Core Surfaces

Michael Magee and Doron Puder

June 20, 2022

Abstract

Let Γ_g be the fundamental group of a closed connected orientable surface of genus $g \geq 2$. We introduce a combinatorial structure of *core surfaces*, that represent subgroups of Γ_g. These structures are (usually) 2-dimensional complexes, made up of vertices, labeled oriented edges, and $4g$-gons. They are compact whenever the corresponding subgroup is finitely generated. The theory of core surfaces that we initiate here is analogous to the influential and fruitful theory of Stallings core graphs for subgroups of free groups.

Contents

1 Introduction

In his influential paper [Sta83], Stallings introduced the simple yet powerful concept of *core graphs*, sometimes known today under the name *Stallings core graphs*. Roughly, core graphs are connected, directed and edge-labeled graphs in one-to-one correspondence with the (conjugacy classes of) subgroups of a given f.g. (finitely generated) free group. We give the exact definition in Section 2 below. Core graphs are especially useful when the corresponding subgroup is f.g., or, equivalently, when the core graph is finite.

Inter alia, core graphs can be used to extract basic information about the subgroup (index, rank) (for these and some of the applications below consult [Sta83] and the surveys [KM02, DV22]). They provide simple proofs to classical theorems, such as Howson's theorem that the intersection of two f.g. subgroups is f.g., Hall's theorem that every f.g. subgroup is a free factor in a finite index subgroup, or Takahasi theorem that given a f.g. subgroup H of the free group $\mathbf{F}$, every supergroup $H \leq J \leq \mathbf{F}$ is a free extension of one of finitely many supergroups of H, to name a few. Core graphs also give rise to algorithms for various natural problems: for instance, determine the subgroup generated by a given set of words and the membership of other words in it, or determine whether a given word is primitive (a basis element) in a given subgroup. Finally, core graphs take part in the proofs of more involved results such as in [PP15].

In the current paper we wish to define an analogous notion, we call *core surfaces*, when a free group and its subgroups are replaced by a surface group and its subgroups. Here, a surface group is the fundamental group of Σ_g, a closed connected orientable surface of genus $g \geq 2$. We denote this group by Γ_g:

$$\Gamma_g \overset{\text{def}}{=} \pi_1(\Sigma_g) \cong \langle a_1, b_1, \ldots, a_g, b_g \,|\, [a_1, b_1] \cdots [a_g, b_g] \rangle . \tag{1.1}$$

In order to motivate our definition of a core surface, we first recall one of the (equivalent) definitions of a core graph. Let B_r be a bouquet consisting of a single vertex and r petals, namely, a wedge of r copies of S^1. Denote the wedge point by o. We identify the fundamental group $\pi_1(B_r, o)$ with $\mathbf{F}_r$, the free group of rank r. Given a subgroup $\{1\} \neq H \leq \mathbf{F}_r$, consider the connected covering space $p\colon \Upsilon \to B_r$ corresponding to the conjugacy class of[1] H. The core graph of H is then the subgraph of Υ which is the union of all non-backtracking cycles in Υ, together with the restriction of the covering map p. In other words, we remove from the covering space Υ all the "hanging trees", which do not affect its fundamental group. Equivalently, this is the unique smallest deformation retract of Υ. A key advantage of core graphs over the original covering spaces is that whenever H is $f.g.$ but not of finite index, the covering space of H is an infinite graph, while the core graph is a finite one.

Now let Γ_g be as in (1.1). Given a subgroup $J \leq \Gamma_g$, consider the covering space Υ of Σ_g corresponding to the conjugacy class of J. We would like to take a "topological core" of Υ. Naturally, when J is of finite index, the covering Υ is a closed compact surface and it makes sense to take it as the core surface of J. But what is the appropriate definition when J has infinite index in Γ_g?

In particular, consider the case where J is f.g. but of infinite index (in particular, J is a f.g. free group). Let us go through some possible definitions of a core surface which we do *not* find appealing.

Smallest retract Defining the core surface as a minimal retract would not work: Υ admits a minimal retract which is a finite graph, but this graph is far from canonical.

Geodesic boundary Another option is to "trim" pieces of Υ which, like hanging trees, do not affect the homotopy type of Υ. The best analog of hanging trees in Υ are "funnels": non-compact pieces that can be cut from Υ by a simple closed curve and are then homeomorphic to a once-punctured disc. Funnels make Υ non-compact even when J is f.g., and cutting outside a certain simple closed curve around every funnel leaves us with a compact retract of Υ. The question is, though, which curve should be used for that. One possibility is to give Σ_g a hyperbolic structure (so a Riemannian geometry with constant curvature -1). This geometry can be pulled back to the cover Υ. In hyperbolic surfaces, the homotopy class of every closed curve has a unique geodesic representative, and one can cut along the unique geodesic representing the simple closed curve around every funnel of Υ. Although this definition is natural and appealing, the construction is not combinatorial and therefore loses many of the flexibility we have with Stallings core graphs.

The definition we do give is a combinatorial construction which is close in spirit to the "geodesic boundary" definition. Consider the construction of Σ_g from a $4g$-gon by identifying its edges in pairs according to the pattern $a_1 b_1 a_1^{-1} b_1^{-1} \ldots a_g b_g a_g^{-1} b_g^{-1}$. This gives rise to a CW-structure on Σ_g consisting of one vertex (denoted o), $2g$ oriented $1-$cells (denoted $a_1, b_1, \ldots, a_g, b_g$) and one 2-cell which is the $4g$-gon glued along $4g$ 1-cells[2]. See Figure 1.1 (in our running examples with $g = 2$, we denote the generators of Γ_2 by a, b, c, d instead of a_1, b_1, a_2, b_2). We identify Γ_g with $\pi_1(\Sigma_g, o)$, so that in the presentation (1.1), words in the generators $a_1, \ldots, b_g$ correspond to the homotopy class of the corresponding closed paths based at o along the 1-skeleton of Σ_g. Note that every covering

[1] Namely, Υ is the unique connected covering space such that for some (and therefore every) vertex v of Υ, the image $p_*(\pi_1(\Upsilon, v))$ in $\pi_1(B_r, o)$ is conjugate to H.

[2] We use the terms vertices and edges interchangeably with 0-cells and 1-cells, respectively.

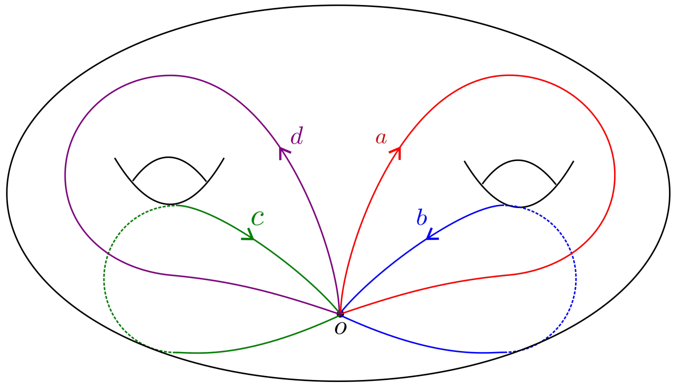

Figure 1.1: The CW-structure we give to the surface Σ_2 with fundamental group $\langle a,b,c,d \,|\, [a,b]\,[c,d] \rangle$: it consists of a single vertex (0-cell), four edges (1-cells) and one octagon (a 2-cell).

space $p\colon \Upsilon \to \Sigma_g$ inherits a CW-structure from Σ_g: the vertices are the pre-images of o, and the open 1-cells (2-cells) are the connected components of the pre-images of the open 1-cells (2-cell, respectively) in Σ_g.

Definition 1.1 (Core surface)**.** Given a subgroup $J \le \Gamma_g = \pi_1(\Sigma_g, o)$, consider the covering space $p\colon \Upsilon \to \Sigma_g$ corresponding to J. Define the **core surface of** J, denoted $\mathrm{Core}(J)$, as a sub-covering space of Υ as follows: (i) take the union of all shortest-representative cycles in the 1-skeleton $\Upsilon^{(1)}$ of every free-homotopy class of essential closed curve in Υ, and (ii) add every connected component of the complement which contains finitely many 2-cells.

For completeness define the core surface of the trivial subgroup to be the 0-dimensional complex consisting of a single vertex mapped to o.

We define $\mathrm{Core}(J)$ as a subcomplex of a covering space of Σ_g, but we usually think of it as an at-most 2-dimensional CW-complex with 1-cells that are directed and labeled by $a_1, b_1, \ldots, a_g, b_g$. These directions and labels on every 1-cell completely determine the restricted covering map. Three core surfaces are illustrated in Figure 1.2 and two others in Figure 4.2.

After having fixed the representation (1.1) for Γ_g, the core surfaces are unique for every conjugacy class of subgroups. Below, we give an intrinsic description of a core-surface which allows one to identify a core surface without knowledge of the full covering space it originates from (Proposition 5.9), we show how to construct the core surface of J from a set of generators using a “folding” process (Theorem 5.10), and prove a one-to-one correspondence between core surfaces and conjugacy classes of subgroups of Γ_g (Section 5.1). We also show some basic properties of core surface. For instance, we prove (Proposition 5.3) that $\mathrm{Core}(J)$ is connected and that it is a retract of the covering space Υ, and show (Proposition 5.8) that it is compact whenever J is a f.g. subgroup. We also prove (Lemma 5.4) that whenever $H \le J$ there is a natural morphism $\mathrm{Core}(H) \to \mathrm{Core}(J)$.

Random coverings of surfaces

We were led to the concept of core surfaces by our work on random homomorphisms from Γ_g to the symmetric group S_N [MP20], as part of a project on spectral gaps in random covering spaces of a fixed hyperbolic surface (see [MNP22]). Within this work we use core surfaces and the other concepts in the current paper to prove a theorem which is parallel to some extent to Takahasi's theorem for free groups.

Theorem 1.2. *[MP20] Let $J \le \Gamma_g$ be finitely generated. Then there is a finite list of subgroups $H_1, \ldots, H_r \le \Gamma_g$ with a fixed pointed sub-surface (Y_i, y_i) which is a deformation retract of the*

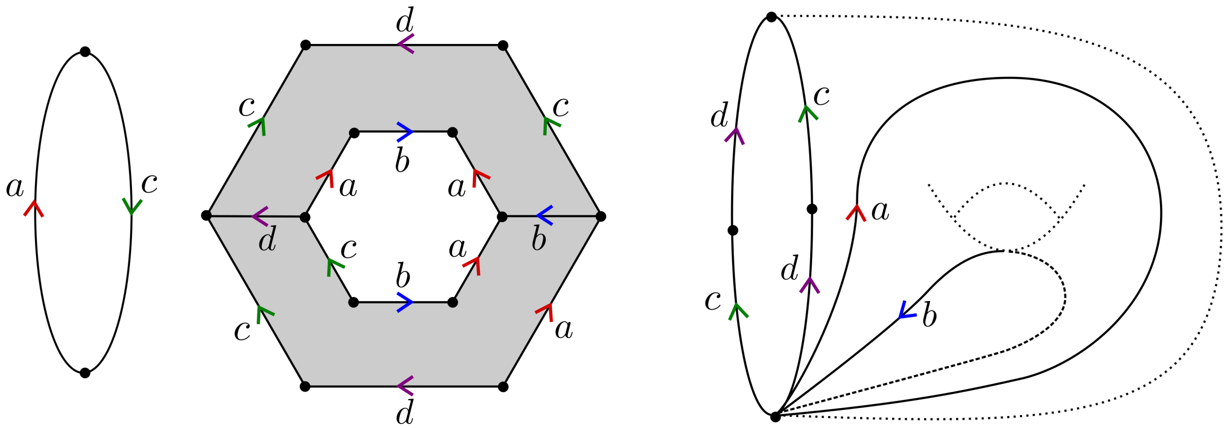

Figure 1.2: Consider $\Gamma_2 = \langle a, b, c, d \,|\, [a, b]\, [c, d] \rangle$. On the left is the core surface Core $(\langle ac \rangle)$. It consists of two vertices, two edges and no 2-cells. The middle object is Core $\left(\langle aba^{-2}b^{-1}c \rangle\right)$, consisting of 12 vertices, 14 edges and two octagons. Topologically it is an annulus. On the right is the core surface Core $(\langle a, b \rangle)$. It consists of four vertices, six edges and one octagon, and topologically it is a genus-1 torus with one boundary component.

covering space of Σ_g corresponding to H_i, so that (i) $J \le H_i$ for all i, and (ii) for every subgroup $J \le L \le \Gamma_g$, there is exactly one $1 \le i \le r$ such that (Y_i, y_i) is embedded in the pointed covering of Σ_g corresponding to L. Moreover, this embedding is π_1-injective.

This theorem is essentially Theorem 2.14 and Proposition 2.15 in [MP20], but see also some clarifications in [PZ22, Section 3.3]. The main goal of [MP20] is to study the average number of elements in $\{1, \ldots, N\}$ which are fixed by all permutations in $\theta(J)$ when θ is a random homomorphism $\theta \colon \Gamma_g \to S_N$. Theorem 1.2 is used in [MP20] to find the asymptotics of this number as $N \to \infty$. The main ingredient here, which hints to how Theorem 1.2 is used, is to show that if $J \le \Gamma_g$ is f.g., then the expected number of embeddings of Core (J) into a random N-sheeted covering space of Σ_g is $N^{\chi(J)} \left(1 + O\left(N^{-1}\right)\right)$.

Remark 1.3. Some of the content of this paper first appeared as part of the first version of [MP20]. Later on we decided to split that paper into two in order to make it shorter and as we believe the current paper is interesting for its own sake and to a potentially different audience. We also significantly expanded the content of the current paper.

Remark 1.4. There have been different successful attempts at generalizing the concepts of Stallings core graphs and of Stallings folds. A non-exhaustive list includes [AO96, Arz98, KMW17, Bro17, BL18, DL21, BZKL22] (see [DV22, Section 1] for a more exhaustive one). Some of these works include surface groups as special cases: [KMW17] introduces Stallings graphs for a large family of groups which includes surface groups using certain types of languages of representatives of the elements of the groups; and [BL18, DL21, BZKL22] introduce Stallings-like techniques for fundamental groups of CAT(0) cube complexes, a family that includes, again, hyperbolic surface groups. Moreover, the folding process we suggest in Section 5.2 has the same rough structure as the one introduced in [DL21, BZKL22]. However, because our generalization of Stallings core graph is much more specific and hands-on, we believe that our approach is more natural for surface groups. In particular, as mentioned above, our core surfaces appear quite naturally in the study of random coverings of surfaces.

Notation

We denote by $Y^{(1)}$ the 1-skeleton of a CW-complex Y. We let g^G denote the conjugacy class of the element g in a group G. Below, $\mathbb{Y}$ will sometimes denote the thick version of a tiled surface, ∂Y

its boundary and $|\partial Y|$ the length of ∂Y (Definition 3.2). The notation Y_+ refers to a tiled surface with hanging half-edges (see page 9). The universal cover of the surface Σ_g is denoted $\widetilde{\Sigma_g}$, and for a path $\mathcal{P}$ or cycle $\mathcal{C}$ in $Y^{(1)}$, we let $\mathcal{P}^*$ and $\mathcal{C}^*$, respectively, denote their inverses.

Paper outline

In Section 2 we recall the notion of Stallings core graphs in a more precise manner than above. In the following two sections, before studying core surfaces per se, we introduce three more general types of combinatorial surfaces which are sub-complexes of covering spaces of Σ_g. The most general concept is that of *tiled surfaces* which is described in Section 3. This is followed by the more restricted classes of boundary reduced and strongly boundary reduced tiled surfaces, which are defined and analyzed in Section 4. These three classes are natural in our context and are also important in our work on random covering surfaces. Finally, Section 5 returns to core surfaces and proves many of their properties.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 850956 and grant agreement No 949143). It was also supported by the Israel Science Foundation: ISF grant 1071/16.

2 Stallings Core graphs

Let $X = \{x_1, \ldots, x_r\}$ be a basis of the free group $\mathbf{F}_r$, and consider the bouquet B_r of r circles with distinct labels from X and arbitrary orientations and with a wedge point o. Then $\pi_1(B_r, o)$ is naturally identified with $\mathbf{F}_r$. An X-graph G is then a graph equipped with a graph morphism $G \to B_r$ which is an immersion, namely, it is locally injective. Equivalently, G is a directed graph with edges labeled by the elements of X, such that no vertex admits two outgoing edges with the same label nor two incoming edges with the same label. An X-graph G is an (X-labeled) core graph if it is connected and if every vertex belongs to some cyclically reduced cycle. If G is finite, the latter is equivalent to G being connected and having no leaves. (One usually also considers the isolated vertex graph to be a core graph.) Multiple edges between two vertices and loops at vertices are allowed.

There is a natural one-to-one correspondence between finite X-labeled core graphs and conjugacy classes of f.g. subgroups[3] of $\mathbf{F}_r$. Indeed, given a core graph G as above, pick an arbitrary vertex v and consider the "labeled fundamental group" $\pi_1^{\mathrm{lab}}(G, v)$: closed paths in a graph with oriented and X-labeled edges correspond to words in the elements of X. In other words, if $p\colon G \to B_r$ is the immersion, then $\pi_1^{\mathrm{lab}}(G, v)$ is the subgroup $p_*(\pi_1(G, v))$ of $\pi_1(B_r, o) = \mathbf{F}_r$. The conjugacy class of $\pi_1(G, v)$ is independent of the choice of v and is the conjugacy class corresponding to G.

Conversely, if $H \le \mathbf{F}_r$ is a non-trivial f.g. subgroup, the conjugacy class $H^{\mathbf{F}_r}$ corresponds to a finite core graph, denoted $G_X(H)$, which can be obtained in several equivalent manners. For example, let Υ be the topological covering space of B_r corresponding to $H^{\mathbf{F}_r}$, which is equal in this case to the Schreier graph depicting the action of $\mathbf{F}_r$ on the right cosets of H with respect to the generators X. Then $G_X(H)$ is obtained from Υ by 'pruning all hanging trees', or, equivalently, as the union of all non-backtracking cycles in Υ.[4] One can also construct $G_X(H)$ from any finite

[3] Sometimes core graphs are defined with a basepoint, which is allowed to be a leaf. Then, the correspondence is between core graphs and subgroups, not conjugacy classes of subgroups. We present here the non-based version because this is the version we think is more elegant in the definition of core surfaces.

[4] The only exception is when H is the trivial group, in which case Υ is a tree and we define Γ_B to consist of a single vertex and no edges.

generating set of H using "Stallings foldings". Finally, a core graph morphism is a graph morphism which commutes with the immersions to B_r. Given two subgroups $H, J \leq \mathbf{F}_r$, there is a core graph morphism $G_X(H) \to G_X(J)$ if and only if some conjugate of H is a subgroup of J. See [Sta83, KM02, Pud14, PP15, DV22] for more details about foldings and about core graphs in general. See also [HP22] for the category of not-necessarily-connected core graphs. As we show in the coming sections, many of the basic properties of core graphs hold for core surfaces as well.

3 Tiled surfaces

3.1 Basic definitions

We first define an object called a "tiled surface" which is the analog of an "X-graph" from Section 2. Recall that Σ_g is equipped with a CW-structure with one vertex, $2g$ edges and a single $4g$-gon, and that this CW-structure can be pulled back to every covering space of Σ_g. A sub-complex of a CW-complex is a subspace consisting of cells such that if some cell is in the subcomplex, than so are the cells of smaller dimension at its boundary.

Definition 3.1 (Tiled surface)**.** A *(g-) tiled surface* Y is a sub-complex of a (not-necessarily-connected) covering space of Σ_g. In particular, a tiled surface is equipped with the restricted covering map $p\colon Y \to \Sigma_g$ which is an immersion.

Alternatively, instead of considering a tiled surface Y to be a complex equipped with a restricted covering map, one may consider Y to be a complex as above with directed and labeled edges: the directions and labels $(a_1, b_1, \ldots, a_g, b_g)$ are pulled back from Σ_g via p. These labels uniquely determine p as a combinatorial map between complexes. Figures 1.1, 1.2 and 4.2 feature examples of tiled surfaces.

Note that a tiled surface is neither necessarily compact nor necessarily connected. Also note that a tiled surface is not always a surface: it may contain, for example, vertices or edges with no 2-cells incident to them. However, as Y is a sub-complex of a covering space of Σ_g, namely, of a surface, any neighborhood of Y inside the covering is a surface, and it is sometimes beneficial to think of Y as such.

Definition 3.2 (Thick version of a tiled surface)**.** Given a tiled surface Y which is a subcomplex of the covering space Υ of Σ_g, consider a small, closed, regular neighborhood of Y in Υ. The resulting closed surface, possibly with boundary, is referred to as the *thick version of* Y, and is occasionally denoted $\mathbb{Y}$. We let ∂Y denote the boundary of the thick version of Y and $|\partial Y|$ denote the number of edges along ∂Y (so if an edge of Y does not border any $4g$-gon, it is counted twice).

In particular, $|\partial Y| = 2e - 4gf$ where e is the number of edges and f the number of $4g$-gons in Y. We stress that we do not think of Y as a sub-complex, but rather as a complex for its own sake, which happens to have the capacity to be realized as a subcomplex of a covering space of Σ_g. Namely, the underlying proper covering space in Definition 3.1 is not part of the data possessed by Y. Indeed, one can also give a direct combinatorial definition of a tiled surface. Moreover, the thick version of a tiled surface is, too, independent of the covering it is a subcomplex of. This is shown by the following claims.

Proposition 3.3 (An intrinsic definition of a tiled surface)**.** *The following definition is equivalent to Definition 3.1: Fix an integer $g \geq 2$. A (g-) tiled surface Y is an at-most two-dimensional* CW*-complex with an assignment of both a direction and a label in $\{a_1, b_1, \ldots, a_g, b_g\}$ to each edge, such that:*

- ***P1:*** *Every vertex of Y has at most one incoming ℓ-labeled edge and at most one outgoing ℓ-labeled edge, for each $\ell \in \{a_1, b_1, \ldots, a_g, b_g\}$.*

- ***P2:*** *Every path in the* 1*-skeleton* $Y^{(1)}$ *of* Y *reading a word in* $\left\{a_1^{\pm 1}, b_1^{\pm 1}, \ldots, a_g^{\pm 1}, b_g^{\pm 1}\right\}$ *which equals the identity in* Γ_g *must be closed*[5].

- ***P3:*** *Every 2-cell in* Y *is a* $4g$*-gon glued along a closed path reading the relation* $[a_1, b_1] \ldots [a_g, b_g]$ *of* Γ_g*, and every such closed path is the boundary of at most one such* $4g$*-gon.*

Note that property **P1** is equivalent to that $Y^{(1)}$ is an $\{a_1, \ldots, b_g\}$-graph, in the sense of Section 2.

Proof. It is straightforward that a tiled surface from Definition 3.1 satisfies the assumptions in the statement of the proposition. Conversely, let Y be a CW-complex as in the statement of the proposition. It is enough to show that every connected component of Y is a tiled surface à la Definition 3.1, so assume without loss of generality that Y is connected and non-empty. Define a map $p\colon Y \to \Sigma_g$ by sending every vertex of Y to o, every edge of Y to the corresponding edge of Σ_g (with the same label and corresponding direction), and every $4g$-gon of Y to the single $4g$-gon of Σ_g. This map is a well-defined map of CW-complexes.

Now pick an arbitrary vertex v in Y. Denote $J \stackrel{\text{def}}{=} p_*\left(\pi_1\left(Y, v\right)\right) \leq \pi_1\left(\Sigma_g, o\right) = \Gamma_g$. Let $q\colon \Upsilon \to \Sigma_g$ be the connected covering space corresponding to J, and let u be a vertex of Υ so that $q_*\left(\pi_1\left(\Upsilon, u\right)\right) = J$. By standard facts from the theory of covering spaces [Hat05, Propositions 1.33 and 1.34], as $p_*\left(\pi_1\left(Y, v\right)\right) \subseteq q_*\left(\pi_1\left(\Upsilon, u\right)\right)$, there is a unique lift $r\colon (Y, v) \to (\Upsilon, u)$ of p so that $q \circ r = p$.

$$\begin{tikzcd} & (\Upsilon, u) \arrow[d, "q"] \\ (Y, v) \arrow[ur, dashed, "\exists ! r"] \arrow[r, "p"'] & (\Sigma_g, o) \end{tikzcd} \tag{3.1}$$

Clearly, r respects the CW-structure of Y and Υ as well as the labels and directions of edges. It remains to show that r is injective, for then Y is indeed a subcomplex of Υ and therefore a tiled surface.

So assume towards contradiction that r is not injective. By properties **P1** and **P3**, there must be two distinct vertices of Y with the same r-image. Without loss of generality we may assume that one of them is v (otherwise replace v with this vertex and replace u accordingly: this does not change q nor r). Call the other vertex v'. But then, any path in $Y^{(1)}$ from v to v' is mapped to a closed path at u and therefore reads a word w_1 in $\left\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\right\}$ representing some element $j \in J$. As $p_*\left(\pi_1\left(Y, v\right)\right) = J$, there is a closed path at v reading w_2 with w_2 also representing the same j. Then there exists a path from v to v' reading $w_2^{-1} w_1$ which is the identity in Γ_g, a contradiction to Property **P2**. $\square$

The combinatorial definition of a tiled surface given in Proposition 3.3 is not very satisfactory as it involves a "global" condition – Property **P2** – which is not easy to check, at least not complexity-wise[6]. In Section 4 we present a restricted version of tiled surfaces – boundary reduced tiled surfaces – where **P2** can be replaced with a local property. Before that, we show that the thick version of a tiled surface can also be extracted from the combinatorial, covering-space-free, definition.

Proposition 3.4 (An intrinsic definition of the thick version of a tiled surface)**.** *The thick version from Definition 3.2 of a tiled surface* Y *can be extracted from the combinatorial data of* Y*, namely, from the* CW*-structure with directions and labels of edges as in Proposition 3.3. This can be done as follows:*

[5] Here, we read a_i if we traverse an a_i-labeled edge in its correct direction, and a_i^{-1} if we traverse an a_i-labeled edge against its direction, and so on.

[6] Although if Y is a finite complex, it is possible to check condition **P2** by the classical solution of Dehn to the word problem in Γ_g [Deh12]. See also Theorem 4.8.

- *Around every vertex, give the half-edges incident to the vertex the cyclic order induced from the cyclic order of half-edges around the vertex o in Σ_g. Namely, this is the cyclic order which is (say, clockwise) a cyclic subsequence of a_1-outgoing, b_1-incoming, a_1-incoming, b_1-outgoing, a_2-outgoing, b_2-incoming, a_2-incoming, b_2-outgoing and so on.*

- *The cyclic ordering of half-edges at each vertex makes $Y^{(1)}$ into a* ribbon graph *that yields an orientable surface with boundary*[7].

- *Every $4g$-gon of Y corresponds to some boundary component of the ribbon graph (which reads the relation $[a_1, b_1] \cdots [a_g, b_g]$), and the thick version of Y is then obtained from the ribbon graph by attaching a $4g$-gon along every boundary component corresponding to a 2-cell of Y.*

Proof. If a tiled surface Y is defined via Definition 3.1 as a sub-complex of a covering space Υ of Σ_g, then every vertex inherits a cyclic ordering of the incident half-edges from the ordering in Υ, which is always the ordering specified in the statement of the proposition. The thick version of $Y^{(1)}$ is therefore precisely the ribbon graph described in the statement of the proposition, and the $4g$-gons are attached as described there too. Finally, because the description in the statement of the proposition is well defined and therefore unique, it must indeed recover the thick version from Definition 3.2. ☐

Example 3.5 (Universal cover of Σ_g)**.** Let $\widetilde{\Sigma_g}$ denote the universal cover of Σ_g endowed with the CW-structure pulled back from Σ_g via the covering map. As a topological space, $\widetilde{\Sigma_g}$ is an open disc. There is a natural action of Γ_g on $\widetilde{\Sigma_g}$ by isomorphisms of tiled surfaces such that $\Gamma_g \backslash \widetilde{\Sigma_g} = \Sigma_g$. We fix, once and for all, an arbitrary vertex u in $\widetilde{\Sigma_g}$, to obtain a pointed tiled surface $(\widetilde{\Sigma_g}, u)$. Note that $\left(\widetilde{\Sigma_g}, u\right)$ is the Cayley complex of Γ_g and its 1-skeleton is the Cayley graph of Γ_g with respect to the generators $a_1, \ldots, b_g$. For every $J \leq \Gamma_g$, the covering space of Σ_g corresponding to J can be also defined as $J \backslash \widetilde{\Sigma_g}$.

Morphisms of tiled surfaces

If Y_1 and Y_2 are tiled surfaces, a *morphism* from Y_1 to Y_2 is a map of CW-complexes which maps i-cells to i-cells for $i = 0, 1, 2$ and respects the directions and labels of edges. Equivalently, this is a morphism of CW-complexes which commutes with the restricted covering maps $p_j \colon Y_j \to \Sigma_g$ ($j = 1, 2$). In particular, the restricted covering map from a tiled surface to Σ_g is itself a morphism of tiled surfaces. It is an easy observation that every morphism of tiled surfaces is an immersion (locally injective).

3.2 Boundary cycles, hanging half-edges, blocks, and chains

In the current Section 3.2 we define some notions related to tiled surfaces which will play an important role in the coming sections, where we define (strongly) boundary reduced tiled surfaces and analyze the properties of core surfaces.

Given a tiled surface Y, a ***path*** in Y is a sequence $\mathcal{P} = (\vec{e_1}, \ldots, \vec{e}_k)$ of directed edges $\vec{e_1}, \ldots, \vec{e}_k$ in $Y^{(1)}$, where for each $1 \leq i \leq k-1$ the terminal vertex of $\vec{e_i}$ is the initial vertex of $\vec{e}_{i+1}$. The direction of the edges in the cycle is not necessarily the same as the direction dictated in the definition of the tiled surface. A ***cycle*** in Y is a path where, in addition, the terminal vertex of $\vec{e}_k$ is the initial vertex of $\vec{e}_1$. We write $\mathcal{P}^*$ or $\mathcal{C}^*$ for the oppositely oriented path or cycle, respectively, $(\vec{e}_k^*, \ldots, \vec{e}_1^*)$ where $\vec{e}_i^*$ is $\vec{e}_i$ with the opposite direction.

[7] A ribbon graph is also known as a fat-graph. The surface is obtained by thickening every vertex to a disc and every edge to a strip.

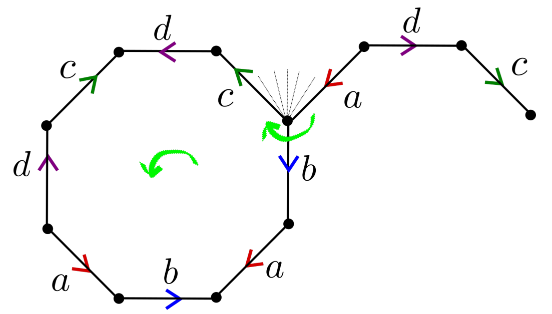

Figure 3.1: Fix $g = 2$. The figure shows a piece of a tiled surface containing one octagon and three more edges outside it. Blocks always follow the boundary of an octagon along the preset orientation (counter-clockwise in our figures), while half-edges around a vertex are ordered clockwise. For instance, any path reading $c^{-1}d^{-1}a$ is a block. In the figure there is a block $c^{-1}d^{-1}a$ on the right and a block $cdc^{-1}d^{-1}a$ on the left (the latter one can be extended on both ends). These two blocks are consecutive: there is exactly one half-edge (outgoing b) between the last edge (incoming a) of the first block and the first edge (outgoing c) of the second block. Therefore, their concatenation forms a chain $c^{-1}d^{-1}acdc^{-1}d^{-1}a$.

Every cycle $\mathcal{C}$ yields a cyclic word $w(\mathcal{C})$ in $\left\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\right\}$ by reading the label $\ell_i \in \{a_1, \ldots, b_g\}$ of the edge $\vec{e}_i$, $1 \le i \le k$ in order and writing (from left to right) ℓ_i if the direction of $\vec{e_i}$ is the same as the given direction in the tiled surface Y, and ℓ_i^{-1} otherwise. Every word w in $\left\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\right\}$ corresponds to an element $\gamma_w \in \Gamma_g$ via the presentation (1.1). We write γ^{Γ_g} for the conjugacy class of γ in Γ_g. We say that w represents the conjugacy class $\gamma_w{}^{\Gamma_g}$.

We denote by $\ell(w)$ the length of a word w, and if $\gamma \in \Gamma_g$, we write $\ell(\gamma^{\Gamma_g})$ for the minimal length of a word w for which $\gamma_w{}^{\Gamma_g} = \gamma^{\Gamma_g}$. For $\gamma \in \Gamma_g$, we say that γ is ***cyclically shortest*** if $\gamma = \gamma_w$ for some word w with $\ell(w) = \ell(\gamma^{\Gamma_g})$. We say that a word w is ***cyclically shortest*** if $\ell(w) = \ell(\gamma_w{}^{\Gamma_g})$.

In the rest of the paper, we wish to use some notions of Birman and Series from [BS87]. However, we make one small adjustment: what Birman and Series call a cycle, we will call a block[8].

We wish to augment the tiled surface Y by adding some new half-edges. Here, formally, a half-edge is a copy of the interval $[0, \frac{1}{2})$. Every edge of Y gives rise to two half-edges which cover the entire edge except for the midpoint. We may also add a new half-edge to a vertex, in which case the point 0 will be identified with the vertex. For every vertex v of Y, we add at most 8 half-edges to v to form a new surface Y_+. The half-edges are added in such a way that the morphism $p\colon Y \to \Sigma_g$ extends to a map $p_+ : Y_+ \to \Sigma_g$ that gives a local homeomorphism of 1-skeleta at each vertex of Y_+. We call a half-edge of Y_+ a ***hanging half-edge*** if it is added in this way: the remaining half-edges in Y_+ are contained in proper edges of Y.

As a result, there are now exactly $4g$ half-edges incident to every vertex in Y_+. The hanging half-edges of Y_+ inherit both a label in $\{a_1, \ldots, b_g\}$ and a direction from the corresponding half-edges in Σ_g. Moreover, at each vertex of Y_+, the incident half-edges have a cyclic ordering given by the (clockwise) cyclic ordering of the half-edges in Σ_g. We fix these labels, directions, and the cyclic ordering of half-edges at each vertex as part of the data of Y_+.

Given two directed edges $\vec{e}_1$ and $\vec{e}_2$ of Y, with the terminal vertex v of $\vec{e}_1$ equal to the source vertex of $\vec{e}_2$, we refer to the m half-edges of Y_+ incident to v between $\vec{e}_1$ and $\vec{e}_2$ in the given cyclic order of the $4g$ half-edges at v as *the **half-edges between** $\vec{e}_1$ and $\vec{e}_2$*. Here $0 \le m \le 4g - 1$.

A path in a tiled surface Y is a ***block*** if it is non-empty and each pair of successive edges have

[8]This is so that we can reserve the term cycle to be used in the usual way as we have above.

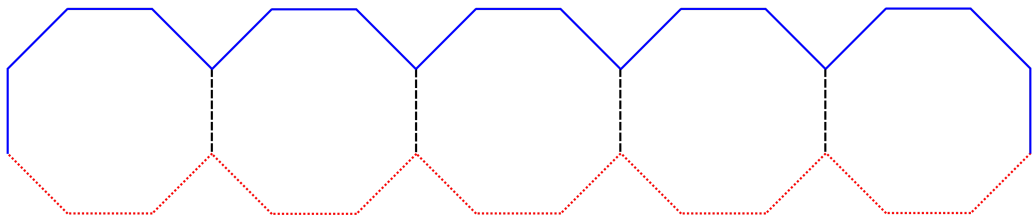

Figure 3.2: Fix $g = 2$. The figure shows a long chain of total length 17 (blocks of sizes $4, 3, 3, 3, 4$, in blue) and its complement of length 15 which is the inverse of a chain made of blocks of lengths $3, 3, 3, 3, 3$ (in red).

no half-edges between them. A block runs along the boundary of a single $4g$-gon in Y or a single "phantom-$4g$-gon" which exists in the covering space of Σ_g of which Y is a sub-complex. In other words, a block is a path that reads a subword of the cyclic word $[a_1, b_1] \dots [a_g, b_g]$. Note that the inverse of a block of length ≥ 2 is *not* a block.

A ***half-block*** is a block of length $2g$, and a ***long block*** is a block of length at least $2g + 1$, including the case of a "full block" of length $4g$. If a (non-cyclic) block of length b sits along the boundary of a $4g$-gon O, the ***complement* of the block** is the inverse of the block of length $4g - b$ consisting of the complement set of edges along O (so the block and its complement share the same starting point and the same terminal point).

We say that two blocks $(\vec{e}_i, \dots, \vec{e}_j)$ and $(\vec{e}_k, \dots, \vec{e}_\ell)$ are ***consecutive*** if $(\vec{e}_i, \dots, \vec{e}_j, \vec{e}_k, \dots, \vec{e}_\ell)$ is a path and there is exactly one half-edge between $\vec{e}_j$ and $\vec{e}_k$. A ***chain*** is a (possibly cyclic) sequence of consecutive blocks. This is illustrated in Figure 3.1. A ***cyclic chain*** is a chain whose blocks pave an entire cycle. A ***long chain*** is a chain consisting of consecutive blocks of lengths

$$2g, 2g - 1, 2g - 1, \dots, 2g - 1, 2g.$$

A ***half-chain***[9] is a cyclic chain consisting of consecutive blocks of length $2g - 1$ each.

The **complement of a long chain** is the inverse of a chain with blocks of lengths $2g - 1, 2g - 1, \dots, 2g - 1$ which sits along the other side of the $4g$-gons bordering the long chain. Note that the complement of a long chain shares the same starting point and terminal point as the long chain, and is shorter by two edges from the long chain. See Figure 3.2.

The **complement of a half-chain** is defined as follows. If the half-chain sits along the boundary of the $4g$-gons $O_1, \dots, O_r$, its complement is the inverse of the half-chain sitting along the other sides of these $4g$-gons: a block (of length $2g - 1$) of the half-chain along O_i is replaced by the path of length $2g - 1$ along O_i, with starting and terminal points one edge away from the starting and terminal points, respectively, of the block. The complement of a half-chain has the same length as the original half-chain. The middle part of Figure 1.2 illustrates two complementing half-chains of length 6 each (with two octagons in between).

The following lemma shows, in particular, that there are no cyclic chains consisting of one block of length $2g$ nor cyclic chains consisting of blocks of lengths $2g, 2g - 1, 2g - 1, \dots, 2g - 1$.

Lemma 3.6. *In every cyclic chain, the number of blocks of even length is even.*

(This excludes the special case of a cyclic chain consisting of a single block of length $4g$, which may be excluded from the definition of a chain.)

Proof. In the defining relation $[a_1, b_1] \dots, [a_g, b_g]$, every letter is at distance 2 from its inverse. In two consecutive blocks $(\vec{e}_i, \dots, \vec{e}_j)$ and $(\vec{e}_k, \dots, \vec{e}_\ell)$, there are thus two possible cases. Either the letter

[9] This notion "half-chain" is ours – it does not appear in [BS87]. While this name does not capture the essence of these objects, we chose it because half-chains are related to half-blocks in roughly the same manner as long chains are related to long blocks. This will be apparent in Section 4.

associated with $\vec{e}_k$ is identical to the letter associated with $\vec{e}_j$, as in $\left(a_1, b_1, a_1^{-1}\right), \left(a_1^{-1}, b_1^{-1}, a_2, b_2\right)$ with an incoming half b_1-edge hanging in between. Or $\vec{e}_k$ comes four places after $\vec{e}_j$ in the defining relation, as in $\left(b_g^{-1}, a_1\right), \left(a_2, b_2, a_2^{-1}, b_2^{-1}\right)$ with an outgoing half b_1-edge hanging in between. Hence the parity of the location in the defining relation of the first letter in a block alternates after an even-length block. As the defining relation has even length, this proves the lemma. $\square$

For every directed edge $\vec{e}$ in a tiled surface Y and every $4g$-gon O in Y that meets $\vec{e}$ at its boundary, we say that O is on the left (resp. right) of $\vec{e}$ if for a small neighborhood N of $\vec{e}$ in O, N is on the left (resp. right) of $\vec{e}$ as $\vec{e}$ is traversed in its given direction, where left/right is defined with respect to an orientation inherited from a fixed orientation of Σ_g. Note that a $4g$-gon O can be both on the left and right of a directed edge if that edge appears twice in the boundary of O: this is the case, for instance, with the a- and b-edges in the right core surface in Figure 1.2.

A ***boundary cycle*** of Y is a cycle $(\vec{e_1}, \ldots, \vec{e_k})$ in Y corresponding to an oriented boundary component of the thick version of Y (see Definition 3.2). We always choose the orientation so that there are no $4g$-gons to the immediate **left** of the boundary component as it is traversed. Therefore, boundary components of Y correspond to unique cycles. Note that $|\partial Y|$ is equal to the sum over boundary cycles of Y of the number of edges in each such cycle.

4 Boundary reduced and strongly boundary reduced tiled surfaces

In this section we describe a restricted class of tiled surfaces called "boundary reduced" and its subclass of "strongly boundary reduced" tiled surfaces. These are tiled surfaces with "nice" boundary, which turns the global property **P2** from Proposition 3.3 into a simpler one. As we show in the next section, this class also contains all core surfaces, and as such the properties of its elements are important for our main object of study. Moreover, these notions are also important for the analysis in [MP20]: for a compact, (strongly) boundary reduced tiled surface we are able to give a rather precise estimate for the expected number of times it is embedded in a random N-sheeted covering of Σ_g [MP20, Propositions 5.26 and 5.26].

Definition 4.1 (Boundary reduced)**.** A tiled surface Y is *boundary reduced* if no boundary cycle of Y contains a long block or a long chain.

In particular, if Y is boundary reduced, then every path that reads $[a_1, b_1] \ldots [a_g, b_g]$ is not only closed, but there is also a $4g$-gon attached to it. We also need a stronger version of this property.

Definition 4.2 (Strongly boundary reduced)**.** A tiled surface Y is *strongly boundary reduced* if no boundary cycle of Y contains a half-block or a half-chain.

Because a long block contains (at least two) half-blocks and a long chain contains (two) half-blocks, a strongly boundary reduced tiled surface is in particular boundary reduced.

We now show that in the combinatorial definition of a tiled surface (Proposition 3.3), if the boundary is reduced, then property **P2** holds automatically. Note that the combinatorial definition of the thick version of a tiled surface, as in Proposition 3.4, does not depend on the validity of property **P2**, namely, the thick version is well-defined even when the complex Y satisfies the assumptions of Proposition 3.3 excluding **P2**.

Proposition 4.3. *Let Y be an at-most two-dimensional* CW*-complex with an assignment of both a direction and a label in $\{a_1, \ldots, b_g\}$ to each edge, such that properties **P1** and **P3** from Proposition 3.3 hold. It the thick version of Y (as in Proposition 3.4) is boundary reduced, then Y is a tiled surface.*

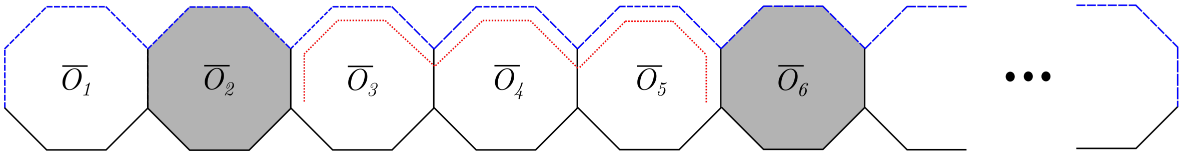

Figure 4.1: Assume $g = 2$, and let Z be a boundary reduced tiled surface containing a long chain c (broken blue line) bordering the octagons $O_1, \ldots, O_r$ (not necessarily distinct). If Y is sub-tiled surface of Z containing c but not all the octagons $O_1, \ldots, O_r$, then Y is not boundary reduced. For example, if $O_3, O_4, O_5 \notin Y$ is a longest subsequence of octagons not in Y (so $O_2, O_6 \in Y$), then there is a long chain (dotted red line) at ∂Y along these three octagons.

Proof. It is enough to prove that Y satisfies **P2**. Assume toward contradiction that there is a non-closed path in $Y^{(1)}$ reading a word which equals the identity in Γ_g. Let $\mathcal{P}$ be such a path of minimal length. By the classical results of Dehn [Deh12], $w(\mathcal{P})$ must then contain a subword which is a cyclic piece of the relation $[a_1, b_1] \ldots [a_g, b_g]$ or its inverse of length more than half (so at least $2g + 1$) (this also follows, for example, from Greendlinger's lemma – see [LS77, Theorem V.4.5]). But this corresponds to a long block in $\mathcal{P}$ or in $\mathcal{P}^*$, and as Y is boundary reduced, the complement of this long block is also in Y. By replacing the long block with its complement, we obtain a strictly shorter path reading a word which equals the identity – a contradiction. □

Definition 4.4 (BR- and SBR-closure)**.** Let Y be a tiled surface embedded in a boundary reduced tiled surface Z. The *boundary reduced closure* of Y in Z, or BR-closure, denoted $\mathsf{BR}(Y \hookrightarrow Z)$, is the intersection of all intermediate tiled surfaces $Y \hookrightarrow Y' \hookrightarrow Z$ which are boundary reduced.

Likewise, if Z is strongly boundary reduced, the *strongly boundary reduced closure* of Y in Z, or SBR-closure, denoted $\mathsf{SBR}(Y \hookrightarrow Z)$, is the intersection of all intermediate strongly boundary reduced tiled surfaces $Y \hookrightarrow Y' \hookrightarrow Z$.

Proposition 4.5. *The* BR*-closure of Y is boundary reduced and contains Y. The* SBR*-closure of Y is strongly boundary reduced and contains Y.*

Proof. By assumption, Z itself is (strongly) boundary reduced, and so the intersection is over a non-empty set of tiled surfaces. It trivially contains Y. We claim that the intersection of every family of boundary reduced tiled sub-surfaces of Z is boundary reduced. Indeed, if the intersection $X = \cap Y'$ has some long block b at its boundary ∂X, then b is also a long block at the boundary of some Y', which is impossible. If c is a long chain at ∂X and O some $4g$-gon of $Z \setminus X$ sitting along c, then there is some Y' not containing O but containing c. The intersection $\partial Y' \cap \partial O$ must then either include a long block or a block which belongs to a long chain of $\partial Y'$ (see Figure 4.1), which is impossible.

Similarly, we claim that the intersection $X = \cap Y'$ of every family of strongly boundary reduced tiled sub-surfaces of Z is strongly boundary reduced. If b is a half-block at ∂X, it must also belong to $\partial Y'$ for some Y', which is impossible. If ∂X contains a half-chain c, then every $4g$-gon O of $Z \setminus X$ sitting along c does not belong to some Y'. But then O sits along a half block, a long chain or the same half-chain along $\partial Y'$, which is impossible. □

A useful property we now prove is that the boundary-reduced closure $\mathsf{BR}(Y \hookrightarrow Z)$ of a compact tiled surface Y is compact too. The result in Proposition 4.6 is not true for SBR-closure. This is illustrated in Figure 4.2.

Proposition 4.6. *Let Y be a compact tiled surface embedded in a boundary reduced tiled surface Z. Then $\mathsf{BR}(Y \hookrightarrow Z)$ is compact too.*

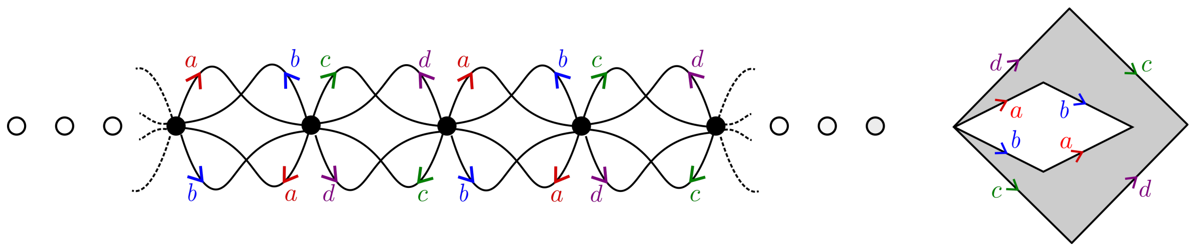

Figure 4.2: Let $g = 2$ and $\Gamma_2 = \langle a, b, c, d \,|\, [a, b]\, [c, d] \rangle$. On the left hand side is a piece of an the infinite 1-skeleton of a tiled surface Z with no boundary. The full surface extends to both sides infinitely with the same fixed pattern, and every possible octagon is included in it. This is the covering space of Σ_2, as well as the core surface, corresponding to the normal subgroup $\ll a^2, ab, ab^{-1}, c^2, cd \gg \trianglelefteq \Gamma_2$. On the right is a tiled surface Y consisting of a single octagon with two of its vertices identified. This is the core surface Core $(\langle [a, b] \rangle)$. Because of symmetry, every morphism $Y \to Z$ looks the same. The image of Y in such a morphism is a tiled surface Y' consisting of three vertices, eight edges and one octagon. It is not hard to see that in this case, although Y' is compact, $\mathsf{SBR}\,(Y' \hookrightarrow Z)$ is the entire of Z and, in particular, not compact.

Proof. Our definition of the BR-closure is from the top down: by taking the intersection of boundary reduced tiled surfaces. But when Y is compact, one can instead construct its BR-closure in Z from the bottom up by adding closed $4g$-gons to Y. We describe this process and show that one needs to add only finitely many $4g$-gons, which will prove the claim.

Indeed, let $Y' = Y$. We add (closed) $4g$-gons to Y' while keeping it a sub-tiled surface of $\mathsf{BR}\,(Y \hookrightarrow Z)$. By (the proof of) Proposition 4.5, as $Y' \subseteq \mathsf{BR}\,(Y \hookrightarrow Z)$, if $\partial Y'$ contains a long block, then the $4g$-gon along it must too belong to $\mathsf{BR}\,(Y \hookrightarrow Z)$, and we may add this $4g$-gon to Y'. In doing so we removed at least $2g + 1$ edges from $\partial Y'$, and added to $\partial Y'$ at most $2g - 1$ new edges, thus reducing $|\partial Y'|$ by at least 2. Likewise, if $\partial Y'$ contains a long chain, then all $4g$-gons along it must belong to $\mathsf{BR}\,(Y \hookrightarrow Z)$ and we may add them to Y'. In this step $|\partial Y'|$ is reduced again by at least two. But $|\partial Y'|$ is always non-negative hence this process must end after finitely many steps. The resulting tiled surface Y' is boundary reduced and thus equals $\mathsf{BR}\,(Y \hookrightarrow Z)$. Because we added a finite number of $4g$-gons to a compact tiled surface, $\mathsf{BR}\,(Y \hookrightarrow Z)$ is compact too. □

We need the following lemma when we construct resolutions in [MP20, Section 2.3].

Lemma 4.7. *Let $f\colon Z_1 \to Z_2$ be a morphism between two strongly boundary reduced tiled surfaces and let Y be a sub-surface of Z_1. Then*

$$f\,(\mathsf{SBR}\,(Y \hookrightarrow Z_1)) \subseteq \mathsf{SBR}\,(f\,(Y) \hookrightarrow Z_2)\,.$$

Proof. Denote $W = \mathsf{SBR}\,(f\,(Y) \hookrightarrow Z_2)$. Let $X = f^{-1}\,(W) \cap \mathsf{SBR}\,(Y \hookrightarrow Z_1)$ be the sub-surface of $\mathsf{SBR}\,(Y \hookrightarrow Z_1)$ consisting of all cells which are mapped by f to W. Assume towards contradiction that X is a proper sub-surface of $\mathsf{SBR}\,(Y \hookrightarrow Z_1)$. This means that X is not strongly boundary reduced, and so ∂X contains a half-block or a half-chain b, and the $4g$-gons along b belong to $\mathsf{SBR}\,(Y \hookrightarrow Z_1)$. Let O be one such $4g$-gon. But then $f\,(b) \subseteq W$ while $f\,(O) \notin W$. As in the proof of Proposition 4.5, $f\,(O)$ lies along a half-block, a long chain or a half-chain in ∂W, which is a contradiction. □

To analyze the properties of (strongly) boundary reduced tiled surfaces, and later on of core surfaces, we need a result of Birman and Series that strengthens classical results of Dehn [Deh12]. This result deals with shortest representatives of conjugacy classes in surface groups. The paper [BS87] concerns a wide class of presentations of Fuchsian group, which includes, in particular, the presentations $\Gamma_g = \langle a_1, b_1, \ldots, a_g, b_g \,|\, [a_1, b_1] \cdots [a_g, b_g] \rangle$ for every $g \geq 2$ as in (1.1). We state here one of the main results for this case.

Let $J \leq \Gamma_g$ be a subgroup. Consider the covering space $\Upsilon = J \backslash \widetilde{\Sigma_g}$ corresponding to J (here $\widetilde{\Sigma_g}$ is the universal cover from Example 3.5). This Υ is a tiled surface without boundary that may be compact or not (depending on whether J has finite index in Γ_g or not). Conjugacy classes in J are in one-to-one correspondence with free homotopy classes of oriented closed curves in Υ, and each such class has representatives contained in $\Upsilon^{(1)}$.

In particular, for an arbitrary $1 \neq \gamma \in \Gamma_g$, consider the tiled surface $\Upsilon = \langle \gamma \rangle \backslash \widetilde{\Sigma_g}$. Topologically, this is a two punctured sphere. The conjugacy class of γ in Γ_g corresponds to the free-homotopy class of the essential simple closed curve[10] in Υ (with an appropriate orientation). The set of cyclically reduced cyclic words in $\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\}$ representing the conjugacy class of γ in Υ is identical to the set of cyclic words coming from non-backtracking cycles in $\Upsilon^{(1)}$ representing the same free-homotopy class of curves. Given a cycle $\mathcal{C}$ in $\Upsilon^{(1)}$, a "half-block switch" consists of identifying a half-block in $\mathcal{C}$ or in $\mathcal{C}^*$, and replacing it with the complement half-block (around the same $4g$-gon). A "half-chain switch" can take place if one of $\mathcal{C}$ or $\mathcal{C}^*$ is a half-chain, in which case it refers to replacing this half-chain with its complement. For example, in the middle part of Figure 1.2, there is a cycle $\mathcal{C}$ reading $aba^{-2}b^{-1}c$, which is a half-chain. Its complement reads $cd^{-1}c^{-1}a^{-1}dc$.

Theorem 4.8 (Birman-Series)**.** *[BS87, Thm. 2.12]*

1. *The cyclically reduced cyclic word w in $\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\}$ is a shortest representative of the conjugacy class in Γ_g it represents if and only if the corresponding bi-infinite periodic path $\mathcal{C}$ in $\widetilde{\Sigma_g}^{(1)}$ and its inverse $\mathcal{C}^*$ do not contain any long block or long chain.*

2. *Assume that the cyclic words w_1, w_2 are both shortest representatives of the conjugacy class γ^{Γ_g} for some $1 \neq \gamma \in \Gamma_g$. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be corresponding cycles in $\Upsilon = \langle \gamma \rangle \backslash \widetilde{\Sigma_g}$. Then either $\mathcal{C}_2$ can be obtained from $\mathcal{C}_1$ by a finite number of half-block switches, or $\mathcal{C}_2$ can be obtained from $\mathcal{C}_1$ by a single half-chain switch.*

For example, the element $\gamma = aba^{-2}b^{-1}c \in \Gamma_2$ has exactly two different cyclic words which are shortest representatives of its conjugacy class: the cyclic words $aba^{-2}b^{-1}c$ and $cd^{-1}c^{-1}a^{-1}dc$. These two words correspond to two disjoint cycles in the 1-skeleton of $\langle \gamma \rangle \backslash \widetilde{\Sigma_2}$, and the complement of their union consists of three components: two with infinitely many $4g$-gons, and one component, an annulus bounded by both cycles, containing exactly two octagons.

Corollary 4.9. *Let Y be a tiled surface and let $\mathcal{C}$ be a (non-backtracking) cycle in $Y^{(1)}$. If $\mathcal{C}$ and $\mathcal{C}^*$ do not contain any long block or long chain, then $\mathcal{C}$ is a shortest representative of its free-homotopy class in Y.*

Proof. If $\mathcal{C}$ and $\mathcal{C}^*$ contain no long block nor long chain, then, by Theorem 4.8, they are shortest representatives of $w(\mathcal{C})^{\Gamma_g}$, the conjugacy class of $w(\mathcal{C})$ in Γ_g. Every other cycle representing the same free-homotopy class in Y, also represents $w(\mathcal{C})^{\Gamma}$, so it cannot be shorter. □

If Y is boundary reduced, the converse also hold.

Corollary 4.10. *Let Y be a boundary reduced tiled surface, and let $\mathcal{C}$ be a (non-backtracking) cycle in $Y^{(1)}$. Then the following holds.*

1. *There is a shortest representative cycle in $Y^{(1)}$ for the free homotopy class of $\mathcal{C}$.*

2. *If $\mathcal{C}$ is a shortest representative of its free-homotopy class in Y then $\mathcal{C}$ and $\mathcal{C}^*$ do not contain any long block or long chain.*

[10] We call a closed curve in a surface *essential* if it is not null-homotopic.

Proof. As Y is boundary reduced, every long block contained in $Y^{(1)}$ lies at the boundary of a $4g$-gon contained in Y, and therefore so does the complement of this long block. Similarly, if $\mathcal{P}$ is a long chain contained in $Y^{(1)}$, then every sequence of $4g$-gons along $\mathcal{P}$ which are not in Y gives rise to a long block or a long chain along ∂Y, which is impossible. Hence all $4g$-gon along $\mathcal{P}$ belong to Y, and so the complement of $\mathcal{P}$ belongs to Y.

For an arbitrary $\mathcal{C} \subseteq Y^{(1)}$, we can greedily shorten it by replacing every long block or long chain in $\mathcal{C}$ or in $\mathcal{C}^*$ by their complement, and this leads to a shortest representative by Theorem 4.8. This proves (1). Now assume that $\mathcal{C} \subseteq Y^{(1)}$ is cyclically shortest. If $\mathcal{C}$ or $\mathcal{C}^*$ contains a long block or a long chain, then replacing this long block/chain with its complement, which is still in Y, reduces the length of $\mathcal{C}$. This proves (2). $\square$

Corollary 4.11. *Every morphism from a boundary reduced tiled surface is π_1-injective.*

Proof. Let $f\colon Y \to Z$ be a morphism of tiled surfaces where Y is boundary reduced, and let $\mathcal{C}$ be a cycle in Y which is not null-homotopic. By Corollary 4.10 there is a representative in $Y^{(1)}$ for the same free-homotopy class in Y, which contains no long blocks nor long chains. But then the f-image of this representative is also shortest, and in particular non-nullhomotopic, in Z. $\square$

5 Properties and construction of core surfaces

5.1 Properties of core surfaces

Recall Definition 1.1: the core surface $\mathrm{Core}\,(J)$ of a subgroup $J \le \Gamma_g$ is the sub-complex of $\Upsilon = J \backslash \widetilde{\Sigma_g}$ obtained as the union of all shortest representative cycles in $\Upsilon^{(1)}$ of non-trivial conjugacy classes of J, together with the connected components of the complement which contain finitely many $4g$-gons. In this section we prove some basic properties of this object. Among these, we show that a core surface is strongly boundary reduced, that it is compact whenever J is f.g., and that whenever $H \le J \le \Gamma_g$, the natural morphism $H\backslash\widetilde{\Sigma_g} \to J\backslash\widetilde{\Sigma_g}$, restricts to a map between the corresponding core surfaces.

We start with analyzing the special case of the core surface of a cyclic subgroup.

Lemma 5.1. *Let $1 \ne \gamma \in \Gamma_g$ be a non-trivial element. Then the core surface* $\mathrm{Core}\,(\langle\gamma\rangle)$ *is connected and compact with its thick version homeomorphic to an annulus. Furthermore, both boundary cycles of* $\mathrm{Core}\,(\langle\gamma\rangle)$ *are of length $\ell\left(\gamma^{\Gamma_g}\right)$.*

Proof. That $\mathrm{Core}\,(\langle\gamma\rangle)$ is connected follows immediately from Theorem 4.8. Let $\mathcal{C}$ be some (simple) cycle in the twice-punctured sphere $\Upsilon = \langle\gamma\rangle\backslash\widetilde{\Sigma_g}$ which is a shortest representative of γ^{Γ_g}. If $\mathcal{C}$ or $\mathcal{C}^*$ is a half-chain, denote by $\mathcal{C}'$ its complement. Then $\mathrm{Core}\,(\langle\gamma\rangle)$ is precisely the compact annulus made of $\mathcal{C}$, $\mathcal{C}'$ and the narrow annulus separating them. Indeed, none of $\mathcal{C}$, $\mathcal{C}'$ and their inverses contain a half-block, so by Theorem 4.8, $\mathcal{C}$ and $\mathcal{C}'$ are the only shortest representatives of γ^{Γ_g}.

Now assume neither $\mathcal{C}$ nor $\mathcal{C}^*$ is a half-chain. By Theorem 4.8, $\mathrm{Core}\,(\langle\gamma\rangle)$ may be obtained from the tiled sub-surface $\mathcal{C}$ of Υ by repeatedly annexing any $4g$-gon sitting along a half-block at the boundary. We only need to show this process must end after finitely many steps.

Let $Y_0 = \mathcal{C}, Y_1, Y_2, \ldots$ denote the sub-surfaces we construct in this process, so Y_{i+1} is obtained from Y_i by annexing a $4g$-gon of Υ bordering a half-block in ∂Y_i. For every i, as $\mathcal{C}$ is a shortest representative, the boundary component of Y_i around each of the two punctures of Υ is of length at least $|\mathcal{C}|$. Clearly, $|\partial Y_{i+1}| \le |\partial Y_i|$, so we get that every Y_i has exactly two boundary components, and each of them is of length $|\mathcal{C}|$. In particular, the $4g$-gon annexed to Y_i to obtain Y_{i+1} contains at its boundary a half block $\mathcal{B}$ which did not belong to Y_i, and such that $\mathcal{B}^*$ is a path in ∂Y_{i+1}.

Finally, denote by s the number of hanging half-edges at the boundary of $(Y_0)_+$. At every step, the boundary of Y_i has constant length $2\,|\mathcal{C}|$, while s increases by $(2g-1)\,(4g-2) - 2 = 8g\,(g-1)$ — see Figure 5.1. As the number of hanging half-edges in a compact tiled surface Y without

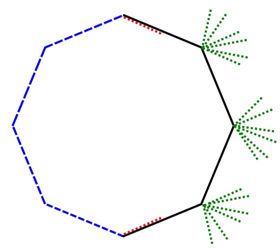

Figure 5.1: Let $g = 2$. Assume Y is a tiled sub-surface of a tiled surface Z with no boundary. Assume that Y has some half-block at its boundary, marked here in broken blue, and sitting along the octagon O. Denote by Y' the union of Y with the closure of O in Z, and assume that the inverse of the other half block along O is an interval along $\partial Y'$. Then, the number of hanging half-edges in $(Y')_+$ is larger by 16 than their number in Y_+: two hanging half-edges of Y_+ (marked in red) are no longer hanging in $(Y')_+$, but $18 = (2g-1)(4g-2)$ new hanging half-edges (marked in green) belong to $(Y')_+$.

isolated vertices nor leaves is at most $(4g-2)|\partial Y|$, this process must terminate after finitely many steps. □

As Core (J) is a closed sub-surface of $\Upsilon = J\backslash\widetilde{\Sigma_g}$, every component of its complement $\Upsilon\backslash$Core (J) is open. Hence every component is a surface with punctures. Each of these punctures corresponds either to a puncture of Υ, in which case we call it a *funnel*, or to a component of ∂Core (J), in which case we call it a *fake-puncture*. In particular, a funnel is of infinite distance (measured in paths of adjacent $4g$-gons, say) from any given $4g$-gon in $\Upsilon \setminus$ Core (J), while a fake-puncture has certain (open) $4g$-gons adjacent to it.

Lemma 5.2. *Let X be a connected tiled surface, which is a proper surface (so every edge is incident with one or two $4g$-gons, and every vertex is incident with a single sequence of m $4g$-gons, $1 \le m \le 4g$). Assume that X has genus g and a total of b boundary components and punctures. If $g \ge 1$ or $b \ge 3$, then $X^{(1)}$ contains a cycle, not contained in ∂X, which is cyclically shortest among the cycles representing its free homotopy class in X.*

Proof. First assume the genus of X is positive. Then it contains two non-homotopic non-separating (and thus essential) simple closed curves α and β away from its boundary with intersection number one (not necessarily contained in $X^{(1)}$). Because every representative of the free homotopy class $[\alpha]$ should intersect β, we get that the shortest representative of this class is not contained in ∂X.

Now assume that X is a sphere with $b \ge 3$ boundary components and/or punctures. By the assumption that X is a proper surface, ∂X consists of disjoint connected components, each of which homeomorphic to S^1. A figure-eight curve around two of the punctures/boundary components is not homotopic in X to any power of a loop around one of the boundary components. This proves the lemma. □

Proposition 5.3 (Basic properties of core surfaces)**.** *Let J be a non-trivial subgroup of* Γ *and let* $\Upsilon \overset{\text{def}}{=} J\backslash\widetilde{\Sigma_g}$ *be the corresponding covering space of* Σ_g*. Then the following properties hold.*

1. *Every boundary cycle δ of* Core (J) *is an essential curve of* Υ*.*

2. *For every boundary cycle δ of* Core (J)*, $w(\delta)$ is a cyclically shortest word.*

3. Core (J) *is strongly boundary reduced.*

4. *Every connected component of the complement of* Core (J) *in* Υ *is homeomorphic to a twice-punctured sphere, with one funnel and one fake-puncture. In particular,* Core (J) *is a deformation retract of* $\Upsilon = J\backslash\widetilde{\Sigma_g}$*.*

5. $\text{Core}(J)$ *is connected.*

6. *The embedding* $\text{Core}(J) \hookrightarrow \Upsilon$ *induces an isomorphism in the level of fundamental groups.*

7. *In step (ii) of Definition 1.1, the connected components with finitely many $4g$-gons that are added to* $\text{Core}(J)$ *are either open discs or open annuli (twice-punctured spheres).*

Proof. Let δ be a boundary cycle of $\text{Core}(J)$. If δ is null-homotopic in Υ, then it bounds a disc in one of its sides. This side cannot be external to $\text{Core}(J)$, because then it should have been annexed to $\text{Core}(J)$ by part (ii) of Definition 1.1. If the disc is on the internal side of δ, then the connected component of δ in $\text{Core}(J)$ does not contain any essential curve of Υ. This is impossible. Hence δ is essential and (1) is proved.

For (2), by Corollary 4.9, it is enough to show that δ and δ^* contain no long blocks nor long chains. We begin with δ^*. Suppose that δ^* contains a long block b, and let $\overline{b}$ denote its complement (along the same $4g$-gon O_b of Υ). Consider the 1-skeleton of $\text{Core}(J)$. All the internal vertices in b (vertices contained in b but not at its endpoints) have degree two in $\text{Core}(J)^{(1)}$, so any non-backtracking cycle traversing one edge of b must traverse all of b, and can be shortened in Υ by replacing b with $\overline{b}$. So there is no shortest cyclic representative using any edge of b, and after step (i) of Definition 1.1, O_b belongs to the same connected component of $\Upsilon \backslash \text{Core}(J)$ as the $4g$-gons on the other side of b. But in step (ii) of Definition 1.1, O_b can only be annexed to $\text{Core}(J)$ if the entire component is, which is not the case. This is a contradiction. A similar argument shows that δ^* cannot contain any long chain c: indeed, any non-backtracking cycle in the 1-skeleton of $\text{Core}(J)$ that intersects the interior of c must contain a long block or a long chain, so no shortest cyclic representative intersects the interior of c. Hence δ^* contains no long block nor long chain.

We still need to show that δ contains neither long blocks nor long chains. Denote by $\mathbb{CORE}(J)$ a realization of the thick version of $\text{Core}(J)$ in Υ. In particular, replacing $\text{Core}(J)$ with $\mathbb{CORE}(J)$ does not alter the topology of the complement $\Upsilon \backslash \text{Core}(J)$. Let C be the connected component of $\Upsilon \backslash \mathbb{CORE}(J)$ bordering δ, and let $\overline{C}$ denote the closure of C in Υ. So the difference between C and $\overline{C}$ is that every fake-puncture of C becomes a closed connected component of $\partial \overline{C}$. We think of $\overline{C}$ as a tiled surface. Formally, every vertex or edge of Υ that belongs to two (or more) different boundary pieces $\overline{C}$, is duplicated in $\overline{C}$. Now $\overline{C}$ is a 2-complex with directed and labeled edges which satisfies properties **P1** and **P3** of Proposition 3.3. Because for every boundary cycle δ of $\text{Core}(J)$ we have that δ^* contains neither long blocks nor long chains, we deduce that $\overline{C}$ is boundary reduced. It now follows from Proposition 4.3 that $\overline{C}$ is indeed a legitimate tiled surface, and a boundary reduced one. By Corollary 4.10, as $\delta \subseteq \overline{C}$, we get that δ can be shortened to a shortest cyclic representative $\delta' \subseteq \overline{C}$ of its free-homotopy class in Υ.

Now suppose that δ contains a long block or a long chain. Then δ' is different from δ, and it is in $\text{Core}(J)$ by definition. So $\delta' \subseteq \overline{C} \cap \text{Core}(J)$. Because δ and δ' are isotopic, different, and lie in $\partial \overline{C}$, we get that C must be a two-punctured sphere containing finitely many $4g$-gons, and therefore should have been part of $\text{Core}(J)$ by part (ii) of Definition 1.1. This is a contradiction, and (2) is proven.

If the boundary component δ of $\text{Core}(J)$ contains a half-block, then a half-block switch yields another shortest representative and should be in $\text{Core}(J)$ together with the $4g$-gon along which the half-block lies. A similar argument works if δ is a half-chain. This proves (3).

Let C be (again) a connected component of $\Upsilon \setminus \mathbb{CORE}(J)$. As C is open, it is a surface and so $\overline{C}$ is a boundary reduced tiled surface which is a proper surface. Because $\overline{C}$ is boundary reduced, every free homotopy class of curves in $\overline{C}$ has a cyclically shortest representatives in C without long blocks or long chains (Corollary 4.10). But such a cycle also belongs to $\text{Core}(J)$ by the definition of a core surface, and so is contained in $\partial \overline{C}$. By Lemma 5.2, C must be a sphere with at most two punctures. By the fact that Υ is connected and by part (ii) of Definition 1.1, C must have two punctures: one fake and one a funnel. This settles item (4). Items (5) and (6) follow immediately.

Finally, let Y denote the union of cyclically shortest representatives in Υ as in the first step of Definition 1.1. Let C be a connected component of the complement of the thick version of Y in Υ with finitely many $4g$-gons. As C is open, it is a surface, and therefore $\overline{C}$ is a proper surface. As above, we think of $\overline{C}$ as a 2-complex (while duplicating vertices and edges of Υ appearing in different boundary pieces of $\overline{C}$), its boundary is reduced because Y is made of shortest cycles only, and therefore $\overline{C}$ is a tiled surface by Proposition 4.3. By Lemma 5.2, unless $\overline{C}$ is a disc or an annulus, it contains a cyclically shortest cycle not contained in $\partial\overline{C}$. This is a contradiction to the definition of Y, and item (7) is proven. $\square$

Consider two subgroups $J_1, J_2 \leq \Gamma_g$. It follows from a standard fact in the theory of covering spaces that there is a morphism of tiled surfaces $J_1\backslash\widetilde{\Sigma_g} \to J_2\backslash\widetilde{\Sigma_g}$ commuting with the quotient maps from $\widetilde{\Sigma_g}$, if and only if ${J_1}^{\gamma} \leq J_2$ for some conjugate ${J_1}^{\gamma} = \gamma J_1 \gamma^{-1}$ of J_1. In this case, any morphism $J_1\backslash\widetilde{\Sigma_g} \to J_2\backslash\widetilde{\Sigma_g}$ restricts to a morphism of the corresponding core surfaces:

Lemma 5.4. *Let $J_1 \leq J_2 \leq \Gamma_g$ and let $f\colon J_1\backslash\widetilde{\Sigma_g} \to J_2\backslash\widetilde{\Sigma_g}$ be the natural morphism. Then,*

1. *f restricts to a morphism* $\mathrm{Core}\,(J_1) \to \mathrm{Core}\,(J_2)$, *and*
2. *for $1 \neq \gamma \in J_1$, every shortest representative cycle of γ^{J_2} in* $\mathrm{Core}\,(J_2)$ *is an f-image of a shortest representative cycle of γ^{J_1} in* $\mathrm{Core}\,(J_1)$.

Proof. For $i = 1, 2$, denote $\Upsilon_i = J_i\backslash\widetilde{\Sigma_g}$. By definition, the morphism f preserves the orientation and labels of edges. So it follows from Corollaries 4.9 and 4.10 that it maps shortest representative cycles of free-homotopy classes to shortest representative cycles of free-homotopy classes. We now show that the connected components we add to the core surface of J_1 in part (ii) of Definition 1.1 are also mapped to the core surface of J_2.

Let T be such a connected component, namely, a connected component of the complement of the union of shortest cycles in Υ_1 which is added to $\mathrm{Core}\,(J_1)$ in part (ii) of Definition 1.1. Consider a $4g$-gon O in T. Let T' denote the connected component of $f\,(O)$ in the complement of the union of shortest cycles in Υ_2. We claim that T' contains finitely many $4g$-gons, and therefore, by Definition 1.1, must be contained in $\mathrm{Core}\,(J_2)$. In fact, all the $4g$-gons in T' are images of $4g$-gons in T, and therefore there are finitely many of them. Indeed, for every $4g$-gon O' in T', there is a "path of $4g$-gons" inside T' from $f\,(O)$ to O', where each $4g$-gon shares an edge with the previous one. If we lift this path to a path of $4g$-gons from O in Υ_1, it cannot leave the connected component T. Hence O' is an image of some $4g$-gon in T. This proves (1).

Now let $1 \neq \gamma \in J_1$ and let $\mathcal{C}$ be a shortest representative cycle of γ^{J_1} in Υ_1, so its image $f\,(\mathcal{C})$ is a shortest representative for γ^{J_2} in Υ_2, as noted above. For any other cycle $\mathcal{C}'$ which is a shortest representative of γ^{J_2} in Υ_2, the (free) homotopy between $f\,(\mathcal{C})$ and $\mathcal{C}'$ in Υ_2 can be lifted to Υ_1 (this follows from the general theory of covering spaces, e.g., [Hat05, Page 30]) and therefore, in particular, $\mathcal{C}'$ is an f-image of a cycle in Υ_1 which represents γ^{J_1} and of the same length as $\mathcal{C}$. This proves (2). $\square$

Lemma 5.5. *Let Y be a strongly boundary reduced tiled surface embedded in $\Upsilon = J\backslash\widetilde{\Sigma_g}$ and let $\mathcal{C} \subseteq Y^{(1)}$ be a non-nullhomotopic cycle. Then* every *shortest representative in Υ of the free homotopy class of $\mathcal{C}$ is contained in Y.*

Proof. By Corollary 4.10(1), there is a shortest representative of the free homotopy class of $\mathcal{C}$ in Y, and without loss of generality, assume that $\mathcal{C}$ is shortest. Further assume that $\mathcal{C}$ represents the conjugacy class γ^J in J, and consider the morphism $f\colon \langle\gamma\rangle\backslash\widetilde{\Sigma_g} \to \Upsilon$. Lemma 5.4(2) shows that every shortest representative of γ^J in Υ is an f-image of a shortest representative in $\mathrm{Core}\,(\langle\gamma\rangle)$, and Lemma 5.1 shows that are finitely many such representatives in $\mathrm{Core}\,(\langle\gamma\rangle)$. As the f-image of a half block (a half-chain) in $\mathrm{Core}\,(\langle\gamma\rangle)$ is a half-block (half-chain respectively) in Υ, we get that all shortest representatives of γ^J are obtained from $\mathcal{C}$ by half-block switches or a half-chain switch. As Y is strongly boundary reduced, it contains the complement of every half-block or half-chain in it. $\square$

Lemma 5.6. *Let Y be a connected boundary reduced tiled surface embedded in $\Upsilon = J \backslash \widetilde{\Sigma_g}$ such that $p_* (\pi_1 (Y, y)) = J \leq \Gamma_g$. Then every component of $\Upsilon \setminus Y$ is a twice-punctured sphere with one funnel and one fake-puncture.*

Proof. As $p_* (\pi_1 (Y, y)) = J$, Y contains a representative of every free homotopy class in Υ. Let C be a connected component of $\Upsilon \setminus Y$. As in the proof of Lemma 5.2, C cannot have positive genus. It cannot be a once-punctured sphere because if this puncture is a funnel, Υ is not connected, and if it is a fake-puncture, the boundary component of Y along this fake-puncture is not boundary reduced (by [Deh12]). In addition, C cannot have two fake-punctures, because Y is connected and so there would be a free homotopy class of essential curves such that any representative must go through C (between these two punctures). If C contains a funnel, then any cycle in C representing a loop around this funnel has a shortest representative in Y. These two cycles are isotopic in Υ and therefore bound an annulus. So C cannot contain two different funnels. We conclude that C is a twice-punctured sphere with one funnel and one fake-puncture. $\square$

Lemma 5.7. *Let Y be a strongly boundary reduced tiled surface embedded in $\Upsilon = J \backslash \widetilde{\Sigma_g}$ such that $p_* (\pi_1 (Y, y)) = J \leq \Gamma_g$. Then $Y \supseteq \mathrm{Core}\,(J)$.*

Proof. By Lemma 5.5, Y contains every shortest representative of every non-trivial free homotopy class, and so contains the subcomplex from part (i) of Definition 1.1. By Lemma 5.6, every component of the complement of Y contains a funnel and thus infinitely many $4g$-gons. In particular, it cannot be contained in one of the components added to $\mathrm{Core}\,(J)$ in part (ii) of Definition 1.1. This completes the proof. $\square$

Proposition 5.8. *If $J \leq \Gamma$ is finitely generated then* $\mathrm{Core}\,(J)$ *is compact.*

Proof. Suppose that $J \leq \Gamma$ is finitely generated and let $S = \{w_1, \ldots, w_k\}$ be a finite generating set represented as words in $\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\}$. Let $(\Upsilon, q) = J \backslash \left(\widetilde{\Sigma_g}, u\right)$ be the pointed quotient of $\widetilde{\Sigma_g}$ with the base point q being the image of the base point u. Then $w_1, \ldots, w_k$ correspond to unique, possibly not cyclically reduced, cycles $\mathcal{C}_1, \ldots, \mathcal{C}_k$ based at q. Consider the sub-surface Y of Υ consisting of the union $\bigcup_{i=1}^k \mathcal{C}_i$, and let $Y' = \mathsf{BR}\,(Y \hookrightarrow \Upsilon)$. By Propositions 4.5 and 4.6, Y' is a compact boundary reduced tiled surface containing Y.

Finally, enlarge Y' to obtain a slightly larger sub-surface $Y'' \subseteq \Upsilon$ by repeatedly adding any $4g$-gon which borders some half-block or half-chain in $\partial Y'$. As in the proof of Lemma 5.5, the $4g$-gons added next to a boundary cycle representing γ^J are all images of the finitely many $4g$-gons in $\mathrm{Core}\,(\langle\gamma\rangle)$, and so there are finitely many steps near this boundary component. As Y' is compact, it has finitely many boundary components, and hence Y'' is constructed in finitely many step and is compact too. Moreover, by the way it is constructed, Y'' is strongly boundary reduced. By Lemma 5.7, the compact Y'' contains $\mathrm{Core}\,(J)$. This proves the proposition. $\square$

We can now give an intrinsic definition for a core-surface, not relying on a given subgroup of Γ_g.

Proposition 5.9 (Intrinsic definition of a core surface)**.** *A (non-empty) tiled surface is a core surface if and only if it is (i) connected, (ii) strongly boundary reduced, (iii) every boundary cycle is a cyclically shortest representative of its free homotopy class, and (iv) it contains no funnels.*

Note that if Y is a *compact* tiled surface, then condition (iv) in the proposition automatically holds.

Proof. That a core surface satisfies properties (i), (ii) and (iii) is the content of items (5), (3) and (2), respectively, in Proposition 5.3. Now let $Y = \mathrm{Core}\,(J) \subseteq \Upsilon = J\backslash\widetilde{\Sigma_g}$ for some $J \leq \Gamma_g$. Let p be a funnel-puncture in Υ and let $\mathcal{C} \subseteq \Upsilon^{(1)}$ be a cyclically shortest representative of a simple closed curve around p. As in the proof of Lemma 5.5, there are finitely many shortest representatives of the free homotopy class of $\mathcal{C}$, and thus the connected component C of p in the complement of these representatives in Υ, contains infinitely many $4g$-gons. Moreover, $\Upsilon \setminus C$ is strongly boundary reduced and a retract of Υ, and therefore contains every shortest representative of free homotopy classes in Υ. By Definition 1.1, C cannot belong to $\mathrm{Core}\,(J)$ and thus (iv) holds.

Conversely, assume that Y is a tiled surface satisfying these four assumptions. Let $p\colon Y \to \Sigma_g$ be the immersion. Choose some vertex $y \in Y$ and let $J = p_*\left(\pi_1\left(Y, y\right)\right) \leq \Gamma_g$. As in the proof of Proposition 3.3, there is an embedding $r\colon Y \hookrightarrow \Upsilon = J\backslash\widetilde{\Sigma_g}$, so we may think of Y as a sub-complex of Υ. Since Y has the same fundamental group as Υ and is strongly boundary reduced, it contains every shortest representatives of free homotopy classes in Υ, by Lemma 5.5. As in the proof of Proposition 5.8, every connected component of $\Upsilon \setminus Y$ is a twice-punctured sphere with one funnel and one fake-puncture and, in particular, contains infinitely many $4g$-gons. Hence $Y \supseteq \mathrm{Core}\,(J)$.

Finally, by Proposition 5.3(4), every connected component C of $\Upsilon\setminus\mathrm{Core}\,(J)$ is a twice-punctured sphere, with one funnel and one fake-puncture. By (iv), Y does not contain the whole of C, and by (iii), Y does not contain any point of C. We conclude that $Y = \mathrm{Core}\,(J)$. $\square$

Using the intrinsic definition of core surfaces from Proposition 5.9, we conclude that we have a one-to-one bijection

$$\left\{\begin{matrix}\text{conjugacy classes of}\\ \text{subgroups of } \Gamma_g\end{matrix}\right\} \longleftrightarrow \left\{\begin{matrix}\text{core surfaces}\\ \text{labeled by } \{a_1, \ldots, b_g\}\end{matrix}\right\} \tag{5.1}$$

which restricts to a one-to-one correspondence

$$\left\{\begin{matrix}\text{conjugacy classes of}\\ \text{f.g. subgroups of } \Gamma_g\end{matrix}\right\} \longleftrightarrow \left\{\begin{matrix}\text{compact core surfaces}\\ \text{labeled by } \{a_1, \ldots, b_g\}\end{matrix}\right\}. \tag{5.2}$$

5.2 Foldings and construction of core surfaces

One of the most useful concepts introduced in [Sta83] is that of "foldings", now known as Stallings foldings. In graphs, a folding is a process in which one merges two equally-labeled oriented edges with the same head-vertex or with the same tail-vertex. Occasionally, one also trims leaves from the graph. This process allows one to construct the core graph of a f.g. subgroup H of the free group $\mathbf{F}_r$ from a finite set $\{w_1, \ldots, w_k\}$ of generators as follows: create a bouquet with k petals, one for every generator. Then fold until no more folding steps are possible. The resulting graph is the core graph of H (e.g. [KM02, Proposition 3.8]).

We now present an analogous folding process for f.g. subgroups of surface groups and their core surfaces.

Theorem 5.10 (Foldings)**.** *Let $J \leq \Gamma_g$ be a f.g. subgroup and let $\{w_1, \ldots, w_k\}$ be a generating set consisting of words in $\left\{a_1^{\pm 1}, b_1^{\pm 1}, \ldots, a_g^{\pm 1}, b_g^{\pm 1}\right\}$. Then* $\mathrm{Core}\,(J)$ *can be constructed via the following finite process:*

1. ***Preparation:*** *Without loss of generality, we may assume all words $w_1, \ldots, w_k$ represent non-trivial elements in Γ_g (this can be easily and efficiently checked using Dehn's algorithm, and trivial elements may be removed from the set).*

   - *Consider a cycle $\mathcal{C}$ representing w_1 and shorten it* cyclically *until one obtains a shortest representative of $w_1^{\Gamma_g}$ (as in Theorem 4.8). Assume the new cyclically shortest word represents the same element as $w_1{}^s = sw_1s^{-1}$ for some word s in $\left\{a_1^{\pm 1}, \ldots, b_g^{\pm 1}\right\}$.*

- *Replace w_1 by the new cyclically shortest representative of $w_1^{\Gamma_g}$, replace $w_2, \ldots, w_k$ by ${w_2}^s, \ldots, {w_k}^s$, and shorten the latter $k-1$ words using Dehn's algorithm (namely, consider the path representing w_i and repeatedly replace long blocks or long chains with their complements). Rename the new words $w_1, \ldots, w_k$, and replace J with J^s (recall that this does not change the corresponding core surface).*
- *Now construct a wedge of k petals, where petal i consists of $|w_i|$ directed edges labeled by $\{a_1, \ldots, b_g\}$ so that it reads the word w_i. Call the resulting directed and edge-labeled graph Y_1.*

2. ***Folding and boundary reduction:*** *Let $Y = Y_1$. Perform the following two steps alternately until none of them is possible, always beginning with folding:*

   - ***Folding edges and $4g$-gons:*** *Fold the 1-skeleton of Y (in the sense of Stallings: so repeatedly merge together pairs of equally-labeled edges with the same head or the same tail), and remove multiplicities of $4g$-gons sharing the same oriented boundary, so that there is at most one $4g$-gon attached to any closed $[a_1, b_1] \ldots [a_g, b_g]$ path.*
   - ***Boundary reduction:*** *If ∂Y contains a long block, choose one such block and add a new $4g$-gon along it so that this block in ∂Y is replaced by its complement. Otherwise, if ∂Y contains a long chain, choose one such long chain and add $4g$-gons along it so that this long chain in ∂Y is replaced by its complement.*

   Call the resulting complex Y_2.

3. ***Strong boundary reduction:*** *Finally, as long as ∂Y contains a half-block or a half-chain, add new $4g$-gon along them so that this piece of ∂Y is replaced by the complement of the half-block or half-chain. Call the resulting complex Y_3.*

Proof. We need to show that the process described is well-defined, i.e., that the boundary ∂Y of Y is well defined whenever we use it (notice that Y may not even be a tiled surface along the way), that it terminates after finitely many steps, and that the resulting complex is indeed $\mathrm{Core}\,(J)$ (and so, in particular, independent of the choices made along the way). We analyze the three parts of the process one by one. Let $(\Upsilon, q) = J \backslash \left(\widetilde{\Sigma_g}, u\right)$ and let $\pi\colon (\Upsilon, q) \to (\Sigma_g, o)$ be the covering map.

In the first part, denote by v the wedge point of the k petals of Y_1. Trivially, there is a map $p\colon (Y_1, v) \to (\Sigma_g, o)$ and $p_*(\pi_1(Y_1, v)) = J \leq \Gamma_g$. As in (3.1), there is a (unique) lift $r\colon (Y_1, v) \to (\Upsilon, q)$ with $\pi \circ r = p$.

$$\begin{tikzcd} & (\Upsilon, q) \arrow[d, "\pi"] \\ (Y_1, v) \arrow[ur, dashed, "\exists! r"] \arrow[r, "p"'] & (\Sigma_g, o) \end{tikzcd}$$

By definition, $\mathrm{Core}\,(J)$ is a subcomplex of Υ. Because w_1 is a cyclically shortest cycle, so is its r-image in Υ (by Theorem 4.8), and therefore its image is contained in $\mathrm{Core}\,(J)$. In particular, $q \in \mathrm{Core}\,(J)$. Now for $i = 2, \ldots, k$, the element $\gamma_i \in J$ represented by w_i has a representative in $\pi_1(\mathrm{Core}\,(J), q)$. Because $\mathrm{Core}\,(J)$ is boundary reduced, $\mathrm{Core}\,(J)^{(1)}$ contains also a shortest representative of γ_i based at q (because one can perform Dehn's algorithm inside $\mathrm{Core}\,(J)$). Call this path p_i. So now $r(w_i)$ and p_i are two closed paths at $\Upsilon^{(1)}$, based at q, representing the same element. They lift to two paths starting at u with the same endpoint in $\widetilde{\Sigma_g}$. By [BS87, Thm 2.8], any two shortest paths with the same endpoints in $\widetilde{\Sigma_g}^{(1)}$ differ by a finite sequence of half-block switches. This sequence of half-block switches descends to a sequence of half-block switches in Υ which turns p_i into $r(w_i)$. Because $\mathrm{Core}\,(J)$ is strongly boundary reduced (by Proposition 5.3(3)), these half-block switches all take place inside $\mathrm{Core}\,(J)$. We conclude that $r\,(Y_1) \subseteq \mathrm{Core}\,(J)$.

Now consider the "folding and boundary reduction" part of the process. In every folding step (of an edge or of removing $4g$-gons), the total number of cells in Y decreases, so every iteration of "folding" must terminate. At the end of such an iteration, properties **P1** and **P3** from Proposition 3.3 hold. As mentioned in the paragraph preceding Proposition 4.3, this guarantees that Y admits a well-defined thick version and therefore that ∂Y is well defined. Thus every "boundary reduction" iteration is well defined. Clearly, in every non-empty iteration of boundary reduction, the length of ∂Y strictly decreases, and in every folding iteration, this length does not increase. Therefore, the second part of the process is well defined and finite.

In addition, along the second part of the process, the map $p\colon (Y,v) \to (\Sigma_g, o)$ is defined and $p_*\,(\pi_1\,(Y,v)) = J$ at every step, so there is a corresponding lift $r\colon (Y,v) \to (\Upsilon, q)$ at every step. A folding step does not alter the image of r. Moreover, every $4g$-gon added to Y in a boundary reduction step, is added along some long block (chain) at ∂Y, which is mapped by r to a long block (chain, respectively) in Υ. But $\mathrm{Core}\,(J)$ is boundary reduced, so for every long block (chain) in $\mathrm{Core}\,(J)$, the $4g$-gons along it also belong to $\mathrm{Core}\,(J)$. By induction we thus see that $r\,(Y) \subseteq \mathrm{Core}\,(J)$ throughout part 2 of the process.

At the end of the second part, Y_2 is a complex with edges directed and labeled by $\{a_1, \ldots, b_g\}$, which satisfies **P1** and **P3** and which is also boundary reduced. By Proposition 4.3, it is a tiled surface, and as in the proof of Proposition 3.3, the unique lift $r\colon (Y_2, v) \to (\mathrm{Core}\,(J)\,, q)$ must be an embedding. So Y_2 is a boundary reduced subcomplex of $\mathrm{Core}\,(J)$ with the same fundamental group.

Finally, we can think of the third part of the process as taking place in Υ: as in the proof of Lemma 5.5, it is a finite process in which the length of the boundary does not change. As $\mathrm{Core}\,(J)$ is strongly boundary reduced, the third part never leaves $\mathrm{Core}\,(J)$. So at the end of the third part, Y_3 is a strongly boundary reduced subcomplex of $\mathrm{Core}\,(J)$ with the same fundamental group. By Lemma 5.7, we actually have $Y_3 = \mathrm{Core}\,(J)$. □

6 Epilogue

This paper came out as a side of our work [MP20, MNP22] on random coverings of compact surfaces. We tried to elaborate here on some basic properties of core surfaces which we use in ibid, as well as some basic properties that illustrate the resemblance of core surfaces to Stallings core graphs. However, we made no systematic attempt to (re-)prove results about subgroups of the surface group Γ_g using core surfaces. We believe core surfaces should be useful here, and think that a more systematic attempt in this direction should be taken in the future.

References

[AO96] G. N. Arzhantseva and A. Y. Ol'shanskii. Generality of the class of groups in which subgroups with a lesser number of generators are free. *Mat. Zametki*, 59(4):489–496, 1996. 4

[Arz98] G. N. Arzhantseva. Generic properties of finitely presented groups and Howson's theorem. *Comm. Algebra*, 26(11):3783–3792, 1998. 4

[BL18] B. Beeker and N. Lazarovich. Stallings' folds for cube complexes. *Israel J. Math.*, 227(1):331–363, 2018. 4

[Bro17] S. T. Brown. *Geometric structures on negatively curved groups and their subgroups*. PhD thesis, University College London, 2017. 4

[BS87] J. S. Birman and C. Series. Dehn's algorithm revisited, with applications to simple curves on surfaces. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 451–478. Princeton Univ. Press, 1987. 9, 10, 13, 14, 21

[BZKL22] M. Ben-Zvi, R. Kropholler, and R. A. Lyman. Folding-like techniques for CAT(0) cube complexes. *Math. Proc. Cambridge Philos. Soc.*, 173(1):227–238, 2022. 4

[Deh12] M. Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Math. Ann.*, 72(3):413–421, 1912. 7, 12, 13, 19

[DL21] P. Dani and I. Levcovitz. Subgroups of right-angled coxeter groups via Stallings-like techniques. *J. Comb. Algebra*, 5(3):237–295, 2021. 4

[DV22] J. Delgado and E. Ventura. A list of applications of Stallings automata. *Trans. Comb.*, 2022. to appear, available at arXiv:2109.01268. 1, 4, 6

[Hat05] A. Hatcher. *Algebraic topology*. Cambridge University Press, 2005. 7, 18

[HP22] L. Hanany and D. Puder. Word measures on symmetric groups. *Int. Math. Res. Not. IMRN*, 2022+. to appear; also at arXiv:2009.00897. 6

[KM02] I. Kapovich and A. Miasnikov. Stallings foldings and subgroups of free groups. *J. Algebra*, 248:608–668, 2002. 1, 6, 20

[KMW17] O. Kharlampovich, A. Miasnikov, and P. Weil. Stallings graphs for quasi-convex subgroups. *J. Algebra*, 488:442–483, 2017. 4

[LS77] R. C. Lyndon and P. E. Schupp. *Combinatorial group theory*, volume 188. Springer, 1977. 12

[MNP22] M. Magee, F. Naud, and D. Puder. A random cover of a compact hyperbolic surface has relative spectral gap $\frac{3}{16} - \varepsilon$. *Geom. Funct. Anal. (GAFA)*, 32:595–661, 2022. 3, 22

[MP20] M. Magee and D. Puder. The asymptotic statistics of random covering surfaces. arXiv preprint arXiv:2003.05892, 2020. 3, 4, 11, 13, 22

[PP15] D. Puder and O. Parzanchevski. Measure preserving words are primitive. *J. Amer. Math. Soc.*, 28(1):63–97, 2015. 1, 6

[Pud14] D. Puder. Primitive words, free factors and measure preservation. *Israel J. Math.*, 201(1):25–73, 2014. 6

[PZ22] D. Puder and T. Zimhoni. Local statistics of random permutations from free products. arXiv preprint arXiv:2203.12250, 2022. 4

[Sta83] J. R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983. 1, 6, 20

Michael Magee, Department of Mathematical Sciences, Durham University, Lower Mountjoy, DH1 3LE Durham, United Kingdom
`michael.r.magee@durham.ac.uk`

Doron Puder, School of Mathematical Sciences, Tel Aviv University, Tel Aviv, 6997801, Israel
`doronpuder@gmail.com`