Order automatic mapping class groups

Colin Rourke
Bert Wiest

Mathematics Institute
University of Warwick
Coventry, CV4 7AL, UK

Email: `cpr@maths.warwick.ac.uk` and `bertold@maths.warwick.ac.uk`

Abstract

We prove that the mapping class group of a compact surface with a finite number of punctures and non-empty boundary is order automatic. More precisely, the group is right-orderable, has an automatic structure as described by Mosher, and there exists a finite state automaton that decides, given the Mosher normal forms of two elements of the group, which of them represents the larger element of the group. Moreover, the decision takes linear time in the length of the normal forms.

AMS Classification numbers Primary: 20F60, 06F15, 20F10

Secondary: 57M07, 57R50

Keywords: mapping class group, right invariant order, automatic structure, Mosher normal form, right-orderable group, order automatic group

0 Introduction

This paper is a sequel to [5] and contains the proof of the linear time algorithm which was announced in [5, remark 5.2]. In addition to giving the full proofs, this paper extends the results of [5] to a considerably larger class of mapping class groups.

Let S be a compact surface, not necessarily orientable, with boundary ∂S and a finite number of distinguished points in its interior, called the *punctures*. The *mapping class group* $\mathcal{MCG}(S)$ is defined to be the group of homeomorphisms of S mapping ∂S identically and permuting the punctures, up to isotopies fixing ∂S and the punctures. The group multiplication is given by composition and written *algebraically*, ie given Φ, $\Psi \in \mathcal{MCG}(S)$ define $\Phi\Psi := \Psi \circ \Phi \colon S \to S$. For instance, if S is a disk with n punctures then $\mathcal{MCG}(S)$ is isomorphic to B_n, the braid group on n strings [1].

In [2, 3] P Dehornoy defined a total ordering on the braid group B_n which is *right invariant*, ie for β, γ, $\delta \in B_n$, $\beta > \gamma$ implies $\beta\delta > \gamma\delta$. This ordering was reinterpreted in [5] in more geometrical terms.

Right invariant orderings are important because of a long-standing conjecture in group theory: that the group ring of a torsion-free group has no zero divisors. The conjecture is true for the smaller class of right-orderable groups (groups admitting a right invariant order). For more detail see [8, chapter 2].

The mapping class group of a surface S is torsion-free if and only if ∂S is non-empty.[1] In section 1 we generalise the construction in [5] to prove that for any surface S with non-empty boundary, $\mathcal{MCG}(S)$ is right-orderable.[2]

In [6, 7] L Mosher constructed an automatic structure (in the sense of [4]) for mapping class groups of surfaces. More precisely, suppose that a surface S with boundary and possibly some punctures is equipped with a triangulation whose vertex set consists of the set of punctures and some points in the boundary, and whose edges are ordered and oriented. Then Mosher associates with every

[1] This result appears to have the status of a “folk-lore” theorem. Although clearly known for a long time and implicit in early work on mapping class groups, we have failed to locate an explicit reference.

[2] There are two well-established conventions for products in mapping class groups, the *algebraic* convention used here and the *functional* convention $\Phi\Psi := \Phi \circ \Psi$. Using the functional convention, the order constructed in this paper is left invariant. However there is no real distinction between left and right invariant orders because any group admitting a left invariant order $\sim$ also admits a right invariant order by $g < h \iff h^{-1} \sim g^{-1}$ and vice-versa.

element Φ of $\mathcal{MCG}(S)$ a unique finite sequence of pairs of triangulations of S, which we call the *automatic normal form* of Φ. In section 2 and 3 we construct a finite state automaton that detects the order from the automatic normal forms, ie given the automatic normal forms of two elements, the automaton decides which element is larger. The resulting algorithm takes linear time in the length of the normal forms; indeed the only unbounded part of the algorithm is reading the sequences to detect the first discrepancy. If the two sequences of triangulations first differ in the ith term, then the order can be read from the ith to $i+3$rd elements of the sequences. Section 2 contains the proof of a weaker result — that the algorithm takes quadratic time in the length of the normal forms and section 3 contains the technical details which sharpen this to linear time.

It is worth noting that our results extend to subgroups of mapping class groups defined by restricting the allowed permutation of the punctures to lie in some subgroup of the permutation group. It is also worth observing that there is no greater generality in considering homeomorphisms which are the identity on just some of the boundary components. This is because a boundary component which is not fixed is essentially the same as another puncture.

1 A right invariant order on $\mathcal{MCG}(S)$

In this section we construct a right invariant order on $\mathcal{MCG}(S)$. The construction involves a lot of choices, and there are many other such orderings apart from the ones we exhibit. The material in this section is a straightforward generalisation of sections 1–3 of [5]. We shall leave out some details, and refer the reader to [5] instead.

Ideal arc systems and their reductions

Suppose S is a compact surface with non-empty boundary and a finite number of punctures (possibly none). The set of punctures is denoted p. Following [6] we define an *ideal arc* to be the image h of a map

$$(I, \partial I, int\, I) \to (S, \partial S \cup p, S - (\partial S \cup p))$$

which is injective in $int\, I$. Two ideal arcs are *isotopic* if there exists an isotopy of S fixed on $\partial S \cup p$ deforming one into the other. An *ideal arc system* Γ is a collection of non-isotopic ideal arcs with disjoint interiors, such that each component of $S - \Gamma$ is a disk.

Two ideal arc systems can be *reduced with respect to each other* or *pulled tight*. We shall sketch this, details can be found in [7] or [5].

Suppose Γ and Δ are two ideal arc systems. We define Γ and Δ to be *transverse* if every arc of Γ either coincides precisely with some arc of Δ, or intersects the arcs of Δ transversely.

Now suppose Γ and Δ are transverse. A D–disk between Γ and Δ is a subset of S homeomorphic to a closed disk which contains no punctures in its interior, and which is bounded by one segment of an ideal arc of Γ and one of Δ; the two intersection points of the two ideal arcs may be in the interior or in the boundary of the arcs.

We say two ideal arc systems Γ_0, Γ_1 are *equivalent* with respect to Δ if there is an isotopy of S that fixes $\partial S \cup p$, leaves the subset $\Delta \subset S$ invariant, and carries Γ_0 to Γ_1.

For proofs of the following two fundamental results see [7] or [5]:

Proposition 1.1 *Any two ideal arc systems Γ, Δ can be reduced with respect to each other, ie after an appropriate isotopy of Γ there are no D–disks between Γ and Δ. Moreover, this reduction process gives a unique result: if two curve diagrams Γ_0, Γ_1 are both transverse to Δ and reduced with respect to Δ, then they are equivalent with respect to Δ.* □

Proposition 1.2 (Triple reduction lemma) *Suppose Γ, Δ and Σ are three ideal arc systems such that Γ and Δ are both reduced with respect to Σ. Then there exists an isotopy between Γ and an ideal arc system Γ', which is an equivalence with respect to Σ, such that Γ', Δ and Σ are pairwise reduced.* □

Curve diagrams

Next we construct special examples of ideal arc systems, with some additional labelling. We shall call any diagram on S which can be obtained in this way a *curve diagram* (as in [5]).

First we construct an ideal arc system Γ of S with the following properties: all ideal arcs are embedded and disjoint (even in their endpoints), all endpoints of ideal arcs lie in ∂S, and $S - \Gamma$ consists of precisely one disk, possibly punctured. If Γ has say k ideal arcs, then we label these arcs $\Gamma_1, \ldots, \Gamma_k$, in any order. Moreover, we equip each of them with an orientation. Next, if S has n punctures, then we label these punctures $1, \ldots, n$, in any order, and add to the diagram Γ an additional n embedded ideal arcs with disjoint interiors as follows. We have one oriented ideal arc from ∂S to the first puncture, labelled Γ_{k+1}, disjoint from all the previous arcs; then one oriented ideal arc from the first to the second puncture, and so on up to the arc Γ_{k+n}, from the $n-1$st to the nth puncture.

This finishes the construction of a general curve diagram. For an example see figure 1. Here S is a torus with one disk removed and three punctures (so $k=2$, $n=3$).

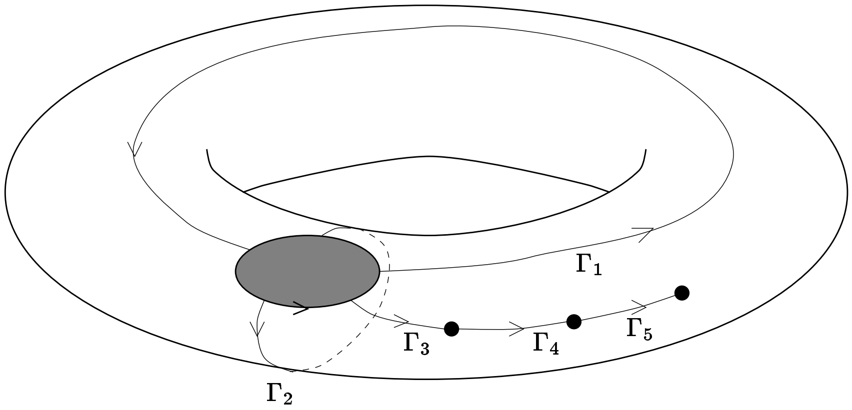

Figure 1: A curve diagram on a torus with one boundary component and 3 punctures

We now fix, once and for all, a curve diagram E on S, which we call the *basic curve diagram*. The choice of E is arbitrary, but it will serve as a 'basepoint' for the constructions that follow. We also choose, once and for all, an orientation for ∂S. (Recall that we are not assuming that S is itself orientable. Even if S is orientable, nothing that follows depends on a choice of orientation for S.)

The right invariant ordering of $\mathcal{MCG}(S)$

Curve diagrams are the main building block in the construction of a right invariant ordering of $\mathcal{MCG}(S)$. If Φ is an element of $\mathcal{MCG}(S)$, then applying Φ to E yields another curve diagram $\Phi(E)$. Note that $\Phi(E)$ is only determined up to isotopy.

Given $\Phi,\ \Psi \in \mathcal{MCG}(S)$, we can assume that the curve diagrams $\Phi(E)$ and $\Psi(E)$ are reduced with respect to each other (recall that $\Phi(E)$ and $\Psi(E)$ have a unique reduced form with respect to each other, by Proposition 1.1). Suppose that the ideal arcs $\Phi(E_l)$ and $\Psi(E_l)$ coincide for $l = 1, \ldots, i-1$, where $1 \leqslant$

$i \leqslant k+n$, and that $\Phi(E_i)$ and $\Psi(E_i)$ are the first non-coincident pair. We note that $\Phi(E_i)$ and $\Psi(E_i)$ start at the same point of ∂S or puncture, are oriented, and that both $\Phi(E)$ and $\Psi(E)$ are locally separating. So it makes sense to define $\Phi > \Psi$ if $\Phi(E_i)$ branches off $\Psi(E_i)$ to the left, ie if an initial segment of $\Phi(E_i)$ lies to the left of $\Psi(E_i)$. Here we are using the chosen orientation of ∂S to define "left" and "right". Left means in the direction of the orientation of ∂S. Figure 2 illustrates this.

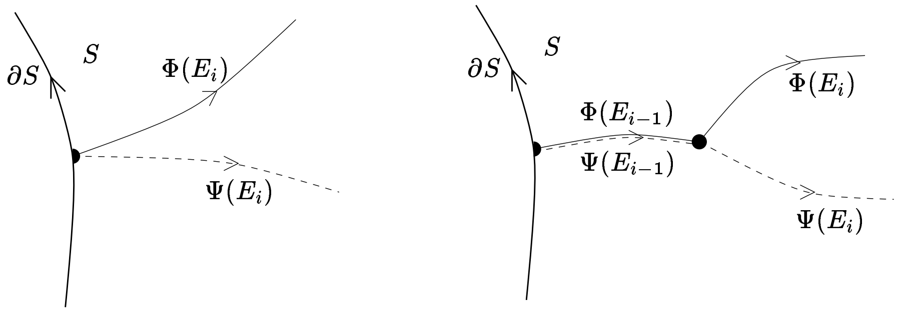

Figure 2: $i \in \{1, \ldots, k+1\}$ $\qquad$ $i = k+2$

The transitivity of this relation follows from the triple reduction lemma: Suppose that $\Phi, \Psi, \Upsilon \in \mathcal{MCG}(S)$, then by the triple reduction lemma 1.2 we can assume that the three curve diagrams $\Phi(E)$, $\Psi(E)$, $\Upsilon(E)$ are pairwise reduced. Then if $\Upsilon(E)$ branches off $\Phi(E)$ to the left, and $\Phi(E)$ branches off $\Psi(E)$ to the left, then $\Upsilon(E)$ branches off $\Psi(E)$ to the left. So '$<$' is indeed an order.

The fact that this order is right invariant is also immediate from the definitions: if $\Phi, \Psi, \Upsilon \in \mathcal{MCG}(S)$, then in order to compare $\Phi\Upsilon$ with $\Psi\Upsilon$, we have to apply a homeomorphism of S representing Υ to the two diagrams $\Phi(E)$ and $\Psi(E)$. This leaves the diagrams reduced with respect to each other, and therefore preserves the order.

2 Automatic ordering

In this section and the next we prove that the mapping class group $\mathcal{MCG}(S)$ is order automatic. More precisely, we construct an algorithm that is executable by a finite state automaton, has as its input the automatic normal forms (in the sense of Mosher [6, 7]) of two elements of $\mathcal{MCG}(S)$, and as its output the decision which of the two elements is larger. Furthermore, we prove that the algorithm takes linear time in the length of the automatic normal forms. A weak algorithm (taking quadratic time) is constructed in this section. The next section contains the necessary technical detail to sharpen this to linear time.

Review of Mosher's normal form for elements of $\mathcal{MCG}(S)$

An *ideal triangulation* of S is a triangulation whose edges are ideal arcs and segments of ∂S, and whose vertices are punctures and endpoints of ideal arcs in ∂S. Moreover, the non-boundary edges are oriented, and numbered $1, \ldots, e$, where e is the number of non-boundary edges. We consider two ideal triangulations to be equivalent if they differ by a boundary- and puncture-fixing isotopy. We say two ideal triangulations are in the same *triangulation class* if they are combinatorially equivalent relative to the boundary, or equivalently, if they differ by the action of an element $\mathcal{MCG}(S)$. Note that the edges and the triangles of the triangulations being considered are allowed to be *bent*, ie edges may have both ends coincident and triangles may have one or more edges or vertices coincident.

Suppose now that S is equipped with an ideal triangulation B, which we call the *base triangulation.* We consider only triangulations which agree with B on ∂S. We now consider the groupoid $\mathcal{G}$ which has for objects the set of triangulation classes and for morphisms the set of ordered pairs (T, T') of ideal triangulations, where (T, T') is identified with $(h(T), h(T'))$ if $h \in \mathcal{MCG}(S)$. The morphism goes from the class of T to the class of T'. If T and T' are in the same class, then there is a unique boundary fixing isomorphism from T' to T up to isotopy, ie an element of $\mathcal{MCG}(S)$. This determines an isomorphism between the vertex group of $\mathcal{G}$ and the mapping class group $\mathcal{MCG}(S)$ with multiplication now written *functionally* ie $\Phi\Psi := \Phi \circ \Psi$.

We consider a particular type of morphism in $\mathcal{G}$.

Definition *Flipping an edge* Let T be an ideal triangulation and α an edge of T adjacent to two triangles δ and δ'. The triangulation T' is obtained by removing α, so that δ and δ' combine to form a quadrilateral of which α is a diagonal, and then cutting the quadrilateral back into two triangles by inserting the opposite diagonal. We call the morphism (T, T') "flipping α" and denote it f_α, see figure 3.

Every morphism (B, T) in $\mathcal{G}$ from the base vertex to another vertex is a product of a canonical sequence of flips. To see this, picture (B, T) as given by superimposing B and T, and *comb* T along B. To be precise, first reduce T with respect to B and then consider edge 1 of B. Suppose that the first edge of T that edge 1 crosses is α. Flip α as in figure 3, where edge 1 is drawn dashed,

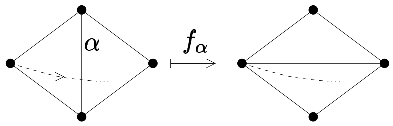

Figure 3: Flipping an edge

and reduce. Repeat this process until there are no more crossings of edge 1 with T. The fact that this process is finite follows by counting the number of interections of edge 1 with T (excluding the next-to-be-flipped edge). For more detail here see [7, pages 321–322]. Now do the same for edge 2 and continue in this way, using the chosen ordering and orientation of edges of B, until T has been converted into a copy of B.

The *Mosher normal form* of (B, T) is the inverse of the sequence of flips described above.[3] Notice that unlike the general case described in [7] there is no relabelling morphism required here. (This is because ∂S is fixed throughout.) Mosher proves that this normal form defines an automatic structure on $\mathcal{G}$ and hence, using results from [4], on the vertex group $\mathcal{MCG}(S)$.

Now consider an element (B, T) of the vertex group at the class of B. As remarked earlier (B, T) determines an element $\Phi \in \mathcal{MCG}(S)$ unique up to isotopy such that $\Phi(T)$ is equivalent to B. (Conversely given $\Phi \in \mathcal{MCG}(S)$ the corresponding triangulation pair is $(B, \Phi^{-1}B)$.)

We observe that if we comb B along $\Phi(B)$ this is combinatorially identical to combing $\Phi^{-1}(B)$ along B and we call the sequence of flips defined by this combing the *combing sequence of* Φ. (The reverse of the combing sequence is the Mosher normal form of Φ.)

The base triangulation

For our purposes we need a base triangulation B with certain special properties. We shall construct B from the basic curve diagram E in two steps. In the first step we construct an ideal arc system which has three edges $B_{3i-2}, B_{3i-1}, B_{3i}$

[3]Strictly speaking the Mosher normal form is not this flip sequence, which only defines an asynchronous automatic structure, but is derived from it by clumping flips together into blocks called "Dehn twists", "partial Dehn twists" and "dead ends" (see [7] pages 342 et seq). This technicality does not affect any of the results proved here. We prove that order can be detected in linear time from the flip sequence. Since the clumped flip sequence can be unclumped in linear time, this implies that order can be detected in linear time from the strict Mosher normal form.

for every arc E_i of E. This step needs care; by contrast, the second step is quite unsubtle: we simply add some more edges, to turn the arc system into an ideal triangulation.

For example, figure 4 shows a base triangulation obtained from the basic curve diagram in figure 1. Here the edges constructed in step 1 are drawn with solid lines and the edges constructed in step 2 with dashed lines.

Step 1 We recall that E has k ideal arcs $E_1, \ldots, E_k$ with endpoints in ∂S, and n arcs $E_{k+1}, \ldots, E_{k+n}$ with a puncture at at least one of their endpoints. We define $B_{3i-2} = E_i$ for $i = 1, \ldots, k+n$. Here B_{3i-2} also carries the same orientation as E_i and we label the starting point of E_i as V_i for future reference.

For $i = 1, \ldots, k+1$ we choose two points of ∂S on either side of V_i and closer to V_i than any other boundary vetex. We define the edge B_{3i-1} to start at the finishing point of E_i, run 'on the left of' and close to E_i, but with opposite orientation, and end on ∂S, near the starting point of E_i, at one of the points just chosen. The edge B_{3i} is defined to sit on the other side of $E_i = B_{3i-2}$, in a similar fashion. So $B_{3i-2}, B_{3i-1}, B_{3i}$ and a segment of ∂S together bound two 2–simplices, with common edge B_{3i-2}. We call this ensemble Q of two 2–simplices with edges $B_{3i-2}, \ldots, B_{3i}$ the ith *beak* of B.

The edge B_{3i-1} for $i = k+2, \ldots, k+n$ is defined to start at the $i-k$th puncture, run close to B_{3i-2} and B_{3i-4}, and to end at the same point of ∂S as B_{3k+2}. So the edge B_{3i-1} sits 'to the left of B_{3i-2}'. The edge B_{3i} for $i = k+2, \ldots, k+n$ are defined similarly, but 'on the right of B_{3i-2}'. So the edges $B_{3i-4}, \ldots, B_{3i}$ together bound two 2–simplices, and again we call this ensemble the ith *beak*.

Let Q_i be the ith beak, $i \in \{1, \ldots, k+n\}$. We call the edge B_{3i-2} the *leading edge* of Q_i, and B_{3i-1} and B_{3i} the *left* and *right outlying edges* of Q_i. We also call the starting point V_i of the leading edge the *principal vertex* of Q_i. For example, in figure 4, the 5th beak Q_5 is shaded and has principal vertex V_5, leading edge 13 and outlying edges 14 and 15.

Step 2 In general, the ideal arc system defined in step 1 is not an ideal triangulation of S, because some of the 2–cells are not triangles. We now add a finite number of ideal arcs so as to create an ideal triangulation. Orientation or labelling of these extra edges will not play any role in the proof, and we shall think of them as unoriented and unlabelled.

Detecting order from the Mosher normal form

We are now ready to state the main result of this section.

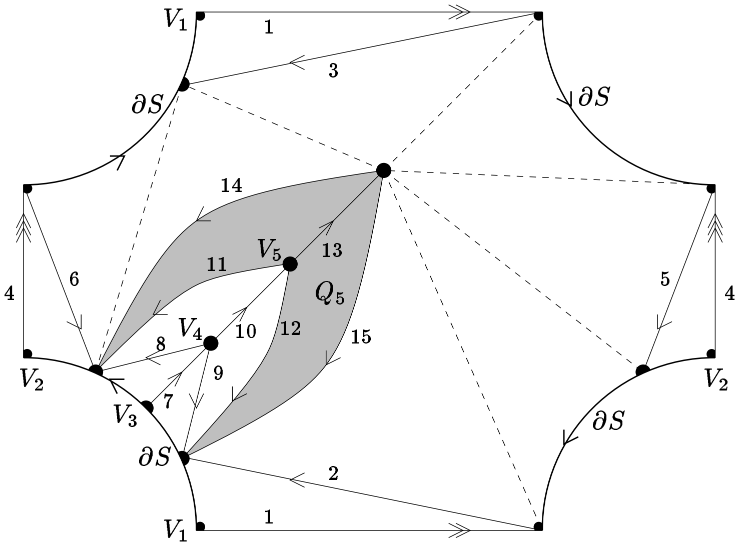

Figure 4: A base triangulation obtained from the basic curve diagram in figure 1

Theorem 2.1 (Main Theorem) *There exists an algorithm that has as its input the combing sequences of two elements Φ, Ψ of $\mathcal{MCG}(S)$, and as its output the decision which of the two elements is larger. Furthermore, if the two combing sequences of Φ, Ψ first differ in the ith term, then the order can be read by inspecting the ith to $i+3$rd terms of the sequences and consulting a finite catalogue of possibilities.*

Notice that the theorem implies that decision takes linear time in the length of the shortest combing sequence, indeed the only unbounded part of the process is the inspection of the two sequences to locate the first discrepancy.

Furthermore the decision can clearly be carried out by a finite-state automaton which has the necessary catalogue built in. Thus we have:

Corollary 2.2 *The mapping class group $\mathcal{MCG}(S)$ is order automatic.*

Before starting on the detailed proof of the theorem it is worth explaining the strategy. The relative order of two elements Φ, Ψ of $\mathcal{MCG}(S)$ is determined by the images of the curve diagrams, which are part of the base triangulation. The combing sequences are obtained by combing along the images of the base triangulation so it is plausible that the position of the curve diagram up to a certain point can be reconstructed from the combing sequence up to the corresponding point. However there are problems. Combing along a given edge of the curve diagram may not cause any flips in the combing sequence for the simple reason that the edge is already part of the triangulation at that point. But notice that we have chosen the triangulation so that each edge in the curve diagram is the leading edge of a beak and the corresponding outlying edges can be used to determine the position of the leading edge in this case. There are further complications due to the fact that, in the middle of the sequence, the triangulation which is being combed may well have bent edges or triangles. If a bent edge is flipped it is often ambiguous where the edge which caused the flip is located. The whole process can be likened to observing a bubble chamber (the bubbles are the flips) and trying to work out the position of the particles (edges which we are combing along) which caused the bubbles.

Remark 2.3 There is a very short proof of a somewhat weaker result, along the lines of the proof given in [5]. By inspecting the combing sequence we can immediately read whether an element of $\mathcal{MCG}(S)$ is *positive* (ie greater than the identity element) or *negative*. More precisely, define $\Phi \in \mathcal{MCG}(S)$ to be i–positive (respectively i–negative) if the first $i-1$ curves of $\Phi(E)$ coincide (after reduction) with the first $i-1$ of E and that $\Phi(E_i)$ branches off E_i to the left (respectively right). Now the curve diagram $\Phi(E)$ of Φ is part of the triangulation $\Phi(B)$ namely the edges numbered $1, 4, 7, \ldots$. Suppose that Φ is i–positive, then it can be assumed to fix the first $i-1$ of these edges (ie $1, 4, \ldots, 3i-5$) and, after reduction, the corresponding outlying edges (ie $2, 5, \ldots, 3i-4$ and $3, 6, \ldots, 3i-3$). But edge $3i-2$ is carried to the left and must meet edge $3i-1$ of B. Thus the first flip in the combing sequence of Φ is f_{3i-1}, ie a flip of the edge numbered $3i-1$. Similarly if Φ is i–negative then the first flip in the combing sequence is f_{3i}. We have proved the following:

Algorithm 2.4 (To decide from the Mosher normal form whether an element of $\mathcal{MCG}(S)$ is i–positive or negative and provide the correct value of i) *Inspect the combing sequence (the reverse of the Mosher normal form). The first flip is either f_{3i-1} or f_{3i} for some i. In the first case the element is i–positive and in the second it is i–negative.*

From general properties of automatic structures (see [4, Theorem 2.3.10]), the combing sequence of ab^{-1} can be computed in quadratic time from the combing sequences of a and b, hence algorithm 2.4 gives a quadratic time algorithm to determine order and implies corollary 2.2. The final section of the paper comprises the technical detail necessary to sharpen this into a linear time algorithm.

3 Proof of the main theorem

We start by observing that in order to detect order we only need to examine the combing of the beak part (ie the first $3(k+n)$ arcs of the triangulation created in 'step 1' above). If combing along the beak part of $\Phi(B)$ and $\Psi(B)$ gives the same result, then the curve diagrams $\Phi(E)$ and $\Psi(E)$ agree, so that $\Phi = \Psi \in \mathcal{MCG}(S)$, and the complete combing sequences of Φ and Ψ coincide. So in order to compare Φ and Ψ, we need not comb along the edges created in 'step 2' above (and this is the reason why their labelling and orientation is irrelevant).

Now consider a definite combing sequence $f_1, f_2, \dots f_i, f_{i+1}, f_{i+2}, \dots$ and consider the triangulation T at the start of f_i. Assume that f_i comes from combing along a particular beak Q. The principal vertex, leading edge etc of Q will be called *the* principal vertex etc at that stage.

As the combing process proceeds we construct a sequence of beaks which coincide with beaks of $\Phi(B)$. These beaks are never touched again throughout the combing process and we call the union of the beaks constructed thus far the *stable set*, denoted U.

Observation 3.1 At any stage in the combing process we can identify the stable set and hence the principal vertex. Indeed the next edge to be flipped is always in the *star* of V (ie the collection of triangles of which V is a vertex, together with their faces), where V is the principal vertex. So we look for the maximal set of beaks of the triangulation T at that stage which contain vertices $V_1, V_2, \dots, V_{j-1}$ such that the next flip is not in the star of V_i for $i < j$. This is the stable set and $V = V_j$.

The next flip is caused by combing along some beak Q. The leading edge q of Q is either (1) already an edge of T (in this case we are combing along one of the outlying edges of $\Phi(B)$) or (2) crosses a triangle A of T from the leading vertex V to the opposite edge (which is the next edge flipped). If we can determine q in case (1) or A in case (2), then we say that we have *located* the leading edge q.

Lemma 3.2 (Main technical lemma) *We can locate q by inspecting T and the four flips $f_i, f_{i+1}, f_{i+2}, f_{i+3}$ and consulting a finite catalogue of possibilities.*

The main theorem follows from the lemma as follows.

If the combing sequences of Φ and Ψ first differ in the ith term then the leading edges at stage i must differ. This is because the outlying edges are determined by the leading edge (they follow the leading edge back and out as indicated in figure 5). Thus if the leading edges of the two beaks coincide then so do the entire beaks and combing along the current beaks would not cause a discrepancy.

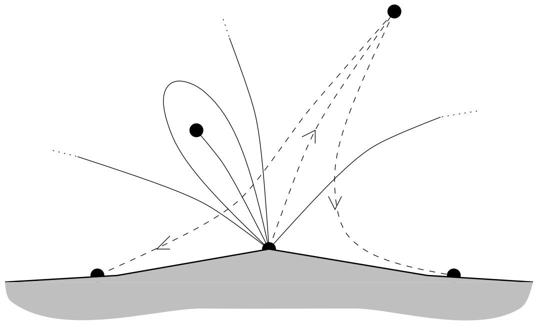

Figure 5: The outlying edges are close to the leading edge

Now suppose that the two principal vertices differ. For definiteness suppose that Φ and Ψ have principal vertices V_k and V_l respectively, with $k < l$ (note that k and l can be determined by observation 3.1). Then T contains a beak with principal vertex V_k and leading edge q say which is already part of $\Psi(B)$ but not of $\Phi(B)$. Thus there is only one edge of T not in the stable set of Φ emanating from V_k, and this must be q. Further the leading edge for Φ must be different. Thus by locating the latter we can read off the order.

Now suppose that the two principal vertices are the same. Then observe that the two leading edges cannot both cross the same triangle (or else the next flip f_i would be the same in both sequences). Thus by locating the two leading edges we can read off the relative order and the theorem is proved.

Proof of the main technical lemma

We shall how to locate the leading edge from T and the four flips. It will be clear that the decision is equivalent to consulting a finite catalogue. There are two possibilities. We can be combing along the leading edge or along one of the outlying edges. We suppose first that we know which of these possibilities holds and we deal with them in cases A and B. In case C we deal with the general case, which includes the possibility that we may not immediately know whether we are in cases A or B.

Case A : combing along the leading edge

By observation 3.1 we can read off the principal vertex, V say. Since we are combing along the leading edge, the first flip f_i is the flip of an edge α such that there is a triangle A with vertex V and opposite edge α. If there is only one such triangle, then the leading edge must cross this triangle and we have located it. However there are a number of situations where this triangle is not unique, the simplest of which is illustrated in figure 6. (In this and later diagrams the letter 'X' indicates that the triangulation is continued in some way in the marked region.)

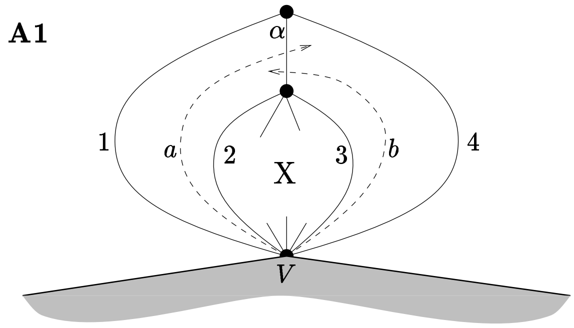

Figure 6: Case A1 : standard ambiguity

There are two possible positions for the leading edge, labelled a and b in the figure. But if the leading edge is in position a then the second flip f_{i+1} must be edge 3 or 4 and in position b it must be 1 or 2. Thus we can resolve the ambiguity and locate the leading edge by inspecting the second flip.

There is a degenerate version of case A1, when edges 2 and 3 coincide, case A2 in figure 7. In this case edge 2 cannot be the second flip for either possible position of the leading edge, because otherwise the leading edge would self-intersect; thus if the leading edge is in position a then the second flip must be edge 4, and in position b it must be 1.

Next come three bent versions, cases A3, A4 and A5 (figure 8). In case A3 the inner end of α from case A1 coincides with V and in case A4 the outer end. In case A5 the outer end of α from case A2 coincides with V. There is no case where the inner end from case A2 coincides with V since this would produce a 2–gon.

In cases A3 and A4 the ambiguity is resolved exactly as in case A1, and in case A5 the ambiguity is resolved exactly as in case A2.

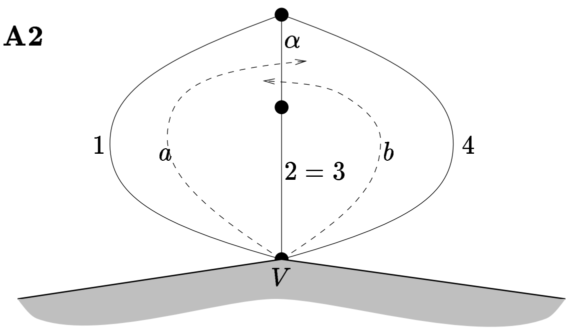

Figure7: Case A2 : degenerate version of case A1

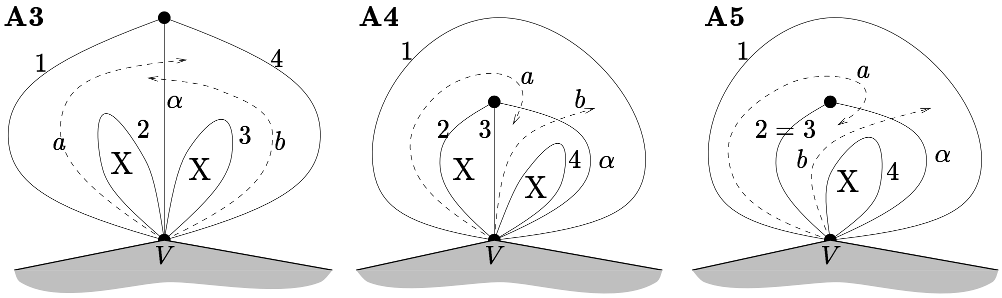

Figure 8: Cases A3, A4 and A5 : bent version of A1 and A2

In case A6 (figure 9) both ends of α from case A1 have been deformed to V (and the ambiguity is again resolved as in case A1) and finally to complete the catalogue of cases where two different triangles have the same vertex V and base α we have included case A7 (figure 9). There is no position for the leading edge in case A7 so this case cannot occur; (the "two different" triangles are in fact the same triangle in two different guises).

Case B : combing along an outlying edge

We use the same notation as in case A: the principal vertex is V and the first flip f_i is the flip of an edge α. This time however α must have one endpoint at V (cf figure 5). First assume that α is straight (ie not bent). Case B1 (figure 10) is the unambiguous case. This is the case where there are either $\geqslant 2$ or no edges with vertex V trapped to the left of α between α and the stable set. The leading edge must be the next edge to the right of α. Case B2 (figure 10) is the ambiguous case. This is the case where there is precisely one edge (labelled a) with vertex V trapped to the left of α and the leading edge could be this

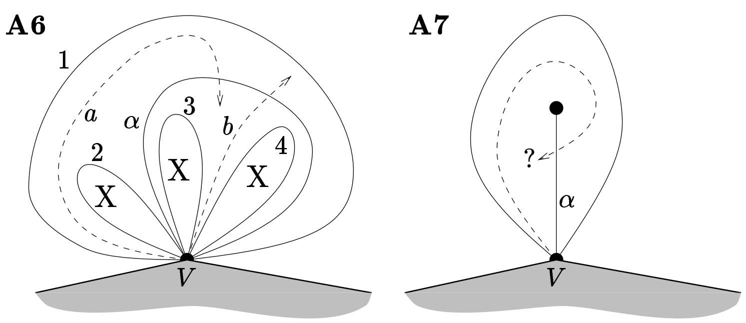

Figure 9: Cases A6 and A7 : very bent versions

edge (α is flipped by the right outlying edge) or the edge labelled b to the right of α, if such an edge exists (if not this case is also unambiguous: the leading edge is a). However if b is the leading edge then the second flip is edge a (and vice-versa); thus the second flip resolves the ambiguity.

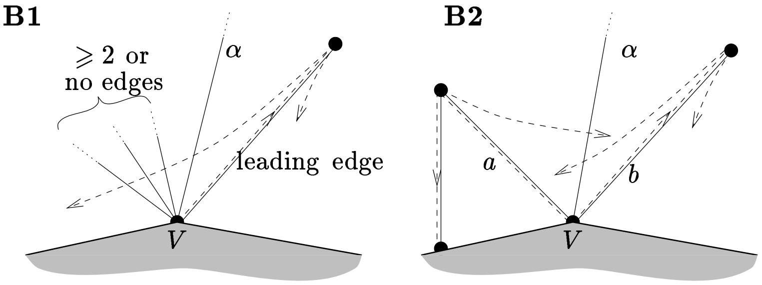

Figure 10: Cases B1 and B2 : α straight

Next assume that α is bent. In this case (figure 11) there are three possible positions for the leading edge, labelled a, b and c in the figure (case B3) — note that a is possible only if there is precisely one edge to the left of α, and that b and c are the next edges to the right of the two branches of α, assuming that c exists. Assume first that there is at least one edge to the left of α. If the leading edge is in positions a or c, then the second flip is of an edge inside α, whilst in position b the second flip is outside. Thus the second flip distinguishes position b. Finally positions a and c are distinguished by the second flip (inside α — the flip is next to the left side for a and next to the right side for c) except in the special case B4. Here there is just one edge b inside α, which is the second flip for both positions a and c; however in this case the third flip f_{i+2} distingushes the cases (it is edge a for leading edge in position c and vice-versa).

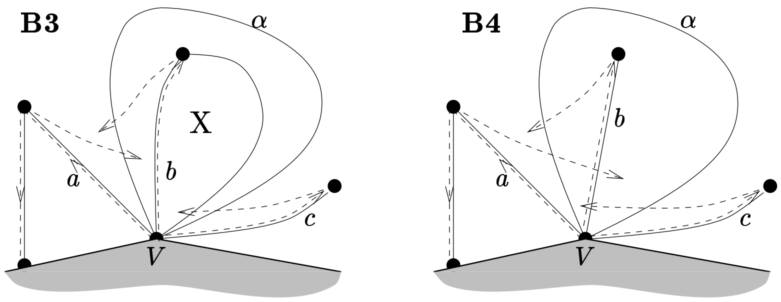

Figure 11: Cases B3 and B4 : α bent, at least one edge to the left of α

Now assume that there are no edges to the left of α (cases B5, B6 and B7 in figure 12). The possible positions for the leading edge are b and c (assuming c exists). Case B5 is the case when there are at least two edges inside α to the right of b. In this case the second flip distingushes b and c; it is the edge closest to b (edge 1) if the leading edge is in position b and the edge closest to α (edge 2) for c. Cases B6 and B7 show the special cases when there are one or no edges (respectively) inside α to the right of b. In case B6 the third flip distingushes (it is b for position c and vice-versa) and in case B7 the second flip distingushes (it is again b for position c and vice-versa).

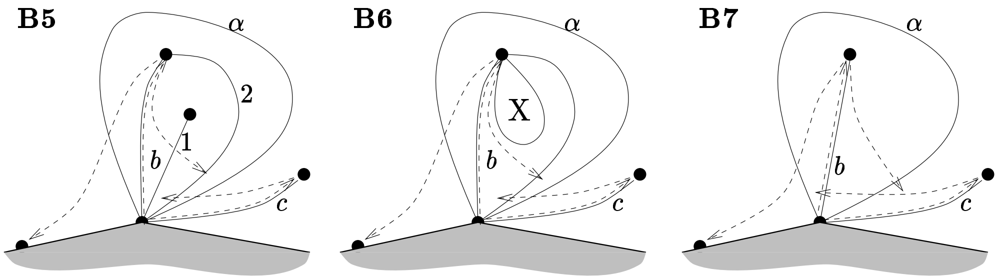

Figure 12: Cases B5, B6 and B7 : α bent, no edges to the left of α

Case C : the general case

Finally we turn to the general case. If the first flip f_i is caused by combing along the leading edge, then the edge α which is flipped is an edge of a triangle A with opposite vertex the principal vertex V. If f_i is caused by combing along one of the outlying edges then α has one endpoint at V. Thus unless α satisfies both these conditions we can immediately decide whether we are in case A or B. However, if the triangle A is appropriately bent, it is possible for

α to satisfy both conditions. The three possible shapes are illustated in figure 13.

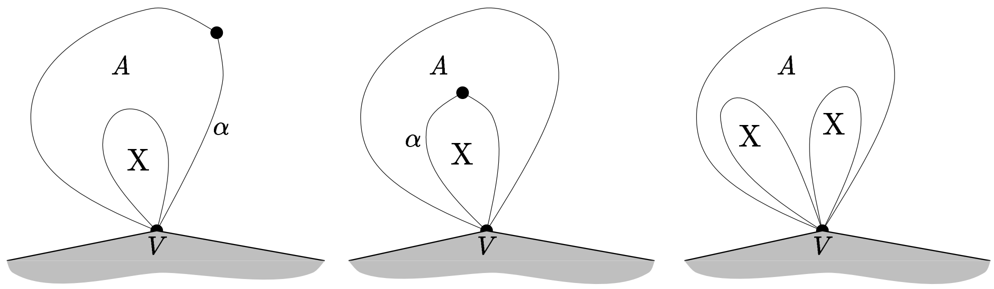

Figure 13: Case C : Three possible shapes for A

One of the edges of A other than α must be bent. The possibilities are that α is outside this bent edge (left in figure 13), inside (middle) or also bent (right — in this case α can be any of the three edges of A). Before discussing each of these cases in detail it is worth remarking that the ambiguous cases for combing the leading edge (cases A1 to A7) cannot occur here because in these cases it can be checked that there is no position for the leading edge so that an outlying edge will flip α. Thus we never meet the case A ambiguity together with ambiguity between cases A and B.

We now discuss each of the three shapes in figure 13 in turn.

Figure 14 shows the two possible orientations for the case when α is outside the bent edge of A. In case C1 the possible positions for the leading edge are a and b and the second flip distinguishes them (it is edge 2 for position b but not for a). In case C2 the two possible positions for the leading edge are again a and b and the second flip for a is edge 1. This may also be the second flip

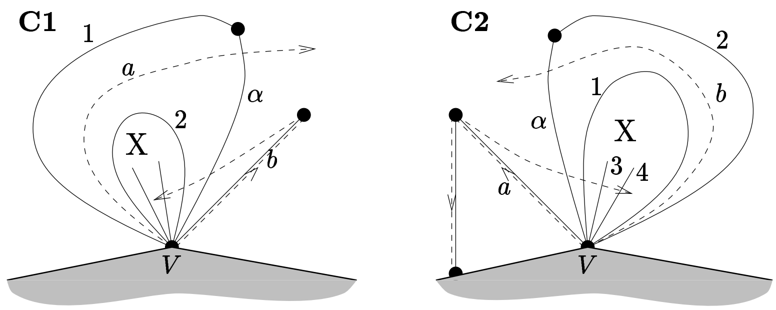

Figure 14: Cases C1 and C2 : α outside

for position b — in the case that b does not cross any more edges of T, but in this case the third flip distinguishes the two posibilities. It is the leftmost edge inside 1 (labelled 3) for a and the rightmost (labelled 4) for b. There is one special case (like case B4 above) when 3 and 4 coincide and this is case C3 shown in figure 15.

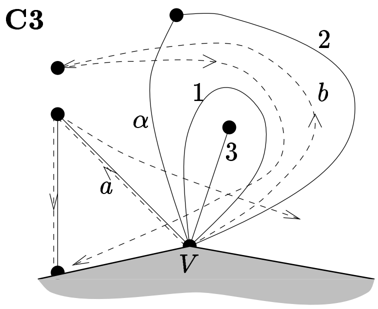

Figure 15: Case C3 : the worst case

In case C3 the leading edge in position b does not cross any more edges of T and the first three flips for position a and b coincide. However the fourth flip distinguishes the cases — it is a for position b and 2 for position a. This case (and the bent version, case C9 below) is the only case when it is necessary to inspect the fourth flip f_{i+3}.

Figure 16 shows the only possible orientation for the case when α is inside the bent edge of A (the middle case in figure 13). Case C4 shows the general case, and case C5 the special case when the leading edge in position a coincides with edge 2. In both cases the second flip distinguishes — it is 1 for leading edge in position a, but not in position b.

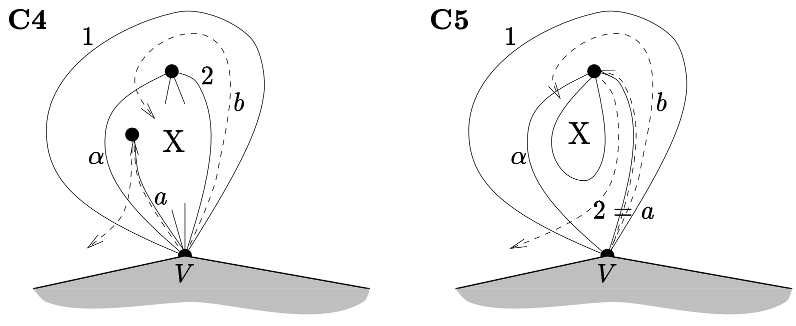

Figure 16: Cases C4 and C5 : α inside

Figure 17 shows the two possibilities for the totally bent case (the right hand case of figure 13) when α is inside. In both cases the second flip distinguishes the possible positions for the leading edge — in case C6 it is 1 for position a but not for b, and in case C7 it is 2 for position b but not for a.

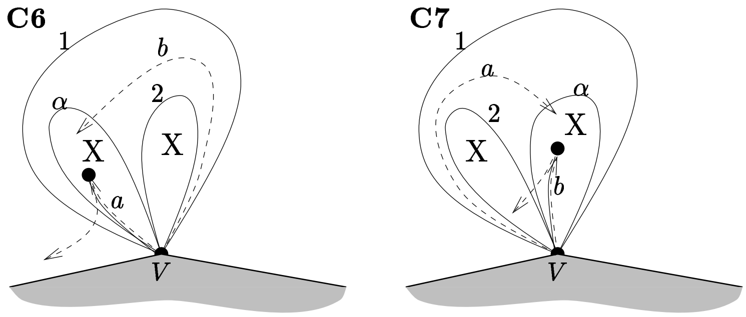

Figure 17: Cases C6 and C7 : α bent and inside

Finally case C8, figure 18, shows the last possibility for the totally bent case — when α is outside. In this case there are three possible positions for the leading edge (labelled a, b and c). The second flip detects c (the flip is 2 for c but not for a and b) and finally distinguishing between a and b is exactly the same as in case C2 above (indeed the figure is obtained from C2 by deforming the top end of α round to the right to coincide with V). In particular there is a deformed copy of the worst case C3 (drawn as case C9) with identical analysis.

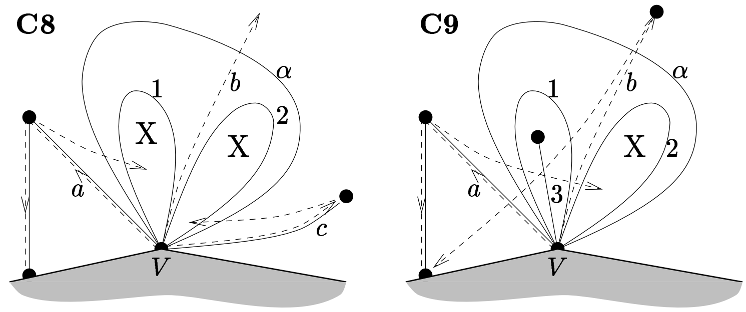

Figure 18: Cases C8 and C9 : α bent and outside

References

[1] **J Birman,** *Braids, links, and mapping class groups*, Annals of Math. Studies, 82, Princeton University Press, Princeton (1975)

[2] **P Dehornoy,** *Braid groups and left distributive operations*, Trans. AMS 343 (1994) 115–150

[3] **P Dehornoy,** *A fast method of comparing braids*, Adv. in Math. 125 (1997) 200–235

[4] **D B A Epstein et al,** *Word processing in groups*, Jones & Bartlett (1992)

[5] **R Fenn, M T Greene, D Rolfsen, C Rourke, B Wiest,** *Ordering the braid groups*, Warwick preprint (1998)

[6] **L Mosher,** *Mapping class groups are automatic*, Math. Research Letters 1 (1994) 249–255

[7] **L Mosher,** *Mapping class groups are automatic*, Annals of Math. 142 (1995) 303–384

[8] **D S Passman,** *Algebraic structure of group rings*, Wiley (1977)