

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/125946>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

On planar Cayley graphs and Kleinian groups

Agelos Georgakopoulos*

Mathematics Institute
University of Warwick
CV4 7AL, UK

May 17, 2019

Abstract

Let G be a finitely generated group acting faithfully and properly discontinuously by homeomorphisms on a planar surface $X \subseteq \mathbb{S}^2$. We prove that G admits such an action that is in addition co-compact, provided we can replace X by another surface $Y \subseteq \mathbb{S}^2$.

We also prove that if a group H has a finitely generated Cayley (multi-)graph C covariantly embeddable in \mathbb{S}^2 , then C can be chosen so as to have no infinite path on the boundary of a face.

The proofs of these facts are intertwined, and the classes of groups they define coincide. In the orientation-preserving case they are exactly the (isomorphism types of) finitely generated Kleinian function groups. We construct a finitely generated planar Cayley graph whose group is not in this class.

In passing, we observe that the Freudenthal compactification of every planar surface is homeomorphic to the sphere.

2010 Mathematics Subject Classification: 05C10, 57M60, 57M07, 57M15.

Keywords: Planar Cayley graphs, covariant embedding, Kleinian groups, properly discontinuous actions, planar surface, Freudenthal compactification.

1 Introduction

The study of discrete groups of isometries, or Möbius transformations, of \mathbb{R}^2 and \mathbb{H}^2 is classic. It is at the heart of Klein's Erlangen program [15] as well as some of Poincaré's most famous work [14]. According to [14], these groups "... are tools used constantly today in various areas of mathematics and even in physics ...". For example, they play an important role in the study of 3-manifolds [32, 45, 46]. The groups alluded to here are generally called *Kleinian* groups, although the precise meaning of the term differs slightly depending on the author and era. Some survey material and many further references can be found e.g. in [27, 29, 32, 36, 41].

*Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 639046).

A well-studied sub-family of the Kleinian groups are the *function groups*, defined by the existence of an invariant component in their domain of discontinuity, see Section 2.5 for definitions, and [28, 29, 31, 32] for some literature. In this paper we show that every finitely generated function group is isomorphic to a group acting faithfully, properly discontinuously, and co-compactly on a planar surface.

Levinson & Maskit [26] proved that these groups are exactly the ones admitting a finitely generated Cayley graph that embeds in \mathbb{S}^2 with a fixed cyclic ordering of the labels of the edges around each vertex. Thus function groups form a subfamily (proper, as shown in this paper) of the *planar groups*, i.e. the groups having planar Cayley graphs, which are studied in recent work by the author and others [1, 10, 11, 13, 17, 18, 19, 24, 34, 42, 47]. We will prove that each finitely generated function group admits a planar Cayley graph having no infinite facial path, by possibly allowing loops and parallel edges. This was of interest to the author since [17], which provides constructions of 3-connected planar Cayley graphs in which no face is bounded by a finite cycle, previously thought to be impossible. These results apply more generally to groups containing orientation-reversing elements.

An embedding $\sigma : G \rightarrow \mathbb{S}^2$ of a Cayley graph G of a group Γ is *covariant*, if the canonical action of Γ on G maps every facial path onto a facial path; see Section 2 for more precise definitions. (This is equivalent to saying that the canonical action $\Gamma \curvearrowright \sigma(G)$ extends into an action $\Gamma \curvearrowright \mathbb{S}^2$ by homeomorphisms, see Corollary 5.5.)

A *planar surface* is a connected 2-manifold homeomorphic to an (open) subset of the sphere \mathbb{S}^2 . Our main result is

Theorem 1.1. *For a finitely generated group Γ , the following are equivalent:*

- (A) Γ admits a faithful, properly discontinuous action by homeomorphisms on a planar surface;
- (B) Γ has a Cayley graph admitting a covariant embedding;
- (C) Γ has a Cayley multi-graph (see Section 2.3) admitting a covariant embedding every facial path of which is finite;
- (D) Γ admits a faithful, properly discontinuous, **co-compact** action by homeomorphisms on the sphere, the plane \mathbb{R}^2 , the open annulus, or the Cantor sphere.

In the orientation preserving case, the groups of Theorem 1.1 coincide, as abstract groups, with the Kleinian function groups mentioned above:

Corollary 1.2. *A finitely generated group Γ admits a faithful, properly discontinuous (co-compact) action by orientation-preserving homeomorphisms on a planar surface if and only if it is isomorphic to a Kleinian function group.*

Corollary 1.2 can be deduced from [26, THEOREM 4], which essentially says that the groups of (B) coincide in the orientation-preserving case with the Kleinian function groups. We will give an alternative proof bypassing Ahlfors' finiteness theorem.

Our proof of Corollary 1.2 makes use of a classical theorem of Maskit, saying that if $p : \tilde{S} \rightarrow S$ is a regular covering of a topologically finite surface S ,

where \tilde{S} is planar, then the group of deck transformations is isomorphic to a function group. By exploiting the equivalence of (A) and (D) of Theorem 1.1 we can strengthen this statement by dropping the topological finiteness condition (Corollary 8.2); I do not know if this result is new. (Maskit's formulation however is stronger than the one above, and it is not possible to strengthen it this way, see Section 8 for details.)

In Corollary 1.2 we cannot replace 'Kleinian function group' by just 'Kleinian group': Section 10 provides an explicit example of a finitely generated Kleinian group that does not admit a planar Cayley graph (in fact most Kleinian groups have this property).

A description of the isomorphism types of Kleinian function groups in terms of fundamental groups of graphs of groups with simpler building blocks can be found in [26]. Dunwoody [13] extends this to groups with planar Cayley graphs.

I suspect that Corollary 1.2 extends to the orientation reversing case, by using orientation reversing Kleinian groups, i.e. discrete subgroups of $PSL(2, \mathbb{C})$. In support of this conjecture, we will also prove the following (in Section 12):

Theorem 1.3. *Every group Γ as in Theorem 1.1 admits a faithful properly discontinuous action on \mathbb{R}^3 .*

We show, in Section 10, that the family of such groups is a proper subfamily of the groups admitting a planar Cayley graph.

I suspect that the restriction of Γ being finitely generated can be dropped in all the above, except that in (D) we would have to extend the list of possible surfaces; see the remark in Section 1.1. Faithfulness is however necessary e.g. for the implication (A) \rightarrow (B).

The groups of Theorem 1.1 can be described by a certain kind of group presentation, which allows them to be effectively enumerated [18]. Thus we also obtain an effective enumeration of the isomorphism types of Kleinian function groups.

I do not know a proof of the implication (B) \rightarrow (C) of Theorem 1.1 that does not go via (D). I also do not know a proof of the implication (A) \rightarrow (D) that does not go via (B). However, if we relax (D) by not requiring co-compactness, then it follows from the following more general and perhaps well-known statement (Section 9):

Proposition 1.4. *Let $\Gamma \curvearrowright X$ be a properly discontinuous action on a metrizable, arc-connected, locally compact space X . Then the canonical extension $\Gamma \curvearrowright (X \cup \Omega^V(X))$ of the action to the non-accumulation ends $\Omega^V(X)$ of X is properly discontinuous.*

We remark that although any action $\Gamma \curvearrowright X$ as in Theorem 1.1 (D) defines a quotient orbifold $O := X/\Gamma$, understanding these groups goes beyond understanding 2-orbifold fundamental groups, because X is not simply connected in the cases we are most interested in, see Section 11.

Any action $\Gamma \curvearrowright X$ as in (D) can be 'geometrised', to turn it into an action by isometries on a smooth manifold homeomorphic to one of those four spaces, see Section 8.

In passing, we observe the following purely topological statement (Section 6)

Corollary 1.5. *Let $X \subseteq \mathbb{S}^2$ be a 2-manifold. Then the Freudenthal compactification of X is homeomorphic to \mathbb{S}^2 .*

Part of the motivation behind this paper was to understand which planar surfaces admit an action by an infinite group. Proposition 1.4, and the discussion in Section 9 sheds some light into this question. One way to produce such surfaces is to start from an action as in Theorem 1.1 (D), remove any totally disconnected subspace of the quotient space, and lift those punctures back to the original space. But this does not account for all such surfaces, see the Remark after Corollary 9.3. I think that it is possible to give a full list of those surfaces by pursuing these ideas further, but this will be rather tedious and will not add much to our understanding. Much more interesting would be to extend Theorem 1.1 (D) to higher dimensions:

Problem 1.1. *Is there for every $n \geq 3$ a finite list M_n of n -manifolds such that if a (finitely generated) group Γ acts faithfully and properly discontinuously on any n -manifold contained in \mathbb{R}^n , then Γ acts faithfully, properly discontinuously and co-compactly on an element of M_n ?*

Theorem 1.3 motivates

Problem 1.2. *Let Γ be a group acting faithfully and properly discontinuously on an n -manifold contained in \mathbb{R}^n . Does Γ act faithfully and properly discontinuously on \mathbb{R}^{n+1} ? On $\mathbb{R}^{n+f(n)}$ for some function $f : \mathbb{N} \rightarrow \mathbb{N}$?*

In this paper we proved the case $n = 2$, and this needs all the implications (A) \rightarrow (B) \rightarrow (D) \rightarrow (C) \rightarrow Theorem 1.3, so it is not straightforward to adapt our proof to higher dimensions.

The proof of Theorems 1.1 and 1.2 span Sections 5–8. Theorem 1.3 needs the full strength of Theorem 1.1 as it is based on (C). We prove it in Section 12.

1.1 Proof ideas

We now sketch the main ideas behind the proof of Theorem 1.1.

A central notion in the study of groups acting on surfaces, going back at least as far as Poincaré’s polyhedron theorem [32], is that of a *fundamental domain*. For the purposes of this sketch, let us say that a fundamental domain of the action $\Gamma \curvearrowright X$ is a subset $D \subset X$ containing exactly one point from each orbit of $\Gamma \curvearrowright X$. Thus Γ is in bijection with the set of translates of D , and the nicer D is, the easier it makes it to understand the action. In general, one wants D to be connected, and the ideal situation is when the closure of D is a polygon, with its translates giving a locally finite tessellation of X . This way one obtains e.g. regular tessellations of the hyperbolic and euclidean plane when Γ is a crystallographic group.

When Γ acts by arbitrary homeomorphisms rather than isometries however, and especially when X is the Cantor sphere, then it is not easy to find useful fundamental domains. Instead, we will work with fundamental domains of *graphs* embedded in X upon which Γ acts: a *fundamental domain* in this sense will be a connected subgraph containing exactly one vertex from each Γ -orbit. We will make extensive use of an observation of Babai [2] (Section 3.2), that if Γ acts freely on a connected graph H , then contracting each translate of a fundamental domain into a vertex turns H into a Cayley graph of Γ .

Let me explain how this helps to prove Theorem 1.1. Suppose Γ is a finite group, and $\Gamma \curvearrowright \mathbb{S}^2$ a faithful action by homeomorphisms. As a warm-up exercise, let me sketch a proof that Γ has a planar Cayley graph, which is the easiest special case of the implication (A) \rightarrow (B). Easily, there is a point $p \in \mathbb{S}^2$ with trivial stabiliser in $\Gamma \curvearrowright \mathbb{S}^2$. In other words, the orbit V of p is in bijection with Γ . Let G be any Cayley graph of Γ , for example the complete graph on Γ . We identify the vertex set of G with V using the aforementioned bijection, and we map each edge of G to an arc in \mathbb{S}^2 between its end-vertices such that the action $\Gamma \curvearrowright \mathbb{S}^2$ permutes these arcs. This can easily be achieved by choosing the arcs corresponding to edges of G going out of a reference vertex, and defining the rest of the arcs as the images of the former under $\Gamma \curvearrowright \mathbb{S}^2$. The map of G into \mathbb{S}^2 thus defined is not necessarily an embedding, as these arcs may intersect each other. However, it is easy to choose the arcs so that they intersect in at most finitely many points. Treating these intersection points as vertices defines a new graph G' . This G' is planar by definition. It is not a Cayley graph, although Γ acts on it by restricting $\Gamma \curvearrowright \mathbb{S}^2$ to G' . With a slight modification, only needed if Γ contains involutions, we can ensure that the latter action $\Gamma \curvearrowright G'$ is free. By Babai's aforementioned result, G' contracts into a Cayley graph C of Γ . Since the contracted sets are connected, and G' was planar, so is its minor C . We have found a planar Cayley graph of Γ , proving in particular that it is one of Maschke's groups [30].

The same technique works when Γ is infinite but finitely generated and acts properly discontinuously on a proper sub-surface of $X \subset \mathbb{S}^2$. Proper discontinuity is important for ensuring that each arc is intersected at most finitely often when mapping G into X , but the rest of the proof is essentially the same, and it yields the implication (A) \rightarrow (B).

The requirement that Γ be finitely generated was useful here in order to guarantee that each arc is intersected at most finitely often. If we drop it, and some arcs are intersected infinitely often, perhaps we can still control those intersections so as to get a graph-like space in the sense of [44]. I suspect that Babai's result can be extended to such spaces, and that this can be used to generalise our proofs to the infinitely generated case, but this will require additional work.

For (B) \rightarrow (D) the main machinery is the work of Thomassen & Richter [39], showing that if G is a 3-connected planar graph, then its Freudenthal compactification $|G|$ embeds in \mathbb{S}^2 , and this embedding is unique up to modifying it by a homeomorphism of \mathbb{S}^2 (the corresponding statement for finite graphs is a classical result of Whitney), see Section 3.1 for details. Given a Cayley graph G as in (B), we extend $\Gamma \curvearrowright G$ to the faces of an embedding of $|G|$ into \mathbb{S}^2 , and after removing the images of the ends of G we are left with a properly discontinuous action on one of the four spaces in (D), because every finitely generated Cayley graph has either at most two or a Cantor set of ends. Co-compactness is a byproduct of this construction. Some technical difficulties arise from the fact that our graphs are not necessarily 3-connected, and are overcome by embedding them into 3-connected graphs on which Γ acts freely but not transitively (Lemma 5.3).

To go from (D) to (C), we revisit the above proof of (A) \rightarrow (B). We are given a plane Cayley graph G of Γ some faces of which have infinite paths of G in their boundaries. We extend G by adding some further generators, and map the corresponding edges into arcs in \mathbb{S}^2 that cut up all such faces into smaller

faces bounded by cycles, when again we treat intersection points of these arcs as new vertices in an auxiliary graph G' . The fact that this is possible with only finitely many additional generators is not obvious: it requires Dunwoody's [12] result that finitely presented groups are accessible, combined with Droms' [10] result that planar groups are finitely presented, see Lemma 7.2. We then apply Babai's contraction result as above to G' to obtain a Cayley graph of Γ , and show that the property that all faces are bounded by finite cycles is hereby preserved.

For the proof of Theorem 1.3, we use (C) and the main result of [18]. The latter states that every covariantly planar Cayley graph G is the 1-skeleton of a Cayley complex Z , which complex can be mapped into \mathbb{S}^2 in such a way that (a) the restriction of the map to G is covariant, and (b) the images of any two 2-cells of Z are *nested*, i.e. either disjoint, or one contained in the other (Lemma 12.1). The latter property allows us to map the 2-cells of Z injectively into the inside B^3 of \mathbb{S}^2 in \mathbb{R}^3 , so that the image of Z separates B^3 into 3-dimensional 'chambers'. Using (C) and some refinements on Lemma 12.1 proved in Section 12, we can control the boundary of those chambers, so that we can use them to extend the action $\Gamma \curvearrowright Z$ into a properly discontinuous action on B^3 .

2 Definitions

2.1 Graphs

We follow the terminology of [9].

A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*.

Two rays R, L in G are *equivalent* if no finite set of vertices separates them; we denote this fact by $R \approx_G L$, or simply by $R \approx L$ if G is fixed. The corresponding equivalence classes of rays are the *ends* of G . We denote the set of these ends by $\Omega = \Omega(G)$.

2.2 Embeddings in the plane

An *embedding* of a graph G will always mean a topological embedding of the corresponding 1-complex in the sphere \mathbb{S}^2 ; in simpler words, an embedding is a drawing in \mathbb{S}^2 with no two edges crossing.

More generally, an embedding of a topological space X in a topological space Y , is a map $\sigma : X \rightarrow Y$ which is a homeomorphism of X to its image $\sigma(X)$.

A *plane* graph is a graph endowed with a fixed embedding. A graph is *planar*, if it admits an embedding.

A *face* of an embedding $\sigma : X \rightarrow \mathbb{S}^2$, where X is a topological space, is a component of $\mathbb{S}^2 \setminus \sigma(X)$. If X is a graph, or the Freudenthal compactification of a graph (see Section 2.7), we will say that the face F has *finite boundary*, if ∂F contains the images of only finitely many vertices and edges of G . Note that in this case ∂F is a cycle of G .

A walk or path in G is called *facial* with respect to σ if it is contained in the boundary of some face of σ .

2.3 Group actions

Given a group Γ and a generating set $S \subset \Gamma$, we define the (right) *Cayley graph* $G = \text{Cay}(\Gamma, S)$ to be the graph with vertex set $V(G) = \Gamma$ and edge set $E(G) = \{g(gs) \mid g \in \Gamma, s \in S\}$. We consider G as a directed, labelled graph, with the edge $g(gs)$ being directed from g to gs and labelled by the generator s . Here, we are assuming that S generates Γ , so that all Cayley graphs in this paper are connected. The group Γ acts on G by automorphisms, by multiplication on the left. By allowing S to be a multi-set, possibly containing the identity, in the above definition we obtain a *Cayley multi-graph*, with parallel edges and loops. For most of the paper parallel edges and loops, but they can matter in (C) of Theorem 1.1.

The *Cayley complex* $C_{\mathcal{P}}$ corresponding to a group presentation $\mathcal{P} = \langle S \mid \mathcal{R} \rangle$ is the 2-complex obtained from the Cayley graph G of \mathcal{P} by glueing a 2-cell along each closed walk of G induced by a relator $R \in \mathcal{R}$. Here, a *walk* W is a sequence v_0, v_1, \dots, v_k of vertices, such that each v_i is joined to v_{i+1} with an edge of G ; it is *closed* when $v_k = v_0$. We say that W is *induced* by a relator R , if R has exactly k letters and the label of the edge from v_i to v_{i+1} coincides modulo k with the $(i+n)$ th letter of R for every i and some fixed n .

Given a topological space X and a group Γ acting on it, the images of a point $x \in X$ under the action of Γ form the *orbit* of x . A *fundamental domain* is a subset of X which contains exactly one point from each of these orbits.

An action $\Gamma \curvearrowright X$ is *properly discontinuous*, if it satisfies any of the following equivalent conditions:

- (i) for every compact subspace K of X , the set $\{g \in \Gamma \mid gK \cap K \neq \emptyset\}$ is finite;
- (ii) for every two compact subspaces K, K' of X , the set $\{g \in \Gamma \mid gK \cap K' \neq \emptyset\}$ is finite;
- (iii) for every $x, y \in X$ there are open neighbourhoods $U_x \ni x, U_y \ni y$ such that U_x intersects gU_y for at most finitely many $g \in \Gamma$.

To see the equivalence of (i) and (ii) it suffices to notice that $K \cup K'$ is compact. A proof of the equivalence of (i) and (iii) can be found in [23], which offers many more equivalent definitions.

An action $\Gamma \curvearrowright X$ is *faithful*, if for every two distinct $g, h \in G$ there exists an $x \in X$ such that $gx \neq hx$; or equivalently, if for each $g \neq e \in G$ there exists an $x \in X$ such that $gx \neq x$. It is *free* if $gx \neq hx$ for every $g, h \in \Gamma$ and $x \in X$. It is *transitive* if for every $x, y \in X$ there is $g \in \Gamma$ with $gx = y$ (we will only encounter transitive actions on discrete spaces X). Finally, $\Gamma \curvearrowright X$ is *regular* if it is free and transient.

An action $\Gamma \curvearrowright X$ is *co-compact*, if the quotient space X/Γ is compact. If X is locally compact, then an equivalent condition is that there is a compact subset K of X such that $\bigcup \Gamma K = X$.

2.4 Covariant embeddings

Let G be a graph and $\sigma : G \rightarrow \mathbb{S}^2$ an embedding into the sphere. We say that σ is *covariant* with respect to a group action $\Gamma \curvearrowright G$, if every element of Γ maps each facial path of σ into a facial path. To simplify notation, if G is a

Cayley graph of Γ then we just say that σ is covariant in this case. If G is a plane graph, then we say that G is *covariant* if its embedding is covariant. If G admits a covariant embedding then we say that it is *covariantly planar*.

2.5 Kleinian groups

A *Kleinian group* is a discrete subgroup of $PSL(2, \mathbb{C})$. Since $PSL(2, \mathbb{C})$ is isomorphic to the group of Möbius transformations of the Riemann sphere, every Kleinian group Γ comes with a canonical action $\Gamma \curvearrowright \mathbb{S}^2$. The set of accumulation points of orbits of $\Gamma \curvearrowright \mathbb{S}^2$ is the *limit set* $\Lambda(\Gamma)$, and the *domain of discontinuity* is $\mathbb{S}^2 \setminus \Lambda(\Gamma)$. It is easy to check that Γ acts properly discontinuously on the latter.

2.6 Topology

The *boundary* ∂U of a subset U of a topological space X comprises the points $x \in X$ such that every open neighbourhood of x intersects both U and $X \setminus U$. The closure \overline{U} of U is the set $U \cup \partial U$.

The *Cantor sphere* is the topological space \mathcal{C} obtained by removing from \mathbb{S}^2 any subspace S homeomorphic to the Cantor set. The well-known fact that \mathcal{C} does not depend on the particular choice of S follows from Richards' classification of noncompact surfaces [38]. The following well-known fact provides some explanation why the Cantor set is important for us.

Proposition 2.1 ([6]). *Every nonempty totally disconnected perfect compact metrizable space is homeomorphic to the Cantor set*

A topological space is *k-connected* for some $k \in \mathbb{N}$, if it is connected and remains so upon the deletion of any $k - 1$ points.

A topological space is *arc-connected*, if it contains a homeomorph of the real unit interval joining any two of its points.

An *n-manifold* is a metric space X such that each point $x \in X$ has an open neighbourhood homeomorphic to \mathbb{R}^n . Smooth manifold structures are irrelevant in this paper, except shortly in Section 8.

2.7 The Freudenthal compactification

Even more than the above, this subsection is to be used as a reminder and for fixing notation; readers unfamiliar with the Freudenthal compactification are advised to consult a textbook for Topology.

Let X be a topological space, and suppose that $K_1 \subseteq K_2 \subseteq K_3 \dots$ is a sequence of compact subsets of X whose interiors cover X .

An *end* of X is an equivalence class of nested sequences $U_1 \supseteq U_2 \supseteq U_3 \dots$, where each U_i is a (connected) component of $X \setminus K_i$, and two such sequences $(U_n)_{n \in \mathbb{N}}, (V_n)_{n \in \mathbb{N}}$ are declared to be *equivalent*, if each U_i contains V_j for sufficiently large j and conversely, each V_i contains U_k for sufficiently large k .

A space X admitting a sequence $(K_n)_{n \in \mathbb{N}}$ as above is called *hemicompact*. Examples include all connected manifolds, all connected locally-finite graphs viewed as 1-complexes, and more generally, all connected locally-finite simplicial complexes.

The set of ends $\Omega(X)$ of X is used to define a compactification $|X| := X \cup \Omega(X)$, called the *Freudenthal compactification* or the end compactification of X . A basis of open neighbourhoods for $|X|$ can be obtained from one for X by declaring $U \cup \Omega(U)$ to be a basic open set whenever U is a component of $X \setminus K_i$ for some K_i as above, and $\Omega(U)$ is the set of equivalence classes of sequences $U_1 \supseteq U_2 \supseteq U_3 \dots$ where $U_i \subseteq U$ for large enough i .

It is straightforward to check that when X is a connected, locally finite graph, then this definition of $\Omega(X)$ coincides with the combinatorial one from Section 2.1.

3 Preliminaries

3.1 The extensions of Whitney's theorem by Thomassen & Richter

Lemma 3.1 ([39, Proposition 3]). *Let K be a compact, 2-connected, locally connected subset of \mathbb{S}^2 . Then every face of K is bounded by a simple closed curve (contained in K , since K is closed).*

Lemma 3.2 ([39, Lemma 12]). *Let G be a locally finite 2-connected planar graph. Then the Freudenthal compactification $|G|$ of G embeds in \mathbb{S}^2 .*

The following classical result of Whitney [48, Theorem 11] (which easily extends to infinite graphs by compactness, see e.g. [22]) says that 3-connected planar graphs have an essentially unique embedding.

Theorem 3.3 (Whitney's theorem). *Let G be a 3-connected graph embedded in the sphere. Then every automorphism of G maps each facial path to a facial path.*

Thomassen & Richter extended Whitney's theorem to the Freudenthal compactification of an infinite graph:

Lemma 3.4 ([39, Theorem 2]). *Let G be a 3-connected planar graph. For every $g \in \text{Aut}(G)$, and every embedding $\sigma : |G| \rightarrow \mathbb{S}^2$, there is a homeomorphism $h_g : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $h_g \circ \sigma = g$.*

Even more, given any face F of σ , and any homeomorphism $f : \overline{F} \rightarrow \overline{h_g(F)}$, we may assume that h_g coincides with f on \overline{F} .¹

In fact, Thomassen & Richter proved a more general version of Lemma 3.4, which we will also use:

Lemma 3.5 ([39, Theorem 2]). *Let X be a 3-connected, compact, locally connected Hausdorff space, admitting an embedding in \mathbb{S}^2 . Then X is uniquely planar.*

3.2 Babai's contraction lemma

Lemma 3.6 ([2]). *Let Γ be a group acting freely on a connected graph G . Then there is a connected subgraph D of G meeting each Γ -orbit at exactly one vertex, such that the contraction G/D is a Cayley graph of Γ .*

¹The second statement of Lemma 3.4 is not stated explicitly in [39, Theorem 2] but is implied by its proof.

Here, G/D is the graph obtained from G by contracting each Γ -translate of D into a vertex. One can think of D as the graph-theoretic analogue of a fundamental domain.

3.3 Existence of a regular orbit

Given a group action $\Gamma \curvearrowright X$, the *orbit* Γx of a point $x \in X$ is the set $\{gx \mid g \in \Gamma\}$.

The following is rather an exercise:

Lemma 3.7. *For every faithful, properly discontinuous action $\Gamma \curvearrowright X$ on a connected manifold X , there is an orbit $O = \Gamma x$ such that the restriction of the action to O is regular.*

One way to prove this for example is using the well-known facts that the quotient space of every properly discontinuous action on a manifold is an orbifold (see e.g. [4, Proposition 20]), and that the singular locus of any orbifold has empty interior (see e.g. [4, Proposition 26]); in particular, there is at least one point of the quotient that does not lie in the singular locus, which means by definition that its preimages have trivial stabilisers.

3.4 Richards' classification of non-compact 2-manifolds

The following is the special case of Richards' [38]² classification of non-compact 2-manifolds when restricting to subspaces of \mathbb{S}^2 .

Theorem 3.8 ([38]). *Let X and X' be two 2-manifolds contained in \mathbb{S}^2 . Then X and X' are homeomorphic if and only if their Freudenthal boundaries are homeomorphic.*

4 Separating faces and ends of planar graphs with cycles

In this section we prove some basic facts about planar graphs that we will need later, which have nothing to do with group actions.

If $\{V_1, V_2\}$ is a bipartition of the vertex set $V(G)$ of a graph G , then the set of edges with exactly one endvertex in each V_i is called a *cut* of G .

The following is a rather trivial consequence of the definition of the Freudenthal compactification

Lemma 4.1 ([9, Lemma 8.5.5 (ii)]). *Let G be a connected and locally finite graph, and B a finite cut of G . Then for every arc A in $|G| \setminus B$, the endpoints of A lie in the same component of $G \setminus B$.*

The next two lemmas allow us to separate faces and ends of a plane graph by cycles.

²Richards calls this 'Kerékjártó's theorem', but mentions that 'Kerékjártó's proof seems to contain certain gaps'.

Lemma 4.2. *Let G be a 2-connected graph and $\phi : |G| \rightarrow \mathbb{S}^2$ an embedding. For any two faces F, H of ϕ , there is a cycle C in G such that $\phi(C)$ separates F from H . Moreover, for every end $\omega \in \Omega(G)$ such that $\phi(\omega) \notin \partial F$, there is a cycle D in G such that $\phi(D)$ separates F from $\phi(\omega)$.*

Proof. If any of F, H is bounded by a cycle C of G then we are done, so assume this is not the case. Let P be a (possibly trivial) ∂F - ∂H path in G , and let $p = P \cap \partial F$ be its starting vertex. Let u, v be the two neighbours of p on ∂F , which are distinct because we are assuming F is not bounded by two parallel edges. By Lemma 3.1, there is a u - v arc in $|G| \setminus P$, namely ∂F with p and its two incident edges removed. Thus, easily, there is also a u - v path R in $G - P$. We claim that the cycle $C := Rvp$ separates F from H . Indeed, the path P , concatenated with any two arcs inside F and H respectively connects F to H and crosses C exactly once (at p), and so F, H lie in distinct sides of $\mathbb{S}^2 \setminus \phi(C)$.

To separate F from ω , we follow the same idea, replacing H with a basic open neighbourhood $O \ni \omega$, such that the closures $\overline{F}, \overline{O}$ are disjoint. Let P be a ∂F - ∂O path in G , and notice that $P \cap O = \emptyset$. As above, let u, v be the two neighbours of $p := P \cap \partial F$ on ∂F . There is now a u - v path R' in $G \setminus (P \cup \partial O)$ because ∂F contains a u - v arc in $|G| \setminus (P \cup \partial O)$. Then $D := R'vp$ is a cycle separating F from $\phi(\omega)$, because appending a ray of ω inside O to P we can obtain an F - $\phi(\omega)$ arc crossing D exactly once. \square

Lemma 4.3. *Let G be a connected locally finite graph and $\phi : |G| \rightarrow \mathbb{S}^2$ an embedding with no infinite face-boundaries. For any two ends $\omega, \chi \in \Omega(G)$, there is a cycle C in G such that $\phi(\omega)$ and $\phi(\chi)$ lie in distinct components of $\mathbb{S}^2 \setminus \phi(C)$.*

Proof. Let B be a finite cut of G separating ω from χ in $|G|$, which exists by the definitions.

Let $U := \bigcup \{ \partial F \mid F \text{ is a face of } \phi \text{ with } \partial F \cap B \neq \emptyset \}$. Note that $B \subseteq U$. Let H be the (finite) subgraph of G induced by the edges in U . Note that H is connected. We claim that $H \setminus B$ is disconnected. Indeed, H meets both sides of the cut B of G , and any path in H between those sides would be a path in G between the sides of B .

Next, we claim that $\phi(\omega)$ and $\phi(\chi)$ lie in distinct faces of H , where we think of H as a plane graph with embedding $\phi(H)$. To see this, let A be an ω - χ arc in $|G|$. Easily, A crosses B an odd number of times because B separates ω from χ (see Lemma 4.1). Since B separates H , the endpoints of $\phi(A)$ lie in faces F_ω, F_χ of H separated by B , and so $F_\omega \neq F_\chi$.

Then ∂F_ω contains a cycle C such that $\phi(C)$ separates F_ω from F_χ , and in particular $\phi(\omega)$ from $\phi(\chi)$ as desired. \square

For a graph K we let $E(K)$ denote its edge-set. The following says that if no element of a set of cycles in a plane graph separates two given points, then neither does the sum (in simplicial homology) of those cycles.

Lemma 4.4. *Let $G \subset \mathbb{S}^2$ be a plane graph and $\omega, \chi \in \mathbb{S}^2 \setminus G$. Let K be a cycle of G such that $E(K) = \sum_{1 \leq i \leq k} E(C_i)$ where C_i is a cycle of G which does not separate ω from χ , and the summation takes place in the cycle space of G . Then K does not separate ω from χ .*

Proof. Choose an ω - χ arc $A \subset \mathbb{S}^2$ that meets each edge of G in at most one point. Note that any Jordan curve in \mathbb{S}^2 , and in particular any cycle of G , separates ω from χ if and only if it crosses A an odd number of times. Thus each C_i crosses A an even number of times, hence so does K since adding edge sets of cycles preserves the parity of the number of crossings of A . We conclude that K does not separate ω from χ . \square

5 From planar Cayley graphs to actions on 2-manifolds

In this section we prove the implication (B) \rightarrow (D) of Theorem 1.1. This is mainly done in Lemma 5.6. We first collect a couple of lemmas.

Lemma 5.1. *Let X be a connected, locally finite graph, and suppose $\Delta \subset \text{Aut}(X)$ fixes a non-empty set $Y \subset V(X)$ and the stabiliser of each vertex of X in Δ is finite. Then all accumulation points of any orbit of $\Delta \curvearrowright |X|$ lie in the closure \overline{Y} of Y in the Freudenthal compactification $|X|$ of X .*

Here, we say that Δ fixes Y if $\Delta Y = Y$.

Proof. For $v \in Y$ it is clear that $\partial(\Delta v) \subseteq \overline{Y}$. For any other $w \in V(X)$, pick a w - v path P in X with $v \in Y$. Suppose, for a contradiction, that $x \in \partial(\Delta w) \setminus \overline{Y}$, and let $O \ni x$ be a basic open set disjoint from \overline{Y} . Since vertex stabilisers are finite, gP meets the finite set ∂O for only finitely many $g \in \Delta$. Thus almost all of $\{gP \mid g \in \Delta\}$ lie in O . Since this holds for an arbitrarily small neighbourhood O of x , we deduce $x \in \partial(\Delta v)$, and so $x \in \overline{Y}$ by our first remark. This contradiction proves that $\partial(\Delta w) \subseteq \overline{Y}$ for every $w \in V(X)$.

It remains to show that $\partial(\Delta \omega) \subseteq \overline{Y}$ for every end $\omega \in \Omega(X)$. For this, let $z \in \partial(\Delta \omega)$, and suppose to the contrary that $z \notin \overline{Y}$. Let $Z \ni z$ be a basic open set with $Z \cap \overline{Y} = \emptyset$. Let R be a ray of ω . We may assume that $R \subseteq Z$, for otherwise we can replace ω by another element ω' in its orbit $\Delta \omega$ for which this is true. Now for every vertex r of R , we have $\partial(\Delta r) \not\subseteq Z$ because we proved $\partial(\Delta r) \subseteq \overline{Y}$ above. Thus we can find a sequence $(g_n)_{n \in \mathbb{N}}$ of elements of Δ with $g_n r \notin Z$ such that $g_n R$ meets Z . Since $g_n R$ is connected, we can find a sequence $(r_n)_{n \in \mathbb{N}}$ of vertices of R such that $g_n r_n \in \partial Z$. As ∂Z is finite, we may even achieve $g_n r_n = w \in \partial Z$ for every n . But then $g_n^{-1} w \in V(R) \subseteq Z$, contradicting the fact that $\partial(\Delta w) \subseteq \overline{Y}$ for every $w \in V(X)$. \square

The following is a general tool for proving that a group action on a topological space is properly discontinuous.

Lemma 5.2. *Let X be a metrizable topological space, $\Gamma \curvearrowright X$ an action by homeomorphisms, and \mathcal{U} an open cover of X such that for every $U \in \mathcal{U}$, the orbit ΓU has no accumulation point in X . Then the action $\Gamma \curvearrowright X$ is properly discontinuous.*

Proof. By the definitions, $\Gamma \curvearrowright X$ is properly discontinuous unless there is a compact subspace K of X and an infinite sequence $(g_n)_{n \in \mathbb{N}}$ of elements of Γ such that $g_n K \cap K \neq \emptyset$ for every $n \in \mathbb{N}$.

In this case, let $\mathcal{U}_K := \{U \in \mathcal{U} \mid U \cap K \neq \emptyset\}$ be the subset of \mathcal{U} intersecting K . Since K is compact, \mathcal{U}_K has a finite subset \mathcal{U}_{K^∞} covering K . Since $g_n K \cap K \neq \emptyset$,

we have $g_n U_n \cap K \neq \emptyset$ for every $n \in \mathbb{N}$ and some $U_n \in \mathcal{U}_{K^\infty}$. As \mathcal{U}_{K^∞} is finite, we deduce that some $U_n \in \mathcal{U}_{K^\infty}$ has an accumulation point in K because K is metrizable, hence sequentially compact. This contradicts our assumptions. \square

Given a planar graph G and $\Gamma \subseteq \text{Aut}(G)$, we call an embedding $\sigma : G \rightarrow \mathbb{S}^2$ Γ -covariant, if every element of Γ maps each face-boundary of σ to a face-boundary of σ .

The following lemma was inspired by an idea of Dunwoody [13].

Lemma 5.3. *Let G be a connected, locally finite graph, $\Gamma \subseteq \text{Aut}(G)$, and $\sigma : G \rightarrow \mathbb{S}^2$ a Γ -covariant embedding. Then there is a 3-connected, locally finite supergraph H of G endowed with an extension $\Gamma \curvearrowright H$ of $\Gamma \curvearrowright G$ and a Γ -covariant embedding $\sigma' : H \rightarrow \mathbb{S}^2$ extending σ . Moreover, the maximum order of a vertex stabiliser of $\Gamma \curvearrowright H$ coincides with that of $\Gamma \curvearrowright G$. Furthermore, the identity map on $V(G)$ gives rise to a canonical homeomorphism from $\Omega(G)$ to $\Omega(H)$.*

Proof. We simultaneously construct H and its embedding σ' as follows. For every facial double-ray R of σ , we embed two copies R', R'' of R in the face F_R of σ incident with R , we join each vertex of R to the corresponding vertex of R' with an arc, and we join each vertex of R' to the corresponding vertex of R'' with an arc. Moreover, for every facial cycle C of σ that contains more than two edges, we embed a copy C' of C in the face of σ incident with R , and join each vertex of C to the corresponding vertex of C' with an arc. Easily, we can embed all those copies and arcs so that they never meet each other except at common vertices. This defines the supergraph H of G and its embedding σ' . Since σ was Γ -covariant, each element of Γ extends canonically into an automorphism of H , still preserving face-boundaries, and the finiteness of vertex stabilisers. It is clear by the construction that H has the same space of ends as G .

It remains to check that H is 3-connected. For this, suppose $\{x, y\} \subset V(H)$ disconnects H . Note that the subgraph of H spanned by a triple of double-rays $R \cup R' \cup R''$ as in the above construction is 3-connected. Similarly, the subgraph of H spanned by a couple $C \cup C'$ as above is 3-connected too. It follows that $\{x, y\} \subset V(G)$.

Let w, z be vertices in distinct components of $H - \{x, y\}$. By the above argument, we may assume $w, z \in V(G)$. Let P be a shortest w - z path in G . Let $F_1 F_2 \dots F_k$ be a sequence of faces of σ such that ∂F_1 contains the first edge of P , and ∂F_k contains the last edge of P , and for every $i < k$, F_i shares an edge e_i with F_{i+1} where e_i is incident with P but not contained in P . In other words, $F_1 \dots F_k$ is a F_1 - F_k path in the dual G^* that does not cross P ; it can be obtained as the sequence of faces visited by an arc 'parallel' to $\sigma(P)$ lying close enough to it. We will construct a w - z path Q in $H - \{x, y\}$, contradicting our assumption that $\{x, y\}$ disconnects w from z .

For this, note that each of the edges e_i from above has exactly one endvertex x_i in P , because P was chosen to be a geodesic path. It is now easy to construct Q as a concatenation of x_i - x_{i-1} paths of H contained in F_i and an appropriate initial and final path in F_1 and F_k , respectively, see Figure 1. \square

Using Lemma 5.3 we can immediately drop the condition of 2-connectedness in Lemma 3.2 for covariantly planar Cayley graphs:




Figure 1: The path Q in the proof of Lemma 5.3 (dashed).

Corollary 5.4. *Let G be a locally finite Cayley graph, and $\tau : G \rightarrow \mathbb{S}^2$ a covariant embedding. Then there is an embedding $\sigma : |G| \rightarrow \mathbb{S}^2$ such that every walk of G is facial in τ if and only if it is facial in σ (in particular, the restriction of σ to G is covariant).*

Proof. We use Lemma 5.3 to embed G into a 3-connected supergraph H . We then apply Lemma 3.2 to H , and restrict the resulting embedding from $|H|$ to its subspace $|G|$. Here, we used the fact that $\Omega(H)$ is canonically homeomorphic to $\Omega(G)$. \square

It should be possible to extend this even further to every locally finite planar graph using the construction in the proof of Lemma 5.3, i.e. to just drop the 2-connectedness condition in Lemma 3.2, but we will not need this. I suspect that Lemma 3.2 generalises even further as follows:

Problem 5.1. *Let G be a planar graph, and $\ell : E(G) \rightarrow \mathbb{R}_+$ an assignment of lengths to its edges. Let \hat{G} denote the metric completion of the metric space defined by ℓ (see [16] for the precise definition). Is \hat{G} always homeomorphic to a subspace of \mathbb{S}^2 ?*

Likewise, we can use Lemma 5.3 to relax the 3-connectedness condition in Lemma 3.4:

Corollary 5.5. *Let G be a connected, locally finite graph, and $\Gamma \subseteq \text{Aut}(G)$. Let $\sigma : |G| \rightarrow \mathbb{S}^2$ be a Γ -covariant embedding. Then for every $g \in \Gamma$, there is a homeomorphism $h_g : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $h_g \circ \sigma = g$.*

Even more, given any face F of σ , and any homeomorphism $f : \overline{F} \rightarrow \overline{h_g(F)}$, we may assume that h_g coincides with f on \overline{F} .

Proof. Let H be the 3-connected supergraph of G and $\sigma' : H \rightarrow \mathbb{S}^2$ its Γ -covariant embedding extending $\sigma(G)$, as provided by Lemma 5.3. As $\Omega(H)$ is canonically homeomorphic to $\Omega(G)$, we can in fact consider σ' to be an embedding of $|H|$. Applying Lemma 3.4 to H we obtain a homeomorphism $h_g : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $h_g \circ \sigma' = g$. Hence $h_g \circ \sigma = g$ as σ' extends σ .

For the second statement, given any face F of σ , and a homeomorphism $f : \overline{F} \rightarrow \overline{h_g(F)}$, we may assume that f maps $F \cap \sigma'(H)$ homeomorphically onto $h_g(F) \cap \sigma'(H)$ by constructing σ' appropriately in the proof of Lemma 5.3. Applying the second statement of Lemma 3.4 to each face F_i of σ' with $F_i \subseteq F$ and the restriction of f to F_i , we easily achieve that h_g coincides with f on \overline{F} . \square

Let G be a Cayley graph of a group Γ with an embedding $\sigma : G \rightarrow \mathbb{S}^2$. We say that σ is *orientation-preserving*, if the clockwise cyclic ordering in which the labels of the edges of a vertex x of G appear in σ is independent of the choice of x . By *labels* here we mean the corresponding element of the generating set of Γ used to define G . (If G is 3-connected, then Whitney's Theorem 3.3 implies that this cyclic ordering is the same for each $x \in V(G)$ up to orientation. See e.g. [17] for more.)

Lemma 5.6. *Let G be a connected, locally finite graph, and suppose $\Gamma \subseteq \text{Aut}(G)$ has finite vertex stabilisers. Let $\sigma : |G| \rightarrow \mathbb{S}^2$ be a Γ -covariant embedding. Then the action of Γ on G extends into a (faithful) properly discontinuous action of Γ on $\mathbb{S}^2 \setminus \sigma(\Omega(G))$ by homeomorphisms.*

Moreover, if Γ has finitely many orbits of vertices, then the latter action can be chosen to be co-compact.

Furthermore, if G is a Cayley multi-graph of Γ , and σ is orientation-preserving, and all faces of σ have finite boundary, then the latter action can be chosen so that, in addition, it stabilises at most one point of each face of σ .

Proof. Given $\Gamma \subseteq \text{Aut}(G)$, we want to extend each $g \in \Gamma$ into a homeomorphism $h_g : \mathbb{S}^2 \rightarrow \mathbb{S}^2$. We obtain h_g by applying Corollary 5.5. In order for this map $g \mapsto h_g$ to extend the action $\Gamma \curvearrowright G$ into an action $\Gamma \curvearrowright \mathbb{S}^2$, we need it to be a homomorphism from Γ to $\text{Aut}(\mathbb{S}^2)$. To achieve this, we will exploit the second statement of Corollary 5.5 appropriately.

For this, given any map $\eta : \partial F \rightarrow \partial F$ which is the restriction of one or more $g \in \Gamma$ to the boundary ∂F of some face F of σ , we fix a homeomorphism $f_\eta : \overline{F} \rightarrow \overline{F}$ such that f_η coincides with η on ∂F . The existence of f_η is a consequence of Corollary 5.5, which unsurprisingly makes use of the Jordan-Schönflies theorem.

Similarly, if for two faces F, F' of σ there is $g \in \Gamma$ mapping $\partial F'$ to ∂F , then we fix a map $\zeta = \zeta_{F',F} : \partial F' \rightarrow \partial F$ which is the restriction of such a g to $\partial F'$, and a homeomorphism $f_\zeta : \overline{F'} \rightarrow \overline{F}$ such that f_ζ coincides with ζ on $\partial F'$.

By compositions of these f_η 's and f_ζ 's we obtain, for any $g \in \Gamma$ and any face F , a unique map f_F extending the restriction of g to ∂F to \overline{F} . When we apply Corollary 5.5 to define the map $g \mapsto h_g$ as above, we can, by its second statement, assume that the restriction of h_g to each face F coincides with this f_F .

This immediately implies that $h_{gg'} = h_g h_{g'}$, hence our map $\phi : \Gamma \rightarrow \text{Aut}(\mathbb{S}^2)$ defined by $g \mapsto h_g$ is an injective homomorphism. By identifying Γ with its image in $\text{Aut}(\mathbb{S}^2)$ under ϕ we obtain the action $\Gamma \curvearrowright \mathbb{S}^2$.

Let $X := \mathbb{S}^2 \setminus \sigma(\Omega(G))$. Since $\Gamma(\Omega(G)) = \Omega(G)$, we can define an action $\Gamma \curvearrowright X$ by homeomorphisms just by restricting the above action from \mathbb{S}^2 to X .

In the special case where all faces of σ have finite boundary, this action turns out to be properly discontinuous and, if $V(G)$ has finitely many Γ -orbits, co-compact. We will first handle this special case as a warm-up towards the more involved general case.³

³In fact this special case can be handled by noting that X is a CW-complex and $\Gamma \curvearrowright X$ a cellular action, and using the equivalence between (2) and (10) in [23, Theorem 9]. But we provide the following proof in order to ease the understanding of the proof of the general case.

So let us first assume that all faces of σ have finite boundary. To show that $\Gamma \curvearrowright X$ is properly discontinuous we apply Lemma 5.2 with the following choice of the cover \mathcal{U} of X :

- (i) for every $x \in \sigma(V(G))$, we choose an open neighbourhood U_x of x in \mathbb{S}^2 meeting only edges of G and faces of σ incident with x , and containing the image of no other vertex of G ;
- (ii) for every point $x \in \sigma(e)$ with $e \in E(G)$, we choose an open neighbourhood U_x of x contained in the union of $\sigma(e)$ with the faces of σ incident with e ;
- (iii) for every x in a face F of σ , we choose $U_x = F$;

Let $\mathcal{U} := \{U_x \mid x \in X\}$ be the resulting cover of X . We claim that \mathcal{U} satisfies the requirement of Lemma 5.2 that ΓU has no accumulation point in X for every $U \in \mathcal{U}$. To see this for $U = U_x$ of type (iii), note that Γ maps every face of σ to a face of σ , and that $U_x = F$ has a finite stabiliser in Γ because it is determined by the finite set of vertices in ∂F and we are assuming that every vertex has a finite stabiliser in Γ . For U_x of type (i) or (ii) we remark that U_x is contained in the closure of the union of finitely many U 's of type (iii), and repeat the same argument.

To show that $\Gamma \curvearrowright X$ is co-compact when $V(G)$ has finitely many orbits, pick $D \subset V(G)$ containing exactly one vertex from each orbit, and let K be the union of the closures of the faces incident with vertices in D . Then K is compact, and it is easy to see that $\bigcup \Gamma K = X$, which means that $\Gamma \curvearrowright X$ is co-compact.

To prove the final statement, we modify G into a triangulation G' by adding a new vertex v_F inside each face F of σ , and joining v_F to each vertex incident with F with a new edge, which edge we draw as an arc inside F to obtain an embedding σ' of G' . Since σ is Γ -covariant, $\Gamma \curvearrowright G$ extends to an action on G' . Since Γ stabilises no vertex of G , and it preserves the orientation of \mathbb{S}^2 , the only vertices of G' it stabilises are the v_F 's. Repeating the above construction with G' instead of G yields the desired action, where the only points of $\mathbb{S}^2 \setminus \sigma(\Omega(G))$ fixed are the images of the v_F 's.

We now consider the general case, where some faces of σ may have infinite boundary. In this case the above argument fails because the stabiliser of such a face may be infinite. To circumvent this difficulty, we will subdivide such faces by embedding appropriate graphs inside them, thus embedding G into a plane supergraph G' , in such a way that the faces of G' have finite stabilisers in $\Gamma \curvearrowright G'$ even if they have *infinite boundary*⁴.

To simplify the exposition, we assume below that G is 2-connected (in order to be able to apply Lemmas 3.1 and 4.2). If not, then we apply the following arguments to the 3-connected supergraph H of G provided by Lemma 5.3 to obtain the desired action for H , hence also for G (the final sentence has already been proved in the above special case of finite face boundaries).

Our auxiliary supergraph G' of G is defined as follows. For every face F of σ (with infinite boundary), we 'reflect' our embedding $\sigma : |G| \rightarrow \mathbb{S}^2$ from the complement of F into F ; that is, we choose a homeomorphism $r : (\mathbb{S}^2 \setminus F) \rightarrow \overline{F}$

⁴We say that a face F has *finite boundary*, if ∂F contains finitely many edges. Otherwise we say that F has *infinite boundary*.

that coincides with the identity on \overline{F} , which exists by the Jordan-Schoenflies theorem [43] and the fact that ∂F is a simple closed curve (Lemma 3.1). Then $r \circ \sigma$ embeds a copy of G into \overline{F} . We call this copy of G the *shadow* of G in F . We define G' to be the plane graph obtained as the union of G with all those shadows, one for each face of σ , and let $\sigma' : G' \rightarrow \mathbb{S}^2$ denote the corresponding embedding. In fact, we can extend σ' into an embedding of $|G'|$ in \mathbb{S}^2 .

It is not hard to see that the action $\Gamma \curvearrowright G$ extends to an action $\Gamma \curvearrowright G'$: as each $g \in \Gamma$ maps each face F of σ to a face gF of σ via the map $g \mapsto h_g$, we can let g map the shadow of G in F to the shadow of G in gF via the unique automorphism of G determined by the restriction of g to ∂F . This defines the action $\Gamma \curvearrowright G'$, and therefore an action $\Gamma \curvearrowright |G'|$. Easily, σ' is Γ -covariant too.

We claim that

$$\text{every orbit of } \Gamma \curvearrowright |G'| \text{ has all its accumulation points in } \Omega(G) \subset \Omega(G'). \quad (1)$$

Indeed, this follows from Lemma 5.1, applied with $X = G'$ and $Y = V(G)$.

Our next claim is that G' is still 2-connected. Indeed, after removing any vertex $x \in V(G')$, the graph G as well as any of its closed shadows is still connected since G was assumed to be 2-connected. Moreover, any shadow S is still connected to G as G and S have several vertices in common by construction, and our claim easily follows.

Since σ' is Γ -covariant, we can apply Corollary 5.5 to it. As we did in the special case where all faces of G have finite boundary above, we exploit the second statement of Corollary 5.5 (the maps f_η now correspond to faces of σ') so as to extend $\Gamma \curvearrowright |G'|$ into an action $\Gamma \curvearrowright \mathbb{S}^2$ and, letting $X := \mathbb{S}^2 \setminus \sigma(\Omega(G))$ (and not $\mathbb{S}^2 \setminus \sigma(\Omega(G'))$), into an action $\Gamma \curvearrowright X$ by restriction.

As above, we use Lemma 5.2 to show that $\Gamma \curvearrowright X$ is properly discontinuous. The cover \mathcal{V} of X we use is similar to \mathcal{U} from above; the only difference is that now X contains points that are images of ends $\omega \in \Omega(G') \setminus \Omega(G)$, for which we need to choose open sets in our cover. For every such end ω , we choose a cycle D_x in G' separating ω from ∂F where F is the face of σ containing $\sigma'(\omega)$, which cycle exists by the second part of Lemma 4.2, and will be used in (iv) below. Our cover \mathcal{V} of X is defined as follows:

- (i) for every $x \in \sigma'(V(G'))$, we choose an open neighbourhood V_x of x in \mathbb{S}^2 meeting only edges of G and faces of σ' incident with x , and containing the image of no other vertex of G' ;
- (ii) for every point $x \in \sigma'(e)$ with $e \in E(G')$, we choose an open neighbourhood V_x of x contained in the union of $\sigma'(e)$ with the faces of σ' incident with e ;
- (iii) for every x in a face F of σ' , we choose $V_x = F$;
- (iv) for every $x \in \sigma'(\Omega(G') \setminus \Omega(G))$, we let V_x be the component of $\mathbb{S}^2 \setminus D_x$ containing x , where D_x is the cycle defined above.

Let $\mathcal{V} := \{V_x \mid x \in X\}$ be the resulting cover of X .

Again, we claim that ΓU has no accumulation point in X for every $U \in \mathcal{U}$. To prove this for $U = V_x$ of type (iii), we recall that Γ maps every face of σ' to a face of σ' , and show that $V_x = F$ has a finite stabiliser in Γ . If F coincides with an original face of σ with finite boundary, then we showed this above.

Otherwise, F lies inside a face H of σ in which σ' embeds a shadow of G . In this case, applying Lemma 4.2 to that shadow yields a cycle C_F separating F from ∂H . By (1), the orbit ΓC_F of C_F has all its accumulation points in $\Omega(G)$ in the topology of $|G'|$. Since $\sigma'(|G'|)$ is a homeomorph of $|G'|$, this means that the orbit of the circle $\sigma'(C_F)$ under our action $\Gamma \curvearrowright \mathbb{S}^2$ has all its accumulation points in $\sigma'(\Omega(G))$. Since F is contained in one of the sides of C_F , this easily implies that the orbit of $V_x = F$ also has all its accumulation points in $\sigma'(\Omega(G))$. Thus V_x has no accumulation point in $X = \mathbb{S}^2 \setminus \sigma'(\Omega(G))$ as desired.

For $U = V_x$ of type (iv) we use the same argument, replacing the cycle C_F by D_x .

For V_x of type (i) or (ii) we argue as earlier: we remark that V_x is contained in the closure of the union of finitely many V 's of type (iii), and repeat the same argument.

Thus \mathcal{V} satisfies the requirement of Lemma 5.2, and we deduce that $\Gamma \curvearrowright X$ is properly discontinuous.

To show that $\Gamma \curvearrowright X$ is co-compact when $V(G)$ has finitely many orbits, we argue as above, except that we now employ Lemma 4.2 as follows. Pick $D \subset V(G)$ containing exactly one vertex from each orbit. For every face F of G' incident with a vertex of D , let C_F be a cycle of G' separating F from the complement of the face F' of G containing F , which exists by Lemma 4.2 applied to the shadow of G in F' . (If F has finite boundary, we can just let $C_F = \partial F$.) Let K be the union of the closures of the interiors of all these cycles C_F . Then K is compact, and again we have $\bigcup \Gamma K = X$, which means that $\Gamma \curvearrowright X$ is co-compact. □

6 From actions on 2-manifolds to planar Cayley graphs

In this section we prove the implication (A) \rightarrow (B) of Theorem 1.1 (Lemma 6.3).

We say that a space $X \subseteq \mathbb{S}^2$ is *uniquely planar*, if for every two embeddings $\sigma, \sigma' : |X| \rightarrow \mathbb{S}^2$ there is a homeomorphism $h : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $h \circ \sigma = \sigma'$, and there is at least one such embedding σ . In other words, every automorphism of $|X|$ extends into a homeomorphism of \mathbb{S}^2 .

Lemma 6.1. *Every connected 2-manifold $X \subseteq \mathbb{S}^2$ is uniquely planar.*

Proof. Our first aim is to show that $|X|$ admits an embedding in \mathbb{S}^2 . We will do this by applying Lemma 3.2 on a triangulation of X .

Let D be a locally finite triangulation of X , which exists by a well-known theorem of Radó [37, 49], and let T be its 1-skeleton. Then T is a 2-connected graph by definition. Therefore, there is an embedding $\sigma : |T| \rightarrow \mathbb{S}^2$ by Lemma 3.2. In fact, we may assume T to be 3-connected: we can subdivide each edge of D by putting a new vertex at its midpoint, and triangulate each original triangle Δ of D into 4 smaller triangles by adding the three edges joining the midpoints of the three edges of Δ . If T is the 1-skeleton of the resulting triangulation, then it is indeed 3-connected because the neighbourhood of Δ is still connected after removing any 2 vertices from Δ .

Assuming, as we now can, that T is 3-connected, we deduce that our embedding $\sigma : |T| \rightarrow \mathbb{S}^2$ is essentially unique when restricted to T by (the infinite version of) Whitney's Theorem 3.3. Therefore, the 1-skeleton of each triangle of D bounds a face of σ , because it does so in the embedding of T induced by the identity on $X \subseteq \mathbb{S}^2$. Thus we can extend σ from T to D by mapping each 2-cell into the corresponding face of σ . Since T and X have the same ends by the definitions, we have thus extended σ into an embedding $\phi : |X| \rightarrow \mathbb{S}^2$.

We can now apply Lemma 3.5 to deduce that $|X|$ is uniquely planar. \square

As a consequence of the first claim of this proof, we obtain Corollary 1.5, which we restate for convenience:

Corollary 6.2. *Let $X \subseteq \mathbb{S}^2$ be a 2-manifold. Then $|X|$ is homeomorphic to \mathbb{S}^2 .*

Proof. We showed above that $|X|$ admits an embedding σ in \mathbb{S}^2 . If σ is not surjective, pick $x \in \mathbb{S}^2 \setminus \sigma(|X|)$, and let F be the face of σ containing x . Applying Lemma 3.1 with $K := \sigma(|X|)$, we obtain that F is bounded by a simple closed curve $S \subset |X|$. Since $\Omega(X)$ is totally disconnected, $\sigma^{-1}(S)$ contains a point y of X . We obtain a contradiction as y has a neighbourhood U homeomorphic to \mathbb{R}^2 , and every neighbourhood of $\sigma(y)$ meets F . \square

We remark that Corollary 6.2 does not extend to 3 dimensions: let M be the inside of a torus embedded in \mathbb{R}^3 . Clearly M is a 3-manifold embeddable in \mathbb{S}^3 . But $|M|$ is not homeomorphic to \mathbb{S}^3 , and in fact it is not homeomorphic to a 3-manifold. Indeed, the boundary $|M| \setminus M$ consists of a single point x , and x has arbitrarily small open neighbourhoods U such that $U \setminus x$ is homeomorphic to M , hence not simply connected.⁵

We can now prove the main result of this section. Recall the definition of an orientation-preserving embedding σ of a Cayley graph given before Lemma 5.6.

Lemma 6.3. *Let Γ be a finitely generated group with a faithful and properly discontinuous action $\Gamma \curvearrowright X$ on a connected 2-manifold X . Then there is a finitely generated Cayley graph G of Γ and an embedding $\sigma : G \rightarrow X$ such that $\sigma(V(G))$ has no accumulation points in X , and $\Gamma \curvearrowright G$ coincides with the restriction of $\Gamma \curvearrowright X$ to $\sigma(G)$.*

Moreover, if $X \subseteq \mathbb{S}^2$, then σ defines a covariant embedding of G into \mathbb{S}^2 . Furthermore, if $\Gamma \curvearrowright X$ is orientation-preserving, then so is σ .

Proof. Pick a finite generating set S of Γ . Let o be a point of X such that the orbit Γo is regular, which exists by Lemma 3.7. For every $s \in S$, pick an o - so arc R_s in X . By straightforward topological manipulations we may assume that these arcs and their translates do not intersect too much: we can assume that for every $s, t \in S$, and any $g \in \Gamma$, the intersection $R_s \cap gR_t$ is either empty or just one point. This means that the union $\bigcup_{s \in S} \Gamma R_s$ of all these arcs and their translates defines a (plane) graph H embedded in X , the vertex set of which consists of our regular orbit Γo and all the intersection points $R_s \cap gR_t$ of our arcs. (We could assume that no three arcs meet at a point, but we will not need to.) To show that H is a graph, we need to check that no arc R_s is intersected by infinitely many other arcs. This is the case because Γ acts properly discontinuously, and the 'star' $\bigcup_{s \in S} R_s$ of o is compact.

⁵I thank Max Pitz for this observation.

Since Γo is an orbit of Γ , and S generates Γ , it follows easily that H is connected. Moreover, $\Gamma \curvearrowright X$ defines an action $\Gamma \curvearrowright H$ by restriction. We would like to apply Babai's contraction Lemma 3.6 to contract H onto a Cayley graph of Γ , but the freeness condition is not necessarily satisfied because some involution $g \in \Gamma$ might exchange two crossing arcs R_s, gR_s , hence stabilising their intersection point. But this is easy to amend by a slight modification of H : we blow up each vertex v of H arising from an intersection point into a circle that intersects H at the edges incident with v only, intersecting each of them exactly one. Thus v is replaced in H by a cycle C_v all vertices of which have degree 3. Let H' denote this modification of H . It is now straightforward to check that Γ acts freely on H' . (If some involution reverses an arc R_s , then we treat its midpoint as an intersection vertex of H and apply to it the aforementioned blow-up operation.)

After this modification, we can indeed apply Lemma 3.6 to H' : we choose a connected subgraph D meeting each orbit at exactly one vertex, and contract each translate $gD, g \in \Gamma$ of D into a vertex to obtain a Cayley graph G of Γ .

Next, we remark that G embeds in X . For this, pick a spanning tree T of D , and a neighbourhood U of T in X homeomorphic to \mathbb{R}^2 meeting no other translate of D . Note that $o \in U$ by the definitions. For each edge xy incident with (the contraction of) D in G , with $x \in V(D)$, say, we pick an o - x arc R_x in U , making sure that R_x is disjoint from R_z for $z \neq x$. We can now think of o rather than D as a vertex of G . We repeat this at every other translate gD of D by mapping U to gU . The union of the translates of the arcs R_x with the appropriate subarcs of the original arcs R_s defining H defines an embedding of G in X .

If $x \in X$ were an accumulation point of $\sigma(V(G))$, then any neighbourhood $U_x \ni x$ would contain infinitely many elements of the orbit of o , violating our definition (iii) of a properly discontinuous action. Thus $\sigma(V(G))$ has no accumulation point in X .

Easily, the canonical action $\Gamma \curvearrowright E(G)$ extends to the original action $\Gamma \curvearrowright H$ obtained by restricting $\Gamma \curvearrowright X$. Thus $\Gamma \curvearrowright G$ extends to $\Gamma \curvearrowright X$.

Moreover, we claim that if $X \subseteq \mathbb{S}^2$ then our embedding σ , which we can then think of as an embedding into \mathbb{S}^2 , is covariant. For this, recall that each $g \in \Gamma$ is a homeomorphism of X , which by Lemma 6.1 extends into a homeomorphism of \mathbb{S}^2 . Since $\Gamma \curvearrowright H$ is defined by restricting $\Gamma \curvearrowright X$, it follows that g maps every facial walk of H into a facial walk.

As G was defined by contracting Γ -translates of a subgraph of H , it is now easy to see that g maps every facial walk of G with respect to σ into a facial walk, and so σ is covariant.

Similarly, if $\Gamma \curvearrowright X$ was orientation-preserving, then by considering the orbit of the edges of o in G under $\Gamma \curvearrowright X$ it is easy to see that σ is orientation-preserving too. \square

Remark: Lemma 7.1 below is proved using similar ideas, and the reader is advised to read that lemma right after reading the proof of Lemma 6.3.

7 Finiteness of face-boundaries

In this section we prove the implication (B) \rightarrow (C) of Theorem 1.1:

Lemma 7.1. *If a group Γ admits a finitely generated, covariantly planar Cayley graph, then Γ admits a finitely generated, covariantly planar Cayley multi-graph with no infinite face-boundaries.*

I do not know how to prove this without using Lemma 5.6, which is the essence of the implication (B) \rightarrow (D).

We will need the following terminology. If $\{V_1, V_2\}$ is a bipartition of the vertex set $V(G)$ of a graph G , then the set of edges with exactly one endvertex in each V_i is called a *cut* of G . The *cut space* $\mathcal{B}(G)$ is the vector space over the 2-element field generated by the finite cuts of G ; see [9, §1.9] for a precise definition.

Lemma 7.2. *Let $G = \text{Cay}(\Gamma, S)$ be a finitely generated planar Cayley graph. Then there is a finite set D of cuts of G such that $\bigcup \Gamma D$ generates the cut space of G .*

Proof. By a theorem of Droms [10, Theorem 5.1], every group Γ that admits a finitely generated planar Cayley graph is finitely presented. By a theorem of Dunwoody [12], every finitely presented Γ is accessible. The fact that every accessible group satisfies the conclusion of our statement is proved in [8, Corollary IV.7.6]. \square

Proof of Lemma 7.1. Let G be a finitely generated, covariantly planar Cayley graph of Γ . To simplify our exposition, let us first assume that G is 3-connected; we will later employ Lemma 5.3 to treat the general case.

By Lemma 5.4 we obtain an embedding $\sigma : |G| \rightarrow \mathbb{S}^2$, and by Lemma 5.6 the canonical action $\Gamma \curvearrowright G$ extends into a properly discontinuous action $\Gamma \curvearrowright \mathbb{S}^2 \setminus \sigma(\Omega(G))$.

Our strategy for modifying G into a planar Cayley graph of Γ with no infinite face-boundaries can be sketched as follows. In each face F of G with infinite boundary we embed some arcs in a Γ -invariant way in order to split F into faces with finite boundaries. We choose those arcs so that treating their intersection points as vertices defines a plane supergraph G^+ of G using the ideas of Lemma 6.3. Then we apply Babai's contraction Lemma 3.6 to the action of Γ on G^+ , to contract the latter onto a Cayley graph H of Γ inheriting the property that all faces have finite boundaries.

To make this sketch precise, let D be a finite set of cuts of G such that $\bigcup \Gamma D$ generates $\mathcal{B}(G)$, provided by Lemma 7.2. If σ has no infinite face-boundary then there is nothing to prove. Otherwise, for every $B \in D$, and any two edges $e, f \in B$ lying in a common face-boundary ∂F of a face F of σ , choose an arc $A_{e,f,F}$ in \overline{F} joining an endvertex of e to an endvertex of f , which arc exists because ∂F is a simple closed curve by Lemma 3.1 and so \overline{F} is homeomorphic to a closed disc by the Jordan-Schönflies theorem. (If e, f lie in more than one such ∂F , choose such an arc in each F .) Let \mathcal{A} denote the union of the Γ -orbits of all such arcs $A_{e,f,F}$ with $e, f \in B \in D$. If these arcs are chosen appropriately, then as in the proof of Lemma 6.3, the union of \mathcal{A} with G is a plane graph G^+ , which is a supergraph of G , embedded in \mathbb{S}^2 via an extension ρ of $\sigma(G)$. Here, we used again the fact that any $A_{e,f,F} \in \mathcal{A}$ intersects only finitely many elements of \mathcal{A} since $\Gamma \curvearrowright \mathbb{S}^2 \setminus \sigma(\Omega(G))$ is properly discontinuous. This fact easily implies that the identity map defines a homeomorphism $h : \Omega(G) \rightarrow \Omega(G^+)$,

and furthermore the extension $\rho' : |G^+| \rightarrow \mathbb{S}^2$ of ρ obtained by mapping each $\omega \in \Omega(G^+)$ to $\sigma(h^{-1}(\omega))$ is an embedding.

By construction, the action $\Gamma \curvearrowright \mathbb{S}^2 \setminus \sigma(\Omega(G))$ defines an action $\Gamma \curvearrowright G^+$. As in Lemma 6.3, by blowing up each vertex of $G^+ \setminus G$ into a circle if needed, we may assume that the latter action is free. By Babai's contraction Lemma 3.6, there is a connected subgraph Δ of G^+ such that $H := G^+/\Delta$ is a Cayley graph of Γ . Here, we keep any parallel edges and loops of H resulting from these contractions, that is, H is a multi-graph, because they could be needed to retain the finiteness of face boundaries.

It is easy to see that Δ is finite, because as noted above every arc $A_{e,f,F}$ as in the definition of G^+ is intersected by only finitely many translates of such arcs.

As in Lemma 6.3, we can modify σ into an embedding $\tau : H \rightarrow \mathbb{S}^2$ so that every face boundary of H can be obtained from one of G^+ by contracting each maximal subpath contained in a translate of Δ into a vertex. We claim that all face-boundaries of H in τ are finite, which proves our statement. By the previous remark, this claim will follow if we can show that G^+ has no infinite face-boundary, or equivalently, that G^+ has no facial double-ray.

Suppose, to the contrary, that R is a facial double-ray of G^+ , and let R' be the corresponding facial double-ray of H , obtained by contracting each maximal subpath contained in a translate of Δ . We distinguish the following two cases.

Case 1: If R comprises two disjoint sub-rays belonging to the same end of G^+ , then as Δ is finite, this situation passes on to H : the two tails of R' belong to the same end of H . This however is ruled out by a result of Krön [24], stating that in any almost transitive, plane graph with finite vertex degrees, disjoint tails of a facial double-ray lie in distinct ends.

Case 2: If R has two disjoint sub-rays belonging to distinct ends ω, χ of G^+ , then let F denote the face of $|G|$ with respect to σ containing R , which face exists since G^+ is a supergraph of G , and $\rho(R)$ avoids $\sigma(\Omega(G))$. Then $\omega, \chi \in \partial F$ because ρ' is an embedding of $|G^+|$ into \mathbb{S}^2 as we saw earlier. We claim that

$$\text{there is an arc } A \in \mathcal{A} \text{ such that } \omega, \chi \text{ lie in distinct components of } \overline{F} \setminus A. \quad (2)$$

Before proving this claim, let us see why it leads to a contradiction. The closure \overline{D} of D in $|G^+|$ is an ω - χ arc by the definitions, and so $\rho'(\overline{D})$ is an ω - χ arc in \overline{F} . Thus claim (2) means that D crosses A , and so it cannot be facial. This contradiction proves that Case 2 cannot occur either.

It remains to prove (2). For this, note that the cutspace $\mathcal{B}(G)$ contains a cut B separating ω from χ . Thus its generating set $\bigcup \Gamma D$ must contain such a cut gC with $C \in D$. We claim that gC contains edges e_1, e_2 in ∂F such that ω, χ lie in distinct components of $\partial F \setminus \{e, f\}$. This is true because ∂F is a simple closed curve contained in $\sigma(|G|)$ by Lemma 3.1, and so its preimage $Z := \sigma^{-1}(\partial F)$ is a homeomorph of \mathbb{S}^1 in $|G|$. In particular, Z contains two internally disjoint ω - χ arcs Z_1, Z_2 . But then the cut gC must meet both those arcs in order to separate ω from χ , because for every ω - χ arc Y and every cut J separating ω from χ we have $Y \cap J \neq \emptyset$ by Lemma 4.1. Thus we can let e_i be any edge in $gC \cap \sigma(Z_i)$ for $i = 1, 2$. Finally, by the definition of G^+ , the latter contains an arc A in \overline{F} from an endvertex of e_1 to an endvertex of e_2 . This arc disconnects ω, χ in \overline{F} because e_1, e_2 disconnect ω, χ in ∂F .

This completes our proof in the case where G is 3-connected. If it is not, then

we apply Lemma 5.3 to embed G into a planar 3-connected plane supergraph G' to which the action of Γ extends covariantly. We then repeat the above construction verbatim with G replaced by G' . The only point that requires some care is to extend the conclusion of Lemma 7.2 to G' . This is straightforward: we apply Lemma 7.2 to G to obtain a set of cuts D . For each $B \in D$, we let B' be the cut of G' obtained by adding, for every $e \in B$, the up to two ‘parallel’ copies of e from any ladders we attached to face-boundaries containing e in the proof of Lemma 5.3. Moreover, for each of the finitely many Γ -orbits O of vertices of G' , we pick a representative $v_O \in O$, and let B_O be the cut comprising the edges incident with v_O . Then the collection $D' := \{B' \mid B \in D\} \cup \bigcup_{O \in V(G')/\Gamma} B_O$ of all these cuts of G' is finite, and their translates generate $\mathcal{B}(G')$ as the reader can easily check. \square

I do not know if Lemma 7.1 remains true if one forbids loops and/or parallel edges:

Problem 7.1. *Suppose Γ admits a finitely generated, covariantly planar Cayley graph. Must Γ admit a finitely generated, covariantly planar Cayley graph with no infinite face-boundaries?*

8 Proof of Theorem 1.1 and Corollary 1.2

We now put the above results together to prove the main results of the paper:

Proof of Theorem 1.1. For the implication (B) \rightarrow (D), let G be a covariantly planar, finitely generated, Cayley graph of a group Γ . By Lemma 5.4 there is a covariant embedding $\sigma : |G| \rightarrow \mathbb{S}^2$. Applying Lemma 5.6 we obtain a properly discontinuous action $\Gamma \curvearrowright (\mathbb{S}^2 \setminus \sigma(\Omega(G)))$ extending $\Gamma \curvearrowright G$. Since $\Gamma \curvearrowright G$ is faithful, so is $\Gamma \curvearrowright (\mathbb{S}^2 \setminus \sigma(\Omega(G)))$. We claim that $\mathbb{S}^2 \setminus \sigma(\Omega(G))$ is one of the four 2-manifolds in the statement of (D). This is indeed the case, because of the well-known fact, first observed by Hopf [21], that $\Omega(G)$ is either homeomorphic to the Cantor set or it contains at most 2 points.

For the implication (A) \rightarrow (B), let $\Gamma \curvearrowright X$ be a properly discontinuous faithful action where $X \subseteq \mathbb{S}^2$ is a 2-manifold. Then X is uniquely planar by Lemma 6.1, and so Lemma 6.3 yields the desired Cayley graph G .

The implication (B) \rightarrow (C) is Lemma 7.1.

As (D) \rightarrow (A) and (C) \rightarrow (B) are trivial, we have proved the equivalence of all four conditions. \square

Theorem 8.1 ([32, p. 299]). *Let $p : \tilde{S} \rightarrow S$ be a regular covering of the (topologically) finite (Riemann) surface S , where \tilde{S} is planar. Let F be the group of deck transformations on \tilde{S} . There is a Koebe group Γ , with invariant component D , and there is a homeomorphism $h : \tilde{S} \rightarrow D$, so that $h_* : F \rightarrow \Gamma$ is an isomorphism.*

Here, a surface S is *topologically finite*, if it is homeomorphic to the interior of a compact 2-manifold with or without boundary. I will not repeat all the details (which can be found e.g. in [32]) of the definition of *Koebe group*, as they are not needed in this paper. What matters for us is that every Koebe group is a function group by definition.

Proof of Corollary 1.2. The backward direction is obvious: if Γ is a function group, then its action on an invariant component of its domain of discontinuity is as desired.

For the forward direction, let $\Gamma \curvearrowright X$ be a faithful, properly discontinuous action by orientation-preserving homeomorphisms on the surface $X \subseteq \mathbb{S}^2$. We would like to apply Theorem 8.1 with $\tilde{S} = X$ and $S = X/\Gamma$, but $\Gamma \curvearrowright X$ need not be free, and X/Γ need not be topologically finite. Therefore, we will use the above results to find a better action of Γ .

Indeed, the implication (A) \rightarrow (C) of Theorem 1.1 yields a planar Cayley graph G of Γ , with a Γ -covariant embedding $\sigma : |G| \rightarrow \mathbb{S}^2$ every face of which is bounded by a cycle. Moreover, when applying Lemma 6.3 to prove the implication (A) \rightarrow (B), we use the last statement of Lemma 6.3 to ensure that σ is orientation-preserving. We then apply the last statement of Lemma 5.6 to obtain a co-compact, properly discontinuous action $\Gamma \curvearrowright Y$ with $Y := \mathbb{S}^2 \setminus (\sigma(\Omega(G)))$, where the set Z of points of Y with non-trivial stabiliser contains at most one point from the interior of each face of σ . Restricting $\Gamma \curvearrowright Y$ to $X' := Y \setminus Z$ thus yields a free action, in other words, X' regularly covers $S := (Y \setminus Z)/\Gamma$. Since $\Gamma \curvearrowright Y$ is co-compact, it is easy to see that S is topologically finite. Thus we can apply Theorem 8.1 to deduce that Γ is a Koebe group, and in particular a function group. \square

As a corollary, we deduce that the finiteness condition in Theorem 8.1 can be dropped if one only wants an algebraic rather than a ‘geometric’ isomorphism between F and some Koebe group (the finiteness condition cannot be dropped in general without this relaxation):

Corollary 8.2. *Let $p : \tilde{S} \rightarrow S$ be a regular covering of a surface S , where \tilde{S} is planar. Let F be the group of deck transformations on \tilde{S} . Then F is isomorphic to a Koebe group (equivalently, to a function group).*

Proof. By the definition of deck transformations, F acts freely by homeomorphisms on \tilde{S} , and by the definition of a covering this action is properly discontinuous. Thus we can apply Corollary 1.2 to deduce that F is isomorphic to a Koebe group. \square

8.1 Geometrizing the action

As mentioned in the introduction, any action $\Gamma \curvearrowright X$ as in (D) can be ‘geometrised’ to obtain an action by isometries on a smooth manifold with the same properties. This can be proved as follows. Recall that every 2-manifold is homeomorphic to a smooth manifold [35]. Thus it just remains to endow X with a Γ -invariant metric d . A standard way to do this (which I learnt from a mathoverflow post⁶ by Misha Kapovich.) is by endowing X with a Γ -invariant Riemannian metric g , and letting d be the corresponding distance function. This g can be constructed by first constructing an arbitrary Riemannian metric on the quotient orbifold X/Γ , e.g. using a partition of unity, and then lifting g to X . Here, we used the well-known observation of Thurston that for every properly discontinuous action on a manifold, the quotient space is an orbifold, see e.g. [4, Proposition 20].

⁶<https://mathoverflow.net/questions/251627/proper-discontinuity-and-existence-of-a-fundamental-domain>

9 Determining X

The main aim of this section is to prove Theorem 1.4.

We say that an end $\omega \in \Omega(X)$ is an *accumulation end* with respect to an action $\Gamma \curvearrowright X$, if ω is an accumulation point of some orbit $\Gamma x, x \in X$ for the extension $\Gamma \curvearrowright |X|$ of the action to the Freudenthal compactification of X .

Given topological spaces $Y \subset X$, we write $\lim Y$ for the set of accumulation points of Y in X .

Proposition 9.1. *Let X be an arc-connected locally compact space, $\Gamma \curvearrowright X$ a properly discontinuous action, and $\omega \in \Omega(X)$. Then the following are equivalent:*

- (i) *there is $x \in X$ such that $\omega \in \lim \Gamma x$;*
- (ii) *there is a compact $K \subset X$ such that $\omega \in \lim \Gamma K$;*
- (iii) *for every compact $L \subset X$, we have $\omega \in \lim \Gamma L$; and*
- (iv) *for every $x \in X$, we have $\omega \in \lim \Gamma x$.*

Proof. The implications (i) \rightarrow (ii) and (iii) \rightarrow (iv) are trivial since points are compact. So is (iv) \rightarrow (i). To show (ii) \rightarrow (iii), let L be any compact subset of X , let A be an K - L arc, and let $K' := K \cup A \cup L$. Note that K' is compact. Let U be a basic open neighbourhood of ω in $|X|$, and recall that $U \cap X$ is a component of $X \setminus S$ for some compact $S \subset X$. Since $\omega \in \lim \Gamma K$, there are infinitely many elements of ΓK , and hence of $\Gamma K'$ meeting U . By definition (ii) of a properly discontinuous action, at most finitely many of these elements meet the compact set S , and therefore infinitely many of them are contained in U . In particular, U meets ΓL . As this holds for any basic open neighbourhood U of ω , we deduce that $\omega \in \lim \Gamma L$ as claimed. \square

Let $\Omega^\Lambda(X) \subseteq \Omega(X)$ denote the set of accumulation ends of X with respect to $\Gamma \curvearrowright X$. In other words, $\Omega^\Lambda(X)$ is the limit set of the extension of the action $\Gamma \curvearrowright X$ to $|X|$.

The proof of the following statement is standard; the main idea goes back to Hopf [21].

Theorem 9.2. *Let $\Gamma \curvearrowright X$ be a properly discontinuous action on an arc-connected metrizable space X . If the space $\Omega^\Lambda(X) \subseteq \Omega(X)$ of accumulation ends contains more than 2 points, then it is homeomorphic to the Cantor set.*

Proof. Suppose $\Omega^\Lambda(X)$ contains three distinct accumulation ends ω, χ, ψ . We will show that none of them is an isolated point of $\Omega^\Lambda(X)$. Let $K \subset X$ be a compact set separating X into subspaces $U_\omega \ni \omega, U_\chi \ni \chi, U_\psi \ni \psi$. We may assume that K is connected, for otherwise we can enlarge it by a set of arcs joining its finitely many components.

Suppose, for a contradiction, that ω has a neighbourhood U containing no other end in $\Omega^\Lambda(X)$. By choosing a larger K if needed, we may assume that $U = U_\omega$. Pick some $x \in K$ and $g \in \Gamma$ such that $gx \in U$ and $gK \cap K = \emptyset$, which exists since U contains an accumulation end and $\Gamma \curvearrowright X$ is properly discontinuous. Then $gK \subset U$ because gK is connected and it meets U . Therefore, gK disconnects no two ends of X living outside U_ω , because any two such ends live

in the connected space $X \setminus U$. But gK disconnects $g\omega, g\chi, g\psi$ from each other, and so at least two of them live in U , contrary to our assumption.

This proves that if $\Omega^\Lambda(X)$ contains more than two points, then it contains no isolated point. It is easy to check that $\Omega^\Lambda(X)$ is closed (therefore compact since $\Omega^\Lambda(X) \subset |X|$). As it is totally disconnected, it is homeomorphic to the Cantor set by Proposition 2.1. \square

Corollary 9.3. *Let $\Gamma \curvearrowright X$ be a faithful, properly discontinuous action on a 2-manifold $X \subseteq \mathbb{S}^2$ such that all ends of X are accumulation ends. Then X is homeomorphic to the sphere, the plane, the open annulus, or the Cantor sphere.*

Proof. By Lemma 9.2, either $\Omega(X)$ contains at most 2 ends, or it is homeomorphic to the Cantor set. By Theorem 3.8, X is homeomorphic to the sphere, the plane, the open annulus, or the Cantor sphere, if it has 0, 1, 2, or a Cantor space of ends, respectively. \square

Remark: The assumption that all ends of X are accumulation ends does not imply that the action is co-compact, or that $|\Omega(X)| = |\Omega(\Gamma)|$; consider for example the action of \mathbb{Z} on \mathbb{R}^2 by addition in one coordinate.

This example is also relevant in Theorem 1.4 from the introduction, which we restate here for convenience:

Theorem 9.4. *Let $\Gamma \curvearrowright X$ be a properly discontinuous action on a metrizable, arc-connected, locally compact space X . Then the canonical extension $\Gamma \curvearrowright (X \cup \Omega^V(X))$ of the action to the non-accumulation ends of X is properly discontinuous.*

For the proof of this we need to introduce the following notions and a lemma.

We say that a subspace $K \subset X$ is a *separator* of a topological space X , if $X \setminus K$ can be written as the union $U \cup U^c$ of two disjoint, non-empty, open subspaces. Note that U, U^c are also closed in $X \setminus K$ in this case. In this case, we say that U, U^c are *sides* of K . We remark that K does not uniquely determine its sides if $X \setminus K$ has more than two components. The following lemma is a variant of [25, Theorem 1, Chapter V, §46, VIII]⁷. We provide a proof for convenience.

Lemma 9.5. *Let X be an arc-connected, metrizable topological space, and $K_Y, K_Z \subset X$ two disjoint closed and connected separators of X with sides Y, Y^c and Z, Z^c , respectively. Then at least one of the sets $Y \cap Z, Y \cap Z^c, Y^c \cap Z, Y^c \cap Z^c$ is empty.*

Proof. Since K_Y, K_Z are connected, and disjoint, each of them must be contained in a side of the other. We may assume without loss of generality that $K_Z \subseteq Y$ and $K_Y \subseteq Z$ as the two sides of each separator are interchangeable.

Suppose, for a contradiction, that all our sets $Y \cap Z, Y \cap Z^c, Y^c \cap Z, Y^c \cap Z^c$ are non-empty, and pick two points $p \in Y^c \cap Z^c, q \in Y \cap Z$. Let A be a p - q arc in X , and let x be the first point of A in $K_Y \cup K_Z$, which exists since $K_Y \cup K_Z$ is closed. If $x \in K_Y$, then the sub-arc of A from p to x is a Z^c - Z arc in $X \setminus K_Z$. Otherwise, we have $x \in K_Z$, and the sub-arc of A from p to x is a Y^c - Y arc in $X \setminus K_Y$. In both cases we obtain a contradiction as we have split an arc as a union of two disjoint open sets. \square

⁷I would like to thank Max Pitz for this reference.

Proof of Theorem 9.4. Consider the subspace $X' := X \cup \Omega^V(X)$ of $|X|$. To show that the canonical extension $\Gamma \curvearrowright X'$ of $\Gamma \curvearrowright X$ is properly discontinuous, we will show that for every $x, y \in X'$ there are open neighbourhoods $O_x \ni x, O_y \ni y$ such that $\{g \in \Gamma \mid gO_y \cap O_x \neq \emptyset\}$ is finite and apply definition (iii) of a properly discontinuous action.

We claim that there is a sequence $(K_n^x)_{n \in \mathbb{N}}$ of compact, connected separators of X' with corresponding sides (U_n, U_n^c) such that $(U_n)_{n \in \mathbb{N}}$ is a base of open neighbourhoods of x (which we will later plug into Lemma 9.5). Indeed, if $x \in \Omega^V(X)$ we can let $(U_n)_{n \in \mathbb{N}}$ consist of basic open neighbourhoods of x , and let K_n^x their associated compact sets. We can assume each K_n^x is connected because otherwise we can join its finitely many components with arcs of X . If $x \in X$, then there is a local base $(C_n)_{n \in \mathbb{N}}$ of compact neighbourhoods of x since X is locally compact. Let $\{\partial C_n\}$ be their frontiers, let K_n^x be a compact connected set containing ∂C_n , and let $U_n := C_n \setminus K_n^x$. Then U_n is open in $X' \setminus K_n^x$ as desired, and so is $U_n^c := X' \setminus (K_n^x \cup U_n)$. Replacing x by y throughout, we analogously define $(K_n^y)_{n \in \mathbb{N}}$ and the sides $\{V_n, V_n^c\}$.

As $\Gamma \curvearrowright X$ is properly discontinuous, the set $\{g \in \Gamma \mid gK_i^x \cap K_j^y \neq \emptyset\}$ is finite for every $i, j \in \mathbb{N}$. Thus we can apply Lemma 9.5 to the separators gK_i^x, K_j^y of X' for every $i, j \in \mathbb{N}$ and all but finitely many $g \in \Gamma$. If the transporter set (U_i, V_j) is finite for some i, j , then we are done by (iii). If not, then either $|(U_i \mid K^y)| = \infty$ for every i , or $|(V_i \mid K^x)| = \infty$ for every i . If the former is the case, then $x \in \lim K^y$, and if the latter is the case, then $y \in \lim K^x$. This leads to a contradiction as X' contains no accumulation end, and no point of X is an accumulation point of an orbit under $\Gamma \curvearrowright X$ (here we tacitly used Proposition 9.1). □

10 Relationship to planar groups and planar discontinuous groups

Droms et. al. [11] provided an example of a Cayley graph G which is planar but admits no embedding in \mathbb{S}^2 for which the natural action on G by its group is realized by homeomorphisms of \mathbb{S}^2 . Here, we provide such an example with the stronger property that its group does not admit any faithful action on \mathbb{S}^2 (or any 2-manifold $X \subseteq \mathbb{S}^2$) by homeomorphisms.

For this, we will find finite groups R, P of homeomorphism of \mathbb{S}^2 , containing involutions b_r, b_p respectively, such that in any faithful action $R \curvearrowright \mathbb{S}^2$ the homeomorphism b_r reverses the orientation of \mathbb{S}^2 , and in any faithful action $P \curvearrowright \mathbb{S}^2$ the homeomorphism b_p preserves the orientation. In addition, we will display planar Cayley graphs G_R, G_P of the two groups in which b_r and b_p are contained in the generating set, and thus appear as edges. By amalgamating any two such groups R, P along the 2-element subgroups spanned by b_r, b_p respectively, we obtain a group with a planar Cayley graph G , which can be obtained by embedding infinitely many copies of G_R and G_P glued along their amalgamated edges corresponding to b_r, b_p .

We can choose R to be the alternating group A_4 . This group has presentations $\langle r, g \mid r^3, g^3, (rg)^2 \rangle$ and $\langle k, r \mid k^2, r^3, (kr)^3 \rangle$, where $r = (1, 2, 3), g = (2, 3, 4), k = (1, 3)(2, 4)$ in permutation cycle notation. The first presentation

shows that R is generated by elements of odd order. Easily, every element of odd order preserves the orientation in any action $R \curvearrowright \mathbb{S}^2$. These two observations combined yield that every element of R preserves the orientation in any action $R \curvearrowright \mathbb{S}^2$.

The second presentation gives rise to a planar Cayley graph of R , namely the truncated tetrahedron⁸, and its generator k is an involution. We can thus let G_R be this Cayley graph and $b_r = k$.




Figure 2: The standard Cayley graph of $\mathbb{Z}_4 \times \mathbb{Z}_2$.

We can choose P to be the group $\mathbb{Z}_4 \times \mathbb{Z}_2$. We claim that every planar Cayley graph $G = \text{Cay}(P, S)$ of this group has a subgraph $G' = \text{Cay}(P, S')$ with $S' \subseteq S$ isomorphic to the *prism* depicted in Figure 2. To see this, note that S must contain the element $(1, 0)$ or its inverse $(3, 0)$ because these elements are not contained in the span of the remaining elements of P by a parity argument. If S also contains any of the elements $(1, 1)$ or $(3, 1)$, then G contains the graph of Figure 3 as a subgraph. But that graph is not planar, because it is isomorphic to the complete bipartite graph $K_{4,4}$, which contains the Kuratowski graph $K_{3,3}$ as a subgraph. Thus S also contains one of the elements $(0, 1)$ or $(2, 1)$. It is easy to check that any choice S' of two generators as above yields a Cayley graph $G' \subseteq G$ isomorphic to the graph of Figure 2. (In fact, it is not hard to show that $G = G'$, but we will not need this.)

Applying Lemma 6.3 to any action $P \curvearrowright \mathbb{S}^2$, we obtain a planar Cayley graph H of P such that the canonical action $P \curvearrowright H$ extends to $P \curvearrowright \mathbb{S}^2$. By the above claim, H contains the Cayley graph G' of Figure 2 as a subgraph. Note that $P \curvearrowright G'$ extends to $P \curvearrowright H$, and from there to $P \curvearrowright \mathbb{S}^2$. Recall that we would like to choose an involution $b_P \in P$ that reverses orientation in any action $R \curvearrowright \mathbb{S}^2$. Note that P contains exactly 3 involutions, namely $(2, 0)$, $(0, 1)$, $(2, 1)$. Of these involutions, only the first one is contained in a subgroup of P isomorphic to \mathbb{Z}^4 (P has exactly two such subgroups). Let $b_P = (0, 1)$. Then no matter which generating set of P gives rise to the Cayley graph G' , our involution b_P will correspond to a vertex of G' that is not in the same monochromatic 4-cycle as the identity. But then the action of b_P on G' reverses the cyclic order of the edges incident with any vertex. Since $P \curvearrowright G'$ extends to $P \curvearrowright \mathbb{S}^2$, this means that b_P reverses the orientation of \mathbb{S}^2 in $P \curvearrowright \mathbb{S}^2$, which was an arbitrary action of P on \mathbb{S}^2 . Finally we can just let $G_P := G'$ be our choice of Cayley graph of P .

⁸See <http://weddslist.com/groups/cayley-plat/index.html> for a figure.




Figure 3: A non-planar Cayley graph of $\mathbb{Z}_4 \times \mathbb{Z}_2$. This graph is the complete bipartite graph $K_{4,4}$: the vertices at the same horizontal/vertical level form the two partition classes.

Let now $\Gamma := R \star_{b_r=b_p} P$ be the amalgamation product of R and P over the subgroups spanned by b_r, b_p , respectively. Then Γ cannot act faithfully on \mathbb{S}^2 , because by the above discussion its element $b_r = b_p$ would have to both preserve and reverse the orientation in any such action. Moreover, we can obtain a planar Cayley graph of Γ by recursively glueing copies of G_R and G_P along their edges corresponding to b_r, b_p . This proves in particular

Proposition 10.1. *The groups satisfying the conditions of Theorem 1.1 form a proper subfamily of the groups admitting finitely generated Cayley graphs.*

Next, we will show that the groups satisfying the conditions of Theorem 1.1 form a proper superfamily of the finitely generated planar discontinuous groups, i.e. the finitely generated groups acting faithfully and properly discontinuously on \mathbb{R}^2 . For the expert reader this may be a straightforward consequence of Corollary 1.2.

The example we will consider is $\Gamma := \mathbb{Z} \times \mathbb{Z}_3$. The standard Cayley graph of Γ is easily seen to be planar and 3-connected. Thus Γ satisfies condition (B) of Theorem 1.1. It acts faithfully and properly discontinuously on the open annulus but not, as we will now show, on \mathbb{R}^2 .

Suppose, for a contradiction, that $\Gamma \curvearrowright \mathbb{S}^2$ is such an action, and apply Lemma 6.3 to obtain a planar Cayley graph G of Γ and an embedding $\sigma : G \rightarrow \mathbb{R}^2$ such that $\sigma(V(G))$ has no accumulation points in \mathbb{R}^2 , and the canonical action $\Gamma \curvearrowright G$ extends to $\Gamma \curvearrowright \mathbb{S}^2$.

Our first claim is that there is a cycle C of G fixed by the action of the subgroup Δ corresponding second factor \mathbb{Z}_3 in the definition of Γ . To see this, let e, g, g^2 denote the elements of Δ , and let A be an e - g path in G . Let $X := A \cup gA \cup g^2A$ be the subgraph of G comprising its Δ -translates. By Babai's Lemma 3.6 applied to the (free) action of Δ on X , there is a contraction X/F of X , with $F \subseteq X$ connected, with X/F being a Cayley graph of Δ . The only Cayley graph of the group Δ with 3 elements is the triangle. Let B be a path in F joining the endvertices of the two edges of X/F incident with F . Then the 3 edges of X/F combined with the 3 Δ -translates of B form the desired cycle C .

Next, we prove that there is a further cycle D of G fixed by the action of Δ such that the closures of the insides of $\sigma(C)$ and $\sigma(D)$ are disjoint. For this, note that there are infinitely many translates $hC, h \in \Gamma$ of C , and only finitely

many of them intersect C by local finiteness. If $C, D = hC$ are disjoint, and any of them is embedded inside the other by σ , then by iterating h or h^{-1} we can obtain an infinite set of translates embedded inside C because $\Gamma \curvearrowright G$ extends to $\Gamma \curvearrowright \mathbb{S}^2$. But this would contradict the fact that $\sigma(V(G))$ has no accumulation points in \mathbb{R}^2 , and so our claim is proved.

Let P be a path from C to D in G such that the interior of P avoids $C \cup D$, which exists since G is connected. If the three translates P, gP, g^2P of P via Δ are pairwise disjoint, then combined with C and D they decompose \mathbb{R}^2 into 6 domains (Figure 4), exactly one of which is non-compact and delimited by two of P, gP, g^2P and a subpath of each of C, D (here, we are tacitly interpreting the aforementioned subgraphs of G as subspaces of \mathbb{S}^2 by using our embedding σ). But the three domains of this form are permuted by the action of Δ , leading to a contradiction as no homeomorphism can map a compact set onto a non-compact set.




Figure 4: The situation when P, gP, g^2P (displayed in green, if colour is shown) are pairwise disjoint. We show that this situation can be achieved by choosing P appropriately.

Thus it remains to show that we can choose P above so that its translates P, gP, g^2P are pairwise disjoint. To see this, assume they are not, and notice that the subgraph $H := P \cup gP \cup g^2P$ of G is connected in this case. Since $\Delta \subset \Gamma$ acts freely on G , it does so on H . By Babai’s Lemma 3.6 again, there is a contraction H/F of F , with $F \subseteq H$ connected, which is a Cayley graph of Δ , and again H/F can only be a triangle. Notice that F must contain exactly one of the six endpoints of the three paths P, gP, g^2P incident with C , and exactly one incident with D by the definitions. Let x, y be those endpoints, and let P' be a x - y path contained in F . Then P' is a path from C to D in G , and its translates gP', g^2P' are contained in the translates gF, g^2F of F , and are hence disjoint to each other and to P' by the choice of F . This completes our proof.

We finish this section with an example of a finitely generated Kleinian group that is not one of the groups of Theorem 1.1; even more, it does not admit any planar Cayley graph. In fact most Kleinian groups have this property, but we present the following explicit example—for which I thank B. Bowditch (private communication)—for the non-expert reader.

Proposition 10.2. *There is a finitely generated Kleinian group that does not admit a planar Cayley graph.*

Proof (sketch). Let S be a closed orientable surface of genus at least 1, and let S' be obtained from S by removing a topological disc bounded by a (contractible) simple closed curve C . Let D be the space obtained by identifying four copies of S' along C . It is shown in [5, Proposition 5.1] that D admits a metric d such that (D, d) embeds (as a metric subspace) in a complete hyperbolic 3-manifold W and W retracts onto D . Thus W and D have the same fundamental group, and we let $H := \pi_1(D) \cong \pi_1(W)$. It is well-known that the fundamental group of every complete hyperbolic 3-manifold, in particular H , is Kleinian.

By the Seifert–van Kampen theorem, H is the amalgamation product of four free groups of rank 2 (or two copies of the fundamental group of the double-torus) along an infinite cyclic subgroup. This remark allows us to visualise its canonical Cayley graph as a union of regular tilings of the plane, glued along certain common lines. It is not hard to deduce from this that H is 1-ended, and torsion-free. Since 1-ended Cayley graphs are 3-connected [3, Lemma 2.4], if H admits a planar Cayley graph G , then G has an essentially unique embedding into \mathbb{S}^2 by Whitney’s Theorem 3.3. Then (D) of Theorem 1.1 implies that H is the fundamental group of a compact 2-orbifold M , and since H is torsion-free M is a manifold, i.e. a closed orientable surface. It is an exercise to show that $\pi_1(D)$ is not the fundamental group of a closed orientable surface, for example by noticing that the aforementioned Cayley graph is not quasi-isometric to any regular tiling of the plane. \square

11 Relationship to orbifold fundamental groups

Let $\Gamma \curvearrowright X$ be an action as in Theorem 1.1 (D). Then the quotient space $O := X/\Gamma$ is a good compact 2-orbifold (see e.g. [4, Proposition 20]). It follows from the standard theory of covering spaces ([20, Proposition 1.40]) that

$$\Gamma \cong \mathring{\pi}_1(O)/p_*(\pi_1(X)), \tag{3}$$

where $\mathring{\pi}_1(O)$ denotes the *orbifold fundamental group* of O (see e.g. [40]) and p_* the canonical projection from X to O . Does (3), combined with the classification of compact 2-orbifolds, provide information about the groups of Theorem 1.1? I do not think so. If $X = \mathbb{S}^2$ or $X = \mathbb{R}^2$, then Γ is one of Maschke’s finite groups of isometries of the sphere, or the Fuchsian and Kleinian groups respectively; these cases are already well-understood. The most interesting case is where X is the Cantor sphere \mathcal{C} . It is not too hard to prove that $\Pi := \pi_1(\mathcal{C})$ is the free group of countably infinite rank. Then (3) tells us that if Γ acts faithfully and properly discontinuously on \mathcal{C} , then $\Gamma \cong \mathring{\pi}_1(O)/\Pi$ for some good compact 2-orbifold.

Have we learned anything about Γ ? Note that every countable group G can be written as a quotient F/F' where both F, F' are free, because defining G via a group presentation $\langle \mathcal{S} \mid \mathcal{R} \rangle$ provides such an expression, with F being the free group with generating set \mathcal{S} and F' its smallest normal subgroup spanned by \mathcal{R} . With easy modifications, we may always assume that both F, F' have infinite rank, and are hence isomorphic to Π .

Here is a concrete example. Let $\Gamma = \mathbb{Z}$, let X be the open annulus, and let \mathbb{Z} act on X by a shift, so that $O = X/\mathbb{Z}$ is the torus T . Then (3) says that $\mathbb{Z} = \pi_1(T)/\pi_1(X) = \mathbb{Z}^2/\mathbb{Z}$. However, not every quotient of \mathbb{Z}^2 by a normal subgroup H isomorphic to \mathbb{Z} yields a group as in Theorem 1.1: the quotient

$\mathbb{Z}^2 / \langle (3, 3) \rangle$ is a finite abelian group of rank 2, which can act by isometries on the torus but not on \mathbb{S}^2 ; this can be proved using the techniques of Section 10.

12 Acting on \mathbb{R}^3

The aim of this section is to prove Theorem 1.3, which states that every group Γ as in Theorem 1.1 admits a faithful properly discontinuous action on \mathbb{R}^3 . We remark that in the orientation-preserving case Theorem 1.3 immediately follows from Corollary 1.2. The proof below follows a different method.

We will use the notion of planar presentations, and the associated almost planar Cayley complexes, from [18, 19], about which we need to prove a couple of additional facts. We start by recalling some terminology.

We say that a Cayley complex Z is *almost planar*, if it admits a map $\rho : Z \rightarrow \mathbb{S}^2$ in which the 2-simplices of Z are *nested* in the following sense. We say that two 2-simplices of Z are nested, if the images of their interiors under ρ are either disjoint, or one is contained in the other, or their intersection is the image of a 2-cell bounded by two parallel edges corresponding to an involution s in the generating set \mathcal{S} defining Z .⁹ We call such a ρ an *almost planar map* of Z .

Every Cayley complex Z in this section is *finitely presented*, i.e. its defining presentation is finite.

The following is proved in [19, Theorem 5.6]:¹⁰

Lemma 12.1. *Let $G = \text{Cay}(\Gamma, \mathcal{S})$ be a finitely generated Cayley graph and $\sigma : G \rightarrow \mathbb{S}^2$ a covariant embedding. Then there is a (finitely presented) Cayley complex Z of Γ with 1-skeleton G , and an almost planar map $\rho : Z \rightarrow \mathbb{S}^2$ such that $\rho(G)$ coincides with σ .*

We call a Cayley complex *standard*, if the closure of each of its 2-cells is homeomorphic to a disc. The Cayley complexes provided by Lemma 12.1 are standard by their construction.

We let Z^1 denote the 1-skeleton of a Cayley complex Z , which is a Cayley graph by the definitions. We write Z^2 for the set of 2-cells of Z . Given $C \in Z^2$, the boundary ∂C of C is a subgraph of Z^1 , and if Z is standard then ∂C is always a cycle.

We call two 2-cells B, C of a Cayley complex Z *equivalent*, if $\partial B = \partial C$. We obtain the corresponding *simplified* Cayley complex by removing all but one representative from each equivalence class of 2-cells of Z .

Let Z be a standard, simplified Cayley complex, and let $\rho : Z \rightarrow \mathbb{S}^2$ be an almost planar map. Then every 2-cell $C \in Z^2$ defines two *sides* C_1, C_2 , namely the two components of $\mathbb{S}^2 \setminus \rho(\partial C)$. Given $B, C \in Z^2$, we say that B is *maximally nested* in the side C_1 of C , if $\rho(\partial B) \subset \overline{C_1}$ and there is no 2-cell $B' \neq B$ of Z such that $\rho(\partial B')$ separates $\rho(\partial B)$ from $\rho(\partial C)$ in \mathbb{S}^2 .

Lemma 12.2. *Let $\rho : Z \rightarrow \mathbb{S}^2$ be an almost planar map of a standard, simplified Cayley complex Z such that $\rho(Z^1)$ is covariant, let $C \in Z^2$, and let $M(C)$ be*

⁹The third option can be dropped by considering the *modified* Cayley complex in the sense of [27], i.e. by representing involutions in \mathcal{S} by single, undirected edges.

¹⁰The difference between this formulation and the one of [1, Theorem 6.6] is only in the terminology used

the set of 2-cells of Z maximally nested in a side I of C . Then either $M(C)$ is finite, or there is exactly one accumulation point of $M(C)$ in $|Z|$.

Proof. Easily, we may assume that for every $D \in Z^2$, $\rho(D)$ is a topological disc bounded by $\rho(\partial D)$. Thus if $\rho(Z^1)$ is given, then $\rho(D)$ is one of the two components of $\mathbb{S}^2 \setminus \rho(\partial D)$. Note that the almost planarity of ρ is not affected if we modify it so as to map $\rho(D)$ to the other component of $\mathbb{S}^2 \setminus \rho(\partial D)$. Therefore, we may assume without loss of generality that

$$\rho(D) \subset I \text{ for every } D \in M(C). \quad (4)$$

Suppose, to the contrary, that there are distinct ends $\omega, \chi \in \Omega(Z)$ in the closure of $\bigcup M(C)$ in $|Z|$. By Corollary 5.4, we may assume that ρ extends to an embedding $\rho' : |Z| \rightarrow \mathbb{S}^2$. Let K be a cycle in Z^1 such that $\rho'(K)$ separates $\rho'(\omega)$ from $\rho'(\chi)$, which exists by Lemma 4.3. Write its edge-set $E(K)$ as a sum (with respect to addition in the cycle space of Z^1) $E(K) = \sum_{1 \leq i \leq k} E(R_i)$, where each R_i is a cycle of Z^1 induced by a relator in the presentation defining Z . (We can choose R_i to be cycle rather than a closed walk because Z is standard.)

We claim that for every cycle R of Z^1 induced by a relator, $\rho'(R)$ does not separate $\rho'(\omega)$ from $\rho'(\chi)$. For if it does, then both components R_1, R_2 of $\mathbb{S}^2 \setminus \rho'(R)$ contain infinitely many images of elements of $M(C)$. In particular each of R_1, R_2 contains a boundary of an element of $M(C)$ not equal to R . This contradicts the fact that the elements of $M(C)$ are maximally nested in I , as $\rho'(R)$ separates one of them from C .

Combining this claim with Lemma 4.4 implies that $\rho'(K)$ does not separate $\rho'(\omega)$ from $\rho'(\chi)$, and we have reached a contradiction to the existence of ω, χ . \square

Lemma 12.3. *Let $\rho : Z \rightarrow \mathbb{S}^2$ be an almost planar map of a standard, simplified Cayley complex Z , such that every face boundary of Z^1 bounds a 2-cell of Z , and $\rho(Z^1)$ extends to an embedding of $|Z^1|$. Let I be a side of a 2-cell C of Z , and let M be the set of 2-cells of Z maximally nested in I . If M is finite, then $C \cup \bigcup M$ is homeomorphic to \mathbb{S}^2 , and if M is infinite, then $C \cup \bigcup M$ is homeomorphic to \mathbb{R}^2 .*

Proof. Let H be the 1-skeleton of $C \cup \bigcup M$, and notice that H is a subgraph of Z^1 , and therefore planar. We claim that

$$\begin{aligned} &\text{the boundary of each face } F \text{ of } H \text{ (with respect to the embedding } \rho(H)) \\ &\text{coincides with the boundary of exactly one 2-cell in } \{C\} \cup M. \end{aligned} \quad (5)$$

To see this, note first that ∂C bounds a face of H by the definitions. Let $B \in M$, and suppose ∂B does not bound a face of H . This means that there is an edge e of H inside $\rho(B)$, and so there is some $A \in M$ containing e in its boundary. Then A, B will contradict the almost planarity of ρ unless $\rho(\partial B)$ separates $\rho(\partial A)$ from $\rho(\partial C)$. The latter however contradicts the assumption that A is maximally nested in I . Thus ∂B must bound a face of H .

Conversely, consider a face F of H . If F is a face of Z^1 too, then ∂F bounds a 2-cell of Z by our assumptions. Otherwise, pick an edge $e \in E(Z^1)$ with $\rho(e) \subset F$. Note that e lies on the boundary of some 2-cell A of Z if and only if its label appears in at least one of the defining relators.

If this is the case, let A' be the 2-cell maximally nested in I such that $\rho(A) \subseteq \rho(A')$, which exists because there are only finitely many 2-cells A' with $\rho(A) \subseteq \rho(A') \subseteq I$. Then $A' \in M$, and so $\partial A'$ has no edges in F , and $\rho(A')$ meets F . This implies that $\partial A' = \partial F$.

If, on the other hand, e does not lie on the boundary of any 2-cell, then at least one of its endvertices v does not lie on ∂F . This is true because an edge of the Cayley graph Z^1 that lies on no relator cycle separates Z^1 , and so it cannot have both end-vertices on the cycle ∂F . In this case, we let A be a 2-cell of Z incident with v and repeat the above arguments to find $A' \in M$ with $\partial A' = \partial F$.

Since our Cayley complex is simplified, no other $K \in M$ can satisfy $\partial K = \partial F$.

This completes the proof of (5). If M is finite, then H is a finite plane graph, hence attaching a 2-cell along each of its faces yields a homeomorph S of \mathbb{S}^2 . As Z is standard, S is homeomorphic to $C \cup \bigcup M$ as desired. If M is infinite, then H has exactly one end by Lemma 12.3. Recall that we are assuming that $\rho(Z^1)$ extends to an embedding $\sigma : |Z^1| \rightarrow \mathbb{S}^2$. This induces an embedding $\sigma' : |H| \rightarrow \mathbb{S}^2$, and attaching, as above, a 2-cell along each of the faces of σ' yields a homeomorph S of $\mathbb{S}^2 \setminus \sigma'(\omega) \cong \mathbb{R}^2$, where ω is the unique end of H , and again S is homeomorphic to $C \cup \bigcup M$. □

We can now prove the main result of this section.

Proof of Theorem 1.3. Let G be a covariantly planar Cayley graph of Γ with no infinite face-boundaries, as provided by (C) of Theorem 1.1. Let $\sigma : |G| \rightarrow \mathbb{S}^2$ be an embedding such that $\sigma(G)$ is covariant, provided by Corollary 5.4. Let Z' be an almost planar Cayley complex with $G = Z^1$, and let $\rho : Z' \rightarrow \mathbb{S}^2$ be an almost planar map, provided by Lemma 12.1, such that $\rho(G)$ coincides with $\sigma(G)$. Recall that we may assume that Z' is standard, and, by adding the corresponding relators to the presentation defining Z' if necessary, that

$$\text{every face boundary of } G \text{ with respect to } \rho \text{ bounds a 2-cell of } Z'. \quad (6)$$

Let Z be the corresponding simplified Cayley complex.

We now modify the almost planar embedding $\rho(Z)$ of Z into an embedding $\tau : Z \rightarrow \mathbb{D}^3$, where \mathbb{D}^3 is the closed Euclidean ball of radius 1 in \mathbb{R}^3 , such that τ coincides with ρ when restricting both maps to G , where we think of \mathbb{S}^2 as the boundary of \mathbb{D}^3 . Let \mathring{Z}^2 denote the set of interiors of 2-cells of Z . To define this τ , we just need to specify the image $\tau(C)$ of each $C \in \mathring{Z}^2$. By the definition of the almost planarity of ρ , we can easily choose τ so that $\tau(C) \cap \mathbb{S}^2 = \emptyset$, and $\tau(C) \cap \tau(D) = \emptyset$ for each two distinct $C, D \in \mathring{Z}^2$. Indeed, we can define $\tau(C_1), \tau(C_2), \dots$ inductively for any enumeration $(C_n)_{n \in \mathbb{N}}$ of \mathring{Z}^2 , exploiting the fact that the circle $\rho(\partial(C_i)) \subset \mathbb{S}^2$ does not cross $\rho(\partial(C_j))$. Moreover, by making $\tau(C_i)$ sufficiently small (e.g. contained in a ball of radius $10 \text{diam}(\rho(\partial(C_i)))$ around $\rho(\partial(C_i))$), we can assume that the images $\tau(C_i)$ have no accumulation points in the interior of \mathbb{D}^3 .

Since Z is just a locally finite 2-complex, it is easy to see that our τ is an embedding of Z into \mathbb{D}^3 . Our plan is to extend Z , and our action $\Gamma \rightarrow Z$, to a 3-complex homeomorphic to $\mathbb{D}^3 \setminus \rho(\Omega(G))$. Let us first consider the case where the first alternative of Lemma 12.2 holds for every $C \in Z^2$, i.e. $M(C)$ is always finite. We will later extend our construction to the general case.

Under this assumption, Lemma 12.3 says that for every 2-cell C of Z , and each of the two sides I_1, I_2 of C , the set M_i of 2-cells maximally nested in I_i together with C form a homeomorph $S_i = S_i(C)$ of \mathbb{S}^2 unless ∂C bounds a face of G in the embedding σ . In the latter case, M_i is empty for one of the sides, I_1 say, of C . Letting Z_F be the 2-complex obtained from Z by adding the faces of G as 2-cells (which we can since $\partial F \subset Z^1$ for every face F of G), we observe that $C \cup F$ is homeomorphic to \mathbb{S}^2 whenever ∂C bounds the face F of G , and we let $S_1(C) = C \cup F$ in this case. Note that if $D \in Z^2$ lies on the boundary of $S_i(C)$, then C lies on the boundary of $S_{2-i}(D)$ by the definitions. This implies that each $D \in Z^2$ lies on the boundary of exactly two elements of $\mathcal{S} := \bigcup_{C \in Z^2} \{S_1, S_2\}$, where we also used (6).

Let T be the 3-complex obtained from Z_F by adding a 3-cell $T(S)$ with simple boundary S for every $S \in \mathcal{S}$. Then τ extends into an embedding ϕ of T into \mathbb{D}^3 , because S bounds a homeomorph S' of \mathbb{R}^3 in \mathbb{D}^3 for every $S \in \mathcal{S}$ by (a rather easy version of) the generalised Schoenflies theorem [7, 33], and we can let $\phi(T(S)) = S'$. Easily, $\Gamma \curvearrowright Z$ extends to a cellular action $\Gamma \curvearrowright T$. Then $\Gamma \curvearrowright T$ is (faithful and) properly discontinuous, because every cellular action on a CW-complex with finite stabilisers of cells is properly discontinuous [23, Theorem 9, (2)=(10)].

Moreover, by the construction of ϕ , we have $Y := \phi(T) = \mathbb{D}^3 \setminus \sigma(\Omega(G))$, and so we can think of $\Gamma \curvearrowright T$ as an action $\Gamma \curvearrowright Y$ since ϕ is an embedding, and so Y is homeomorphic to T (here we used our assumption that $\tau(C_i)$ have no accumulation points in the interior of \mathbb{D}^3). By restricting that action to $\mathbb{D}^3 \setminus \mathbb{S}^2$ we thus obtain a faithful, properly discontinuous action on the interior of \mathbb{D}^3 , hence on its homeomorph \mathbb{R}^3 .

We remark that $\Gamma \curvearrowright Y$ is co-compact, but its restriction to $\mathbb{D}^3 \setminus \mathbb{S}^2$ is not.

We now consider the general case, where $M(C)$ in Lemma 12.2 is possibly infinite for some of the 2-cells $C \in Z^2$. For every such C , and every side I_i of C for which $M(C) = M_i(C)$ is infinite, Lemma 12.2 yields a unique end $\omega_i(C)$ accumulating $M_i(C)$. We are going to use $\omega_i(C)$ in order to triangulate the interior of the sphere formed by $\tau(\bigcup M_i(C))$ and $\rho(\omega_i(C))$. For this, given any edge e in the boundary of an element of $M_i(C)$, we add to Z_F two new 1-cells e_0, e_1 , each joining a distinct endvertex of e to $\omega_i(C)$; we also add $\omega_i(C)$ to Z_F as a 0-cell. In addition, we add to Z_F a 2-cell bounded by the triangle e, e_0, e_1 . We let W denote the 2-complex obtained from Z_F after adding all those cells (for every side I_i of a $C \in Z^2$ with infinite $M_i(C)$).

Note that for every $D \in \{C\} \cup M_i(C)$ where $M_i(C)$ is infinite, the set $N \subset W \setminus Z$ of newly added 2-cells sharing an edge with D combined with D forms a homeomorph S_D of \mathbb{S}^2 . We define \mathcal{S} as above, except that we now let $S_i(C) = S_C$ whenever one or both sides I_1, I_2 of $C \in Z^2$ has infinite $M_i(C)$. As above, we construct a 3-complex T by adding a 3-cell $T(S)$ with simple boundary S for every $S \in \mathcal{S}$. Easily, we can extend τ into an embedding $\tau' : W \rightarrow \mathbb{D}^3$, and from there to an embedding ϕ of T into \mathbb{D}^3 , with image $Y = \phi(T) = \mathbb{D}^3 \setminus \sigma(\Omega(G) \setminus \Omega')$, where Ω' is the set of 0-cells $\omega_i(C)$ in $W^0 \setminus Z_F^0$. The extension $\Gamma \curvearrowright T$ of our action is still cellular, but it now fails to have finite vertex stabilisers because of the 0-cells in Ω' . In fact, $\Gamma \curvearrowright T$ is not properly discontinuous because any point in Ω' accumulates orbits. Therefore, we restrict our action to the topological subspace $T' \subseteq T$ obtained from T by removing Ω' . This T' is not a cell complex anymore, so we need a different argument to

prove that $\Gamma \curvearrowright T'$ is properly discontinuous. But this is not hard: we apply Lemma 5.2 as in the proof of Lemma 5.6. The cover \mathcal{U} of T we choose for this comprises the sets $U_v, v \in V(G)$ obtained as the union of all open cells of any dimension in T' that have a vertex v of G in their boundary. Restricting $\Gamma \curvearrowright Y$ to $\mathbb{D}^3 \setminus \mathbb{S}^2$ again we obtain the desired action. □

Acknowledgements

I thank Caroline Series for suggesting studying the connection between planar groups and Kleinian groups some years ago. I thank Misha Kapovich for suggesting using Maskit's Theorem 8.1 to prove Corollary 1.2. I thank Matthias Hamann for the proof of Lemma 7.2, and Max Pitz for numerous remarks.

I am very grateful to Brian Bowditch for many helpful discussions.

References

- [1] G. N. Arzhantseva and P.-A. Cherix. On the Cayley graph of a generic finitely presented group. *Bulletin of the Belgian Mathematical Society - Simon Stevin*, 11(4):589–601, December 2004.
- [2] L. Babai. Some applications of graph contractions. *J. Graph Theory*, 1(2):125–130, 1977.
- [3] L. Babai. The growth rate of vertex-transitive planar graphs. In *Proceedings of the eighth ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 564–573. Soc. for Industrial and Applied Mathematics, 1997.
- [4] J. E. Borzellino. *Riemannian Geometry of Orbifolds, PhD thesis*. University of California, Los Angeles, 1992.
- [5] B. H. Bowditch and G. Mess. A 4-Dimensional Kleinian Group. *Trans. Am. Math. Soc.*, 344(1):391–405, 1994.
- [6] L. E. J. Brouwer. On the structure of perfect sets of points. *Proc. Koninklijke Akademie van Wetenschappen*, 12:785–794, 1910.
- [7] M. Brown. A proof of the generalized Schoenflies theorem. *Bull. Amer. Math. Soc.*, 66:74–76, 1960.
- [8] Warren Dicks and Martin J. Dunwoody. *Groups acting on graphs*. Cambridge University Press, 1989.
- [9] Reinhard Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005. Electronic edition available at: <http://www.math.uni-hamburg.de/home/diestel/books/graph.theory>.
- [10] C. Droms. Infinite-ended groups with planar Cayley graphs. *J. Group Theory*, 9(4):487–496, 2006.
- [11] C. Droms, B. Servatius, and H. Servatius. Connectivity and planarity of Cayley graphs. *Beitr. Algebra Geom.*, 39(2):269–282, 1998.

- [12] M.J. Dunwoody. The accessibility of finitely presented groups. *Invent. math.*, 81:449–457, 1985.
- [13] M.J. Dunwoody. Planar graphs and covers. Preprint available at <http://www.personal.soton.ac.uk/mjd7/>, 2009.
- [14] Eric Charpentier, Etienne Ghys, Annick Lesne (eds.). *The Scientific Legacy of Poincaré*, volume 36 of *History of mathematics*. American Mathematical Soc, 2010.
- [15] Felix Klein. *Vergleichende Betrachtungen über neuere geometrische Forschungen*. Verlag von Andreas Deichert, Erlangen, 1872.
- [16] A. Georgakopoulos. Graph topologies induced by edge lengths. In *Infinite Graphs: Introductions, Connections, Surveys. Special issue of Discrete Math.*, volume 311 (15), pages 1523–1542, 2011.
- [17] A. Georgakopoulos. The planar cubic Cayley graphs. *Memoirs of the AMS*, 250(1190), 2017.
- [18] A. Georgakopoulos and M. Hamann. The planar Cayley graphs are effectively enumerable I: consistently planar graphs. To appear in *Combinatorica*.
- [19] A. Georgakopoulos and M. Hamann. The planar Cayley graphs are effectively enumerable II. Preprint 2018.
- [20] A. Hatcher. *Algebraic Topology*. Cambridge Univ. Press, 2002.
- [21] H. Hopf. Enden offener Räume und unendliche diskontinuierliche Gruppen. *Comment. Math. Helv.*, 16:81–100, 1944.
- [22] W. Imrich. On Whitney’s theorem on the unique embeddability of 3-connected planar graphs. In *Recent Adv. Graph Theory, Proc. Symp. Prague 1974*, pages 303–306. 1975.
- [23] M. Kapovich. A note on properly discontinuous actions. “<https://www.math.ucdavis.edu/~kapovich/EPR/prop-disc.pdf>”.
- [24] B. Krön. Infinite faces and ends of almost transitive plane graphs. Preprint.
- [25] Kazimierz Kuratowski. *Topology, Volume II*. Academic Press, 1968.
- [26] H. Levinson and B. Maskit. Special embeddings of Cayley diagrams. *J. Combin. Theory (Series B)*, 18(1):12–17, 1975.
- [27] Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory*. Springer Science & Business Media, January 2001.
- [28] A. Marden. Geometrically finite Kleinian groups and their deformation spaces. In *Discrete groups and automorphic functions, W. Harvey (ed.)*, pages 259–293. Academic Press, London, 1977.
- [29] Albert Marden. *Outer Circles*. Cambridge University Press, 2007.

- [30] H. Maschke. The representation of finite groups, especially of rotation groups of three and four dimensional space, by Cayley's color diagrams. *Amer. J. Math.*, 18:156–194, 1896.
- [31] B. Maskit. On the classification of Kleinian groups: I—Koebe groups. *Acta Math.*, 135:249–270, 1975.
- [32] Bernard Maskit. *Kleinian Groups*, volume 287 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1988.
- [33] B. Mazur. On embeddings of spheres. *Bull. Amer. Math. Soc.*, 65:59–65, 1959.
- [34] B. Mohar. Tree amalgamation of graphs and tessellations of the cantor sphere. *Journal of Combinatorial Theory. Series B*, 96(5):740–753, 2006.
- [35] Edwin E. Moise. *Geometric Topology in Dimensions 2 and 3*. Graduate Texts in Mathematics. Springer, 1977.
- [36] Ken'ichi Ohshika. *Discrete Groups*, volume 207 of *Iwanami Series In Modern Mathematics*. American Mathematical Society, 2002.
- [37] T. Radö. Über den Begriff der Riemannschen Fläche. *Acta Sci. Math. (Szeged)*, 2(1):96–114, 1925.
- [38] I. Richards. ON THE CLASSIFICATION OF NONCOMPACT SURFACES. *Trans. Am. Math. Soc.*, 106(2):259–269, 1963.
- [39] R.B. Richter and C. Thomassen. 3-connected planar spaces uniquely embed in the sphere. *Trans. Am. Math. Soc.*, 354:4585–4595, 2002.
- [40] P. Scott. The Geometries of 3-Manifolds. *Bull. London Math. Soc.*, 15(5):401–487, 1983.
- [41] C. Series. A crash course on kleinian groups. *Rend. Istit. Mat. Univ. Trieste*, 37((1-2)):1–38, 2006.
- [42] C. Thomassen. Tilings of the Torus and the Klein Bottle and Vertex-Transitive Graphs on a Fixed Surface. *Trans. Am. Math. Soc.*, 323(2):605–635, 1991.
- [43] C. Thomassen. The Jordan-Schoenflies Theorem and the Classification of Surfaces. *American Mathematical Monthly*, 99:116–130, 1992.
- [44] C. Thomassen and A. Vella. Graph-like continua, augmenting arcs, and Menger's theorem. *Combinatorica*, 29:595–623, 2008.
- [45] W. P. Thurston. Three dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.
- [46] William Thurston. *The geometry and topology of three-manifolds*. Princeton lecture notes, 1980.
- [47] T. W. Tucker. Finite Groups Acting on Surfaces and the Genus of a Group. *J. Combin. Theory (Series B)*, 34:82–98, 1983.

- [48] H. Whitney. Congruent graphs and the connectivity of graphs. *American J. of Mathematics*, 54(1):150–168, 1932.
- [49] H. Zieschang, E. Vogt, and H.-D. Coldewey. *Surfaces and planar discontinuous groups. Revised and expanded transl. from the German by J. Stillwell.* Lecture Notes in Mathematics 835. Springer-Verlag, 1980.