A PRIMER ON HANDLEBODY GROUPS

SEBASTIAN HENSEL

ABSTRACT. We survey properties of the *handlebody group* – the mapping class group of a 3–dimensional handlebody. The text is aimed at readers who are familiar with the modern theory of surface mapping class groups, and focuses on similarities and differences of handlebody groups to surface mapping class groups and $\mathrm{Out}(F_n)$. We also briefly discuss topological applications and open questions.

CONTENTS

Date: July 24, 2018.

1. Introduction

A *handlebody* V_g of genus g is the three-manifold with boundary obtained from the 3–ball by attaching g one-handles (in any way). The *handlebody group* $\mathcal{H}_g$ is the mapping class group of the three-manifold V_g, i.e. the group of orientation preserving self-homeomorphisms of V_g up to isotopy. Alternatively, $\mathcal{H}_g$ can be identified with a subgroup of the mapping class group of the boundary surface of V_g.

This survey is concerned with the group theory, dynamics and geometry of $\mathcal{H}_g$. The group $\mathcal{H}_g$ has important applications in topology and therefore has been studied in some depth in the latter half of the twentieth century, in particular in the context of Heegaard splittings (see e.g. [31, 47] for good sources of the classical theory).

On the other hand, the theory of mapping class groups $\mathrm{Mcg}(\Sigma_g)$ of surfaces has recently become a widely known topic in low-dimensional topology and geometric group theory – to no small extent due to the excellent book [10]. In addition, many new and powerful tools to study surface mapping class groups have been developed over the last years which allow to answer geometric and algebraic questions on surface mapping class groups $\mathrm{Mcg}(\Sigma_g)$.

This leads to a specific audience which this survey is aimed at: readers who are familiar with the modern theory of surface mapping class groups, but not its three-dimensional cousin. We will discuss similarities and differences both in results and methods – and aim to advertise the study of $\mathcal{H}_g$ as a natural "next step" in the study of mapping class groups.

Three themes will recur throughout the text. The first is the *membership* problem, i.e. how to detect if a mapping class of a surface Σ_g is an element of, or conjugate into, the handlebody group. This question is particularly relevant for topological applications of $\mathcal{H}_g$.

The second theme is *(intrinsic) similarity*: many group-theoretic properties are shared between surface mapping class groups and handlebody groups, even though the underlying reasons may be different.

The third and possibly most intriguing theme is *(extrinsic) incompatibility*: in almost any sense the structure of handlebody subgroup is not compatible with its ambient mapping class group. It is a "very small" subgroup which behaves differently both from a geometric or dynamic perspective.

During this survey we will develop tools and results from the ground up. We cannot make the text completely self-contained, but we will give detail and (sketches of) proofs whenever possible. We only assume basic low-dimensional topology and familiarity with surface mapping class groups as discussed in [10]. In fact, we will roughly follow the sequence of topics covered in [10] (with some handlebody-specific interludes), and arrange topics as in a textbook, discovering and proving results as methods become available. This is done for the benefit of a reader who wants to discover and learn about the structure of the handlebody group leisurely and with motivating examples. For readers who want to quickly get an overview of the most important results first, we summarise them in the next section.

1.1. **A User's Guide to the Handlebody Group.** We can define the handlebody group $\mathcal{H}_g$ in two equivalent ways: as the group of isotopy classes of orientation preserving self-homeomorphisms of a 3–dimensional handlebody V_g of genus g, or as the subgroup of the mapping class group of a surface Σ_g of genus g formed by those elements which extend to a (fixed) handlebody with boundary Σ_g (compare Lemma 3.1 and the discussion following it). In other words, restriction to the boundary induces an inclusion

$$\mathrm{Mcg}(V_g) \to \mathrm{Mcg}(\Sigma_g)$$

whose image is exactly $\mathcal{H}_g$ (see Section 3 for details). This "surface point of view" is the main tool to study $\mathcal{H}_g$. From this perspective, the handlebody group $\mathcal{H}_g$ is an infinite index subgroup of $\mathrm{Mcg}(\Sigma_g)$, which is not normal (Corollary 5.4).

There are a few immediate ways to characterise when an element $\phi \in \mathrm{Mcg}(\Sigma_g)$ extends to the chosen handlebody V_g (and therefore defines an element in $\mathcal{H}_g$). All of these criteria eventually rely on the basic idea that homeomorphisms of V_g preserve the set of those essential curves on ∂V_g which bound disks in V_g (the *meridians*). Namely, the following are equivalent (Corollary 5.11):

i) $\phi \in \mathcal{H}_g$.
ii) For every meridian δ, the image $\phi(\delta)$ is also a meridian.
iii) For some nonseparating system $\alpha_1, \dots, \alpha_g$ of meridians, the image $\phi(\alpha_1), \dots, \phi(\alpha_g)$ also consists of meridians.
iv) The induced map $\phi_* : \pi_1(\Sigma_g) \to \pi_1(\Sigma_g)$ preserves the kernel

$$\ker(\pi_1(\Sigma_g) \to \pi_1(V_g))$$

For topological applications it is more natural to ask if ϕ is conjugate into $\mathcal{H}_g$ – or, equivalently, if ϕ extends to *some* handlebody, rather than the specific one we chose. While testing membership in $\mathcal{H}_g$ is fairly straightforward using i)-iv), the conjugacy problem is much more intricate. Intuitively, the problem is that for this question we do not know in advance which curves are the meridians.

Nevertheless, testing if ϕ is conjugate into $\mathcal{H}_g$ can be decided algorithmically (see the end of Section 5) and there are obstructions and criteria for specific types of elements.

Concerning membership and conjugacy, we will cover the following results:

- A Dehn twist T_α is contained in $\mathcal{H}_g$ if and only if α is a meridian (Theorem 5.6).
- A multitwist is contained in $\mathcal{H}_g$ exactly if it is (in the obvious way) a product of Dehn twists about meridians and embedded annuli (Theorem 5.6, compare also Example 5.1 for the definition of annulus twist).

– On the other hand, any multitwist is conjugate into $\mathcal{H}_g$ (Lemma 5.7).
– A torsion element (or finite subgroup) is contained in $\mathcal{H}_g$ exactly if it preserves a collection of disjoint disks cutting V_g into balls. It is conjugate into $\mathcal{H}_g$ if and only if it preserves a multicurve with complementary regions which are genus 0 (Lemma 8.4).
– If a pseudo-Anosov is contained in $\mathcal{H}_g$, then its stable and unstable laminations are limits of meridians (in $\mathcal{PML}$). The converse is not true (but partial results are possible, see Theorem 9.3).

We warn the reader that, in general, reducible elements which are contained in $\mathcal{H}_g$ can have reducing systems which do not consist of meridians, and therefore the irreducible components may not even be conjugate into a smaller handlebody group (compare Example 5.8). Hence, an intrinsic theory of irreducible handlebody mapping classes will in general be very different from the Nielsen-Thurston classification of the ambient mapping class group (see e.g. [51] for such an approach). In this survey, *reducible* will always mean that the surface mapping class group preserves a multicurve.

Criterion iv) above immediately implies that under the standard action of $\mathrm{Mcg}(\Sigma_g)$ on $H_1(\Sigma_g;\mathbb{Z})$ the handlebody group $\mathcal{H}_g$ preserves the Lagrangian subspace

$$L = \ker(H_1(\Sigma_g;\mathbb{Z}) \to H_1(V_g;\mathbb{Z})).$$

In fact, any such matrix in $\mathrm{Sp}(2g,\mathbb{Z})$ is actually induced by an element of $\mathcal{H}_g$ (Theorem 7.1).

On the other hand, in the Torelli group (and in fact any term of the Johnson filtration) there are elements which do not extend to any handlebody – so homology alone is unable to characterise $\mathcal{H}_g$ (compare Theorem 7.2).

In addition to the surface mapping class group, the handlebody group is also intimately connected to another favourite of geometric group theorists: the outer automorphism group of free groups. Namely, the fundamental group of V_g is free, and we thus obtain a surjection (compare Theorem 6.3)

$$\mathcal{H}_g \to \mathrm{Out}(F_g).$$

The kernel of this map is generated by Dehn twists about meridians (Theorem 6.4) and is not finitely generated (Theorem 6.5). In general it is hard (but not impossible; see e.g. Lemma 6.7) to relate properties of handlebody mapping classes and outer automorphisms. Understanding this relation more fully might be interesting in the study of $\mathrm{Out}(F_g)$ as well.

The intrinsic algebraic structure of handlebody groups is similar to that of surface mapping class groups. A few important results are:

– $\mathcal{H}_g$ is finitely presented (Theorem 8.1).
– There is a generating set consisting of mapping classes supported in annuli and pairs of annuli (Corollary 6.6).
– Free Abelian subgroups have ranks up to at most $3g-3$ (just as for the surface mapping class group).

– It is virtually torsion free (being a subgroup of the mapping class group), and its virtual cohomological dimension is $4g - 5$ (Theorem 8.6), just as for the full mapping class group.
– The first Betti number of $\mathcal{H}_g$ is zero, but there is always 2–torsion in the first homology (Theorem 8.5).
– $\mathcal{H}_g$ enjoys homological stability properties akin to those of mapping class groups [19].

From a geometric perspective however, the handlebody group is much less understood. The surjection to $\mathrm{Out}(F_g)$ has infinitely generated kernel (Theorem 6.5) and therefore does not seem to be amenable to methods from geometric group theory. The inclusion into surface mapping class groups is exponentially distorted (Theorem 10.5) – so the ambient geometry also does not directly imply anything about the geometry of $\mathcal{H}_g$. Initial results indicate that $\mathcal{H}_g$ resembles $\mathrm{Out}(F_g)$ geometrically much more than $\mathrm{Mcg}(\Sigma_g)$, making the word geometry of $\mathcal{H}_g$ seem a fruitful, if ambitious, target for geometric group theory. We will discuss some questions and results in Section 10.

1.2. **Limitations of this survey.** As with any survey, the choice of topics and perspective are influenced by the author's point of view, and there are always some important developments which cannot be touched upon. In the case of this article, we chose only to very briefly touch upon *applications* of handlebody groups (maybe most importantly their connections to Heegaard splittings), partly since there exist already many good resources covering these topics, but more so because we want to advertise the study of handlebody groups on their own, as a variant of mapping class groups and outer automorphisms of free groups.

Two further important topics which did not fit into the main flow of the discussion are an intrinsic Nielsen-Thurston theory for handlebody mapping classes (see e.g. [51, 39]) and Hain's algebraic perspective [14].

Most likely there are more accidental omissions – but we hope that our choice of topics is a useful introduction anyway.

Acknowledgements I would like to thank Federica Fanoni, Benson Farb, Dan Margalit, Thomas Ng, Bram Petri, José Pedro Quintanilha and Katie Vokes for interest and many, very helpful comments on drafts of this survey. I would also like to thank Susumu Hirose for pointing out a mistake in a previous version and helpful references.

2. Basic Handlebody Notions

This section develops some basic notions on handlebodies and their boundaries which will feature throughout. We will not review the theory of simple closed curves on surfaces (see e.g. [10] for that), but only discuss the connection to handlebodies. We make the convention that all *curves* on surfaces are assumed to be simple, closed and essential unless explicitly stated otherwise.

We will usually consider curves and their isotopy classes interchangeably, and indicate when it is important to choose a specific representative. We do not develop the tools of three-dimensional (differential) topology from the ground up here, compare e.g. [30] for background. We will however highlight differences to the two-dimensional situation in this section.

By a *handlebody V_g of genus g* we will always mean the oriented 3–manifold with boundary obtained from a 3–ball by attaching g one-handles[1]. The boundary ∂V_g of a handlebody is homeomorphic to a surface genus g. A *spotted handlebody* is a handlebody V_g together with a disjoint collection of embedded disks $D_i \subset \partial V_g$. Slightly abusing notation, we define the boundary surface of a spotted handlebody to be the complement of the interior of the spots on ∂V_g (so that the boundary of a spotted handlebody is a surface Σ_g^n with some number n of boundary components). The importance of this notion comes from the fact that cutting a handlebody V_g at a system of disjoint disks naturally yields a spotted handlebody Y, and the boundary surface of Y (with the above definition) then corresponds to a subsurface of ∂V_g.

A curve α on ∂V_g is a *meridian*, if it bounds an embedded disk in V_g. By Dehn's lemma, this is equivalent to asserting that α is a simple closed curve which defines an element of the kernel

$$K = \ker(\pi_1(\partial V_g, p) \to \pi_1(V_g, p))$$

of the map induced by inclusion of the boundary.

The intersection pattern between meridians is much more constrained than the one between simple closed curves in general. To begin with, we note the following observation, which follows from naturality of the intersection form.

Lemma 2.1. *The algebraic intersection number between any two meridians is zero. In particular, meridians always intersect in an even number of points.*

If α and α' are two meridians (or, in fact, systems of disjoint meridians), then we can find disks D, D' bounded by α, α' which intersect minimally in the following sense.

Lemma 2.2 (compare e.g. [45, Lemma 5.1])**.** *If D, D' are two properly embedded disks in V_g, then they can be isotoped to intersect in $\frac{1}{2}i(\partial D, \partial D')$ arcs. In particular, the boundary meridians intersect minimally.*

Using the fact that handlebodies are aspherical and a standard surgery argument, we also have the following uniqueness statement.

Lemma 2.3. *Any two embedded disks bounded by a meridian are isotopic.*

[1]In particular, all of our handlebodies are oriented.

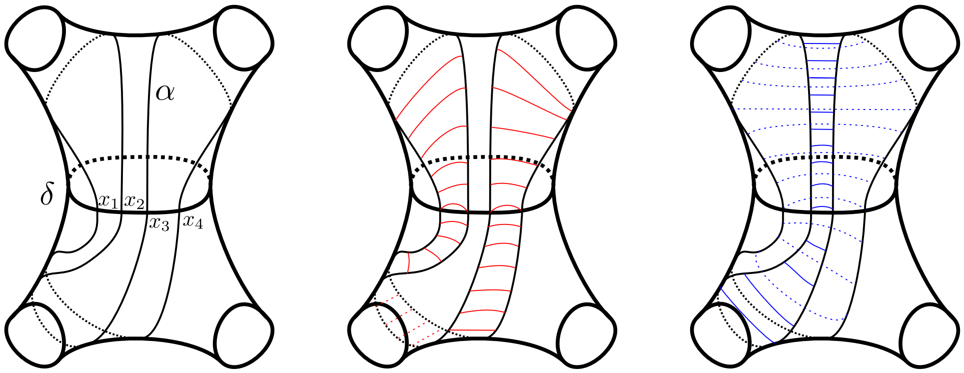

FIGURE 1. The problem with normal forms.

Sketch of proof. Suppose that D_1, D_2 are two disks bounded by α. Up to a small isotopy we may assume that D_1, D_2 intersect transversely. Thus, $D_1 \cap D_2$ is a finite disjoint union of embedded circles. Consider an innermost such circle in the interior of D_1. That is, we assume that there is a subdisk $D' \subset D_1$ which intersects D_2 exactly in $\partial D'$. Then D' and a suitable subdisk $D'' \subset D_2$ form an embedded S^2 in the handlebody. Since handlebodies are aspherical, and by the sphere theorem, $D' \cup D''$ bounds an embedded 3–ball B. Hence, we can modify D_2 by an isotopy ("pushing D'' through B") to reduce the number of intersection circles in $D_1 \cap D_2$. Inductively, we may therefore assume that D_1 and D_2 intersect only in their boundaries. But then they bound an embedded ball, and are isotopic as claimed. □

A word of warning is pertinent here. Suppose α, α' are two simple closed curves which intersect minimally, and consider the arcs $\alpha' \cap (\Sigma - \alpha)$ defined by α' in the surface obtained by cutting[2] at α. These are determined up to isotopy by the isotopy classes of α, α'. If we now assume that the curves in question are meridians, we can choose disks D, D' as in Lemma 2.2, and these disks are unique up to isotopy by Lemma 2.3. One might now hope that the intersections $D' \cap (V_g - D)$ are also determined by the isotopy classes of α, α'. As the following example shows, this is not the case. In particular, which intersection points of $\alpha \cap \alpha'$ are joined by an arc in $D \cap D'$ depends on the choice of disks D, D' and not just the isotopy classes of α, α'.

Example 2.4. *Consider a* 3*–ball B with four spots on the boundary. Let δ be a meridian separating ∂B into two pairs of pants. This is the "horizontal" meridian in Figure 1. We choose once and for all a disk D bounded by δ (horizontal in the Figure), cutting B into a "lower" and "upper" half-ball.*

A second meridian α will intersect δ in four points x_1, x_2, x_3, x_4 (shown on the left of the Figure). The middle and right pictures in Figure 1 show two

[2]For an oriented manifold M and a oriented codimension-1 submanifold X, we denote by $M - X$ the result of cutting M at X; i.e. for each component x of X, the manifold $M - X$ has two boundary components, corresponding to the two sides of x

disks D_1, D_2 bounded by α. Both are glued from one rectangular region, and two caps (bigons). For one of them, say D_1 (middle picture) the rectangular band is contained in the lower half-ball, whereas for the other (right picture) the rectangular band is contained in the upper half-ball.

Note that in the intersection $D \cap D_1$ there are intersection arcs between x_1, x_2 and between x_3, x_4. In the intersection $D \cap D_2$ on the other hand there are intersection arcs between x_1, x_4 and x_2, x_3.

Even without uniqueness, the intersection pattern between disks bounded by meridians guaranteed by Lemma 2.2 can be very useful. For example, we have the following standard tool. In its formulation, a *multimeridian* is a collection of pairwise disjoint, nonhomotopic meridians.

Proposition 2.5. *Suppose that α, β are two multimeridians. Then either α, β are disjoint, or there is a* wave, *i.e. an arc $a \subset \alpha$ so that*

i) a intersects β exactly in its endpoints.
ii) a approaches β from the same side at both of its endpoints.
iii) For any embedded arc $b \subset \beta$ with the same endpoints as a, the curve $a \cup b$ is a meridian.

Proof. Choose disks D_α, D_β bounded by α, β and intersecting as in Lemma 2.2. Consider the union of arcs $D_\alpha \cap D_\beta \subset D_\alpha$. Choose an *outermost* arc c. By definition there is then an arc $a \subset \alpha = \partial D_\alpha$ which, together with c, bounds a subdisk of D_α which intersects D_β exactly in c. In particular, a satisfies i) and ii). Any arc b as in iii) is homotopic relative to its endpoints in V to the arc c (they are both subarcs of D_β with the same endpoints), and so $a \cup b$ is indeed nullhomotopic in V. □

In particular, we can define the *meridian surgery* β' obtained by replacing the component of β containing b by $a \cup b$. The meridian β' is disjoint from β, and has at least two fewer intersection points with α. Such meridian surgery is a very common tool in the study of handlebody groups (see e.g. [41, 15, 45]).

At this point we also want to introduce the *disk complex*, the analogue of the curve complex in the setting of handlebodies. For a brief introduction to curve complexes, compare [10, Section 4.1]. We define $\mathcal{D}(V_g)$ to be the (flag) simplicial complex whose k–cells correspond to $k+1$ disjoint, distinct meridians up to isotopy. Using meridian surgery one can show

Lemma 2.6. *$\mathcal{D}(V_g)$ is connected. In fact, the distance between any two meridians α, β is bounded by $\frac{1}{2}i(\alpha, \beta)$.*

Just as the corresponding distance estimate for the curve graph $\mathcal{C}(\Sigma_g)$, intersection numbers are a very crude upper bound. In fact, meridians of distance two may have arbitrarily large intersection numbers.

A *cut system* is a collection $\alpha_1, \ldots, \alpha_g$ of g disjoint, non-isotopic meridians, so that the complement $\partial V_g - (\alpha_1 \cup \cdots \cup \alpha_g)$ is connected. Equivalently, there are disjointly embedded disks $D_1, \dots, D_g$ so that $\partial D_i = \alpha_i$ for all i, and so that $V_g - (D_1 \cup \dots \cup D_g)$ is a 3–ball. Using meridian surgery, one can show:

Lemma 2.7 (e.g. [20, Lemma 1.2], [15, Lemma 5.4])**.** *Suppose A, A' are two cut systems. Then there is a sequence A_i of cut systems*

$$A = A_1, A_2, \dots, A_n = A'$$

so that the multimeridians A_i and A_{i+1} are disjoint for all i (possibly having meridians in common).

3. Two Perspectives: Intrinsic and Surface

In this section we will develop two equivalent definitions of the handlebody group. The first definition is intrinsic, as the mapping class group $\mathrm{Mcg}(V_g)$ of the handlebody V_g. Here, the *mapping class group* $\mathrm{Mcg}(M)$ of a closed manifold is the group of isotopy classes of orientation-preserving self-homeomorphisms of M.

Identify the boundary ∂V_g of a genus g handlebody V_g with a surface of genus g. Since homeomorphisms and isotopies of V_g preserve the boundary, we then have a restriction homomorphism

$$\iota : \mathrm{Mcg}(V_g) \to \mathrm{Mcg}(\Sigma_g).$$

Our second definition will rest on the following standard lemma.

Lemma 3.1. *Two homeomorphisms F, F' of a handlebody V_g are isotopic if and only if their restrictions $\partial F, \partial F'$ to the boundary are isotopic.*

Sketch of proof. One direction is clear. We thus need to show that if F is an orientation-preserving homeomorphism of V_g whose restriction to the boundary is isotopic to the identity then F is isotopic to the identity as well. We produce the desired isotopy in several steps. Choose a cut system $\alpha_1, \dots, \alpha_g$ for V_g. Since ∂F is isotopic to the identity, we may modify F by an isotopy so that $F(\alpha_i) = \alpha_i$ for all i (apply the isotopy from ∂F to id in a collar neighbourhood of the boundary).

Choose disks D_i bounded by α_i. Then $F(D_i)$ are also disks bounded by α_i. By Lemma 2.3, we can modify F by an isotopy so that $F(D_i) = D_i$ for all i. In fact, we may assume that F restricts to the identity on D_i because any homeomorphism of a disk which is the identity on the boundary is isotopic to the identity (Alexander's trick).

At this point F is the identity on ∂V_g, and induces a homeomorphism of the 3–ball $V_g - (D_1 \cup \dots \cup D_g)$, whose boundary map is the identity. Hence, we can isotopy F to be the identity by Alexander's trick again. □

A immediate consequence of the lemma is that the restriction homomorphism ι is injective. We are thus led to our second definition, that of the *handlebody subgroup.* To this end, choose an identification of Σ_g with ∂V_g, and define

$$\mathcal{H}_g = \{\phi \in \mathrm{Mcg}(\Sigma_g) \mid \exists F : V_g \to V_g \text{ homeomorphism}, \phi = [\partial F]\}.$$

Lemma 3.1 implies that $\mathcal{H}_g$ is isomorphic to the mapping class group $\mathrm{Mcg}(V_g)$ of the handlebody via the restriction homomorphism ι.

A word should be said here about choices. The identification of Σ_g with the boundary ∂V_g depends on the choice of an homeomorphism. Two different choices will yield subgroups of $\mathrm{Mcg}(\Sigma_g)$ which are conjugate, but not necessarily equal (as we will see below).

Strictly speaking, "the" handlebody group is therefore only defined up to conjugacy. For intrinsic (algebraic or geometric) properties, this indeterminacy is not relevant, but in many topological applications it is crucial to distinguish the different conjugates of $\mathcal{H}_g$ (compare Section 4).

Convention 3.2. To ease notation, we will from now on call $\mathcal{H}_g$ (or, equivalently, $\mathrm{Mcg}(V_g)$) the *handlebody group*. We reserve the term *mapping class group* for $\mathrm{Mcg}(\Sigma_g)$.

Sometimes it is useful to consider handlebody groups of spotted handlebodies and handlebodies with marked points. From the intrinsic point of view it is clear how to define these: consider homeomorphisms which restrict to the identity on each spot and fix the marked points, up to isotopy doing the same. The analogue of Lemma 3.1 remains true in this setting, and we can therefore identify the handlebody groups again with subgroups of the corresponding surface mapping class groups. In fact, we can say a little bit more. For simplicity we will consider the case of one marked point or spot.

We begin by recalling the *Birman exact sequence* for the mapping class group of a surface $\Sigma_{g,1}$ of genus g with one marked point. Namely, there is a short exact sequence

$$1 \to \pi_1(\Sigma_g) \to \mathrm{Mcg}(\Sigma_{g,1}) \to \mathrm{Mcg}(\Sigma_g) \to 1.$$

where the map $\mathrm{Mcg}(\Sigma_{g,1}) \to \mathrm{Mcg}(\Sigma_g)$ is the forgetful map, and the image of $\pi_1(\Sigma_g)$ consists of "point pushing maps" (compare [10, Chapter 4.2]). The core observation for us is the following.

Lemma 3.3. *If Σ_g is identified with the boundary of V_g, then* any *element in the image of $\pi_1(\Sigma_g)$ extends to V_g.*

Proof. Let $F_t : \Sigma_g \to \Sigma_g$ be an isotopy from $F_0 = \mathrm{id}$ to a representative F_1 of the kernel. Consider U a closed regular neighbourhood of $\partial V_g = \Sigma_g$ in V_g, and choose an identification $U = \Sigma_g \times [0,1]$, where $\Sigma_g \times \{1\} = \partial V_g$. We define a map F on U as

$$F(x,t) = F_t(x).$$

By definition, F restricts to F_1 on ∂V_g and to the identity on the other boundary component of U. Thus, we can extend F by the identity to a homeomorphism of V_g restricting to F_1 on the boundary. □

As a consequence, if $f \in \mathcal{H}_g$ is any mapping class in the handlebody group, then any element of $\mathrm{Mcg}(\Sigma_{g,1})$ in the preimage of f lies in $\mathcal{H}_{g,1}$. As a consequence, $\mathcal{H}_{g,1}$ is simply the preimage of $\mathcal{H}_g$ in $\mathrm{Mcg}(\Sigma_{g,1})$, and we have a sequence

$$1 \to \pi_1(\Sigma_g) \to \mathcal{H}_{g,1} \to \mathcal{H}_g \to 1.$$

For the mapping class group of a surface with one boundary component Σ_g^1 and the surface with one marked point $\Sigma_{g,1}$ there is the following well-known sequence

$$1 \to \mathbb{Z} \to \mathrm{Mcg}(\Sigma_g^1) \to \mathrm{Mcg}(\Sigma_{g,1}) \to 1$$

where the map $\mathrm{Mcg}(\Sigma_g^1) \to \mathrm{Mcg}(\Sigma_{g,1})$ is defined by gluing a punctured disk to the boundary component of Σ_g^1, and the kernel is generated by the Dehn twist about the boundary component [10, Chapter 3.6.2]. Note here that in $\mathrm{Mcg}(\Sigma_g^1)$ we consider homeomorphisms which restrict to the identity on the boundary component of Σ_g^1 up to isotopies with the same property. Again, it is easy to see that the Dehn twist about the boundary of a spot extends to the handlebody, and therefore $\mathcal{H}_g^1$ is simply the preimage of $\mathcal{H}_{g,1}$ in $\mathrm{Mcg}(\Sigma_g^1)$. We thus have the sequence

$$1 \to \mathbb{Z} \to \mathcal{H}_g^1 \to \mathcal{H}_{g,1} \to 1.$$

4. A Topological Interlude

Before continuing to discuss properties of handlebody groups in earnest, this section is concerned with two applications which highlight the role of handlebody groups in topology.

We begin with recalling the standard way of constructing three-manifolds using handlebodies: *Heegaard splittings*. Let V_g be a handlebody of genus g. The boundary of V_g is then a surface Σ_g of genus g. If F is a representative of a mapping class $\varphi \in \mathrm{Mcg}(\Sigma_g)$, then one can obtain a closed three-manifold M_φ by gluing two copies of V_g along their boundary according to F (to be precise, we need to glue a copy of V_g to one with the opposite orientation). It is not hard to see that, up to homeomorphism, the resulting manifold $M_\varphi = V_g \cup_\varphi V_g$ depends only on the mapping class φ and not F. From the fact that every closed oriented three–manifold admits a triangulation one can quickly conclude that this construction in fact yields all closed oriented three-manifolds.

The description of M as M_φ is called a *Heegaard splitting of M*. All information about M is then encoded in φ, and so this point of view allows us, in principle, to study topological and geometric properties of three–manifolds by studying the mapping class group of the surface Σ_g.

However, in a description $M = M_\varphi$, neither the genus g nor the gluing mapping class φ are determined by the three-manifold M. To describe this ambiguity, we can use the handlebody group. We begin with the gluing map. Intuitively, if we modify φ by an element which extends to one of the handlebodies, then this extension (together with the identity on the other handlebody) will define a homeomorphism between the resulting three-manifolds. Formally, there is an element $f_g \in \mathrm{Mcg}(\Sigma_g)$ so that if $\psi = a\varphi b$ for $a \in \mathcal{H}_g, b \in f_g\mathcal{H}_g f_g^{-1}$, then M_ψ is homeomorphic to M_φ. For a quick and readable discussion of this, compare [53, Section 2.2] or any textbook on 3-manifold topology. This defines an equivalence relation $\sim$ on $\mathrm{Mcg}(\Sigma_g)$: two elements are equivalent exactly if they are in the same double coset of the handlebody subgroup $\mathcal{H}_g$ and the conjugate $f_g\mathcal{H}_g f_g^{-1}$. We then get a map

$$\mathrm{Mcg}(\Sigma_g)/\sim \to \mathcal{M}(3)$$

where $\mathcal{M}(3)$ denotes the set of homeomorphism classes of closed oriented three-manifolds.

To describe the ambiguity in genus, we need to use surfaces with boundary. Namely, consider a surface Σ_g^1 of genus g with one boundary component. Gluing a disc to the boundary component defines a map $\mathrm{Mcg}(\Sigma_g^1) \to \mathrm{Mcg}(\Sigma_g)$, and for a $\varphi \in \mathrm{Mcg}(\Sigma_g^1)$ we let M_φ to be the manifold defined by the image of φ in $\mathrm{Mcg}(\Sigma_g)$.

The surface of genus $g+1$ with one boundary component can be obtained from Σ_g^1 by gluing a torus with two boundary components to the boundary component of Σ_g^1. Extending mapping classes by the identity then defines a stabilisation map

$$s_g : \mathrm{Mcg}(\Sigma_g^1) \to \mathrm{Mcg}(\Sigma_{g+1}^1),$$

equipping the set of groups $\{\mathrm{Mcg}(\Sigma_g^1), g \geq 0\}$ with the structure of a directed system. The stabilisation map is compatible with the handlebody subgroups $\mathcal{H}_g^1$ in the sense that $s_g(\mathcal{H}_g^1) \subset \mathcal{H}_{g+1}^1$.

For any g and any $\varphi \in \mathrm{Mcg}(\Sigma_g^1)$, the manifolds $M_{s_g(\varphi)}$ and M_φ are diffeomorphic, and we thus have a map

$$\lim_{g\to\infty} \mathrm{Mcg}(\Sigma_g^1) \to \mathcal{M}(3)$$

where $\mathcal{M}(3)$ denotes the set of diffeomorphism classes of closed oriented three-manifolds.

In fact, the stabilisation maps s_g are also compatible with the stabilisation equivalence relation $\sim$ described above. The central result on Heegaard splittings can now be phrased as follows

Theorem (Reidemeister-Singer)**.** *The map*

$$\lim_{g\to\infty} \left((\mathcal{H}_g^1 \backslash \mathrm{Mcg}(\Sigma_g^1)/f_g\mathcal{H}_g^1 f_g^{-1} \right) \to \mathcal{M}(3)$$

is a bijection.

By this theorem, topological invariants of three-manifolds can be seen as invariants of mapping classes which are constant on double cosets of the handlebody group and do not change under stabilisation.

We want to give one example of this point of view. Namely, consider a Heegaard splitting of the three-sphere $S^3 = M_f$. As above, changing the splitting by an element of the handlebody group does not change the resulting manifold. On the other hand, if we modify the splitting by an element in the Torelli group, the resulting manifold need not be S^3 – but it is easy to see that it will still be an integral homology sphere (and in fact all of them arise in this way). Using the action of $\mathcal{H}_g$ on the homology of the surface (compare Section 7), Morita [49] showed that one can in fact obtain any homology sphere already using an element of the Johnson kernel, an infinite index subgroup of the Torelli group. Pitsch [53] improved this further, showing that maps in the next term of the Johnson filtration also suffice. This has consequences for the study of the Casson invariant of homology spheres. For details, we refer the interested reader to [49, 53].

There is one other important topological application of handlebody groups which we want to highlight. Namely, let k be a knot in S^3. We call k *fibered* if the complement in S^3 of a regular neighbourhood K of k is homeomorphic to a mapping torus

$$\Sigma_g^1 \times [0,1]/(x,0) \sim (\varphi(x),1)$$

where Σ_g^1 is a genus g surface with one boundary component. The mapping class $\varphi \in \mathrm{Mcg}(\Sigma_g^1)$ is then called the *monodromy* of the knot k. In this setting, we have the following theorem.

Theorem ([8])**.** *A fibered knot k is homotopically ribbon if and only if its monodromy φ is conjugate into the handlebody group $\mathcal{H}_g^1$.*

The precise definition of homotopically ribbon is not important here; it may suffice to remark that it is stronger than being slice and weaker than being ribbon – and the exact relation between these properties is a major open question in knot theory. The point we want to make is that the algebraic condition of being conjugate into the handlebody group has a purely topological characterisation in terms of k.

As a quick application, note that this theorem has been used [5] to show that certain knots are not ribbon, by showing that their monodromies cannot be conjugate into the handlebody group.

5. Examples and the Membership Problem

The purpose of this section is twofold: first we will construct three important types of elements in the handlebody group (which will allow us to learn the about first algebraic properties of $\mathcal{H}_g$), and then we will develop methods to detect when a surface mapping class is contained in $\mathcal{H}_g$.

We begin with examples. For a simple closed curve δ on Σ_g denote by T_δ the left handed Dehn twist about δ (compare [10]).

Example 5.1 (Meridional Dehn Twists)**.** *Let δ be a meridian. Then the Dehn twist T_δ is contained in $\mathcal{H}_g$. We call these* meridional Dehn twists*.*

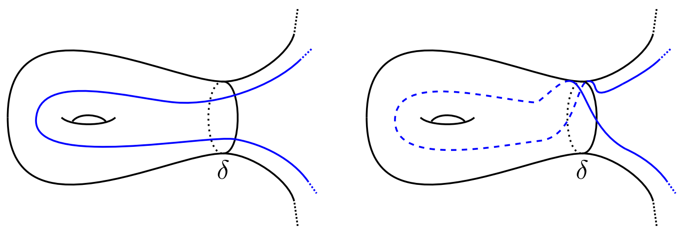

FIGURE 2. A half-twist and its effect on a curve.

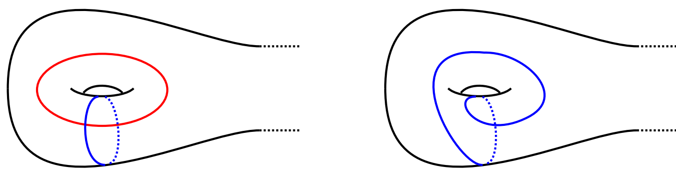

FIGURE 3. A twist not in the handlebody group

Intuitively, we can cut V_g at a disk bounded by δ and twist the handle by a full turn. Formally, let D be a disk bounded by δ, and let $U \subset V$ be a closed regular neighbourhood of D. Choose an identification $U = D^2 \times [0,1]$ so that $\partial D^2 \times [0,1]$ is a closed regular neighbourhood of δ. Define a map on U by

$$F(x,t) = (r_{t2\pi}(x), t)$$

where r_θ denotes the counterclockwise rotation by an angle of θ. F restricts to the identity on $D^2 \times \{0\}, D^2 \times \{1\}$ and therefore extends to a homeomorphism of V_g, and clearly restricts to a Dehn twist about δ on the boundary.

Example 5.2. *If δ is a separating meridian, then the (left) Dehn twist about δ has a square root in the handlebody group, called a* (left) half-twist about δ. *To construct it, we can cut the handle at δ and twist by a half-turn (compare Figure 2). Formally, suppose that δ cuts ∂V into Y_1, Y_2. The half-twist will be the identity on Y_1, a mapping class of order 2 on Y_2 (a rotation by π), and half of the twist as described in Example 5.1.*

In fact, Dehn twists can also be used to give examples of elements not contained in the handlebody group.

Example 5.3. *Consider a handlebody V_g, a meridian δ, and a curve α which intersects δ in a single point. Then $T_\alpha \notin \mathcal{H}_g$.*

Namely, consider $T_\alpha(\delta)$. This curve intersects δ in a single point, and therefore cannot be a meridian (compare Figure 3 and Lemma 2.2). However, handlebody elements clearly preserve the property of being a meridian, and therefore $T_\alpha \notin \mathcal{H}_g$.

As a conclusion we get the first indication that the handlebody group is a "complicated" subgroup of the mapping class group.

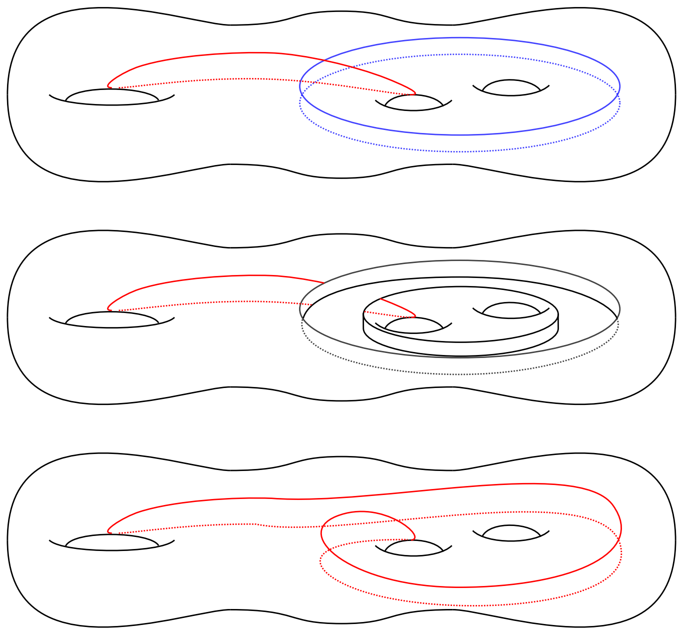

FIGURE 4. An annulus twist. In the top figure, there are two curves in the handlebody, which bound an annulus. Cutting at the annulus yields a manifold with two preferred annuli on the boundary (middle figure). The annulus twist is the homeomorphism obtained by twisting one of them by a full turn and then regluing. The effect on a meridian is shown on the bottom. Note that with the orientation of the surface, the twists on the boundary have opposide handedness.

Corollary 5.4. *The handlebody group $\mathcal{H}_g$ is an infinite, infinite-index subgroup of* $\mathrm{Mcg}(\Sigma_g)$*, which is not normal.*

Proof. $\mathcal{H}_g$ is clearly infinite, since it contains a Dehn twist T_δ about some non-separating meridian δ. Let α be as in the previous example. Then no power of T_α is contained in $\mathcal{H}_g$. Namely, the algebraic intersection number of $T^n_\alpha(\delta)$ with δ is n, and therefore $T^n_\alpha(\delta)$ is not a meridian. This shows that $\mathcal{H}_g$ has infinite index. Finally, since α and δ are non-separating, and all non-separating Dehn twists are conjugate in the surface mapping class group we conclude that $\mathcal{H}_g$ is not normal. □

Next, we turn to products of several disjoint Dehn twists. The most important construction is given by the following example.

Example 5.5 (Annulus Twists)**.** *Suppose that α_1, α_2 are two disjoint curves, so that there is a properly embedded annulus $A \subset V_g$ with $\partial A = \alpha_1 \cup \alpha_2$. Then the* annulus twist *$T_{\alpha_1} T^{-1}_{\alpha_2}$ is an element of $\mathcal{H}_g$.*

Intuitively, we cut V_g at the annulus A and twist one side by a full turn (compare Figure 5). Looking from above, at both curves the annulus turns the same way, so in the orientation of boundary we see a product of a left and a right twist. Formally, consider a neighbourhood U of A which is homeomorphic to $S^1 \times [0,1] \times [0,1]$, where $S^1 \times [0,1] \times \{0,1\} = U \cap \partial V_g$. Define a map on U by

$$F(x,t,s) = (r_{t2\pi}(x), t, s)$$

F restricts to the identity on $S^1 \times \{0\} \times [0,1], S^1 \times \{1\} \times [0,1]$ and therefore extends to a homeomorphism of V_g. It restricts to the boundary as a product of a left and a right Dehn twist about α_1, α_2, since one of the induced inclusions $S^1 \times [0,1] \times \{0,1\}$ is orientation preserving, and one orientation reversing.

In fact, as the following theorem shows, these examples already yield all multitwists which are contained in the handlebody group.

Theorem 5.6 ([51, Theorem 1.11], [46, Theorem 1])**.** *Suppose that $\alpha_1, \dots, \alpha_k$ are disjoint simple closed curves on Σ_g. Then*

$$T_{\alpha_1}^{n_1} \dots T_{\alpha_k}^{n_k} \in \mathcal{H}_g$$

if and only if, up to reordering,

i) $\alpha_1, \dots, \alpha_r$ are meridians, and
ii) For all $i = r+2, r+4, \dots, r+2l = k$, α_{i-1}, α_i bound an embedded annulus, and $n_{i-1} = -n_i$.

In particular, this theorem implies that the only (single) Dehn twists which are contained in the handlebody group are the meridional Dehn twists.

The theorem also severely constrains which multitwists are contained in the handlebody group. In constrast, the following easy lemma shows that the conjugacy problem is trivial for multitwists.

Lemma 5.7. *Every multitwist is conjugate into $\mathcal{H}_g$.*

Proof. Let $\Delta \subset \Sigma_g$ be any multicurve. We begin by observing that since any surface with boundary is the boundary of a handlebody with spots, there is an identification of Σ_g with the boundary of a handlebody so that Δ consists of meridians (namely, glue handlebodies with spots according to the complementary components of Δ). But this implies that there is a mapping class group element ϕ so that $\phi(\Delta)$ consists of meridians for our preferred identification of Σ_g with the boundary of a handlebody. Thus, conjugating any multitwist about Δ by ϕ is a product of twists about meridians, and thus contained in $\mathcal{H}_g$. □

There is a third class of examples which is crucial to understanding handlebody groups. To describe them, we need a slightly different construction of handlebodies. To this end, let S be a surface of genus $g \geq 0$ with $n \geq 1$ boundary components. The product $V = S \times [0,1]$ is a handlebody (see

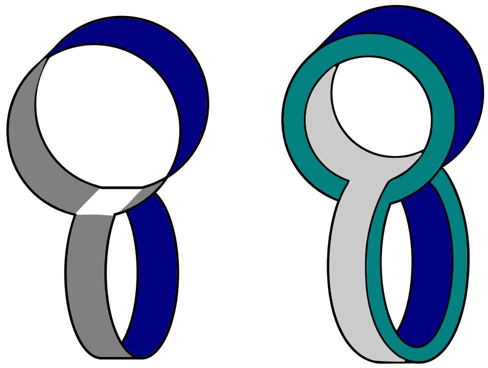

FIGURE 5. The interval bundle over a bordered torus is a genus 2 handlebody.

Figure 5 for an example). Namely, consider disjoint arcs $a_1, \dots, a_k \subset S$ so that $S - (a_1 \cup \dots \cup a_k)$ is a single polygon. Then $D_i = a_i \times [0,1] \subset V$ are disjoint disks, and $V - (D_1 \cup \dots \cup D_k)$ is a single 3–ball. Hence, V is a 3–manifold obtained from a 3–ball by identifying disjoint disks on the boundary in pairs, and thus is a handlebody.

Example 5.8 (I-bundle maps)**.** *Let S be a surface with boundary and $V = S \times [0,1]$ the trivial interval bundle. If f is any homeomorphism of S, we obtain a homeomorphism $f \times \mathrm{Id}$ of V. The isotopy class of this homeomorphism defines an element of $\mathrm{Mcg}(V)$. We call the elements obtained in this fashion* interval bundle maps *or* I-bundle maps *for short.*

We can make this construction a bit more formal.

Corollary 5.9. *Suppose S is a surface of genus $g \geq 0$ with $n \geq 1$ boundary components, and let k be such that $\pi_1(S)$ is a free group of rank k. Let S' be the surface obtained from S by gluing a punctured disk to each boundary component of S. Then the construction from Example 5.9 induces a homomorphism*

$$I : \mathrm{Mcg}(S') \to \mathcal{H}_k$$

We will later see that this map is in fact injective (compare Corollary 6.2).

Proof. The construction from Example 5.8 yields a homomorphism

$$i : \mathrm{Mcg}(S) \to \mathrm{Mcg}(V_k).$$

We have the well-known sequence (compare e.g. [10, Proposition 3.19])

$$1 \to \mathbb{Z}^n \to \mathrm{Mcg}(S) \to \mathrm{Mcg}(S') \to 1$$

where the kernel $\mathbb{Z}^n$ is generated by Dehn twists about the boundary components of S. To prove the corollary, we therefore only have to show that these Dehn twists map to the identity under i. However, this follows since the twist about the boundary component β maps to a homeomorphism which restricts on the boundary to a product of a left and a right Dehn twist about curves parallel to β (compare the discussion in Example 5.5 to see why the handedness is opposite). □

Suppose now that $\phi \in \mathrm{Mcg}(\Sigma_g)$ is given to us – how can we detect if ϕ is in fact an element of the handlebody group?

Maybe the most straightforward and useful criterion is the following.

Lemma 5.10. *$\phi \in \mathrm{Mcg}(\Sigma_g)$ is an element of $\mathcal{H}_g$ if and only if for some cut system $\alpha_1, \ldots, \alpha_g$ the image $\phi(\alpha_1), \ldots, \phi(\alpha_g)$ is also a cut system.*

Proof. Choose disjointly embedded disks D_i bounded by the α_i. If ϕ is in the handlebody group, then clearly the images $\phi(\alpha_i)$ are also meridians, and the claim holds. The converse is the Alexander trick again: if $\phi(\alpha_1), \ldots, \phi(\alpha_g)$ is a cut system, then choose disks D'_i bounded by $\phi(\alpha_i)$. The homeomorphism $\Sigma_g - (\alpha_1 \cup \cdots \cup \alpha_g) \to \Sigma_g - (\phi(\alpha_1) \cup \cdots \cup \phi(\alpha_g))$ extends to a homeomorphism $V_g - (D_1 \cup \cdots \cup D_g) \to V_g - (D'_1 \cup \cdots \cup D'_g)$ since every homeomorphism of a sphere extends to the ball. This implies that ϕ extends to V_g, and therefore $\phi \in \mathcal{H}_g$. □

As a corollary we have the following.

Corollary 5.11. *Let $\phi \in \mathrm{Mcg}(\Sigma_g)$ be given. The following are equivalent*

i) $\phi \in \mathcal{H}_g$.
ii) The outer automorphism of $\pi_1(\Sigma_g)$ induced by ϕ preserves the kernel

$$K = \ker(\pi_1(\Sigma_g, p) \to \pi_1(V_g, p))$$

induced by the inclusion of the boundary into V_g.
iii) For any meridian α, the curve $\phi(\alpha)$ is a meridian.
iv) For some cut system $\alpha_1, \ldots, \alpha_g$ the image $\phi(\alpha_1), \ldots, \phi(\alpha_g)$ is also a cut system.

At this point we want to mention that the question if an element ϕ is conjugate into $\mathcal{H}_g$ is much more challenging but, in a sense, more natural: being conjugate into $\mathcal{H}_g$ exactly means that ϕ can be extended to a handlebody, whereas being contained in $\mathcal{H}_g$ means that ϕ can be extended to a handlebody using the chosen, non-canonical identification $V_g \cong \Sigma_g$. We will see later (Lemma 7.2) that there are elements in $\mathrm{Mcg}(\Sigma_g)$ which are not conjugate into $\mathcal{H}_g$.

We can also try to characterise when an element is conjugate into $\mathcal{H}_g$ using variants of the conditions of Corollary 5.11. Conditions ii) is particularly amenable for this:

Corollary 5.12. *An element $\phi \in \mathrm{Mcg}(\Sigma_g)$ is conjugate into $\mathcal{H}_g$ if and only if there is a surjection $p : \pi_1(\Sigma_g) \to F_g$ so that ϕ_* preserves* $\ker(p)$.

Proof. In light of Corollary 5.11 it suffices to show that any surjection p as in the assumption can be realised by an identification of Σ_g with the boundary of a handlebody. This is proved e.g. in [36, Lemma 2.2]. □

However, applying this corollary is not very straightforward, since it is not clear how to predict which p can work in advance. Similarly, to apply Corollary 5.11 iii) or iv) one needs to know in advance which simple closed curves are the meridians.

Nevertheless, the conjugation problem has an algorithmic answer. Namely,

Theorem 5.13 ([9])**.** *There is an algorithm which detects, given a pseudo-Anosov map ϕ, if it is conjugate into $\mathcal{H}_g$. In case ϕ is not conjugate into $\mathcal{H}_g$, the algorithm can describe the maximal compression body to which ϕ extends.*

A *compression body* a 3–manifold obtained from a trivial surface bundle $\Sigma \times [0, 1]$ in two steps. First, one attaches 2-handles (think of them as thickened disks) along disjoint simple closed curves on $\Sigma \times \{0\}$ to obtain a manifold M. In a second step, one glues a ball to each boundary component of M which is homeomorphic to a sphere. The resulting C has one boundary component which was the original surface $\Sigma \times \{1\}$, and possibly several other boundary components. A handlebody is a compression body where the system of curves chosen in the first step has complements which are spheres with boundary. Compare e.g. [6] for details on compression bodies and their relation to extending surface homeomorphisms.

The basic idea behind the algorithm is the following: fix a hyperbolic metric on the surface. The main step is to show that there is a constant $L = L(\phi)$ so that if ϕ extends to some handlebody, then there is a meridian α of length at most L. Once this is established, the fact that there are only finitely many simple closed curves on Σ_g up to any given length can be used to test if ϕ extends. To prove the main step, suppose that ϕ extends, and consider any meridian α. The images $\phi^n(\alpha)$ are then meridians, and so for large n have waves as in Section 2. In fact, one can show that any long enough segment $a \subset \alpha$ (depending on the dilatation of ϕ) already has an image $\phi(a)$ which forms a wave with a. Using meridian surgery produces then the short meridian.

We close this section with a brief discussion of how the different conjugates of the handlebody group interact. On the one hand, that there are many conjugates of $\mathcal{H}_g$ which intersect in infinite, nontrivial subgroups. Namely, we can use a mapping class ϕ fixing a separating meridian δ, so that on one complementary component, ϕ restricts to the identity, and on the other ϕ restricts to an element that cannot be conjugate into a handlebody group (compare Section 7). Now, the conjugate of $\mathcal{H}_g$ by ϕ will intersect $\mathcal{H}_g$ in the desired way.

The next theorem shows that in some sense this is the only way that two conjugates can intersect in a complicated way.

Theorem 5.14 ([1, 3, 9])**.** *Let ϕ be any pseudo-Anosov map. Then ϕ extends to at most finitely many handlebodies. In other words, ϕ can be contained in at most finitely many conjugates of $\mathcal{H}_g$.*

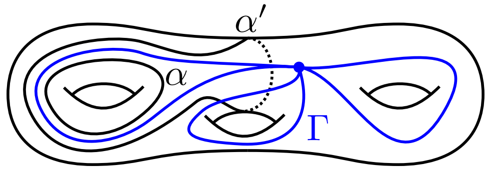

FIGURE 6. A Nielsen twist realised by a annulus twist

6. A Third Perspective: Free Groups

If we consider $\mathcal{H}_g$ as the mapping class group of a handlebody, then we also obtain a natural map

$$A : \mathcal{H}_g \to \mathrm{Out}(F_g)$$

induced by the action of handlebody homeomorphisms on the fundamental group of V_g, which is simply a free group. One way to understand the action on $\pi_1(V_g)$ is the following. Consider $\delta_1, \ldots, \delta_g$ a cut system. Consider an embedded graph $\Gamma \subset \partial V_g$ with one vertex and g edges, so that each edge of Γ intersects one of the δ_i in exactly one point. We call such a graph a *dual graph* to the cut system. The map $\pi_1(\Gamma) \to \pi_1(V_g)$ induced by the inclusion is then an isomorphism.

For a class $\phi \in \mathcal{H}_g$, we can now consider the image graph $\phi(\Gamma)$, and read off the elements defined by its edges according to the intersection pattern with the meridians δ_i, in order to determine the action of ϕ on $\pi_1(V)$.

To get a feeling for this action, let us consider again the example elements we built in Section 5.

Example 6.1. *i) If δ is a meridian, then $A(T_\delta) = 1$. This can for example easily be seen from the explicit description given in Example 5.1: the extension F of T_δ to V_g is the identity outside an embedded ball in V_g (the regular neighbourhood of the disk). For any loop α the image $F(\alpha)$ is therefore homotopic to α, showing the result.*

ii) Suppose that δ is a meridian such that $\partial V_g - \delta$ consists of one torus and one genus $g-1$ subsurface. We will then say that δ cuts off a handle*. We can choose a free basis $x_1, \ldots, x_g$ for $\pi_1(V_g)$ so that the image of the torus is $\langle x_1 \rangle$ and the image of the other subsurface is $\langle x_2, \ldots, x_g \rangle$. The image of the half-twist H_δ under A then maps x_1 to x_1^{-1}, fixing all other x_i.*

iii) Annulus twists can act nontrivially on $\pi_1(V_g)$. Intuitively, every time a loop crosses the annulus A, it is modified by inserting a copy of the element defined by the annulus.

We make this precise in a concrete example, which we will use later.

Consider a cut system $\delta_1, \ldots, \delta_g$ and a dual graph Γ as above. Take α a curve disjoint from Γ intersecting δ_1 once, and α' a curve bounding an annulus with α (in V_g) which intersects Γ in a single point, on the edge *dual to δ_2 (compare Figure 6). We consider the annulus twist $T_\alpha T_{\alpha'}^{-1}$. Every edge except for the one dual to δ_2 is fixed (since the twist curves are disjoint), whereas the edge dual to δ_2 is mapped (up to homotopy) to a concatenation of itself and the edge dual to δ_1. Hence we have*

$$A(T_\alpha T_{\alpha'}^{-1})(x_i) = \begin{cases} x_i & \text{if } i \neq 2 \\ x_1 x_2 & \text{if } i = 2 \end{cases}$$

iv) *Let S be a surface of genus $g \geq 0$ with $n \geq 1$ boundary components, and $\phi \in \mathrm{Mcg}(S)$ any element. Recall the construction of the I–bundle map $I(\phi)$ from Corollary 5.9. We then have*

$$A(\phi) = \phi_*$$

where ϕ_* denotes the induced map on $\pi_1(S) = \pi_1(S \times [0,1])$.

Part iv) also shows:

Corollary 6.2. *Suppose S is a surface of genus $g \geq 0$ with $n \geq 1$ boundary components, and let k be such that $\pi_1(S)$ is a free group of rank k. Then the map*

$$I : \mathrm{Mcg}(S') \to \mathcal{H}_k$$

from Corollary 5.9 is injective.

Proof. The concatenation $A \circ I$ agrees with the action of $\mathrm{Mcg}(S')$ on $\pi_1(S')$, and it is known that this is injective, e.g. [10, Theorem 8.8]. □

We can also use these examples for the following.

Theorem 6.3 ([12, 56, 48])**.** *The map $A : \mathcal{H}_g \to \mathrm{Out}(F_g)$ is surjective.*

Sketch of proof. Let $x_1, \ldots, x_g$ be a free basis of F_g. It is well-known that $\mathrm{Out}(F_g)$ is generated by automorphisms of the following form

$$\sigma_{ij}(x_k) = x_k \text{ if } k \neq i, \quad \sigma_{ij}(x_i) = x_i x_j$$

and

$$\varepsilon_i(x_k) = x_k \text{ if } k \neq i, \quad \varepsilon_i(x_i) = x_i^{-1}.$$

Thus, we only need to realise them as handlebody homeomorphisms. For the first class, this is possible using annulus twists as in the example iii) above. The second class is realised by half-twists about meridians cutting off the handle corresponding to x_i as in the example ii). □

As we have seen, Dehn twists about meridians are contained in the kernel of A – showing in particular that the handlebody group $\mathcal{H}_g$ is *not* generated by the Dehn twists contained in it.

In fact, the meridional Dehn twists exactly generate the kernel.

Theorem 6.4 ([37])**.** *The kernel of $A : \mathcal{H}_g \to \mathrm{Out}(F_g)$ is the* twist group, *i.e. the group generated by Dehn twists about meridians.*

Sketch of proof. We have a short exact sequence

$$1 \to \ker(A) \to \mathcal{H}_g \to \mathrm{Out}(F_g) \to 1.$$

It is known that $\mathrm{Out}(F_g)$ is a finitely presented group, and $\mathcal{H}_g$ is finitely generated (compare Section 8). In fact, Luft considers a generating set $\mathcal{G}$ of $\mathcal{H}_g$ a part of which lies in the kernel of A, and the rest of which maps bijectively under A to a specific generating set $\mathcal{G}'$ of $\mathrm{Out}(F_g)$. Using an explicit presentation $\langle \mathcal{G}' | \mathcal{R} \rangle$ of $\mathrm{Out}(F_g)$ with respect to $\mathcal{G}'$, one can now easily give a normal generating set of $\ker(A)$: those words in $\mathcal{G}$ defined by the relators $\mathcal{R}$, and the elements $\ker(A) \cap \mathcal{G}$. To prove the theorem, one now only has to check that these finitely many elements of $\mathcal{H}_g$ are indeed products of meridional twists. □

The twist group is a somewhat unpleasant group, as shown by the following

Theorem 6.5 ([44]). *The twist group is not finitely generated. In fact, its Abelianisation contains* $\mathbb{Z}^\infty$*.*

On the other hand, the proof of Theorem 6.3 together with Luft's theorem also yields the following natural (infinite) generating set for the handlebody group.

Corollary 6.6. *The handlebody group* $\mathcal{H}_g$ *is generated by meridional Dehn twists, meridional half-twists and annulus twists.*

The connection to $\mathrm{Out}(F_n)$ can sometimes be used to say something about handlebody mapping classes. For example, we note the following result, which is a source of interesting pseudo-Anosov mapping classes in $\mathcal{H}_g$. For its formulation, recall that an outer automorphism Θ of the free group is *geometric* if there is an identification $F_g = \pi_1(\Sigma)$ of the free group with the fundamental group of a surface, so that Θ is induced by a homeomorphism of Σ. It is *fully irreducible* if no power Θ^k preserves a free factor (up to conjugacy). For background on these terms, compare e.g. [2].

Lemma 6.7. *Suppose that* Θ *is a non-geometric, fully irreducible outer automorphism of* F_g*. Then any element of* $A^{-1}(\Theta)$*, i.e. any handlebody group element* ϕ *acting as* Θ *on* $\pi_1(V_g)$*, is pseudo-Anosov.*

Proof. Suppose that some power ϕ^n fixes a curve α. If α is a meridian, let D be a disk bounded by α. By the van-Kampen theorem, D defines a nontrivial free splitting of $\pi_1(V_g)$ as a free product $\pi_1(V_g) = B_1 * B_2$ or a HNN extension $\pi_1(V_g) = B_1 *$. Since ϕ^n preserves α, and thus D up to homotopy, we would then get that Θ^{2n} preserves the free factor B_1 up to conjugacy. This is impossible by full irreducibility.

If α however is not a meridian, then it defines a nontrivial conjugacy class z in $\pi_1(V_g) = F_g$ which is then preserved by Θ. By [2, Theorem 4.1] this is only possible if Θ is geometric (and z corresponds to the boundary of a surface realising it) which contradicts our assumption. □

Question 6.8. *Can one relate geometric properties of* Θ *to geometric properties of a (suitable) element in* $A^{-1}(\Theta)$*?*

A concrete variant of this would be: consider the action of Θ *on Culler-Vogtmann Outer Space. How does the translation length with respect to the Lipschitz metric relate to possible translation lengths of* $\varphi \in A^{-1}(\Theta)$ *acting on Teichmüller space with the Thurston metric?*

Let us also remark that geometric, fully irreducible elements can in fact be realised by handlebody group elements which are reducible (surface) mapping classes: they are in the image of the interval bundle map I from Corollary 5.9, all elements of which preserve the boundary of the surface S. Also note that since the twist group contains pseudo-Anosov maps, the fibre $A^{-1}(\Theta)$ for *any* nontrivial Θ will contain pseudo-Anosov maps.

A somewhat more careful analysis of possible stabilisers of elements in the preimage of Θ can also be used to study group-theoretic properties of Θ. Namely, we can show

Proposition 6.9 ([22])**.** *For* $g \geq 4$ *and any finite index subgroup* $\Gamma < \mathrm{Out}(F_g)$ *there is no homomorphism* $s : \Gamma \to \mathcal{H}_g$ *so that* $A \circ s$ *is the identity.*

7. Symplectic Representation

For the surface mapping class group, we have the following well-known short exact sequence induced by the action of mapping classes on $H_1(\Sigma_g; \mathbb{Z})$:

$$1 \to \mathcal{I}_g \to \mathrm{Mcg}(\Sigma_g) \to \mathrm{Sp}(2g, \mathbb{Z}) \to 1.$$

Here, $\mathcal{I}_g$ is the *Torelli group* (compare [10, Theorem 6.4 and Chapter 6.5]). In this section we discuss the interaction of $\mathcal{H}_g$ with the quotient and the kernel of this sequence. We begin with the image in the symplectic group. There is one obvious constraint for the handlebody group acting on homology. Namely, let

$$L = \ker(H_1(\Sigma_g; \mathbb{Z}) \to H_1(V_g; \mathbb{Z}))$$

be the kernel of the map induced by inclusion of the boundary. L is a Lagrangian subspace, i.e. half-dimensional and totally isotropic with respect to the algebraic intersection pairing. Isotropy follows from Lemma 2.1, and a cut system spans a g–dimensional subspace of homology. In fact, any cut system therefore defines a basis of L. We can choose a standard symplectic basis $a_1, \ldots, a_g, b_1, \ldots, b_g$ of $H_1(\Sigma_g; \mathbb{Z})$ so that L is generated by $a_1, \ldots, a_g$ and we will do so throughout this section.

Elements of the handlebody group clearly preserve L. By constructing enough explicit homeomorphisms of V_g one can show that this obstruction is the only one.

Theorem 7.1 ([4, Lemma 2.2], [26])**.** *The image of* $\mathcal{H}_g \to \mathrm{Sp}(2g, \mathbb{Z})$ *consists exactly of those symplectic matrices which preserve* L*. Explicitly, a matrix* *in* $\mathrm{Sp}(2g,\mathbb{Z})$ *is realised by a handlebody group element if it has the following block form with respect to the basis* $a_1,\ldots,a_g,b_1,\ldots,b_g$ *of* $H_1(\Sigma_g;\mathbb{Z})$*:*

$$\begin{pmatrix} A & B \\ 0 & D \end{pmatrix}$$

with $g\times g$ *matrices* A,B,C *satisfying*

$$A^t = D^{-1}, B^tD = D^tB$$

If a matrix M preserves a subspace, then its characteristic polynomial is reducible – and it is easy to construct symplectic matrices with irreducible characteristic polynomial. As a consequence, there are elements of the symplectic group which do not leave any Lagrangian invariant and we conclude:

Lemma 7.2. *There are elements* $\phi\in\mathrm{Mcg}(\Sigma_g)$ *which are not conjugate into* $\mathcal{H}_g$*.*

Next, we turn to the interactions between the handlebody group and the Torelli group $\mathcal{I}_g$. The first important result is that homology alone is unable to detect membership in the handlebody group.

Theorem 7.3. *i)* [36] *There are elements of* $\mathcal{I}_g$ *no power of which is conjugate into the handlebody group.*
ii) [32] *There are elements of any term in the Johnson filtration which are not conjugate into the handlebody group.*

The proofs of both of these results use the connection of the handlebody group to topological questions, and we can only very briefly touch upon the methods.

For i), one uses Heegaard splittings as in Section 4. The idea is to show that if ϕ extends to some handlebody, then there is an embedding of the surface into S^3 so that modifying the standard Heegaard splitting of S^3 by ϕ still yields S^3. On the other hand, if $\phi\in\mathcal{I}_g$, modifying by ϕ will yield a homology sphere. To prove i), one now exploits the connection of Birman-Craggs homomorphisms to Rokhlin invariants to construct a ϕ so that the resulting homology spheres are never actual spheres. For details, we refer to [36, Section 6].

The proof of ii) uses mapping tori and Johnson homomorphisms. The latter can be defined in terms of pushing forward fundamental classes of mapping tori. If the monodromy of the mapping torus M is conjugate into the handlebody group, M is the boundary of a 4–manifold fibering over the circle. In this case, the Johnson homomorphism applied to the monodromy satisfies an additional vanishing result. A dimension argument is then used to show the existence of Torelli group elements which do not satisfy this obstruction.

The intersection of the Torelli group with the handlebody group can in fact be described a bit more explicitly. Namely, we have

Theorem 7.4 ([52])**.** *The intersection $\mathcal{I}_{g,1} \cap \mathcal{H}_{g,1}$ is generated by annular twists that cut off a genus* 1 *subsurface of Σ_g.*

This should be compared to the classical theorem of Johnson that the usual Torelli group is generated by genus 1 *boundary pair maps*, i.e. twist products $T_\alpha T_\beta^{-1}$ where α, β cut off a genus 1 subsurface. In other words, the intersection $\mathcal{I}_{g,1} \cap \mathcal{H}_{g,1}$ is generated by the subset of the canonical generating set of $\mathcal{I}_{g,1}$ which is contained in $\mathcal{H}_{g,1}$.

Question 7.5. *Can one develop a Johnson theory for the handlebody group?*

The first stage of this question – an (infinite) generating set for the Johnson kernel – is ongoing work in progress of the author and Andy Putman.

8. Algebraic Properties

In this section, we collect some algebraic properties of the handlebody group. As a first step, note that simply being a subgroup of a surface mapping class group automatically implies many interesting properties for $\mathcal{H}_g$. For example

- $\mathcal{H}_g$ satisfies the Tits alternative.
- $\mathcal{H}_g$ is virtually torsion free.
- $\mathcal{H}_g$ is residually finite.
- Non-elementary subgroups of $\mathcal{H}_g$ cannot be lattices in higher rank semi-simple Lie groups.

However, even many algebraic properties which are not automatically inherited by subgroups are shared between $\mathcal{H}_g$ and $\mathrm{Mcg}(\Sigma_g)$.

8.1. Generation.

We begin with the following classical result.

Theorem 8.1 ([54, 55])**.** *For any genus $g \geq 2$, the handlebody group is finitely generated and finitely presented.*

Suzuki and Wajnryb produce explicit generating sets and presentations, using actions on connected and simply connected complexes respectively. A fairly simple proof of finite generation can be given inductively, following the same strategy as for the mapping class groups. To this end, one uses the disk graph defined in Section 2. The quotient of $\mathcal{D}(V)$ by $\mathcal{H}_g$ is finite, and so finite generation of $\mathcal{H}_g$ is implied by finite generation of stabilisers of vertices in $\mathcal{D}(V)$ – in other words, stabilisers of meridians. If δ is a meridian, let V' be the (spotted, possibly disconnected) handlebody obtained by cutting V at a disk bounded by δ. There is a surjective map

$$\mathrm{Mcg}(V') \to \mathrm{Stab}_{\mathrm{Mcg}(V)}(\delta)$$

of the handlebody group of V' to the stabiliser of δ. V' is a handlebody, each component of which has genus strictly smaller than g, but more spots. As we discussed in Section 3, the kernel of the spot-forgetting map is the same as the kernel of the boundary-forgetting map in the case of surface mapping class groups, and therefore finitely generated. Hence we may, by induction on the genus, assume that $\mathrm{Mcg}(V')$ is finitely generated, and conclude that the same is true for $\mathrm{Stab}_{\mathrm{Mcg}(V)}(\delta)$.

8.2. **Subgroups.** In the study of mapping class groups, one frequently uses free Abelian and finite subgroups. We understand the structure of both classes in the handlebody group.

We begin with Abelian subgroups. Using meridional Dehn twists about a pants decomposition, we find that $\mathcal{H}_g$ contains subgroups isomorphic to $\mathbb{Z}^{3g-3}$. Since this is the maximal rank even for the surface mapping class group, we have

Theorem 8.2. *The maximal rank of a free Abelian subgroup in $\mathcal{H}_g$ is $3g-3$.*

Torsion on the other hand is more restrictive in the handlebody group. Maybe the most important difference is that torsion in the handlebody group has much smaller order.

Theorem 8.3 ([57, 58])**.** *Every finite subgroup of $\mathcal{H}_g$ has order at most $12(g-1)$.*

In fact, the finite groups of maximal size can also be determined more explicitly (compare [57]). A geometric reason for the smaller size of torsion can be seen in the following characterisation.

Lemma 8.4 (e.g. [24, Corollary 6.2])**.** *If $G < \mathcal{H}_g$ is finite, then there is a multicurve consisting of meridians which is preserved by G, and whose complements are spheres with boundary components (in other words, planar subsurfaces).*

Conversely, any finite $G < \mathrm{Mcg}(\Sigma_g)$ of this form is conjugate into the handlebody group.

We note that in surface mapping class groups that are finite order elements which are non-reducible, and thus there are finite cyclic subgroups of surface mapping class groups which are not conjugate into the handlebody group.

8.3. **Homological Properties.** We begin with the first homology of handlebody groups, where already a further difference to the mapping class group appears. Namely, the first homology of handlebody groups is always non-vanishing.

Theorem 8.5 ([55, 28, 27][3])**.** *We have*

$$H_1(\mathcal{H}_g, \mathbb{Z}) = \begin{cases} \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z} & \text{for } g = 1 \\ \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z} & \text{for } g = 2 \\ \mathbb{Z}/2\mathbb{Z} & \text{for } g > 2 \end{cases}$$

[3] The computation of the first homology in [55] contains a mistake, compare [28, Remark 3.5]

Recall that the *virtual cohomological dimension* is the cohomological dimension of a finite index torsion free subgroup. It is known that the virtual cohomological dimension of the mapping class group is $4g-5$ – in fact, the same is true for the handlebody group:

Theorem 8.6 ([25]). *The virtual cohomological dimension of the handlebody group is* $4g-5$.

The homology groups of $\mathcal{H}_g$ also stabilise, similar to the ones for surface mapping class groups, as shown by Hatcher and Wahl [19]. In fact, Hatcher announced a variant of the Madsen-Weiss theorem, computing the stable homology of $\mathcal{H}_g$ completely (as generated by the even Morita-Mumford-Miller classes).

From Theorem 8.5, it is immediate that the first Betti number of handlebody groups of genus ≥ 2 is zero. Just as for mapping class groups and outer automorphisms of free groups, it is a very interesting question to ask if this property remains true virtually, i.e. for finite index subgroups. It is known that $\mathrm{Out}(F_3)$ has a finite index subgroup which surjects to $\mathbb{Z}$ [13], and therefore $\mathcal{H}_3$ also has virtual positive first Betti number. However, we have the following, likely very hard question.

Question 8.7. *Does* $\mathcal{H}_g$, $g \geq 4$ *admit a finite index subgroup which surjects to* $\mathbb{Z}$*?*

8.4. **Homomorphisms.** Surface mapping class groups are known, by the work of Ivanov [29] and many others, to be rigid in the following sense: if $\Gamma_1, \Gamma_2 < \mathrm{Mcg}(\Sigma_g)$ are any finite index subgroups, and $\psi : \Gamma_1 \to \Gamma_2$ is any isomorphism of groups, then ψ is simply a conjugation.

A way to phrase this succinctly uses the notion of the *(abstract) commensurator* of a group G: $\mathrm{Comm}(G)$ consists of isomorphisms between arbitrary finite index subgroups $G_1, G_2 < G$, where two isomorphisms $\phi, \phi' : G_1 \to G_2$ are deemed equivalent if they coincide on a subgroup $G_3 < G_1 \cap G_2$ which is finite index in G_1, G_2[4].

In spite of $\mathcal{H}_g$ being an infinite index subgroup of $\mathrm{Mcg}(\Sigma_g)$, we still have rigidity for the handlebody groups as well.

Theorem 8.8 ([21]). *The abstract commensurator of* $\mathcal{H}_g$, $g \geq 3$ *is the handlebody group* $\mathcal{H}_g$ *(via its conjugation action on itself).*

In particular, $\mathrm{Out}(\mathcal{H}_g) = 1$. This consequence was known before, by a result of Korkmaz-Schleimer [35].

It might be interesting to study maps from handlebody groups to other groups, in particular mapping class groups. For example, there is the following "superrigidity type" question

[4]Note that since the intersection of finite index subgroups is finite index, the composition of two isomorphisms between finite index subgroups is well-defined up to the equivalence relation

Question 8.9. *Is there a finite index subgroup* $\Gamma < \mathcal{H}_g$ *and an inclusion* $f : \Gamma \to \mathrm{Mcg}(\Sigma')$ *which does not (virtually) extend to* $\mathrm{Mcg}(\Sigma_g)$*?*

9. Actions on measured laminations

In this section we summarise briefly some results on the action of the handlebody group on the sphere of projective measured laminations $\mathcal{PML}$. In a similar spirit as before, we will mainly focus on the differences between the actions of the full mapping class group and the handlebody group – showing that the handlebody group is a very "small" subgroup from a dynamical point of view.

We refer the reader to [11] for background on measured laminations. For our purposes it suffices to recall that $\mathcal{PML}(\Sigma_g)$ is homeomorphic to a sphere of dimension $6g-7$. The mapping class group acts on $\mathcal{PML}(\Sigma_g)$ by homeomorphisms. The action is minimal, i.e. every orbit is dense. Every simple closed curve on Σ_g also defines a point in $\mathcal{PML}(\Sigma_g)$, and the resulting set is dense. Also, geometric intersection number extends to a continuous function $i : \mathcal{ML} \times \mathcal{ML} \to \mathbb{R}$. For projective laminations λ, λ' the intersection number is not defined, but having $i(\lambda, \lambda') = 0$ is a well-defined property.

Define $\Lambda(\Sigma) \subset \mathcal{PML}$ to be the closure of the set of meridians, and

$$Z(\Lambda) = \{\mu \in \mathcal{PML} \mid \exists \lambda \in \Lambda : i(\mu, \lambda) = 0\}.$$

We then have

Theorem 9.1 ([41])**.** *i)* Λ *is the* limit set *of the action of the handlebody group on* $\mathcal{PML}(\Sigma_g)$*, i.e. the smallest nonempty closed subset of* $\mathcal{PML}(\Sigma_g)$ *invariant under* $\mathcal{H}_g$*.*
ii) The handlebody group acts properly discontinuously on $\mathcal{PML} - Z(\Lambda)$*.*

The set $\mathcal{PML} - Z(\Lambda)$ is sometimes called the *Masur domain*, and is of importance in the study of degenerations of Kleinian groups. Giving details on this connection would lead us too far afield; we refer the interested reader to e.g. [34] and the references therein for details.

The sphere of projective measured laminations carries a natural (Lebesgue) measure class which is preserved by the action of the mapping class group. With respect to this class, we have the following.

Theorem 9.2 ([41, 33] [38, Appendix][5])**.** *The set* Λ *has measure zero.*

The lamination point of view also lets us give yet another perspective on the membership problem. Namely, if $\phi \in \mathcal{H}_g$ is a pseudo-Anosov element, then the stable and unstable laminations of ϕ are contained in Λ: take any meridian δ, and consider $\phi^n(\delta)$ as $n \to \pm\infty$. These converge, by north-south-dynamics of pseudo-Anosov maps, to the stable and unstable laminations of ϕ, and are clearly sequences of meridians.

[5]Masur showed this result in genus $g = 2$, and Kerckhoff extended it to any genus – but his proof contained a gap which Gadre closed.

It is natural to ask if the converse is true as well. This turns out to be not quite true.

Theorem 9.3 ([3, 1]). *A pseudo-Anosov ϕ has the stable or unstable lamination in Λ if and only if it admits a power which partially extends to V_g. This means there is a compression body $C \subset V$ sharing a boundary with V_g, so that ϕ^n extends to C.*

In [3], Biringer, Johnson and Minsky show that the theorem is in fact optimal – there are elements whose powers extend, but they themselves do not, and there are elements whose (un)stable laminations lie in Λ, but no power extends to all of V. Also, we note that the theorem works in higher generality than handlebodies.

10. Geometric Properties

By Theorem 8.1, the handlebody group is finitely generated, and therefore it carries a metric which is unique up to quasi-isometry (any word metric defined by a finite generating set will do). In this section we discuss what is known about this coarse geometry for handlebody groups.

10.1. **Geometry of the disk graph.** In the case of the surface mapping class group, the corresponding coarse geometry of $\mathrm{Mcg}(\Sigma_g)$ is by now very well understood. A core tool is here the foundational work of Masur-Minsky, which uses the curve graphs of Σ_g and its subsurfaces to construct explicit quasi-geodesics in $\mathrm{Mcg}(\Sigma_g)$ and derive a distance formula in terms of subsurface projections.

To study handlebody groups, one is therefore first led to trying to understand the geometry of disk graphs which we have already encountered in Section 2. Their geometry is understood to a certain degree. The most naive approach to study $\mathcal{D}(V)$ is to consider it as a sub-graph of the curve graph $\mathcal{C}(\partial V)$. We have the following

Theorem 10.1 ([43]). *As a subset of $\mathcal{C}(\partial V)$, the disk graph $\mathcal{D}(V)$ is quasi-convex.*

However, recall that the curve (and disk) graphs are locally infinite, and therefore not proper metric spaces. Hence, quasi-convexity does not suffice to show that the inclusion is a quasi-isometric embedding, and tells us very little about the intrinsic geometry of $\mathcal{D}(V)$. In fact, the disk graph is arbitrarily badly distorted in the curve graph.

To describe this issue, will use the notion of subsurface projection developed in [40]. We refer the reader to the article by Masur and Minsky for details, and only briefly recall the necessary concepts here.

Let Y be a subsurface of a surface S, and let α be a simple closed curve in S which intersects Y nontrivially. Then there is a procedure called *subsurface projection* which associates to α an arc or simple closed curve in Y. The result of this projection is well-defined up to distance 2 in the arc or curve graph of Y. For the arc version, $\pi_Y(\alpha)$ is simply one of the arcs in $Y \cap \alpha$. To obtain the curve projection, one considers any intersection arc of α with Y, and completes this arc to a simple closed curve with a segment in ∂Y. We refer to Section 2 of [40] for details.

Example 10.2. *Consider a description of the handlebody of genus $2g$ as a trivial interval bundle $V = \Sigma_g^1 \times [0,1]$. The boundary of Σ_g^1 defines a simple closed curve δ on ∂V which is* disk-busting, *i.e. every meridian intersects δ (since, as we have seen before, the inclusion of Σ_g^1 into V is π_1–injective). We therefore have a subsurface projection map*

$$\pi : \mathcal{D}(V) \to \mathcal{A}(\Sigma_g^1)$$

from the disk complex of V to the arc complex of Σ_g^1 which associates to a meridian α the intersection of α with $\Sigma_g^1 \times \{1\}$ (which can be identified with one of the complementary components of δ). This map is clearly Lipschitz, as disjoint meridians map to disjoint arc systems. Since pseudo-Anosov mapping classes ψ of Σ_g^1 act with infinite diameter orbits on $\mathcal{A}(\Sigma_g^1)$, the corresponding handlebody group elements defined by such ψ act with infinite order orbits on $\mathcal{D}(V)$.

However, as mapping classes of ∂V, they are reducible (they fix δ!) and therefore have finite diameter orbits.

In spite of the problems of the embedding, the intrinsic geometry of the disk graph shows similarities to that of the curve graph.

Theorem 10.3 ([42])**.** *For any $g \geq 2$, the disk graph $\mathcal{D}(V)$ is Gromov hyperbolic.*

In fact, we can describe distances in $\mathcal{D}(V)$ using a Masur-Minsky type distance formula:

Theorem 10.4 ([42])**.** *We have, for large enough K,*

$$d_{\mathcal{D}(V)}(\alpha, \alpha') =_K \sum_{Y \text{ witness}} [d_{\mathcal{AC}(Y)}(\pi_Y(\alpha), \pi_Y(\alpha'))]_K.$$

Here, "$=_K$" denotes equality up to uniform additive and multiplicative constants. π_Y denotes subsurface projection, $d_{\mathcal{AC}(Y)}$ is the distance in the curve-and-arc-graph of Y, and $[\cdot]_K$ is the function which is x if the argument is at least K, and zero otherwise. For the definition of "witness", we refer to [42] (where they are called "holes", but the terminology "witness" is becoming more common). Briefly, a *witness* is subsurface for which the subsurface projection of meridians is always nonempty, and the diameter of the projection is large enough.

10.2. **Geometry of $\mathcal{H}_g$.** Even though the geometry of the "ambient" mapping class group is well-understood, the following theorem shows that this does not suffice to understand the geometry of $\mathcal{H}_g$. For the formal version, note that the *distortion function* of a finitely generated subgroup H of a finitely generated G is the function

$$D(n) = \max\{\|h\|_H \mid \|h\|_G \leq n\}.$$

The distortion function depends on the choice of word norms, but its growth type does not. We say that a subgroup is *undistorted* if the distortion function is linear, and distorted otherwise. In that case, the growth type of the distortion function gives an indication of how different the word norms are.

Theorem 10.5 ([15])**.** *For any $g \geq 2$, the handlebody group $\mathcal{H}_g$ is exponentially distorted in the mapping class group* $\mathrm{Mcg}(\Sigma_g)$*.*

One reason why the geometry of handlebody groups is very different to that of its ambient mapping class group is the following phenomenon, which shows that subsurface projections of meridians are (usually) not meridians.

Example 10.6. *Consider a handlebody $V = V_3$ of genus 3, and a separating meridian δ. Denote by Y the complementary component of δ which has genus 2, and by T the other component. The torus T contains a unique meridian, and let δ_1, δ_2 be two parallel, disjoint copies of this meridian in T. We choose two disjoint arcs a_1, a_2 so that one endpoint of a_i lies on δ_i, and the other on δ. Now, let $a \subset Y$ be any arc whose endpoints are also endpoints of a_1, a_2. Then*

$$\Gamma = \delta_1 \cup a_1 \cup a \cup a_2 \cup \delta_2$$

is connected, and intersects Y exactly in a. The boundary of a regular neighbourhood of Γ consists of three simple closed curves, two of which are parallel to δ_1, and the third of which we call γ_a.

The important observations are that γ_a is a meridian for any *choice of a, and $\pi_Y(\gamma_a) = a$. In particular, the subsurface projection of a meridian to Y can be* any arc whatsoever*.*

As a consequence of Theorem 10.5, there is no reason to expect that the word geometry of $\mathcal{H}_g$ should be comparable to that of the surface mapping class group. Direct applications of geometric methods from the mapping class group will usually yield, if they can be used, rather coarse results. In fact, it seems as if the word geometry of $\mathcal{H}_g$ is more similar to that of $\mathrm{Out}(F_g)$. For example we have,

Theorem 10.7 ([16, 17])**.** *For $g \geq 3$, the Dehn function of $\mathcal{H}_g$ is exponential.*

Surface mapping class group have quadratic Dehn functions since it they are automatic [50], while it is known that $\mathrm{Out}(F_n)$ has an exponential Dehn function for $n \geq 3$ [18, 7].

More intricate geometric questions about $\mathcal{H}_g$ are currently completely open. To mention one basic question, note that it is an immediate consequence of Masur-Minsky's distance formula [40] that the stabiliser of any simple closed curve is undistorted in the mapping class group. Correspondingly, let us ask

Question 10.8. *Which stabilisers of curves on ∂V_g are undistorted?*

It is very likely that stabilisers of meridians are undistorted – while stabilisers of a curve α mapping in $\pi_1(V_g)$ to a generator of a rank 1 free factor are exponentially distorted [23].

The reason these questions about stabilisers are interesting is due to work of Handel and Mosher [18]: they show that the stabiliser of a rank 1 free factor is exponentially distorted in $\mathrm{Out}(F_g)$ for $g \geq 3$, while stabilisers of free splittings are undistorted – which again suggests that $\mathcal{H}_g$ is geometrically more akin to outer automorphism groups of free groups rather than surface mapping class groups.

References

[1] Robert Ackermann. An alternative approach to extending pseudo-Anosovs over compression bodies. *Algebr. Geom. Topol.*, 15(4):2383–2391, 2015.

[2] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math. (2)*, 135(1):1–51, 1992.

[3] Ian Biringer, Jesse Johnson, and Yair Minsky. Extending pseudo-Anosov maps into compression bodies. *J. Topol.*, 6(4):1019–1042, 2013.

[4] Joan S. Birman. On the equivalence of Heegaard splittings of closed, orientable 3-manifolds. pages 137–164. Ann. of Math. Studies, No. 84, 1975.

[5] F. Bonahon. Ribbon fibred knots, cobordism of surface diffeomorphisms and pseudo-Anosov diffeomorphisms. *Math. Proc. Cambridge Philos. Soc.*, 94(2):235–251, 1983.

[6] Francis Bonahon. Cobordism of automorphisms of surfaces. *Ann. Sci. École Norm. Sup. (4)*, 16(2):237–270, 1983.

[7] Martin R. Bridson and Karen Vogtmann. The Dehn functions of $Out(F_n)$ and $Aut(F_n)$. *Ann. Inst. Fourier (Grenoble)*, 62(5):1811–1817, 2012.

[8] A. J. Casson and C. McA. Gordon. A loop theorem for duality spaces and fibred ribbon knots. *Invent. Math.*, 74(1):119–137, 1983.

[9] A. J. Casson and D. D. Long. Algorithmic compression of surface automorphisms. *Invent. Math.*, 81(2):295–303, 1985.

[10] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.

[11] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's work on surfaces*, volume 48 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 2012. Translated from the 1979 French original by Djun M. Kim and Dan Margalit.

[12] H. B. Griffiths. Automorphisms of a 3-dimensional handlebody. *Abh. Math. Sem. Univ. Hamburg*, 26:191–210, 1963/1964.

[13] Fritz Grunewald and Alexander Lubotzky. Linear representations of the automorphism group of a free group. *Geom. Funct. Anal.*, 18(5):1564–1608, 2009.

[14] Richard Hain. Relative weight filtrations on completions of mapping class groups. In *Groups of diffeomorphisms*, volume 52 of *Adv. Stud. Pure Math.*, pages 309–368. Math. Soc. Japan, Tokyo, 2008.

[15] Ursula Hamenstädt and Sebastian Hensel. The geometry of the handlebody groups I: distortion. *J. Topol. Anal.*, 4(1):71–97, 2012.

[16] Ursula Hamenstädt and Sebastian Hensel. Isoperimetric inequalities for the handlebody groups. *Trans. Amer. Math. Soc.*, 365(10):5313–5327, 2013.

[17] Ursula Hamenstädt and Sebastian Hensel. The geometry of the handlebody groups II: Dehn function. In preparation., 2017.

[18] Michael Handel and Lee Mosher. Lipschitz retraction and distortion for subgroups of $\mathrm{Out}(F_n)$. *Geom. Topol.*, 17(3):1535–1579, 2013.

[19] Allen Hatcher and Nathalie Wahl. Stabilization for mapping class groups of 3-manifolds. *Duke Math. J.*, 155(2):205–269, 2010.

[20] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001.

[21] Sebastian Hensel. Rigidity and flexibility for the handlebody group. *Comment. Math. Helv.*, (to appear).

[22] Sebastian Hensel. $\mathrm{Out}(f_n)$ does not virtually lift to the handlebody group. note available at www.math.uni-bonn.de/people/hensel/research, 2017.

[23] Sebastian Hensel. Stabilisers in the handlebody group. In preparation., 2017.

[24] Sebastian Hensel, Damian Osajda, and Piotr Przytycki. Realisation and dismantlability. *Geom. Topol.*, 18(4):2079–2126, 2014.

[25] Susumu Hirose. Some homological invariants of the mapping class group of a three-dimensional handlebody. *Tohoku Math. J. (2)*, 55(4):543–549, 2003.

[26] Susumu Hirose. The action of the handlebody group on the first homology group of the surface. *Kyungpook Math. J.*, 46(3):399–408, 2006.

[27] Susumu Hirose. Abelianization and Nielsen realization problem of the mapping class group of a handlebody. *Geometriae Dedicata*, 157(1):217225, Apr 2012.

[28] Tomohiko Ishida and Masatoshi Sato. A twisted first homology group of the handlebody mapping class group. *Osaka J. Math.*, 54(3):587–619, 07 2017.

[29] Nikolai V. Ivanov. Automorphism of complexes of curves and of Teichmüller spaces. *Internat. Math. Res. Notices*, (14):651–666, 1997.

[30] William Jaco. *Lectures on three-manifold topology*, volume 43 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I., 1980.

[31] Klaus Johannson. *Topology and combinatorics of 3-manifolds*, volume 1599 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1995.

[32] Jamie Bradley Jorgensen. *Surface Homeomorphisms That Do Not Extend to Any Handlebody and the Johnson Filtration*. PhD thesis, Rice University, 2008.

[33] Steven P. Kerckhoff. The measure of the limit set of the handlebody group. *Topology*, 29(1):27–40, 1990.

[34] Inkang Kim, Cyril Lecuire, and Ken'ichi Ohshika. Convergence of freely decomposable Kleinian groups. *Invent. Math.*, 204(1):83–131, 2016.

[35] Mustafa Korkmaz and Saul Schleimer. Automorphisms of the disk complex. *arXiv:0910.2038*, 2009.

[36] Christopher J. Leininger and Alan W. Reid. The co-rank conjecture for 3-manifold groups. *Algebr. Geom. Topol.*, 2:37–50, 2002.

[37] E. Luft. Actions of the homeotopy group of an orientable 3-dimensional handlebody. *Math. Ann.*, 234(3):279–292, 1978.

[38] Martin Lustig and Yoav Moriah. Are large distance Heegaard splittings generic? *J. Reine Angew. Math.*, 670:93–119, 2012. With an appendix by Vaibhav Gadre.

[39] Martin Lustig and Ulrich Oertel. Invariant laminations for diffeomorphisms of handlebodies. *Q. J. Math.*, 51(4):485–507, 2000.

[40] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.

[41] Howard Masur. Measured foliations and handlebodies. *Ergodic Theory Dynam. Systems*, 6(1):99–116, 1986.

[42] Howard Masur and Saul Schleimer. The geometry of the disk complex. *J. Amer. Math. Soc.*, 26(1):1–62, 2013.

[43] Howard A. Masur and Yair N. Minsky. Quasiconvexity in the curve complex. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 309–320. Amer. Math. Soc., Providence, RI, 2004.

[44] Darryl McCullough. Twist groups of compact 3-manifolds. *Topology*, 24(4):461–474, 1985.

[45] Darryl McCullough. Virtually geometrically finite mapping class groups of 3-manifolds. *J. Differential Geom.*, 33(1):1–65, 1991.

[46] Darryl McCullough. Homeomorphisms which are Dehn twists on the boundary. *Algebr. Geom. Topol.*, 6:1331–1340, 2006.

[47] Darryl McCullough and Andy Miller. Homeomorphisms of 3-manifolds with compressible boundary. *Mem. Amer. Math. Soc.*, 61(344):xii+100, 1986.

[48] D. R. McMillan, Jr. Homeomorphisms on a solid torus. *Proc. Amer. Math. Soc.*, 14:386–390, 1963.

[49] Shigeyuki Morita. Casson's invariant for homology 3-spheres and characteristic classes of surface bundles. I. *Topology*, 28(3):305–323, 1989.

[50] Lee Mosher. Mapping class groups are automatic. *Ann. of Math. (2)*, 142(2):303–384, 1995.

[51] Ulrich Oertel. Automorphisms of three-dimensional handlebodies. *Topology*, 41(2):363–410, 2002.

[52] Genki Omori. A small normal generating set for the handlebody subgroup of the Torelli group. *arXiv:1607.06553*, 2016.

[53] Wolfgang Pitsch. Integral homology 3-spheres and the Johnson filtration. *Trans. Amer. Math. Soc.*, 360(6):2825–2847, 2008.

[54] Shin'ichi Suzuki. On homeomorphisms of a 3-dimensional handlebody. *Canad. J. Math.*, 29(1):111–124, 1977.

[55] Bronisł aw Wajnryb. Mapping class group of a handlebody. *Fund. Math.*, 158(3):195–228, 1998.

[56] Heiner Zieschang. über einfache Kurven auf Vollbrezeln. *Abh. Math. Sem. Univ. Hamburg*, 25:231–250, 1961/1962.

[57] Bruno Zimmermann. über Homöomorphismen n-dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen. *Comment. Math. Helv.*, 56(3):474–486, 1981.

[58] Bruno Zimmermann. Finite group actions on handlebodies and equivariant Heegaard genus for 3-manifolds. *Topology Appl.*, 43(3):263–274, 1992.