Geometry & Topology

Volume 27 (2023)

Large-scale geometry of big mapping class groups

Kathryn Mann
Kasra Rafi

msp

Geometry & Topology **27**:6 (2023) 2237–2296
DOI: 10.2140/gt.2023.27.2237
Published: 25 August 2023

Large-scale geometry of big mapping class groups

KATHRYN MANN
KASRA RAFI

We study the large-scale geometry of mapping class groups of surfaces of infinite type, using the framework of Rosendal for coarse geometry of non-locally-compact groups. We give a complete classification of those surfaces whose mapping class groups have local *coarse boundedness* (the analog of local compactness). When the end space of the surface is countable or *tame*, we also give a classification of those surfaces where there exists a coarsely bounded generating set (the analog of finite or compact generation, giving the group a well-defined quasi-isometry type) and those surfaces with mapping class groups of bounded diameter (the analog of compactness).

We also show several relationships between the topology of a surface and the geometry of its mapping class groups. For instance, we show that *nondisplaceable subsurfaces* are responsible for nontrivial geometry and can be used to produce unbounded length functions on mapping class groups using a version of subsurface projection; while *self-similarity* of the space of ends of a surface is responsible for boundedness of the mapping class group.

57K20, 57M07

1 Introduction

Mapping class groups of surfaces of infinite type (with infinite genus or infinitely many ends) form a rich class of examples of non-locally-compact Polish topological groups. These "big" mapping class groups can be seen as natural generalizations of, or limit objects of, the mapping class groups of finite type surfaces, and also arise naturally in the study of laminations and foliations, and the dynamics of group actions on finite type surfaces.

Several recent papers have studied big mapping class groups through their actions on associated combinatorial structures such as curve or arc complexes; see for instance Aramayona, Fossas and Parlier [1], Bavard, Dowdall and Rafi [4] and Durham, Fanoni and Vlamis [7]. From this perspective, an important problem is to understand whether a given mapping class group admits a *metrically nontrivial* action on such a space, namely, an action with unbounded orbits. It is our observation that this should be framed as part of a larger question, one of the *coarse* or *large-scale geometry* of big mapping class groups. This is the goal of the present work.

However, describing the large-scale structure of big mapping class groups — or even determining whether this notion makes sense — is a nontrivial problem, as standard tools of geometric group theory apply only to locally compact, compactly generated groups, and big mapping class groups do not fall in this category. Instead, we use recent work of Rosendal [18] that extends the framework of geometric group theory to a broader class of topological groups, using the notion of *coarse boundedness*.

Definition 1.1 Let G be a topological group. A subset $A \subset G$ is *coarsely bounded*, abbreviated *CB*, if every compatible left-invariant metric on G gives A finite diameter. A group is *locally CB* if it admits a CB neighborhood of the identity, and *CB generated* if it admits a CB generating set.[1]

To give an example, in a locally compact group, the CB sets are precisely the compact ones. As is well known, among locally compact groups, those who admit a CB (ie compact) generating set have a well-defined quasi-isometry type, namely that given by the word metric with respect to any compact generating set (the discrete, finitely

[1] In Rosendal [17] and much earlier work, this condition is called (OB), for *orbites bornées*, as it is equivalent to the condition that for any continuous action of G on a metric space X by isometries, the diameter of every orbit $A \cdot x$ is bounded. *Coarsely bounded* appears in [18]; we prefer this terminology as it is more suggestive of the large-scale geometric context.

generated groups are a special case of this). Extending this notion, one says that a left-invariant metric d on a group G is said to be *maximal* if for any other left-invariant metric d' there exist constants C and K such that

$$d'(f,g) \le Kd(f,g) + C$$

holds for all $f, g \in G$. If G admits a maximal metric, then the coarse equivalence class of this metric gives G a well-defined quasi-isometry type. Rosendal shows the following.

Theorem 1.2 (Rosendal [18, Theorem 1.2]) *Let G be a Polish group. The following are equivalent*:

(i) *G is generated by a CB subset.*

(ii) *G admits a maximal left-invariant metric*, *among the left-invariant metrics which generate its topology.*

(iii) *G has a CB neighborhood of the identity and cannot be expressed as the union of a countable chain of proper open subgroups.*

Furthermore, the word metric from any CB generating set is in the quasi-isometry class of the maximal metric, giving a concrete description of the geometry of the group [18, Proposition 2.5].

In this work, we show that among the big mapping class groups there is a rich family of examples to which Rosendal's theory applies, and give the first steps towards a classification of such groups up to quasi-isometry.

1.1 Main results

For simplicity, we assume all surfaces are oriented and have empty boundary, and all homeomorphisms are orientation-preserving. (The cases of nonorientable surfaces, and those with finitely many boundary components can be approached using essentially the same tools.)

Summary We give a complete classification of surfaces Σ for which $\mathrm{Map}(\Sigma)$ is *locally CB* (Theorem 1.4). By Theorem 1.2, this is necessary for the group to be generated by a CB subset, but these are not equivalent. Under mild hypotheses, we give a full classification of those surfaces which are *CB generated* and therefore have a well-defined quasi-isometry type (Theorem 1.6), as well as those which are *globally CB*, ie have trivial QI type (Theorem 1.7).

To give the precise statements, we need to recall the classification of surfaces and state two key definitions.

End spaces Recall that topological spaces admit a standard compactification by a *space of ends*. By a theorem of Richards [16] orientable, boundaryless, infinite-type surfaces are completely classified by the following data: the genus (possibly infinite), the *space of ends* E, which is a totally disconnected, separable, metrizable topological space, and the subset of ends E^G that are accumulated by genus, which is a closed subset of E. Every such pair (E, E^G) occurs as the space of ends of some surface, with $E^G = \varnothing$ if and only if the surface has finite genus. We call a pair (E, E^G) *self-similar* if for any decomposition $E = E_1 \sqcup E_2 \sqcup \cdots \sqcup E_n$ of E into pairwise disjoint clopen sets, there exists a clopen set D contained in some E_i such that the pair $(D, D \cap E^G)$ is homeomorphic to (E, E^G).

Complexity A key tool in our classification is the following ranking of the "local complexity" of an end, which (as we show) gives a partial order on equivalence classes of ends.

Definition 1.3 For $x, y \in E$, we say $x \preccurlyeq y$ if every neighborhood of y contains a homeomorphic copy of a neighborhood of x. We say x and y are *equivalent* if $x \preccurlyeq y$ and $y \preccurlyeq x$.

We show that this order has maximal elements (Proposition 4.7), and for A a clopen subset of E, we denote the maximal ends of A by $\mathcal{M}(A)$.

The following theorem gives the classification of locally CB mapping class groups. While the statement is technical, it is easy to apply in specific examples. For instance, the surfaces in Figure 1, left, satisfy the conditions, while those on the right fail to have CB mapping class group.

Theorem 1.4 (classification of locally CB mapping class groups) Map(Σ) *is locally CB if and only if there is a finite-type surface $K \subset \Sigma$ with the following properties*:

(i) *Each complementary region of K has one or infinitely many ends and infinite or zero genus.*

(ii) *The complementary regions of K partition E into clopen sets, indexed by finite sets $\mathcal{A}$ and $\mathcal{P}$ such that*

- *each $A \in \mathcal{A}$ is self-similar, with $\mathcal{M}(A) \subset \mathcal{M}(E)$ and $\mathcal{M}(E) \subset \bigsqcup_{A \in \mathcal{A}} \mathcal{M}(A)$,*

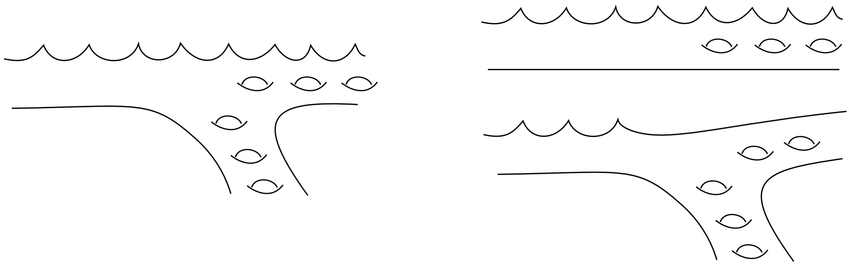

Figure 1: By Theorem 1.4, the surface on the left has a locally CB mapping class group and those on the right do not. All have $\mathcal{P} = \varnothing$.

- *each* $P \in \mathcal{P}$ *is homeomorphic to a clopen subset of some* $A \in \mathcal{A}$, *and*
- *for any* $x_A \in \mathcal{M}(A)$ *and any neighborhood* V *of the end* x_A *in* Σ, *there is* $f_V \in \mathrm{Homeo}(\Sigma)$ *such that* $f_V(V)$ *contains the complementary region to* K *with end set* A.

Moreover, in this case the set $\mathcal{V}_K := \{g \in \mathrm{Homeo}(\Sigma) : g|_K = \mathrm{id}\}$ *is a CB neighborhood of the identity.*

In order to illustrate Theorem 1.4 and motivate the conditions in the next two classification theorems, we now state results in the much simpler special case when Σ has genus zero and countable end space.

Special case: E countable, genus zero If E is a countable set and $E^G = \varnothing$, a classical result of Mazurkiewicz and Sierpinski [13] states that there exists a countable ordinal α such that E is homeomorphic to the ordinal $\omega^\alpha n + 1$ equipped with the order topology. Thus, any $x \in E$ is locally homeomorphic to $\omega^\beta + 1$ for some $\beta \le \alpha$ (here β is the Cantor–Bendixon rank of the point x). In this case, our partial order $\prec$ agrees with the usual ordering of the ordinal numbers, points are equivalent if and only if they are locally homeomorphic, and we have the following.

Theorem 1.5 (special case of Theorems 1.4, 1.6 and 1.7) *Suppose* Σ *is an infinite-type surface of genus zero with* $E \cong \omega^\alpha n + 1$. *Then*:

(i) $\mathrm{Map}(\Sigma)$ *is CB if and only if* $n = 1$; *in this case* E *is self-similar.*

(ii) *If* $n \ge 2$ *and* α *is a successor ordinal*, *then* $\mathrm{Map}(\Sigma)$ *is locally CB and generated by a CB set*, *but admits a surjective homomorphism to* $\mathbb{Z}$, *so is not globally CB.*

(iii) *If* $n \ge 2$ *and* α *is a limit ordinal*, *then* $\mathrm{Map}(\Sigma)$ *is locally CB*, *but not generated by any CB set.*

Classification: general case One cannot hope for such a clean statement as that of Theorem 1.5 to hold in general, since there is no similarly clean classification of end spaces. In fact, even in the genus-zero case, classifying possible end spaces E (ie closed subsets of Cantor sets) up to homeomorphism is a difficult and well-studied problem, equivalent to the classification problem for countable Boolean algebras.[2] Ketonen [10] gives some description and isomorphism invariants. In practice these invariants are difficult to use, and yet they are in some sense an optimal classification, as Carmelo and Gao show in [5] that the isomorphism relation is Borel complete. Our definition of the partial order $\preccurlyeq$ allows us to sidestep the worst of these issues.

For technical reasons, the order is better behaved under a weak hypothesis on the topology of the end space, which we call "tameness". See Section 6 for motivation and the definition. To our knowledge, tame surfaces include all concrete examples studied thus far in the literature, including the mapping class groups of some specific infinite-type surfaces in Aramayona, Patel and Vlamis [2], Bavard [3] and Fanoni, Hensel and Vlamis [8], and the discussion of geometric or dynamical properties of various translation surfaces of infinite type in Chamanara [6], Hooper [9] and Randecker [15]. Although nontame examples do exist (see Example 6.13), there are no known nontame surface that have a well-defined quasi-isometry type (Problem 6.12). Under this hypothesis, we can give a complete classification of surfaces with a well-defined QI type, and those with a trivial QI type, as follows.

Theorem 1.6 (classification of CB generated mapping class groups) *For a tame surface Σ with locally (but not globally) CB mapping class group,* $\mathrm{Map}(\Sigma)$ *is CB generated if and only if E is* ***finite rank*** *and not of* ***limit type****.*

Theorem 1.7 (classification of globally CB mapping class groups) *Suppose Σ is either tame or has countable end space. Then* $\mathrm{Map}(\Sigma)$ *is CB if and only if Σ has infinite or zero genus and E is self-similar or a variant of this called "telescoping". The telescoping case occurs only when E is uncountable.*

Finite rank, loosely speaking, means that finite-index subgroups of $\mathrm{Map}(\Sigma)$ do not admit surjective homomorphisms to $\mathbb{Z}^n$ for arbitrarily large n. *Limit type* refers to behavior of equivalence classes for the partial order that mimics the behavior of limit ordinals in the special countable case stated above; see Section 6.2. *Telescoping* is a slightly broader notion of homogeneity or local similarity of an end space. Informally

[2] By Stone duality, totally disconnected, separable, compact sets are in one-to-one correspondence with countable Boolean algebras.

speaking, *self-similar* sets either appear very homogeneous (eg a Cantor set) or may have one "special" point, any neighborhood of which contains a copy of the whole set — for instance, a countable set with a single accumulation point is self-similar. Telescoping is a generalization that allows for two special points. Further motivation and a precise definition are given in Section 3.2.

Key tool: nondisplaceable subsurfaces The following tool is of independent interest and provides an easily employable criterion to certify that a surface has non-CB mapping class group (or, equivalently, admits a continuous isometric action on a metric space with unbounded orbits).

Definition 1.8 A connected, finite-type subsurface S of a surface Σ is said to be *nondisplaceable* if $f(S) \cap S \neq \varnothing$ for each $f \in \mathrm{Homeo}(\Sigma)$. A nonconnected surface is nondisplaceable if, for every $f \in \mathrm{Homeo}(\Sigma)$, there are connected components S_i and S_j of S such that $f(S_i) \cap S_j \neq \varnothing$.

Theorem 1.9 *If* Σ *is a surface that contains a nondisplaceable finite-type subsurface, then* $\mathrm{Map}(\Sigma)$ *is not globally CB.*

A key ingredient of the proof is *subsurface projection*, a familiar tool from the study of mapping class groups of finite-type surfaces, introduced by Masur and Minsky [12].

Theorem 1.9 immediately gives many examples of surfaces whose mapping class groups are not CB, and hence admit unbounded orbits on combinatorial complexes. For instance, any surface with finite but nonzero genus has this property. (See Theorem 1.5 below for a number of other easily described examples.) Theorem 1.9 also recovers, with a new proof, some of the work of Bavard in [3] and Durham, Fanoni and Vlamis in [7].

Outline

- Section 2 contains background information on standard mapping class group techniques, and the proof of Theorem 1.9.
- Section 3 gives two criteria for CB mapping class groups: self-similarity and telescoping end spaces. This is used later in the proof of the local and global CB classification theorems.
- Section 4 introduces the partial order on the end space and proves key properties of this relation, and a characterization of self-similar end spaces in terms of the partial order.

- Section 5 contains the proof of Theorem 1.4. This and the following section form the technical core of this work.
- Section 6 contains the proof of Theorem 1.6.
- Section 7 gives the proof of Theorem 1.7.

Acknowledgements Mann was partially supported by NSF grant DMS-1844516. Rafi was partially supported by NSERC Discovery grant RGPIN 06486. Part of this work was completed at the 2019 AIM workshop on surfaces of infinite type. We thank Camille Horbez, Justin Lanier, Brian Udall and Ferran Valdez for helpful comments on earlier versions of this paper, and the thoughtful work of the referee, which greatly improved the paper.

2 Proof of Theorem 1.9

In this section we prove that nondisplaceable finite-type subsurfaces of a surface Σ are responsible for nontrivial geometry in $\mathrm{Map}(\Sigma)$. We begin by introducing some notions from large-scale geometry and setting some conventions that will be useful throughout.

A criterion for coarse boundedness Recall that a subset $A \subset G$ of a metrizable, topological group is said to be *coarsely bounded* or *CB* if it has finite diameter in every compatible left-invariant metric on G. The following result gives an equivalent condition that is often easier to use in practice.

Theorem 2.1 (Rosendal [18, Proposition 2.7(5)]) *Let A be a subset of a Polish group G. The following are equivalent*:

(i) *A is coarsely bounded.*

(ii) *For every neighborhood $\mathcal{V}$ of the identity in G, there is a finite subset $\mathcal{F}$ and some $k \geq 1$ such that $A \subset (\mathcal{F}\mathcal{V})^k$.*

While Rosendal's theory is quite broadly applicable, mapping class groups (of any manifold) fall into the nicest family to which it applies, namely the completely metrizable or *Polish* groups. For any manifold M, the homeomorphism group $\mathrm{Homeo}(M)$ endowed with the compact-open topology is Polish, and hence also for any closed subset of M, the closed subgroups $\mathrm{Homeo}(M, X)$ and $\mathrm{Homeo}(M \,\mathrm{rel}\, X)$ of homeomorphisms, respectively preserving and pointwise fixing X. (In the mapping class groups context, X is typically taken to be the boundary of M or a set of marked points.) Thus, since the identity component $\mathrm{Homeo}_0(M, X)$ is a closed, normal subgroup, the quotient $\mathrm{Homeo}(M, X)/\mathrm{Homeo}_0(M, X)$ is also a Polish group.[3]

One useful tool for probing the geometry of a topological group is the following concept of a length function.

Definition 2.2 A *length function* on a topological group G is a continuous function $\ell\colon G \to [0, \infty)$ satisfying $\ell(g) = \ell(g^{-1})$, $\ell(\mathrm{id}) = 0$ and $\ell(gh) \leq \ell(g) + \ell(h)$ for all $g, h \in G$.

If ℓ is any length function, then for any $\epsilon > 0$ the set $\ell^{-1}([0, \epsilon))$ is a neighborhood of the identity in G. It follows from the criterion in Theorem 2.1 that ℓ is bounded on any CB subset.

Our strategy for the proof of Theorem 1.9 is to use the presence of a nondisplaceable subsurface to construct an unbounded length function. In order to do this, we introduce some notation and conventions which will also be used in later sections.

Surfaces: conventions The following conventions will be used throughout this work. Infinite-type surfaces, typically denoted by Σ, are assumed to be connected and orientable, and unless otherwise specified will be assumed to have empty boundary. By a *curve* in Σ we mean a free homotopy class of a nontrivial, nonperipheral, simple closed curve. In the first part of this section, when we talk about a subsurface $S \subset \Sigma$, we always assume that S is connected, has finite type and is *essential*, meaning that every curve in ∂S is nontrivial and nonperipheral in Σ. (Later we will broaden our discussion to include nonconnected subsurfaces.) As is standard, the *complexity* of a finite-type surface S is defined to be $\xi(S) = 3g_S + b_S + p_S$, where g_S is the genus, p_S is the number of punctures and b_S is the number of boundary components of S. *Finite type* simply means that all these quantities are finite.

The *intersection number* between two curves γ_1 and γ_2 is the usual geometric intersection number $\mathrm{i}(\gamma_1, \gamma_2)$, defined to be the minimal intersection number between representatives in the free homotopy classes of γ_1 and γ_2. To simplify the exposition going forward, we will fix a complete hyperbolic structure on Σ. Then every curve has a unique geodesic representative and the homotopy class of every subsurface has a unique representative that has geodesic boundary. A pair of curves γ_1 and γ_2 have

[3] For the case where M is a surface, that mapping class groups are Polish was also observed in [2] using the property that these groups are the automorphism groups of the curve complex of the surface.

disjoint representatives if and only if their geodesic representatives are disjoint. In this case, we say that $\mathrm{i}(\gamma_1, \gamma_2) = 0$. Otherwise, we say γ_1 intersects γ_2 and in this case, the intersection number $\mathrm{i}(\gamma_1, \gamma_2)$ is the cardinality of the intersection of their geodesic representatives.

Similarly, two subsurfaces R and S (or a subsurface R and geodesic γ) intersect if every subsurface homotopic to R intersects every subsurface homotopic to S (or analogously for γ), and this is equivalent to saying that the representatives of R and S with geodesic boundaries intersect each other. Hence, from now on, every time we consider a curve we assume it is a geodesic and every time we consider a subsurface we assume it has geodesic boundary. This allows us to unambiguously speak of intersections.

Definition 2.3 A finite-type, connected subsurface $S \subset \Sigma$ is *nondisplaceable* if $S \cap f(S) \neq \varnothing$ for all $f \in \mathrm{Map}(\Sigma)$.

Example 2.4 When Σ has positive, finite genus, any subsurface S whose genus matches that of Σ is nondisplaceable. This is because S contains nonseparating curves but $\Sigma - S$ does not. Since every image of S under a homeomorphism of Σ will also contain a nonseparating curve, it must intersect S.

Example 2.5 (nondisplaceable subsurfaces) It is also easy to construct examples of nondisplaceable surfaces using the topology of the end space. Suppose Σ has infinite end space, and Z is an invariant, finite set of ends of cardinality at least 3. Then any surface S which separates all the points of Z into different complementary regions will be nondisplaceable.

To give another prototypical example, if X and Y are disjoint, closed invariant sets of ends, with X homeomorphic to a Cantor set, then a subsurface homeomorphic to a pair of pants which contains points of X in two complementary regions, and all of Y in the third complementary region, will also be nondisplaceable.

Curve graphs and subsurface projections We recall some basic material on curve graphs. A reader unfamiliar with this machinery may wish to consult the introductory notes [19] or paper [11] for more details. As in the previous paragraph, we continue to assume here that surfaces are connected.

The *curve graph* $\mathcal{C}(S)$ of a finite-type surface S is a graph whose vertices are curves in S and whose edges are pairs of disjoint curves. We give each edge length one and denote the induced metric on $\mathcal{C}(S)$ by d_S. With this metric, as soon as $\xi(S) \geq 5$,

$(\mathcal{C}(S), d_S)$ has infinite diameter and is Gromov hyperbolic [11]. One can define curve graphs analogously for infinite-type surfaces, but these no longer have infinite diameter and we will use only the classical finite-type setting.

If Σ is any surface and $S \subset \Sigma$ a subsurface, there is a *projection map* π_S from the set of curves in Σ that intersect S to the set of subsets of $\mathcal{C}(S)$, defined as follows: for a curve γ, the intersection $\gamma \cap S$ of the geodesic γ with the subsurface S is either equal to γ (if $\gamma \subset S$) or is a union of arcs with endpoints in ∂S. For every such arc ω, one may perform a surgery between γ and ∂S to obtain in curve in S disjoint from ω, possibly in two different ways (the curve is a concatenation of one or two copies of ω and one or two arcs in ∂S). We define the projection $\pi_S(\gamma)$ to be γ if $\gamma \subset S$ and otherwise to be the *union* of curves associated to each arc on $\gamma \cap S$ obtained by surgery as above. When $\xi(S) \geq 5$, the set $\pi_S(\gamma)$ has diameter at most 2 in $\mathcal{C}(S)$; in fact, we have

$$\mathrm{i}(\gamma_1, \gamma_2) = 0 \implies \mathrm{diam}_S\, \pi_S(\gamma_1 \cup \gamma_2) \leq 2. \tag{1}$$

See [12, Lemma 2.2] for more details. In general, if μ is a subset of $\mathcal{C}(S)$, we define

$$\pi_S(\mu) = \bigcup_{\gamma \in \mu} \pi_S(\gamma).$$

The natural distance d_S on $\mathcal{C}(S)$ can be extended to a distance function on curves in Σ that intersect S via

$$d_S(\gamma_1, \gamma_2) = \max_{\alpha_i \in \pi_S(\gamma_i)} d_S(\alpha_1, \alpha_2).$$

The following result states that a bound on the intersection number between two curves gives a bound on their projection distance in any subsurface. This principle is well known and there are many similar results in the literature. We give a short proof with a suboptimal bound.

Lemma 2.6 *Let γ_1 and γ_2 be curves in Σ that intersect S. Then*

$$d_S(\gamma_1, \gamma_2) \leq 2 \log_2(\mathrm{i}(\gamma_1, \gamma_2) + 1) + 6. \tag{2}$$

Proof Let ω_1 be an arc in S that is a component of the restriction of γ_1 and let $\alpha_1 \in \pi_S(\gamma_1)$ be the curve in $\mathcal{C}(S)$ that is obtained by doing a surgery between ω_1 and the boundary of S. Then α_1 is a concatenation of one or two copies of ω_1 (depending on whether the endpoints of ω_1 are on the same boundary or different boundary components of S) and some arcs in ∂S. Similarly, let ω_2 be an arc in S that is a restriction of γ_2

and let α_2 be the associated curve in $\pi_S(\gamma_2)$. Then every intersection point between ω_1 and ω_2 results in 1, 2 or 4 intersection points between α_1 and α_2. Also, applying surgery between ω_2 and ∂S can result in two intersection points between α_2 and α_1 at each end of ω_2. Therefore,

$$\mathrm{i}(\alpha_1, \alpha_2) \le 4\,\mathrm{i}(\omega_1, \omega_2) + 4.$$

On the other hand, from [19, Lemma 1.21], we have

$$d_S(\alpha_1, \alpha_2) \le 2 \log_2(\mathrm{i}(\alpha_1, \alpha_2)) + 2.$$

Therefore,

$$d_S(\alpha_1, \alpha_2) \le 2 \log_2(4\,\mathrm{i}(\omega_1, \omega_2) + 4) + 2 \le 2 \log_2(\mathrm{i}(\gamma_1, \gamma_2) + 1) + 6,$$

which is as we claimed. □

The notions of distance d_S and intersection number can also be extended further to take finite sets of curves as arguments. If μ_i are finite sets of curves, we define

$$d_S(\mu_1, \mu_2) = \max_{\gamma_1 \in \mu_1, \gamma_2 \in \mu_s} d_S(\gamma_1, \gamma_2) \quad \text{and} \quad \mathrm{i}(\mu_1, \mu_2) = \max_{\gamma_1 \in \mu_1, \gamma_2 \in \mu_s} \mathrm{i}(\gamma_1, \gamma_2).$$

Using equation (2), for any finite subsets μ_1 and μ_1 of $\mathcal{C}(S)$, we have

$$d_S(\mu_1, \mu_2) \le 2 \log_2(\mathrm{i}(\mu_1, \mu_2) + 1) + 6. \tag{3}$$

Note that the triangle inequality still holds for this generalized distance d_S.

Construction of an unbounded length function We now proceed with the proof of Theorem 1.9. Let Σ be any surface, and let S be a nondisplaceable subsurface. Enlarge S if needed so that $\xi(S) \ge 5$ and so that S is connected. (In Section 2.1, we give an alternative modification for nonconnected subsurfaces that will be useful in later work.)

Let $\mathcal{I}$ denote the set of (isotopy classes of) subsurfaces of the same type as S, ie

$$\mathcal{I} = \{ f(S) \mid f \in \mathrm{Map}(\Sigma) \}.$$

As before, while $f(S)$ denotes only an isotopy class of a surface when $f \in \mathrm{Map}(\Sigma)$, the reader may identify it with an honest subsurface by taking the representative with geodesic boundary. Let μ_S be a filling set of curves in $\mathcal{C}(S)$, ie a set of curves with the property that every curve in S intersects some curve in μ.

For $R \in \mathcal{I}$ let $\mu_R = \pi_R(\mu_S)$. Note that this is always defined since μ_S fills S, and R intersects S because S was assumed nondisplaceable.

Now, define

$$\ell\colon \mathrm{Map}(\Sigma) \to \mathbb{Z} \quad \text{by} \quad \ell(\phi) = \max_{R \in \mathcal{I}} d_{\phi(R)}(\phi(\mu_R), \mu_{\phi(R)}).$$

Equivalently, we have

$$\ell(\phi) = \max_{T \in \mathcal{I}} d_T(\phi(\mu_{\phi^{-1}(T)}), \mu_T). \tag{4}$$

Note that ℓ is finite because, for every ϕ, the intersection number $\mathrm{i}(\mu_S, \phi(\mu_S))$ is a finite number. Hence, by equation (3), their projections to $\phi(R)$ lie at a bounded distance in $\mathcal{C}(R)$, with a bound that depends on ϕ alone, not on R.

The latter definition also makes it clear that $\ell(\phi) = \ell(\phi^{-1})$, since

$$\begin{aligned}
\ell(\phi^{-1}) &= \max_{T \in \mathcal{I}} d_T(\phi^{-1}(\mu_{\phi(T)}), \mu_T) \\
&= \max_{T \in \mathcal{I}} d_{\phi(T)}(\mu_{\phi(T)}, \phi(\mu_T)) \\
&= \max_{R = \phi(T) \in \mathcal{I}} d_R(\mu_R, \phi(\mu_{\phi^{-1}(R)})) = \ell(\phi).
\end{aligned}$$

We now check the triangle inequality. Let ψ and ϕ be given, and let $R \in \mathcal{I}$ be a surface such that $\ell(\psi\phi) = d_{\psi\phi(R)}(\psi\phi(\mu_R), \mu_{\psi\phi(R)})$. Then we have

$$\begin{aligned}
\ell(\psi\phi) &= d_{\psi\phi(R)}(\psi\phi(\mu_R), \mu_{\psi\phi(R)}) \\
&\le d_{\psi\phi(R)}(\psi\phi(\mu_R), \psi(\mu_{\phi(R)})) + d_{\psi\phi(R)}(\psi(\mu_{\phi(R)}), \mu_{\psi\phi(R)}) \\
&= d_{\phi(R)}(\phi(\mu_R), \mu_{\phi(R)}) + d_{\psi(Q)}(\psi(\mu_Q), \mu_{\psi(Q)}) \quad (\text{where } Q = \phi(R)) \\
&\le \ell(\phi) + \ell(\psi).
\end{aligned}$$

Continuity of ℓ as a function on $\mathrm{Map}(\Sigma)$ is a consequence of the following observation.

Observation *If ϕ and ϕ' agree on S, then $\ell(\phi) = \ell(\phi')$.*

Proof First note that for any $T \in \mathcal{I}$, we have $\mu_{\phi^{-1}(T)} \subset S \cap \phi^{-1}(T)$, hence

$$\phi(\mu_{\phi^{-1}(T)}) \subset \phi(S) \cap T.$$

Similarly,

$$\phi'(\mu_{\phi'^{-1}(T)}) \subset \phi'(S) \cap T.$$

But $\phi(S) \cap T = \phi'(S) \cap T$. In fact, $\phi(\mu_{\phi^{-1}(T)})$ is the projection of $\phi(\mu_S)$ to T and $\phi'(\mu_{\phi'^{-1}(T)})$ is the projection of $\phi'(\mu_S)$ to T. Since ϕ and ϕ' agree on S,

$$\phi(\mu_{\phi^{-1}(T)}) = \phi'(\mu_{\phi^{-1}(T)}),$$

from which it follows from (4) that $\ell(\phi) = \ell(\phi')$. □

Thus, the preimage of $\ell(\phi)$ under ℓ contains the open set consisting of mapping classes agreeing with ϕ on S. The remaining condition on a length function is that the length of identity should be zero. This is not a consequence of our definition, however we may simply redefine $\ell(\phi) = 0$ for all ϕ which restrict to the identity on S, without affecting the validity of the triangle-inequality computation above, as can be checked easily by hand.

To see that ℓ is unbounded, let $\phi \in \mathrm{Map}(\Sigma)$ be a homeomorphism that preserves S and such that the restriction $\phi|_S$ of ϕ to S is a pseudo-Anosov homeomorphism of S. Then, for any curve γ in S,

$$d_S(\gamma, \phi^n(\gamma)) \to \infty \quad \text{as } n \to \infty. \tag{5}$$

See eg [11] for details. Thus, ℓ is an unbounded length function, and so $\mathrm{Map}(\Sigma)$ is not coarsely bounded. □

2.1 Disconnected subsurfaces

While we have so far worked only with connected nondisplaceable subsurfaces, there is a natural generalization of the work above to nonconnected subsurfaces. This will be useful when we need to find a nondisplaceable subsurface that is disjoint from a given compact subset of Σ to determine if $\mathrm{Map}(\Sigma)$ is locally CB. The extension to this broader framework requires a little care since, if we simply take the definitions above verbatim, then the diameter of the curve graph $C(S)$ is finite as soon as S is not connected. However, the following minor adaptations allow our work above to carry through in this case.

Definition 2.7 A *disconnected finite-type subsurface* is a finite union of pairwise disjoint finite-type surfaces. We say such a subsurface $\bar{S}$ is *nondisplaceable* if, for any $f \in \mathrm{Map}(\Sigma)$ and any connected component S_i of $\bar{S}$, there is a connected component S_j of $\bar{S}$ such that $S_j \cap f(S_i) \neq \varnothing$.

We now show how to use such a disconnected surface $\bar{S}$ to construct a length function on $\mathrm{Map}(\Sigma)$. As before, let $\mathcal{I}$ denote the set of images of $\bar{S}$ under mapping classes, ie

$$\mathcal{I} = \{f(\bar{S}) \mid f \in \mathrm{Map}(\Sigma)\}.$$

If $\bar{S} = \bigsqcup_{i=1}^k S_i$, where S_i are the connected components, then an element $\bar{R}$ of $\mathcal{I}$ is simply the disjoint union of a set $\{R_1, \dots R_k\}$, where $R_i = f(S_i)$. Let $\mu_{\bar{S}}$ be a set of curves in $\bigcup_i \mathcal{C}(S_i)$ that fill every S_i. Keeping the notation from before, note that $\pi_{R_i}(\mu_{\bar{S}})$ is always defined since R_i intersects some S_j, and curves in $\mu_{\bar{S}}$ fill S_j. Now, define $\ell_{\bar{S}}\colon \mathrm{Map}(\Sigma) \to \mathbb{Z}$ by

$$\ell_{\bar{S}}(\phi) = \max_{\bar{R}\in\mathcal{I}} \max\{d_{\phi(R_i)}(\phi(\mu_{R_i}), \mu_{\phi(R)}) \mid R_i \text{ a component of } \bar{R}\}.$$

The same computation as in the connected case shows that $\ell_{\bar{S}}$ is finite, is continuous as a function on $\mathrm{Map}(\Sigma)$, and satisfies the triangle inequality with the same adjustment that $\ell_{\bar{S}}(\phi) = 0$ when ϕ is identity on $\bar{S}$. To see that $\ell_{\bar{S}}$ is unbounded, let $\phi \in \mathrm{Map}(\Sigma)$ be a homeomorphism that preserves $\bar{S}$ and such that the restriction $\phi|_{S_1}$ of ϕ to S_1 is a pseudo-Anosov homeomorphism of S_1. Since $\ell_{\bar{S}}$ is defined as a maximum of distances in various curve graphs, if ϕ has a positive translation length in $\mathcal{C}(S_1)$ (or in any $\mathcal{C}(S_i)$) then $\ell_{\bar{S}}(\phi^n) \to \infty$ as $n \to \infty$. This gives an alternative proof of Theorem 1.9 in the disconnected case, and the following more general statement:

Proposition 2.8 *If Σ contains a connected or disconnected, nondisplaceable, finite-type subsurface S such that each connected component of S has complexity at least* 5, *then there exists a length function ℓ defined on* $\mathrm{Map}(\Sigma)$ *such that the restriction of ℓ to mapping classes supported on S is unbounded.*

3 Self-similar and telescoping end spaces

In this section we give two topological conditions (in Propositions 3.1 and 3.5) that imply coarse boundedness of the mapping class group: *self-similarity* and *telescoping*.

3.1 Self-similar end spaces

Recall that a space of ends (E, E^G) is said to be *self-similar* if for any decomposition $E = E_1 \sqcup E_2 \sqcup \cdots \sqcup E_n$ of E into pairwise disjoint clopen sets, there exists a clopen set D in some E_i such that $(D, D \cap E^G)$ is homeomorphic to (E, E^G). There are many examples of such sets; a few basic ones are:

- E equal to a Cantor set, and E^G either empty, equal to E, or a singleton.
- E a countable set homeomorphic to $\omega^\alpha + 1$ with the order topology, for some countable ordinal α, and E^G the set of points of type $\omega^\beta + 1$ for all ordinals $\beta \geq \beta_0$, where β_0 is a some fixed ordinal.
- E the union of a countable set Q and a Cantor set where the sole accumulation point of Q is a point in the Cantor set, and $E^G = \bar{Q}$.

Convention Going forward, we drop the notation E^G, assuming that E comes with a designated closed subset of ends accumulated by genus, empty if the genus of Σ is finite, and that all homeomorphisms between sets or subsets of end spaces preserve (setwise) the ends accumulated by genus.

As E and E^G are totally disconnected spaces, we also make the following convention.

Convention For the remainder of this work, when we speak of a *neighborhood* in an end space E, we always mean a *clopen neighborhood*.

Proposition 3.1 (self-similar implies CB) *Let Σ be a surface of infinite or zero genus. If the space of ends of Σ is self-similar, then* $\mathrm{Map}(\Sigma)$ *is CB.*

Note that finite, nonzero-genus surfaces cannot have CB mapping class groups by Example 2.4, so Proposition 3.1 is optimal in this sense. Note also that the proposition holds for finite-type surfaces as well, but the only applicable example is the once-punctured sphere, which has trivial mapping class group.

Proof of Proposition 3.1 Let Σ be an infinite-type surface satisfying the hypotheses of the proposition, and let $\mathcal{V}$ be a neighborhood of the identity in $\mathrm{Map}(\Sigma)$. Then there exists some finite-type subsurface S such that $\mathcal{V}$ contains the open set $\mathcal{V}_S$ consisting of mapping classes of homeomorphisms that restrict to the identity on S. By Theorem 2.1, it suffices to find a finite set $\mathcal{F} \subset \mathrm{Map}(\Sigma)$ and $k \in \mathbb{N}$ (which are allowed to depend on $\mathcal{V}_S$, hence on S) such that $\mathrm{Map}(\Sigma) = (\mathcal{F}\mathcal{V}_S)^k$. Enlarging S (and therefore shrinking $\mathcal{V}_S$) if needed, we may assume that each connected component of $\Sigma - S$ is of infinite type.

Since the proof is somewhat technical, we begin with an outline. The first step is to find a suitable homeomorphism f of Σ so that $f(S) \cap S = \varnothing$, and declare $\mathcal{F}$ to be the finite set consisting of f and f^{-1}. Now suppose one is given $g \in \mathrm{Map}(\Sigma)$. Obviously if g restricts to the identity on S, then $g \in \mathcal{V}_S$ and we are done (in fact $k = 1$ would work). If instead g restricted to the identity on $f(S)$, then we would have $g \in f\mathcal{V}_S f^{-1}$, and again are done, and could have taken $k = 2$. The general philosophy of the proof is to cleverly choose f so that *every* mapping class g can be written as a product of at most three elements which are either the identity on S or on $f(S)$, and use this to get the desired bound on k. In practice, we do this by finding an additional homeomorphic copy of S in Σ. Now we provide the details.

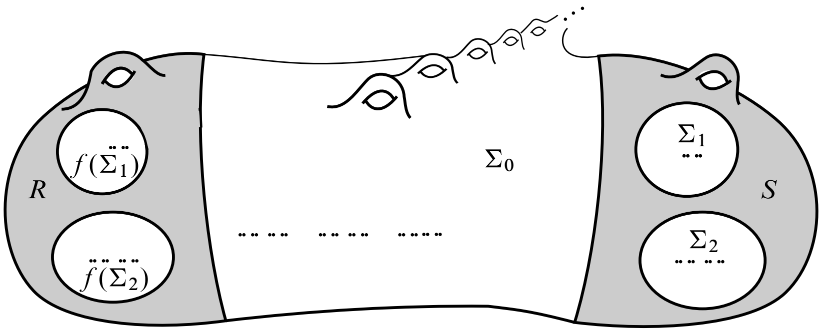

Figure 2: A homeomorphic copy R of S contained in the complementary region Σ_0.

The connected components of $\Sigma - S$, together with the finite set P of punctures of S, partition E into clopen sets. Let

$$E = E_0 \sqcup E_1 \sqcup \cdots \sqcup E_n \sqcup P$$

denote this decomposition, and let Σ_i denote the connected component of $\Sigma - S$ containing E_i. Since S is of finite type, $E^G \cap P = \varnothing$. Since E is self-similar, one of the E_i contains a copy of E. Without loss of generality, we assume this is E_0, the set of ends of Σ_0; thus we may write $E_0 = E' \sqcup D$, where E' is homeomorphic to E. The next lemma asserts that we may find a surface $R = f(S)$ as depicted in Figure 2.

Lemma 3.2 *With the notation above, there exists* $f \in \mathrm{Homeo}(\Sigma)$ *such that*

(i) $R = f(S) \subset \Sigma_0$,

(ii) $S \subset f(\Sigma_0)$, *and*

(iii) *the end set of* $f(\Sigma_0) \cap \Sigma_0$ *contains a homeomorphic copy of* E.

Proof Since E' is homeomorphic to E we can write E' as the disjoint union of sets E_i', with $i = 0, 1, \ldots, n$, and P', where $E_i' \cong E_i$ and $P' \cong P$. (Of course, by $\cong$ we mean homeomorphic via a homeomorphism which respects E^G.) We can further write $E_0' = E'' \sqcup D'$, where $E'' \cong E$ and $D' \cong D$.

Consider a subsurface R disjoint from S with puncture set P' and $n+1$ complementary regions, one with end space E_i' for each $i = 1, 2, \ldots, n$, and the final one containing the remaining ends, namely $D' \sqcup E'' \sqcup D \sqcup \left(\bigsqcup_{1 \le i \le n} E_i\right)$. Now we have

$$E'' \sqcup D \sqcup \left(\bigsqcup_{1 \le i \le n} E_i \right) \cong E,$$

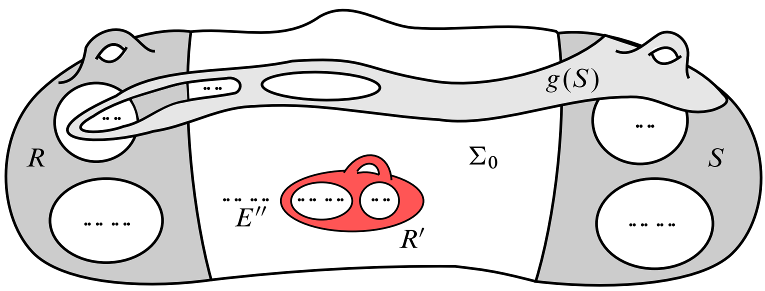

Figure 3: The surface $g(S)$ may intersect R and S in a complicated way, but R' lies in the "big" complementary region of at least one of them (in this illustration, it is in their intersection Z).

therefore,

$$D' \sqcup E'' \sqcup D \sqcup \Bigl(\bigsqcup_{1 \le i \le n} E_i \Bigr) \cong D' \sqcup E \cong E_0.$$

Thus, we may apply Richards' classification of surfaces and conclude that there is a homeomorphism f of Σ such that $f(S) = R$ and for $i \ge 1$ we have $f(E_i) = E_i'$, and $f(E_0) = D' \sqcup E'' \sqcup D \sqcup \bigl(\bigsqcup_{1 \le i \le n} E_i\bigr)$. □

Now fix R and f as in Lemma 3.2 and let $\mathcal{F} = \{f, f^{-1}\}$. We will show

$$\operatorname{Map}(\Sigma) = (\mathcal{F}\mathcal{V}_S)^5.$$

Let $g \in \operatorname{Map}(\Sigma)$. Let E' be a homeomorphic copy of E in the end space of $\Sigma_0 \cap f(\Sigma_0)$, and consider the set $g(E')$.

Since the clopen sets $Z := (f(E_0) \cap E_0)$, $(E_0 - Z)$ and $(E - E_0)$ partition E, their intersections with $g(E')$ partition $g(E')$. Since $g(E') \cong E$ is a self-similar set, one of these three sets in the partition contains a homeomorphic copy of E; call this E''. Thus, E'' lies either in $g(E') \cap f(E_0)$ or in $g(E') \cap E_0$ (or both). If the first case occurs, then we have $f^{-1}g(E') \cap f^{-1}(f(E_0))$. This means that, at the cost of replacing g by $f^{-1}g$, and therefore using one more letter from $\mathcal{F}$, we can assume that we are in the second case, ie where $E'' \subset g(E') \cap E_0$. So it suffices to show that in this case, we have $g \in (\mathcal{F}\mathcal{V}_S)^4$. This situation is illustrated in Figure 3. (For simplicity, we did not draw infinite genus on this image.)

Assuming that $E'' \subset g(E') \cap E_0$, the next step is to find another copy of S in a small neighborhood of E'', and hence in $g(\Sigma_0) \cap \Sigma_0$. In detail, just as in Lemma 3.2, but using E'' instead of E', and working with the subsurface R of the surface Σ_0 instead of the subsurface S of Σ, we may find a surface $R' \subset \Sigma_0 \cap g(\Sigma_0)$ homeomorphic to R, and a homeomorphism v of Σ_0 mapping R to R' that satisfies $R \subset v(f(\Sigma_0) \cap \Sigma_0)$. Extend v to a homeomorphism of Σ by declaring it to be the identity on $\Sigma - \Sigma_0$; abusing notation slightly, denote this homeomorphism also by v, and so we have $v \in \mathcal{V}_S$. Then R, S and $g(S)$ are all contained in $v(f(\Sigma_0))$. See Figure 3 for a schematic.

The same argument as that in Lemma 3.2 using the classification of surfaces now shows that we may find u restricting to the identity on R', with $ug(S) = S$ and ug equal to identity on S. (The details are a straightforward exercise.) Since u is the identity on R', it follows that $(vf)^{-1}u(vf)$ is the identity on $(vf)^{-1}(R') = S$, which implies that $u \in (\mathcal{F}\mathcal{V}_S)^3$, hence $g \in (\mathcal{F}\mathcal{V}_S)^4$. This concludes the proof of Proposition 3.1. □

3.2 Telescoping end spaces

Motivation Recall from Example 2.5 that, if Σ is a surface such that there exists a finite, $\mathrm{Map}(\Sigma)$–invariant set $F \subset E$ of cardinality at least three, then $\mathrm{Map}(\Sigma)$ is not CB: any finite-type subsurface S such that the elements of F each lie in different complementary regions of S is easily seen to be nondisplaceable. The definition of *telescoping* below was motivated by the question: *Under what conditions is a two-element* $\mathrm{Map}(\Sigma)$*–invariant subset of E compatible with global coarse boundedness*? As will follow from our work in Section 7, this never happens if E is countable: every surface with countable end space and coarsely bounded mapping class group is self-similar. However, in the uncountable case, surfaces with telescoping end spaces provide additional examples (and are the only additional examples among tame surfaces). Informally speaking, telescoping spaces of ends have two "special" points with the property that neighborhoods of each point can be expanded an arbitrary amount, and can also be expanded a fixed amount relative to a neighborhood of the other point.

Convention In the following definition, and for the remainder of this work, we wish to work only with specific neighborhoods of ends in Σ, not every open subset of the surface containing this end. Thus, going forward, a *neighborhood of an end x in Σ* means a connected subsurface with a single boundary component that has x as an end.

Definition 3.3 A surface Σ is *telescoping* if there are ends $x_1, x_2 \in E$ and disjoint clopen neighborhoods V_i of x_i in Σ such that for all clopen neighborhoods $W_i \subset V_i$ of x_i, there exist homeomorphisms f_i and h_i of Σ, both pointwise fixing $\{x_1, x_2\}$, with

$$f_i(W_i) \supset (\Sigma - V_{3-i}), \quad h_i(W_i) = V_i, \quad h_i(V_{3-i}) = V_{3-i}.$$

When we wish to make the points x_1, x_2 explicit, we say also *telescoping with respect to* $\{x_1, x_2\}$. We may equivalently require h_i to restrict to the identity on V_{3-i}.

Note that this definition implies that Σ has infinite or zero genus, as does $\Sigma - (V_1 \cup V_2)$.

While the complement of a Cantor set in S^2 is both self-similar and telescoping with respect to any pair of points, there are many examples of telescoping sets that are *not* self-similar, for instance:

- E^G a Cantor set, and E the union of E^G and another Cantor set which intersects E^G at exactly two points.
- E the union of two copies of the Cantor set, C_1 and C_2, which intersect at exactly two points, and a countable set Q such that the accumulation points of Q are exactly C_1. E^G could be empty, equal to the closure of Q, or equal to E.

Note that, in Definition 3.3, f_i and h_i are required to be homeomorphisms of the surface, not merely the end space.

Remark 3.4 An equivalent definition of telescoping may be given by replacing "there exist disjoint neighborhoods V_i of x_i" with "for all sufficiently small neighborhoods V_i of x_i". The proof is an immediate consequence of the definition.

The telescoping condition also implies that all neighborhoods of x_i in $\Sigma - \{x_{3-i}\}$ are homeomorphic. With this fact, one can use a standard back-and-forth argument to show that there is a homeomorphism of Σ taking x_i to x_{3-i}. We omit the proof as it is not needed for what follows.

Proposition 3.5 (telescoping implies CB) *Let Σ be a surface that is telescoping with respect to $\{x_1, x_2\}$. Then the pointwise stabilizer of $\{x_1, x_2\}$ in* $\mathrm{Map}(\Sigma)$ *is CB.*

In particular, if $\{x_1, x_2\}$ is a $\mathrm{Map}(\Sigma)$–invariant set, then $\mathrm{Map}(\Sigma)$ is itself CB.

Remark 3.6 In fact, it will follow easily from the tools developed in the next section (see Proposition 4.8) that if $\{x_1, x_2\}$ is not invariant, then the end space of Σ is self-similar and so $\mathrm{Map}(\Sigma)$ is CB in this case as well.

Proof of Proposition 3.5 Suppose that Σ is telescoping and let x_i and V_i be as in the definition. To simplify notation, let G denote the pointwise stabilizer of $\{x_1, x_2\}$ in $\mathrm{Map}(\Sigma)$. Fix a neighborhood of the identity in $\mathrm{Map}(\Sigma)$; shrinking this if needed we may take it to be the set $\mathcal{V}_S$ of mapping classes that restrict to the identity on some finite-type subsurfaces S. By Remark 3.4, we may assume that $S \subset \Sigma - (V_1 \cup V_2)$. Let

$\mathcal{V} \subset \mathcal{V}_S$ be the set of mapping classes that restrict to the identity on $\Sigma' := \Sigma - (V_1 \cup V_2)$. We will exhibit a finite set $\mathcal{F}$ such that $G \subset (\mathcal{F}\mathcal{V})^{10} \subset (\mathcal{F}\mathcal{V}_S)^{10}$. This is sufficient to show that G is CB, by Theorem 2.1.

Fix neighborhoods $W_i \subset V_i$ of x_i in Σ and homeomorphisms f_i with $f_i(W_i) \supset (\Sigma - V_{3-i})$, as given by the definition of telescoping. Let $\mathcal{F} = \{f_1^{\pm 1}, f_2^{\pm 1}\}$. This is our finite set. Note that any homeomorphism which restricts to the identity on V_i lies in $f_{3-i}\mathcal{V} f_{3-i}^{-1}$.

Given $g \in G$, let W_i' be a neighborhood of x_i small enough that $W_i' \subset g^{-1}(V_i) \cap g(V_i)$. By definition of telescoping, there exist homeomorphisms h_1 and j_1, both restricting to the identity on V_2, with $h_1(g(W_1')) = V_1$ and $j_1(W_1) = V_1$. Then $g_1 := j_1^{-1} h_1$ is the identity on V_2, hence lies in $f_1 \mathcal{V} f_1^{-1}$, and satisfies $g_1 g(W_1') = W_1$.

Similarly, we can find $g_2 \in f_2 \mathcal{V} f_2^{-1}$ restricting to the identity on V_1, and satisfying $g_2 g(W_2') = W_2$. Thus,

$$g_2 g_1 g(W_i') = W_i \quad \text{for } i = 1, 2.$$

It follows that $g_2 g_1 g(\Sigma') \subset (\Sigma - W_1 \cup W_2)$, so $f_1 g_2 g_1 g(\Sigma') \subset V_2$ and

$$f_2^{-1} f_1 g_2 g_1 g(\Sigma') \subset W_2.$$

For notational convenience, let $\phi = f_2^{-1} f_1 g_2 g_1 g$. Since $\phi(\Sigma')$ and $f_2^{-1} f_1 \Sigma'$ both lie in W_2, as a consequence of the definition of telescoping there exists a homeomorphism ψ restricting to the identity on V_1, with $\psi\phi(\Sigma') = f_2^{-1} f_1(\Sigma')$. Precomposing ψ with a homeomorphism that is also the identity on V_1, we can also ensure that $(f_2^{-1} f_1)^{-1} \psi\phi$ restricts to the identity on Σ'. Thus, we have shown that $(f_2^{-1} f_1)^{-1} \psi\phi = (f_2^{-1} f_1)^{-1} \psi f_2^{-1} f_1 g_2 g_1 g \in \mathcal{V}$. Since $\psi^{-1} \in (\mathcal{F}\mathcal{V})^2$, and $g_i^{-1} \in (\mathcal{F}\mathcal{V})^2$, we conclude that $g \in (\mathcal{F}\mathcal{V})^{10}$. Since $\mathcal{F}$ and the exponent are independent of g, we have proved the desired result. □

We conclude this section with a result whose proof serves as a good warm-up for the technical work to come.

Proposition 3.7 *No telescoping surface has countable end space.*

Proof Suppose that Σ has countable end space E. Recall in this case $E \cong \omega^\alpha n + 1$ by [13], and $E^G \subset E$ is some closed subset. Assume for contradiction that E is telescoping with respect to some pair of ends x_1, x_2. For each point $x \in E$, there exists

$\beta = \beta(x) \leq \alpha$ such that every sufficiently small neighborhood of x is homeomorphic to $\omega^\beta + 1$ (this ordinal $\beta(x)$ is simply the Cantor–Bendixon rank of x). It follows from the definition of telescoping that every clopen neighborhood U of x_i disjoint from x_{3-i} is homeomorphic to every other such neighborhood. In particular, necessarily $n = 2$ and x_1 and x_2 are points of equal and maximal rank α. Suppose as a first case that α is a successor ordinal and let η denote its immediate predecessor. Then the set of points of rank η accumulates only at x_1 and x_2. If V_i is any neighborhood of x_i, then $\Sigma - (V_1 \cup V_2)$ contains finitely many points of rank η. Thus, if $W_1 \subset V_1$ satisfies that $V_1 - W_1$ contains exactly one point of rank η, then no homeomorphism fixing V_2 can send W_1 to V_1, and the definition of telescoping fails.

The case where α has limit type is similar. Given neighborhoods V_i of x_i, let $\eta < \alpha$ be the supremum of the ranks of points in $E - (V_1 \cup V_2)$. Let $W_1 \subset V_1$ be a set such that $V_1 - W_1$ contains a point of rank α where $\eta < \alpha$. Then no homeomorphism fixing V_2 can send W_1 to V_1, and the definition of telescoping fails. □

As we will see in the next sections, this limit type phenomenon is closely related to the failure of the mapping class group to be generated by a CB set. However, to treat this in the case where E is uncountable, we will need to develop a more refined ordering on the space of ends.

4 A partial order on the space of ends

Let Σ be an infinite-type surface with set of ends (E, E^G). As in the previous section, we drop the notation E^G and, by convention, all homeomorphisms of an end space E of a surface Σ are required to preserve E^G, so to say that $A \subset E$ is homeomorphic to $B \subset E$ means that there is a homeomorphism from $(A, A \cap E^G)$ to $(B, B \cap E^G)$. It follows from Richards' classification of surfaces in [16, Theorem 1] that each homeomorphism of (E, E^G) is induced by a homeomorphism of Σ.[4] Thus, we will pass freely between speaking of homeomorphisms of the end space and the underlying surface.

Observe also that, if U and V are two disjoint, clopen subsets of E, then any homeomorphism f from U onto V can be extended to a globally defined homeomorphism $\bar{f}$ of E by declaring $\bar{f}$ to agree with f^{-1} on V and to pointwise fix the complement

[4] While this is not in the statement of [16, Theorem 1], the proof gives such a construction. This was originally explained to the authors by J Lanier following work of S Afton.

of $U \cup V$. Thus, to say points x and y are locally homeomorphic is equivalent to the condition that there exists $\bar{f} \in \mathrm{Map}(\Sigma)$ with $\bar{f}(x) = y$. We will use this fact frequently. In particular, we have the following equivalent rephrasing of Definition 1.3:

Definition 4.1 Let $\preccurlyeq$ be the binary relation on E given by $y \preccurlyeq x$ if, for every neighborhood U of x, there exists a neighborhood V of y and $f \in \mathrm{Map}(\Sigma)$ such that $f(V) \subset U$.

Note that this relation is transitive.

Notation 4.2 For $x, y \in E$ we say that $x \sim y$ or "x and y are of the same type" if $x \preccurlyeq y$ and $y \preccurlyeq x$, and write $E(x)$ for the set $\{y \mid y \sim x\}$ of "ends of type x".

It is easily verified that $\sim$ defines an equivalence relation: symmetry and reflexivity are immediate from the definition, while transitivity follows from the transitivity of $\preccurlyeq$. From this it follows that the relation $\prec$, defined by $x \prec y$ if $x \preccurlyeq y$ and $x \nsim y$, gives a partial order on the set of equivalence classes under $\sim$. For any homeomorphism f of Σ, we have $x \succ y$ (resp. $x \succcurlyeq y$) if and only if $f(x) \succ f(y)$ (resp. $f(x) \succcurlyeq f(y)$).

Proposition 4.3 *If E is countable, then $x \sim y$ if and only if x and y are locally homeomorphic. If additionally $E^G = \varnothing$, then the Cantor–Bendixon rank gives an order isomorphism between equivalence classes of points and countable ordinals.*

Proof Suppose that E is countable. Consider first the case where $E^G = \varnothing$. Then every point $x \in E$ has a neighborhood U_x homeomorphic to the set $\omega^{\alpha(x)} + 1$, where $\alpha(x)$ is the Cantor–Bendixon rank of x. If $x \preccurlyeq y$ and $y \preccurlyeq x$ both hold, it follows that $\alpha(x) = \alpha(y)$, and so any homeomorphism from a neighborhood of x into a neighborhood of y necessarily takes x to y. Thus, x and y are locally homeomorphic. In particular, these points also have the same rank.

In the general case where $E^G \neq \varnothing$, let $\bar{E}$ denote the topological space of ends (with no distinction between those accumulated by genus or not). Any homeomorphism of E induces one of $\bar{E}$ by simply forgetting that E^G is preserved. Thus, the argument above shows that if $x \sim y$ in E, then they admit neighborhoods U_x and U_y in $\bar{E}$ which are homeomorphic. Moreover, such a homeomorphism necessarily takes x to y, and in fact no homeomorphism of E can take x to another point of U_y. Thus, $x \preccurlyeq y$ implies that there is a homeomorphism of E taking a neighborhood of x in E to one of y. The converse statement is immediate. □

Remark 4.4 We do not know if Proposition 4.3 holds in the uncountable case. This appears to be an interesting question. However, it is quite easy to construct large families of examples for which it does hold.

Remark 4.5 Despite the above remark, there are indeed some marked differences between the behavior of $\prec$ when E is countable and when E is uncountable. In the countable case, it follows from Proposition 4.3 that $x \prec y$ if and only if y is an accumulation point of $E(x)$, giving a convenient alternative description of $\prec$. In general, a weaker statement holds: we show below that if y is an accumulation point of $E(x)$, then $x \preccurlyeq y$. However, if E is a Cantor set and $E^G = \varnothing$, for example, then all points are equivalent and all are accumulation points of their equivalence class.

We now prove some general results on the structure of $\preccurlyeq$.

Lemma 4.6 *For every $y \in E$, the set $\{x \mid x \succcurlyeq y\}$ is closed.*

Proof Consider a sequence $x_n \to x$ where $x_n \succcurlyeq y$ holds for all n. Let U be a neighborhood of x. Then, for large n, U is also a neighborhood of x_n and hence contains homeomorphic copies of some neighborhood of y. □

Proposition 4.7 *The partial order $\succ$ has maximal elements. Furthermore, for every maximal element x, the equivalence class $E(x)$ is either finite or a Cantor set.*

Proof To show that E has maximal elements, by Zorn's lemma it suffices to show that every chain has an upper bound. Suppose that $\mathcal{C}$ is a totally ordered chain. Consider the family of sets $\{x \mid x \succcurlyeq y\}$, for $y \in \mathcal{C}$. Then, by Lemma 4.6, this is a family of nested, closed, nonempty sets and hence

$$\mathcal{C}_M = \bigcap_{y \in \mathcal{C}} \{x \mid x \succcurlyeq y\}$$

is nonempty. By definition, any point of this intersection is an upper bound for $\mathcal{C}$.

To see the second assertion, consider a maximal element x. If $E(x)$ is an infinite set, then it has an accumulation point, say z. Then $z \succcurlyeq x$, but since x is maximal, we have $z \sim x$. Since any neighborhood of any other point in $E(x)$ contains a homeomorphic copy of a neighborhood of z, it follows that all points of $E(x)$ are accumulation points and hence $E(x)$ is a Cantor set. □

Going forward, we let $\mathcal{M} = \mathcal{M}(E)$ denote the set of maximal elements for $\succ$.

4.1 Characterizing self-similar end sets

The remainder of this section consists of a detailed study of the behavior of end sets using the partial order. We will develop a number of tools for the classification of locally CB and CB generated mapping class groups that will be carried out in the next sections.

Proposition 4.8 *Let Σ be a surface with end space E and no nondisplaceable subsurfaces. Then E is self-similar if and only if $\mathcal{M}$ is either a singleton or a Cantor set of points of the same type.*

One direction is easy and does not require the assumption that Σ has no nondisplaceable subsurfaces: if $\mathcal{M}$ contains two distinct maximal types x_1 and x_2, then a partition $E = E_1 \sqcup E_2$, where $E(x_i) \subset E_i$, fails the condition of self-similarity. Similarly, if $\mathcal{M}$ is a finite set of cardinality at least two, then any partition separating points of $\mathcal{M}$ similarly fails the condition. By Proposition 4.7, the only remaining possibility is that $\mathcal{M}$ is a Cantor set of points of the same type. This proves the first direction. The converse is more involved, so we treat the singleton and Cantor set case separately. We will need the following easy observation.

Observation 4.9 ("shift maps") *Suppose $U_1, U_2, \dots$ are disjoint, pairwise homeomorphic clopen sets which Hausdorff converge to a point x. Then $\bigcup_{i=1}^{\infty} U_i \cup \{x\}$ is homeomorphic to $\bigcup_{i=2}^{\infty} U_i \cup \{x\}$.*

Proof For each i, fix a homeomorphism $f_i \colon U_i \to U_{i+1}$. Since the U_i Hausdorff converge to a point, the union of these defines a global homeomorphism $\bigcup_{i=1}^{\infty} U_i \to \bigcup_{i=2}^{\infty} U_i$ that extends continuously to x. □

We will also use the following alternative characterization of self-similarity:

Lemma 4.10 *Self-similarity is equivalent to the following condition: if $E = E_1 \sqcup E_2$ is a decomposition into clopen sets, then some E_i contains a clopen set homeomorphic to E.*

Proof Self-similarity implies the condition by taking $n = 2$. For the converse, suppose the condition holds and let $E = E_1 \sqcup E_2 \sqcup \dots \sqcup E_n$ be a decomposition into clopen sets. Grouping these as $E_1 \sqcup (E_2 \sqcup \dots \sqcup E_n)$, by assumption one of these subsets contains a clopen set E' homeomorphic to E. If it is E_1, we are done. Else, the sets $E' \cap E_i$ with $i = 2, 3, \dots, n$ form a decomposition of $E' \cong E$ into clopen sets; so by the same reasoning either $E_2 \cap E'$ contains a clopen set homeomorphic to E, or the union of the sets $E' \cap E_i$, for $i \geq 3$, does. Iterating this argument eventually produces a set homeomorphic to E in one of the E_i. $\square$

The next three lemmas give the proof of Proposition 4.8.

Lemma 4.11 *Suppose Σ has no nondisplaceable subsurfaces and $\mathcal{M}$ is a singleton. Let $E = A \sqcup B$ be a decomposition into clopen sets. If $\mathcal{M} \subset A$, then A contains a homeomorphic copy of B.*

Proof Let $E = A \sqcup B$ be a decomposition of E into clopen sets with $\mathcal{M} = \{x\} \subset A$. Since A is a neighborhood of x, every point $y \in B$ has a neighborhood homeomorphic to a subset of A. Since B is compact, finitely many of these cover B, say $U_1, U_2, \dots, U_k$. Without loss of generality, we may assume all the U_i are disjoint and their union is equal to B. For each i, let V_i be a homeomorphic copy of U_i in A; note that $x \notin \bigcup_i V_i$. Let S be a three-holed sphere subsurface such that the disjoint sets $\{x\}$, $\bigcup_i V_i$ and B all lie in different connected components of the complement of S. Let f be a homeomorphism displacing S. Since $f(x) = x$, up to replacing f with its inverse, we have either $f(B) \subset A$, in which case we are done, or A contains a homeomorphic copy of $A \sqcup \left(\bigcup_i V_i\right)$. In this latter case, by iterating f we can find k disjoint copies of $\bigcup_i V_i$ inside A. Since each contains a copy of U_i, this gives a subset of A homeomorphic to $\bigsqcup U_i = B$. $\square$

As a consequence, we can prove the first case of Proposition 4.8.

Lemma 4.12 *Suppose Σ has no nondisplaceable subsurfaces and $\mathcal{M}$ is a singleton. Then E is self-similar.*

Proof Let $E = E_1 \sqcup E_2$ be a decomposition of E into clopen sets. Without loss of generality, suppose $\mathcal{M} = \{x\} \subset E_1$. Lemma 4.11 says that there is a homeomorphic copy U_2 of E_2 inside E_1, necessarily this is disjoint from $\{x\}$. Let A be a small neighborhood of x, disjoint from U_2. Lemma 4.11 again gives a homeomorphic copy U_3 of E_2 inside A. Proceeding in this way, we may find $E_2 = U_1, U_2, U_3, \dots$, each homeomorphic to E_2 and Hausdorff converging to x. Define $f \colon E_1 \sqcup E_2 \to E_1$ to be the homeomorphism where the restriction of f to $\bigcup_{i=1}^{\infty} U_i \cup \{x\}$ is constructed as in Observation 4.9, and the restriction of f to the rest of E is the identity. $\square$

The second case is covered by the following:

Lemma 4.13 *Suppose Σ has no nondisplaceable subsurfaces and $\mathcal{M}$ is a Cantor set of points all of the same type. Then E is self-similar.*

Proof Let $E = E_1 \sqcup E_2$ be a decomposition of E into clopen sets. If $\mathcal{M}$ is contained in only one of the E_i, then one may apply the argument from Lemma 4.12, by letting x be any point of $\mathcal{M}$. Thus, we assume that both E_1 and E_2 contain points of $\mathcal{M}$.

For concreteness, fix a metric on E. For each $n \in N$, fix a decomposition $A_1^{(n)}, \dots, A_{j_n}^{(n)}$ of E into clopen sets of diameter at most 2^{-n}, such that E_1 and E_2 are each the union of some number of these sets. Let S_n be a subsurface homeomorphic to a j_n–holed sphere, with complementary regions containing the sets $A_k^{(n)}$. Since S_n is displaceable, there exists some k such that $A_k^{(n)}$ contains a copy of all but one of the sets $A_j^{(n)}$; in particular, it contains a copy of either E_1 or E_2. Passing to a subsequence, we conclude that for either $i = 1$ or $i = 2$ there exist homeomorphic copies of E_i of diameter less than 2^{-n}, for each n. Without loss of generality, say that this holds for E_1. Passing to a further subsequence, we can assume these copies of E_1 Hausdorff converge to a point x, so in particular every neighborhood of x contains a copy of E_1.

It follows from the definition of $\preccurlyeq$ that each $y \in \mathcal{M}$ therefore also has this property: every neighborhood of y contains a homeomorphic copy of E_1. Let $y_2, y_3, y_4, \dots$ be a sequence of points in E_2 converging to $y \in E_2$, and let $U_1 = E_1$. Fix disjoint neighborhoods N_i of y_i converging to y, and let U_i be a homeomorphic copy of E_1 in N_i. Now apply Observation 4.9. □

This completes the proof of Proposition 4.8.

4.2 Stable neighborhoods

Motivated by the behavior of maximal points in the proposition above, we make the following definition:

Definition 4.14 For $x \in E$, call a neighborhood U of x *stable* if for any smaller neighborhood $U' \subset U$ of x, there is a homeomorphic copy of U contained in U'.

Our use of the terminology "stable" is justified by Lemma 4.17 below, which says that all such neighborhoods of a point are homeomorphic. (Recall that, by convention, neighborhood always means clopen neighborhood.)

Remark 4.15 Stable neighborhoods are automatically self-similar sets, and if U is a stable neighborhood of x, then $x \in \mathcal{M}(U)$. Our work in the previous section shows that when $\prec$ has a unique maximal type and all subsurfaces are displaceable, each maximal point has a stable neighborhood.

It follows immediately from the definition that if x has one stable neighborhood, then *every* sufficiently small neighborhood of x is also stable. More generally, we have the following.

Lemma 4.16 *If x has a stable neighborhood, and $y \sim x$, then y has a stable neighborhood.*

Proof Let U be a stable neighborhood of x. Since $y \prec x$, there is a neighborhood V of y such that U contains a homeomorphic copy of V. Suppose $V' \subset V$ is a smaller neighborhood of y. Since $x \prec y$, there is some neighborhood U' of x (without loss of generality, we may assume that $U' \subset U$) such that V' contains a homeomorphic copy of U'. By definition of stable neighborhoods, U' contains a homeomorphic copy of U, thus V' contains a homeomorphic copy of U and hence a homeomorphic copy of V. □

Lemma 4.17 *If x has a stable neighborhood U, then for any $y \sim x$, all sufficiently small neighborhoods of y are homeomorphic to U via a homeomorphism taking x to y.*

Proof The proof is a standard back-and-forth argument. Suppose $x \prec y$ and $y \prec x$. Let V_x be a stable neighborhood of x and V_y a stable neighborhood of y. Take a neighborhood basis $V_x = V_0 \supset V_1 \supset V_2 \supset \cdots$ of x consisting of nested neighborhoods, and take a neighborhood basis $V_y = V_0' \supset V_1' \supset V_2' \supset \cdots$ of y. Since $y \prec x$ and $x \prec y$, each V_i contains a homeomorphic copy of V_0' and each V_i' a copy of V_0.

Let f_1 be a homeomorphism from $V_0 - V_1$ into V_0'. Note that we may assume the image of f_1 avoids y: if y is the unique maximal point of V_0', then this is automatic, otherwise, $E(y)$ is a Cantor set of points, each of which contains copies of V_0 in every small neighborhood. Let g_1 be a homeomorphism from the complement of the image of f_1 in $V_0' - V_1'$ onto a subset of $V_1 - \{x\}$. Iteratively, define f_i to be a homeomorphism from the complement of the image of g_{i-1} in $V_{i-1} - V_i$ onto a subset of $V_{i-1}' - \{y\}$, and g_i a homeomorphism from the complement of the image of f_i in $V_{i-1}' - V_i'$ onto a subset of $V_i - \{x\}$. Then the union of all f_i and g_i^{-1} is a homeomorphism from $V_0 - \{x\}$ to $V_0' - \{y\}$ that extends to a homeomorphism from V_0 to V_0' taking x to y. □

The following variation on Lemma 4.11 uses stable neighborhoods as a replacement for displaceable subsurfaces.

Lemma 4.18 *Let $x, y \in E$, and assume x has a stable neighborhood V_x and that x is an accumulation point of $E(y)$. Then for any sufficiently small neighborhood U of y, $U \cup V_x$ is homeomorphic to V_x.*

Proof If $x \sim y$, then let U be a stable neighborhood of y disjoint from V_x. Let $V_1 \supset V_2 \supset V_3 \supset \cdots$ be a neighborhood basis for x consisting of stable neighborhoods. Since x is an accumulation point of $E(y)$, for any sufficiently small neighborhood U_0 of y (and hence for any stable neighborhood U), there is a homeomorphic copy U_1 of U_0 in $V_1 - \{x\}$. Shrinking neighborhoods if needed, we may take U_1 to be disjoint from V_{i_1} for some $i_1 \in \mathbb{N}$. Since V_{i_1} is homeomorphic to V_1, there is also a homeomorphic copy of U_2 of U_0 in V_{i_1}, disjoint from some V_{i_2}. Iterating this process we can find disjoint sets $U_n \subset V_1$, each homeomorphic to U, and Hausdorff converging to x. Define $f : V_1 \cup U_0 \to V_1$ to be the identity on the complement of $\bigcup_n U_n$ and send U_i to U_{i+1} by a homeomorphism as in Observation 4.9.

If instead $y \prec x$, then take any neighborhood U of y disjoint from V_x and small enough that V_x contains a homeomorphic copy of U. Since $y \prec x$, this copy lies in $V_x - \{x\}$, and we may repeat the same line of argument above. □

5 Classification of locally CB mapping class groups

We now prove properties of locally CB mapping class groups, building towards our general classification theorem. Recall that we have the following notational convention.

Notation 5.1 If $K \subset \Sigma$ is a finite-type subsurface, we denote by $\mathcal{V}_K$ the identity neighborhood consisting of mapping classes of homeomorphisms that restrict to the identity on K.

Lemma 5.2 *Let $K \subset \Sigma$ be a finite-type subsurface such that each component of $\Sigma - K$ has infinite type. If there exists a finite-type, nondisplaceable (possibly disconnected) subsurface S in $\Sigma - K$, then $\mathcal{V}_K$ is not CB. If this holds for every such finite-type $K \subset \Sigma$, then* $\mathrm{Map}(\Sigma)$ *is not locally CB.*

Proof Let K be a surface as in the statement of the proposition, with a nondisplaceable subsurface $S \subset \Sigma - K$. Since each complementary region to K was assumed to have infinite type, by enlarging S if needed we may assume that S still remains in the complement of K, but is such that each component of S has high enough complexity that the length function ℓ_S defined in Section 2 will be unbounded. As in Proposition 2.8, this gives a length function which is unbounded on $\mathcal{V}_K$, hence on $\mathcal{V}$, so $\mathrm{Map}(\Sigma)$ is not locally CB.

As remarked above, the sets $\mathcal{V}_L$, where L ranges over finite-type subsurfaces, form a neighborhood basis of the identity in $\mathrm{Map}(\Sigma)$. But one may in fact restrict this to range over finite-type surfaces whose complementary regions are all of infinite type, since if L is finite type, then the union of L and its finite-type complementary regions is again a compact surface, say K, and $\mathcal{V}_K \subset \mathcal{V}_L$. Thus, $\mathrm{Map}(\Sigma)$ is locally CB if and only if some such set $\mathcal{V}_K$ is CB. □

Going forward, we reference the partial order $\prec$ defined in Section 4.

Lemma 5.3 *If* $\mathrm{Map}(\Sigma)$ *is locally CB, then the number of distinct maximal types under* $\prec$ *is finite.*

Proof We prove the contrapositive. Suppose that there are infinitely many distinct maximal types. Let K be any subsurface of finite type. By Lemma 5.2, it suffices to find a nondisplaceable subsurface contained in $\Sigma - K$, which we do now.

To every end $x \in E$ of maximal type, let $\sigma(x)$ denote the set of connected components of $\Sigma - K$ which contain ends from $E(x)$. Since $\Sigma - K$ has finitely many connected components, by the pigeonhole principle there are two ends x and y with $x \not\sim y$ but $\sigma(x) = \sigma(y)$. That is, each complementary region of $\Sigma - K$ that has an end from $E(x)$ also contains ends from $E(y)$, and vice versa. Fix any $z \in E$ with $z \not\sim x$ and $z \not\sim y$.

Construct a surface S as follows. For each component τ of $\sigma(x)$, take a three-holed sphere subsurface contained in τ so that the complementary regions of the three-holed sphere separate $E(x)$ from $E(y)$ and $E(z)$ in τ. That is to say, one complementary region contains only ends from $E(x)$ and none from $E(y)$ or $E(z)$, while another contains only ends from $E(y)$ and none from $E(x)$ or $E(z)$, and the third contains at least some points of $E(z)$ (possibly those from another complementary region of K). Let S be the union of these three holed spheres, one in each component of $\sigma(x)$. Thus, each end from $E(x)$ is the end of some complementary region of S which has no ends of type y, and vice versa.

We claim that S is nondisplaceable. For if S_i is a connected component of S, then one complementary region of S_i contains ends from $E(x)$, but none from $E(y)$. By invariance of $E(x)$ and $E(y)$, if some homeomorphic image $f(S_i)$ were disjoint from S, then we would have to have $f(S_i)$ contained in one of the complementary regions of S containing points of $E(x)$. However, this region contains no points of $E(y)$ or $E(z)$, contradicting our construction of S_i. Hence, S is nondisplaceable and, by Lemma 5.2, $\mathrm{Map}(\Sigma)$ is not locally CB. □

We now state the first structure theorem for end spaces of surfaces with locally CB mapping class groups.

Proposition 5.4 *If* $\mathrm{Map}(\Sigma)$ *is locally CB, then there is a partition*

$$E = \bigsqcup_{A \in \mathcal{A}} A,$$

where $\mathcal{A}$ is finite, each $A \in \mathcal{A}$ is clopen and self-similar, and $\mathcal{M}(A) \subset \mathcal{M}(E)$. Moreover, this decomposition can be realized by the complementary regions to a finite-type surface $L \subset \Sigma$ with $|\mathcal{A}|$ boundary components, either of zero genus or of finite genus equal to the genus of Σ.

This will be a quick consequence of the following stronger result:

Proposition 5.5 *Suppose that* $\mathrm{Map}(\Sigma)$ *is locally CB. Then there exists a CB neighborhood $\mathcal{V}_K$ of the identity, where K is a finite-type surface with the following properties*:

(i) *Each connected component of $\Sigma - K$ has one or infinitely many ends and zero or infinite genus.*

(ii) *The connected components of $\Sigma - K$ partition E as*

$$E = \bigsqcup_{\hat{A} \in \mathcal{A}} \hat{A} \sqcup \bigsqcup_{P \in \mathcal{P}} P,$$

where each $\hat{A} \in \mathcal{A}$ is self-similar, and for each $P \in \mathcal{P}$, there exists some $\hat{A} \in \mathcal{A}$ such that P is homeomorphic to a clopen subset of $\hat{A}$.

(iii) *For all $\hat{A} \in \mathcal{A}$, the maximal points $\mathcal{M}(\hat{A})$ are maximal in E, and $\mathcal{M}(E) = \bigsqcup_{\hat{A} \in \mathcal{A}} \mathcal{M}(\hat{A})$.*

Our choice of $\mathcal{A}$ as the notation for the index set in both propositions is because they may be canonically identified. In fact, the proof of Proposition 5.4 consists of showing that each of the sets A is a union of one set $\hat{A}$ from Proposition 5.5 and some number of the sets in $\mathcal{P}$, and that A is homeomorphic to $\hat{A}$.

Proof of Proposition 5.5 Suppose that $\mathcal{V}$ is a CB neighborhood of the identity in $\mathrm{Map}(\Sigma)$. Let K be a finite-type surface such that $\mathcal{V}_K \subset \mathcal{V}$, so $\mathcal{V}_K$ is also CB. Enlarging K if needed (and hence shrinking $\mathcal{V}_K$), we may assume that each complementary region to K has either zero or infinite genus. Since $\prec$ has only finitely many maximal types, enlarging K further, we may assume that its complementary regions separate the different maximal types, and moreover, if for some maximal x the set $E(x)$ is finite, then all the ends from $E(x)$ are separated by K. Thus, complementary regions to K have either no end from $\mathcal{M}(E)$, a single end from $\mathcal{M}(E)$ or a Cantor set of ends of a single type from $\mathcal{M}(E)$.

Our goal is to show that the complementary regions containing ends from $\mathcal{M}(E)$ are all self-similar sets, and the end sets of the remaining regions have the property desired of the sets $P \in \mathcal{P}$ described above. It will be convenient to introduce some terminology for the set of ends of a complementary region to K, so call such a subset of E a *complementary end set*.

For simplicity, assume as a first case that for each maximal type x, the set $E(x)$ is *finite*. Fix a maximal type point $x \in E$, and let $B_1, B_2, \dots, B_k \subset E$ be the complementary end sets whose maximal points lie in $E(x)$. We start by showing that at least one of the sets B_i is self-similar. Let x_i denote the maximal point in B_i. Let U_i be any clopen neighborhood of x_i in B_i. Since $x_i \in E(x)$, we may find smaller neighborhoods $V_i \subset U_i$ such that each U_i contains a homeomorphic copy of V_j for all $j = 1, 2, \dots, k$. Let $S \subset \Sigma - K$ be a subsurface, homeomorphic to the disjoint union of k pairs of pants, such that the complementary regions of the i^{th} pair of pants partitions the ends of Σ into V_i, $B_i - V_i$ and $E - B_i$.

Since $\mathcal{V}_K$ is assumed CB, the surface S is displaceable by Lemma 5.2, so at least one of the connected components of S can be moved to be disjoint from S by a homeomorphism. Since $E(x)$ is homeomorphism invariant, we conclude that there is a copy of B_j in some V_i, possibly with $i \neq j$. Our choice of V_i now implies that there is in fact a homeomorphic copy of B_j in U_j. Thus, we have shown that, for any neighborhoods U_i of x_i, there exists j such that U_j contains a copy of B_j. Applying this conclusion to each of a nested sequence of neighborhoods of the x_i which give a neighborhood basis, we conclude that some j must satisfy this conclusion infinitely often (ie has a homeomorphic copy contained in every neighborhood of x_j), giving us some B_j which is self-similar.

Since x_i are the unique maximal points of B_i, this implies that each x_i has a neighborhood M_i homeomorphic to B_j, ie a self-similar set. Repeating this process for all of the distinct maximal types, we conclude that each maximal point has a self-similar neighborhood. Fix a collection of such neighborhoods. Since this collection is finite we may enumerate them $A_1, A_2, \dots, A_n$.

For each nonmaximal point y, Lemma 4.18 implies that there exists a neighborhood P_y of y such that $P_y \cup A_i$ is homeomorphic to some A_i, a neighborhood of a maximal point that is a successor (though not necessarily an immediate successor) of y. Since $E - \bigsqcup_i A_i$ is compact, finitely many such neighborhoods P_y cover it. Enlarging K, we may assume that it partitions the end sets into the disjoint union of such sets of the form P_y and A_i. This concludes the proof in the case where $\mathcal{M}$ is finite.

Now we treat the general case where, for some maximal types, the set $E(x)$ is a Cantor set. The strategy is essentially the same. We use the following lemma, which parallels the argument just given above.

Lemma 5.6 *Keeping the hypotheses of the proposition, let x be a maximal type with $E(x)$ a Cantor set. Then x has a neighborhood which is self-similar.*

Proof Let $A_1, \dots, A_k$ be the complementary end sets which contain points of $E(x)$, and fix a maximal end x_i in each A_i. As before, we start by showing that, for some j, every neighborhood of x_j contains a homeomorphic copy of A_j, so in particular A_j is self-similar. Let U_i be a neighborhood of x_i. For each $z \in E(x)$, let V_z be a neighborhood of z such that each of the sets U_i contains a homeomorphic copy of V_z. Since $E(x)$ is compact, finitely many such V_z cover $E(x)$, so from now on we consider only a finite subcollection that covers. Let $S \subset \Sigma - K$ be a subsurface homeomorphic to the union of k disjoint n–holed spheres, where n is chosen large enough that each complementary region of S has its set of ends either contained in one of the finitely many V_z, or containing all but one of the sets A_i.

Again, since $E(x)$ is invariant, and S is displaceable, this means that there is some V_z and some A_j such that V_z contains a homeomorphic copy of A_j. Thus, by definition of V_z, we have that U_j contains a homeomorphic copy of A_j. Repeating this for a nested sequence of neighborhoods of the x_i, we conclude that some x_j satisfies this infinitely often. This means that A_j is a stable neighborhood of x_j, hence by Lemma 4.17, each point of $E(x)$ has a stable neighborhood, which is necessarily a self-similar set. □

Now we can finish the proof as in the case where all $E(x)$ are finite, by fixing a finite cover of $\bigcup_{x \in \mathcal{M}(E)} E(x)$ by stable neighborhoods, and using Lemma 4.18 as before. □

Proof of Proposition 5.4 Let $E = \bigsqcup_{\widehat{A} \in \mathcal{A}} \widehat{A} \sqcup \bigsqcup_{P \in \mathcal{P}} P$ be the decomposition given by Proposition 5.5. By construction of the sets P and Lemma 4.18, for each $P \in \mathcal{P}$, there exists $\widehat{A} \in \mathcal{A}$ such that $P \sqcup \widehat{A} \cong \widehat{A}$. Applying this to each P iteratively, we conclude that E is homeomorphic to the disjoint union $\bigsqcup_{A \in \mathcal{A}} \widehat{A}$. Relabeling $\widehat{A}$ as A gives the desired result, and we may take L to be a subset of K. □

With this groundwork in place, we can prove Theorem 1.4. We restate it in slightly different form, for convenience.

Theorem 5.7 $\mathrm{Map}(\Sigma)$ *is locally CB if and only if there is a finite-type surface K such that the complementary regions of K each have one or infinitely many ends and zero or infinite genus, and partition of E into finitely many clopen sets*

$$E = \Big(\bigsqcup_{\widehat{A} \in \mathcal{A}} \widehat{A} \Big) \sqcup \Big(\bigsqcup_{P \in \mathcal{P}} P \Big)$$

with the following properties:

(i) *Each $\widehat{A} \in \mathcal{A}$ is self-similar, $\mathcal{M}(\widehat{A}) \subset \mathcal{M}(E)$ and $\mathcal{M}(E) = \bigsqcup_{\widehat{A} \in \mathcal{A}} \mathcal{M}(\widehat{A})$.*

(ii) *Each $P \in \mathcal{P}$ is homeomorphic to a clopen subset of some $\widehat{A} \in \mathcal{A}$.*

(iii) *For any $x_A \in \mathcal{M}(\widehat{A})$ and any neighborhood V of the end x_A in Σ, there is $f_V \in \mathrm{Homeo}(\Sigma)$ such that $f_V(V)$ contains the complementary region to K with end set $\widehat{A}$.*

Moreover, in this case $\mathcal{V}_K := \{g \in \mathrm{Homeo}(\Sigma) : g|_K = \mathrm{id}\}$ is a CB neighborhood of the identity, and K may always be taken to have genus zero if Σ has infinite genus, and genus equal to that of Σ otherwise, and if the number of isolated planar ends of Σ is finite, we may additionally take all of these ends to be punctures of K.

Note that the case where $K = \varnothing$ implies that Σ has zero or infinite genus and self-similar end space, in which case we already showed that $\mathcal{V}_\varnothing = \mathrm{Map}(\Sigma)$ is CB. In this case, conditions (ii) and (iii) are vacuously satisfied. The reader may find it helpful to refer to Figure 1 for some very basic examples, all with $\mathcal{P} = \varnothing$, and keep this in mind during the proof.

Proof of Theorem 5.7 ($\Longrightarrow$) The forward direction is obtained by a minor improvement of Proposition 5.5. Assume $\mathrm{Map}(\Sigma)$ is locally CB. Let $K \subset \Sigma$ be a finite-type surface with $\mathcal{V}_K$ a CB neighborhood of the identity and the properties given in Proposition 5.5. We may enlarge K if needed so that each of its boundary curves are separating, and so that whenever some maximal type x has $E(x)$ homeomorphic to a Cantor set, then $E(x)$ is contained in at least two complementary regions to K. This latter step can be done as follows: if $\widehat{A}$ is the unique complementary region of K containing the Cantor set $E(x)$, then glue a strip to K that separates $\widehat{A}$ into two clopen sets, each containing points of $E(x)$. Since $\widehat{A}$ is self-similar, each point of $\mathcal{M}(\widehat{A})$ has a stable neighborhood by Lemma 4.16, and so the two clopen sets of our partition are again each self-similar and each homeomorphic to $\widehat{A}$. Enlarging K further if needed, we may assume it also contains all isolated punctures if this number is finite.

Thus, we assume K now has these properties, and let $E = \big(\bigsqcup_{\widehat{A}\in\mathcal{A}} \widehat{A}\big) \sqcup \big(\bigsqcup_{P\in\mathcal{P}} P\big)$ be the resulting decomposition of E, with $\Sigma_{\widehat{A}}$ and Σ_P denoting the connected component of K with end space $\widehat{A}$ or P, respectively. We need to establish that the third condition holds. Fix $\widehat{A}$, let $x_A \in \mathcal{M}(\widehat{A})$, let $V \subset \Sigma$ be a neighborhood of the end x_A, and let $E(V) \subset \widehat{A}$ denote the end space of V. We may without loss of generality assume that V has a single boundary component. Recall that our goal is to show that the pair $V, (\Sigma - V)$ is homeomorphic to the pair $\Sigma_{\widehat{A}}, (\Sigma - \Sigma_{\widehat{A}})$.

First consider the case where $|E(x_A)| > 1$. By construction there exists $\widehat{B} \neq \widehat{A} \in \mathcal{A}$ with $E(x) \cap \widehat{B} \neq \varnothing$. Since points of $E(x_A)$ have stable neighborhoods, $\widehat{B} \cup (\widehat{A} - E(V))$ is homeomorphic to $\widehat{B}$. Moreover, if $\Sigma_{\widehat{A}}$ has infinite genus, then V and $\Sigma - \Sigma_{\widehat{A}}$ and $\Sigma - V$ all do as well, while if $\Sigma_{\widehat{A}}$ has genus 0, then so does V, and both complementary regions are of the same genus as well (equal to the genus of Σ). Thus, by the classification of surfaces, the pair $V, \Sigma - V$ is homeomorphic to $\Sigma_{\widehat{A}}, \Sigma - \Sigma_{\widehat{A}}$ and so there is some $f_V \in \mathrm{Map}(\Sigma)$ taking V to $\Sigma_{\widehat{A}}$. This is what we needed to show.

Now suppose instead $|E(x_A)| = 1$. Here we will use the displaceable subsurfaces condition to find the desired f_V. Let S be a pair of pants in the complement of K, with one boundary component equal to ∂V and another homotopic to $\partial \Sigma_{\widehat{A}}$. Since $S \subset (\Sigma - K)$, it is displaceable, so let f be a homeomorphism displacing S. Since $E(x_A) = x_A$ is an invariant set, up to replacing f with its inverse, we have $f(S) \subset V$. If, as a first case, there exists a maximal end $y \nsim x$, then $E(y)$ is also an invariant set. Thus, $f(\Sigma_{\widehat{A}}) \subset V$, hence we may take $f_V = f^{-1}$ and have $f_V(V) \supset \widehat{A}$.

If, as a second case, Σ has finite genus, then $f(\Sigma - \Sigma_A)$ necessarily contains all the genus of Σ, hence again we have $f(\Sigma_{\widehat{A}}) \subset V$. Finally, if neither of these two cases holds, then Σ has infinite or zero genus, and a unique maximal end, so $|\mathcal{A}| = 1$ and E is self-similar. Thus, $\mathrm{Map}(\Sigma)$ is CB by Proposition 3.1.

($\Longleftarrow$) For the converse, the case where $K = \varnothing$, we have that Σ has zero or infinite genus and a self-similar end space is covered by Proposition 3.1.

So suppose Σ is not zero or infinite genus with a self-similar end space, but instead we have a finite-type surface K with the properties listed. We wish to show that $\mathcal{V}_K$ is CB. Let $T \subset \Sigma$ be a finite-type surface with $\mathcal{V}_T \subset \mathcal{V}_K$, ie $T \supset K$. We need to find a finite set F and some n such that $(F\mathcal{V}_T)^n$ contains $\mathcal{V}_K$.

For each $\widehat{A} \in \mathcal{A}$, fix $x_A \in \mathcal{M}(\widehat{A})$ and let V_A be the connected component of T containing x_A. Let f_V be the homeomorphism provided by our assumption. Also, for each $P \in \mathcal{P}$, choose a homeomorphism f_P of Σ that exchanges P with a clopen subset of some $\widehat{A} \in \mathcal{A}$ which is homeomorphic to P. Let F be the set of all such $f_V^{\pm 1}$ and $f_P^{\pm 1}$.

Now suppose $g \in \mathcal{V}_K$. We can write g as a product of $|\mathcal{A}| + |\mathcal{P}|$ homeomorphisms, where each one is supported on a surface of the form Σ_A for $\widehat{A} \in \mathcal{A}$ or Σ_P for $P \in \mathcal{P}$ (adopting our notation from the previous direction of the proof).

If some such homeomorphism g_A is supported on Σ_A, then $f_V^{-1} g_A f_V$ restricts to the identity on T, so $g_A \in F\mathcal{V}_T F$. For a homeomorphism g_P supported on Σ_P, we have that $f_P^{-1} g_P f_P$ is supported in Σ_A, so $g_P \in F^2\mathcal{V}_T F^2$. This shows that $g \in (F^2\mathcal{V}_T F^2)^{|\mathcal{A}|+|\mathcal{P}|}$, which is what we needed to show. □

5.1 Examples

While the statement of Theorem 5.7 is somewhat involved, it is practical to apply in specific situations. Below are a few examples illustrating some of the subtlety of the phenomena at play. The first is an immediate consequence:

Corollary 5.8 *If Σ has finite nonzero genus and countable self-similar end space, then Σ is not locally CB.*

As another example, one could take Σ to have finite nonzero genus, and end space equal to the union of cantor set and a countable set of isolated points, accumulating on the Cantor set at exactly one point. Many other variations are possible. As a more involved example, we have the following:

Corollary 5.9 *Suppose that Σ has finite nonzero genus and self-similar end space, with a single maximal end x, but infinitely many distinct immediate precursors to x. Then* $\mathrm{Map}(\Sigma)$ *is not locally CB.*

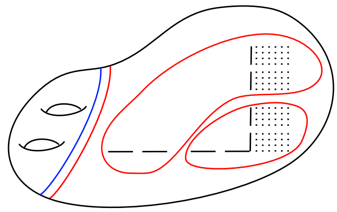

Figure 4: The subsurface with red boundary defines a CB neighborhood, while the smaller subsurface with blue boundary does not.

As a concrete example, one could construct E by taking countably many copies of a Cantor set indexed by $\mathbb{N}$, all sharing a single point in common and Hausdorff converging to that point, with the n^{th} copy accumulated everywhere by points locally homeomorphic to $\omega^n + 1$.

Proof If $\mathrm{Map}(\Sigma)$ were locally CB, then we would have a finite-type surface K as in Theorem 5.7. Since $\mathcal{M}$ is a singleton, $\mathcal{A} = \{\hat{A}\}$, $x_A = x$ and Σ_A is some neighborhood of the end x. However, by construction, $E - \hat{A}$ contains ends of only finitely many types of immediate precursors. Thus, we may choose a smaller neighborhood V of x so that $\Sigma - V$ has more distinct types of ends. Then $\Sigma - V$ cannot possibly be homeomorphic to $\Sigma - \Sigma_A$, so no such f_V exists. □

By contrast, if Σ is finite genus with end space equal to a Cantor set, or attained by the construction in Corollary 5.9 but replacing $\mathbb{N}$ with a finite number, then $\mathrm{Map}(\Sigma)$ is locally CB. We draw attention to a specific case of this to highlight the role played by $\mathcal{A}$ and $\mathcal{P}$.

Example 5.10 Let Σ be a surface of finite nonzero genus g, with E homeomorphic to the union of a Cantor set C and a Cantor set D, and a countable set Q, with $C \cap D = \{x\}$ and the accumulation points of Q equal to D, as illustrated in Figure 4. Then by Theorem 5.7, a CB neighborhood of the identity in $\mathrm{Map}(\Sigma)$ can be taken to be $\mathcal{V}_K$ where K is a finite-type subsurface of genus g with two boundary components, with one complementary region to K having x as an end, and the other containing points of both C and D. In this case, $\mathcal{A}$ and $\mathcal{P}$ are both singletons, with one complementary region in each.

The set E itself is self-similar, and the decomposition into self-similar sets given by Proposition 5.4 is trivial. However, if K' is a finite-type subsurface realizing this decomposition (with a single complementary region), then $\mathcal{V}_{K'}$ is not a CB set. Indeed, one may find a nondisplaceable subsurface in the complement homeomorphic to a three-holed sphere, where one complementary region has x as an end, one contains all the genus of Σ but no ends, and the third contains points of C, for example.

6 CB generated mapping class groups

In this section we give general criteria for when mapping class groups are CB generated, building towards the proof of Theorem 1.6.

6.1 Two criteria for CB generation

Notation 6.1 For a subset $X \subset E$, we say a family of neighborhoods U_n in E *descends to* X if U_n are nested, meaning $U_{n+1} \subseteq U_n$, and if $\bigcap_{n \in \mathbb{N}} U_n = X$. As a shorthand, we write $U_n \searrow X$. If $X = \{x\}$ is a singleton, we abuse notation slightly and write $U_n \searrow x$ and say U_n descends to x.

Definition 6.2 (limit type) We say that an end set E is *limit type* if there is a finite-index subgroup G of $\mathrm{Map}(\Sigma)$, a G–invariant set $X \subset E$, points $z_n \in E$, indexed by $n \in \mathbb{N}$ which are pairwise inequivalent, and a family of neighborhoods $U_n \searrow X$ such that

$$E(z_n) \cap U_n \neq \varnothing, \quad E(z_n) \cap U_0^c \neq \varnothing, \quad E(z_n) \subset (U_n \cup U_0^c).$$

Here $U_0^c = E - U_0$ denotes the complement of U_0 in E.

The following example explains our choice of the terminology "limit type":

Example 6.3 Suppose that α is a countable limit ordinal, and $E \cong \omega^\alpha \cdot n + 1$, with $n \geq 2$ and $E^G = \varnothing$. To see that this end space is limit type, take G to be the finite-index subgroup pointwise fixing the n maximal ends. Fix a maximal end x and a clopen neighborhood U_0 of x disjoint from the other maximal ends, and let $U_n \subset U_0$ be nested clopen sets forming a neighborhood basis of x. Since $U_n - U_{n+1}$ is closed, there is a maximal ordinal β_n such that U_n contains points locally homeomorphic to $\omega^{\beta_n} + 1$. Passing to a subsequence we may assume that all of these are distinct, and one may choose $z_n \in U_n$ to be a point locally homeomorphic to $\omega^{\beta_n} + 1$. Note that necessarily the sequence β_n converges to the limit ordinal α. The assumption that $n \geq 2$ ensures that the sets $E(z_n)$ contain points outside of U_0, and we require this in the definition to ensure that E is not self-similar.

Lemma 6.4 (limit-type criterion) *If an end set E has limit type, then* $\mathrm{Map}(\Sigma)$ *is not CB generated.*

Proof Let G, X, U_n and z_n be as in the definition of limit type. We will show G is not CB generated. Since G is finite index, this is enough to show that $\mathrm{Map}(\Sigma)$ is not CB generated. Furthermore, since $\mathrm{Map}(\Sigma)$ (and hence G) is assumed to be locally CB, it suffices to show that there is some neighborhood $\mathcal{V}_G$ of the identity in G such that for any finite set F, the set $F\mathcal{V}_G$ does not generate G.

Let $\mathcal{V}_G$ be a neighborhood of the identity in G, chosen small enough that, for every $g \in \mathcal{V}_G$ and all $n > 0$, we have $g(U_n) \subset U_0$ and $g(U_0^c) \cap U_n = \varnothing$.

Let F be any finite subset of G. Since G preserves both the set X and the set $E(z_n) \subset U_n \sqcup U_0^c$, there exists $N \in \mathbb{N}$ such that for all $n > N$ and all $f \in F$, we have

$$f(E(z_n) \cap U_n) \subset U_n.$$

The same holds for elements of $\mathcal{V}_G$.

Fix such an $n > N$, and let $x_n \in E(z_n) \cap U_n$ and $y_n \in E(z_n) \cap U_0^c$. Since $x_n \sim y_n$, there is a homeomorphism h with $h(x_n)$ lying in a small neighborhood of y_n contained in U_0^c. By our observation above, h is not in the subgroup generated by $F\mathcal{V}_G$, which shows that G is not CB generated, as desired. □

A second obstruction to CB generation is the following "rank" condition:

Definition 6.5 (infinite rank) We say $\mathrm{Map}(\Sigma)$ has *infinite rank* if there is a finite-index subgroup G of $\mathrm{Map}(\Sigma)$, a closed G–invariant set X, neighborhood U of X and points z_n, for $n \in \mathbb{N}$, each with a *stable neighborhood* (see Definition 4.14) such that

- $z_n \notin E(z_m)$ if $m \neq n$,
- for all n, $E(z_n)$ is countably infinite and has at least one accumulation point in both X and in $E - U$, and
- the set of accumulation points of $E(z_n)$ in U is a subset of X.

If the above does not hold, we say instead that $\mathrm{Map}(\Sigma)$ has *finite rank.*

Example 6.6 A simple example of such a set is as follows. Let C_n be the union of a countable set and a Cantor set, with Cantor–Bendixson rank n, and n^{th} derived set equal to the Cantor set. For each C_n, select a single point z_n of the Cantor set to be an end accumulated by genus. Now create an end space E by taking $\mathbb{N}$ copies of each C_n, arranged so that they have exactly two accumulation points x and y (and these accumulation points are independent of n). Then $X = \{x\}$ and the points z_n satisfy the definition.

Examples of surfaces with *countable* end spaces and infinite-rank mapping class groups are much more involved to describe. (Note that these necessarily must have infinite genus.) It would be nice to see a general procedure for producing families of examples.

Lemma 6.7 (infinite-rank criterion) *If* $\mathrm{Map}(\Sigma)$ *has infinite rank, then it is not CB generated.*

Proof Let G, X, U and z_n be as the definition of infinite rank. For every z_n, we define a function $\ell_n \colon G \to \mathbb{Z}$ as follows. For $\phi \in G$, define

$$\ell_n(\phi) = |(E(z_n) \cap U) - \phi^{-1}(U)| - |(E(z_n) \cap \phi^{-1}(U)) - U|.$$

That is, $\ell_n(\phi)$ is the difference between the number of points in $E(z_n)$ that ϕ maps out of U, and the number of points in $E(z_n)$ that ϕ maps into U.

Since X is G–invariant and contains all of the accumulation points of $E(z_n)$ in U, the value of ℓ_n is always finite. It is also easily verified that ℓ_n is a homomorphism. Moreover, as each z_n has a stable neighborhood (all of which are pairwise homeomorphic), for any finite collection $n_1, \dots, n_k$ one may construct, for each i, a "shift" homomorphism ϕ_i supported on a union of disjoint stable neighborhoods of $E(z_{n_i})$, taking one stable neighborhood to the next, so that $\ell_{n_i}(\phi_i) = 1$ and $\ell_{n_j}(\phi_j) = 0$ for $j \neq i$. Finally, ℓ_n is continuous; in fact for any neighborhood $\mathcal{V}$ of the identity in G which is small enough that elements of $\mathcal{V}$ fix the isotopy class of a curve separating U from $E - U$, we will have $\ell_n(\mathcal{V}) = 0$.

Thus, we have for each $k \in \mathbb{N}$ a surjective, continuous homomorphism

$$(\ell_{n_1}, \dots, \ell_{n_k}) \colon G \to \mathbb{Z}^k,$$

which restricts to the trivial homomorphism on the neighborhood $\mathcal{V}$ of the identity described above.

By Theorem 2.1, any CB set is contained in a set of the form $(F\mathcal{V})^k$ for some finite set F and $k \in \mathbb{N}$. Given any such F, choose $j > |F|$. Then restriction of $(\ell_{n_1}, \dots, \ell_{n_j})$ to the subgroup generated by $(F\mathcal{V})^k$ cannot be surjective, as $\mathcal{V}$ lies in its kernel. It follows that no CB set can generate G. Since G is finite index in $\mathrm{Map}(\Sigma)$, the same is also true for $\mathrm{Map}(\Sigma)$. □

6.2 End spaces of locally CB mapping class groups

For the remainder of this section, we assume that $\mathrm{Map}(\Sigma)$ is locally CB, our ultimate goal being to understand which such groups are CB generated. Recall that Proposition 5.4 gave a decomposition of E into a disjoint union of self-similar sets homeomorphic to $A \in \mathcal{A}$, realized by a finite-type subsurface $L \subset K$. However, as shown in Example 5.10, the neighborhood $\mathcal{V}_L$ might not be CB. We now show that $\mathcal{V}_L$ is CB generated.

Lemma 6.8 *Assume that* $\mathrm{Map}(\Sigma)$ *is locally CB. Let L be a finite-type surface whose complementary regions realize the decomposition $E = \bigsqcup_{A \in \mathcal{A}} A$ given by Proposition 5.4. Then $\mathcal{V}_L$ is CB generated.*

Furthermore, we may take L to have genus zero if Σ has infinite genus, and genus equal to that of Σ otherwise; and a number of punctures equal to the number of isolated planar (not accumulated by genus) ends of Σ if that number is finite, and zero otherwise.

For the proof, we need the following observation, which follows from well-known results on standard generators for mapping class groups of finite-type surfaces.

Observation 6.9 *Let Σ be an infinite-type surface, possibly with finitely many boundary components, and $S \subset \Sigma$ a finite-type subsurface. Then there is a finite set of Dehn twists D such that for any finite-type surface S',* $\mathrm{Map}(S')$ *is generated by D and $\mathcal{V}_S$.*

In fact, akin to Lickorish's Dehn twist generators for the mapping class group of a surface of finite type, one can find a set $\mathcal{D}$ of simple closed curves in Σ such that every curve in $\mathcal{D}$ intersects only finitely many other curves in $\mathcal{D}$, and such that the set of Dehn twists around curves in $\mathcal{D}$ generates the subgroup of $\mathrm{Map}(\Sigma)$ consisting of mapping classes supported on finite-type subsurfaces of Σ; see [14]. One can then take the set D of Observation 6.9 to be the set of Dehn twists around the curves in $\mathcal{D}$ that intersect S.

Proof of Lemma 6.8 Let K be the surface given by Theorem 5.7. For each $P \in \mathcal{P}$, there exists $\hat{A} \in \mathcal{A}$ such that $\hat{A} - \{x_A\}$ contains a homeomorphic copy of P. Choose one such $\hat{A}$ for each $P \in \mathcal{P}$, and for $\hat{A} \in \mathcal{A}$ let P_A denote the union of the elements of $\mathcal{P}$ assigned to $\hat{A}$. Let $L \subset K$ be a connected, finite-type surface with $|\mathcal{A}|$ boundary components, and such that the complementary regions of L partition E into the sets $\hat{A} \cup P_A$ as $\hat{A}$ ranges over $\mathcal{A}$. We take L to have the same number of punctures and genus as K. For each $\hat{A}$, let Σ_A denote the complementary region to L with end space $\hat{A} \cup P_A$.

If $f \in \mathcal{V}_L$, then f can be written as a product of $|\mathcal{A}|$ homeomorphisms, one supported on each surface Σ_A (and hence identifiable with an element of $\mathrm{Map}(\Sigma_A)$). So it suffices to show, for each $\hat{A} \in \mathcal{A}$, that $\mathrm{Map}(\Sigma_A)$ is generated by $\mathcal{V}_K \cap \mathrm{Map}(\Sigma_A)$, which is a CB subset of $\mathrm{Map}(\Sigma)$, together with a finite set.

Fix $\hat{A}$, let K' denote $K \cap \Sigma_A$, let $\Sigma_1, \Sigma_2, \dots, \Sigma_n$ denote the connected components of $\Sigma_A - K'$ with end spaces elements of $\mathcal{P}$, and let Σ' be the connected component with end space $\hat{A}$. Let

$$\mathcal{G} = \mathcal{V}_K \cap \mathrm{Map}(\Sigma_A) = \mathrm{Map}(\Sigma') \times \mathrm{Map}(\Sigma_1) \times \cdots \times \mathrm{Map}(\Sigma_n).$$

In view of Observation 6.9, we can find a finite set of Dehn twists D_A whose support is contained in Σ_A such that, for any finite-type surface $S' \subset \Sigma_A$, $\mathrm{Map}(S')$ is contained in the group generated by D_A and $\mathcal{G}$.

Recall from Proposition 5.5 that P_A contains no maximal points, that $A = \hat{A} \cup P_A$ is a self-similar set (and homeomorphic to $\hat{A}$), and in particular we can find a copy of P_A in any neighborhood of x_A. This implies there is some homeomorphism g_A of Σ_A with $g_A(P_A) \subset \mathrm{End}(\Sigma')$, where $\mathrm{End}(\Sigma)$ denotes the space of ends of the surface Σ'. We now set our desired finite set to be

$$\mathcal{F} = D_A \cup \{g_A\}.$$

We now show that $\mathrm{Map}(\Sigma_A)$ is generated by

$$\mathcal{G}' = \mathcal{G} \cup \mathcal{F}.$$

Let $f \in \mathrm{Map}(\Sigma_A)$. Since $\mathcal{M}(\hat{A})$ is an invariant set, we may find a neighborhood U of $\mathcal{M}(\hat{A})$ in Σ_A, which we may take to be a (infinite-type) subsurface of Σ' with a single boundary component, such that $f(U) \subset \Sigma'$. Let P'_A be a homeomorphic copy of P_A contained in $\mathrm{End}(U)$. Thus, $f(P'_A) \subset \mathrm{End}(\Sigma')$, and so there exists $h \in \mathrm{Map}(\Sigma')$ with $hf(P'_A) = g_A(P_A)$. This means $g_A^{-1}hf(P'_A) = P_A$ and therefore,

$$g_A^{-1}hf(P_A) \subset \mathrm{End}(\Sigma').$$

Thus, there exists $h' \in \mathrm{Map}(\Sigma')$ interchanging $g_A^{-1}hf(P_A)$ with $g_A(P_A)$, such that the map $h' \circ (g_A^{-1}hf)$ agrees with g_A on P_A. It follows that

$$g_A^{-1} \circ h' \circ g_A^{-1}hf|_{P_A} = \mathrm{id}.$$

Applying another element $h'' \in \mathrm{Map}(\Sigma')$, we can ensure that

$$f' = h'' \circ g_A^{-1} \circ h' \circ g_A^{-1}hf$$

is the identity on $\mathrm{End}(\Sigma_A)$ — that is, it is an element of the pure mapping class group of Σ_A. Since h, h', h'' and g_A are in $\mathcal{G}'$, it is sufficient to show that f' is also contained in the group generated by $\mathcal{G}'$.

Let S' be a genus-zero surface of finite type that contains $K' \cup f'(K')$. Since f' is a pure mapping class, for each boundary curve α of K' the curves α and $f'(\alpha)$ cut out the same subset of $\mathrm{End}(\Sigma_A)$. Hence they also cut out the same set of boundary curves of S'. But S' has genus zero, therefore the component of $S' - K'$ associated to α is homeomorphic to the component of $S' - f'(K')$ associated to $f'(\alpha)$. That is, there is a homeomorphism $g' \in \mathrm{Map}(S')$ such that

$$g' f'(K') = K'.$$

But, as mentioned above, g' (which has finite support) is in the group generated by $\mathcal{G}'$. Also, $g' f'$ fixes Σ' and hence is contained in $\mathcal{V}_K \cup \mathrm{Map}(K')$. But $\mathcal{V}_K \subset \mathcal{G}'$ and K' has finite type, so $\mathrm{Map}(K')$ is also contained in the group generated by $\mathcal{G}'$. This finishes the proof. □

Going forward, we will ignore the surface K produced earlier that defined the CB neighborhood $\mathcal{V}_K$, and instead use the surface L, which gives a simpler decomposition of the end space. The sets $P \in \mathcal{P}$ play no further role, and we focus on the decomposition $E = \bigsqcup_{A \in \mathcal{A}} A$ given by the end spaces of complementary regions to the surface L. This is the reason for our choice of notation $\widehat{A}$ for the smaller sets of the finer partition of E, for we may now abandon the cumbersome hats.

Further decompositions of end sets Now we begin the technical work of the classification of CB generated mapping class groups. As motivation for our next lemmas, consider the surface depicted in Figure 1 on the left. This surface has a mapping class group which is both locally CB and CB generated — we have not proved CB generation yet, but the reader may find it an illustrative exercise to attempt this case by hand. Here, the decomposition of E given by the surface L is $E = A \sqcup B \sqcup C$, where A and C are accumulated by genus, A and B are homeomorphic to $\omega + 1$, and C is a singleton. As well as a neighborhood of the identity of the form $\mathcal{V}_L$, any generating set must include a "handle shift" moving genus from A into C (see Definition 6.20 below), as well as a "puncture shift" that moves isolated punctures out of A and into B. If each handle was replaced by, say, a puncture accumulated by genus, one would need a shift moving these end types in and out of neighborhoods of A and C instead.

To generalize this observation to other surfaces with more complicated topology, we need to identify types of ends of Σ that accumulate at the maximal ends of the various sets in the decomposition. The sets $W_{A,B}$ defined in Lemma 6.10 and refined in Lemma 6.17 below pick out blocks of ends that can be shifted between elements A and B in $\mathcal{A}$. Ultimately, we will have to further subdivide these blocks to distinguish different ends that can be independently shifted; this is carried out in Section 6.4.

Lemma 6.10 *Assume that* $\operatorname{Map}(\Sigma)$ *is locally CB and that* E *does not have limit type. Then*:

- *For every* $A \in \mathcal{A}$, *there is a neighborhood* $N(x_A) \subset A$ *containing* x_A *such that* $A - N(x_A)$ *contains a representative of every type in* $A - \{x_A\}$.
- *For every pair* $A, B \in \mathcal{A}$, *there is a clopen set* $W_{A,B} \subset (A - N(x_A))$ *with the property that* $E(z) \cap W_{A,B} \neq \varnothing$ *if and only if*
$$E(z) \cap (A - \{x_A\}) \neq \varnothing \quad \text{and} \quad E(z) \cap (B - \{x_B\}) \neq \varnothing.$$
- *For every* $A \in \mathcal{A}$, *there is a clopen set* $W_A \subset (A - N(x_A))$ *with the property that if* $E(z) \cap (A - \{x_A\}) \neq \varnothing$ *and, for all* $B \neq A$, $E(z) \cap (B - \{x_B\}) = \varnothing$ *then* $E(z) \cap W_A \neq \varnothing$.

In other words, $W_{A,B}$ contains representatives of every type of end that appears in both $A - \{x_A\}$ and $B - \{x_B\}$, and W_A contains representatives of every type that appears only in A.

We declare $W_{A,B} = \varnothing$ if $A - \{x_A\}$ and $B - \{x_B\}$ have no common types of ends, and similarly take $W_A = \varnothing$ if each type of end in A appears also in some $B \neq A$.

Proof We start with the first assertion. If $\mathcal{M}(A)$ is a Cantor set then we can take $N(x_A)$ to be any neighborhood of x_A that does not contain all of $\mathcal{M}(A)$, and the first assertion follows since $\mathcal{M}(A)$ is the set of maximal points. Otherwise, $\mathcal{M}(A) = \{x_A\}$. Let G be the finite-index subgroup of $\operatorname{Map}(\Sigma)$ that fixes $E(x_A)$ (which we know is finite). Also recall that $A = \widehat{A} \cup P_A$. If such a neighborhood $N(x_A)$ does not exist, then there is a nested family of neighborhoods $U_n \subset \widehat{A}$ descending to x_A and points $z_n \in U_n$ where $(E(z_n) \cap \widehat{A}) \subset U_n$. We also have that $E(z_n)$ has nontrivial intersection with the complement of $\widehat{A}$, in fact if we choose V to be a neighborhood of x_A excluding z_n, then for f_V as in part (iii) of Theorem 5.7, $f_V(z_n)$ is not in A. Then, letting $X = \{x_A\}$ and assuming $U_0 = \widehat{A}$, we see that E has limit type. The contradiction proves the first assertion.

For the second assertion, fix A and $B \in \mathcal{A}$ and let

$$X = \{x \in E \mid E(x) \cap A \neq \varnothing \text{ and } E(x) \cap B = \varnothing\}.$$

Then $X \cap A$ is closed: this follows since A is closed, and if x_n is a sequence of points in $X \cap A$ converging to x_∞ but there is some point $z \in E(x_\infty) \cap B$, then any neighborhood of z would contain homeomorphic copies of neighborhoods of x_n, for sufficiently large n, contradicting the fact that $E(x_n) \cap B = \varnothing$.

Now consider a family of neighborhoods U_n of $X \cap A$ with $U_n \searrow X$ and $U_0 \cap B = \varnothing$. Let $W_n = A - (U_n \cup N(x_A))$. Since we have removed the neighborhood U_n of X, every point in W_n has a representative in B. We claim that, for some $N \in \mathbb{N}$, W_N contains a representative of all points that appear in both A and B, that is to say, $W_{A,B}$ can be taken to be W_N. To prove the claim, suppose for contradiction that it fails. Then after passing to a subsequence, we may find points z_n, all of distinct types, such that $z_n \in U_n$, $E(z_n) \cap A \neq \varnothing$ and $E(z_n) \cap B \neq \varnothing$. Since $E(z_n)$ intersects $U_0^c \supset B$, this implies that E has limit type. The contradiction proves the second assertion.

For the third assertion, consider the closed set

$$X = \{x \in E \mid E(x) \cap A \neq \varnothing \text{ and } E(x) \cap B = \varnothing \text{ for all } B \neq A\}.$$

Let U be any clopen neighborhood of $X \cap A$ in A, and let $W_A = U - N(x_A)$. Then by definition of $N(x_A)$, $(X \cap A) - N(x_A)$ contains a representative of every type appearing only in A, so this remains true of its clopen neighborhood W_A. □

6.3 Tame end spaces

Definition 6.11 An end space E is *tame* if, for every $A \in \mathcal{A}$, the point x_A has a stable neighborhood (as in Definition 4.14), and for any $A, B \in \mathcal{A}$, every maximal point in $W_{A,B}$ has a stable neighborhood.

If Σ has locally CB mapping class group, then Theorem 1.4 implies that maximal points have stable neighborhoods, so half of the tameness condition is satisfied. The other half is an assumption that will be used in the next two sections. While this seems like a restrictive hypothesis, the class of tame surfaces is very large. In fact, the following problem seems to be challenging, as the examples of nontame surfaces (excluding those which are self-similar; see Example 6.13 below) which we can easily construct all seem to have infinite-rank or limit-type like behavior.

Problem 6.12 Does there exist an example of a nontame surface whose mapping class group has nontrivial, well-defined quasi-isometry type (ie is locally, but not globally, CB and CB generated)?

Example 6.13 (nontame surfaces) Suppose $\{z_n\}_{n\in\mathbb{N}}$ is a sequence of points in an end space which are not comparable, ie for all $i \neq j$ we have neither $z_i \preccurlyeq z_j$ nor $z_j \preccurlyeq z_i$. An end space containing such a sequence may be constructed, for instance, as in Example 6.6, and even (as in that example) have the property that each z_n admits a stable neighborhood V_n. Let D denote a set consisting of the disjoint union of one copy of each stable neighborhood V_n and a singleton x, so that the sets V_n Hausdorff converge to x. Then x is a maximal point in D, but fails the stable neighborhood condition in the definition of tame, since the homeomorphism types of small neighborhoods of x do not eventually stabilize.

A surface with end space D fails the condition of Theorem 1.4 so is not locally CB, but one can easily modify this construction to provide locally, and even globally, CB examples. For instance, let E be the disjoint union of countably many copies of D, arranged to have exactly k accumulation points. If $k = 1$, the end space constructed is self-similar, with the sole accumulation point the unique maximal point. If $k > 1$, the end space may be partitioned into finitely many self-similar sets satisfying the condition of Theorem 1.4, but has immediate predecessors to the maximal points with no tame neighborhood. (However, we note that this example is infinite rank, so the mapping class group of a surface with this end type is not CB generated.)

The main application of the tameness condition is that it allows us to give a standard form to other subsets of E. We begin with a definition and some preliminary lemmas.

Definition 6.14 When $E(z)$ is countable, we will say that z is a point *of countable type*. Define $E_{\mathrm{cp}}(A, B)$ (the *countable predecessor set*) to be the subset of $W_{A,B}$ consisting of points z where z is maximal in $W_{A,B}$ and of countable type. Since $W_{A,B}$ is clopen, it has maximal points as in Proposition 4.7.

Observation 6.15 *If z is any point of countable type, then any accumulation point p of $E(z)$ satisfies $z \prec p$. Thus, if $z \in E_{\mathrm{cp}}(A, B)$, then $E(z)$ does not have any accumulation points in $W_{A,B}$ and hence $E(z) \cap W_{A,B}$ is a finite set.*

Lemma 6.16 *Suppose E is tame and* $\mathrm{Map}(\Sigma)$ *has neither limit type nor infinite rank. Then, for any $A, B \in \mathcal{A}$, the set $E_{\mathrm{cp}}(A, B)$ contains only finitely many different types.*

Proof As a first case, suppose that $\mathcal{M}(A)$ is a single point. Let G be the finite-index subgroup of $\mathrm{Map}(\Sigma)$ that fixes x_A; recall that $E(x_A)$ is finite. Now $X = \{x_A\}$ is G–invariant and since $\mathrm{Map}(\Sigma)$ does not have infinite rank, we can take $U = A$ and conclude that $E_{\mathrm{cp}}(A, B)$ has finitely many different types.

Otherwise, $\mathcal{M}(A)$ is a Cantor set. If $E(x_A)$ does not intersect B, we can take $X = E(x_A)$ and $U = B^c$. Then X is $\mathrm{Map}(\Sigma)$–invariant and again the fact that $\mathrm{Map}(\Sigma)$ does not have infinite rank implies that $E_{\mathrm{cp}}(A, B)$ has finitely many different types.

If $\mathcal{M}(A)$ is a Cantor set and $E(x_A)$ intersects B, then $E(x_A)$ intersects $W_{A,B}$ and thus $E_{\mathrm{cp}}(A, B)$ is empty. □

Lemma 6.17 *Suppose that Σ has tame end space. Then, under the hypotheses of Lemma 6.16, the sets $W_{A,B}$ from Lemma 6.10 can be chosen so that for any z in $E_{\mathrm{cp}}(A, B)$, the set $E(z) \cap W_{A,B}$ is a singleton. Such a choice specifies a set which is unique up to homeomorphism, and in this case $W_{A,B}$ is homeomorphic to $W_{B,A}$.*

Proof Fix a choice of set $W_{A,B}$ as given by Lemma 6.10. For each $z \in E_{\mathrm{cp}}(A, B)$, choose disjoint stable neighborhoods around every point in the finite set $E(z) \cap W_{A,B}$ (this set is finite by Observation 6.15) and remove all but one neighborhood, leaving the rest of $W_{A,B}$ unchanged. Denote this new set by $W'_{A,B}$. Since one such neighborhood remains, any type that was represented in $W_{A,B}$ is still represented there, so it satisfies the conditions of Lemma 6.10. We wish to show that the homeomorphism type of $W'_{A,B}$ is independent of our choices of stable neighborhoods, and that $W'_{A,B}$ is homeomorphic to $W'_{B,A}$. We prove both assertions simultaneously, by showing that $W'_{A,B}$ is homeomorphic to any choice of set $W'_{B,A}$ as defined by the same procedure.

Let $z_1, \dots, z_k \in W'_{A,B}$ be the points of $E_{\mathrm{cp}}(A, B)$; recall there is one of each type. Let $V_1, \dots, V_k$ be the chosen disjoint stable neighborhoods of these points in $W'_{A,B}$, which exist by the tameness assumption. Let $W = W'_{A,B} - \bigcup_i V_i$. Similarly, choose $V'_1, \dots, V'_k$ to be disjoint stable neighborhoods of points of countable predecessor type in $W_{B,A}$ so that V_i is homeomorphic to V'_i, and let $W' = W'_{B,A} - \bigcup_i V'_i$. We start by showing that

$$W \cup W'_{B,A} \cong W'_{B,A}.$$

This is because, for any point in $x \in W$, there is a point $y \in W'_{B,A}$ that is maximal in $W'_{A,B}$, where y is an accumulation point of $E(x)$. Hence, by Lemma 4.18, there is a neighborhood U_x of x and stable neighborhood V_y of y such that $U_x \cup V_y$ is homeomorphic to V_y. Since W is compact, finitely many such neighborhoods are enough to cover W and, shrinking these neighborhoods if needed, we can write W as the disjoint union of finitely many such neighborhoods. Thus, W can be absorbed into $W'_{B,A}$.

Similarly we have that $W' \cup W'_{A,B}$ is homeomorphic to $W'_{A,B}$. That is,

$$\begin{aligned} W'_{A,B} &\cong W'_{A,B} \cup W' \cong W \cup W' \cup \Big(\bigcup_i V_i\Big) \cong W \cup W' \cup \Big(\bigcup_i V'_i\Big) \\ &\cong W \cup W'_{B,A} \cong W'_{B,A}. \end{aligned} \qquad \square$$

Going forward, we will use $W_{A,B}$ to denote the (well-defined up to homeomorphism) sets constructed in the lemma, each containing a single representative of each of its countable predecessor types.

6.4 Classification of CB generated mapping class groups

The purpose of this section is to prove Theorem 1.6, namely, the statement that the necessary conditions for CB generation introduced in Section 6.1 are also sufficient for tame surfaces.

We continue with the notation and conventions introduced in the previous section, in particular the following.

Convention Going forward, we let L denote the finite-type surface furnished by Proposition 5.4, so that the complementary regions to L produce a decomposition $E = \bigsqcup_{A \in \mathcal{A}} A$, where each A is self-similar, and we have $\bigsqcup \mathcal{M}(A) = \mathcal{M}(E)$.

The next proposition is the main technical ingredient in the proof of Theorem 1.6. It says that, by using elements from a CB set, one may map any neighborhood U of x_A in E homeomorphically onto A while pointwise fixing any set $B \in \mathcal{A}$ which shares no end types with $A - U$.

Proposition 6.18 *Assume that E is tame and not of limit type, and that* $\mathrm{Map}(\Sigma)$ *does not have infinite rank. Then there is a finite set $F \subset \mathrm{Map}(\Sigma)$ such that the following holds*:

Let $A \in \mathcal{A}$, and let $U \subset A$ be a neighborhood of x_A. If $\mathcal{B}_U \subset \mathcal{A}$ is a subset that satisfies $E(y) \cap \big(\bigcup_{B \in \mathcal{B}_U} B\big) \neq \varnothing$ for all $y \in A - U$, then there is an element f in the group generated by F and $\mathcal{V}_L$ with $f(U) = A$, and $f|_C = \mathrm{id}$ for all $C \in (\mathcal{A} - \mathcal{B}_U)$.

Proof The proof consists of several preliminary structural results on end spaces, carried out in Steps 1–4; the set U and $\mathcal{B}_U$ are introduced in the final step.

Step 1: decomposition of the sets $A \in \mathcal{A}$ Fix $A \in \mathcal{A}$. For every $B \in \mathcal{A}$, consider a copy of $W_{A,B} \subset A$ as in Lemma 6.17, as well as a homeomorphic copy of W_A. A short argument shows that we may choose these sets to be pairwise disjoint, so that we have $W_{A,B} \cap W_{A,B'} = \varnothing$ whenever $B \neq B'$ and $W_{A,B} \cap W_A = \varnothing$ for all B. This is as follows: enumerate the sets $B_1, B_2, \dots, B_k$ of $\mathcal{A} - \{A\}$ and perform our original construction to obtain W_{A,B_1}. This set is disjoint from $N(x_A)$. By self-similarity, there is a homeomorphic copy of A inside $N(x_A)$, hence we may find a set W_{A,B_2} disjoint from W_{A,B_1} and also disjoint from a smaller copy of $N(x_A)$. Continuing in this manner, we may produce the desired sets. Doing this one more time, we also find a disjoint copy of W_A. We keep these sets (and refer to them to by this notation, $W_{A,B}$ and W_A) for the remainder of the proof.

Let

$$T_0 = W_A \sqcup \Bigl(\bigsqcup_{B \in \mathcal{A} - \{A\}} W_{A,B} \Bigr) \subset A.$$

By construction, for every $y \in A - \{x_A\}$, $E(y)$ intersects T_0 by Theorem 1.4.

Let $V_1 = A - T_0$ and consider a family of neighborhoods $V_k \searrow x_A$. Each V_k contains a copy of A and hence a copy T_k of T_0. After dropping some of the sets V_k from the nested sequence and reindexing, we can assume $T_1 \subset (V_1 - V_2)$. Continuing in this way, we find a new nested sequence of neighborhoods, which we again denote by V_k, so that $(V_k - V_{k+1})$ contains a copy T_k of T_0. In particular, the sets T_k are disjoint.

Our next goal is to modify this construction so that we in fact have $(V_k - V_{k+1}) \cong T_k$, ie we obtain a nested family of neighborhoods such that the annular regions between them are homeomorphic to the sets T_k above. To do this, we first show that we can distribute the set

$$Q = (V_1 - V_2) - T_1$$

among finitely many of the other sets T_k, with $k > 1$, while preserving the homeomorphism class of the T_k; and then proceed iteratively.

For each point $y \in Q$, $E(y)$ intersects T_0 and hence y has a neighborhood $V_y \subset Q$ that has a homeomorphic copy inside T_0. Since Q is compact, finitely many such neighborhoods are sufficient to cover Q. Making some of these neighborhoods smaller, we can write $Q = Q_1 \sqcup \cdots \sqcup Q_m$, where every Q_i has a copy in T_0 and hence in every T_k. For $j = 1, \dots, m$ and $k \equiv j \mod m$ let Q'_k be the copy of Q_j in V_k. For $k = 1, \dots, m$ define

$$T'_k = (T_k - Q'_k) \cup Q_k,$$

and for $k > m$ define

$$T'_k = (T_k - Q'_k) \cup Q'_{k-m}.$$

Each T'_k is still homeomorphic to T_0, the sets T'_k are disjoint and every point in $(V_1 - V_2)$ is contained in some T'_k. Note that T_0 is not modified.

Similarly to the above, we can distribute the points in

$$Q' = (V_2 - V_3) - \bigcup_{k \geq 1} T'_k$$

among the sets T'_k, with $k = 2, 3, \dots$, without changing their topology. That is, we obtain a family T''_k of disjoint sets homeomorphic to T_0 whose union covers $A - V_3$, without modifying T_0 or T'_1. Continuing in this way, every T_k is modified finitely many times and stabilizes after k steps. Thus, $\{T_k^{(k)} \mid k \in \mathbb{N}\}$ is a family of disjoint copies of T_0 that covers $A - \{x_A\}$. To simplify notation, denote $T_k^{(k)}$ by $T_k(A)$. To summarize,

$$A - \{x_A\} = \bigsqcup_{k \geq 0} T_k(A),$$

and, defining

$$U_n := \bigsqcup_{k \geq n} T_k(A),$$

we have

$$U_n \searrow \{x_A\}.$$

Since $T_0 = W_A \sqcup \left(\bigsqcup_{B \neq A} W_{A,B}\right)$, we have a similar decomposition of each homeomorphic set $T_k(A)$ into sets homeomorphic to W_A and $W_{A,B}$, which we notate by

$$T_k(A) = W_A^k \sqcup \left(\bigsqcup_{B \in \mathcal{A} - \{A\}} W_{A,B}^k \right),$$

where, for $k \in \mathbb{N}$, W_A^k is a set homeomorphic to W_A and $W_{A,B}^k$ is a set homeomorphic to $W_{A,B}$.

We also have the above decomposition for every $B \in \mathcal{A} - \{A\}$. For notational convenience, when $k < 0$, we define

$$W_{A,B}^k := W_{B,A}^{-k-1}.$$

Step 2: a first shift map Using the decomposition above, we define the first homeomorphism (of several) that shifts points between A and B. Since the sets $W^k_{A,B}$ for $k \in \mathbb{Z}$ are disjoint and homeomorphic and Hausdorff converge to the points x_A and x_B as k approaches ∞ and $-\infty$, respectively, there exists a homeomorphism $\eta_{A,B}$ such that

$$\eta_{A,B}(W^k_{A,B}) = W^{k-1}_{A,B} \quad \text{for all } k \in \mathbb{Z},$$

and restricts to the identity elsewhere in E. Fix one such map for each (unordered) pair $A, B \in \mathcal{A}$. Visually, the map $\eta_{A,B}$ pushes a copy of $W_{A,B}$ out of A and into B.

Step 3: shifting countable predecessor ends independently Now we define homeomorphisms allowing one to shift the countable predecessor ends one by one. As motivation, consider, for instance, a surface with $E \cong \omega \cdot 2 + 1$, such that E^G and the closure of $E - E^G$ are both homeomorphic to $\omega \cdot 2 + 1$, as shown in Figure 5. There are two maximal ends, $\mathcal{A} = \{A, B\}$, and we have the simple situation where $W_{A,B} = T_0$ consists of one of each type of nonmaximal end. The map $\eta_{A,B}$ shifts ends of both types towards B, simultaneously. However, there is evidently a homeomorphism of Σ which pointwise fixes $E - E^G$ and shifts the nonmaximal ends of E^G.

For $z \in E_{\mathrm{cp}}(A, B)$, let $W^k_{A,B}(z) \subset W^k_{A,B}$ be a stable neighborhood of the unique intersection point of $E(z)$ with $W^k_{A,B}$. By making these neighborhoods smaller, we can assume that the $W^k_{A,B}(z)$ for different $z \in E_{\mathrm{cp}}(A, B)$ are disjoint. (This is a very slight abuse of notation since $W^k_{A,B}(z)$ depends only on the equivalence class of z under $\sim$, not the point itself.) Define $\eta_{A,B,z}$ to be a homeomorphism of Σ such that

$$\eta_{A,B,z}(W^k_{A,B}(z)) = W^{k-1}_{A,B}(z) \quad \text{for } k \in \mathbb{Z}$$

and acts by the identity elsewhere in E. Note that the actions of $\eta_{A,B,z}$ on E commute with each other and have support in $A \cup B$.

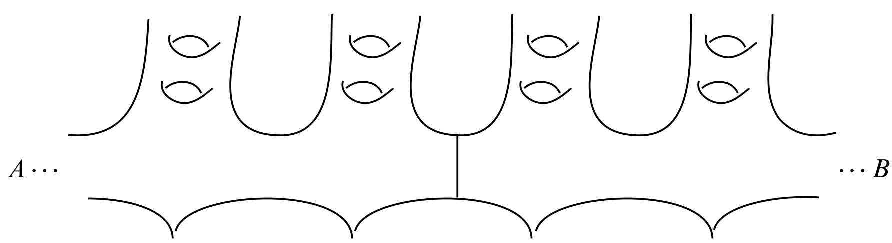

Figure 5: $E - E^G$ and the nonmaximal ends of E^G can be shifted independently.

Step 4: standard decomposition for sets of shared ends Define

$$E(A,B)=\bigsqcup_{k=0}^{\infty} W^k_{A,B}.$$

The following claim shows that clopen subsets of $E(A,B)$ have a standard form:

Claim *Let $W\subset E(A,B)$ be any clopen set in $E(A,B)$ containing $W_{A,B}$ and disjoint from x_A. For $z\in E_{\mathrm{cp}}(A,B)$, let $p_z(W)=|E(z)\cap W|$. Then W is homeomorphic to the set*

$$W_{A,B}\sqcup\Bigg(\bigsqcup_{z\in E_{\mathrm{cp}}(A,B)}\ \bigsqcup_{k=1}^{p_z(W)-1} W^k_{A,B}(z)\Bigg).$$

Recall that $W_{A,B}\subset T_0$ was a fixed set, chosen in Step 1. However, note that this structure theorem also applies to any clopen subset of $E(A,B)$ which contains a *homeomorphic copy* of $W_{A,B}$.

Proof of claim For $z\in E_{\mathrm{cp}}(A,B)$ and $y\in E(z)\cap(W-W_{A,B})$, choose a stable neighborhood V_y of y in W. Making the neighborhoods small enough, we can assume they are disjoint from each other and from $W_{A,B}$. Since stable neighborhoods are canonical, we can map the union of these neighborhoods homeomorphically to

$$\bigsqcup_{z\in E_{\mathrm{cp}}(A,B)}\ \bigsqcup_{k=1}^{p_z(W)-1} W^k_{A,B}(z).$$

It remains to show that if $p_z(W)=1$ for every $z\in E_{\mathrm{cp}}(A,B)$, then W is homeomorphic to $W_{A,B}$.

For every point in $y\in(W-W_{A,B})$, there is a point $x\in W_{A,B}$ that is maximal in $W_{A,B}$ where x is an accumulation point of $E(y)$. By the tameness assumption, x has a stable neighborhood and by Lemma 4.18, for any stable neighborhoods V_x of x and any neighborhood V_y of y, $V_x\cup V_y$ is homeomorphic to V_x. Taking a cover of $W-W_{A,B}$ by such neighborhoods, we conclude that

$$W=(W-W_{A,B})\cup W_{A,B}\cong W_{A,B}. \qquad \square$$

Step 5: finishing the proof of Proposition 6.18 Let

$$F=\{\eta^{\pm1}_{A,B},\eta^{\pm1}_{A,B,z}\mid B\in\mathcal{A}-\{A\}\text{ and }z\in E_{\mathrm{cp}}(A,B)\}.$$

Let $U\subset A$ be a neighborhood of x_A and let $\mathcal{B}_U\subset\mathcal{A}-\{A\}$ be as in the statement of the proposition. The homeomorphism $\prod_{B\in\mathcal{B}_U}\eta^{-1}_{A,B}$ shifts the sets $W_{B,A}$ from $\bigsqcup_{B\in\mathcal{B}_U} B$

into A, and, in particular,

$$\bigsqcup_{B\in\mathcal{B}_U} W_{A,B} \subset A - \Big(\prod_{B\in\mathcal{B}_U} \eta_{A,B}^{-1}\Big)(U).$$

Thus, up to applying this homomorphism, we may assume that U is sufficiently small that its complement contains $\bigsqcup_{B\in\mathcal{B}_U} W_{A,B}$, the subset of T_0.

Fix $B_1 \in \mathcal{B}_U$. Since $(A-U)\cap E(A,B_1)$ contains W_{A,B_1}, the claim proved in Step 4 implies that $(A-U)\cap E(A,B_1)$, it is homeomorphic to the standard set

$$W_{A,B_1} \sqcup \Big(\bigsqcup_{z\in E_{\mathrm{cp}}(A,B_1)} \bigsqcup_{k=1}^{p_z(W)-1} W_{A,B_1}^k(z)\Big)$$

in A, and the complements of both this standard set and of $(A-U)\cap E(A,B_1)$ in A are homeomorphic (each is easily seen to be homeomorphic to A). Thus, by the classification of surfaces there is a homeomorphism v_1 supported on the complementary region to L with end space A, hence in $\mathcal{V}_L$, taking $(A-U)\cap E(A,B)$ to this standard set. However, by construction, the image of this standard set under

$$\eta_{A,B_1} \circ \prod_{z\in E_{\mathrm{cp}}(A,B_1)} \eta_{A,B_1,z}^{p_z(W)-1}$$

is disjoint from A, and the image of its complement in A is equal to A. Let

$$U' = \eta_{A,B_1} \circ \prod_{z\in E_{\mathrm{cp}}(A,B_1)} \eta_{A,B_1,z}^{p_z(W)-1} \circ v_1(U).$$

Note that $\mathcal{B}_{U'} = \mathcal{B}_U - \{B_1\}$. We now repeat the process above using $B_2 \in \mathcal{B}_{U'}$ and U' and produce an element of the subgroup generated by F and $\mathcal{V}_L$ which takes U' to a subset of A containing $E(A,B_2)$. Iterating this process for each $B \in \mathcal{B}_U$ achieves the desired result. □

We are almost ready to prove the main result of this section. In order to do so, we need another finite set of mapping classes, the *handle shifts*, which we define now. See also [14, Section 6] for earlier use of this class of maps.

Definition 6.19 An *infinite strip with genus* is the surface $\mathbb{R}\times[-1,1]$ with a handle attached to the interior of each set $[m,m+1]\times[0,1]$ so that $(x,y)\mapsto(x+1,y)$ is a homeomorphism of the surface.

A *handle shift* on the infinite strip with genus is the mapping class of the homeomorphism h which pointwise fixes the boundary, agrees with $(x, y) \mapsto (x+1, y)$ outside an ϵ–neighborhood of the boundary, and on the ϵ–neighborhood agrees with $(x, y) \mapsto (x + (1 - |y|)/\epsilon, y)$.

Definition 6.20 Suppose that Σ has locally CB mapping class group and L is a surface as in Lemma 6.8 We call a (infinite-type) subsurface $R \subset \Sigma$ an *infinite strip with genus* in Σ if it is homeomorphic to an infinite strip with genus, and has the property that the complement of R in each complementary region to L has infinite genus.

A *handle shift* on R is the mapping class of the map h above (under our identification), extended to agree with the identity on the complement of R.

Recall that the *pure mapping class group*, denoted by $\mathrm{PMap}(\Sigma)$, is the subgroup of $\mathrm{Map}(\Sigma)$ which pointwise fixes E. We now prove a lemma on generating pure mapping classes.

For each pair (A, B) such that x_A and x_B are both accumulated by genus, let $R_{AB} \subset \Sigma$ be an infinite strip with genus, with one end in A and one end in B. We may choose these (one at a time) so that they are disjoint subsurfaces of Σ. Fix also a handle shift $h_{AB} \in \mathrm{Homeo}(\Sigma)$ on R_{AB}.

Lemma 6.21 (generating $\mathrm{PMap}(\Sigma)$) *Let G be a subgroup of* $\mathrm{Map}(\Sigma)$ *containing all mapping classes supported on finite-type subsurfaces, all mapping classes that fix each of the boundary components of L and the handle shifts h_{AB} defined above. Then G contains* $\mathrm{PMap}(\Sigma)$.

Proof For $A \in \mathcal{A}$, let Σ_A denote the connected component of $\Sigma - L$ with end space A, and let ∂_A denote its boundary component. Let $g \in \mathrm{PMap}(\Sigma)$. Then $g(\Sigma_A)$ also has end space A, and a single boundary component $g(\partial A)$. Let $T \subset \Sigma$ be a connected, finite-type subsurface large enough to contain $L \cup g(L)$. If, for each $A \in \mathcal{A}$, the surface $\Sigma_A \cap T$ is homeomorphic rel ∂T to $g(\Sigma_A) \cap T$, then there is a mapping class ϕ supported on T such that $\phi g(L) = L$, preserving each of its boundary components, which proves what we needed to show.

So we are reduced to the case where, for some A, the surface $\Sigma_A \cap T$ is not homeomorphic to $g(\Sigma_A) \cap T$. Both are connected surfaces with the same number of boundary components, so we conclude that they must have different genus. In particular, this

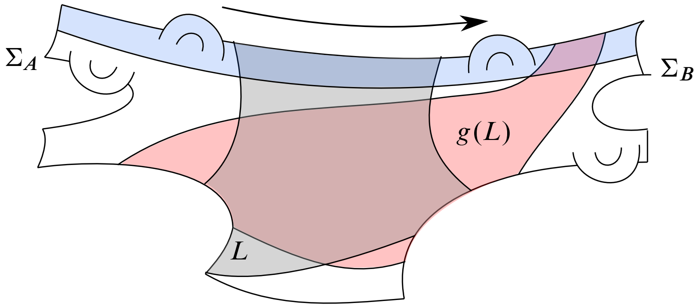

Figure 6: T containing L and $g(L)$, and the domain $\phi(R_{AB})$ of the handle shift.

only occurs if Σ is itself of infinite genus, for otherwise we choose K by convention to contain all the genus of Σ.

Without loss of generality, assume that the genus of $g(\Sigma_A) \cap T$ is larger than that of $\Sigma_A \cap T$. Since T is finite genus, there must also be another $B \in \mathcal{A}$ such that the genus of $g(\Sigma_B) \cap T$ is smaller than that of $\Sigma_B \cap T$. Since L is chosen so that complementary regions have either zero or infinite genus, we conclude that $\mathcal{M}(A)$ and $\mathcal{M}(B)$ must be accumulated by genus.

Consider the handle shift h_{AB} supported on R_{AB}, which has one end in A and one end in B. Let ϕ be a homeomorphism preserving the ends of Σ, preserving each of the boundary components of L, and such that the intersection of $\phi(R_{AB})$ with $T \cap (g(\Sigma_A) - \Sigma_A)$ and with $T \cap (\Sigma_B \cap g(\Sigma_B))$ each have genus one, and $\phi(R_{AB}) \cap T$ has genus two (so there is no genus elsewhere in T), and so that, up to replacing h_{AB} with its inverse, $\phi h_{AB} \phi^{-1}$ shifts the genus from $T \cap g(\Sigma_A)$ into $T \cap \Sigma_B$. See Figure 6 for an illustration in a simple setting. Such a homeomorphism ϕ exists by the classification of surfaces, and our stipulation that the complement of R_{AB} have infinite genus in complementary regions of L.

Then the genus of $\phi h_{AB} \phi^{-1} g(\Sigma_A) \cap T$ is one less than that of $g(\Sigma_A) \cap T$, and the genus of $\phi h_{AB}^{-1} \phi^{-1} g(\Sigma_A) \cap T$ is one more, and there is no change otherwise in the genus of complementary regions. Continuing in this fashion, one may iteratively modify g by composing by elements of G so as to arrive at a homeomorphism g' with the property that $\Sigma_A \cap T$ is homeomorphic to $g'(\Sigma_A) \cap T$ for all $A \in \mathcal{A}$, which is what we needed to show. □

A CB generating set We are now in a position to prove the main theorem on CB generation. Our CB generating set will consist of $\mathcal{V}_K$, together with the finite set consisting of the Dehn twists D from Observation 6.15, the finite set F from Proposition 6.18, the handle shifts h_{AB}, and a finite collection of homeomorphisms g_{AB} (to be specified), one for each pair $A, B \in \mathcal{A}$ such that x_A and x_B are of the same type.

Proof of Theorem 1.6 One direction follows from Lemmas 6.4 and 6.7. We prove the other direction. For this, we show that the generating set described in the paragraph above (after giving precise definitions of g_{AB}) is in fact CB.

Let $\mathcal{V}_K \cup D$ be the CB set given by Observation 6.15 (recall that D is a finite collection of Dehn twists). Let F be the finite set from Proposition 6.18. For each pair of maximal points x_A, x_B in E^G, let h_{AB} be the handle shift defined above Lemma 6.21. Let χ be the CB set consisting of $\mathcal{V}_K \cup D$ together with the homeomorphisms from F and all the h_{AB}. By Lemma 6.21, we already know that this set generates the pure mapping class group, so we start by considering only the action on the end space.

We show first that χ generates the pointwise stabilizer of $\{x_A : A \in \mathcal{A}\}$. After this, we will add finitely many more homeomorphisms g_{AB} to generate $\mathrm{Map}(\Sigma)$.

Suppose that ϕ fixes each of the points x_A. We proceed inductively on the number of elements of $\mathcal{A}$ which are pointwise fixed by the action of ϕ on E. Let $\mathcal{A}_{\mathrm{id}}$ denote the subset (possibly empty) of $\mathcal{A}$ such that, for each $A \in \mathcal{A}_{\mathrm{id}}$, the ends of A are pointwise fixed by ϕ. Let $\mathcal{A}^c = \mathcal{A} - \mathcal{A}_{\mathrm{id}}$. Choose a set $A \in \mathcal{A}^c$. For every $B \neq A \in \mathcal{A}^c$, let $U_B = B - \phi(A)$. Then for every end $z \in (B - U_B) \subset \phi(A)$, there is some end $y \sim z$ which lies in A. Hence, by Proposition 6.18 setting $\mathcal{B}_{U_B} = \mathcal{A} - \{A, B\}$, there is an element g in the group generated by χ with support in $A \cup B$ that sends U_B to B. In particular, $g\,\phi(A) \cap B = \varnothing$ and the restriction of $g\,\phi$ to sets in $\mathcal{A}_{\mathrm{id}}$ is still the identity.

Repeating this for each element of $\mathcal{A}^c$, we may modify ϕ by elements of χ to obtain a map ϕ' such that $\phi'(A)$ is disjoint from every $C \in \mathcal{A} - \{A\}$, ie $\phi'(A) \subset A$, and so that ϕ' restricts to identity on each element of $\mathcal{A}_{\mathrm{id}}$. Letting $U = \phi'(A)$, we see that the conditions of Proposition 6.18 are again satisfied taking $\mathcal{B}_U = \mathcal{A}^c$. Hence, there is a $g' \in \langle \chi \rangle$ that is also the identity on every set in $\mathcal{A}_{\mathrm{id}}$, and sends U to A. Thus, $g'\phi'(A) = A$ and we may take some $\psi \in \mathcal{V}_L$ such that the restriction of $\psi g' \phi'$ to A is the identity.

Continuing in this way, at every step, we increase the number of sets in $\mathcal{A}_{\mathrm{id}}$, and eventually obtain a homeomorphism which pointwise fixes all ends. Since χ generates $\mathrm{PMap}(\Sigma)$, we conclude that $\phi \in \langle \chi \rangle$.

Now we show that there is a finite set F' such that $\chi \cup F'$ generates $\mathrm{Map}(\Sigma)$. Construct F' as follows. For any $A, B \in \mathcal{A}$ such that points in $\mathcal{M}(A)$ and $\mathcal{M}(B)$ are of the same type, choose one element $g_{A,B}$ sending $N(x_A)$ to $N(x_B)$ (recall that these are stable neighborhoods) and restricting to the identity on every set in $\mathcal{A}-\{A, B\}$. Let F' be the set of all such chosen $g_{A,B}$. To see that $\chi \cup F'$ generates, let $\phi \in \mathrm{Map}(\Sigma)$. Suppose $\phi(x_A) \in B$. We modify ϕ to a map ϕ' in one of the following ways.

Case 1 Assume $\phi(x_B) \neq x_B$. There is a $\psi \in \mathcal{V}_L$ with support in B that sends $\phi(x_A)$ to x_B and hence

$$\phi' = g_{A,B}\psi\phi$$

fixes x_A.

Case 2 Assume $\phi(x_B) = x_B$. Then $\mathcal{M}(B)$ has more than one point and hence it is a Cantor set. Take a map $\psi \in \mathcal{V}_L$ with support in B that sends $\phi(x_A)$ to x_B and sends x_B to a point in $B - N(x_B)$. Then

$$\phi' = \psi^{-1} g_{A,B}\psi\phi$$

sends x_A to x_A and still fixes x_B.

The number of points x_A that are fixed by ϕ' is one more than that for ϕ. Hence, after repeating this process finitely many times, we arrive at an element fixing each maximal point, hence generated by χ. This finishes the proof. □

7 Classification of CB mapping class groups

In this section we prove Theorem 1.7 classifying the surfaces Σ for which the group $\mathrm{Map}(\Sigma)$ is CB. In the case where E is uncountable, we will add the hypothesis that Σ is tame. However, we expect the classification theorem to hold without this additional hypothesis, since it is only used in the very last portion of the proof.

Note that the telescoping case occurs only when E is uncountable, by Proposition 3.7.

Proof of Theorem 1.7 If Σ has zero or infinite genus and is either telescoping or has self-similar end space, then it was shown in Propositions 3.1 and 3.5 that $\mathrm{Map}(\Sigma)$ is CB, with no hypothesis on tameness. We prove the other direction. Assume that Σ has a CB mapping class group. By Example 2.4, this implies that Σ has zero or infinite genus. Also, being globally CB, $\mathrm{Map}(\Sigma)$ is in particular locally CB so the end space admits a decomposition $E = \bigsqcup_{A \in \mathcal{A}} A$ into finitely many self-similar sets as in Theorem 1.4. Then Example 2.5 implies that, if we take such a decomposition with $\mathcal{A}$ of minimal cardinality, then $\mathcal{A}$ has either one or two elements. Finally, if $\mathcal{A}$ is a singleton, then E is self-similar. Thus, we only need to take care of the case where $\mathcal{A}$ has exactly two elements.

Example 2.5 also shows that, if $\mathcal{A} = \{A, B\}$, then $\mathcal{M}(A)$ and $\mathcal{M}(B)$ are either both singletons or Cantor sets. A slight variation on the argument there also allows us to eliminate the case where they are both Cantor sets: if points of $\mathcal{M}(A)$ are not of the same type as those in $\mathcal{M}(B)$, then one may construct a nondisplaceable subsurface just as in the example by having $\mathcal{M}(A)$ play the role of the singleton. Otherwise, points of $\mathcal{M}(A)$ and $\mathcal{M}(B)$ are all of the same type and hence

$$\mathcal{M}(E) = \mathcal{M}(A) \cup \mathcal{M}(B) = E(x_A),$$

and Lemmas 5.6 and 4.18 together imply that E is self-similar.

Thus, we can assume that $\mathcal{M}(A) = \{x_A\}$ and $\mathcal{M}(B) = \{x_B\}$. We start by showing in this case that $E_{\mathrm{mc}}(A, B) = \varnothing$. To show this, suppose for contradiction that we have some $z \in E_{\mathrm{mc}}(A, B)$. Then $E(z)$ accumulates to both x_A and x_B and since z is maximal in $E - \{x_A, x_B\}$, the set $E(z)$ has no other accumulation points. As in Lemma 6.7, we can define a continuous homomorphism to $\mathbb{Z}$ on the subgroup that pointwise fixes $\{x_A, x_B\}$ (which is of index at most two in $\mathrm{Map}(\Sigma)$), via

$$\ell(\phi) = |\{x \in E(z) : x \in A, \phi(x) \in B\}| - |\{x \in E(z) : x \in B, \phi(x) \in A\}|.$$

Let $U_0 \subset A$ be a neighborhood of z not containing x_A. Since $z \in E_{\mathrm{cp}}(A, B)$, we can find a homeomorphic copy $U_1 \subset B$ of U_0 in B. Since A and B are self-similar, we may find disjoint homeomorphic copies $U_2, U_3, \dots$ of U_0 in A descending to x_A, and homeomorphic copies $U_{-1}, U_{-2}, \dots$ of U_0 in B descending to x_B. Let η be a homeomorphism that sends U_i to U_{i+1} and restricts to the identity everywhere else. Then $\ell(\eta^n) = n$, so the homomorphism ℓ is unbounded and $\mathrm{Map}(\Sigma)$ is not CB. This gives the desired contradiction, so we conclude that $E_{\mathrm{mc}}(A, B) = \varnothing$. Note that, in particular, this implies E is not countable.

We now show that E is telescoping. Let $N(x_A)$ and $N(x_B)$ be as in Lemma 6.10. Let V_1 and V_2 be subsurfaces with a single boundary component, such that the end space of V_1 is $N(x_A)$ and that of V_2 is $N(x_B)$. We will check the definition of telescoping by using these neighborhoods of $x_1 = x_A$ and $x_2 = x_B$.

Let $W_1 \subset V_1$ and $W_2 \subset V_2$ be neighborhoods of x_A and x_B respectively. Let S be a finite-type subsurface, homeomorphic to a pair of pants, whose complementary regions partition E into W_1, V_2 and the remaining ends. Provided $N(x_A)$ and $N(x_B)$ are chosen small enough, condition (ii) of Theorem 1.4 ensures that either Σ has genus zero, or $\Sigma - (V_1 \cup V_2)$ has infinite genus.

Let f_1 be a homeomorphism displacing S. We may also assume that f_1 fixes x_A and x_B, since existence of a nondisplaceable subsurface in the finite-index subgroup of $\mathrm{Map}(\Sigma)$ stabilizing x_A and x_B is sufficient to show that $\mathrm{Map}(\Sigma)$ is not CB. Then, up to replacing f_1 with its inverse, we have $f_1(\Sigma - W_1) \subset V_2$. A similar argument gives a homeomorphism f_2 with $f_2(\Sigma - W_2) \subset V_1$ and so the second condition in the definition of telescoping is satisfied.

For the first condition, we need to find a homeomorphism of the subsurface $\Sigma - V_2$ that maps W_1 to V_1. By Lemma 4.18, we know that V_1 and W_1 are homeomorphic — their end sets are homeomorphic, and they each have zero or infinite genus and one boundary component — so we need only show that their complements are homeomorphic and apply the classification of surfaces. Since, as remarked above, Σ either has genus zero or $\Sigma - (V_1 \cup V_2)$ has infinite genus, we need only produce such a homeomorphism on the level of end spaces. Here we will finally invoke tameness. Let

$$\Sigma' = \Sigma - (V_1 \cup V_2).$$

By definition of $N(x_A)$, for any end z of $V_1 - W_1$ there exists a maximal point $x \in W_1$ with $z \preccurlyeq x$. Tameness means that x has a stable neighborhood. Since x is not of countable type, it is necessarily an accumulation point of $E(z)$ (even if z and x are of the same type), and hence Lemma 4.18 implies that z has a neighborhood U_z such that $U_z \cup V_x$ is homeomorphic to V_x. Thus, on the level of ends, the end space of Σ' is homeomorphic to that of its union with U_z.

Since the end space of $V_1 - W_1$ is compact, it may be covered by finitely many such neighborhoods U_z (varying z); applying the procedure above to each of them in turn produces the desired homeomorphism on the level of end spaces, showing the two subsurfaces are homeomorphic. □

References

[1] **J Aramayona**, **A Fossas**, **H Parlier**, *Arc and curve graphs for infinite-type surfaces*, Proc. Amer. Math. Soc. 145 (2017) 4995–5006 MR Zbl

[2] **J Aramayona**, **P Patel**, **N G Vlamis**, *The first integral cohomology of pure mapping class groups*, Int. Math. Res. Not. 2020 (2020) 8973–8996 MR Zbl

[3] **J Bavard**, *Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire*, Geom. Topol. 20 (2016) 491–535 MR Zbl

[4] **J Bavard**, **S Dowdall**, **K Rafi**, *Isomorphisms between big mapping class groups*, Int. Math. Res. Not. 2020 (2020) 3084–3099 MR Zbl

[5] **R Camerlo**, **S Gao**, *The completeness of the isomorphism relation for countable Boolean algebras*, Trans. Amer. Math. Soc. 353 (2001) 491–518 MR Zbl

[6] **R Chamanara**, *Affine automorphism groups of surfaces of infinite type*, from "In the tradition of Ahlfors and Bers, III" (W Abikoff, A Haas, editors), Contemp. Math. 355, Amer. Math. Soc., Providence, RI (2004) 123–145 MR Zbl

[7] **M G Durham**, **F Fanoni**, **N G Vlamis**, *Graphs of curves on infinite-type surfaces with mapping class group actions*, Ann. Inst. Fourier (Grenoble) 68 (2018) 2581–2612 MR Zbl

[8] **F Fanoni**, **S Hensel**, **N G Vlamis**, *Big mapping class groups acting on homology*, Indiana Univ. Math. J. 70 (2021) 2261–2294 MR Zbl

[9] **W P Hooper**, *Grid graphs and lattice surfaces*, Int. Math. Res. Not. 2013 (2013) 2657–2698 MR Zbl

[10] **J Ketonen**, *The structure of countable Boolean algebras*, Ann. of Math. 108 (1978) 41–89 MR Zbl

[11] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves, I: Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR Zbl

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

[13] **S Mazurkiewicz**, **W Sierpiński**, *Contribution à la topologie des ensembles dénombrables*, Fund. Math. 1 (1920) 17–27 Zbl

[14] **P Patel**, **N G Vlamis**, *Algebraic and topological properties of big mapping class groups*, Algebr. Geom. Topol. 18 (2018) 4109–4142 MR Zbl

[15] **A Randecker**, *Wild translation surfaces and infinite genus*, Algebr. Geom. Topol. 18 (2018) 2661–2699 MR Zbl

[16] **I Richards**, *On the classification of noncompact surfaces*, Trans. Amer. Math. Soc. 106 (1963) 259–269 MR Zbl

[17] **C Rosendal**, *Global and local boundedness of Polish groups*, Indiana Univ. Math. J. 62 (2013) 1621–1678 MR Zbl

[18] **C Rosendal**, *Coarse geometry of topological groups*, Cambridge Tracts in Mathematics 223, Cambridge Univ. Press (2021) MR Zbl

[19] **S Schleimer**, *Notes on the complex of curves*, lecture notes (2020) Available at `http://homepages.warwick.ac.uk/~masgar/Maths/notes2.pdf`

Department of Mathematics, Cornell University
Ithaca, NY, United States

Department of Mathematics, University of Toronto
Toronto, ON, Canada

`k.mann@cornell.edu`, `rafi@math.toronto.edu`

Proposed: Mladen Bestvina Received: 24 April 2020
Seconded: David Gabai, David Fisher Revised: 29 September 2021

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2023 is US $740/year for the electronic version, and $1030/year (+$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Geometry & Topology

Volume 27 Issue 6 (pages 2049–2496) 2023