_n^- \backslash\backslash 1$ to $T_j \backslash\backslash 1$, and a geodesic $(T_i'')_{0 \le i \le k''}$ from $T_j \backslash\backslash 1$ to $A_n^+ \backslash\backslash 1$. According to Lemma 4.1 and Theorem 2.4, the following inequality holds:

$$k' + k'' \leq d(A_n^-, A_n^+) - 3. \tag{4.11}$$

As portrayed on the left of Fig. 17, the triangles incident to α_n in $A_n^-\backslash\backslash 1$, $T_j\backslash\backslash 1$, and $A_n^+\backslash\backslash 1$ are pairwise distinct. As a consequence, at least one flip must be incident to α_n along each of the geodesics $(T_i')_{0\leq i\leq k'}$ and $(T_i'')_{j\leq i\leq k''}$.

In this case, Theorem 2.4 yields

$$k' \geq d(A_n^-\backslash\backslash 1\backslash\backslash n, T_j\backslash\backslash 1\backslash\backslash n) + 1 \quad \text{and} \quad k'' \geq d(T_j\backslash\backslash 1\backslash\backslash n, A_n^+\backslash\backslash 1\backslash\backslash n) + 1.$$

By the triangle inequality, one obtains

$$k' + k'' \geq d(A_n^-\backslash\backslash 1\backslash\backslash n, A_n^+\backslash\backslash 1\backslash\backslash n) + 2. \tag{4.12}$$

Since $A_n^-\backslash\backslash 1\backslash\backslash n$ and $A_n^+\backslash\backslash 1\backslash\backslash n$ are isomorphic to A_{n-2}^- and A_{n-2}^+ via the same vertex relabeling, the desired result follows from inequalities (4.11) and (4.12).

Now assume that the j-th flip introduces an arc with vertices a_1 and a_2. It follows from Lemma 4.1 and Theorem 2.4 that

$$d(A_n^-, A_n^+) \geq d(A_n^-\backslash\backslash 1, A_n^+\backslash\backslash 1) + 4. \tag{4.13}$$

Observe that the arc α_{n-1} is not incident to the same triangle in $A_n^-\backslash\backslash 1$ and in $A_n^+\backslash\backslash 1$. Therefore, there must be at least one flip incident to α_{n-1} along any geodesic between these triangulations, and by Theorem 2.4,

$$d(A_n^-\backslash\backslash 1, A_n^+\backslash\backslash 1) \geq d(A_n^-\backslash\backslash 1\backslash\backslash n-1, A_n^+\backslash\backslash 1\backslash\backslash n-1) + 1. \tag{4.14}$$

As $A_n^-\backslash\backslash 1\backslash\backslash n-1$ and $A_n^+\backslash\backslash 1\backslash\backslash n-1$ are isomorphic to A_{n-2}^- and A_{n-2}^+ via the same vertex relabeling, the result is obtained by combining (4.13) and (4.14).

Finally, if the j-th flip introduces an arc with vertices a_1 and a_p, where $2 < p < n$, then $n > 3$ and the result follows from Lemmas 4.3 and 4.4. □

We can conclude the following.

Theorem 4.6. *The diameter of $\mathcal{MF}(\Gamma_n)$ is $\lfloor 5n/2\rfloor - 2$ for all $n \geq 1$.*

Proof. Since Γ_1 has a unique triangulation up to homeomorphism, $\mathcal{MF}(\Gamma_1)$ has diameter 0. Moreover, as can be seen in Fig. 5, $\mathcal{MF}(\Gamma_2)$ has diameter 3. The inequality

$$\text{diam}(\mathcal{MF}(\Gamma_n)) \geq \lfloor 5n/2\rfloor - 2$$

therefore follows by induction from Theorem 4.5. Combining this inequality with the upper bound provided by Theorem 3.4 completes the proof. □

5. Lower bounds for Π

We now turn our attention to the triangulations of Π, the filling surface of genus 0, with two marked boundary loops in addition to the privileged boundary, and no marked or unmarked point in its interior. We shall build two triangulations B_n^- and B_n^+ in $\mathcal{MF}(\Pi_n)$ whose flip distance is $3n + K_\Pi$, where K_Π does not depend on n.

First assume that $n > 2$. Recall that A_n^- has an ear at the vertex $a_{\lfloor n/2\rfloor+1}$ (see Fig. 16). One can transform A_n^- into a triangulation that belongs to $\mathcal{MF}(\Pi_n)$ by placing a boundary loop α_+ with a vertex a_+ in this ear and by re-triangulating the pierced ear around the boundary loop as shown in the top row of Fig. 21 (where a_+ is labeled with a $+$). The vertex a_0 and the arc α_0 will be relabeled a_- and α_- (and marked with a $-$ in the figure). The resulting triangulation will be called B_n^-.

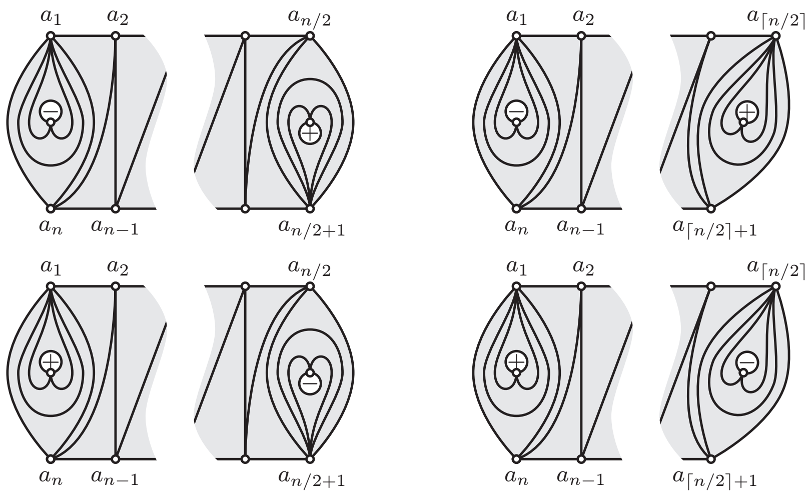

Fig. 21. The triangulations B_n^- (top row) and B_n^+ (bottom row) depicted when n is even (left) and odd (right). The vertices a_- and a_+ are labeled $-$ and $+$, respectively.

Similarly, consider the ear of A_n^+ at a_1. One can obtain a triangulation that belongs to $\mathcal{MF}(\Pi_n)$ by placing a boundary loop α_+ with a vertex a_+ in this ear and by re-triangulating the pierced ear as shown in the bottom row of Fig. 21, where a_+ is labeled with a $+$. The resulting triangulation, wherein the vertex a_0 and the boundary arc α_0 have been respectively relabeled a_- and α_-, will be called B_n^+.

When $1 \leq n \leq 2$, B_n^- and B_n^+ will be the triangulations in $\mathcal{MF}(\Pi_n)$ depicted in Fig. 22. Most of the section is devoted to proving that when $n \geq 3$,

$$d(B_n^-, B_n^+) \geq \min\big(d(B_{n-1}^-, B_{n-1}^+) + 3, d(B_{n-2}^-, B_{n-2}^+) + 6\big). \tag{5.1}$$

The proof consists in finding a geodesic between B_n^- and B_n^+ within which at least a certain number of flips (typically three) are incident to given arcs, and invoking Theorem 2.4 with well chosen vertex deletions. These deletions will be the same as in the case of the triangulations A_n^- and A_n^+. Indeed, when $n \geq 2$, the same vertex relabeling sends $B_n^- \backslash\backslash n$ and $B_n^+ \backslash\backslash n$ to B_{n-1}^- and B_{n-1}^+, respectively. Moreover, if $n \geq 3$ and if i and j are two integers such that $1 \leq i < n$ and $j \in \{n-i, n-i+1\}$, then another vertex relabeling sends $B_n^- \backslash\backslash i \backslash\backslash j$ and $B_n^+ \backslash\backslash i \backslash\backslash j$ to B_{n-2}^- and B_{n-2}^+, respectively.

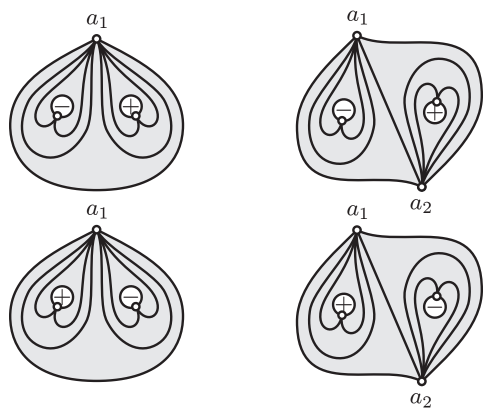

Fig. 22. The triangulations B_n^- (top row) and B_n^+ (bottom row) depicted when $n = 1$ (left) and when $n = 2$ (right). The vertices a_- and a_+ are labeled $-$ and $+$, respectively.

5.1. *When an ear is found along a geodesic*

In this subsection, we consider the geodesics between B_n^- and B_n^+ along which some triangulation has an ear. Ears at a_1 and at a_n are first reviewed separately. The following lemma deals with the case of an ear at a_1. Note that, by symmetry, this also settles the case of an ear at $a_{\lfloor n/2 \rfloor+1}$.

Lemma 5.1. *Assume that $n \geq 2$ and consider a geodesic $(T_i)_{0\leq i\leq k}$ between B_n^- and B_n^+. If there exists $j \in \{0, \ldots, k\}$ such that T_j has an ear at a_1, then*

$$d(B_n^-, B_n^+) \geq d(B_{n-1}^-, B_{n-1}^+) + 4.$$

Proof. Assume that T_j has an ear at a_1 for some integer $j \in \{0, \ldots, k\}$. Call this ear t, and let t^- be the triangle incident to α_n in B_n^-. At least two of the first j flips along $(T_i)_{0\leq i\leq k}$ must be incident to α_n. Indeed, the unique such flip would otherwise replace the triangle t^- by t. This flip would then simultaneously remove two edges of t^- (see the sketch of B_n^- on the left of Fig. 21), which is impossible. By symmetry, at least two of the last $k - j$ flips along the path $(T_i)_{0\leq i\leq k}$ must be incident to α_n. Hence, there are at least four such flips along $(T_i)_{0\leq i\leq k}$, and Theorem 2.4 yields

$$d(B_n^-, B_n^+) \geq d(B_n^- \backslash\backslash n, B_n^+ \backslash\backslash n) + 4.$$

Since an isomorphism sends $B_n^- \backslash\backslash n$ and $B_n^+ \backslash\backslash n$ to B_{n-1}^- and B_{n-1}^+ via the same vertex relabeling, the lemma is proven. □

The next lemma deals with the case of an ear at a_n. By symmetry this also settles the case of an ear at $a_{n/2}$ when n is even and at $a_{\lceil n/2 \rceil+1}$ when n is odd.

Lemma 5.2. *Assume that $n \geq 3$ and consider a geodesic $(T_i)_{0\leq i\leq k}$ between B_n^- and B_n^+. If there exists $j \in \{0, \ldots, k\}$ such that T_j has an ear at a_n, then*

$$d(B_n^-, B_n^+) \geq \min\big(d(B_{n-1}^-, B_{n-1}^+) + 3, d(B_{n-2}^-, B_{n-2}^+) + 6\big).$$

Proof. Assume that T_j has an ear at a_n for some $j \in \{0, \ldots, k\}$. One can see in Fig. 21 that the triangles of B_n^- incident to the arcs α_{n-1} and α_n do not have a common edge. Therefore, it follows from Lemma 2.5 that at least two of the first j flips along $(T_i)_{0 \leq i \leq k}$ are incident to the arc α_r for some $r \in \{n-1, n\}$. Similarly, the triangles of B_n^+ incident to the arcs α_{n-1} and α_n do not have a common edge, and according to the same lemma, at least two of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_s for some $s \in \{n-1, n\}$.

Since the triangles incident to α_n in B_n^- and in B_n^+ are distinct from the ear at a_n, at least one of the first j flips and at least one of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_n. Hence, if r or s is equal to n, then at least three flips along this geodesic are incident to α_n. In this case, the desired result follows from Theorem 2.4 because $B_n^- \backslash\backslash n$ and $B_n^+ \backslash\backslash n$ are isomorphic to B_{n-1}^- and B_{n-1}^+, respectively, via the same vertex relabeling.

Now assume that r and s are both equal to $n-1$. In this case, at least four flips along the path $(T_i)_{0 \leq i \leq k}$ are incident to α_{n-1}, and Theorem 2.4 yields

$$d(B_n^-, B_n^+) \geq d(B_n^- \backslash\backslash n-1, B_n^+ \backslash\backslash n-1) + 4. \tag{5.2}$$

Denote by t^- and t^+ the triangles incident to the arc α_1 in $B_n^- \backslash\backslash n-1$ and in $B_n^+ \backslash\backslash n-1$, respectively. One can see using Fig. 21 that these two triangles separate the two boundary loops in opposite ways. As shown in Fig. 23, a single flip cannot exchange t^- and t^+. Hence, at least two flips are incident to α_1 along any geodesic between $B_n^- \backslash\backslash n-1$ and $B_n^+ \backslash\backslash n-1$, and according to Theorem 2.4,

$$d(B_n^- \backslash\backslash n-1, B_n^+ \backslash\backslash n-1) \geq d(B_n^- \backslash\backslash n-1 \backslash\backslash 1, B_n^+ \backslash\backslash n-1 \backslash\backslash 1) + 2. \tag{5.3}$$

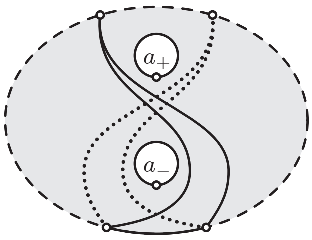

Fig. 23. No flip can replace the triangle t^- (solid lines) by the triangle t^+ (dotted lines) because such a flip would simultaneously remove two edges of t^-.

Since $B_n^- \backslash\backslash n-1 \backslash\backslash 1$ and $B_n^+ \backslash\backslash n-1 \backslash\backslash 1$ are isomorphic to B_{n-2}^- and B_{n-2}^+, respectively, via the same vertex relabeling, combining (5.2) with (5.3) completes the proof. □

When $n \geq 3$, the last two lemmas can be generalized to any ear placement as follows. Note that the proof of this theorem is similar to that of Lemma 4.4.

Theorem 5.3. *Assume that $n \geq 3$ and consider a geodesic $(T_i)_{0 \leq i \leq k}$ between B_n^- and B_n^+. If there exists $j \in \{0, \ldots, k\}$ such that T_j has an ear, then*

$$d(B_n^-, B_n^+) \geq \min\big(d(B_{n-1}^-, B_{n-1}^+) + 3, d(B_{n-2}^-, B_{n-2}^+) + 6\big).$$

Proof. Assume that T_j has an ear at a_q for some $j \in \{0, \ldots, k\}$ and $q \in \{1, \ldots, n\}$. If $q \in \{1, n\}$, then the desired result follows from Lemma 5.1 or Lemma 5.2. Similarly, if $q \in \{\lceil n/2 \rceil, \lceil n/2 \rceil + 1\}$, these two lemmas also provide the desired result because of the symmetries of B_n^- and B_n^+. For the remainder of the proof, we may thus assume that q is distinct from 1, $\lceil n/2 \rceil$, $\lceil n/2 \rceil + 1$, and n.

Denote $r = n - q + 1$. The portion of the triangulation B_n^- placed between the edges α_{q-1}, α_q, and α_r is depicted on the left of Figure 24. Note that if one splits the geodesic $(T_i)_{0 \leq i \leq k}$ at the triangulation T_j, then Lemma 2.5 can be invoked for each of the resulting portions. Doing so, we find that either α_{q-1} and α_q are both incident to exactly three flips along this geodesic, or one of these arcs is incident to at least four flips along it.

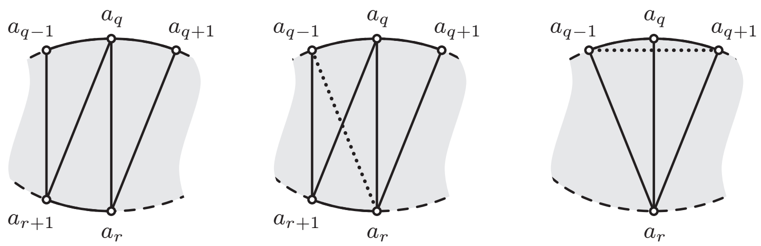

Fig. 24. The portion of the triangulation B_n^- placed between the arcs α_{q-1}, α_q, and α_r (left), the l-th flip along the geodesic used in the proof of Theorem 5.3 (middle), and the j-th flip along this geodesic (right). The arc introduced by each flip is dotted.

First assume that at least four flips are incident to α_s along $(T_i)_{0 \leq i \leq k}$, where s is $q-1$ or q. Denote by t^- and t^+ the triangles incident to α_r in $B_n^- \backslash\backslash s$ and in $B_n^+ \backslash\backslash s$, respectively. Using Fig. 21, one can see that these two triangles separate the two boundary loops in opposite ways. As shown in Fig. 23, a single flip cannot exchange t^- and t^+. Hence, at least two flips are incident to α_r along any geodesic between $B_n^- \backslash\backslash s$ and $B_n^+ \backslash\backslash s$, and invoking Theorem 2.4 twice yields

$$d(B_n^-, B_n^+) \geq d(B_n^- \backslash\backslash s \backslash\backslash r, B_n^+ \backslash\backslash s \backslash\backslash r) + 6.$$

Since $B_n^- \backslash\backslash s \backslash\backslash r$ and $B_n^+ \backslash\backslash s \backslash\backslash r$ are isomorphic to B_{n-2}^- and B_{n-2}^+, respectively, via the same vertex relabeling, the theorem is proven in this case.

Now assume that α_{q-1} and α_q are both incident to exactly three flips along $(T_i)_{0 \leq i \leq k}$. Note that at least one of the first j flips and at least one of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ must be incident to each of these arcs because B_n^- and B_n^+ do not have an ear at a_q. Thanks to the symmetry between B_n^- and B_n^+, one can assume without loss of generality that exactly one of the first j flips and two of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_q, by reversing $(T_i)_{0 \leq i \leq k}$ if needed. Then, by Lemma 2.5, exactly two of the first j flips and exactly one of the last $k - j$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_{q-1}.

Without loss of generality, we may assume that the j-th flip along $(T_i)_{0 \leq i \leq k}$ introduces the ear at a_q. This flip is then both the first flip incident to α_q and the second flip incident to α_{q-1} along the geodesic. In particular, it must replace the triangle of B_n^- incident to α_q by the ear at a_q, as shown on the right of Fig. 24. Now assume that the first flip incident to α_{q-1} along $(T_i)_{0\leq i\leq k}$ is the l-th one. Since there is no other such flip among the first $j-1$ flips along the geodesic, it must be as shown in the middle of Fig. 24.

Consider a geodesic $(T_i')_{0\leq i\leq k'}$ from $B_n^-\backslash\backslash q$ to $T_l\backslash\backslash q$, and a geodesic $(T_i'')_{0\leq i\leq k''}$ from $T_l\backslash\backslash q$ to $B_n^+\backslash\backslash q$. Since three flips are incident to α_q along $(T_i)_{0\leq i\leq k}$, Theorem 2.4 yields

$$k'+k''\leq d(B_n^-,B_n^+)-3. \tag{5.4}$$

Observe that the triangles incident to α_r in $B_n^-\backslash\backslash q$ and in $T_l\backslash\backslash q$ are distinct. Hence, at least one flip is incident to α_r along $(T_i')_{0\leq i\leq k'}$, and by Theorem 2.4,

$$k'\geq d(B_n^-\backslash\backslash q\backslash\backslash r, T_l\backslash\backslash q\backslash\backslash r)+1. \tag{5.5}$$

Now denote by t^- and t^+ the triangles incident to the arc α_r in $T_l\backslash\backslash q$ and $B_n^+\backslash\backslash q$, respectively. By construction, t^- and t^+ separate the two boundary loops in opposite ways. As shown in Fig. 23, a single flip cannot exchange t^- and t^+. Hence, at least two flips are incident to α_r along $(T_i'')_{j\leq i\leq k''}$, and Theorem 2.4 yields

$$k''\geq d(T_l\backslash\backslash q\backslash\backslash r, B_n^+\backslash\backslash q\backslash\backslash r)+2. \tag{5.6}$$

By the triangle inequality, (5.5) and (5.6) yield

$$k'+k''\geq d(B_n^-\backslash\backslash q\backslash\backslash r, B_n^+\backslash\backslash q\backslash\backslash r)+3. \tag{5.7}$$

Since $B_n^-\backslash\backslash q\backslash\backslash r$ and $B_n^+\backslash\backslash q\backslash\backslash r$ are isomorphic to B_{n-2}^- and B_{n-2}^+ by the same vertex relabeling, the desired inequality is obtained by combining (5.4) and (5.7). □

5.2. *When no ear is found along a geodesic*

We call a geodesic between B_n^- and B_n^+ *earless* if none of the triangulations along this geodesic has an ear. We will first show that under mild conditions, one always finds two particular triangulations along any such geodesics. These triangulations are sketched in Fig. 25. The triangulation shown in the top row will be called $C_n^-(p)$, where a_p is the privileged boundary vertex of the triangle of $C_n^-(p)$ incident to the boundary loop α_+. Further note that $C_n^-(p)$ is sketched separately when $p>\lceil n/2\rceil$ (left) and when $p\leq\lceil n/2\rceil$ (right). The triangulation shown in the bottom row of Fig. 25, called $C_n^+(p)$, has a similar structure, but the boundary loop α_- is placed in a different way.

Observe that the triangulations $C_n^-(p)$ and $C_n^+(p)$ do not have an ear. In fact, if at most two flips are incident to either α_n and $\alpha_{\lceil n/2\rceil}$ along an earless geodesic between B_n^- and B_n^+, then these two triangulations are necessarily both found along this geodesic for appropriate values of p. In order to prove this, the following lemma is needed:

Lemma 5.4. *Let $n>2$. If at most two flips are incident to α_n along an earless geodesic from B_n^- to B_n^+, then the first flip incident to α_n along this geodesic either introduces an arc with vertices a_- and a_n, or an arc with vertices a_1 and a_+.*

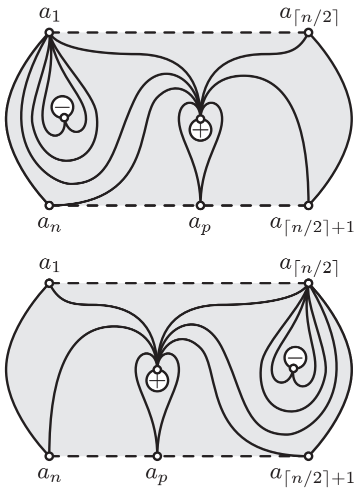

Fig. 25. A sketch of $C_n^-(p)$ (top) and $C_n^+(p)$ (bottom) when $p > \lceil n/2 \rceil$ (left) and when $p \leq \lceil n/2 \rceil$ (right). Not all the interior edges of these triangulations are shown. The omitted edges connect privileged boundary vertices to a_+.

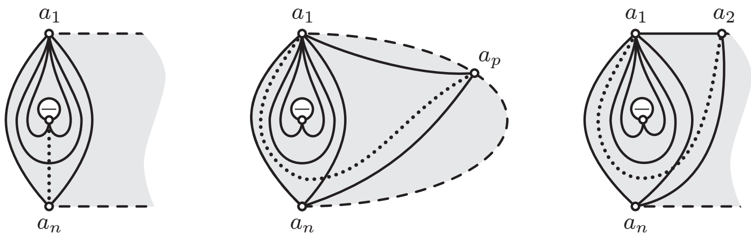

Fig. 26. The j-th flip along the geodesic $(T_i)_{0 \leq i \leq k}$ used in the proof of Lemma 5.4. The arc introduced by this flip (dotted) has vertices a_- and a_n (left), or vertices a_1 and a_p with $2 \leq p < n$ (middle and right).

Proof. Consider a geodesic $(T_i)_{0 \leq i \leq k}$ from B_n^- to B_n^+ and assume that at most two flips are incident to α_n along it. Further assume that the first flip incident to α_n along this geodesic is the j-th one. If this flip removes the loop edge of B_n^- at the vertex a_1, then it introduces the arc with vertices a_- shown on the left of Fig. 26, and the desired result holds. It is therefore assumed in the remainder of the proof that this flip removes the interior arc of B_n^- with vertices a_1 and a_n. The introduced arc is incident to a_1 and its other vertex is either a_+ or a_p where $1 < p < n$. We will use an indirect argument. Assume that the introduced arc is incident to a_p where $1 < p < n$. One can see in the middle of Fig. 26 that, in this case, T_j induces a triangulation U in the portion Σ of Π_n bounded by the dotted arc and by the arcs $\alpha_p, \dots, \alpha_n$. This triangulation cannot be a triangulation of a disk. Indeed, otherwise, one of the ears of U would be an ear of T_j. This shows that the boundary loop with vertex a_+ must be a boundary of Σ. In this case, the j-th flip along $(T_i)_{0 \leq i \leq k}$ must be the one shown on the right of Fig. 26. Indeed, T_j

would otherwise induce a triangulation of a disk in the portion Π_n bounded by the arcs $\alpha_1, \ldots, \alpha_{p-1}$ and by the interior arc with vertices a_1 and a_p shown in the middle of the figure as a solid line. This triangulation would then share one of its ears with T_j.

Now, let t^- and t^+ be the triangles incident to α_n in T_j and B_n^+, respectively. As the j-th flip along $(T_i)_{0\leq i\leq k}$ is the one shown on the right of Fig. 26, t^- and t^+ separate the two boundary loops in opposite ways. As shown in Fig. 23, a single flip cannot exchange these triangles. Hence, at least two of the last $k - j$ flips must be incident to α_n along $(T_i)_{0\leq i\leq k}$, and at least three such flips are found along this geodesic, a contradiction. □

Lemma 5.5. *Let $n > 2$. If both α_n and $\alpha_{\lceil n/2\rceil}$ are incident to at most two flips along an earless geodesic from B_n^- to B_n^+, then there exist a geodesic $(T_i)_{0\leq i\leq k}$ and integers p^-, p^+, j^-, and j^+ such that $j^- < j^+$, and the triangulations T_{j^-} and T_{j^+} are equal to $C_n^-(p^-)$ and $C_n^+(p^+)$, respectively.*

Proof. Assume that α_n and $\alpha_{\lceil n/2\rceil}$ are each incident to at most two flips along an earless geodesic $(T_i)_{0\leq i\leq k}$ from B_n^- to B_n^+. In this case, these arcs are each incident to exactly two flips along this geodesic. Indeed, otherwise the unique such flip would remove two arcs simultaneously, as shown in Fig. 23. Assume that the first flip incident to α_n along $(T_i)_{0\leq i\leq k}$ is the j^--th one. Denote by t^- the triangle of T_{j^-} incident to α_n. From now on, t^- remains incident to α_n in the triangulations along the geodesic until the second flip incident to α_n removes it. Moreover, according to Lemma 5.4, the vertices of t^- are a_1, a_n, and either a_- or a_+. Thanks to the symmetries of B_n^- and B_n^+, one can assume that this vertex is a_+. Indeed, if a_- is a vertex of t^-, then exchanging the labels of a_- and a_+ and reversing the direction of the geodesic $(T_i)_{0\leq i\leq k}$ results in a geodesic from B_n^- to B_n^+ whose first flip incident to α_n introduces an arc with vertices a_1 and a_+.

In particular, T_{j^-} must contain all the arcs of $C_n^-(p)$ shown as solid lines at the top of Fig. 25, except possibly for the edges of the triangles incident to $\alpha_{\lceil n/2\rceil}$ and α_+. However, since T_{j^-} does not have an ear, all its other interior arcs must connect the privileged boundary vertices to a_+. In particular T_{j^-} is necessarily equal to $C_n^-(p^-)$, where a_{p^-} is the privileged boundary vertex of the triangle incident to α_+ in T_{j^-}.

Now consider the triangle t^+ incident to $\alpha_{\lceil n/2\rceil}$ in T_{j^-}, i.e., in $C_n^-(p^-)$. The vertices of t^+ are $a_{\lceil n/2\rceil}$, $a_{\lceil n/2\rceil+1}$, and a_+, as shown at the top of Fig. 25. This triangle must be introduced by the first flip incident to $\alpha_{\lceil n/2\rceil}$ along $(T_i)_{0\leq i\leq k}$, and removed by the second flip incident to $\alpha_{\lceil n/2\rceil}$ along this geodesic. Say j^+ is the index such that the latter flip transforms T_{j^+} into T_{j^++1}. It turns out that t^- must still be a triangle of T_{j^+}. Indeed, otherwise the triangle of B_n^+ incident to α_n would already be a triangle of T_{j^+}, which is impossible because it intersects the interior of t^+. By this argument, t^- and t^+ are not affected by a flip between T_j^- and T_j^+. Therefore, along this portion of the geodesic, α_- must remain in the subsurface Σ of Π_n bounded by the arcs α_1 to $\alpha_{\lceil n/2\rceil-1}$ and by the edges of t^- and t^+ not incident to a_n and to $a_{\lceil n/2\rceil+1}$, respectively.

Now recall that the flip that transforms T_{j^+} into T_{j^++1} replaces t^+ by the triangle incident to $\alpha_{\lceil n/2\rceil}$ in B_n^+. It follows that the triangle of T_{j^+} incident to the boundary loop α_- must already be the same as in B_n^+, and its privileged boundary vertex is either $a_{\lceil n/2\rceil}$ or $a_{\lceil n/2\rceil+1}$ depending on the parity of n. By the argument in the last paragraph, this triangle is contained in Σ, and therefore must be incident to $a_{\lceil n/2\rceil}$ because Σ does not contain $a_{\lceil n/2\rceil+1}$. In particular, T_{j^+} necessarily contains all the arcs of $C_n^+(p)$ shown as solid lines at the bottom of Fig. 25, except possibly for the edges of the triangles incident to α_n and α_+. However, T_{j^+} does not have an ear, and as a consequence it coincides with $C_n^+(p^+)$, where a_{p^+} is the privileged boundary vertex of the triangle incident to α_+. □

Lemma 5.6. *Let $n > 2$. Consider an integer p such that $2 \leq p \leq n$. If at most one flip is incident to α_1 along some geodesic between $C_n^+(p)$ and B_n^+, then at least two flips are incident to α_n along this geodesic.*

Proof. Consider a geodesic $(T_i)_{0\leq i\leq k}$ from B_n^+ to $C_n^+(p)$ and assume that at most one flip along this geodesic is incident to the arc α_1. In this case, there is exactly one such flip, say the j-th one. Denote by β the interior arc of B_n^+ with vertices a_1 and a_n. This arc belongs to $T_0, \ldots, T_{j-1}$ and it is removed by the flip that transforms T_{j-1} into T_j. More precisely, this flip replaces β by an arc with vertices a_2 and a_+. There are exactly two ways to do so, shown in the middle and on the right of Fig. 27.

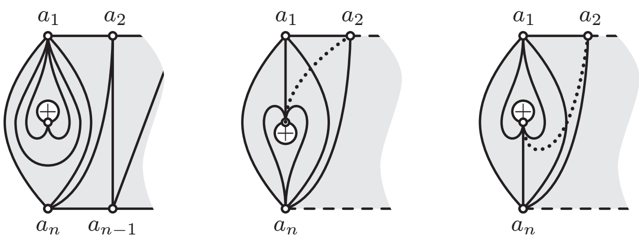

Fig. 27. A sketch of B_n^+ (left) and the two possibilities for the j-th flip along the geodesic $(T_i)_{0\leq i\leq k}$ in the proof of Lemma 5.6 (middle and right), where the arc introduced is dotted.

If the j-th flip along $(T_i)_{0\leq i\leq k}$ is the one shown in the middle of Fig. 27, then at least two flips must have been performed within the portion Σ of Π_n bounded by β and α_n earlier along the path (see $\mathcal{MF}(\Gamma_2)$ in Fig. 5). By Proposition 2.2, these two flips are incident to α_n and the desired result holds.

If the j-th flip along $(T_i)_{0\leq i\leq k}$ is the one shown on the right of Fig. 27, then at least one of the earlier flips along the path modifies the triangulation within Σ. By Proposition 2.2, this flip is incident to α_n. One can see on the right of Fig. 27 that the triangles incident to α_1 and α_n in T_j have no common edge. However, since p is not equal to 1, the triangles incident to these arcs in $C_n^+(p)$ share an edge. Hence, at least one of the last $k - j$ flips along $(T_i)_{0\leq i\leq k}$ must be incident to α_n, thereby proving that there are at least two such flips along the geodesic. □

Lemma 5.7. *Let $n > 2$. If no flip is incident to α_n along a geodesic between $C_n^-(\lceil n/2\rceil)$ and $C_n^+(1)$, then at least two of its flips are incident to α_1.*

Proof. Assume that no flip is incident to α_n along a geodesic between $C_n^-(\lceil n/2\rceil)$ and $C_n^+(1)$. The triangles incident to α_1 in $C_n^-(\lceil n/2\rceil)$ and $C_n^+(1)$ are depicted in Fig. 28,

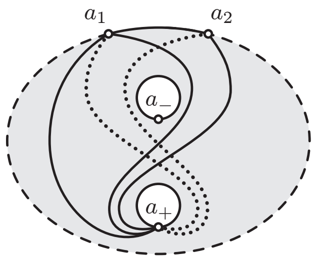

Fig. 28. The triangles incident to α_1 in $C_n^-(\lceil n/2 \rceil)$ (solid lines) and in $C_n^+(1)$ (dotted lines), and an edge of the triangle incident to α_n in these triangulations.

in solid lines and in dotted lines, respectively. In this figure, the leftmost arc with vertices a_1 and a_+ is an edge of the triangle incident to α_1 in both $C_n^-(\lceil n/2 \rceil)$ and $C_n^+(1)$.

By hypothesis, this arc is never removed along our geodesic. Therefore, if there is exactly one flip incident to α_1 along this geodesic, it must remove two edges of the triangle of $C_n^-(\lceil n/2 \rceil)$ incident to α_1, as can be seen in Fig. 28. As a consequence, there are at least two flips incident to α_1 along the geodesic. □

5.3. A lower bound on the diameter of $\mathcal{MF}(\Pi_n)$

Theorem 5.8. *For any $n > 2$,*

$$d(B_n^-, B_n^+) \geq \min\bigl(d(B_{n-1}^-, B_{n-1}^+) + 3, d(B_{n-2}^-, B_{n-2}^+) + 6\bigr).$$

Proof. Assume that $n > 2$. If one of the triangulations along any geodesic between B_n^- and B_n^+ has an ear, then the desired result follows from Theorem 5.3. We may thus assume, for the remainder of the proof, that all the triangulations along the geodesic between B_n^- and B_n^+ are earless. Moreover, if $p \in \{n, \lceil n/2 \rceil\}$ and if at least three flips are incident to α_p along some geodesic between B_n^- and B_n^+, the result follows from Theorem 2.4 because $B_n^- \backslash\backslash p$ and $B_n^+ \backslash\backslash p$ are isomorphic to B_{n-1}^- and B_{n-1}^+, respectively, via the same vertex relabeling. Hence, it will also be assumed that α_n and $\alpha_{\lceil n/2 \rceil}$ are incident to at most two flips along any geodesic between B_n^- and B_n^+. Under these assumptions, Lemma 5.5 provides an earless geodesic $(T_i)_{0 \leq i \leq k}$ from B_n^- to B_n^+ and four integers p^-, p^+, j^-, and j^+ such that $j^- < j^+$, and T_{j^-} and T_{j^+} are equal to $C_n^-(p^-)$ and $C_n^+(p^+)$, respectively.

First assume that $p^+ > 1$. Observe that the triangle incident to α_n in $C_n^+(p^+)$ is distinct from the triangles incident to this arc in B_n^- and in B_n^+. As no more than two flips are incident to α_n along $(T_i)_{0 \leq i \leq k}$, exactly one of the first j^+ flips and exactly one of the last $k - j^+$ flips along this geodesic are incident to α_n. In this case, Lemma 5.6 states that at least two of the last $k - j^+$ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_1. Now observe that the triangle incident to α_1 in $C_n^-(p^-)$ is distinct from the triangles incident to this arc in B_n^- and in $C_n^+(p^+)$. Hence at least two of the first j^+ flips along $(T_i)_{0 \leq i \leq k}$ are incident to α_1, which proves that at least four such flips are found along this geodesic.

Theorem 2.4 then yields

$$d(B_n^-, B_n^+) \geq d(B_n^- \backslash\backslash 1, B_n^+ \backslash\backslash 1) + 4. \quad (5.8)$$

Thanks to the symmetry between B_n^- and B_n^+, the arguments in the last paragraph also prove (5.8) when p^- is distinct from $\lceil n/2 \rceil$. Now assume that $p^- = \lceil n/2 \rceil$ and $p^+ = 1$. We will show that (5.8) still holds in this case. According to Lemma 5.7, at least two flips are incident to α_1 in the portion of $(T_i)_{0 \leq i \leq k}$ between $C_n^-(p^-)$ and $C_n^+(p^+)$. Now observe that the triangles of B_n^- and $C_n^-(\lceil n/2 \rceil)$ incident to α_1 are distinct. Hence at least three of the first p^+ flips along $(T_i)_{0 \leq i \leq k}$ are incident to this arc. In addition, the triangles of $C_n^+(1)$ and B_n^+ incident to α_1 are distinct. Therefore, at least one of the last $k - p^+$ flips along $(T_i)_{0 \leq i \leq k}$ is incident to α_1, which proves that there are at least four such flips along this geodesic, and inequality (5.8) still holds in this case.

Finally, observe that there must be at least two flips incident to α_{n-1} along any geodesic between $B_n^- \backslash\backslash 1$ and $B_n^+ \backslash\backslash 1$. Indeed, the triangles incident to α_{n-1} in these triangulations separate the two boundary loops in opposite ways and, as can be seen in Fig. 23, a single flip cannot exchange them. Hence, at least two flips are incident to α_{n-1} along any geodesic between $B_n^- \backslash\backslash 1$ and $B_n^+ \backslash\backslash 1$, and Theorem 2.4 yields

$$d(B_n^- \backslash\backslash 1, B_n^+ \backslash\backslash 1) \geq d(B_n^- \backslash\backslash 1 \backslash\backslash n-1, B_n^+ \backslash\backslash 1 \backslash\backslash n-1) + 2. \tag{5.9}$$

Since $B_n^- \backslash\backslash 1 \backslash\backslash n-1$ and $B_n^+ \backslash\backslash 1 \backslash\backslash n-1$ are isomorphic to B_{n-2}^- and to B_{n-2}^+ by the same vertex deletion, the result is obtained by combining (5.8) and (5.9). □

We are now able to bound the diameter of $\mathcal{MF}(\Pi_n)$ as follows.

Theorem 5.9. *The diameter of $\mathcal{MF}(\Pi_n)$ is not less than $3n$.*

Proof. One can see in Fig. 22 that at least three of the interior arcs of A_1^- have to be removed in order to transform it into A_1^+. For instance, either all the arcs incident to a_-, or all the arcs incident to a_+ have to be removed. Hence

$$d(B_1^-, B_1^+) \geq 3. \tag{5.10}$$

One can see in the same figure that transforming A_2^- into A_2^+ requires removing the arcs incident to a_- and the arcs incident to a_+. As there are six such arcs,

$$d(B_2^-, B_2^+) \geq 6. \tag{5.11}$$

The lower bound of $3n$ on the diameter of $\mathcal{MF}(\Pi_n)$ therefore follows by induction from Theorem 5.8 and from inequalities (5.10) and (5.11). □

Observe that B_1^- and B_1^+ are exactly three flips distant from each other (flipping all the arcs incident to a_- provides a geodesic). The triangulations B_2^- and B_2^+, however, are at least seven flips apart because all the interior arcs of B_2^- have to be removed in order to transform it into B_2^+. In particular, the bound provided by Theorem 5.9 is not sharp.

Finally, consider the triangulations shown in Fig. 29. In order to transform the triangulation on the left into the other one, the three interior arcs incident to a_1 must be removed as well as the interior arc twice incident to a_- and at least one of the arcs with vertices a_- and a_+. Hence, these triangulations are at least five flips apart. In particular, even already when $n = 1$, the triangulations B_n^- and B_n^+ are not maximally distant.

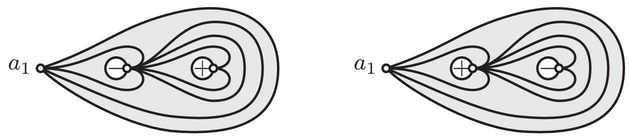

Fig. 29. Two triangulations in $\mathcal{MF}(\Pi_1)$ at least five flips apart. The vertices a_- and a_+ are labeled $-$ and $+$, respectively.

6. Consequences and further questions

As a first consequence of the above theorems, we prove the following.

Theorem 6.1. *Let Σ be a filling surface. If $\Gamma \subset \Sigma$ is an essential embedding, then*

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathcal{MF}(\Sigma_n))}{n} \geq \frac{5}{2}.$$

Proof. If Γ is embedded in Σ, there exists a surface Σ' (possibly empty if Γ is equal to Σ) such that gluing Σ' and Γ results in Σ.

Now we take two diametrically opposite triangulations U and V in $\mathcal{MF}(\Gamma_n)$ and send them to triangulations in $\mathcal{MF}(\Sigma_n)$ by gluing a fixed triangulation of Σ' to U and to V. Denote by U' and V' the resulting triangulations of $\mathcal{MF}(\Sigma_n)$. We claim that

$$d(U', V') = d(U, V).$$

That the distance of U' and V' is at most that of U and V is obvious, as any path in $\mathcal{MF}(\Gamma_n)$ can easily be emulated in $\mathcal{MF}(\Sigma_n)$. To see that $d(U', V') \geq d(U, V)$ we will use Lemma 2.6. By the lemma, if two triangulations in $\mathcal{F}(\Sigma_n)$ have an arc or a set of arcs in common, then any geodesic between them conserves these arcs. Now, of course this property may no longer be true when one quotients by the group of homeomorphisms, but it turns out that it works in this particular case. Indeed, as we consider homeomorphisms that preserve marked points, the isotopy class of a curve parallel to the privileged boundary curve is preserved by any such homeomorphism. This implies that the isotopy class of the embedding of the boundary loop of Γ_n is also preserved. Thus there exists a geodesic between U' and V' such that all triangulations contain this arc. Along this geodesic, any flip in Σ' would be superfluous. As a consequence, it lies entirely in this natural copy of $\mathcal{MF}(\Gamma_n)$, and we are done. □

This theorem implies that the diameter growth rate for all filling surfaces is at least of the order of $5n/2$ except for the disk, the once-punctured disk, and possibly for the filling surfaces of positive genus without interior vertices or non-privileged boundaries. As shown in [13], however, the growth rate is also at least of the order of $5n/2$ in the latter case, and we are left with only the disk and the once-punctured disk (whose diameter of modular flip-graphs grows like $2n$ [14, 16]).

In fact, there are multiple variations and consequences either of the above results or of the methods we use to prove them. For example, one could try to emulate these methods for filling surfaces more complicated than Π, but the combinatorics become more and more difficult to handle. There is reason to believe that increasing the number of marked boundary loops might increase the diameter of the underlying flip-graph. In the case of unmarked boundary loops, we can also expect some form of monotonicity with respect to the number of boundary loops. In fact, we suspect that the following is true.

Conjecture 6.2. *For any $\varepsilon > 0$ there exists a k_ε such that if Σ is a surface with k_ε marked boundary loops, the diameter of its flip-graphs satisfies*

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathcal{MF}(\Sigma_n))}{n} \geq 4 - \varepsilon.$$

In the unmarked case, we conjecture the following.

Conjecture 6.3. *For any $\varepsilon > 0$ there exists a k_ε such that if Σ is a surface with k_ε unmarked boundary loops, the diameter of its flip-graphs satisfies*

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathcal{MF}(\Sigma_n))}{n} \geq 3 - \varepsilon.$$

There are many other questions that we feel could be interesting. A very basic one is to understand the growth of the diameter of the flip-graph when Σ is a torus (with a privileged boundary curve). This problem is studied in [13] with the same methods, but in their current state, these methods are not able to provide sharp estimates.

Other more complicated variations of the above problems include considering surfaces with multiple privileged boundary components and adding points to several of them, whose number is not fixed. We suspect that one could be able to find very different diameter growths by sufficiently varying the problem.

To conclude we now have examples of filling surfaces with $2n$, $\frac{5}{2}n$ and $3n$ growth rates. This begs the question of classifying which numbers can appear as growth rates of these diameters. We suspect that the growth rates continue to change when the topology changes. More precisely we conjecture the following.

Conjecture 6.4. *The number of topological types of filling surfaces whose diameter of flip-graphs has a given growth rate is finite.*

Acknowledgments. H. Parlier was partially supported by Swiss National Science Foundation grant number PP00P2_153024 and acknowledges support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 "RNMS: GEometric structures And Representation varieties" (the GEAR Network). L. Pournin was partially supported by the ANR project IComb (grant ANR-08-JCJC-0011) and by Ville de Paris Émergences project "Combinatoire à Paris".

References

[1] Bridson, M. R., Haefliger, A.: Metric Spaces of Non-Positive Curvature. Grundlehren Math. Wiss. 319, Springer (1999) Zbl 0988.53001 MR 1744486

[2] Brooks, R., Makover, E.: Random construction of Riemann surfaces. J. Differential Geom. **68**, 121–157 (2004) Zbl 1095.30037 MR 2152911

[3] De Loera, J. A., Rambau, J., Santos, F.: Triangulations: Structures for Algorithms and Applications. Algorithms Comput. Math. 25, Springer (2010) Zbl 1207.52002 MR 2743368

[4] Disarlo, V., Parlier, H.: The geometry of flip graphs and mapping class groups. arXiv:1411.4285 (2014)

[5] Edelman, P. H., Reiner, V.: Catalan triangulations of the Möbius band. Graphs Combin. **13**, 231–243 (1997) Zbl 0890.05034 MR 1469823

[6] Fomin, S., Shapiro, M., Thurston, D.: Cluster algebras and triangulated surfaces. Part I: Cluster complexes. Acta Math. **201**, 83–146 (2008) Zbl 1263.13023 MR 2448067

[7] Fomin, S., Thurston, D.: Cluster algebras and triangulated surfaces. Part II: Lambda lengths. arXiv:1210.5569 (2012)

[8] Fomin, S., Zelevinsky, A.: Y-systems and generalized associahedra. Ann. of Math. **158**, 977–1018 (2003) Zbl 1057.52003 MR 2031858

[9] Gel'fand, I. M., Kapranov, M. M., Zelevinsky, A. V.: Discriminants of polynomials of several variables and triangulations of Newton polyhedra. Leningrad Math. J. **2**, 449–505 (1990) Zbl 0741.14033 MR 1073208

[10] Korkmaz, M., Papadopoulos, A.: On the ideal triangulation graph of a punctured surface. Ann. Inst. Fourier (Grenoble) **62**, 1367–1382 (2012) Zbl 1256.32015 MR 3025746

[11] Lee, C. W.: The associahedron and triangulations of the n-gon. Eur. J. Combin. **10**, 551–560 (1989) Zbl 0682.52004 MR 1022776

[12] Mosher, L.: Tiling the projective foliation space of a punctured surface. Trans. Amer. Math. Soc. **306**, 1–70 (1988) Zbl 0647.57005 MR 0927683

[13] Parlier, H., Pournin, L.: Modular flip-graphs of one holed surfaces. arXiv:1510.07664 (2015)

[14] Parlier, H., Pournin, L.: Once punctured disks, non-convex polygons, and pointihedra. Ann. Combin. (to appear)

[15] Penner, R.: The decorated Teichmüller space of punctured surfaces. Comm. Math. Phys. **113**, 299–339 (1987) Zbl 0642.32012 MR 0919235

[16] Pournin, L.: The diameter of associahedra. Adv. Math. **259**, 13–42 (2014) Zbl 1292.52011 MR 3197650

[17] Sleator, D., Tarjan, R., Thurston, W.: Rotation distance, triangulations, and hyperbolic geometry. J. Amer. Math. Soc. **1**, 647–681 (1988) Zbl 0653.51017 MR 0928904

[18] Stasheff, J.: Homotopy associativity of H-spaces. Trans. Amer. Math. Soc. **108**, 275–312 (1963) Zbl 0114.39402 MR 0158400

[19] Stasheff, J.: How I 'met' Dov Tamari. In: Associahedra, Tamari Lattices and Related Structures, Progr. Math. 299, Birkhäuser, 45–63 (2012) Zbl 1269.52014 MR 3221533

[20] Tamari, D.: Monoïdes préordonnés et chaînes de Malcev. Bull. Soc. Math. France **82**, 53–96 (1954) Zbl 0055.01501 MR 0063363