

NOTES ON MASKIT'S PLANARITY THEOREM

BRIAN H. BOWDITCH

ABSTRACT. We give an account of the Planarity Theorem of Maskit. This gives a classification of finitely generated groups acting effectively properly discontinuously by orientation-preserving homeomorphisms on a planar surface. One can also realise such groups as kleinian function groups. We also explain how one can give another proof of the planarity theorem using Dunwoody's theory of tracks.

1. INTRODUCTION

In this paper we give a discussion of the Planarity Theorem of Maskit [Mas1], and some related results. This effectively gives a classification of groups of homeomorphisms of planar surfaces which have finite-type quotients. In particular, they can be described in terms of geometrically finite “function groups” acting by Möbius transformations on the Riemann sphere, which we will view here as the boundary of hyperbolic 3-space (cf. [Mas2]). Topologically they can be described as a class of orbifold fundamental groups of a particular class of 3-orbifolds, which generalise the standard notion of a compression body. They can also be described in terms of finitely generated groups with planar Cayley graphs, as described in [G].

Much of what we do here can be seen as exposition of known results, though the arguments are a bit different, and there are a number of statements which we could not find in the literature. We will also explain how one can give a different proof of the Planarity Theorem using the theory of tracks due to Dunwoody [D].

A “planar surface” is a space homeomorphic to a connected open subset of the 2-sphere \mathbb{S}^2 . A general classification of (infinite type) surfaces is given in [R]. The case of planar surfaces is somewhat simpler. In particular, a planar surface is uniquely determined by its space of ends — a totally disconnected compactum. Of particular interest here is the “Cantor surface”, where the space of ends is a Cantor set.

We begin with a statement of the Planarity Theorem [Mas1].

In the following discussion, all groups will be assumed to act effectively properly discontinuously by orientation preserving homeomorphisms. (Here “effective”

Date: First draft: 5th June 2019. Revised: 25th August 2020.

means that only the identity element fixes the space pointwise, and “properly discontinuous” means that any compact subset meets only finitely many images of itself.)

Let S be a planar surface. To simplify the exposition, we will assume that $S \not\cong \mathbb{S}^2$ (since that case is elementary). Thus, we can write $S = W/N$, where $W \cong \mathbb{R}^2$, and the group $N \cong \pi_1(S)$ acts freely. (In fact, N is a countable free group.) Suppose that a group G acts on S . Let $\Sigma = S/G$ be the quotient orbifold. We can write $\Sigma = W/\Gamma$, where Γ acts on W , and where $N \triangleleft \Gamma$ and $G = \Gamma/N$. In summary, we have $\Sigma = W/\Gamma = (W/N)/(\Gamma/N) = S/G$. We write $C \subseteq \Sigma$ for its set of cone points.

By a *multicurve* $\gamma \subseteq \Sigma$, we mean a closed subset which is a disjoint union of simple closed curves in $\Sigma \setminus C$. We will assume that no component of γ is trivial and that no two are parallel. In other words, each disc or annulus component of $\Sigma \setminus \gamma$ meets C . We write $\tilde{\gamma} \subseteq S$ for its preimage in S . We say that γ is *liftable* if each component of $\tilde{\gamma}$ is compact. In this case, each component of $\tilde{\gamma}$ determines an infinite cyclic subgroup of N , well defined up to conjugacy. Together, the collection of all such components generate a normal subgroup, $N_{\tilde{\gamma}} \triangleleft N$.

In general, we will consider “curves” to be defined up to homotopy (or equivalently isotopy) in $\Sigma \setminus C$. We will sometimes abuse terminology or notation by identifying such a curve with a particular realisation thereof.

We say that Σ is *finite type* if it is orbifold-equivalent to the interior of a compact orbifold with boundary. In such a case, both Γ and $G = \Gamma/N$ are finitely presented groups.

Theorem 1.1. (*Maskit*) *Let G act (effectively and properly discontinuously by orientation preserving homeomorphisms) on a planar surface, S , such that the quotient $\Sigma = S/G$ is an orbifold of finite type. Then there is a liftable multicurve $\gamma \subseteq \Sigma$ such that $N_{\tilde{\gamma}} = N$. (Here $N \cong \pi_1(S)$ and $N_{\tilde{\gamma}}$ are defined as above.)*

Necessarily, γ has finitely many components. It’s not hard to see that it contains all curves of the form $\delta(a)$, where $\delta(a)$ is the boundary of a small disc neighbourhood of a cone point $a \in C$ (considered to be defined up to homotopy in $\Sigma \setminus C$). We refer to such curves as *cone curves*.

In fact, Theorem 1.1 was originally stated with the additional assumption that the action of G is free, so that Σ is a finite-type surface. The general case can be easily deduced from this. We simply remove all cone points from Σ , as well as their preimages in S . We then get a multicurve in $\Sigma \setminus C$, to which we finally adjoin the set of cone curves.

There is also an inverse to this process, which will allow us to simplify some of the discussion.

Suppose we have an isolated end of S . Its stabiliser will be finite cyclic, and in the quotient, it descends to an end of Σ . We can adjoin points respectively to S and Σ in order to compactify these ends. The point we have added to Σ will be a cone point if the end-stabiliser in S is non-trivial, otherwise it is a non-singular

point. For this to work, we need to assume that S is not homeomorphic to \mathbb{R}^2 nor to an annulus. But those cases are again elementary, and can be treated separately. After compactifying the isolated ends in S (and again ruling out the cases where we end up with \mathbb{R}^2 or an annulus), we can assume that S has no isolated ends. In other words its space of ends is perfect, hence a Cantor set. It then follows that S is a *Cantor surface*, that is homeomorphic to \mathbb{S}^2 minus a Cantor set. (See Section 6 for further discussion of this.) In other words, we don't lose much by restricting to Cantor surfaces.

The conclusion of Theorem 1.1 can be interpreted topologically. We can construct a 2-complex by gluing a disc along its boundary to each component of $\tilde{\gamma}$ in S . The conclusion then says that this space is simply connected. If we imagine all the discs to be attached on “one side” of S , then we can thicken up this complex to give a simply connected 3-manifold, with one boundary component identified with S . We can carry out this construction G -equivariantly, so that the quotient, M , is a 3-orbifold with one “outer” boundary component identified with Σ .

One can also realise this geometrically as follows. Suppose that a group, G , acts properly discontinuously by isometry on hyperbolic 3-space, \mathbb{H}^3 . Let $\Omega G \subseteq \partial\mathbb{H}^3$ be its discontinuity domain in the ideal sphere, $\partial\mathbb{H}^3$. We have a quotient hyperbolic orbifold, $M(G) = \mathbb{H}^3/G$, as well an orbifold with boundary: $M_C(G) = (\mathbb{H}^3 \cup \Omega G)/G$. If G is geometrically finite, then $M_C(G)$ has finitely many ends, each a standard cusp region. (See Sections 5 and 6 for more detailed discussion.)

The following is a slight elaboration of the result of [Mas2]. (See Section 5 for a more detailed statement.)

Theorem 1.2. *Suppose a group G acts properly discontinuously on a Cantor surface, S , with finite-type quotient, $\Sigma = S/G$. Then G also admits a geometrically finite action on \mathbb{H}^3 , such that ΩG has a G -invariant component, Ω_0 , equivariantly homeomorphic to S . In this way, Σ can be identified with the boundary component, Ω_0/G , of $M_C(G)$. Moreover, each end of Σ corresponds to a rank-1 cusp region of $M_C(G)$.*

In fact, modulo some simple adjustments, $M_C(G)$ is topologically the same as the orbifold M constructed from Σ as described earlier (see Proposition 5.1).

The 3-orbifolds that arise in this way can be completely classified, and so we get a complete classification of G -actions on planar surfaces with finite-type quotient (up to G -equivariant homeomorphism). We will describe these in terms of “graphs of orbifolds” in Section 3.

In fact, we can drop the assumption that the quotient is of finite type, and instead just assume G to be finitely generated. However, to apply the result, we need to modify S and Σ slightly, by adjoining a certain canonical set of ends to each. (See Proposition 6.5.)

The multicurve, $\gamma \subseteq \Sigma$, given by Theorem 1.1 is not in general unique. However, the ambiguity can be largely understood. For example, if the G -stabiliser of each

component of $\tilde{\gamma}$ is non-trivial, then γ is determined up to isotopy in $\Sigma \setminus C$. (See Proposition 4.3.)

For completeness, we note that Theorem 1.1 also applies if S is \mathbb{S}^2 , \mathbb{R}^2 or an annulus. In the first two cases, N is trivial and $\gamma = \emptyset$. In the last case, $N \cong \mathbb{Z}$, and γ is the core curve of the annulus S . The possible actions of G are easily classified in these cases.

In Section 7, we note how one can give another, more direct, proof of the planarity theorem using the theory of “tracks” as introduced by Dunwoody [D]. In fact, although the hypotheses of the main result of [D] do not hold exactly, one can follow the argument more or less verbatim with some minor reinterpretations.

This paper arose out of discussions with Agelos Georgakopoulos. I thank him for explaining to me his results in this area, and for his comments on a preliminary draft of this paper.

2. ORBIFOLDS

We begin by recalling some basic facts about orbifolds in dimensions 2 and 3. Here we will assume all orbifolds to be orientable. Unless otherwise stated, we will assume them to be connected, and in dimension 2 also to have empty boundary.

Let Σ be a 2-orbifold. This consists of a topological surface with a discrete set, $C \subseteq \Sigma$, of *cone points* and with an integer *degree*, $p(a) \geq 2$, associated to each $a \in C$. We will write Σ_T for the underlying surface (forgetting cone points).

We write $\pi_1^o(\Sigma)$ for the orbifold fundamental group. One way to describe this is as follows. Given $a \in C$, let $\Delta(a)$ be a small disc neighbourhood of a (so that $C \cap \Delta(a) = \{a\}$). Write $\delta(a) = \partial\Delta(a)$ for the cone curve (as described in Section 1). Now remove the interior of $\Delta(a)$, and glue in another disc, $D(a)$, such that $\partial D(a)$ wraps $p(a)$ times around $\delta(a)$. We do this for all $a \in C$, so as to give us a 2-complex, $\hat{\Sigma}$, with a natural map to Σ . This map induces an isomorphism from $\pi_1(\hat{\Sigma})$ to $\pi_1^o(\Sigma)$.

We say that Σ is *topologically finite* if there is a compact subset $\Sigma_0 \subseteq \Sigma$, with $C \subseteq \text{int}(\Sigma_0)$ and with each component of $\Sigma \setminus \Sigma_0$ homeomorphic to $\mathbb{S}^1 \times [0, \infty)$. We refer to Σ_0 as a *core* for Σ . The components of $\Sigma \setminus \Sigma_0$ correspond to the ends of Σ . The orbifold type of Σ is determined by its genus, the number of ends, and the set of degrees, $p(a)$, for $a \in C$ (counting multiplicities). One can check that Σ is topologically finite if and only if $\pi_1^o(\Sigma)$ is finitely generated. In this case, each end of Σ determines a *peripheral* subgroup of $\pi_1^o(\Sigma)$, or more precisely a conjugacy class of such subgroups. We write \mathcal{H} for the collection of all such subgroups. Each element of \mathcal{H} is a maximal subgroup isomorphic to \mathbb{Z} , except for a couple of trivial cases (see (E0) and (E1) below). We note that another way to say that Σ is finite-type is to demand that it is orbifold-equivalent to the interior of a compact orbifold with boundary (which will be orbifold-equivalent to the core).

We say that Σ is *bad* if it is a 2-sphere either with one cone point, or with two cone points of differing degrees. Otherwise it is *good*. The latter is equivalent to saying that it has a manifold cover, i.e. it arises as the quotient of a surface by a properly discontinuous group of homeomorphisms. In this case, the (orbifold) universal cover is homeomorphic to either \mathbb{S}^2 or \mathbb{R}^2 .

Henceforth, we will assume that all our 2-orbifolds are good.

Let Σ be a (good) topologically finite 2-orbifold, and let $G = \pi_1^o(\Sigma)$. If $a \in C$, then the loop $\delta(a)$ encircling a determines a subgroup, $G(a) \leq G$, well defined up to conjugacy, and isomorphic to $\mathbb{Z}_{p(a)}$. In fact, the subgroups, $G(a)$ are precisely the maximal non-trivial finite (cyclic) subgroups of G .

We note that Σ is uniquely determined by the pair (G, \mathcal{H}) . To see this, note that the number of ends is equal to the number of conjugacy classes of \mathcal{H} (except when $G \cong \mathbb{Z}$ and Σ is an annulus). The cone degrees are determined by the orders of the subgroups $G(a)$. After quotienting G by the $G(a)$ and the peripheral subgroups, we get a compact surface group, whose genus is equal to the genus of Σ .

We have the following classification of topologically finite 2-orbifolds according to the (singular) geometric structure they admit.

(S) “spherical”. Here Σ_T is \mathbb{S}^2 . In this case, $\mathcal{H} = \emptyset$, and $|G| < \infty$.

(E) “euclidean”.

We distinguish three subcases:

(E0): $\Sigma_T \cong \mathbb{R}^2$ and $|C| \leq 1$. Here G is finite cyclic (or trivial) and there is just one peripheral subgroup, namely G itself.

(E1): G is virtually \mathbb{Z} . In this case, Σ is either an annulus with no cone points, or a disc with two cone points each of degree 2. In the former case, $G \cong \mathbb{Z}$ and there is just one peripheral subgroup namely G (even though Σ has two ends). In the second case, G is infinite dihedral. There is again one peripheral subgroup, this time of index 2 in G .

(E2): G is virtually \mathbb{Z}^2 . In this case, $\mathcal{H} = \emptyset$, and Σ is finitely covered by $\mathbb{S}^1 \times \mathbb{S}^1$.

(H) “hyperbolic”. In this case Σ admits a finite-area (singular) hyperbolic structure. In other words, it is the quotient of the hyperbolic plane by a finite-coarea fuchsian group, where the peripheral subgroups (if any) are precisely the maximal parabolic subgroups.

We will consider the case (E0) above as “exceptional”, and will want to rule it out from some of the subsequent discussion.

We will also discuss 3-orbifolds. In this case, we will allow non-empty boundary. We do not assume a-priori that the 3-orbifolds are good (i.e. have a manifold covering space) though this will turn out to be true of the 3-orbifolds we consider (see Lemma 3.2).

Let M be a 3-orbifold with (possibly empty) boundary, ∂M . The boundary is a (possibly disconnected) 2-orbifold. Let $\sigma(M) \subseteq M$ be the singular locus. This is a graph whose terminal (valence 1) vertices are precisely those which lie in ∂M , and give rise to the cone points of ∂M . Every other vertex has valence 3 — its link in M is a spherical 2-orbifold. Each edge of $\sigma(M)$ has an associated *degree* at least 2. In general, $\sigma(M)$ may also have circular components. We write $\pi_1^o(M)$ for the orbifold fundamental group.

We say that M is *simple* if $\sigma(M)$ is a disjoint union, $\sigma(M) = \bigsqcup_{e \in E} I(e)$, of compact intervals, $I(e)$, indexed by some set E . Note that $\partial I(e) = I(e) \cap \partial M$ consists of two cone points of the same degree $p(e) \geq 2$.

If M is simple, we can describe $\pi_1^o(M)$ similarly as for surface groups. Namely, we can construct a 3-complex, \hat{M} , with a natural map to M which induces an isomorphism from $\pi_1(\hat{M})$ to $\pi_1^o(M)$. The operation on each $I(e)$ is the same as the 2-orbifold case, just taking a direct product with a compact interval. Note that the preimage, $\partial \hat{M}$, of ∂M , can be identified with the space $\partial \hat{M}$ as defined for 2-orbifolds.

3. GRAPHS OF ORBIFOLDS

In this section, we describe what we mean by a “graph of orbifolds” which leads to a generalisation of the standard notion of a compression body. First, we introduce a “graph of surfaces”.

Definition. A *graph of surfaces* is connected 2-complex,

$$\Phi = \left(\bigsqcup_{v \in V} \Pi(v) \right) \cup \left(\bigsqcup_{e \in E} I(e) \right),$$

where each $\Pi(v)$ is a surface, and each $I(e)$ is an interval with each of its endpoints lying in some $\Pi(v)$. Here V and E are arbitrary indexing sets.

We write $\Pi = \bigcup_{v \in V} \Pi(v)$ and $\tau = \bigcup_{e \in E} I(e)$, so that $\Phi = \Pi \cup \tau$.

Given any graph of surfaces, $\Phi = \Pi \cup \tau$, we can construct a 3-manifold, $M = M(\Phi)$, by “thickening it up”. More formally, we take a copy of $\Pi \times [-1, 1]$ together with τ times a closed 2-disc, and glue them together according to the combinatorial structure dictated by Φ . The handles are all attached to $\Pi \times \{1\}$, so these surfaces are all connected together to form a single boundary component, $\partial^+ M$, of M . We can identify Π with $\Pi \times \{-1\} \equiv \partial M \setminus \partial^+ M$. This is a union of components of ∂M , which we denote by $\partial^- M$. We refer to $\partial^- M$ and $\partial^+ M$ respectively as the *inner* and *outer* boundaries of M . (We can embed M into \mathbb{R}^3 so that $\partial^+ M$ is the boundary of the unbounded component of $\mathbb{R}^3 \setminus M$, and such that each component of $\partial^- M$ bounds a handlebody in the complement of M .) We can view Φ as being embedded in M , so that if we remove an open regular neighbourhood of τ in M , the result is homeomorphic to $\partial^+ M \times [-1, 1]$. In particular, M deformation retracts onto Φ .

It is also convenient to define a 3-manifold, $M_0 \supseteq M$, by capping off each sphere component of $\partial^- M$ by a closed ball. Clearly the inclusions, $\Phi \hookrightarrow M \hookrightarrow M_0$ induce isomorphisms on π_1 . The inner boundary, $\partial^- M_0$, of M_0 is just $\partial^- M$ with the sphere components removed.

As an example, if Π is compact, then M_0 is a compression body in the usual sense. If also $\partial^- M_0 = \emptyset$, then M_0 is a handlebody.

More generally, we make the following definition:

Definition. We say that Φ is of *finite-type* if $|V|, |E| < \infty$ and each $\Pi(v)$ is a finite type surface.

In this case, the ends of Φ , Π , M and $\partial^+ M$ are all in natural bijective correspondence. Note that each end of M is homeomorphic to $\mathbb{S}^1 \times [-1, 1] \times [0, \infty)$, with $\mathbb{S}^1 \times \{\pm 1\} \times [0, \infty) \subseteq \partial^\pm M$. (We will generally use the word “end” to refer either to an end or to a neighbourhood thereof. In other words, we assume we have chosen a particular disjoint set of neighbourhoods of the ends.)

Given a graph of surfaces, Φ , we can construct a connected graph, $\Theta = \Theta(\Phi)$, by collapsing each component of Π to a point. In this way, the indexing sets V and E are respectively identified with the vertex set, $V(\Theta)$, and edge set $E(\Theta)$, of Θ . We refer to Θ as the *underlying graph*.

The following special case will be used later.

Lemma 3.1. *Suppose that Φ is a graph of surfaces such that $\Theta(\Phi)$ is a tree, and each component of Π is homeomorphic either to \mathbb{S}^2 or to \mathbb{R}^2 . Then $M_0 = M_0(\Phi)$ is homeomorphic to an open subset of the closed 3-ball, B , with its interior, $\text{int}(M_0)$, identified with $\text{int}(B)$.*

In particular, it follows that $M = M(\Phi)$ is simply connected, and that $\partial^+ M$ is a planar surface.

Lemma 3.1 can be deduced in a number of ways. Note that if Π were \mathbb{S}^2 then M_0 would be a 3-ball; and if Π were \mathbb{R}^2 then M_0 would be $\mathbb{R}^2 \times [0, 1]$. In general M_0 will be the result of connecting together spaces of this form by handles attached to their boundaries in treelike fashion. It's not hard to see that the result will satisfy the conclusion of the lemma. Another way of viewing this in terms of hyperbolic geometry will be discussed in Section 5.

In fact, we can say a bit more. Note that the inclusion $\partial^+ M \hookrightarrow M$ induces a homeomorphism of the respective spaces of ends. Each end either corresponds to the end of some $\Pi(v)$ which is homeomorphic to \mathbb{R}^2 , or to an ideal boundary point of the tree Θ . In particular, if we assume that every vertex of Θ has degree at least 3 and those for which $\Pi(v) \cong \mathbb{R}^2$ all have infinite degree, then there are no isolated ends. It follows that the space of ends in this case will be a Cantor set, and so $\partial^+ M$ is the Cantor surface.

We want to generalise the above to 2-orbifolds.

Definition. A *graph of orbifolds* is a connected 2 complex,

$$\Phi = \left(\bigsqcup_{v \in V} \Pi(v) \right) \cup \left(\bigsqcup_{e \in E} I(e) \right),$$

where each $\Pi(v)$ a (good) connected 2-orbifold (without boundary) and where each $I(e)$ is a compact real interval with each of its endpoints in one of the $\Pi(v)$. Associated to each $I(e)$ is a natural number $p(e) \geq 1$. If $p(e) = 1$, then neither endpoint of $I(e)$ is a cone point of the orbifold it lies in. If $p(e) \geq 2$, then each endpoint of $I(e)$ is a cone point of degree $p(e)$. We also assume that no $\Pi(v)$ is of type (E0) (a plane with at most one cone point).

As before, we write $\Pi = \bigcup_{v \in V} \Pi(v)$ and $\tau = \bigcup_{e \in E} I(e)$. We write Θ for the underlying graph so that $V(\Theta) \equiv V$ and $E(\Theta) \equiv E$. Note that by taking underlying surfaces, we get a graph of surfaces as before. We write $E = E_1 \sqcup E_2$ where $E_1 = \{e \in E \mid p(e) = 1\}$ and $E_2 = \{e \in E \mid p(e) \geq 2\}$.

Most of the earlier discussion goes through. In particular, we construct $M = M(\Phi) \supseteq \Phi$, so that $\partial^- M$ is identified with Π . We write $C \subseteq \Pi$ for its set of cone points.

In fact, M has a natural structure as a simple 3-orbifold. Its singular set consists of those intervals $I(e)$ for $e \in E_2$ (with both endpoints in $\partial^- M$), together with an interval, $I(a)$, for each cone point $a \in C \setminus \tau$. The interval $I(a)$ has degree $p(a)$. It has one endpoint at $a \in \partial^- M$, and the other in $\partial^+ M$. This gives $\partial^+ M$ the structure of a 2-orbifold with these cone points.

As before, M retracts onto Φ which further retracts onto the the graph Θ . We can assume that the map $M \rightarrow \Theta$ collapses handles onto edges, so that the preimage of the midpoint of any edge e is a disc, $\Delta(e)$, in M , meeting $\tau \cup \sigma(M)$ transversely in a single point of $I(e)$. We can think of $\Delta(e)$ as a 2-orbifold with boundary $\gamma(e) = \partial\Delta(e) = \Delta(e) \cap \partial M \subseteq \partial^- M$, and with a single cone point of order $p(e)$ at this intersection (or no cone point if $p(e) = 1$). We refer to the family $(\Delta(e))_{e \in E}$ as a system of *dual discs*. We write $G(e) = \pi_1^o(\Delta(e)) \cong \mathbb{Z}_{p(e)}$.

We write $\Delta = \bigcup_{e \in E} \Delta(e)$. Note that the closure of $M \setminus \Delta$ in M is a (disconnected) orbifold isomorphic to $\Pi \times [-1, 1]$. In particular, each component has the form $\Pi \times [-1, 1]$, for some $v \in V$. We write $G(v) = \pi_1^o(\Pi(v))$ for its orbifold fundamental group. Note that if e is incident on v , then there is a monomorphism from $G(e)$ to $G(v)$. (Here $G(e)$ corresponds to a cone point of $\Pi(v)$, as discussed in Section 2.) In this way, the family of groups, $(G(v))_{v \in V}$ and $(G(e))_{e \in E}$, together with the above monomorphisms, has the structure of a graph of groups, with underlying graph Θ . The orbifold version of the van-Kampen theorem tells us that its fundamental group is precisely $G = G(\Phi) = \pi_1^o(M(\Phi))$. (See Section 4 for more explanation of this in the present context.)

In terms of Bass-Serre theory, this gives rise to a simplicial action of G on a tree, T , with quotient $\Theta = T/G$. We write $\pi : T \rightarrow \Theta$ for the quotient map.

To each $v \in V(T)$ we can associate a copy of the universal cover, $P(v)$, of $\Pi(\pi v)$, together with an action of the stabiliser, $G(v) \cong G(\pi v)$, of v . Note that G permutes the $P(v)$, so performing these constructions equivariantly, we get a properly discontinuous action of G in $P = \bigsqcup_{v \in V(T)} P(v)$, with $\Pi = P/G$. We can extend this to an action on a tree of surfaces, $F = P \cup t$, with underlying graph T , where $t = \bigsqcup_{e \in E(T)} I(e)$, and with $F/G = \Phi$. Note that each component of P is either \mathbb{S}^2 or \mathbb{R}^2 .

The construction of the manifold $M(F)$ can now be carried out G -equivariantly, with orbifold quotient $M(\Phi) = M(F)/G$. In particular, it follows that $M(\Phi)$ is a good orbifold. Moreover, by Lemma 3.1, $M(F)$ is simply connected and $S = \partial^+ M(F)$ is planar. In fact, from the subsequent discussion, we see that S is the Cantor surface, unless $\tau = \emptyset$ and $\Phi = \Pi$. In the latter case, S is \mathbb{S}^2 or \mathbb{R}^2 .

In summary, we have shown:

Lemma 3.2. *Let $\Phi = \Pi \cup \tau$ be a graph of orbifolds. Then $M(\Phi)$ is a good orbifold whose universal cover has a planar boundary component, S , with $S/G(\Phi) = \partial^+ M(D)$. Moreover, S is homeomorphic to \mathbb{R}^2 , \mathbb{S}^2 , an annulus, or a Cantor surface.*

Definition. We will say that a graph of orbifolds $\Phi = \Pi \cup \tau$ is of *finite type* if $|V|, |E| < \infty$ and each component of Π is a finite-type orbifold.

In particular, we can decompose the vertex set, $V = V_S \sqcup V_E \sqcup V_H$ and $V_E = V_E^1 \sqcup V_E^2$, according to the type of $\Pi(v)$ in the classification in Section 2.

Note again that the ends of Φ , Π , M and $\partial^+ M$ are all in natural bijective correspondence. In particular, this gives a family of *peripheral* subgroups, \mathcal{H} , of $G = G(\Phi)$. Each element of \mathcal{H} is a maximal subgroup isomorphic to \mathbb{Z} , and conjugate into exactly one of the vertex groups, $G(v)$. In other words, we have a splitting of G as a finite graph of groups, over finite cyclic subgroups, relative to the family, \mathcal{H} .

We will describe in the next section how Φ can be recovered from such data.

4. DETERMINING THE ACTION

This section is somewhat orthogonal to the rest of the discussion. We consider how various aspects of finite-type graphs of orbifolds are determined by combinatorial data.

To this end, we view two graphs of orbifolds as isomorphic if there is a homeomorphism between them which respects the orbifold structure on each of the 2-orbifold pieces, and hence also the degrees of the edges. Clearly the thickened-up manifolds we construct from them will then be isomorphic as orbifolds.

Suppose that $\Phi = \Pi \cup \tau$ is a finite-type graph of orbifolds. This is completely determined (up to isomorphism) by a finite set of data. This consists of the underlying graph, Θ , a degree $p(e)$ for each $e \in E(\Theta)$, and to each $v \in V(\Theta)$ we have a genus, a number of ends, and a set of degrees of those cone points

(counting multiplicities) which are not contained in any of the edges $I(e)$. From this information, we can reconstruct each orbifold $\Pi(v)$. (Its cone degrees are those specified in the data, together with the values of $p(e)$ for $e \in E_2$ incident on v .) We can then connect these orbifolds together by intervals as specified by Θ . This gives us our graph of orbifolds, Φ .

There are a few restrictions on the possible data, which we mention below.

The following will be convenient to simplify the discussion.

Definition. We say that a vertex $v \in V$ is *useless* if it has valence at most 2 in Θ and $\Pi(v)$ is a sphere with at most two cone points.

If v has valence 1 in Θ , then $\Pi(v) \cong \mathbb{S}^2$, and we remove it together with the incident edge. If it has valence 2 then we remove it and join together the incident edges. (Since we are assuming that $\Pi(v)$ is good, these edges have the same degree.) The effect on $M(\Phi)$ is to cap off a boundary component with a ball, possibly quotiented out by a cyclic action. Both these operations preserve $G(\Phi)$ and $\partial^+ M(\Phi)$. There is a special case where $V = \{v\}$. In this case, $M(\Phi)$ is topologically $\mathbb{S}^2 \times \mathbb{S}^1$, possibly with a circular singular set (an \mathbb{S}^1 -factor). Then $G(\Phi)$ is finite cyclic and $\mathcal{H} = \emptyset$. After iterating these operations we can eliminate useless vertices, or possibly end up in the special case mentioned above. The special case can be easily understood explicitly.

For this reason, we will assume:

(S1): There are no useless vertices.

In this case, any collection of data will give rise to a graph of orbifolds with the following restriction. If the genus at any $v \in V$ is 0, then the number of ends plus the number of incident edges there must be least 3.

We next explain how much of this data can be recovered from the group, G , together with its peripheral structure, \mathcal{H} . To this end, recall that G acts on a Bass-Serre tree, T , with quotient graph, Θ . The action is “2-acylindrical” meaning that if $e, e' \in E(T)$, then $G(e) \cap G(e')$ is trivial. It is enough to verify this when e, e' are incident on a vertex $v \in V(T)$. But this follows from the earlier discussion on orbifolds. We can assume that $G(e)$ and $G(e')$ are each non-trivial. They are therefore subgroups associated to cone points, and the family of such subgroups in $G(v)$ form a malnormal family. (Here we are using the assumption that v is not useless.) In particular, $G(e) \cap G(e')$ is trivial as claimed. We have also noted that each peripheral subgroup, $H \in \mathcal{H}$, lies in a unique vertex stabiliser $G(v)$. We also note (again from the orbifold description) that no $G(v)$ splits relative to the collection of its peripheral subgroups and incident edge groups.

To proceed, it is convenient to consider first the following case.

Definition. We say that Φ is *indecomposable* if $E_1 = \emptyset$.

In other words, $p(e) \geq 2$ for all $e \in E$. This means that all the edge stabilisers are non-trivial. It is equivalent to saying that G does not split as a free product relative to \mathcal{H} .

In this case (from 2-acylindricity) two vertices $v, w \in V(T)$ will be adjacent if and only if $G(v) \cap G(w)$ is non-trivial. Indeed we can completely recover the graph-of-groups structure as the maximal splitting over finite (cyclic) groups relative to the peripheral subgroups.

To see this, suppose we have two splittings with the above properties, with Bass-Serre trees, T and T' respectively. If $v \in V(T)$, then since $G(v)$ does not split, it must fix some vertex $v' \in V(T')$. For the same reason, $G(v')$ fixes a vertex $v'' \in V(T)$. Therefore $v = v''$ and $G(v) = G(v')$. This gives a natural bijection from $V(T)$ to $V(T')$, preserving stabilisers. By an earlier observation this must preserve adjacency. In other words, we get a G -equivariant isomorphism from T to T' , as required.

As observed in Section 2, the group $G(v)$ together with its peripheral structure determines $\Pi(v)$ up to orbifold isomorphism. Fitting these together as determined by Θ and the edge-degrees, we can completely recover Φ .

In summary we have shown:

Lemma 4.1. *An indecomposable finite-type graph of orbifolds, Φ , satisfying (S1) is completely determined up to isomorphism by the group $G(\Phi)$ together with its peripheral structure.*

In fact, we can say more than this, as we discuss shortly. First, we discuss the decomposable case.

Let Φ be a graph of orbifolds. Let $\tau_0 = \bigcup_{e \in E_2} I(e)$.

Definition. An (indecomposable) factor of Φ is a connected component of $\Pi \cup \tau_0$.

In other words Φ consists of a finite number of factors, Φ_1, \dots, Φ_k , connected together by a finite number of degree-1 intervals. In particular, it follows that $G = G(\Phi)$ splits as a free product $(\ast_{i=1}^k G_i) \ast F_m$ relative to \mathcal{H} , where $G_i = G(\Phi_i)$, and F_m is a free group of rank m not containing any element of \mathcal{H} . (Possibly $m = 0$.) Moreover, none of the G_i splits as a free product relative to \mathcal{H} . It follows by the relative version of the Grushko decomposition theorem that each of the factors G_i is determined up to conjugacy in G , by the pair (G, \mathcal{H}) . So also is the rank m .

We can perform certain sliding operations. In the graph Θ , we can slide an endpoint of an edge $e \in E_1$, across another edge. In Φ , we slide an endpoint of $I(e)$ from one component of Π to another. Note that this operation corresponds to a handle slide on $M(\Phi)$, so its topological type does not change. In the process, we can also remove any useless vertices that arise. This does not change $G(\Phi)$ nor $\partial^+ M(\Phi)$. After a finite number of such steps we can eliminate all (non-singular) sphere components of Π . Unless that is, we end up with just one sphere component, with a number, m , of intervals attached at both endpoints. In that

case, after capping off this sphere, M becomes a handlebody, M_0 , with $G = \pi_1(M_0) = \pi_1^o(M_0) \cong F_m$, and with $\mathcal{H} = \emptyset$. We can view this as a special case.

It is therefore convenient to assume, in addition to (S1), that:

(S2): no component of Π is a 2-sphere with no cone points.

Now, after sliding handles, we can assume that $M(\Phi)$ consists of a handlebody of genus $m \geq 0$, connected to each $M(\Phi_i)$ by a handle attached to $\partial^+ M(\Phi_i)$.

We have observed that each G_i is determined by (G, \mathcal{H}) . By Lemma 4.1, G_i in turn determines Φ_i hence $M(\Phi_i)$. Also, m is determined, and so finally is $M(\Phi)$.

In summary, we have shown:

Lemma 4.2. *Let G be a group of the form $G = G(\Phi)$, where Φ is a finite-type graph of orbifolds satisfying (S1) and (S2), and let \mathcal{H} be its peripheral structure. Then the orbifold $M(\Phi)$ is completely determined up to isomorphism by the pair (G, \mathcal{H}) .*

Returning to the case of indecomposable graphs, we can say a bit more.

Let Φ be a finite-type indecomposable graph of orbifolds satisfying (S1). Write $\Sigma = \partial^+ M$, and let $C \subseteq \Sigma$ be its set of cone points. Let $(\Delta(e))_{e \in E}$ be a family of dual discs as defined in Section 3. This is determined up to an isomorphism of $M(\Phi)$ by the fact that each $\Delta(e)$ intersects $I(e)$ exactly once. This comes from the fact that removing an open regular neighbourhood of τ we get a copy of $\Sigma \times [-1, 1]$.

In fact, we can make a stronger statement. Let $\gamma(e) = \partial\Delta(e)$. This determines an order- $p(e)$ element $g(e) \in G$, well defined up to inverse and conjugacy. We claim:

Proposition 4.3. *Let $e_0 \in E$. Suppose $\beta \subseteq \Sigma \setminus C$ be a simple closed curve representing $g(e_0)$ (up to inverse and conjugacy). Then β is isotopic to $\gamma(e_0)$ in $\Sigma \setminus C$.*

From the earlier discussion, the conclusion is equivalent to asserting that β bounds a dual disc to $I(e)$.

To prove Proposition 4.3, let $\Delta = \bigcup_{e \in E} \Delta(e)$, and let $\gamma = \Delta \cap \Sigma = \bigcup_{e \in E} \gamma(e)$. Given $v \in V$, let $\Psi(v) \subseteq M$ be the closure of the component of $M \setminus \Delta$ corresponding to v . (Thus, $\Psi(v)$ is orbifold-isomorphic to $\Pi(v) \times [-1, 1]$.) Let $\Upsilon(v) = \Sigma \cap \Psi(v)$. Thus, $\Upsilon(v)$ is the closure of a component of $\Sigma \setminus \gamma$.

Recall the construction of \hat{M} from Section 2. This is a 3-complex, with a natural map, $\pi : \hat{M} \rightarrow M$, inducing an isomorphism from $\pi_1(\hat{M})$ to $\pi_1^o(M)$. Let $\hat{\Sigma} = \pi^{-1}\Sigma$. Then the induced map $\pi_1(\hat{\Sigma}) \rightarrow \pi_1^o(\Sigma)$ is also an isomorphism. We can assume that $\pi|_{\pi^{-1}\beta}$ is a homeomorphism, so that $\pi^{-1}\beta$ is also a curve in $\hat{\Sigma}$, which we will also denote by β . Let $\hat{\Delta}(e) = \pi^{-1}\Delta(e)$, $\hat{\Delta} = \pi^{-1}\Delta$, $\hat{\Psi}(v) = \pi^{-1}\Psi(v)$ and $\hat{\Upsilon}(v) = \pi^{-1}\Upsilon(v)$. Thus, $\hat{\Psi}(v)$ is the closure of a component of $\hat{M} \setminus \hat{\Delta}$, and is

homeomorphic to $\hat{\Pi}(v) \times [-1, 1]$. If e is incident on v , then the inclusion of $\hat{\Delta}(e)$ into $\hat{\Psi}(v)$ is injective on π_1 .

We now assume that β is in general position with respect to γ and realised in its homotopy class so that $|\beta \cap \gamma|$ is minimal. This means that no arc of β can bound a bigon in $\Sigma \setminus C$ with any arc of γ .

We claim that $\beta \cap \gamma = \emptyset$.

To see this, note that the $p(e)$ -fold cover of β in \hat{M} is trivial in $\pi_1(\hat{M}) \equiv \pi_1^o(M)$. Therefore there is a map, $f : D \rightarrow \hat{M}$, of the disc, D , such that $f|_{\partial D}$ wraps $p(e)$ times around β . We can assume that $f^{-1}\hat{\Sigma} = \partial D$, and that f is in general position with respect to $\hat{\Delta}$. In particular, $f^{-1}\hat{\Delta}$ is a collection of disjoint arcs and circles. In fact, by the π_1 -injectivity of the maps $\hat{\Delta}(e) \rightarrow \hat{M}$, we can assume that there are no circles.

Suppose for contradiction that $\beta \cap \gamma \neq \emptyset$, so that $f^{-1}\hat{\Delta} \neq \emptyset$. Let $B \subseteq \Delta$ be an outermost component of $D \setminus f^{-1}\hat{\Delta}$. This is a bigon with boundary $\alpha \cup \delta$, where $\alpha \subseteq f^{-1}\hat{\Delta}$ and $\delta \subseteq \partial D$ are arcs. Note that $f(B) \subseteq \hat{\Psi}(v)$ for some $v \in V$. Clearly, $f|_{\delta}$ is injective, so $f(\delta) \subseteq \beta$ is an arc with endpoints in $\gamma(e)$ for some $e \in E$ incident on v . Also $f(\alpha) \subseteq \hat{\Delta}(e)$. We can therefore push $f(\alpha)$ out to the boundary of $\hat{\Delta}(e)$, fixing its endpoints, so as to give us a path ϵ in $\gamma(e)$, and such that $\epsilon \cup f(\delta)$ bound a disc in $\hat{\Psi}(v)$. (We don't know a-priori that ϵ is an arc.) Since $\hat{\Psi}(v) \cong \hat{\Pi}(v) \times [-1, 1]$, $\epsilon \cup f(\delta)$ is also trivial in $\pi_1(\hat{\Pi}(v)) \cong \pi_1^o(\Pi(v))$. But this implies that $f(\delta)$ must bound a bigon with $\gamma(e)$ (so that we can indeed take ϵ to be an arc). (To see this, note that $\pi(f(\delta))$ gives us an injective loop at a cone point of $\Pi(v)$, which represents an element in the corresponding cyclic subgroup of $\pi_1^o(G(v))$.) But now this gives us a contradiction, proving the claim that $\beta \cap \gamma = \emptyset$.

Therefore $\beta \subseteq \Upsilon(v) \subseteq \Pi(v) \times [-1, 1]$ for some $v \in V$. Since β is conjugate to $\gamma(e_0)$ in G , it follows that v is incident on e_0 . (This follows from the fact, used earlier, that the action on the Bass-Serre tree is 2-acylindrical.) Since $\Pi(v)$ is a 2-orbifold, it now follows that β is isotopic to γ in the complement of its cone points, hence also in $\Sigma \setminus C$.

This proves Proposition 4.3.

5. KLEINIAN GROUPS

We describe how one can realise the 3-orbifolds discussed earlier as quotients of kleinian groups.

We elaborate on some basic definitions mentioned in Section 1. By a *kleinian group* we mean a group, G , acting properly discontinuously by isometry on hyperbolic 3-space \mathbb{H}^3 . We write $M = M(G) = \mathbb{H}^3/G$, for the quotient orbifold. The action extends to the ideal boundary, $\partial\mathbb{H}^3$. We partition $\partial\mathbb{H}^3$ into the limit set, ΛG , and the discontinuity domain, ΩG . In fact, G acts properly discontinuously on $\mathbb{H}^3 \cup \Omega G$, and we write $M_C = M_C(G) = (\mathbb{H}^3 \cup \Omega G)/G$. This is a 3-orbifold with

boundary. Note that we can identify $G \equiv \pi_1^o(M) \equiv \pi_1^o(M_C)$. The action is geometrically finite if M_C has a finite number of ends, each isomorphic to a standard parabolic cusp region [Mar]. In particular, in this case, the peripheral subgroups of G (identified with $\pi_1^o(M_C)$) are precisely the maximal parabolic subgroups. A cusp can be “rank-1” or “rank-2” depending on whether the maximal parabolic subgroup is virtually \mathbb{Z} or virtually \mathbb{Z}^2 .

Let Φ be a finite-type graph of orbifolds. We construct a larger manifold, $M^+(\Phi) \supseteq M(\Phi)$ as follows. To ∂^+M , we glue a copy of $\partial^+M \times [0, \infty)$; and to each component, $\Pi(v)$, of ∂^-M we glue a copy of the space $\Pi^+(v)$ defined as follows. If $v \in V_E \cup V_H$, we set $\Pi^+(v) = \Pi(v) \times [0, \infty)$. If $v \in V_S$, we take $\Pi^+(v)$ to be a cone over $\Pi(v)$: in other words a compact 3-orbifold with underlying space a 3-ball, and with boundary, $\Pi(v)$.

We can construct an even bigger 3-orbifold, $M_C^+(\Phi)$, by gluing in $\partial^+M(\Phi) \times [0, \infty)$ to $\partial^+M(\Phi)$, and a copy of $\Pi_C^+(v)$ to each $\Pi(v)$. Here $\Pi_C^+(v) = \Pi^+(v)$ if $v \in V_S \cup V_E$, and $\Pi_C^+(v) = \Pi(v) \times [0, \infty)$ if $v \in V_H$. We again refer to “outer” and “inner” boundaries of $M_C^+(\Phi)$.

Note that the ends of M_C^+ are of two sorts. They can be “rank-2”: that is they correspond to ends, $\Pi^+(v)$, for $v \in V_E^2$. Or else they can be “rank-1”: they correspond to ends of $\Pi_C^+(v)$ for $v \in V_H \cup V_E^1$. When $v \in V_H$, these ends have a neighbourhood of the form $\mathbb{S}^1 \times [-1, 1] \times [0, \infty)$, with $\mathbb{S}^1 \times \{-1\} \times [0, \infty)$ and $\mathbb{S}^1 \times \{1\} \times [0, \infty)$ respectively ends of the inner and outer boundaries. When $v \in V_E^1$, either $\Pi(v)$ is an annulus, and the end again has the form $\mathbb{S}^1 \times [-1, 1] \times [0, \infty)$, though now the sets $\mathbb{S}^1 \times \{\pm 1\} \times [0, \infty)$ both lie in the outer boundary; or else $\Pi(v)$ is a disc with two degree-2 cone points, and we get a quotient of the above by a \mathbb{Z}_2 action.

We claim:

Proposition 5.1. *Let Φ be a finite-type graph of orbifolds. Then there is a geometrically finite action of $G = G(\Phi)$ on \mathbb{H}^3 having the following properties. There is a homeomorphism of $M_C^+(\Phi)$ to $M_C(G)$, inducing the identity on G (as the orbifold fundamental groups). This restricts to a homeomorphism of $M^+(\Phi)$ to $M(G)$. Moreover, we can take the subgroup corresponding to each inner boundary component to be fuchsian (i.e. preserving a plane in \mathbb{H}^3). The rank-1 and rank-2 ends of $M_C^+(\Phi)$ correspond respectively to the rank-1 and rank-2 maximal parabolic subgroups.*

Note that since $G = \pi_1^o(M_C^+(\Phi))$ is supported on ∂^+M , the corresponding component, S , of ΩG is G -invariant. (Such a kleinian group is traditionally called a “function group”.) All other components of ΩG are round discs (since their stabilisers are fuchsian). In the terminology of [Mas2], G is a “Koebe group”. Note that ∂^+M can be naturally identified with S/G .

We remark that an equivalent formulation of Proposition 5.1 could be obtained by replacing $M_C(G)$ with the convex core or M : in other words the quotient by G of the convex hull if the limit set, ΛG . Topologically, this can be obtained from

$M_C(G)$ by removing a regular neighbourhood of the boundary. (The case where G is fuchsian is “degenerate” in the sense that the core collapses onto a totally geodesic subspace.)

We now set about the proof of Proposition 5.1. We have seen in Section 2 that any 2-orbifold carries a certain geometric structure. From this we can construct an action of each group $G(v)$ on \mathbb{H}^3 as described below. To construct the action of G on \mathbb{H}^3 , we need to piece these actions together in a nice way. There are various standard “combination” theorems which allow us to do this. The argument in [Mas2] primarily makes use of the actions on $\partial\mathbb{H}^3$. Here we will use the geometry of \mathbb{H}^3 instead.

To begin, we note that we can put a complete singular hyperbolic metric on each $\Pi^+(v)$ with convex boundary, $\Pi(v)$. In each case we realise $\Pi^+(v)$ as the quotient of a closed convex subset, $P^+(v)$, in \mathbb{H}^3 , by an isometric action of the group, $G(v)$. (For the moment, $P^+(v)$ is only defined up to isometry.) More precisely, if $v \in V_S$, we take $P^+(v)$ to be a closed ball in \mathbb{H}^3 of fixed radius; if $v \in V_E$, we take $P^+(v)$ to be a closed horoball; and if $v \in V_H$, we take $P^+(v)$ to be closed half space. In the induced path metric, $\Pi(v) = \partial P^+(v)$ is respectively, \mathbb{S}^2 , \mathbb{R}^2 and \mathbb{H}^2 . Thus, $\Pi(v) = \partial\Pi^+(v)$ is respectively a spherical, euclidean or hyperbolic orbifold. In the last case, we arrange for it to have finite area. In the euclidean and hyperbolic cases, this involves an arbitrary choice of geometric structure, though the choice is not relevant to the present discussion.

Now let $\Phi^+ \supseteq \Phi$ be the space obtained from Φ by gluing a copy of $\Pi^+(v)$ to $\Pi(v)$ for each $v \in V$. We can put a metric on Φ^+ by giving each $\Pi^+(v)$ the metric described in the previous paragraph. We choose some constant $L \geq 0$ big enough, as required by Lemma 5.2 below, and give each interval $I(e)$ the structure of a real interval of length L . We can then take the induced path metric on Φ^+ .

Recall that we can realise Φ as a quotient F/G as described in Section 3. We can similarly write $\Phi^+ = F^+/G$, where $F^+ \supseteq F$ is obtained by gluing a copy of $P^+(v)$ to $P(v)$ for each $v \in V(T)$. Again, we can give F^+ a metric as described above, so that G acts isometrically, and inducing the given metric on Φ^+ .

In fact, we can endow F^+ with some extra structure. Namely, for each $e \in E(T)$, we choose an (isometric) identification of the tangent spaces to $P(v)$ and $P(w)$ at the endpoints of $I(e)$, where $v, w \in V(T)$ are the endpoints of e . Moreover, we can do this in G -equivariant fashion.

Finally note that (by equivariance) there is some $\eta > 0$, such that if e and e' are distinct edges incident on v , then the endpoints of $I(e)$ and $I(e')$ are distance at least η apart in $P(v)$.

We now claim:

Lemma 5.2. *Suppose L is chosen sufficiently large as a function of $\eta > 0$. Then we can realise F^+ as a closed subset $F^+ \subseteq \mathbb{H}^3$, inducing the original path metric on F^+ , and such that if e is incident on v then $I(e)$ meets $P^+(v)$ orthogonally.*

Moreover, if e is also incident on w , then the specified identification of tangent spaces to $P(v)$ and $P(w)$ is given by parallel transport in \mathbb{H}^3 along $I(e)$.

In fact, all we require of the sets $P^+(v)$ for this to hold is that they be isometric to closed convex subsets of \mathbb{H}^3 with smooth boundary. (The action of G is not relevant to the statement of the lemma.)

It's not hard to see that the embedding of F^+ as described by the conclusion of Lemma 5.2 is unique up to isometry of \mathbb{H}^3 . Therefore the action of G on F^+ extends to an action on \mathbb{H}^3 . We will see that it satisfies the conclusion of Proposition 5.1. First, we explain why Lemma 5.2 holds.

We will proceed by constructing a "thickening" of F^+ which is isometric to a closed convex subset of \mathbb{H}^3 .

We begin with the following lemma. Suppose $B \subseteq \mathbb{H}^3$ is a (round) ball, horoball or half-space. Suppose that $I \subseteq \mathbb{H}^3$ is a geodesic segment meeting B orthogonally at one of its endpoints. We write $N(B, r)$ and $N(I, s)$ respectively for the closed r - and s -neighbourhoods. We will assume that $\text{length}(I) > r$, so that I meets $\partial N(B, r)$ orthogonally at a single point, x . Let H be the closed convex hull of $N(B, r) \cup N(I, s)$, and let $\Xi = H \setminus (N(B, r) \cup N(I, s))$.

Lemma 5.3. *Given any $\eta > 0$, there is some $t = t(\eta)$, such that if $r, s \geq t$ and B, I, x, Ξ are as above, then $\Xi \subseteq N(x, t)$.*

Note that the picture has circular symmetry: rotating about the bi-infinite geodesic containing I . It is therefore sufficient to verify the corresponding statement in \mathbb{H}^2 , which is a simple exercise. We see in fact, that Ξ is a solid torus glued in at the junction of $N(B, r)$ and $N(I, s)$.

To explain what we are aiming for in the proof of Lemma 5.2, let us suppose for the moment that we have already embedded F^+ into \mathbb{H}^3 . Let $t = t(\eta)$. Choose $r > t$ sufficiently large as described below and assume that $L > 2r + 2t$. Given $v \in V(T)$ incident on $e \in E(T)$, we write $x(v, e)$ for the point of intersection of $\partial N(P^+(v), r)$ with $I(e)$. Let $\Xi(v, e)$ be the set featuring in Lemma 5.3, where $s = \eta$, $B = P^+(v)$ and $I = I(e)$. Therefore, $\Xi(v, e) \subseteq N(x(v, e), t)$.

Given the uniform divergence of geodesics in hyperbolic space, we can choose r big enough so that if $e, e' \in E(T)$ are distinct edges incident on v , then the distance between $x(v, e)$ and $x(v, e')$ is greater than $2t$. This implies that $\Xi(v, e) \cap \Xi(v, e') = \emptyset$. Since $L > 2r + 2t$, it also follows that $\Xi(v, e) \cap \Xi(w, e'') = \emptyset$ for any $w \neq v$ and any e'' incident on w . Moreover, if $\Xi(v, e) \cap N(P^+(w), r) \neq \emptyset$, then $w = v$; and if $\Xi(v, e) \cap N(I(e'), \eta) \neq \emptyset$, then $e' = e$.

We now set

$$Q = \bigcup_{v \in V(T)} N(P^+(v), r) \cup \bigcup_{e \in E(T)} N(I(e), \eta) \cup \bigcup_{v, e} \Xi(v, e),$$

where the last union is taken over all $v \in V(T)$ and $e \in E(T)$, with e incident on v . (We have already observed that this last term is a disjoint union.) Note that

the inclusions, $F \hookrightarrow F^+ \hookrightarrow Q$ are both homotopy equivalences. In particular, Q is simply connected.

Of course, this presupposes that $F^+ \subseteq \mathbb{H}^3$. However, if we are given F^+ abstractly, together with the associated identifications of tangent spaces, we can construct Q directly, just by constructing the component pieces, and then gluing them together. (Note that the gluing just involves attaching a set of the form $\Xi(v, e)$ to an incident set of the form $N(P^+(v), r)$ or $N(I(e), \eta)$, and that the gluings are all disjoint.) By construction, Q is simply connected, and locally isometric to a closed convex subset of \mathbb{H}^3 . It therefore follows that Q is globally isometric to a convex subset of \mathbb{H}^3 . In other words, we have embedded Q , hence also F^+ into \mathbb{H}^3 , thereby proving Lemma 5.2.

In fact, we see by construction that the interior of the convex set, Q , is G -equivariantly homeomorphic to the universal cover of $M^+(\Phi)$, and so $\text{int } Q/G$ is homeomorphic to $M^+(G)$.

Proof of Proposition 5.1. Let $\Phi^+ \supseteq \Phi$ be as constructed above. We can write $\Phi^+ = F^+/G$. By Lemma 5.2 we can properly embed $F^+ \subseteq \mathbb{H}^3$. Since the construction is canonical (once we have fixed some constant L) the action of G on F^+ extends to an isometric action on \mathbb{H}^3 . In fact, we also get a G -equivariant embedding, $Q \subseteq \mathbb{H}^3$. In particular, we see that the action of G on \mathbb{H}^3 is properly discontinuous.

As observed above, $\text{int } Q/G$, is homeomorphic to $M^+(\Phi)$, and we see that its closure in $M_C(G) = (\mathbb{H}^3 \cup \Omega G)/G$ is homeomorphic to $M_C^+(\Phi)$. This is also homeomorphic to $M_C(G)$. By construction the ends are in bijective correspondence to the parabolic cusps. In particular, the complement of open neighbourhoods of the ends of $M_C(G)$ is compact, and so G is geometrically finite. \square

6. THE PLANARITY THEOREM

We begin with a general discussion of planar surfaces.

By a *planar surface* we mean a connected space homeomorphic to an open subset $S \subseteq \mathbb{S}^2$. We will assume for the moment that $S \neq \mathbb{S}^2$. Then the universal cover, W , of S is homeomorphic to \mathbb{R}^2 , and we write $S = W/N$, where $N = \pi_1(S)$.

We write $\mathcal{E} = \mathcal{E}(S)$ for the space of ends of S . This is a totally disconnected compact space, and $S \cup \mathcal{E}$ gives a hausdorff compactification of S . It is not hard to see that $S \cup \mathcal{E}$ is homeomorphic to the quotient space obtained by collapsing each component of $\mathbb{S}^2 \setminus S$ to a point. This is in turn homeomorphic to \mathbb{S}^2 . The latter statement is an old result of Moore [Mo], though of course one could give a number of proofs of this fact today. In view of this, there is no loss in assuming that $\mathbb{S}^2 \setminus S$ is totally disconnected, and thus naturally homeomorphic to \mathcal{E} . If \mathcal{E} is perfect, then it is a Cantor set. There is only one embedding of the Cantor set in \mathbb{S}^2 up to homeomorphism, so in that case, the surface S is unique up to homeomorphism. (This is also a direct consequence of the classification theorem

[R].) We refer to this space as the *Cantor surface*. (It has also been called the “Cantor sphere” or the “Cantor tree surface”.)

Let S be planar. Suppose that $\beta \subseteq S$ is a closed subset which is a disjoint union of essential simple closed curves. We have a dual tree, T_β , where each edge $e \in V(T_\beta)$ corresponds to a component, $\beta(e)$, of β , and each vertex $v \in V(T_\beta)$ corresponds to the closure of a component, $Y(v)$, of $\mathbb{S}^2 \setminus S$. Here e is incident on v if and only if $\beta(e) \subseteq Y(v)$. In this way, we can represent $N = \pi_1(S)$ as a graph of groups with underlying graph T_β . Each edge group, $N_\beta(e)$, is infinite cyclic and supported on the curve $\beta(e)$. Each vertex group, $N_\beta(v)$, has the form $\pi_1(Y(v))$.

Let $N_\beta \triangleleft N$ be the (normal) subgroup of N generated by the family all of edge groups, $N_\beta(e)$. Then N/N_β is also represented by a graph of groups with underlying graph T_β , with trivial edge-groups, and with vertex groups of the form $N_\beta(v)/J_\beta(v)$, where $J_\beta(v) \triangleleft N_\beta(v)$ is the normal subgroup of $N_\beta(v)$ generated by the incident edge groups. (Topologically, this is equivalent to gluing a disc to each component of β , and then taking the fundamental group.) Note that $N_\beta(v)/J_\beta(v) \cong \pi_1(Z(v))$, where $Z(v)$ is the surface obtained by gluing a disc to each boundary component of $Y(v)$. It is easily seen that the following statements are all equivalent: $Z(v) \cong \mathbb{R}^2$; $Z(v)$ is one-ended; $Y(v)$ is one-ended; $\pi_1(Z(v))$ is trivial; and $J_\beta(v) = N_\beta(v)$.

Now N/N_β is isomorphic to the free product $\cong *_{v \in V(T_\beta)} (N_\beta(v)/J_\beta(v))$. It follows that N/N_β is trivial if and only if $N_\beta(v)/J_\beta(v)$ is trivial for all $v \in V(T_\beta)$.

In particular, we have shown:

Lemma 6.1. *$N_\beta = N$ if and only if $Y(v)$ is one-ended for all $v \in V(T_\beta)$.*

Now suppose that G acts properly discontinuously on a planar surface S with quotient orbifold, $\Sigma = S/G$. We write $\pi : S \rightarrow \Sigma$ for the quotient map, and $C \subseteq \Sigma$ for the set of cone points.

Write $S = W/N$ and $\Sigma = W/\Gamma$, where $\Gamma = \pi_1^o(\Sigma)$ and W is the orbifold universal cover. Thus, $N \triangleleft \Gamma$ and $G = \Gamma/N$. In other words, $\Sigma = S/G = (W/N)/(\Gamma/N) = W/\Gamma$.

By a *multicurve* γ in Σ , we mean a closed subset of $\Sigma \setminus C$ which is a topological disjoint union of essential closed curves in $\Sigma \setminus C$. We will assume that no two elements of γ are parallel (that is, bound an annulus in $\Sigma \setminus C$). We write $\tilde{\gamma} = \pi^{-1}\gamma \subseteq S$. We say that γ is *liftable* if each component of $\tilde{\gamma}$ is compact. In other words, each such component maps with some finite degree to some component of γ . This determines a normal subgroup, $N_{\tilde{\gamma}} \triangleleft N$, as described above.

Definition. We say that a liftable multicurve, $\gamma \subseteq \Sigma$, is *complete* if $N_{\tilde{\gamma}} = N$.

Note that any isolated end of S has finite stabiliser in G , and we can compactify it by adding in an isolated point. In the quotient, this corresponds to adjoining a point, possibly a cone point. For this reason, it will be convenient to assume, for the moment, that S has no isolated ends.

Let $\gamma \subseteq \Sigma$ be a liftable multicurve. We take a regular neighbourhood of γ in $\Sigma \setminus C$. This is a disjoint union of annuli, and we collapse each annulus to an interval, by collapsing the \mathbb{S}^1 -factor to a point. The result will be a graph of orbifolds, Φ , with each edge assigned the degree of the covering of the corresponding component of γ . Let $M = M(\Phi)$. We can naturally identify Σ with $\partial^+ M$. The inclusion of Σ into M gives us an epimorphism from $\Gamma = \pi_1^o(\Sigma)$ to $G(\Phi) = \pi_1^o(M)$. Its kernel is precisely $N_{\tilde{\gamma}} \triangleleft N$. If γ is complete then $N_{\tilde{\gamma}} = N$, and so we get a natural identification of $G = \Gamma/N$ with $G(\Phi)$.

In these terms, we can state the Planarity Theorem of Maskit [Mas1] as follows.

Theorem 6.2. *If Σ is of finite type, then it admits a complete liftable multicurve.*

If we assume that S has no isolated ends, then the above discussion applies.

In summary, this shows:

Corollary 6.3. *Suppose that the group, G , acts properly discontinuously on a Cantor surface S , and with finite-type quotient, $\Sigma = S/G$. Then there is a finite-type graph of orbifolds, Φ , with $G \cong G(\Phi) = \pi_1^o(M(\Phi))$, and a G -equivariant identification of S with the outer boundary of the universal cover of M (so that Σ gets identified with $\partial^+ M$).*

Using Lemma 6.1 it can now be seen that the action of G on S extends to a properly discontinuous action of G on a 3-ball with S identified as a subset of its boundary, \mathbb{S}^2 . (Of course, this need not coincide with the original embedding of S in \mathbb{S}^2 .) In fact we can realise this geometrically.

Putting Corollary 6.3 together with Proposition 5.1, we deduce:

Corollary 6.4. *Suppose that the group, G , acts properly discontinuously on a Cantor surface, S , and with finite-type quotient. Then there a geometrically finite action of G on \mathbb{H}^3 , and a G -equivariant homeomorphism of S to a G -invariant component, Ω_0 , of the discontinuity domain $\Omega_G \subseteq \partial\mathbb{H}^3$.*

By construction (or directly applying Ahlfors's Finiteness Theorem) the quotient Ω_0/G is a finite-type Riemann surface. It follows that the ends of $\Sigma \cong \Omega/G$ correspond to parabolic elements of the action of G . Of course, one can say more about the action of G on \mathbb{H}^3 , as described by Proposition 5.1. (As remarked after Proposition 5.1, we could alternatively identify S with a boundary component of the convex hull of the limit set of G , so that S/G gets identified with a boundary component of the convex core of M .)

If we allow for isolated ends in S , then the conclusion of Corollary 6.3 holds, except we would need to remove a finite set of points from $\partial^+ M$, as well as its preimage in the universal cover of M .

As noted in Section 1, a group G of this sort is necessarily finitely generated. In fact, if we assume G to be finitely generated, we can drop the assumption on the quotient being finite type, modulo modifying the surface S . (See [G] for a similar statement.)

Proposition 6.5. *Suppose that G is a finitely generated group acting effectively and properly discontinuously on a planar surface, S . Then there is an open G -invariant subset, $\mathcal{E}_0 \subseteq \mathcal{E}(S)$, of the space of ends of S , such that G also acts properly discontinuously on $S \cup \mathcal{E}_0$ with $(S \cup \mathcal{E}_0)/G$ a finite-type orbifold.*

(Recall that $S \cup \mathcal{E}(S) \cong \mathbb{S}^2$, and $S \cup \mathcal{E}_0 \subseteq S \cup \mathcal{E}(S)$ is open, so $S \cup \mathcal{E}_0$ is a planar surface. In fact, from the earlier discussion it must be \mathbb{S}^2 , \mathbb{R}^2 , an annulus or a Cantor surface.)

To begin with, let $\Sigma_0 \subseteq \Sigma$ be a compact core for Σ : that is a compact submanifold with boundary containing all the cone points in its interior, and such that the inclusion $\Sigma_0 \hookrightarrow \Sigma$ induces an isomorphism on π_1^o .

Such a core can be constructed by taking an equivariant map of a Cayley graph of G into $S \setminus \pi^{-1}C$, thickening it up to a manifold, and then filling in any disc components of the complement so as to give us a closed G -invariant subsurface, $S_0 \subseteq S$, with $\Sigma_0 = S_0/G$ compact.

Now let R be the closure of a component of $S \setminus S_0$. Since S is planar, ∂R is connected. Let $G_R \leq G$ be its stabiliser. Then G_R acts freely on ∂R with quotient a circle. Let $\mathcal{E}(R)$ be the space of ends of R . Note that the action of G extends to $R \cup \mathcal{E}(R)$. There are two cases.

Maybe ∂R is a circle, and G_R is cyclic of order p . In this case $R \cup \mathcal{E}(R)$ is a disc and $\mathcal{E}(R)$ is a clopen subset of $\mathcal{E}(S)$. Either $p = 1$, or else G_R has exactly one fixed point of order $p \geq 2$. Thus $(R \cup \mathcal{E}(R))/G_R$ is a disc with at most one cone point. We set $\mathcal{E}_0(R) = \mathcal{E}(R)$.

Or maybe ∂R is the real line. In this case, both ends of ∂R converge on the same end $x \in \mathcal{E}(R)$. Now $R \cup \mathcal{E}(R)$ is again a disc, with boundary $\partial R \cup \{x\}$. The group G_R is infinite cyclic and acts on the disc with unique fixed point, x . We set $\mathcal{E}_0(R) = \mathcal{E}(R) \setminus \{x\}$. This is an open subset of $\mathcal{E}(S)$. Moreover, $(R \cup \mathcal{E}_0(R))/G_R$ is an annulus (with no cone points).

In view of the above we set \mathcal{E}_0 to be the union of the sets $\mathcal{E}_0(R)$ as R ranges over the closures of all components of $S \setminus S_0$.

By construction, $(S \cup \mathcal{E}_0)/G$ consists of Σ_0 , together with a finite number of discs and annuli glued to the boundary components, and such that each such disc has at most one cone point. This is a finite type orbifold, thereby proving Proposition 6.5.

We also note the following result in [G]:

Theorem 6.6. [G] *Suppose that G is finitely generated group which admits an effective properly discontinuous action in a planar surface, S . Then G also admits such an action on a planar surface, S' , with S'/G compact.*

The argument of [G] uses planarity of Cayley graphs.

One can also give an interpretation in terms of graphs of orbifolds, as we outline below.

First, by Proposition 6.5, we can suppose that S/G is of finite type. We can also suppose that S is a Cantor surface. (Otherwise, it is \mathbb{S}^2 , \mathbb{R}^2 or an annulus, and those cases can be easily dealt with.) Corollary 6.3 now gives us a finite-type graph of orbifolds, $\Phi = \Pi \cup \tau$, with $G = G(\Phi) = \pi_1^o(M(\Phi))$. Let $\Pi_0 \subseteq \Pi$ be a compact core of Π . (In other words we remove disjoint neighbourhoods of the ends of Π so that the inclusion $\text{int } \Pi_0 \hookrightarrow \Pi$ is an equivalence of orbifolds.) Let $\Phi_0 = \Pi_0 \cup \tau \subseteq \Phi$. We can thicken up Φ_0 to give us a 3-orbifold, $M(\Phi_0) \subseteq M(\Phi)$, similarly as with the construction of $M(\Phi)$. In other words, $M(\Phi_0)$ is obtained by removing a neighbourhood of each end of $M(\Phi)$. Note that $\partial\Pi$ gives rise to a disjoint collection of annuli in $\partial M(\Phi_0)$, and that the complement of these annuli is a core of $\partial M(\Phi)$. In particular, one component, Σ' , of $\partial M(\Phi_0)$ contains $\partial M(\Phi) \cap M(\Phi_0)$.

Let $F = P \cup t$ be the covering space of $\Phi = \Pi \cup \tau$ constructed as in Section 4. Thus F is a tree of surfaces, each component of P being either \mathbb{S}^2 or \mathbb{R}^2 . Let $F_0 = P_0 \cup t \subseteq P \cup t = F$ be the preimage of $\Phi_0 = \Pi_0 \cup \tau$, so that $\Phi_0 = F_0/G$. We get a quotient map $M(F_0) \rightarrow M(\Phi_0)$, where $M(F_0) \subseteq M(F)$ is the preimage of $M(\Phi_0)$ in $M(F)$. In fact, $M(F)$ can be constructed from F_0 similarly as with $M(F)$. Let $S' \subseteq \partial M(F_0)$ be the preimage of $\Sigma' \subseteq \partial M(\Phi_0)$. This is a G -invariant component of $\partial M(F_0)$, and $\Sigma' = S'/G$ is compact. In fact, S' is planar. This follows by essentially the same argument as Lemma 3.1: it arises as an infinite connected sum of planar surfaces, each arising from some component of F_0 . This therefore proves Theorem 6.6.

We remark that one can also interpret this in terms of kleinian groups. By Corollary 6.4 and the subsequent remark, we can realise G as kleinian function group in such a way that S is equivariantly identified with the a boundary component of the convex hull of the limit set. We can remove an equivariant family of horoballs from the core. The boundary of the resulting space has a G -invariant component, S' , with S'/G compact. Topologically, this is the same construction as that described above.

7. TRACKS

We explain how to give a direct proof of the Planarity Theorem using tracks.

Let G act effectively and properly discontinuously on a planar surface, S , with quotient $\Sigma = S/G$ a finite-type orbifold. We assume for simplicity that S has no isolated ends. Let $\bar{\Sigma}$ be the compactification of Σ obtained by adjoining the space of ends. Thus, $\bar{\Sigma}$ is topologically a closed surface. We triangulate $\bar{\Sigma}$ so that all cone points and all ends are vertices of the triangulation. We write Λ for the resulting simplicial complex. We can lift this to a simplicial complex, L , with $S \subseteq L$, where the G -action of S extends to a simplicial G -action on L with $\Lambda = L/G$. We write $V_\infty(L) \subseteq V(L)$ for the set of vertices of infinite degree. Thus G acts properly discontinuously on $L \setminus V_\infty(L)$ and freely on $L \setminus V(L)$.

The aim now is to construct a complete liftable multicurve, $\gamma \subseteq \Sigma \subseteq \Lambda$. As observed in Section 6, completeness is equivalent to saying the closure of each component of $\Sigma \setminus \tilde{\gamma}$ is one-ended (see Lemma 6.6).

We now follow the argument in [D].

Suppose that K is either L or Λ . First, we define a “track” in K exactly as in Section 2 of [D]. In our case, this is a properly embedded connected 1-submanifold: either a circle or the real line. Since K is orientable, any track for us will be “untwisted”: that is to say, a regular neighbourhood thereof is a direct product with an interval.

The discussion of Section 2 of [D] now goes through more or less verbatim. Here we should interpret K as our complex Λ . In [D] it was assumed that $H^1(K, \mathbb{Z}_2) = 0$. But this was only used there to show that any track separates. In our case, this follows instead from the fact that $\Lambda \setminus V(\Lambda)$ is a planar surface.

Moving on to Section 3, a “minimal track”, $t \subseteq \Lambda$, is now a circle which cuts Λ into two components, each containing infinitely many vertices of Λ . The latter statement is equivalent to saying that the closure of each component of $\Sigma \setminus t$ is non-compact. The “minimal” clause means that the intersection of t with the 1-skeleton of Λ has minimal cardinality for a track with these properties.

We now proceed to Section 4. Again, $K = \Lambda$. Our planarity assumption again substitutes for the cohomological hypothesis, and the statements go through.

We now follow the proof of Theorem 5.1 of [D], to give us a maximal collection of minimal tracks, whose properties we now interpret.

Our hypotheses are slightly different. For us $K = L$. In [D] it was assumed that the action is free. But all that is really required is that it should be free on $L \setminus V(L)$. Instead of one-endedness of L , we substitute one-endedness of $L \setminus V_\infty(L)$. The negation of this hypothesis is equivalent to saying that there is a track which separates L into two components, each of which contains an infinite subset of $V(L)$. In other words, we can write $V(L) = U \sqcup U^*$, where U and U^* are both infinite, and such that there are only finitely many 1-cells of L with one endpoint in each of U and U^* . It is such a partition which is actually used of multiendedness in [D].

At the end of the day, we arrive at a G -invariant disjoint union, \tilde{t} , of minimal tracks in L , which descends to a finite disjoint union, t , of tracks in Λ . No two components of t are parallel (that is, bound an annulus, that is to say an “untwisted band”, in Λ). Moreover (from the maximality) the closure of each component of $L \setminus \tilde{t}$ in L is one-ended.

We now let $\gamma = t \cup \bigsqcup_{a \in C} \delta(a) \subseteq \Sigma \subseteq \Lambda$, where $\delta(a)$ is the cone curve encircling the cone point $a \in C$. Then γ is a complete liftable multicurve in Σ .

REFERENCES

- [D] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449–457.

- [G] A.Georgakopoulos, *On planar Cayley graphs and Kleinian groups* : Trans. Amer. Math. Soc. **373** (2020) 4649–4684.
- [Mar] A.Marden, *The geometry of finitely generated kleinian groups* : Ann. of Math. **99** (1974) 383–462.
- [Mas1] B.Maskit, *A theorem on planar covering surfaces with applications to 3-manifolds* : Ann. of Math. **81** (1965) 341–355.
- [Mas2] B.Maskit, *Kleinian groups* : Grundlehren der Mathematischen Wissenschaften, Vol. 287, Springer-Verlag (1988).
- [Mo] R.L.Moore, *Concerning upper semi-continuous collections of continua* : Trans. Amer. Math. Soc. **27** (1925) 416–428.
- [R] I.Richards, *On the classification of noncompact surfaces* : Trans. Amer. Math. Soc. **106** (1963) 259–269.

MATHEMATICS INSTITUTE, UNIVERSITY OF WARWICK, COVENTRY, CV4 7AL, GREAT BRITAIN