

Connectivity results for surface branched ideal triangulations

RICCARDO BENEDETTI

Dedicato a Bruno Zimmermann per il suo 70esimo compleanno

ABSTRACT. *We consider triangulations of closed surfaces S with a given set of vertices V ; every triangulation can be branched that is enhanced to be a Δ -complex. Branched triangulations are considered up to the b -transit equivalence generated by b -flips (i.e. branched diagonal exchanges) and isotopy keeping V pointwise fixed. We extend a well-known connectivity result for ‘naked’ triangulations; in particular, in the generic case when $\chi(S) < 0$, we show that each branched triangulation is connected to any other if $\chi(S)$ is even, while this holds also for odd $\chi(S)$ possibly after the complete inversion of one of the two branchings. Natural distribution of the b -flips in sub-families gives rise to restricted transit equivalences with nontrivial (even infinite) quotient sets. We analyze them in terms of certain structures of geometric/topological nature carried by each branched triangulation, invariant for the given restricted equivalence.*

Keywords: triangulation of surfaces, branching, branched flips.
MS Classification 2010: 57Q15, 57N05, 57M50, 57M27.

1. Introduction

Let (S, V) be a compact closed connected smooth surface S with a set V of n marked points, $n \geq 1$, and Euler-Poincaré characteristic $\chi(S)$, such that $\chi(S) - n < 0$. It is well-known that (S, V) carries *ideal triangulations*, say T . This means that T is a possibly *loose* triangulation (self and multiple edge adjacency being allowed) of S whose set of vertices coincides with V . The ideal triangulations of (S, V) share the same numbers of edges and triangles, $3(n - \chi(S))$ and $2(n - \chi(S))$ respectively. It is sometimes useful to consider an ideal triangulation T as a way to realize (S, V) by assembling $2(n - \chi(S))$ “abstract” triangles by gluing their “abstract” edges in pairs in such a way that no edge remains unglued. Ideal triangulations of (S, V) are considered up to the *ideal transit equivalence* which is generated by isotopy fixing V pointwise and

the elementary *diagonal exchange move* also called *flip*. Denote by $\mathcal{T}^{id}(S, V)$ the corresponding quotient set. The following is a well-known connectivity result (for a proof see, for instance, [10, 12]).

THEOREM 1.1. *For every (S, V) , $\mathcal{T}^{id}(S, V)$ consists of one point.*

This important result is the starting point for the study of the space of ideal triangulations of (S, V) , optimal geodesic flip paths on it, with application to the study of coarse geometry of mapping class groups (see, for instance, the body and the references of the recent papers [6, 7]). On another hand, different notions of “decorated ideal triangulations” of 3-manifolds considered up to various transit equivalences naturally arise in the developments of 3-dimensional quantum hyperbolic geometry and in several other instances of 3D quantum invariants based on state sums over triangulations. Understanding the corresponding quotient sets is an interesting and usually non-trivial task. We refer to [1] for information about such a 3-dimensional case. In particular, surface ‘branched’ triangulations (see below) emerged within [1], Section 5, in a “holographic” approach to so-called 3D nonambiguous structures.

Every “naked” triangulation T of (S, V) carries some *branchings* (T, b) (see Lemma 2.10) where, by definition, b is a system of edge orientations which lifted to every abstract triangle (t, b) of T is induced by a (local) ordering of the vertices so that each edge goes towards the biggest endpoint. Equivalently, b promotes T to be a Δ -*complex* following [9], Chapter 2. It is easy to see that for every branching (T, b) , every naked flip $T \rightarrow T'$ can be enhanced to a *b-flip* $(T, b) \rightarrow (T', b')$ defined by the property that every “persistent” edge in both (T, b) and (T', b') keeps the orientation. Isotopy relatively to V as above and b -flips generate the so-called *ideal b-transit equivalence* and we denote by $\mathcal{B}^{id}(S, V)$ the corresponding quotient set. We define the *symmetrized relation* by adding to the generators the *complete inversion* that is we stipulate that every (T, b) is equivalent to $(T, -b)$ where $-b$ is obtained by inverting all edge orientations of b , and we denote by $\tilde{\mathcal{B}}^{id}(S, V)$ the corresponding quotient sets. It is not hard to see that $\sigma([T, b]) = [(T, -b)]$ defines an involution on $\mathcal{B}^{id}(S, V)$ and that $\tilde{\mathcal{B}}^{id}(S, V) \sim \mathcal{B}^{id}(S, V)/\sigma$. By the topological homogeneity of every surface, the cardinality of $\mathcal{B}^{id}(S, V)$ only depends on the topological type of S and the number $n = |V|$; sometimes we will write (S, n) instead of (S, V) .

Assuming Theorem 1.1, the following branched version of the connectivity result is the main result of the present note.

THEOREM 1.2. *(1) If S is orientable or is nonorientable and $\chi(S)$ is even and strictly negative, then for every (S, V) , $\mathcal{B}^{id}(S, V)$ consists of one point.*

(2) If S is nonorientable and either $\chi(S) = 0$ or $\chi(S)$ is odd, then for every (S, V) , $\tilde{\mathcal{B}}^{id}(S, V)$ consists of one point.

As $\tilde{\mathcal{B}}^{id}(S, V)$ is a quotient of $\mathcal{B}^{id}(S, V)$ by an involution, it follows that in case (2), $|\tilde{\mathcal{B}}^{id}(S, V)| \leq 2$.

CONJECTURE 1.3. *If S is nonorientable and either $\chi(S) = 0$ or $\chi(S)$ is odd, then for every (S, V) , $|\mathcal{B}^{id}(S, V)| = 2$.*

This will be confirmed at least for $\mathcal{B}^{id}(\mathbf{P}^2(\mathbb{R}), 2)$ (Proposition 3.3).

By adding to the b -flips the positive branched $0 \rightarrow 2$ b -bubble moves (see Section 2.3 below) and their inverse (or equivalently the positive stellar $1 \rightarrow 3$ branched moves and their inverse), we get the *completed b -transit equivalence* with quotient set denoted by $\mathcal{B}(S, V)$. ‘Positive’ means that the number of triangles increases, for example, $1 \rightarrow 3$ means that an initial triangle t is subdivided by three sharing one vertex internal to t . A positive bubble (stellar) move produces an ideal triangulation of (S, V') where V' contains one further marked point of S ; if it is part of a b -transit which connects two ideal branched triangulations of (S, V) , then it must be compensated later by a negative inverse move. We will see a quick direct proof of the following weaker connectivity result (no matter if S is orientable or not).

PROPOSITION 1.4. *For every (S, V) , $\mathcal{B}(S, V)$ consists of one point.*

We will see in Section 2 that b -flips can be naturally organized in some sub-families so that restricted transit equivalences can be defined leading to non-trivial, actually infinite quotient sets. We will study them by pointing out some structures on (S, V) of geometric/topological nature, carried by each branched triangulation (T, b) of (S, V) that are invariant for the given restricted equivalence. Theorem 1.2 itself should be enlightened by the mutations of these structures following an arbitrary ideal b -transit.

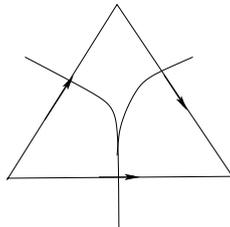


Figure 1: A branched triangle with its dual track

The dual viewpoint. Let S_V be the surface with n boundary components ($n = |V|$) obtained by removing from S a small open disk around each $v \in V$. For every ideal triangulation T of (S, V) , the 1-skeleton $\theta = \theta_T$ of the dual cell decomposition is a generic (internal) spine of S_V . That is, θ is a graph with

3-valent vertices and S_V is a ribbon graph that thickens θ . If (T, b) is branched, this promotes θ to be a *transversely oriented train track* (see Remark 2.2) (θ, b) in the interior of S_V ; for simplicity we keep the same notation “ b ”, see Figure 1. In such a case, the open 1-cells are called the *branches*, the 0-cells the *switch-points* of (θ, b) . If S is oriented, the transverse orientation at each branch can be expressed by the dual orientation of the branch itself, so that it intersects the dual b -oriented edge of (T, b) with intersection number equal to 1. In this way (θ, b) becomes an *oriented train track*. By definition, (θ, b) is a *branched spine* of S_V . Viceversa, every ribbon graph, \tilde{S} say, carried by a (possibly branched) spine θ as above gives rise to a (possibly branched) ideal triangulation $T = T_\theta$ of (S, V) obtained by filling each boundary component of \tilde{S} with a punctured 2-disk. Triangulation moves, possibly branched, can be equivalently rephrased in terms of (possibly branched) spine moves. We will freely adopt both equivalent dual viewpoints.

REMARK 1.5: Although they are equivalent, there is some qualitative difference between spines and triangulations. A flip is a *discrete* transition with a cell decomposition as intermediate “state” which is no longer a triangulation (it includes one quadrilateral). The corresponding spine transition can be realized by a *continuous* deformation passing through a nongeneric spine (with one 4-valent vertex).

2. Generalities on b -transit

An “abstract” b -flip acts on a quadrilateral Q endowed with a branched triangulation $(t_1 \cup t_2, b)$ made by two triangles with one common edge $e = t_1 \cap t_2$ (a diagonal of the quadrilateral). A b -flip produces another branched triangulation $(t'_1 \cup t'_2, b')$ of Q made by two triangles having as common edge $e' = t'_1 \cap t'_2$ the other diagonal of Q , while b and b' coincide on the persistent edges which form the boundary of Q . An abstract b -flip can be applied at every couple of abstract triangles of any branched ideal triangulation (T, b) of any (S, V) , (partially) glued in T along a common edge. When we say that a b -flip verifies a certain property we mean that this holds “universally” for every (S, V) and every triangulation (T, b) at which the flip operates.

2.1. A combinatorial classification of b -flips

For every branched triangulation $(t_1 \cup t_2, b)$ of Q as above, there are either one or two ways to enhance the naked flip $t_1 \cup t_2 \rightarrow t'_1 \cup t'_2$ to a b -flip $(t_1 \cup t_2, b) \rightarrow (t'_1 \cup t'_2, b')$. This last is sometimes denoted by $f_{e, b, b'}$ while the underlying naked flip is denoted by f_e . Then we can distinguish a few families of b -flips.

DEFINITION 2.1. 1. A b -flip $f_{e, b, b'}$ is forced if it is the unique branched flip which enhances f_e , starting from $(t_1 \cup t_2, b)$.

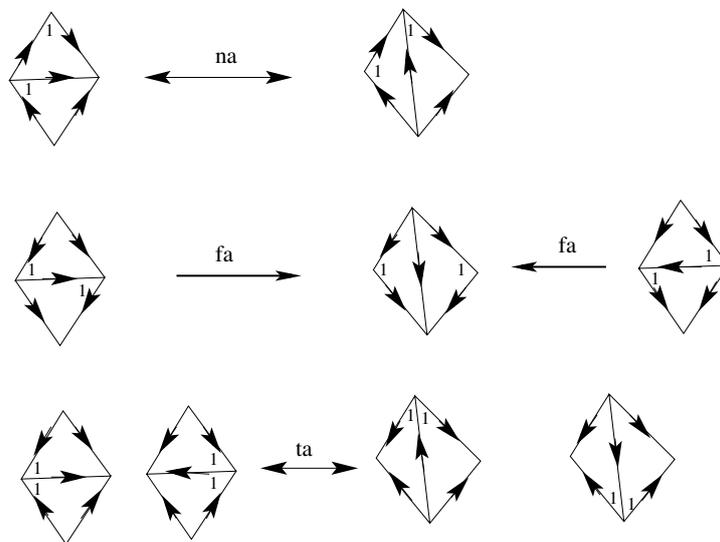


Figure 2: Branched flips.

2. A b -flip $f_{e,b,b'}$ is non-ambiguous if both $f_{e,b,b'}$ and the inverse b -flip $f_{e',b',b}$ are forced.
3. A b -flip $f_{e,b,b'}$ is forced ambiguous if it is forced but the inverse b -flip is not.
4. A b -flip $f_{e,b,b'}$ is totally ambiguous if no one among $f_{e,b,b'}$ and $f_{e',b',b}$ is forced.

In Figure 2, we show typical samples of b -flips following the above classification. Note that it is invariant under the total inversion of the branchings. We have labelled by 1 the corner of each branched triangle formed by the two edges that carry the *prevalent orientation*. Here 1 is just a highlighting label. For every branched triangulation (T, b) of (S, V) , for every vertex v of T , the number of corners at v in its star labelled by 1 is even, say $2d_b(v)$. For the branching b induces an orientation on each (abstract) triangle of T ; given the (abstract) star of v an auxiliary orientation, the 1-labelled corners at v belong to triangles (t, b) whose b -orientations alternate compared with the reference one. It is clear that

$$\chi(S) = |V| - \sum_v d_b(v) = \sum_v (1 - d_b(v)) .$$

REMARK 2.2: Let (T, b) be a branched ideal triangulation of (S, V) and (θ, b) the dual (transversely oriented) track. Strictly speaking, to be a train track in the sense, for instance, of [13], we should require furthermore that for every vertex v of T , the index $1 - d_b(v) < 0$. But this is not relevant for the present text.

A few remarks:

- A (abstract) b -flip $f_{e,b}$ is totally ambiguous if and only if the two vertices of $(t_1 \cup t_2, b)$ opposite to the edge e are both either a source or a pit.
- A b -flip is not totally ambiguous if and only if it preserves (in the ‘universal’ sense said above) the vertex numbers $d_b(v)$.
- Let us consider an oriented surface S . The orientation corresponds to a unique simplicial *fundamental* \mathbb{Z} -2-cycle

$$f(T, b) = \sum_t *_{(t,b)}(t, b)$$

where every $*_{(t,b)} \in \{\pm 1\}$. Denote by $S_{\pm} = S_{\pm}(T, b)$ the union of triangles such that $*_{(t,b)} = \pm 1$. Then S decomposes as $S = S_+ \cup S_-$. Denote by ∂S_{\pm} the boundary 1-cycle of the simplicial 2-cochain supported by S_{\pm} . Then a b -flip $f_{e,b}$ is nonambiguous if and only if it is not totally ambiguous and for every oriented (S, V) and every application of the flip on triangulations of (S, V) , $(T, b) \rightarrow (T', b')$, we have that $S_{\pm}(T, b) = S_{\pm}(T', b')$, hence also ∂S_{\pm} is preserved.

In Section 5, we will consider again this classification of b -flips in a more conceptual way.

2.2. Inversion of an ambiguous edge

DEFINITION 2.3. (1) Let T be a naked ideal triangulation of (S, V) . An edge e of T is said *trapped* if it results by the identification of two edges of one “abstract” triangle. Otherwise, e is said *untrapped*. A trapped edge corresponds to a one-vertex loop in the dual spine θ .

(2) Given a branching (T, b) , a b -oriented edge e is said *ambiguous* in (T, b) if by inverting its orientation we keep a branched triangulation (T, b') .

With the notation of the above point (2), we have:

LEMMA 2.4. *If e is ambiguous and untrapped in (T, b) , then (T, b) and (T, b') are connected by two b -flips.*

Proof. Denote by f_{e,b,b^*} a b -flip that enhances the naked flip f_e with inverse naked flip $f_{e'}$. We easily see that the untrapped edge e is ambiguous if and only if either f_{e,b,b^*} is forced ambiguous or is totally ambiguous. Hence f_{e,b,b^*} followed by $f_{e',b^*,b'}$ convert (T, b) to (T, b') . \square

Then we can add the elementary move of inverting any untrapped ambiguous edge without changing the ideal b -transit equivalence.

We have

LEMMA 2.5. (1) *For every T as above, there is a sequence of flips $T \Rightarrow T'$ such that T' does not contain trapped edges.*

(2) *If T and T' do not contain trapped edges then they can be connected by a sequence of flips through triangulations without trapped edges.*

Proof. The vertex of a loop in the spine θ which is dual to a trapped edge of T is connected by an edge to the rest of the spine. By performing the dual flip at this edge we remove the loop without introducing new ones. If such a loop appears in a sequence of flips connecting T and T' as in (2), then we can follow it till it disappears so that we can eventually remove it from the sequence. \square

2.3. Bad nutshells

A positive naked $0 \rightarrow 2$ bubble produces a so-called *nutshell* made by two triangles identified along two common edges. Not every branched nutshell (N, b) supports a negative $2 \rightarrow 0$ b -bubble.

DEFINITION 2.6. *A branched nutshell (N, b) is bad if the two boundary edges form an oriented circle. Otherwise (N, b) is a good nutshell.*

The following Lemma is immediate.

LEMMA 2.7. (1) *If (N, b) is a bad nutshell, then the central vertex is necessarily either a pit or a source.*

(2) *(N, b) is good if and only if it supports a negative b -bubble move.*

(3) *Two different good nutshells (N, b) and (N, b') sharing the same oriented boundary edges are connected by either one or two consecutive inversions of internal (hence untrapped) ambiguous edges.*

A positive naked $1 \rightarrow 3$ move produces a so-called *triangular star*. Similarly as before, not every branched triangular star, say (\mathfrak{S}, b) , supports a negative b - $(3 \rightarrow 1)$ move.

DEFINITION 2.8. *A branched triangular star (\mathfrak{S}, b) is bad if the three boundary edges form an oriented circle. Otherwise (\mathfrak{S}, b) is good.*

Easily we have

LEMMA 2.9. (1) If (\mathfrak{S}, b) is bad, then the central vertex is necessarily either a pit or a source.

(2) (\mathfrak{S}, b) is good if and only if it supports a negative b - $(3 \rightarrow 1)$ move.

(3) Two good b -triangular stars (\mathfrak{S}, b) and (\mathfrak{S}, b') sharing the same oriented boundary edges are connected by a finite sequence of consecutive inversions of internal (hence untrapped) ambiguous edges.

2.4. Existence of branched triangulations

We have

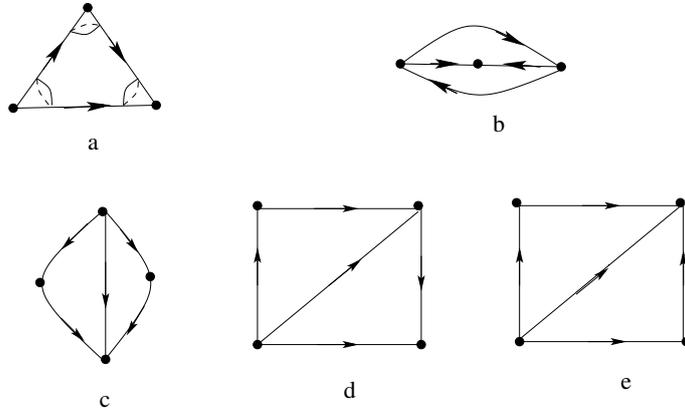


Figure 3: Existence of branched triangulations.

LEMMA 2.10. Every ideal triangulation T of (S, V) can be branched.

Proof. Thanks to Theorem 1.1 and since for every branched triangulation (T, b) , every naked flip $T \rightarrow T'$ can be enhanced to a b -flip $(T, b) \rightarrow (T', b')$, it is enough to show that every (S, V) admits a branched triangulation. Using the b -bubbles, we see that if (S, V) admits such a triangulation, then this holds for every (S, V') such that $|V'| \geq |V|$. Then it is enough to show that for every S , there exists a branched triangulation (T, b) of (S, V) such that $n = |V|$ is the minimum for which $\chi(S) - n < 0$. There are, of course, several ways to do it. We indicate a way that will be suited for further use. If S is a sphere, then $n = 3$ and a branched triangulation is obtained by gluing two copies of a branched triangle along the common boundary (see Figure 3-a). If $S = \mathbf{P}^2(\mathbb{R})$

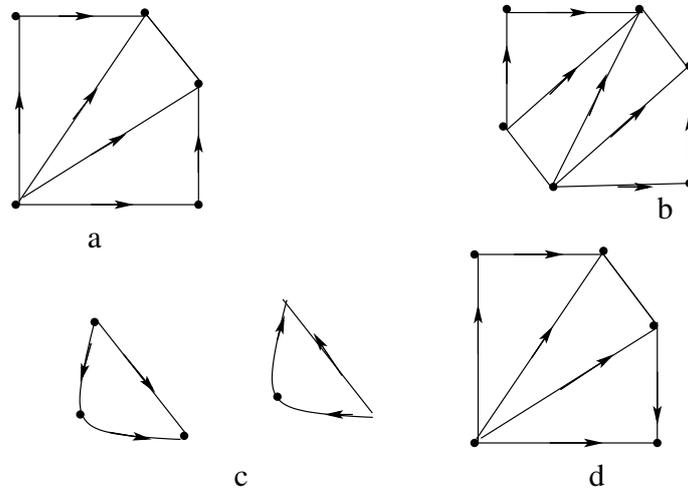


Figure 4: Existence of branched triangulations.

is a projective plane, then $n = 2$ and we can use the realization of $\mathbf{P}^2(\mathbb{R})$ by identifying the two edges of a bigon and triangulating it with one internal vertex (see Figure 3-b). For all other S , $n = 1$. If $S = \mathbf{P}^2(\mathbb{R}) \# \mathbf{P}^2(\mathbb{R})$ is a Klein bottle, that is the connected sum of two projective planes, we get a branched triangulation of (S, V) using the realization of S which identifies the boundary of a quadrilateral obtained by gluing two “truncated bigons” (see Figure 3-c). In Figure 3-d, we show another branched triangulation of the Klein bottle by identifying in pairs the opposite edges of a quadrilateral; similarly, in Figure 3-e, we suggest a triangulation of the torus $S^1 \times S^1$. In Figure 4, we show the elementary bricks to realize all other cases. These bricks are branched triangulations of certain surfaces with boundary. In Figure 4-a, we see a one-pierced torus, in Figure 4-b, a twice-pierced torus, that is a torus from which we have removed respectively one or two open 2-disks. They are obtained using respectively a one-truncated or a twice-truncated quadrilateral with opposite edges identified in pairs. In Figure 4-d, we see a similar realization of a one-pierced Klein bottle. In Figure 4-c, we see a one-pierced projective plane given by a truncated bigon with the two possible branchings. The edges without any indicated orientation are ambiguous so that the orientation can be chosen arbitrarily. If S is orientable of genus $g > 1$, we can realize it by a chain of $g - 2$ twice-pierced tori capped by two one-pierced ones. If S is nonorientable and $\chi(S) = 2 - r < 0$ is odd, set $g = \frac{r - 1}{2}$; then we can obtain S by a chain of $g - 2$ twice-pierced tori capped by a one-pierced torus and a one-pierced projective

plane. If $\chi(S) - r < 0$ is even, set $g = \frac{r-2}{2}$, then we obtain S by a chain of $g-2$ twice-pierced tori capped by a one-pierced torus and a one-pierced Klein bottle. \square

The boundary of every brick as above is the union of loops with one vertex. The corresponding edge of the triangulation is called a *connection edge*. For every triangulation of S obtained so far, the connection edges become separating loops that decompose S by the bricks; in Figure 4, these edges correspond to the ambiguous edges on the boundary of the truncated quadrilaterals or to the nonambiguous edge in the branched truncated bigons.

2.5. A preliminary reduction in studying the b -transit equivalence

We assume Theorem 1.1. Isotopy relative to V will be understood. At the end of this section, we will obtain a quick proof of Proposition 1.4.

LEMMA 2.11. *The following facts are equivalent to each other:*

1. $\mathcal{B}^{id}(S, V)$ consists of one point.
2. For every naked ideal triangulation T of (S, V) , every two branchings (T, b) and (T, b') are connected by a chain of b -flips.
3. There exists a naked ideal triangulation T of (S, V) such that every two branchings (T, b) and (T, b') are connected by a chain of b -flips.

Proof. Obviously, (1) \Rightarrow (2) \Rightarrow (3). To prove (3) \Rightarrow (1) we argue similarly to the proof of Lemma 2.10. Let (T_1, b_1) and (T_2, b_2) be ideal triangulations of (S, V) . By Theorem 1.1, there is a naked ideal transit $T_j \Rightarrow T$, $j = 1, 2$. There is no obstruction to enhance it to sequences of b -flips $(T_j, b_j) \Rightarrow (T, b'_j)$. The Lemma follows immediately. \square

Similar statements hold for both the symmetrized and completed b -transit equivalences. Given (T, b) and (T, b') , denote by $\delta(b, b')$ the set of edges of T at which b and b' are opposite. The previous considerations suggest two possible ways to prove that $\mathcal{B}^{id}(S, V)$ (or $\tilde{\mathcal{B}}^{id}(S, V)$ or $\mathcal{B}(S, V)$) consists of one point.

(A) For a given (S, V) , detect a distinguished naked triangulation T for which we can check directly that (3) of Lemma 2.11 holds.

(B) Point out a few procedures such that for every naked T and every couple (T, b) and (T, b') having non-empty $\delta(b, b')$, we can apply one of them producing b -transits $(T, b) \Rightarrow (T', b_1)$ and $(T, b') \Rightarrow (T', b_2)$ such that $|\delta(b_1, b_2)| < |\delta(b, b')|$.

The **B**-way, suitably adapted, works quickly for $\mathcal{B}(S, V)$.

Proof of Proposition 1.4. Given (T, b) and (T, b') as above, perform on both triangulations a b -(1 \rightarrow 3) move at each triangle of T in such a way that all the three new oriented edges point toward the new internal vertex. We get in this way (T', b_1) and (T', b_2) such that $\delta(b_1, b_2) = \delta(b, b')$. We realize that every $e \in \delta(b, b')$ is untrapped and ambiguous in both (T', b_1) and (T', b_2) . So we conclude by several applications of Lemma 2.4 (by the way, in the present situation every inversion of e is obtained by a sequence of two totally ambiguous b -flips). \square

In the next section, we will prove Theorem 1.2 by implementing the **A**-way.

3. A-Proof of the main Theorem

We obtain stronger results. We have:

THEOREM 3.1. (1) *If S is orientable or is nonorientable and $\chi(S)$ is even and strictly negative, then for every (S, V) , there exists a distinguished ideal triangulation T such that every (T, b) and (T, b') can be explicitly connected by a sequence of inversions of untrapped ambiguous edges.*

(2) *If S is nonorientable and either $\chi(S) = 0$ or $\chi(S)$ is odd, then for every (S, V) , there exists a distinguished ideal triangulation T such that for every couple (T, b) and (T, b') , either (T, b) and (T, b') or $(T, -b)$ and (T, b') can be explicitly connected by a sequence of inversions of untrapped ambiguous edges.*

For every (S, V) , let us said *inversive* any ideal triangulation T which, case by case, satisfies the conclusions of Theorem 3. Finally, we have:

THEOREM 3.2. *For every (S, V) , every triangulation T without trapped edges is inversive.*

Proof of Theorem 3.1. For every S there is a minimum n_S such that $\chi(S) - n_S < 0$. The proof is by induction on $n \geq n_S$.

The initial case: (S, n_S) .

- $(\mathbf{P}^2(\mathbb{R}), 2)$. We use the naked triangulation T of Figure 3-b. Let (T, b) and (T, b') be supported by T . By total inversion, we can assume that b and b' agree on the internal edges so that the internal vertex is a pit. The boundary edges of the nutshell lift an ambiguous edge of both (T, b) and (T, b') ; if $b \neq b'$ we conclude by inverting it in b . Then $\tilde{\mathcal{B}}^{id}(\mathbf{P}^2(\mathbb{R}), V)$ consists of one point.

PROPOSITION 3.3. $|\mathcal{B}^{id}(\mathbf{P}^2(\mathbb{R}), 2)| = 2$.

Proof. Let (T, b) and (T, b') be as above such that $\delta(b, b')$ consists of the two internal edges. If we flip an internal edge of (T, b) we produce a trapped edge; to get the same naked configuration we must flip the same edge in (T, b') and $|\delta|$ is unchanged. If we flip the edge of (T, b) which lifts to the boundary edges of the nutshell, we get a triangulation (T_1, b_1) which is abstractly like (T, b') for *another* nutshell, but the two vertices exchange their role, so (T_1, b_1) cannot be relatively isotopic to (T, b') . \square

- $(S^2, 3)$. Take T made by two triangles glued along the common boundary as in Figure 3-a. Every branched (T, b) is determined by a labelling of the vertices by 0, 1, 2. Fix a (T, b_0) , then all (T, b) are indexed by the elements σ of the symmetric group Σ_3 , so that (T, b_0) corresponds to the identity. This group is generated by the transpositions $\{(01), (12)\}$. If $b = b_\sigma$, write σ as a product of a minimal number of these generators. Every such a sequence of transpositions corresponds to a sequence of inversions of ambiguous edges going from (T, b_0) to (T, b) .

In all other cases $n_S = 1$.

- $S = S^1 \times S^1$. Consider T as in Figure 3-e. Every branching of T is uniquely encoded by a total order of the three vertices of one of the abstract triangles of T . Then we can manage similarly as for $(S^2, 3)$ by checking that also in this case every transposition in a product of the generators corresponds to the inversion of an ambiguous edge.

- Let S be a Klein bottle. Refer to Figure 3-c,d. If we use the triangulation T made by two truncated bigons, it is immediate that it carries exactly two branchings say (T, b) and $(T, -b)$. If we use as T the other triangulation, we see that it carries four branchings, distributed into two pairs $\{(T, b), (T, b')\}$, $\{(T, -b), (T, -b')\}$ such that (T, b') is obtained from (T, b) via the inversions of an ambiguous edge.

Let us face now the remaining *generic cases* such that $\chi(S) < 0$, following the two sub-cases in the statement of Theorem 3.1.

Case (1) Let S be either orientable with $\chi(S) < 0$ or nonorientable with $\chi(S) < 0$ and even. Take the naked triangulation T depicted in the proof of Lemma 2.10; let (T, b_0) be a branched triangulation constructed therein. Let (T, b) any branched triangulation supported by T . This determines a system of orientations on the family of connection edges. Every connection edge is ambiguous in (T, b_0) so that up to some inversions, we can assume that (T, b_0) and (T, b) share such a system of connection orientations. Let us cut S at the connection edges. By restriction, we get a family of pairs of branched triangulated bricks (B, b_0) and (B, b) that coincide at every connection edge. It is enough to show that every (B, b) is connected to (B, b_0) by a sequence of inversions of ambiguous edges, without touching the connection edges. This

can be checked case by case. We have three types of B , the one-pierced torus (Figure 4-a) or Klein bottle (Figure 4-d) and the twice-pierced torus (Figure 4-b). For each one-pierced brick, we have two possible b_0 -orientations at the connection edges; for the twice-pierced torus, there are four. For every brick and every pair of opposite local configurations at the connection edges, we see using the total inversion that the desired result holds for one if and only if it holds for the other. Then we are reduced to study one configuration in the one-pierced cases, two in the twice-pierced case. We organize the discussion as follows, referring to Figure 4:

- Denote by t the top (abstract) triangle of (B, b_0) . Encode (t, b_0) by labelling its vertices by 0, 1, 2, say v_0, v_1, v_2 ; do the same for the bottom triangle (t', b_0) , getting v'_0, v'_1, v'_2 . In the one-pierced bricks, $v_0 = v'_0$. Every (abstract) edge of B has two vertices belonging to $\{v_0, v_1, v_2, v'_0, v'_1, v'_2\}$. For every (B, b) , every oriented edge (e, b) will be denoted by its vertices, $e = vw$, written in the order so that the orientation emanates from the initial vertex v toward the final vertex w . For every one-pierced brick, we stipulate that the connection edge in (B, b_0) has v_2 as initial vertex. For the twice-pierced torus we stipulate that in (B, b_0) the pairs of connection edges is either $(v_2v'_2, v_0v'_0)$ or $(v_2v'_2, v'_0v_0)$. Having fixed the orientation of the connection edges, b_0 is completely determined by (t, b_0) , that is this propagates in a unique way to a global branching.
- For every permutation $\sigma \in \Sigma_3$ we consider the corresponding branched triangulated triangle (t, b_σ) and we list all the extensions to a global branching, generically denoted (B, b_σ) , if any. (B, b_0) corresponds to the identity.
- By varying $\sigma \in \Sigma_3$, the so obtained (B, b_σ) cover all possible branchings (B, b) and we have to connect (B, b_0) with each (B, b_σ) . It is convenient to start with generating transpositions $\sigma = (0, 1), (1, 2)$, and express all other σ as a product of three or two generators.

Let us pass now to the actual verifications.

The one-pierced torus.

$\sigma = (0, 1)$: there a unique extension $(B, b_{(0,1)})$ which differs from (B, b_0) by the inversion of the ambiguous edge v_1v_0 .

$\sigma = (1, 2)$: there are several extensions. There is only one containing the edge v_0v_2 and this differs from (B, b_0) by the inversion of the ambiguous edge v_2v_1 . Two extensions contain the edge v_2v_0 and differ from each other by the inversion of the ambiguous edge $v_0v'_2$. In the one containing $v_0v'_2$, v_2v_0 is ambiguous, hence by inverting it we are in the first case.

$\sigma = (0, 1, 2) = ((1, 2)(0, 1))$: there are two extensions which differ from each other by the inversion of the ambiguous edge $v_0v'_2$. In the one containing $v_0v'_2$,

the edge v_0v_2 is ambiguous, hence possibly by inverting it we reach the case $(B, b_{(1,2)})$.

$\sigma = (0, 2, 1) = (0, 1)(1, 2)$: the discussion is similar to the one for $(0, 1, 2)$; up to some inversion of ambiguous edges we reach the case $(B, b_{(0,1)})$.

$\sigma = (0, 2) = (0, 1)(1, 2)(1, 0)$: there is only one extension. Here, $v_0v'_1$ is ambiguous. By inverting it we reach the case $(B, b_{(0,2,1)})$.

The first verification is complete.

The one-pierced Klein bottle.

$\sigma = (0, 1)$: there a unique extension $(B, b_{(0,1)})$ and it differs from (B, b_0) by the inversion of the ambiguous edge v_0v_1 .

$\sigma = (1, 2)$: there are no extensions.

$\sigma = (0, 1, 2)$: there is only one extension $(B, b_{(0,1,2)})$. The following sequence of inversions of ambiguous edges realizes a transit from this extension to $(B, b_{(0,1)})$ (we indicate the initial orientation before the inversion): $v_0v_1, v'_2v_0, v_2v_1, v_2v_0$.

$\sigma = (0, 2, 1)$: there is only one extension. The following sequence of inversions of ambiguous edges realizes a transit to (B, b_0) : v_2v_0, v_1v_0 .

$\sigma = (0, 2)$: there are two extensions which differ to each other by the ambiguous edge $v_0v'_2$. In the ones containing the oriented v'_2v_0 , the edge v_1v_0 is ambiguous. By inverting it we reach $(B, b_{(0,1,2)})$.

The second verification is complete.

The twice-pierced torus. We have two sub-cases depending on the orientations either $(v_2v'_2, v_0v'_0)$ or $(v_2v'_2, v'_0v_0)$ of the two connection edges.

Sub-case $(v_2v'_2, v_0v'_0)$.

$\sigma = (0, 1)$: There are two extensions that differ by the ambiguous edge $v'_0v'_2$. In the ones containing the oriented edge $v'_0v'_2$, v_1v_0 is ambiguous and possibly inverting it we reach (B, b_0) .

$\sigma = (1, 2)$: this is very similar to the case $(0, 1)$.

$\sigma = (0, 1, 2)$: there are several extensions. There is only one containing v'_2v_0 which is ambiguous. In the ones containing $v_0v'_2$ both v_0v_2 and $v'_0v'_2$ are ambiguous, then after at most two inversions we reach $(B, b_{(1,2)})$.

$\sigma = (0, 2, 1)$: this is very similar to the case $(0, 1, 2)$; via a sequence of inversions we reach now $(B, b_{(0,1)})$.

$\sigma = (0, 2)$: there are four extensions that differ by suitable inversions of the edges v_1v_2 and $v_0v'_2$ which are both ambiguous. Then up to such inversions we reach $(B, b_{(0,1)})$.

Sub-case $(v_2v'_2, v'_0v_0)$. At this point, the fourth verification is a routine, we leave it to the reader.

is connected to (B_{n-1}, b_{n-1}) by a sequence of inversions of ambiguous edges. Finally, we conclude by applying Lemma 2.9 to the innermost triangular star.

(b) The innermost triangular star as above is bad in (B_n, b) . Consider first (B_2, b) ; we readily see that v_1 (resp. v'_1) is necessarily either a pit or a source (either a source or a pit). In any case at least one among v_0v_2 and $v'_0v'_2$ is ambiguous; by inverting it, the triangular star becomes good and we reach the case (a). In general, we can assume by induction that the innermost triangular star in B_{n-1} is good for the restriction of b to B_{n-1} , so that we can apply the above reasoning to the innermost triangular star in (B_n, b) and reach again the case (a).

The proof of Theorem 3.1 in the generic cases is now complete.

For the remaining sporadic cases such that $\chi(S) \geq 0$ we limit ourselves to some indications.

- (S^2, n) , $n \geq 3$. Denote by T_3 the triangulation used above for $n = 3$. Select one triangle t and one edge e . For every $n > 3$, the distinguished triangulation T_n for (S^2, n) is obtained by induction on n by performing a $1 \rightarrow 3$ move on the triangle of T_{n-1} which is contained in t and contains e . In particular, T_4 corresponds to the triangulation of the boundary of a tetrahedron. Every (T_4, b) is determined by a labelling of the vertices by $0, 1, 2, 3$. Fix a (T, b) ; then the branchings are indexed by the elements of the symmetric group Σ_4 . This is generated by the transpositions $(0, 1), (1, 2), (2, 3)$. Write every σ as a product of these generators with a minimal number of terms. This corresponds to a sequence of inversions of ambiguous edges connecting (T, b) and (T, b_σ) . For $n > 4$ we argue by induction on n .

- $(S^1 \times S^1, n)$ or $(\mathbf{P}^2(\mathbb{R}) \# \mathbf{P}^2(\mathbb{R}), n)$, $n \geq 1$. In both cases we start with the triangulation say T_1 used for $n = 1$. Precisely, we refer to Figure 3-e and to Figure 3-d respectively. Then T_n is obtained from T_{n-1} by performing a $1 \rightarrow 3$ move on the triangle contained in the top triangle of T_1 and containing its diagonal edge.

- $(\mathbf{P}^2(\mathbb{R}), n)$, $n \geq 2$. We start with T_2 used for $n = 2$. Referring to Figure 3-b, T_3 is obtained by performing a bubble move at the internal edge on the left side. Denote by t the new triangle contained in the top half of T_2 and by e its edge contained in the interior of this top-half. Then, for $n > 3$, T_n is obtained from T_{n-1} by performing a $1 \rightarrow 3$ move on the triangle of T_{n-1} which contains e .

The proof of Theorem 3.1 is now complete. □

Proof of Theorem 3.2. As usual, we have to consider two cases.

Case (a): $\chi(S)$ is not strictly negative and odd. The distinguished inversive triangulations T_n of (S, n) constructed in the proof of Theorem 3.1 have no

trapped edges. Using Lemma 2.5 , we deduce that for every triangulation T of (S, V) without trapped edges, there exists a (possibly varying) inversive triangulation T^* without trapped edges and a sequence of flips $T^* \Rightarrow T$ through triangulations without trapped edges. Denote by l the number of flips. We work by induction on l . Let T' be obtained by performing the first $l - 1$ flips. By induction, the theorem holds for T' . Hence we are reduced to check the case $l = 1$. We fix the notations as follows: e is the flipping edge, that is a diagonal of a quadrilateral $Q = t_1 \cup t_2$ in T^* , $t_1 \cap t_2 = e$; e' denote the other diagonal of Q , that is the edge of T which replaces e . Let (T^*, b) and (T^*, b') be connected by a sequence of k inversions of ambiguous edges. Let (T, \tilde{b}) , (T, \tilde{b}') be obtained by b -enhancing in some way the flip $(T^*, b) \rightarrow T$ and $(T^*, b') \rightarrow T$. We want to modify the sequence to get one connecting these branchings of T . If a b -flip is not forced then the two possibilities are related by inverting an ambiguous edge, so this is essentially immaterial for our discussion. If an inversion concerns an edge not contained in (Q, b) then it makes sense also on (T, \tilde{b}) . By these remarks and working by induction on k , we are reduced to analyze the inversion of an edge e^* contained in (Q, b) . There are a few possibilities.

- $e^* = e$; then the b -flip is either totally ambiguous or forced ambiguous if the vertices of (Q, b) opposite to e are either both a pit (resp. source) or one is a pit and the other a source. So in the first case, we possibly replace the inversion of e with the inversion of e' .

Assume now that $e^* \neq e$.

- e is ambiguous in (Q, b) as above. If the flip is totally ambiguous, there are two possibilities for e^* . Then the inversion of e^* can be performed on (T, \tilde{b}) , possibly after having inverted the ambiguous edge e' . If the flip is forced ambiguous, again there are two possibilities for e^* and in every case the inversion of e^* can be performed on (T, \tilde{b}) .

- e is ambiguous in one of the two triangles of (Q, b) and nonambiguous in the other. The flip is nonambiguous. There are three possibilities for e^* . We readily check that in every case, the inversion of e^* can be performed on (T, \tilde{b}) .

- e is nonambiguous in both triangles of (Q, b) . The flip is not forced. We check that the inversion of e^* can be performed on (T, \tilde{b}) , possibly after having inverted the ambiguous edge e' .

This completes the proof in **Case (a)**.

Case (b): $\chi(S)$ is strictly negative and odd. The distinguished triangulations T_n of (S, n) have one trapped edge carried by the one-pieced projective plane. Let T_n^* be obtained by flipping its connection edge. T_n^* does not contain any trapped edge and arguing similarly as above, we see that it is inversive. Then the proof is like in **Case (a)**.

Theorem 3.2 is achieved. □

4. On a **B**-proof of the main Theorem

The basic difference between the **A**- and the **B**-way is that the first deals with distinguished triangulations while the second should apply to any couple (T, b) , (T, b') of branchings on the same arbitrary naked triangulation. Using Lemma 2.5 and Lemma 2.4, it is not restrictive to deal under the assumptions that: (1) T does not contain trapped edges; (2) Every edge $e \in \delta(b, b')$ is nonambiguous in both (T, b) and (T, b') . Moreover, we specify the **B**-way requiring that these conditions are preserved along with the process.

Using Theorem 3.2, it is not hard to build an algorithm that connects (T, b) with (T, b') . We can construct a (not unique) oriented tree with root at (T, b) , such that every vertex of the tree corresponds to a branching of T , every branching appears only once, if two vertices $\mathfrak{b} < \mathfrak{b}'$ are connected by an edge then we pass from \mathfrak{b} to \mathfrak{b}' by inverting one ambiguous edge of (T, \mathfrak{b}) , and every vertex has one antecedent. Then (T, b') is a vertex of the tree and the unique path in the tree connecting the root (T, b) to (T, b') corresponds to a sequence of inversions of ambiguous edges connecting the two branchings on T .

On the other hand, genuine implementation of the **B**-way does not require that the naked triangulation T is fixed along with the process. Moreover, the above application of Theorem 3.2 is not really in the spirit of the **B**-way if we make explicit the underlying idea that the modifying procedures should be local, that is ‘universally’ applicable whenever certain local configurations appear in a pair (T, b) , (T, b') . To make concrete these considerations, we will implement a **B**-way. For simplicity, we will deal in the generic ($\chi(S) < 0$) and *orientable* case. Beyond a different proof of (part of) Theorem 1.2, we believe that this is useful to have some insight about the mutations of branched triangulations under b -transit.

At first, we analyze the effects of *flipping* $e \in \delta(b, b')$ in both (T, b) and (T, b') looking for a decrease of $|\delta|$, if any. Sometimes, we say that an edge $e \in \delta(b, b')$ is *disorientated*. Let t_1 and t_2 be the two triangles of T which share e . As e is nonambiguous in both triangulations, then e is nonambiguous in at least one of the branched triangles (t_j, b) and similarly for the (t_j, b') 's. There are two possibilities:

1. There is at least one triangle, say t_1 , such that e is nonambiguous in both (t_1, b) and (t_1, b') , so that necessarily $(t_1, b') = (t_1, -b)$.
2. e is nonambiguous (resp. ambiguous) in (t_1, b) (resp. (t_2, b)) while e is nonambiguous (resp. ambiguous) in (t_2, b') (resp. (t_1, b')).

Let us analyze the first case.

Case (1). Concerning t_2 , there are three possibilities: it can contain either $k = 0, 1$ or 2 further edges belonging to $\delta(b, b')$. Note that $k = 2$ if and only

if $b' = -b$ on the whole of $t_1 \cup t_2$. We say that the disorientated e is *(1)bad* if $k = 2$ and e is ambiguous in (t_2, b) (hence in (t_2, b')). In all other cases, we say that e is *(1)good*. We have

LEMMA 4.1. *Let $e \in \delta(b, b')$ be (1)good. Then by flipping e in both (T, b) and (T, b') we get (T', b_1) and (T', b_2) which still satisfy our assumptions and such that $|\delta(b_1, b_2)| < |\delta(b, b')|$.*

Proof. First, we note that if f_e creates a trapped edge, then necessarily e is internal to some naked nutshell say N in T . We analyze the situation case by case, according to the value of $k = 0, 1, 2$.

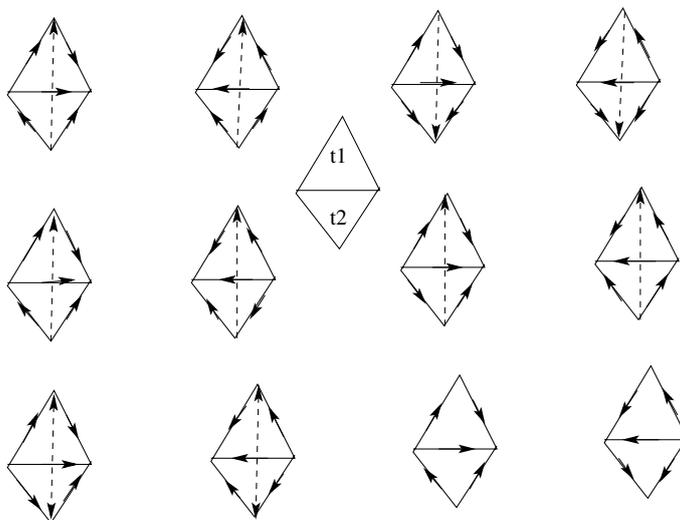


Figure 6: Flipping an untrapped edge, Case (1).

1. If $k = 0$ both f_{e, b, b_1} and f_{e, b', b_2} are nonambiguous and $|\delta|$ decreases by 1. There are not compatible branched nutshells containing e , hence the flip does not create any trapped edge (see the first row of Figure 6).
2. If $k = 1$, we can choose f_{e, b, b_1} and f_{e, b', b_2} in such a way that $|\delta|$ decreases by 1 and the new edge is ambiguous in one of the triangulations obtained so far (see Figure 6, second row). We claim that f_e does not create any trapped edge. Otherwise, one of the branched nutshells (N, b) and (N, b') would be good with internal vertex which is either a pit or a source. Then (recall Lemma 2.7 (3)) it would contain an ambiguous internal edge \tilde{e} belonging to $\delta(b, b')$, against our assumption.

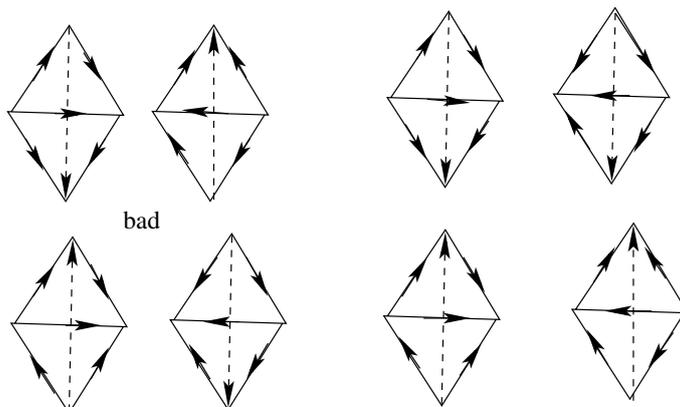


Figure 7: Flipping an untrapped edge, Case (2).

3. If $k = 2$ and e is (1)good, then necessarily e is nonambiguous also in (t_2, b) (hence in (t_2, b')). We can choose f_{e,b,b_1} and f_{e,b',b_2} in such a way that $|\delta|$ decreases by 1 and the new edge is ambiguous in both triangulations obtained so far (see the left side of the third row of Figure 6). If e would belong to a naked nutshell in T , then both (N, b) and (N, b') are good with internal vertex which is either a pit or a source, and we can argue as above. So f_e does not create any trapped edge.

The Lemma is proved. \square

Concerning the (1)bad situation, we readily realize that

LEMMA 4.2. *If e is (1)bad, then by flipping e we keep the same value of $|\delta|$. The new edge is nonambiguous in both triangulations obtained so far (see the right side of the third row of Figure 6). If the flip creates a trapped edge, then both nutshells (N, b) and (N, b') are bad and all edges of N belong to $\delta(b, b')$.*

Let us turn now to the second case.

Case (2). Necessarily, the boundary of $t_1 \cup t_2$ contains exactly one couple $e_1 \subset t_1$, $e_2 \subset t_2$ of edges which do not belong to $\delta(b, b')$. There are two possibilities, see Figure:

- (i) e_1 and e_2 are consecutive edges in the (abstract) quadrilateral $t_1 \cup t_2$ (see the right side of Figure 7). In this case, we can flip e and $|\delta|$ decreases by 1. If this creates a trapped edge, then e is an internal edge of a nutshell N , and both (N, b) and (N, b') are bad with the two boundary edges which do not belong to $\delta(b, b')$.

- (ii) e_1 and e_2 are opposite edges in the (abstract) quadrilateral $t_1 \cup t_2$ and their orientations are necessarily compatible, that is they extend to an orientation of the whole boundary of the quadrilateral (see the left side of Figure 7). By flipping e , we keep the same value of $|\delta|$. The flip does not create any trapped edge.

If we are in case (i) and the flip does not create a trapped edge, then we say that e is *(2)good*. The other cases are *(2)bad*. Let us call generically *good* an edge $e \in \delta(b, b')$ which is either (1)good or (2)good. Otherwise, let us say that it is *bad*. We can successively flip good edges until this is possible so that either δ vanishes and we have done, or δ is nonempty and all disorientated edges are bad. Summarizing we have:

LEMMA 4.3. *To prove Theorem 1.2 it is not restrictive to deal under the following **all-bad assumptions**:*

- (1) T does not contain trapped edges;
- (2) Every edge $e \in \delta(b, b')$ is nonambiguous in both (T, b) and (T, b') ;
- (3) Every edge in $\delta(b, b')$ is bad.

We show now that the all-bad assumptions are quite constraining.

LEMMA 4.4. *Let (T, b) and (T, b') satisfy the all-bad assumptions. Then every $e \in \delta(b, b')$ is (2)bad.*

Proof. Assume that there exists a (1)bad edge $e \in \delta(b, b')$. This propagates to all edges of T so that $b' = -b$ and all edges should be (1)bad. We want to show that this is impossible. We note that there are not adjacent (necessarily bad) nutshells because a common boundary edge should be ambiguous, hence good. Let v be a vertex of T which is not the center of a nutshell and analyze the possible configurations of its (abstract) developed *star* $St(v, b)$ (T, b) (the one in (\tilde{T}, b') is obtained by just reversing the orientations).

Claim: Every such a star $St(v, b)$ has the following qualitative configuration. Every edge in the boundary of the star is *ambiguous* in the respective triangle. $St(v, b)$ can contain an *even* number of bad nutshells sharing the vertex v (necessarily even because otherwise, the star would contain a good edge) and the orientations of their boundaries alternate (compared with the reference orientation of \tilde{S}). The edges of the boundary of the star between two consecutive nutshells have compatible orientations as well as the internal edges of their respective triangles are all either ingoing or outgoing for the central vertex v . Moving along the boundary of the star, the boundary orientations and the “in-out” types switch each time we pass a nutshell. In particular, if there are no nutshells, then the boundary of $St(v, b)$ is an oriented circle.

Proof of the Claim. Assume that there is an edge e in the boundary of $St(v, b)$ which is nonambiguous in the relative triangle. Let us try to complete the star by moving along its boundary in the direction of the orientation of e . Possibly after some boundary edges which are oriented like e and are *ambiguous* in the respective triangle (with internal edges pointing towards the central vertex v), we necessarily find either a boundary edge e' which is nonambiguous in the respective triangle and has opposite orientation compared with e , or a bad nutshell (whose boundary orientation is uniquely determined). We see that in both cases there is an internal edge that is ambiguous in the star, hence good. The claim is proved.

Now we can conclude by noticing that for every $St(v, b)$ with the properties stated in the Claim, there is a vertex v' in the boundary of the star such that the boundary of $St(v', b)$ contains an edge which is nonambiguous in the relative triangle of $St(v', b)$. Lemma 4.4 is proved. \square

REMARK 4.5: The hypothesis that S is orientable has been already employed to limit the way a trapped edge can be produced; it will be important also in the rest of the discussion; the key point is that it prevents that the stars of the disorientated (2)bad edges (in a all-bad configuration) glue each other at edges not belonging to $\delta(b, b')$ producing a Möbius strip. For example, the opposite edges not belonging to $\delta(b, b')$ in a basic (2)bad configuration $(t_1 \cup t_2, b)$, $(t_1 \cup t_2, b')$ cannot be identified.

DEFINITION 4.6. (1) A terminal (2)bad type is either:

- A couple of (2)bad nutshell (N, b) , (N, b') such that $|\delta(b, b')| = 2$ and the boundary edges do not belong to $\delta(b, b')$.
- A couple of triangulated annuli (A, b) , (A, b') obtained from a basic (2)bad configuration $(t_1 \cup t_2, b)$, $(t_1 \cup t_2, b')$ by identifying the opposite boundary edges of the quadrilateral which belong to $\delta(b, b')$. For the resulting triangulations we have $|\delta(b, b')| = 2$ and the boundary of (A, b) (similarly for (A, b')) is formed by two circles each containing one vertex and endowed with opposite orientations.

(2) A couple (T, b) (T, b') is said terminal all-bad if verify the all-bad assumptions and every disorientated edge $e \in \delta(b, b')$ is contained in a terminal (2)bad type.

We have

LEMMA 4.7. To prove Theorem 1.2 it is not restrictive to deal with terminal all-bad couples (T, b) , (T, b') .

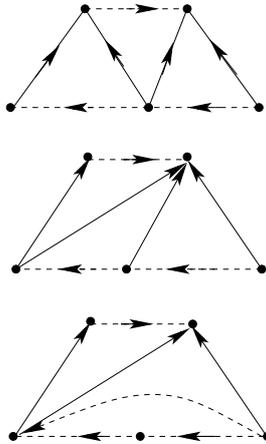


Figure 8: Terminal move.

Proof. Let (T, b) , (T, b') verify the all-bad assumptions. If (T, b) presents a pattern as in the top of Figure 8 (we stipulate that the dashed edges do not belong to $\delta(b, b')$) we can perform the sequence of b -flips suggested by descending the rows of the picture (the corresponding flips on (T, b') are understood). We get (T', b_1) and (T', b_2) such that $|\delta(b_1, b_2)| = |\delta(b, b')| - 1$ and the number of (2)bad $t_1 \cup t_2$ decreases by 1. We stop when we reach a terminal configuration. \square

We can state now the conclusive lemma.

LEMMA 4.8. *Let (T, b) and (T, b') be a terminal all-bad couple. Then we can find a sequence of inversions of ambiguous edges $(T, b) \Rightarrow (T, \tilde{b})$, $(T, b') \Rightarrow (T, \tilde{b}')$ such that $|\delta(\tilde{b}, \tilde{b}')| < |\delta(b, b')|$.*

By iterating all the above procedure starting from (T, \tilde{b}) , (T, \tilde{b}') we eventually get $|\delta| = 0$ and the main theorem follows.

Proof of Lemma 4.8. Let e be an edge not belonging to $\delta(b, b')$ and contained in the star of a disorientated $\bar{e} \in \delta(b, b')$. If e is ambiguous, let us invert it in both (T, b) and (T, b') . If e was in the boundary of a bad nutshell, after the inversion the nutshell becomes good and we can apply Lemma 2.7. If e was in the boundary of a basic (2)bad configuration $(t_1 \cup t_2, b)$, $(t_1 \cup t_2, b')$, then after the inversion, \bar{e} becomes ambiguous (recall Remark 4.5) and we can invert it in (T, b) to decrease $|\delta|$ by 1. So, if e is ambiguous we have done. Assume that e is not ambiguous. Then e is the edge of a (abstract) triangle which is entirely formed by edges not belonging to $\delta(b, b')$. Let v be the vertex of this triangle which does not belong to e . We realize that there is an internal edge

of $St(v, b)$ which is ambiguous. Then by successive inversions of ambiguous edges, we eventually make e ambiguous and we can conclude as above. \square

The **B**-proof in the generic orientable case is complete.

5. The ideal sliding equivalence

Consider again the classification of b -flips in Definition 2.1. Let us call *sliding flip* (s -flip) any b -flip which is not totally ambiguous (sometimes, a totally ambiguous flip is also said a *bump flip*).

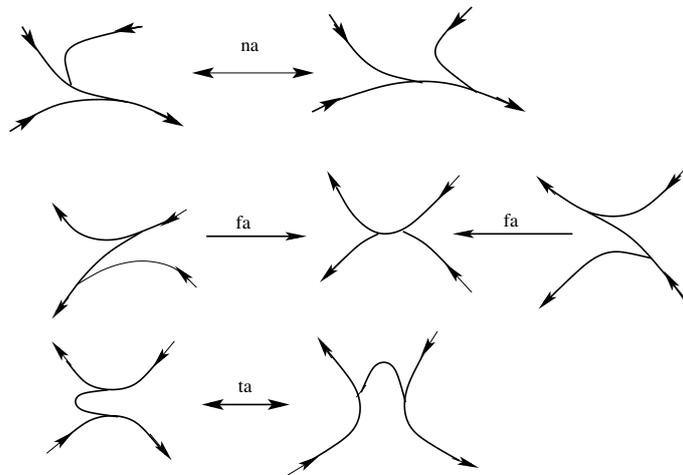


Figure 9: Sliding and bump branched spine flips.

The reason for this terminology is more clear if we look at the picture in terms of dual tracks (Figure 9) corresponding by duality to Figure 2. In the dual picture, instead of the transverse orientations, we prefer to indicate the *local* orientations on the dual tracks, using the counterclockwise planar orientation as reference. If S is oriented, this has a global meaning, as we know that in this case (θ, b) is an oriented train track. Referring to Remark 1.5, we see that the continuous deformations of tracks associated with s -flips are realized by the smooth sliding of track branches. This is not true for bump flips. The sliding transits have been already considered in Section 5 of [1]. After having extended the notion of sliding move to bubble and $1 \rightarrow 3$ stellar moves, in [1] we have been mainly concerned with the completed sliding equivalence, more precisely, with branched triangulations of a given compact closed oriented surface with an arbitrary number of vertices, considered all together.

Here we focus on the ideal setting, we consider also nonorientable surfaces, and we introduce some new constructions (the so-called horizontal foliation). The s -flips together with isotopy relative to V , generate the restricted s -ideal transit equivalence with quotient sets denoted $\mathcal{S}^{id}(S, V)$. We pass from $\mathcal{S}^{id}(S, V)$ to $\mathcal{B}^{id}(S, V)$ adding the bump b -flips to the sliding ones. The basic idea is that every branched triangulation of (S, V) carries some remarkable structures of geometric/topological types that are preserved by the ideal sliding while, accordingly with Theorem 1.2, they can be widely modified by the bump transits.

REMARK 5.1: A version of the sliding equivalence has been widely studied in [13], in a different setting. One considers the set of generic *measured train tracks* (τ, μ) on an orientable surface S , $\chi(S) < 0$. The track τ is not necessarily orientable nor dual to any (naked) triangulation T of S . The measure μ is a function which assigns to each branch e of τ a nonnegative real number $\mu(e)$ and verifies a suitable 3-terms relation at each switch-point of τ . If we forget the orientation of the branches for the sliding flips in Figure 9, we get the generators of the equivalence relation, provided that such (nonoriented) sliding flips are enhanced with the transit of measures. In this setting, a measured flip supported by a nonambiguous flip is usually called a *shift*, a measured flip supported by a forced ambiguous flip is called a *collapse*, the inverse of a collapse is called *splitting*.

5.1. The transverse foliations carried by a branched triangulation

Given (S, V) we recall that S_V denotes the surface with boundary obtained by removing a family of disjoint open 2-disks centred at each $v \in V$. We are going to consider possibly singular foliations on S_V or S . Every such a foliation can be obtained by integration of some field of tangent direction (a tangent vector field if the leaves are oriented). We will say that two foliations are *homotopic (isotopic)* if they are obtained by the integration of homotopic (isotopic) fields. Let (T, b) be a branched triangulation of (S, V) . First, we are going to show that (T, b) carries canonically a pair of regular transverse foliations $(\mathcal{V}, \mathcal{H})$ on S_V , called respectively the *vertical* and the *horizontal* foliations. \mathcal{V} is always oriented. If S is oriented, then also \mathcal{H} can be oriented in such a way that every intersection point has the intersection number equal to 1. \mathcal{V} and \mathcal{H} can be extended to singular foliations \mathfrak{v} and \mathfrak{h} defined on the whole of S . They share the singular set Z which consists of the vertices $v \in V$ where the index $1 - d_b(v) \neq 0$. The two foliations are transverse on $S \setminus Z$. If S is oriented, then both \mathfrak{v} and \mathfrak{h} are oriented; if the singularity indices are all non-positive, so that S is of genus $g \geq 1$, then $(\mathfrak{v}, \mathfrak{h})$ looks like the couple of vertical and horizontal foliations of the square of an Abelian differential on a Riemann surface.

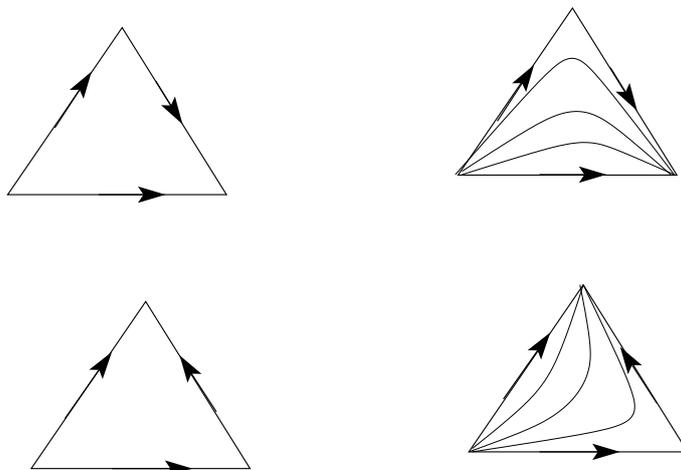


Figure 10: Whitney field tiles

The vertical foliation. Recall a classical way to prove that the Euler-Poincaré characteristic $\chi(S)$ defined through the sum of indices of zeros of any tangent vector fields on S with isolated zeros coincides with its combinatorial definition in terms of any triangulation T of S . Given such a triangulation T , we take its first barycentric subdivision $T^{(1)}$ endowed with a standard Δ -complex structure (i.e. a branching) so that, in particular, each vertex of T is a pit. Then each branched triangle of $T^{(1)}$ carries a so-called *Whitney tangent vector field* which can be defined explicitly in terms of the barycentric coordinates (see [5]). These locally defined vector fields match to define a vector field on the whole of S , with an isolated zero at the barycenter of each iterated face c of T of index $(-1)^{\dim c}$, in such a way that the sum of these indices equals the combinatorial characteristic of T . From this general construction, we retain that every branched triangle (t, b) carries a tile for such a Whitney field. This is illustrated in Figure 10; the two tiles are distinguished from each other when the ambient surface is oriented. Let (T, b) be any branched ideal triangulation of (S, V) . Up to isotopy, the intersection of S_V with every (abstract) triangle t of T is a “truncated triangle” \bar{t} , i.e. a hexagon with 3 internal “long” edges (each one contained in an edge of T) and 3 “short” edges contained in the boundary ∂S_V . The short edges are in bijection with the corners of triangles of T ; some are labelled by 1 like the associated corners. The union of the hexagons forms a cell decomposition of S_V , by the restriction to the long edges of the gluing in pairs of the abstract edges of T . The union of the short edges forms a triangulation of ∂S_V . We endow each hexagon \bar{t} with the oriented foliation

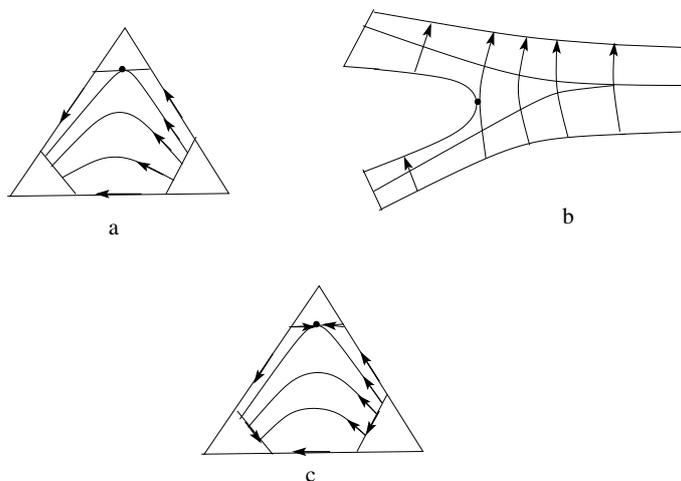


Figure 11: The vertical foliation on S_V .

$\mathcal{V}(\bar{t}, b)$ obtained by the integration of the restriction to \bar{t} of the Whitney field carried by (t, b) . These constitute the tiles of a puzzle that, once assembled, realizes the so-called *vertical foliation* $\mathcal{V} = \mathcal{V}(T, b)$ of S_V carried by (T, b) . A tile $\mathcal{V}(\bar{t}, b)$ is illustrated in Figure 11-a. In the dual viewpoint, recall that the spine (θ, b) of S_V is an embedded transversely oriented train track in S_V ; the foliation \mathcal{V} is positively transverse to it. In Figure 11-b, we see the dual picture corresponding to the puzzle tile. The foliation \mathcal{V} has remarkable properties.

DEFINITION 5.2. *A traversing foliation on S_V is a foliation with oriented leaves such that:*

1. *Every leaf of \mathcal{F} is a closed interval that intersects transversely ∂S_V at its endpoints.*
2. *There is a nonempty finite set of exceptional leaves of \mathcal{F} which are simply tangent to ∂S_V at a finite number of points.*
3. *\mathcal{F} is generic if every exceptional leaf is tangent to the boundary at one point.*

Then it is not hard to see that:

PROPOSITION 5.3. *(1) The vertical foliation \mathcal{V} associated with a branched spine (T, b) of (S, V) is a generic traversing foliation on S_V . The exceptional leaves of \mathcal{V} are in bijection with the 1-labelled short edges of*

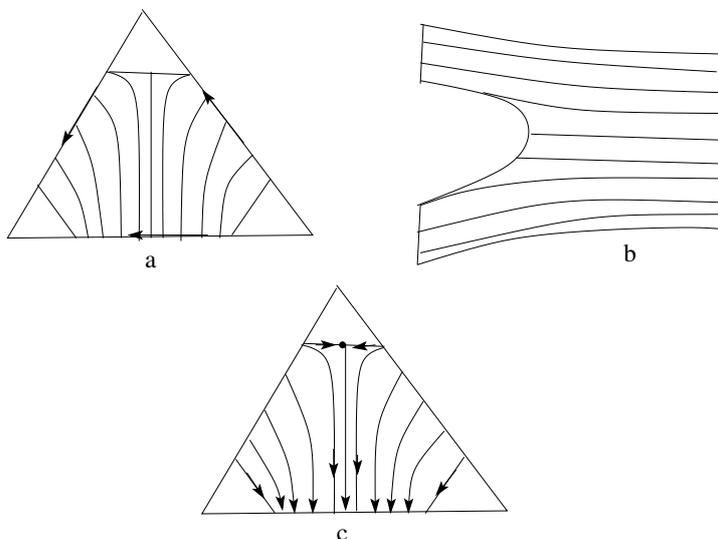
∂S_V ; every exceptional leaf is tangent to the interior of the associated edge. \mathcal{V} is uniquely determined up to isotopy.

- (2) The dual branched spine (θ, b) of S_V intersects transversely all leaves of \mathcal{V} . Every exceptional leaf intersects transversely θ at two points. A leaf passing through a singular point of θ is generic and intersects θ at one point. A generic leaf intersects θ at one or two points. In the second case, it is contained in a quadrilateral in S_V vertically bounded by an exceptional leaf and a leaf passing through a singular point of θ .
- (3) Every generic traversing foliation on S_V can be realized as the vertical foliation of some branched spine of (S, V) .

Boundary bicoloring. Given a traversing foliation \mathcal{F} of S_V , denote by $X = X_{\mathcal{F}} \subset \partial S_V$ the set of tangency points of the exceptional leaves. \mathcal{F} determines a *bicoloring* of the components of $\partial S_V \setminus X$, denoted by $\partial \mathcal{F}$; let us say that a component c is *white* (resp. *black*) if the foliation is ingoing (outgoing) along c . If S is oriented the color can be encoded by an orientation, in the sense that a black component keeps the boundary orientation of ∂S_V (according to the usual rule “first the outgoing normal”), while a white component has the opposite orientation. In Figure 11-c, we see the oriented enhancement of the \mathcal{V} -tile (we stipulate that in the picture the b -orientation of the triangle agrees with the orientation of S_V ; we obtain the picture for the negative branching $-b$ by just inverting all arrows).

The horizontal foliation. Alike \mathcal{V} , we define \mathcal{H} as the result of a puzzle. In Figure 12-a we show the “horizontally” foliated hexagon, in Figure 12-b, the corresponding dual picture. In general, \mathcal{H} is not oriented. If S is oriented then \mathcal{H} is oriented as well; in Figure 12-c, we show the oriented version of the tile. Now we realize that:

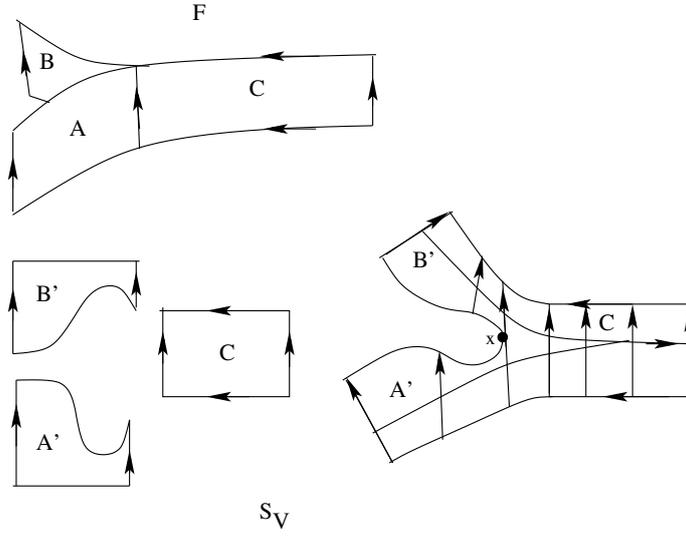
1. \mathcal{V} and \mathcal{H} are transverse foliations.
2. Let Y be the union of the 1-labelled short edges. Then every component of $\partial S_V \setminus Y$ is contained in a leaf of \mathcal{H} , while \mathcal{H} is transverse to the interior of every 1-labelled short edge.
3. If S , hence \mathcal{H} , is oriented then every boundary component in a leaf is oriented like the leaf and is contained in a component of $\partial S_V \setminus X$; these orientations propagate to the whole components of $\partial S_V \setminus X$ and reproduce the bicoloring orientation. \mathcal{V} intersects \mathcal{H} with the intersection number equal to 1 everywhere.
4. The pair $(\mathcal{V}, \mathcal{H})$ is uniquely defined up to isotopy.


 Figure 12: The horizontal foliation on S_V .

If S is oriented, there is another way to realize the foliations $(\mathcal{V}, \mathcal{H})$ by using the 2D case of a result of [8]. This can be described as follows. Take an abstract (non-embedded) copy θ^* of the oriented train track (θ, b) . Consider the oriented branched surface $F := \theta^* \times [-1, 1]$; this carries the vertical foliation \mathcal{V}^* with leaves of the form $\{x\} \times [-1, 1]$ and the horizontal one \mathcal{H}^* with branched leaves of the form $\theta^* \times \{y\}$. Then one can find an embedding of S_V into F which preserves the orientation and such that \mathcal{V} is just the restriction of \mathcal{V}^* to S_V . This is suggested in Figure 13. \mathcal{H}^* restricts as well to a regular foliation of S_V which becomes our final horizontal foliation \mathcal{H} after a suitable homotopy.

Extension to singular foliations. Let (S, V) be as usual. A function $i : V \rightarrow \{n \in \mathbb{Z} | n \leq 1\}$ is said *admissible* if $\chi(S) = \sum_v i(v)$. For every admissible function i , a *vertical foliation of type i* on (S, V) is an oriented singular foliation \mathfrak{v} that verifies by definition the following properties:

1. The singular set Z of \mathfrak{v} consists of the $v \in V$ such that $i(v) \neq 0$.
2. If $i(v) \neq 0, 1$, then the local model of \mathfrak{v} at v is given by the vertical foliation at 0 of the quadratic differential $z^{-2i(v)} dz^2$. If $i(v) = 1$, the local model is given by the integral lines at 0 of the gradient of either the function $|z|^2$ or $-|z|^2$.

Figure 13: Another realization of \mathcal{V} .

A *horizontal foliation of type i* on (S, V) is a nonoriented singular foliation \mathfrak{h} that verifies by definition the following properties:

1. The singular set Z of \mathfrak{h} consists of the $v \in V$ such that $i(v) \neq 0$.
2. If $i(v) \neq 0, 1$, then the local model of \mathfrak{h} at v is given by the horizontal foliation at 0 of the quadratic differential $z^{-2i(v)} dz^2$. If $i(v) = 1$, the local model is given by the level curves at 0 of the function $|z|^2$.

A *transverse pair of foliations of type i* is a pair $(\mathfrak{v}, \mathfrak{h})$ such that

1. \mathfrak{v} and \mathfrak{h} are vertical and horizontal foliations of type i respectively.
2. The two foliations are transverse on $S \setminus Z$ and, case by case, the above local models at the singular points hold simultaneously for both \mathfrak{v} and \mathfrak{h} .
3. If S is oriented, we require furthermore that also \mathfrak{h} is oriented in such a way that \mathfrak{v} and \mathfrak{h} intersect everywhere with intersection number equal to 1.

LEMMA 5.4. *For every admissible function i there are transverse pairs of foliations of type i on (S, V) .*

Proof. It is enough to prove that there exists a vertical foliation of type \mathfrak{i} , for we can take as transverse horizontal foliation the orthogonal one for a suitable auxiliary Riemannian metric on S . Let S_Z be the surface with boundary obtained by removing from S a small 2-disk around every $v \in Z$. Consider the foliation on a neighbourhood of ∂S_Z determined by the \mathfrak{i} -local model at singular points. By a simple variation of Hopf index theorem, we realize that it extends to the whole of S_Z without introducing new singularities. \square

Such transverse pairs of type \mathfrak{i} are considered up homotopy through transverse pairs of type \mathfrak{i} which is locally an isotopy at the singular points. We denote by $\mathcal{TP}(S, V, \mathfrak{i})$ the so obtained quotient set. Set $\mathcal{TP}(S, V) = \cup_{\mathfrak{i}} \mathcal{TP}(S, V, \mathfrak{i})$. Finally, denote by $\mathfrak{T}(S_V)$ the quotient set of the set of generic traversing foliations on S_V considered up to homotopy through traversing (not necessarily generic) foliations.

The following theorem summarizes the main features of the s -transit equivalence.

THEOREM 5.5. *For every (S, V) :*

- (1) *The correspondence $(T, b) \rightarrow \mathcal{V}(T, b)$ induces a well-defined bijection*

$$\tau : \mathcal{S}^{id}(S, V) \rightarrow \mathfrak{T}(S_V) .$$

- (2) *For every (T, b) , consider the admissible function $\mathfrak{i}_b(v) = 1 - d_b(v)$. Then the associated pair of transverse foliations $(\mathcal{V}(T, b), \mathcal{H}(T, b))$ on S_V extends to a transverse pair $(\mathfrak{v}(T, b), \mathfrak{h}(T, b))$ of type \mathfrak{i}_b on (S, V) in such a way that this induces a well-defined map*

$$\mathfrak{p} : \mathcal{S}^{id}(S, V) \rightarrow \mathcal{TP}(S, V) .$$

- (3) *Fix a base point $v_0 \in V$. Assume that the set of admissible functions such that $\mathfrak{i}(v_0) = 0$ is nonempty and denote by $\mathcal{TP}_0(S, V)$ the corresponding subset of $\mathcal{TP}(S, V)$. Then the set of triangulations (T, b) of (S, V) such that $\mathfrak{i}_b(v_0) = 0$ is nonempty and, denoting $\mathcal{S}_0^{id}(S, V)$ the corresponding subset of $\mathcal{S}^{id}(S, V)$, we have that the restricted map $\mathfrak{p} : \mathcal{S}_0^{id}(S, V) \rightarrow \mathcal{TP}_0(S, V)$ is bijective.*

Proof. The fact that the map τ in (1) is well defined and onto follows just by looking at the sliding flips and from (3) of Proposition 5.3.

Once the extension $(\mathcal{V}(T, b), \mathcal{H}(T, b)) \rightarrow (\mathfrak{v}(T, b), \mathfrak{h}(T, b))$ will be established just below, that the map \mathfrak{p} of (2) is well defined follows from the fact that τ is well defined.

The fact that τ in (1) is injective as well as item (3) are simpler 2D versions of results established in [3] and [4] for branched spines of 3-manifolds with

nonempty boundary. In [3], we essentially considered the case of closed manifolds, that is when the boundary consists of one 2-sphere. In [4], we faced the general case with minor changes. We limit here to illustrate the main points, referring to the harder proofs in 3D.

The injectivity of τ is the 2D analogous of Theorem 4.3.3 of [3]. By transversality, we can assume that the homotopy is generic, that is, it contains only a finite number of nongeneric traversing foliations, each one containing one exceptional leaf which is tangent at two points of ∂S_V . Then we analyze how two generic traversing foliations close to a nongeneric one are related to each other and we realize that the sliding b -flips cover all possible configurations.

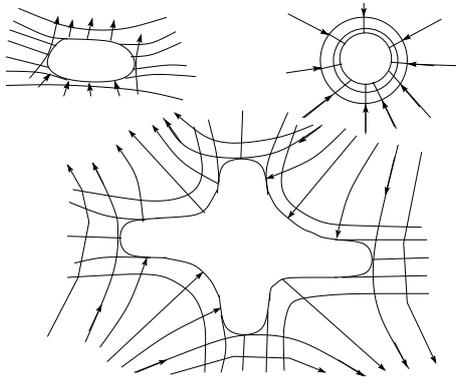


Figure 14: Extendible configurations at ∂S_V

As for the extension of $(\mathcal{V}, \mathcal{H})$, we look at the configuration of this foliations at each component C of ∂S_V which is also the boundary of a small disk D in S centred at one of the vertex $v \in V$ (see Figure 14). If $i_b(v) = 0$, there are exactly two exceptional leaves of \mathcal{V} tangent to C . Then we can extend \mathcal{V} without singularities through D , respecting the bicoloring of C , and manage similarly for \mathcal{H} . In the other cases, we easily realize that, up to isotopy, the configuration of $(\mathcal{V}, \mathcal{H})$ at C is the restriction of the configuration of the local models carried by D . The arbitrary choices in implementing the construction are immaterial as the $(\mathfrak{v}, \mathfrak{h})$ are considered up to kind of homotopy stated above. Note that the extension of \mathcal{V} can be obtained by using the non-truncated tiles of Figure 10.

Item (3) is more demanding. Notice that if $\chi(S) = 0$, then the hypothesis can be satisfied also when $|V| = 1$, in all other cases necessarily $|V| \geq 2$. First, we prove at the same time that $\mathcal{S}_0^{id}(S, V)$ is nonempty and the restricted map $\mathfrak{p} : \mathcal{S}_0^{id}(S, V) \rightarrow \mathcal{TP}_0(S, V)$ is onto. The key point is to prove that every vertical foliation \mathfrak{v} occurring in $\mathcal{TP}_0(S, V)$ is realized by a triangulation (T, b) with the

given distribution of $d_b(v)$'s. This is the 2D counterpart of Proposition 5.1.1 of [3]. The proof is based on Ishii's notion of *flow spines* [11]. Let $A := V \setminus \{v_0\}$ and define S_A as usual, so that $S_V \subset S_A$. In general, \mathfrak{v} is not traversing S_A . The idea is to find an embedded 2-disk D in the interior of S_A , centred at some point v'_0 , such that \mathfrak{v} becomes traversing $S_{V'}$ and generic, where $V' = \{v'_0\} \cup A$, and such that only two exceptional leaves are tangent to ∂D . This is carried by a branched triangulation of (S, V') . Finally, via the homogeneity of S , we get the desired triangulation (T, b) of S_V . The injectivity of the restricted map \mathfrak{p} is the counterpart of Theorem 5.2.1 of [3]. We can 'cover' any homotopy connecting $\mathfrak{v}(T, b)$ with $\mathfrak{v}(T', b')$ with a chain of flow-spines connecting (T, b) with (T', b') in such a way that the traversing foliation associated to one is homotopic through traversing foliation to the traversing foliations of the subsequent. \square

REMARK 5.6: Every $\mathcal{TP}(S, V, \mathfrak{i})$ is an affine space over $H_1(S_Z; \mathbb{Z})$, Z being the singular set prescribed by \mathfrak{i} . So in general $\mathfrak{p} : \mathcal{S}_0^{id}(S, V) \rightarrow \mathcal{TP}_0(S, V)$ is a bijection between *infinite sets*.

5.2. On the nonambiguous transit

The nonambiguous b -flips define, in the usual way, a so-called *ideal na-transit equivalence* with associated quotient sets $\mathcal{NA}^{id}(S, V)$, endowed with natural surjective projections $\mathcal{NA}^{id}(S, V) \rightarrow \mathcal{S}^{id}(S, V)$. In 3D, the notion of nonambiguous structure (defined indeed through the transit of *pre-branchings* rather than of branchings) supports nontrivial examples of intrinsic interest and interesting applications to quantum hyperbolic invariants (see [1], [2]). The intrinsic content of this 2D *na*-equivalence is not so evident. We limit to a few remarks assuming furthermore that S is oriented. We refer to the notations introduced at the end of Section 2.1.

EXAMPLE 5.7: Let S be the torus and $|V| = 1$. Let T be the triangulation of (S, V) as in Figure 3-e. We check by direct inspection that for every branching (T, b) , there is not any edge of T supporting a nonambiguous flip. This holds for *every* triangulation T' of (S, V) because all these triangulations are equivalent to each other up to diffeomorphism of (S, V) . Hence, in this case, the *na*-transit equivalence is nothing else than the identity relation. On the other hand, we check that the branchings on T which share the same decomposition $S = S_+ \cup S_-$ are not *s*-equivalent to each other.

QUESTION 5.8: Are branched triangulations (T, b) and (T', b') of (S, V) *na*-equivalent if and only if they are *s*-equivalent and share the decomposition $S = S_+ \cup S_-$?

REFERENCES

- [1] S. BASEILHAC AND R. BENEDETTI, *Non-ambiguous structures on 3-manifolds and quantum symmetry defects*, Quantum Topology **8** (2017), 749–846.
- [2] S. BASEILHAC AND R. BENEDETTI, *On the quantum Teichmüller invariants of fibred cusped 3-manifolds*, Geometriae Dedicata **197** (2018), 1–32.
- [3] R. BENEDETTI AND C. PETRONIO, *Branched standard spines of 3-manifolds*, Lect. Notes Math. 1653, Springer (1997).
- [4] R. BENEDETTI AND C. PETRONIO, *Reidemeister-Turaev torsion of 3-dimensional Euler structures with simple boundary tangency and pseudo-Legendrian knots*, Manuscripta Math. **106** (2001), 13–61.
- [5] S. HALPERIN AND D. TOLEDO, *Stiefel-Whitney homology classes*, Ann. of Math. (2) **96** (1972), 511–525.
- [6] V. DISARLO AND H. PARLIER, *The geometry of flip graphs and mapping class groups*, Trans. Amer. Math. Soc. **372** (2019), no. 6, 3809–3844.
- [7] V. DISARLO AND H. PARLIER, *Simultaneous flips on triangulated surfaces*, Michigan Math. J. **67** (2018), no. 3, 451–464.
- [8] D. GILLMAN AND D. ROLFSEN, *The Zeeman conjecture for standard spines is equivalent to the Poincaré conjecture*, Topology **22** (1983), no. 3, 315–323.
- [9] A. HATCHER, *Algebraic Topology*, Cambridge University Press, 2001.
- [10] A. HATCHER, *On triangulations of surfaces*, Topology Appl. **40** (1991), no. 2, 189–194.
- [11] I. ISHII, *Flows and spines*, Tokyo J. Math. **9** (1986) 505–525.
- [12] M. LACKENBY, *Taut ideal triangulations of 3-manifolds*, Geom. Topol. **4** (2000) 369–395.
- [13] R. C. PENNER WITH J. L. HARER, *Combinatorics of train tracks*, Annals of Mathematics Studies, vol. 125, Princeton University Press, Princeton, NJ, 1992.

Author’s address:

Riccardo Benedetti
Università degli Studi di Pisa
Dipartimento di Matematica
Largo B. Pontecorvo 5
56127 Pisa, Italy
E-mail: riccardo.benedetti@unipi.it

Received March 3, 2020

Accepted May 26, 2020