



ALGEBRAIC COMBINATORICS


Rémi Bottinelli, Laura Grave de Peralta & Alexander Kolpakov

Telescopic groups and symmetries of combinatorial maps

Volume 3, issue 2 (2020), p. 483-508.

[<http://alco.centre-mersenne.org/item/ALCO_2020__3_2_483_0>](http://alco.centre-mersenne.org/item/ALCO_2020__3_2_483_0)

© The journal and the authors, 2020.
Some rights reserved.

 This article is licensed under the
CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL LICENSE.
<http://creativecommons.org/licenses/by/4.0/>

Access to articles published by the journal *Algebraic Combinatorics* on
the website <http://alco.centre-mersenne.org/> implies agreement with the
Terms of Use (<http://alco.centre-mersenne.org/legal/>).



Algebraic Combinatorics is member of the
Centre Mersenne for Open Scientific Publishing
www.centre-mersenne.org



Telescopic groups and symmetries of combinatorial maps

Rémi Bottinelli, Laura Grave de Peralta
& Alexander Kolpakov

ABSTRACT In the present paper, we show that many combinatorial and topological objects, such as maps, hypermaps, three-dimensional pavings, constellations and branched coverings of the two-sphere admit any given finite automorphism group. This enhances the already known results by Frucht, Cori–Machi, Širáň–Škoviera, and other authors. We also provide a more universal technique for showing that “any finite automorphism group is possible”, that is applicable to wider classes or, in contrast, to more particular sub-classes of said combinatorial and geometric objects. Finally, we show that any given finite automorphism group can be realised by sufficiently many non-isomorphic such entities (super-exponentially many with respect to a certain combinatorial complexity measure).

1. INTRODUCTION

A combinatorial (oriented, labelled) map is a triple $M = (D; R, L)$ where D is a non-empty finite set (called the set of darts) and R and L are two permutations of D with $L^2 = \text{id}$. The orbits of L are conventionally called the edges of M , the orbits of R are its vertices, and the orbits of $R^{-1}L$ are its faces. The map M is called connected if the group $\langle R, L \rangle$ acts transitively on D . Unless otherwise stated, we shall assume all maps to be connected.

A topological (oriented) map $M = (\Sigma_g; \Gamma)$ is an oriented (connected) genus $g \geq 0$ surface with an embedded graph Γ such that the complement $\Sigma_g \setminus \Gamma$ is a collection of disjoint topological discs. By providing a labelling on the half-edges of Γ (thus defining its labelled darts), and thus obtaining a labelled topological map, one can recover the permutations L and R , so that L encodes the identification of half-edges into edges, and R encodes the positive cyclic order of half-edges around each vertex. Vice versa, provided a combinatorial map, one can recover its corresponding topological labelled counterpart by creating the faces (which are discs) by following the cycles of $R^{-1}L$, and then identifying their boundaries by using L . For more details, cf. [8, 13, 23].

One can define a more elaborate class of combinatorial objects (and the corresponding topological objects) such as hypermaps [6]. A triple $H = (D; R, L)$, where D is a non-empty finite set of darts and R, L are permutation of D , is called an (oriented, labelled) hypermap. The orbits of L are called the hyper-edges of H , the orbits of

Manuscript received 30th January 2019, revised 7th November 2019, accepted 14th November 2019.

KEYWORDS. automorphism, free product, free group, hypermap, ribbon graph, symmetry.

ACKNOWLEDGEMENTS. The authors were supported by the Swiss National Science Foundation project no. PP00P2-144681/1.

R are its hyper-faces. A hypermap H is connected whenever the group $\langle R, L \rangle$ acts transitively on D (which will be our standing assumption).

A hypermap also naturally appears in the setting of an orientable genus g surface Σ_g and a graph Γ embedded in Σ_g that satisfies the following properties:

- (1) the complement $\Sigma_g \setminus \Gamma$ is a union of topological discs called faces,
- (2) the faces are properly two-colourable (e.g. into black and white), i.e. faces of the same colour intersect only at vertices of Γ , and
- (3) the corners of the white faces are labelled with the numbers $1, 2, 3, \dots$ in some fashion, and a black face corner label is equal to the adjacent white face corner label, when moving clockwise around their common vertex.

Then $H = (\Sigma_g; \Gamma)$ is an oriented labelled topological hypermap.

The correspondence between the topological and combinatorial definitions is as follows:

- (1) each disjoint cycle of R is obtained from recording the corner labels of a white face in a counter-clockwise direction,
- (2) each disjoint cycle of L is obtained from recording the corner labels of a black face in a counter-clockwise direction,
- (3) each disjoint cycle of $R^{-1}L$ is obtained from recording the labels around a vertex in a counter-clockwise direction.

We remark that condition (3) above is a consequence of (1) and (2).

The set of face labels becomes the set of darts of H , the white faces become hyper-faces of H and the black faces become hyper-edges of H . Thus, the combinatorial and topological definitions of an oriented labelled hypermap agree.

If $L^2 = \text{id}$, then each bigon in the hypermap $H = (D; R, L)$ can be interpreted as a pair of darts pointing in opposite directions, and thus H becomes a map, as defined above.

We say that two oriented labelled (hyper-)maps $M_1 = (D; R_1, L_1)$ and $M_2 = (D; R_2, L_2)$ are isomorphic if, in the combinatorial setting, there exists a permutation T of D such that $TR_1 = R_2T$ and $TL_1 = L_2T$. In the topological setting, two oriented labelled (hyper-)maps $M_1 = (\Sigma_g; \Gamma_1)$ and $M_2 = (\Sigma_g; \Gamma_2)$ are isomorphic if there exists an orientation-preserving homeomorphism $\tau : \Sigma_g \rightarrow \Sigma_g$ such that $\tau(\Gamma_1) = \Gamma_2$ and the labelling of the corresponding half-edges is respected.

A rooted isomorphism will require only the root (a dedicated labelled dart) of one (hyper-)map to be carried to the root of another.

Finally, an isomorphism is not required to respect the dart labelling, nor the roots.

The above definition allows us to generalise the setting of maps to higher-dimensional objects, the so-called pavings. Namely, as defined in [1], a three-dimensional oriented combinatorial map or, simply, a (combinatorial) paving, is a quadruple $P = (D; R, L, V)$, where D is a non-empty set of darts and R, L, V are permutations of D such that $H_P = (D; R, L)$ is a map (not necessarily connected), and

- (1) the product LV is an involution,
- (2) the product VR^{-1} is an involution,
- (3) none of the above involutions have fixed points.

A paving P is connected if the group $\langle L, R, V \rangle$ acts transitively on D . The notion of (labelled, rooted) isomorphism for oriented combinatorial pavings is analogous to the one for combinatorial maps.

We may also think of P as a quadruple $P = (D; L, S, T)$, where D is the set of darts and L, S, T are its involutions without fixed points. In this case it is easy to see that letting $V = LS$ and $R = TLS$ produces the initial definition. As in the case of two-dimensional maps, a combinatorial paving P has a topological realisation

which, however, is not always a three-dimensional manifold (however, it's always a pseudo-manifold).

In order to assemble an oriented cellular complex M_P , as described in [20], we first produce its underlying map $H_P = (D; R, L)$, and realise each connected component of H as a topological map, i.e. as a surface Σ^i with an embedded graph Γ^i , $i = 1, 2, \dots, m$, having labelled half-edges. Each surface Σ^i represents the boundary of a handle-body B^i , and then the handle-bodies B^i become identified along their boundaries in order to produce a labelled oriented cellular complex representing P topologically. Indeed, the faces of Σ^i 's defined by the permutation $R^{-1}L$ are identified in accordance with the permutation V , and the conditions (1), (2), and (3) above ensure that one face cannot be identified to multiple disjoint counterparts (implied by (1) and (2)), and edges or faces cannot bend onto themselves (implied by (3)). Also, conditions (1) and (2) ensure that M_P is an orientable topological space.

There are other generalisations of maps, hypermaps and pavings, such as constellations, cf. the monograph [15] for more information and references.

One of the basic questions is understanding possible symmetries, or automorphisms (i.e. unrooted self-isomorphisms), of any of the above defined objects. Those can be understood by means of building a one-to-one correspondence between a class of rooted (hyper-)maps \mathcal{M} (or isomorphism classes of maps) on a set of dart D and (usually, torsion-free) subgroups (or their conjugacy classes) of a given single group Δ^+ . This correspondence will associate to each map $M \in \mathcal{M}$ a subgroup $H_M \subset \Delta^+$ of index $|D|$. The origins of this technique draw back to the paper by Jones and Singerman [13], and have been developed more in the recent works by Breda, Mednykh and Nedela [3], Mednykh and Nedela [16, 17] for the purpose of solving Tutte's problem of (hyper-)map classification, cf. also [4, 5].

Let us consider the case of maps, as described in [3, 13]. Namely, the rooted maps on n darts (where the root is always supposed to be marked 1) are in a one-to-one correspondence with index n free subgroups of $\Delta^+ = \mathbb{Z} * \mathbb{Z}_2$. Indeed, each free subgroup $H < \Delta^+$ of index n produces a set of cosets $D = \Delta^+ / H$ of cardinality n , which can be considered as a set of darts. The root dart here is the identity coset. A subgroup of Δ^+ is torsion-free if and only if it is free, as a consequence of Kurosh's theorem. Thus $\Delta^+ = \mathbb{Z} * \mathbb{Z}_2 \cong \langle \sigma \rangle * \langle \alpha \rangle$ acts on D transitively, and its generators σ and α give rise to permutations R and L acting transitively on D . Thus, we obtain a map M_H corresponding to a free subgroup $H < \Delta^+$. Vice versa, given a map $(D; R, L)$, we have a homomorphism $S : \Delta^+ \rightarrow \langle R, L \rangle$ by setting $S(\sigma) = R$, $S(\alpha) = L$. The homomorphism S defines an action of Δ^+ on D , and the subgroup corresponding to M is $H_M = \text{Stab}(1) < \Delta^+$.

The above correspondence between the free subgroups of Δ^+ and rooted maps can be extended to the case of hypermaps (with $\Delta = \mathbb{Z} * \mathbb{Z} \cong F_2$), or the so-called (p, q) -hypermaps (with $\Delta^+ = \mathbb{Z}_p * \mathbb{Z}_q$ [5]), or 3-dimensional maps (also called pavings [1, 20], with $\Delta^+ = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ [4]).

The isomorphisms classes of all aforementioned objects correspond to the conjugacy classes of free subgroups of Δ^+ [13, Theorem 3.7]. The symmetries (i.e. unrooted self-isomorphisms) of a (hyper-)map M corresponding to a subgroup $H < \Delta^+$ form a group isomorphic to $N(H)/H$, where $N(H) = \{g \in \Delta^+ \mid gHg^{-1} = H\}$ is the normaliser of H in Δ^+ [13, Theorem 3.8].

In the sequel we shall study a more abstract question, namely the property of free products of cyclic groups being "telescopic", cf. Definition 2.1. Such a free product T being telescopic allows us to realise any finite group Γ as the "symmetry group" $N(H)/H$ of a suitable finite-index subgroup $H \leq T$. Thus, one of our main results is the following statement, cf. Theorem 3.7.

THEOREM. *Any free product of at least two non-trivial cyclic groups is freely telescopic, except for the infinite dihedral group $D_\infty \cong \mathbb{Z}_2 * \mathbb{Z}_2$.*

If we allow the index of H to be sufficiently large, depending on the cardinality of Γ , then a great deal of same index subgroups H with $N(H)/H \cong \Gamma$ can be obtained, cf. Theorem 6.1. More precisely, the following holds.

THEOREM. *Let T be a finite free product of cyclic groups, different from $\mathbb{Z}_2 * \mathbb{Z}_2$. Then for any finite group Γ , there exist constants $A > 1$, $B > 0$ and $M \in \mathbb{N}$ such that for all $d \geq M$ the set $F(T, \Gamma, d) = \{\text{free subgroups } H \leq T \text{ of index } \leq d \text{ with } N_T(H)/H \cong \Gamma, \text{ up to conjugacy}\}$ has cardinality $\geq A^{Bd \log d}$.*

Finally, we translate our group-theoretic statements into the combinatorial language of (hyper-)maps and pavings, cf. Theorems 7.1–7.3. Such a transition from combinatorics to groups, to combinatorics again is an integral part of our approach. First, we want to obtain some information about symmetries of a sufficiently complicated combinatorial object. Next, we translate our questions about symmetries into a question about the existence of (torsion-free) subgroups of a free product of cyclic groups with some condition on their normalisers. This condition is formulated in terms of combinatorial automorphisms of the subgroup’s Schreier graph, by analogy to the approach introduced in [14, 24]. The symmetries of the corresponding Schreier graphs appear more amenable to combinatorial analysis, which finally provides us with the desired results both in group-theoretic and combinatorial terms.

REMARK. Soon after a draft of this paper appeared on the arXiv, the authors were notified by Gareth A. Jones that his paper [12] contains similar results for a wider class of group. In particular, by [12, Theorem 3], all hyperbolic (extended) triangle groups are shown to be “finitely abundant”, which is equivalent to being telescopic for the non-compact ones among them. Also, all subgroups produced in [12, Theorem 3] are, in fact, torsion-free. The methods used in [12] and in our paper differ substantially, as well as the emphasis in our work is on the quantitative aspects, such as counting of combinatorial objects with given symmetries.

2. PRELIMINARIES

We first establish the necessary notation and provide some basic definitions. Let G be a group, and H be a subgroup of G . Let $N_G(H) = \{g \in G : gHg^{-1} = H\}$ denote the normaliser of H in G .

DEFINITION 2.1. *We say that a group T is telescopic if for every finite group Γ there exists a finite-index subgroup $H \leq T$ such that $N_T(H)/H \cong \Gamma$.*

DEFINITION 2.2. *If in the above definition we can always choose H to be a free subgroup of T , we say that T is freely telescopic.*

DEFINITION 2.3. *A (di-)graph is a tuple $(V, E, \iota : E \rightarrow V, \tau : E \rightarrow V)$, where V is the set of vertices, E is the set of directed edges and ι , resp. τ , assigns to each edge e its initial vertex (or origin) $\iota(e)$, resp. its terminal vertex (or terminus) $\tau(e)$. We shall write $\text{St}_+ v := \{e \in E : \iota(e) = v\}$ and $\text{St}_- v := \{e \in E : \tau(e) = v\}$. A morphism of graphs $\phi : (V_1, E_1, \iota_1, \tau_1) \rightarrow (V_2, E_2, \iota_2, \tau_2)$ consists of a pair of maps $\phi_V : V_1 \rightarrow V_2$ and $\phi_E : E_1 \rightarrow E_2$ such that $\iota_2 \phi_E = \phi_V \iota_1$ and $\tau_2 \phi_E = \phi_V \tau_1$. If \mathcal{G} is a (di-)graph, then $V\mathcal{G}$ will denote its set of vertices, and $E\mathcal{G}$ will be its set of edges.*

A graph is labelled by the elements of a set S if a map $\mu : E \rightarrow S$ is given. Let us write S -digraph for an S -labelled digraph. A labelled S -digraph \mathcal{G} is called folded (cf. [21]) if, for any $v \in V\mathcal{G}$, the restrictions of μ to $\text{St}_+ v$ and $\text{St}_- v$ are injective,

and \mathcal{G} is called regular if they are bijective. A morphism of S -digraphs is a morphism of digraphs satisfying $\mu_2\phi_E = \mu_1$.

Recall that if G is a group generated by a set S , and $H \leq G$ is a subgroup of G , then the Schreier graph $\text{Sch}_{G,S}(H)$ of H is an S -labelled regular graph having vertex set the right cosets of H , and an edge $H \cdot g \xrightarrow{s} H \cdot gs$ from $H \cdot g$ to $H \cdot gs$, labelled s , for each element $s \in S$ and each coset $H \cdot g$. The Cayley graph $\text{Cay}(G, S)$ of G is defined as the Schreier graph of the trivial subgroup $\{\text{id}\}$ of G . Both Cayley and Schreier graphs are generally considered with a basepoint: the coset $H \cdot e = H$. When graphs with basepoints are considered, their morphisms are assumed to send basepoints to basepoints.

If a graph \mathcal{G} has label set $S \subset G$, for some group G , and a vertex $v \in V\mathcal{G}$ is specified, then the set

$$L(\mathcal{G}, v) := \{\text{labels of loops at } v, \text{ evaluated in } G\}$$

forms a group, called the language of \mathcal{G} at v (an empty loop gives the identity of G , paths concatenation corresponds to taking products, and reversing paths corresponds to taking inverses). Let $\text{ev} : \{\text{words in } S \sqcup S^{-1}\} \rightarrow G$ denote the evaluation map, so that $L(\mathcal{G}, v) = \text{ev} \circ \mu$ (“loops at v ”).

We also present a few key results that we make use of, most of which can be found in [14, 21, 24]. First, recall Kurosh’s theorem.

THEOREM 2.4 (Kurosh’s Subgroup Theorem). *Let G_1, \dots, G_n be groups, and $G := *_{i=1}^n G_i$ be their free product. Then, any subgroup $H \leq G$ of G has the form:*

$$(*_{i=1}^n *_{j=1}^{m_i} w_{ij} H_{ij} w_{ij}^{-1}) * F(X)$$

where each H_{ij} is a subgroup of G_i , $F(X)$ is a free subgroup generated by a subset X of G , and w_{ij} is an element of G , for $1 \leq i \leq n$ and $1 \leq j \leq m_i$.

Proof. See, for instance, the monograph [18, pp. 56–57]. □

LEMMA 2.5 ([14, Lemma 7.5]). *If \mathcal{G} is an S -digraph with S a subset of a group G , then for any vertices v_1, v_2 connected by a path $p : v_1 \rightsquigarrow v_2$ with $g := \text{ev} \circ \mu(p)$:*

$$L(\mathcal{G}, v_1) = gL(\mathcal{G}, v_2)g^{-1}.$$

Proof. If l is a loop at v_2 , then plp^{-1} is a loop at v_1 , and $\text{ev} \mu(plp^{-1}) = g \text{ev} \mu(l)g^{-1}$, so that the conjugation by g maps elements of $L(\mathcal{G}, v_2)$ to elements of $L(\mathcal{G}, v_1)$, and thus $gL(\mathcal{G}, v_2)g^{-1} \subseteq L(\mathcal{G}, v_1)$. Symmetrically, $g^{-1}L(\mathcal{G}, v_1)g \subseteq L(\mathcal{G}, v_2)$, and the result follows. □

LEMMA 2.6 ([14, Lemma 4.2]). *If \mathcal{A}, \mathcal{B} are S -digraphs and \mathcal{B} is folded, then for any vertices $v \in V\mathcal{A}$ and $u \in V\mathcal{B}$, there exists at most one morphism of S -digraphs $\phi : \mathcal{A} \rightarrow \mathcal{B}$ satisfying $\phi(v) = u$.*

Proof. Follow paths; cf. [14, Lemma 4.2] or [21, 5.1 (c)]. □

Let $\mathbb{Z}_{p_i} = \langle s_i \mid s_i^{p_i} \rangle$ denote the cyclic group of order $p_i \in \mathbb{N} \cup \{\infty\}$, for $p_i \geq 2$, while setting $p_i = \infty$ yields \mathbb{Z} (which will be our standard notation for the rest of the paper).

LEMMA 2.7 ([22, Theorem 1.2]). *Let $G := \mathbb{Z}_{p_1} * \dots * \mathbb{Z}_{p_n}$ be a free product of cyclic groups with generators $S = \{s_1, \dots, s_n\}$, assuming $\mathbb{Z}_{p_i} = \langle s_i \rangle$. There is a bijection between the sets*

- $A :=$ “subgroups of G ” and
- $B :=$ “connected, regular S -digraphs with a basepoint, such that for any s_i , the edges labelled by s_i form cycles of length dividing $p_i < \infty$, up to isomorphism”.

Moreover, free (equivalently, torsion-free) subgroups correspond to graphs with s_i -labelled cycles of length exactly $p_i < \infty$, and under this equivalence, the index of a subgroup equals the number of vertices in the corresponding graph.

If, in a labelled digraph as in the lemma above, a cycle labelled by s_i has length a proper divisor of $p_i < \infty$, such a cycle will be called degenerate, following [22]. An element of B is a (G, S) -Schreier graph, which is called *non-degenerate*, if it contains no degenerate cycle.

Sketch of proof. The Schreier graph of the quotient $\text{Sch}_{G,S}(\cdot)$ provides one direction of the equivalence, while the language at the root $L(\cdot)$ proves the other. In order to show the equivalence “free subgroup” \leftrightarrow “no degenerate cycle”, one uses the fact that a torsion element in G must be conjugate to an element of one of the factors, and vice versa. \square

The equivalence of Lemma 2.7 can actually be generalised to arbitrary finitely generated groups if one does not care about freeness: this is done (using a slightly different language) in [24, Theorem 3.5].

PROPOSITION 2.8. *Let $\mathcal{S}(H) := \text{Sch}_{G,S}(H)$. Then $\text{Aut}(\mathcal{S}(H)) \cong N_G(H)/H$.*

In the above statement, Aut denotes the group of automorphisms of a labelled digraph without basepoint.

Proof. Let $N := N_G(H)$, and let us consider the following map:

$$\begin{aligned} \Phi : N &\longrightarrow \text{Aut}(\mathcal{S}(H)) \\ n &\longmapsto (\phi_n : Hg \mapsto Hng) \end{aligned}$$

Then, Φ is a well-defined surjective group homomorphism with kernel exactly H .

It is routine to check that Φ is well-defined and is a homomorphism: this fact depends on N being the normaliser of H . Let us verify its surjectivity.

Let ϕ be an element of $\text{Aut}(\mathcal{S}(H))$ and let Hg be the image of H under ϕ . Since ϕ is an automorphism, we have that

$$H = L(\mathcal{S}(H), He) = L(\phi(\mathcal{S}(H)), \phi(He)) = L(\mathcal{S}(H), Hg).$$

Also, we know that $L(\mathcal{S}(H), Hg) = g^{-1}L(\mathcal{S}(H), He)g$, since changing the basepoint changes the language by conjugation, as in Lemma 2.5. This implies $H = g^{-1}Hg$, and thus $g \in N$. Since $\mathcal{S}(H)$ is a regular graph, there is a unique morphism sending H to Hg , so that $\phi = \phi_g$. Therefore, Φ is surjective.

Finally, let us verify that $\ker \Phi = H$. Let n be an element of $\ker \Phi$. Then $\phi_n = \text{id}_{\mathcal{S}(H)}$, which implies that $Hn = H$, and $n \in H$. Conversely, for any $h \in H$, ϕ_h is readily seen to be the identity map, since $Hh = H$. Therefore, $\ker \Phi = H$ and the claim follows by the first isomorphism theorem. \square

3. FREE PRODUCTS OF CYCLIC GROUPS

In this section we show that any free product of (non-trivial) cyclic groups with at least two factors is freely telescopic, with the obvious exception of $D_\infty \cong \mathbb{Z}_2 * \mathbb{Z}_2$, the infinite dihedral group.

From now on, let Γ denote a finite group with generating set S . Let T be a finite free product of cyclic groups $\mathbb{Z}_{p_1} * \dots * \mathbb{Z}_{p_n}$, and X be the natural choice of its generators (one per factor). We always assume that $p_i \geq 2$ and, if $T = \mathbb{Z}_p * \mathbb{Z}_q$, also that $p \geq q \geq 2$, while $p \geq 3$.

We will proceed as follows, in order to prove that any T as above is freely telescopic, or equivalently, that any finite group Γ is isomorphic to a quotient $N_T(H)/H$ for H a free finite-index subgroup of T .

3.1. PLAN OF PROOF.

- (1) Some algebraic arguments (Lemmas 3.5–3.6) using Kurosh’s Subgroup Theorem reduce the problem to free products of the form $T = \mathbb{Z}_p * \mathbb{Z}_q$ ($p \geq 3, q \geq 2$) and $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$.
- (2) Then, Lemma 2.7 translates the question of finding a free subgroup H of T into finding a non-degenerate (T, X) -graph \mathcal{G} .
- (3) By Lemma 2.8, the condition that $N_T(H)/H$ be isomorphic to Γ is equivalent to the condition that the automorphism group of \mathcal{G} be isomorphic to Γ . Hence, the initial problem effectively reduces to finding a non-degenerate (T, X) -graph with a given automorphism group.
- (4) Starting with the Cayley graph $\text{Cay}(\Gamma, S)$ of the finite group Γ , we replace its edges and vertices by certain pieces of non-degenerate (T, X) -graphs (defined in Section 4.3), so that the automorphism group is preserved, while obtaining a valid Schreier graph for a finite-index free subgroup H of T .

3.2. BASIC CASES. We start first by proving that $\mathbb{Z}_p * \mathbb{Z}_q$, with $p \geq 3, q \geq 2$, and $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ are freely telescopic. These are the “base cases” for the general statement that follows in Theorem 3.7.

PROPOSITION 3.1. *The free product $\mathbb{Z}_p * \mathbb{Z}_q$ is freely telescopic for any $p \geq 3$ and $q \geq 2$.*

PROPOSITION 3.2. *The free product $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ is freely telescopic.*

The proofs of these two results rely on a LEGO-like construction using pieces of non-degenerate (T, X) -graphs, that we produce below. Once this is done, and the necessary properties of the construction hold, the proofs will follow easily, cf. Section 3.2.6.

3.2.1. *Vertex splitting and gluing.* Let us consider a Schreier graph $\mathcal{S} = \text{Sch}_{T,X}(H)$ of a subgroup $H \leq T$, where T has generating set X . If $Y \subsetneq X$, a vertex v of \mathcal{S} is *split* along Y if v is replaced by two vertices v_Y, v_{X-Y} , where v_Y keeps the Y -coloured edges of v , and v_{X-Y} keeps its $(X - Y)$ -coloured edges, as shown in Figure 1. We shall call v_Y a *dangling* Y -coloured vertex. Observe that splitting vertices breaks the X -regularity of the graph. If u_Y and v_{X-Y} are, respectively, Y - and $(X - Y)$ -coloured dangling vertices we say that u_Y and v_{X-Y} are *complementary*, and were we to identify them, we would gain regularity back at the newly created vertex. With this idea in mind, dangling vertices are seen as “connection points” for our graphs: an Y -coloured dangling vertex can only be connected to an $(X - Y)$ -dangling vertex, and once all dangling vertices of a graph are connected, the resulting graph is X -regular. Let us call the identification of complementary vertices *gluing*.

Finally, if \mathcal{S} is a Schreier graph of a free subgroup of T , it has no degenerate cycles. Since the operations of splitting vertices and gluing complementary ones do not change the lengths of cycles of any given colour, as soon as regularity is gained back by gluing all dangling vertices of some split graph, one gets the Schreier graph of a *free* subgroup once again.

The following is now essentially obvious from the above considerations.

PROPOSITION 3.3. *Let us choose $Y \subsetneq X$ and a finite number of Schreier graphs $\mathcal{S}_i = \text{Sch}_{T,X}(H_i)$ of free, finite-index subgroups H_i of T . Consider their disjoint union $\sqcup \mathcal{S}_i$, in which we split a certain number of vertices along Y , and glue them, in complementary pairs, so that the resulting graph is connected. Then we obtain a Schreier graph for a free, finite-index subgroup of T .*




FIGURE 1. Before and after splitting an X -regular graph, with $X = \{\text{red, green, blue}\}$, at a vertex v along $Y = \{r, b\}$

3.2.2. *Sketch of the construction.* Below we explain the main idea of the construction. Let T be either $\mathbb{Z}_p * \mathbb{Z}_q$ or $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$, with X its natural set of generators, and Γ be a finite group generated by a set S . Let us choose two X -coloured graphs, say \mathcal{L}_e and \mathcal{L}_v obtained by splitting, respectively, one and two vertices in the Schreier graph of a finite-index free subgroup of T , i.e. a non-degenerate (T, X) -graph. Call \mathcal{L}_e an *edge-link* and \mathcal{L}_v a *vertex-link*. We shall also choose an easily identifiable and *unique*, as we shall see in the sequel, vertex of \mathcal{L}_e to be its *root*, denoted $r(\mathcal{L}_e)$.

We shall connect edge-links by gluing complementary vertices, so as to connect them into chains, and call the result *edge-graphs*. Similarly, we shall connect vertex-links but in a way to produce cycles of them, and call the result *vertex-graphs*. Finally, in the Cayley graph $\mathcal{C} := \text{Cay}(\Gamma, S)$ of Γ , we shall replace the edges by edge-graphs and the vertices by vertex-graphs, following the procedure of Section 4.3. Our construction will ensure that

- (1) the automorphism group $\text{Aut } \mathcal{C}^*$ of the resulting graph \mathcal{C}^* is the same as the automorphism group of \mathcal{C} , which is exactly Γ (by Proposition 4.1); and
- (2) the graph \mathcal{C}^* is actually the Schreier graph of a finite-index free subgroup of T (by Proposition 3.3).

From now on, if v is a vertex of a *folded* X -coloured graph, we write $v \cdot x_1 \dots x_n$ for the terminus of the unique path labelled x_1, \dots, x_n and starting at v , if it exists. Then, an equality of the form $v \cdot x_1 \dots x_n = u \cdot y_1 \dots y_n$ holds if and only if both paths in question exist and their termini are equal.

3.2.3. *Constructing the links.* Our construction differs slightly for $T = \mathbb{Z}_p * \mathbb{Z}_q$ and $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$. In the former case, one uses a relatively generic construction, while the latter is mostly ad-hoc.

CASE $T = \mathbb{Z}_p * \mathbb{Z}_q$. Let us fix any $p \geq 3$ and $q \geq 2$, let the corresponding generators of each free factor of T be red and cyan and consider q copies of a p -cycle with edge labels r and vertices $v_{0,i}, \dots, v_{p-1,i}$ for each i -th copy, where $0 \leq i \leq q-1$. Next, add two “special” q -cycles labelled c :

$$v_{0,0} \rightarrow v_{1,0} \rightarrow v_{1,1} \rightarrow v_{1,2} \rightarrow \dots \rightarrow v_{1,q-2} \rightarrow v_{0,0},$$

and

$$v_{0,1} \rightarrow v_{0,2} \rightarrow \dots \rightarrow v_{0,q-1} \rightarrow v_{1,q-1} \rightarrow v_{0,1}.$$

Then, for each $j \geq 2$, draw an extra q -cycle labelled c :

$$v_{j,0} \rightarrow v_{j,1} \rightarrow \dots \rightarrow v_{j,q-1} \rightarrow v_{j,0}.$$

In the case $q = 2$, the extra q -cycles have the form $v_{0,0} \rightarrow v_{1,0} \rightarrow v_{0,0}$ and $v_{0,q-1} \rightarrow v_{1,q-1} \rightarrow \dots \rightarrow v_{0,q-1}$. Let $\mathcal{G}_{p,q}$ denote the resulting non-degenerate $(\mathbb{Z}_p * \mathbb{Z}_q, \{r, c\})$ -graph.

Now, split the vertex $v_{0,0}$ in order to produce an *edge-link* denoted $\mathcal{L}_{p,q}^e$ and, subsequently, split the vertex $v_{0,q-1}$ to get a *vertex-link* $\mathcal{L}_{p,q}^v$. The vertices obtained by splitting $v_{0,0}$ will be denoted v_+, v_- , and those obtained by splitting $v_{0,q-1}$ will be called u_+, u_- . In the sequel, we swap the assignment of v_+ and v_- vertices for edge- and vertex-links, as shown in Figures 2–4. This allows us to keep a consistent and clear notation for all associated objects.

In the edge-link $\mathcal{L}_{p,q}^e$, the vertex $v_{0,q-1}$ is unique in the following sense: this is the only vertex $v \in V\mathcal{L}_{p,q}^e$ that satisfies $v \cdot r = v \cdot c$. Observe that the assumption $p \geq 3$ is important here: if $p = q = 2$, then once $v \cdot r = v \cdot c$, the vertex $w = v \cdot r$ also satisfies $w \cdot r = w \cdot c$. Let then $r(\mathcal{L}_{p,q}^e) := v_{0,q-1}$ be called the *root*⁽¹⁾ of $\mathcal{L}_{p,q}^e$.

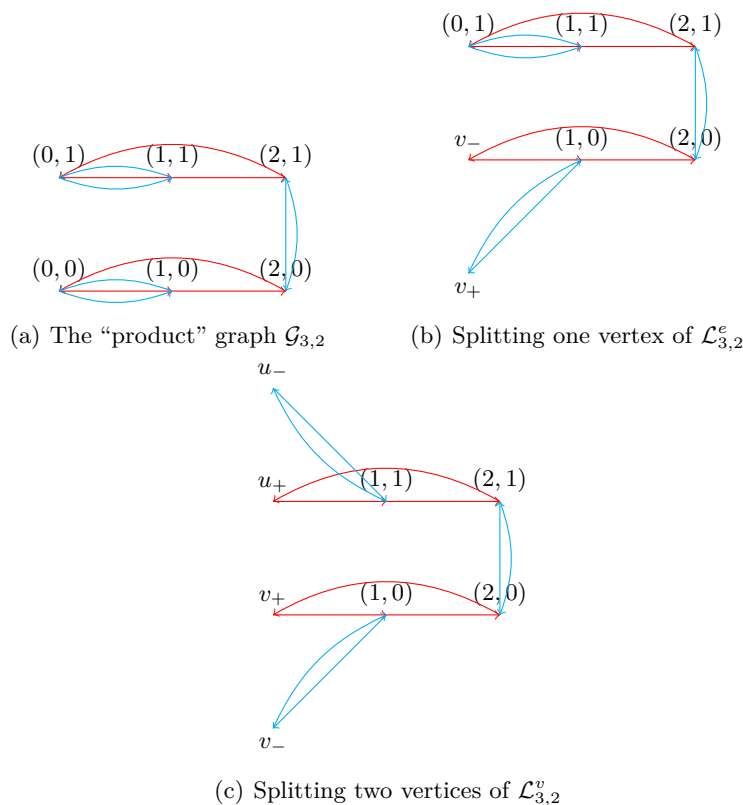


FIGURE 2. The “link” graphs for $T = \mathbb{Z}_3 * \mathbb{Z}_2$: (a) the original graph; (b) the result of splitting $(0, 0)$; (c) the result of splitting $(0, 0)$ and $(0, 1)$. The root of the edge link is $(0, 1)$.

Some examples of “product graphs” and their splitting at one and two vertices that generate edge- and vertex-link graphs are depicted in Figures 2–3.

CASE $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$. Let the generators of the group T be red, green and blue. Let $\mathcal{G}_{2,2,2}$ denote the graph (a) in Figure 4. In $\mathcal{G}_{2,2,2}$, first split the vertex v_7 to get an edge-link that we call $\mathcal{L}_{2,2,2}^e$, and then split v_0 to get a vertex-link called $\mathcal{L}_{2,2,2}^v$.

⁽¹⁾In the case of an edge- or vertex-link the notion of a root is practically opposed to the notion of a root in a graph (map, hypermap, etc.). Indeed, the former is intrinsic to the respective combinatorial structure, while the latter is a matter of choice and can be assigned arbitrarily

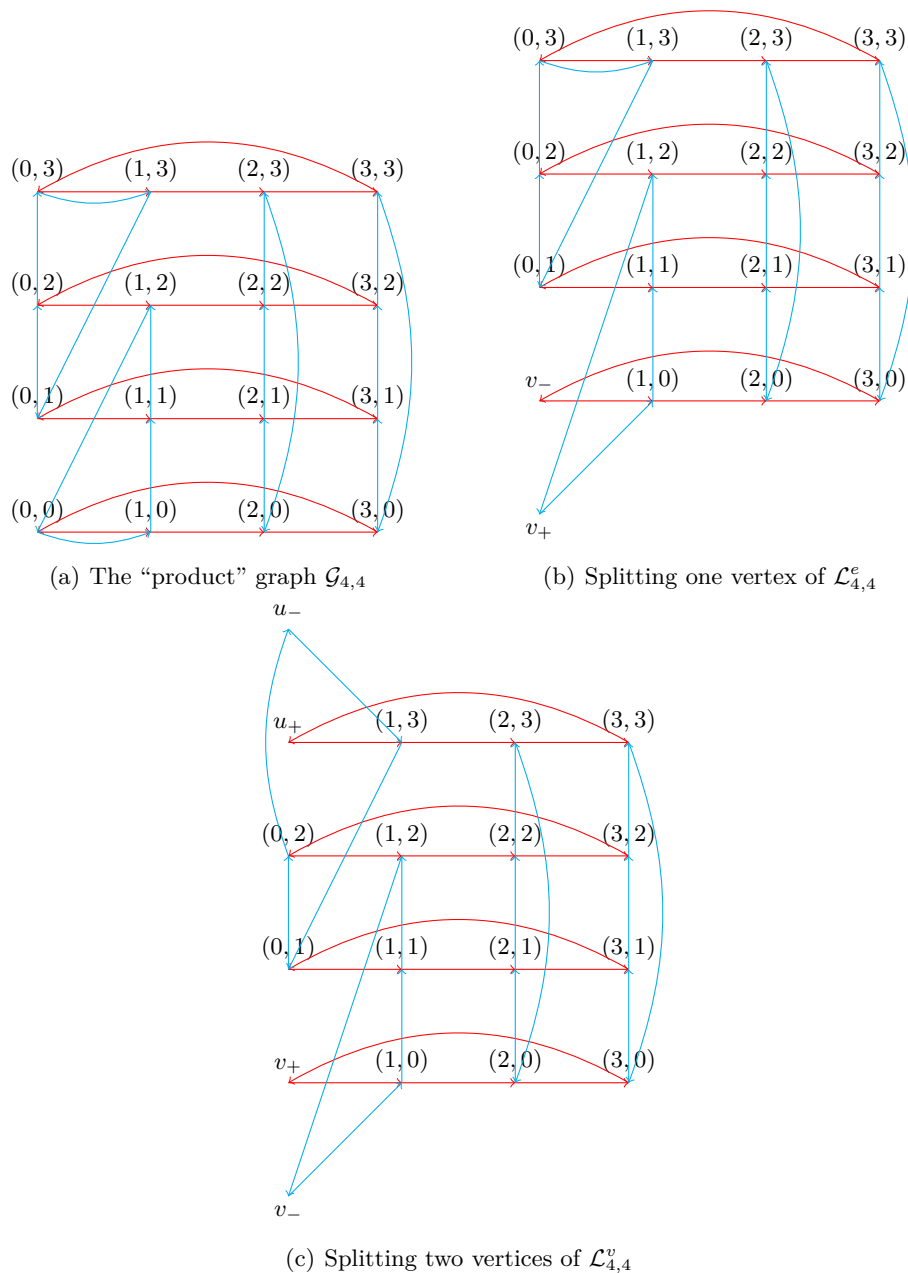


FIGURE 3. The “link” graphs for $T = \mathbb{Z}_4 * \mathbb{Z}_4$: (a) the original graph; (b) the result of splitting $(0,0)$; (c) the result of splitting $(0,0)$ and $(0,3)$. The root of the edge link is $(0,3)$.

The vertices obtained by splitting the vertex v_7 will be denoted v_+, v_- , and those obtained by splitting the vertex v_0 will be called u_+, u_- .

Observe that the vertices v_0 and v_1 are unique in $\mathcal{L}_{2,2,2}^e$ in the following sense: they are the only vertices $v \in V\mathcal{L}_{2,2,2}^e$ satisfying $v \cdot b = v \cdot g$. Furthermore, one can distinguish v_0 from v_1 as follows: while $v_1 \cdot rgbgr = v_1$, it is not the case for v_0 . In other words,

the path labelled $rgbgrgr$ and starting at v_1 is a loop, but the one identically labelled and starting at v_0 is not. Let then $r(\mathcal{L}_{2,2,2}^e) := v_1$ be the *root* of $\mathcal{L}_{2,2,2}^e$.




FIGURE 4. The “link” graphs for $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$: (a) the original graph; (b) the result of splitting 7; (c) the result of splitting 0 and 7. The root of the edge link is encircled.

Let us write \mathcal{L}^v for any of $\mathcal{L}_{p,q}^v$ and $\mathcal{L}_{2,2,2}^v$, and \mathcal{L}^e for any of $\mathcal{L}_{p,q}^e$ and $\mathcal{L}_{2,2,2}^e$. Since the following constructions do not depend on the exact nature of those graphs, but rather on their abstract properties, this ambiguity is harmless.

ON THE BEHAVIOUR OF ROOTS. Let us define the following conditions

$$P_{p,q}(v) := "v \in V\mathcal{L}_{p,q}^e \text{ satisfies } v \cdot r = v \cdot c",$$

$$P_{2,2,2}(v) := "v \in V\mathcal{L}_{2,2,2}^e \text{ satisfies } v \cdot rgbgrgr = v \text{ and } v \cdot b = v \cdot g",$$

and observe that the following statements hold for $P = P_{p,q}$ or $P_{2,2,2}$, whichever is appropriate.

- G.1 No vertex of a vertex-link \mathcal{L}^v satisfies P , and exactly one vertex $r(\mathcal{L}^e)$ of an edge-link \mathcal{L}^e does.
- G.2 Gluing edge- and vertex-links together by identifying complementary dangling vertices does not create new vertices satisfying P , as long as the gluing is done on vertices with disjoint neighbourhoods (where the neighbourhood of a vertex v is the set of adjacent vertices.)
- G.3 If $\iota : \mathcal{A} \hookrightarrow \mathcal{B}$ is an embedding of folded graphs (in our case \mathcal{A} and \mathcal{B} will be obtained by gluing vertex- and edge-links), then the image of a vertex satisfying P also satisfies P .

Of the above, G.1 holds by construction of the vertex- and edge-links and G.3 is evident from the fact that "following the labels" and "passing to the image under an embedding" are commuting operations. Only G.2 is not as direct. Let u_Y and v_{X-Y} be complementary dangling vertices with disjoint neighbourhoods. Then, letting x be a label in $X - Y$, y be a label in Y , and w be the result of gluing u_Y to v_{X-Y} , we get $w \cdot x = v_{X-Y} \cdot x$ and $w \cdot y = u_Y \cdot y$, which are distinct by the hypothesis. Since both conditions $P_{p,q}$ and $P_{2,2,2}$ involve equalities of the form $v \cdot x = v \cdot y$, it follows that no glued complementary vertices can satisfy them.

3.2.4. *Constructing vertex- and edge-graphs.* First of all, let us introduce some necessary notation, which will also be used in Section 4.3, later on.

A *vertex-graph* is a graph \mathcal{V} , along with an injection $\chi : S \times \{+, -\} \hookrightarrow V\mathcal{V}$, an example of which is depicted in Figure 5. Let the *boundary* of a vertex-graph be $\partial\mathcal{V} = \text{im } \chi$.

An *edge-graph* for a label s is a graph \mathcal{E}_s , along with two distinguished vertices $h^+(\mathcal{E}_s)$ and $h^-(\mathcal{E}_s)$, as shown in Figure 6. We shall provide the general definitions of the graphs \mathcal{V} and \mathcal{E}_s below, which will be case-specific for different choices of T .

In the Cayley graph $\mathcal{C} := \text{Cay}(\Gamma, S)$ of Γ , each vertex v will be replaced by an isomorphic copy \mathcal{V}_v of \mathcal{V} , and each edge e labelled s will be replaced by an isomorphic copy \mathcal{E}_e of \mathcal{E}_s . We shall also make the following identifications:

$$\mathcal{E}_e \ni h^+(\mathcal{E}_s) \sim \chi(s, +) \in \mathcal{V}_u, \quad \mathcal{E}_e \ni h^-(\mathcal{E}_s) \sim \chi(s, -) \in \mathcal{V}_v,$$

whenever e has label s , origin u and terminus v , where $h^\pm(\mathcal{E}_s) \in \mathcal{E}_e$ means a copy of the vertex $h^\pm(\mathcal{E}_s)$ inside \mathcal{E}_e , and similarly for other instances of vertex- and edge-graphs.

Let \mathcal{C}^* denote the resulting graph, and let $\iota_e : \mathcal{E}_s \hookrightarrow \mathcal{C}^*$ and $\iota_v : \mathcal{V} \hookrightarrow \mathcal{C}^*$ be the embeddings corresponding to an edge $e \in VC$ (with $\mu(e) = s$) and a vertex $v \in VC$, respectively. The image $\iota_e(\mathcal{E}_s)$ will be denoted by \mathcal{E}_e and the image $\iota_v(\mathcal{V})$ will be called \mathcal{V}_v . Let $h^\pm(e)$ be the image of $h^\pm(\mathcal{E}_s)$ under ι_e .

Finally, if an ordering on the labels $s_1 < \dots < s_n$ is given, then let

$$\mathcal{C}_{(i)}^* := \left(\bigcup_{v \in VC} \mathcal{V}_v \right) \cup \left(\bigcup_{\substack{j \leq i \\ e \in EC \\ \text{s.t. } \mu(e) = s_j}} \mathcal{E}_e \right),$$

which means that $\mathcal{C}_{(i)}^*$ is the subgraph of \mathcal{C}^* consisting of the vertex-graphs and *only* the edge-graphs corresponding to the edges with labels s_j for $j \leq i$.




FIGURE 5. The vertex-graph \mathcal{V} for $|S| = 6$. The squares represent the connected vertex-links \mathcal{L}^v , the gray dots represent the vertices glued along the way, while the cyan and red dots represent the remaining dangling vertices, which are exactly the images of χ .

From now on, assume that the generating set $S = \{s_1, \dots, s_n\}$ is ordered: $s_1 < s_2 < \dots < s_n$. First, we consider the vertex- and edge-graphs that we use in the case $T = \mathbb{Z}_p * \mathbb{Z}_q$, with $p \geq q$, $p \geq 3$, $q \geq 2$.

Vertex-graphs. Let us take $|S|$ disjoint copies of the vertex-link \mathcal{L}^v from Section 3.2.3, and call them $\mathcal{L}_v^{(0)}, \dots, \mathcal{L}_v^{(|S|-1)}$, respectively. Observe that each graph $\mathcal{L}_v^{(i)}$ has four dangling vertices $v_+^{(i)}, v_-^{(i)}, u_+^{(i)}, u_-^{(i)}$. In order to create the vertex graph \mathcal{V} , identify each $v_+^{(i)}$ to $u_-^{((i+1) \bmod |S|)}$, for $0 \leq i \leq |S| - 1$. Let us also define $\chi : S \times \{+, -\} \hookrightarrow V\mathcal{V}$ as $\chi((s_i, +)) = u_+^{(i)}$ and $\chi((s_i, -)) = v_-^{(i)}$. For a sketch of the resulting graph, see Figure 5.

Note that in the case $|S| = 1$, only one vertex-link is used, and the fact that $v_+^{(0)}$ is glued to $u_-^{(0)}$, and not to $v_-^{(0)}$, ensures that the hypothesis G.2 holds, so that no extra root appears.




FIGURE 6. An edge-graph of length 4. The squares represent the connected edge-links \mathcal{L}^e , and the cyan and red dots represent the remaining dangling vertices, which are exactly $h^+(e)$ and $h^-(e)$.

Edge-graphs. For each $s_i \in S$, take i copies of the edge-link \mathcal{L}^e from Section 3.2.3, and call them $\mathcal{L}_e^{(1)}, \dots, \mathcal{L}_e^{(i)}$, respectively. Observe that each graph $\mathcal{L}_e^{(j)}$ has two dangling vertices $v_+^{(j)}, v_-^{(j)}$. Now, glue each $v_-^{(j)}$ to $v_+^{(j+1)}$, for $1 \leq j \leq i - 1$. Let \mathcal{E}_i denote the

resulting edge-graph, a sketch of which is depicted in Figure 6. The vertex $v_+^{(1)}$ will be referred to as $h^+(e)$ and $v_-^{(i)}$ as $h^-(e)$.

Note that the exact same construction works for $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$, when using the corresponding objects, i.e. when $\mathcal{L}_{p,q}^v$ is replaced by $\mathcal{L}_{2,2,2}^v$ and $\mathcal{L}_{p,q}^e$ is replaced by $\mathcal{L}_{2,2,2}^e$.

In both cases, letting P be the appropriate condition, i.e. either $P_{p,q}$ or $P_{2,2,2}$, each edge graph \mathcal{E}_i contains exactly i vertices satisfying P : they are exactly the copies of the root vertex $r(\mathcal{L}^e)$ satisfying P in the edge-link \mathcal{L}^e . Let $r_1(s_i), \dots, r_i(s_i)$, where $r_j(s_i) \in \mathcal{L}^{(j)}$, denote those roots. Observe that the vertex-graphs do not contain any vertex satisfying P and that $\text{Aut}(\mathcal{E}_i) = \{\text{id}\}$, where the former follows by construction and the latter follows from foldedness.

3.2.5. *Substitution.* Let $\mathcal{C} = \text{Cay}(\Gamma, S)$ be the Cayley graph of a finite group Γ with respect to a generating set $S = \{s_1, s_2, \dots, s_n\}$. Our goal is to prove that $\text{Aut}(\mathcal{C}^*) \cong \text{Aut}(\mathcal{C})$ which we can do by using the argument of Proposition 4.1 (cf. Section 4 for its proof). Thus, we only have to show that the following conditions hold:

- S.1 $\{\phi \in \text{Aut } \mathcal{V} : \phi|_{\partial\mathcal{V}} = \text{id}_{\partial\mathcal{V}}\} = \{\text{id}_{\mathcal{V}}\}$;
- S.2 one can order the labels s_1, \dots, s_n in such a way that for any $i \in \{0, \dots, n\}$, $\mathcal{C}_{(i)}^*$ contains no subgraph isomorphic to \mathcal{E}_{s_i} except for the subgraphs \mathcal{E}_e , for $e \in EC, \mu(e) = s_i$.

In the edge-link \mathcal{L}^e , choose a shortest path from v_+ to $r(\mathcal{L}^e)$ and from $r(\mathcal{L}^e)$ to v_- , and let w_1 and w_2 be their respective labels.

- S.1 Recall that \mathcal{V} is a folded graph. Let $\phi \in \text{Aut } \mathcal{V}$ be an automorphism with $\phi|_{\partial\mathcal{V}} = \text{id}_{\partial\mathcal{V}}$, and let $v \in \partial\mathcal{V}$ be a vertex in the boundary of \mathcal{V} . Then, since ϕ and $\text{id}_{\mathcal{V}}$ agree at v , and \mathcal{V} is folded, we have $\phi = \text{id}_{\mathcal{V}}$. Thus S.1 holds.
- S.2 Fix $i \in \{0, \dots, n\}$ and consider $\mathcal{C}_{(i)}^*$. Our goal is to show that $\mathcal{C}_{(i)}^*$ contains no other copies of \mathcal{E}_i than those of the form \mathcal{E}_e , for e an edge with label $\mu(e) = s_i$. Let $r := r_1(s_i)$ be the root of the first edge-link in \mathcal{E}_i . Suppose that there exists an embedding $\iota : \mathcal{E}_i \hookrightarrow \mathcal{C}_{(i)}^*$, and let us show that $\iota(\mathcal{E}_i)$ is equal to $\iota_e(\mathcal{E}_i)$ for some e with label $\mu(e) = s_i$.

We know that $\iota(r)$ lies in an edge-graph because $\iota(r)$ satisfies property P by G.3 and no vertex of a vertex-graph satisfies property P . Let $e \in \mathcal{E}_e$ be the edge such that $\iota(r)$ lies in \mathcal{E}_e , and $s_k := \mu(e)$. It is essential to notice that in \mathcal{E}_i , for any $1 \leq j \leq i - 1$, we have $r_j(s_i)w_2w_1 = r_{j+1}(s_i)$. Similarly, in \mathcal{E}_e , for any $1 \leq j \leq k - 1$, it holds that $r_j(e)w_2w_1 = r_{j+1}(e)$. A sketch of how the roots and boundary vertices are arranged within an edge graph is given in Figure 7.

We shall show that e has label s_i and $\iota(r) = r_1(e)$. This is enough to conclude that $\iota(\mathcal{E}_i)$ equals $\iota_e(\mathcal{E}_i)$, since, by assumption, the embeddings ι and ι_e agree at a vertex $r \in \mathcal{E}_i$, and thus must coincide by foldedness.

To this end, assume that e has label s_k with $k \leq i$, and $\iota(r) = r_j(e)$, necessarily with $1 \leq j \leq k$. Then, we prove that $k - j \geq i - 1$, which is equivalent, by using the constraints $k \leq i$ and $1 \leq j \leq k$, to $k = i$ and $j = 1$.

Let us assume, contrary to the above, that $k - j < i - 1$. Following the path labelled $(w_2w_1)^{k-j+1}$ and starting at $r = r_1(s_i)$ in \mathcal{E}_i , we reach the vertex $r_{k-j+2}(s_i)$ in \mathcal{E}_i , which is a root. We claim that $r_j(e)(w_2w_1)^{k-j+1}$ is *not* a root, which thus contradicts ι being an embedding, since $r_j(e)(w_2w_1)^{k-j+1} = \iota(r_{k-j+2}(s_i))$ is the image of a root.

First, let us consider the vertex $r_j(e)(w_2w_1)^{k-j} = r_k(e)$. This is the last root of \mathcal{E}_e , so that $r_k(e)w_2$ is in a vertex-graph: in fact, $r_k(e)w_2 = h^-(e)$.

Now, if $\hat{r} := r_j(e)(w_2w_1)^{k-j+1} = r_k(e)(w_2w_1)$ were a root somewhere in $C_{(i)}^*$, then $r_k(e)w_2 = \hat{r}w_1^{-1}$ would either be in the interior of an edge-graph, or of the form $h^+(e')$ for some edge e' . Indeed, if \hat{r} were the first root $r_1(e')$ of some $\mathcal{E}_{e'}$, then $r_1(e')w_1^{-1} = h^+(e')$ would hold. If it were a subsequent root, then $\hat{r}w_1^{-1}$ would lie in the interior of $\mathcal{E}_{e'}$. In neither case can it be equal to $r_k(e)w_2 = h^-(e)$, and the desired contradiction is reached.




FIGURE 7. Edge-links assembled into an edge-graph, their respective roots and boundary vertices.

3.2.6. Free telescopicity of $\mathbb{Z}_p * \mathbb{Z}_q$ and $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$.

Proof of Propositions 3.1 and 3.2. Given T either $\mathbb{Z}_p * \mathbb{Z}_q$ ($p \geq q, p \geq 3, q \geq 2$) or $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ with its natural generating set X (one generator for each cyclic factor), and Γ a finite group with generators S , let us substitute the edges and vertices of $\mathcal{C} := \text{Cay}(\Gamma, S)$ by the vertex- and edge-graphs as described in Sections 3.2.4-3.2.5.

By Proposition 3.3, the resulting graph \mathcal{C}^* is a non-degenerate (T, X) -graph with automorphism group

$$\text{Aut } \mathcal{C}^* \cong \text{Aut } \mathcal{C} \cong \Gamma,$$

as verified in Section 3.2.5.

Let H be a finite-index free subgroup of T with $\text{Sch}_{T,X}(H) \cong \mathcal{C}^*$. Then, by Lemma 2.8,

$$N_T(H)/H \cong \Gamma.$$

Since the procedure above can be performed for arbitrary Γ , both $\mathbb{Z}_p * \mathbb{Z}_q$ and $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ are freely telescopic. \square

3.3. ADDING FACTORS. The following lemma allows us to pull back telescopicity by surjections. The only property not guaranteed in this case is freeness.

LEMMA 3.4. *Let $f : A \rightarrow B$ be an epimorphism. Then, if B is telescopic, A is also telescopic, although not necessarily freely.*

Proof. Let Γ be a finite group. Since B is telescopic, there exists a finite index subgroup $H \leq B$ such that $N_B(H)/H \cong \Gamma$. Consider $\tilde{H} := f^{-1}(H) \leq A$, the preimage of H in A . Clearly, \tilde{H} is a finite index subgroup of A since it is a preimage of a finite index subgroup under an epimorphism.

We have that $N_A(\tilde{H}) = f^{-1}(N_B(H))$. Indeed, let x be an element of $f^{-1}(N_B(H))$. By definition, this means that $f(x) \in N_B(H)$, and thus $f(x)H = Hf(x)$. By applying f^{-1} to both sides, we obtain $x\tilde{H} = \tilde{H}x$. The latter yields that x is an element of $N_A(\tilde{H})$. The reverse inclusion is analogous.

Thus, $N_A(\tilde{H})/\tilde{H}$ equals $f^{-1}(N_B(H))/f^{-1}(H)$. It remains to check that the quotient group $f^{-1}(N_B(H))/f^{-1}(H)$ is isomorphic to $N_B(H)/H \cong \Gamma$. Let us consider the map $\tilde{f} : f^{-1}(N_B(H)) \rightarrow N_B(H)/H$ defined by $\tilde{f}(x) = f(x) \cdot H$. Since \tilde{f} is a surjective homomorphism with kernel $\ker(\tilde{f}) = f^{-1}(H)$, the desired result follows from the first isomorphism theorem. \square

The following lemmas are key to making an inductive step and proving the main result stated as Theorem 3.7.

LEMMA 3.5. *Assume that a product of the form $T * \mathbb{Z}_m$, $m \geq 2$, is freely telescopic. Then $T * \mathbb{Z}_a * \mathbb{Z}_b$ is also freely telescopic, for all $a, b \geq 2$ such that $\text{lcm}(a, b) = m$.*

Proof. Let $c := m/a$, $d := m/b$ and consider the following morphisms:

$$\begin{aligned} \iota_a : \mathbb{Z}_a &\rightarrow \mathbb{Z}_m, \\ [k]_a &\mapsto [ck]_m \end{aligned}$$

and

$$\begin{aligned} \iota_b : \mathbb{Z}_b &\rightarrow \mathbb{Z}_m \\ [k]_b &\mapsto [dk]_m. \end{aligned}$$

Since $\iota_a([1]_a) = [d]_m$ and $\iota_b([1]_b) = [c]_m$, and c and d are coprime, any element of \mathbb{Z}_m can be written as a sum of elements in the images of ι_a and ι_b by Bézout’s theorem. Let us consider the surjective morphism $\phi : T * \mathbb{Z}_a * \mathbb{Z}_b \rightarrow T * \mathbb{Z}_m$ induced by the maps:

$$\iota_a : \mathbb{Z}_a \rightarrow \mathbb{Z}_m, \quad \iota_b : \mathbb{Z}_b \rightarrow \mathbb{Z}_m, \quad \text{id} : T \rightarrow T.$$

Indeed, ι_a, ι_b and id all extend to $T * \mathbb{Z}_m$, while the universal property of the free product yields ϕ .

Since the generator $[1]_m$ of \mathbb{Z}_m and all of T are in the image of ϕ , the latter is surjective. Also, the restrictions of ϕ to each of the subgroups T , \mathbb{Z}_a , and \mathbb{Z}_b are injective since they correspond exactly to the post-compositions of id , ι_a , and ι_b , respectively, with their inclusions in $T * \mathbb{Z}_m$.

Now, fix a finite group Γ to be realised as a quotient “normaliser/subgroup” of $T * \mathbb{Z}_a * \mathbb{Z}_b$. We know, by the hypothesis, that there exists a free subgroup $H \leq T * \mathbb{Z}_m$, of finite index, such that $N(H)/H \cong \Gamma$. Since ϕ is surjective, $\phi^{-1}(H)$ is the desired subgroup by Lemma 3.4, once we verify that it is free. To this end, assume that $\phi^{-1}(H)$ is not. Then, by Kurosh’s theorem, $\phi^{-1}(H)$ contains a conjugate of a non-free subgroup of T , \mathbb{Z}_a or \mathbb{Z}_b . In the first case, there exists $R \leq T$ non-free such that $\phi^{-1}(H)$ contains wRw^{-1} , for some w , as a free factor. Then

$$H \geq \phi(wRw^{-1}) = \phi(w)\phi(R)\phi(w^{-1}) = \phi(w)R\phi(w^{-1}),$$

where the third equality stems from the fact that ϕ restricts to “the identity + inclusion” on T , by definition. This implies that H contains a non-free subgroup, which is a contradiction.

Similarly, assume that $\phi^{-1}(H)$ contains a free factor of the form wCw^{-1} , with C a non-free subgroup of \mathbb{Z}_a . Then we have

$$H \geq \phi(wCw^{-1}) = \phi(w)\phi(C)\phi(w^{-1}) \cong \phi(w)\iota_a(C)\phi(w^{-1}),$$

where, once again, the third equality stems from the fact that ϕ , when restricted to \mathbb{Z}_a , is just ι_a , plus the inclusion of \mathbb{Z}_m in $T * \mathbb{Z}_m$. Since ι_a is injective, and C is not free, then $\iota_a(C)$ is not free either, so that H contains a non-free subgroup, which is again a contradiction. An analogous reasoning applies if we assume that C is a non-free subgroup of \mathbb{Z}_b , and the lemma follows. \square

LEMMA 3.6. *Assume that a product of the form $T * \mathbb{Z}_m$, $m \geq 2$, is freely telescopic. Then $T * \mathbb{Z}$ is freely telescopic.*

Proof. As above, consider the morphisms $\text{id} : T \rightarrow T$ and $q : \mathbb{Z} \rightarrow \mathbb{Z}_m$, together with the induced surjective “composite” morphism

$$\phi : T * \mathbb{Z} \rightarrow T * \mathbb{Z}_m.$$

In this case, the injectivity of ϕ on the \mathbb{Z} -factor is not required: we already know that any subgroup of \mathbb{Z} is free. \square

THEOREM 3.7. *Any free product of at least two non-trivial cyclic groups is freely telescopic, except for $\mathbb{Z}_2 * \mathbb{Z}_2$.*

Proof. If $T = \mathbb{Z}_2 * \mathbb{Z}_2$, the infinite dihedral group, then the only Schreier graphs associated with finite-index free subgroups of T are cycles, and T cannot be freely telescopic (obviously, it cannot be telescopic at all). If $T = *_{i=1}^n \mathbb{Z}_{p_i}$ is a finite free product of non-trivial cyclic groups with $n \geq 2$ and $p_1 \geq 3$ or $n \geq 3$, the proof proceeds by induction on n .

If $n = 2$, then Proposition 3.1 yields that $T = \mathbb{Z}_{p_1} * \mathbb{Z}_{p_2}$ is freely telescopic, assuming that $p_1 \geq 3$ and $p_2 \geq 2$, without loss of generality. If $n = 3$, either each of p_1, p_2, p_3 equals 2, in which case Proposition 3.2 yields free telescopicity of $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$, or, without loss of generality, we have $p_3 \neq 2$. In the latter case, $\mathbb{Z}_{p_1} * \mathbb{Z}_{\text{lcm}(p_2, p_3)}$ is already freely telescopic, and Lemma 3.5 shows that $\mathbb{Z}_{p_1} * \mathbb{Z}_{p_2} * \mathbb{Z}_{p_3}$ is so, as well.

The inductive step towards $n \geq 4$ is made by using Lemma 3.5, with p_i 's being finite. Setting $p_i = \infty$, for any i , also yields freely telescopic groups, by Lemma 3.6. This concludes the proof of the theorem. \square

4. GRAPH SUBSTITUTION

In this section we shall always consider directed labelled graphs. Our goal is to define a reasonable condition that allows replacing vertices and edges of a Cayley graph by other “chunks of graphs” in a way that preserves the automorphism group. All edges having a fixed given label will be replaced by the same “chunk” of a suitable graph, for each label, and a similar procedure takes place for all vertices. In what follows, let Γ be a finite group generated by a finite set S , and let $\mathcal{C} := \text{Cay}(\Gamma, S)$ be its Cayley graph.

4.1. VERTEX GRAPHS. If v is a vertex in the Cayley graph $\mathcal{C} = \text{Cay}(\Gamma, S)$, let $\text{Adj } v$ be the set of tuples consisting of edges adjacent to v and their orientations relative to v , i.e. $\text{Adj } v = (\text{St}_+ v) \times \{+\} \sqcup (\text{St}_- v) \times \{-\}$. Let $\Sigma := S \times \{+, -\}$, then each set $\text{Adj } v$ is naturally in bijection with Σ : one can easily identify the corresponding pairs of edge labels and their orientations relative to v . Let $\tau_v : \text{Adj } v \rightarrow \Sigma$ be this bijection, to which we shall refer as *the signature* of v .

If \mathcal{V} is a connected graph and $\chi : \Sigma \leftrightarrow V\mathcal{V}$ is an injective map, then we call the pair (\mathcal{V}, χ) a *signed graph*. Looking ahead, the map χ will tell us how one should connect the edges adjacent to a vertex $v \in V\mathcal{C}$ to the vertices of \mathcal{V} when replacing v in \mathcal{C} by \mathcal{V} . The injectivity of χ also comes useful later.

Let $\partial\mathcal{V} := \text{im } \chi$, which we shall call the boundary of \mathcal{V} . Fix a signed graph (\mathcal{V}, χ) , and let

$$\text{Aut}(\mathcal{V}, \chi) := \{\phi \in \text{Aut } \mathcal{V} : \phi|_{\partial\mathcal{V}} = \text{id}_{\partial\mathcal{V}}\}$$

be the group of automorphisms of \mathcal{V} which restrict to the identity map on the boundary $\partial\mathcal{V}$. We shall call $\text{Aut}(\mathcal{V}, \chi)$ the group of *signed automorphisms* of (\mathcal{V}, χ) .

Now let us consider the following condition, that will play an important role in Section 4.4.

S.1 $\text{Aut}(\mathcal{V}, \chi) = \{\text{id}\}$.

In other words, S.1 states that any non-trivial automorphism of \mathcal{V} has to move some vertex of its boundary $\partial\mathcal{V}$.

For a given vertex $v \in V\mathcal{C}$, a signed graph \mathcal{V} can be inserted in place of v by connecting each edge $e \in \text{Adj } v$ to the vertex $\chi(\tau_v(e)) \in \partial\mathcal{V}$. Broadly speaking, condition S.1 forbids any automorphism local to \mathcal{V} to appear when v is replaced by \mathcal{V} .

A graph \mathcal{V} described above will be called the *vertex-graph* associated with v .

4.2. EDGE GRAPHS. The case of edge substitution is simpler: if s is a label, an *edge-graph* for s is a connected graph \mathcal{E}_s with distinct distinguished vertices h_s^+ and h_s^- . We substitute an edge e labelled s by first removing e , and then identifying the origin of e with h_s^+ and its terminus with h_s^- .

4.3. VERTEX AND EDGE SUBSTITUTION. Below we describe the complete substitution procedure. Recall that we start by considering the Cayley graph \mathcal{C} of a group Γ with respect to generators S , as well as

- a connected signed vertex graph (\mathcal{V}, χ) ,
- a connected edge graph \mathcal{E}_s , for each edge label s , with two distinct distinguished vertices h_s^+ and h_s^- and trivial automorphism group $\text{Aut } \mathcal{E}_s = \{\text{id}\}$.

Let \mathcal{C}' be the result of replacing each vertex v of \mathcal{C} with an instance of \mathcal{V} as explained in Section 4.1. More precisely, if v is a vertex with signature $\tau_v : \text{Adj } v \rightarrow \Sigma$, we remove v , insert a copy of \mathcal{V} and connect each edge $e \in \text{Adj } v$ to the vertex $\chi(\tau_v(e))$ of \mathcal{V} .

Now, let \mathcal{C}^* be the result of replacing each “old” edge (i.e. an edge that is not in any of the vertex graphs) labelled s in \mathcal{C}' by a copy of \mathcal{E}_s , with h_s^+ identified with the origin of e , and h_s^- with its terminus, as described in Section 4.2.

Here and below, copies of \mathcal{V} will be always called “vertex-graphs”, and copies of \mathcal{E}_s will be “edge-graphs”.

For $v \in V\mathcal{C}$, write \mathcal{V}_v for the instance of the vertex-graph \mathcal{V} inserted in place of v , and let $\iota_v : \mathcal{V} \hookrightarrow \mathcal{C}^*$ be the corresponding graph embedding. Similarly, for $e \in E\mathcal{C}$ with label $\mu(e) = s$, we shall write \mathcal{E}_e for the instance of the edge-graph labelled s which is inserted in place of e , and $\iota_e : \mathcal{E}_s \hookrightarrow \mathcal{C}^*$ will be the corresponding graph embedding.

If e is an edge with origin u , resp. terminus v , and label s , then we identify $\iota_e(h_s^+)$ with $\iota_u(\chi(s, +))$, resp. $\iota_e(h_s^-)$ with $\iota_v(\chi(s, -))$. The vertex that we obtain after such identification is shared between \mathcal{E}_e and \mathcal{V}_u , resp. \mathcal{V}_v . We shall call these vertices $h^+(e)$ and $h^-(e)$, respectively, in order to distinguish them in \mathcal{C}^* .

Observe that our construction implies the following:

- The vertex-graphs are disjoint from each other, since h_s^+ and h_s^- are distinct in any \mathcal{E}_s , and so are the edge-graphs, by injectivity of χ .
- The vertices of $\partial\mathcal{V}_v$ are exactly those of the form $h^\pm(e)$ for an edge e adjacent to v in the initial graph \mathcal{C} .
- $h^+(e) \in \mathcal{V}_u$ if and only if u is the origin of e , and $h^-(e) \in \mathcal{V}_v$ if and only if v is the terminus of e .
- $\chi^{-1}\iota_u^{-1}(h^+(e)) = (\mu(e), +)$, resp. $\chi^{-1}\iota_v^{-1}(h^-(e)) = (\mu(e), -)$, if u is the origin, resp. the terminus, of e .

Finally, consider an ordering $s_1 < \dots < s_n$ on the edge labels. Let, for each $i = 0, \dots, n$, $\mathcal{C}_{(i)}^*$ be the subgraph of \mathcal{C}^* that consists of the vertex-graphs and only the edge-graphs corresponding to the labels s_j with $j \leq i$. In other words, this is a subgraph obtained by removing the “interiors” (i.e. everything but the vertices $h^\pm(e)$) of all the edge-graphs corresponding to the labels s_j for $j > i$.

4.4. MAIN STATEMENT. Below we formulate the main statement regarding our graph substitution procedure.

PROPOSITION 4.1. *With the notation above, if S.1 holds, as well as if*

- S.2 *one can order the labels s_1, \dots, s_n in such a way that for any $i \in \{0, \dots, n-1\}$, the graph $\mathcal{C}_{(i)}^*$ contains no subgraph isomorphic to \mathcal{E}_{s_i} except for the edge-subgraphs \mathcal{E}_e , for $e \in E\mathcal{C}$ with label $\mu(e) = s_i$,*

then

$$\text{Aut } \mathcal{C}^* \cong \text{Aut } \mathcal{C}.$$

Here, condition S.2 ensures that under any automorphism of \mathcal{C}^* the instances of \mathcal{V} , resp. the instances of any \mathcal{E}_s , must be sent to each other. As can be understood from the proof, any other condition ensuring this fact can be used instead of S.2. Then, it is enough to define, given an automorphism of \mathcal{C}^* , a corresponding automorphism of \mathcal{C} by looking at the correspondence

$$\{\text{vertex / edge of } \mathcal{C}\} \leftrightarrow \{\text{vertex-graph / edge-graph in } \mathcal{C}^*\}.$$

Condition S.1 ensures that not too much liberty is gained by making the aforementioned substitutions.

Proof of Proposition 4.1. For $v \in V\mathcal{C}$, recall that \mathcal{V}_v stands for the instance of \mathcal{V} inserted in place of v , and $\iota_v : \mathcal{V} \hookrightarrow \mathcal{C}^*$ is the corresponding graph embedding. For $e \in EC$, \mathcal{E}_e stands for the instance of the edge-graph with label $\mu(e)$, inserted in place of e , while $\iota_e : \mathcal{E}_{\mu(e)} \hookrightarrow \mathcal{C}^*$ is the corresponding embedding.

Constructing a morphism $\mathfrak{B} : \text{Aut } \mathcal{C}^ \rightarrow \text{Aut } \mathcal{C}$.* Fix an automorphism $\phi \in \text{Aut } \mathcal{C}^*$. We claim that, for any edge $e \in EC$, there exists a unique $e' \in EC$ such that ϕ restricts to an isomorphism $\phi|_{\mathcal{E}_e} : \mathcal{E}_e \rightarrow \mathcal{E}_{e'}$ and $\mu(e) = \mu(e')$. From this follows that $\phi : h^+(e) \mapsto h^+(e')$ and $\phi : h^-(e) \mapsto h^-(e')$, since $\mathcal{E}_{\mu(e)}$ has trivial automorphism group.

Let, without loss of generality, $s_1 < \dots < s_n$ be the ordering on the edge labels required by condition S.2. We shall verify by reverse induction on $1 \leq i \leq n$ that

- (*) if e has label $\mu(e) = s_i$, then there exists a unique e' such that $\phi|_{\mathcal{E}_e}$ is an isomorphism from \mathcal{E}_e to $\mathcal{E}_{e'}$, and $\mu(e') = s_i$ (as noted above, this implies that we have $\phi|_{\mathcal{E}_e} : h^\pm(e) \mapsto h^\pm(e')$).

If $i = n$, we know by the hypothesis (since condition S.2 is satisfied) that the only subgraphs of $\mathcal{C}^* = \mathcal{C}_{(n)}^*$ isomorphic to \mathcal{E}_{s_n} are the graphs \mathcal{E}_e with e an edge labelled s_n . Fix such an edge: then $\phi|_{\mathcal{E}_e}$ defines an isomorphism onto its image, which must therefore be of the form $\mathcal{E}_{e'}$, for some e' of label s_n . This provides the induction base.

Now, fix $i \leq n - 1$ and assume that property (*) holds for any $i + 1 \leq j \leq n$. Since, for all $i + 1 \leq j \leq n$, ϕ sends instances of \mathcal{E}_{s_j} to instances of \mathcal{E}_{s_j} , and their vertices $h^\pm(e)$ to themselves, it restricts to an automorphism of $\mathcal{C}_{(i)}^*$. Then, by applying S.2, we know that ϕ must send an instance of \mathcal{E}_{s_i} to another one. Indeed, $\phi|_{\mathcal{E}_{s_i}}$ has range in $\mathcal{C}_{(i)}^*$, while S.2 guarantees that its image must then be an instance of \mathcal{E}_{s_i} . This proves the induction step and our claim is thus verified.

Now, let us show that for each vertex $v \in V\mathcal{C}$, there exists a unique $v' \in V\mathcal{C}$ such that ϕ restricts to an isomorphism $\phi|_{\mathcal{V}_v} : \mathcal{V}_v \rightarrow \mathcal{V}_{v'}$. Indeed, it follows from the above that ϕ sends edge-graphs to edge-graphs, and boundary vertices (those of the form $h^\pm(e)$) to boundary vertices, hence it restricts to an automorphism of the subgraph of \mathcal{C}^* consisting only of the vertex-graphs employed in the construction. The latter is just a disjoint union of all vertex-graphs. Since $\phi(\mathcal{V}_v)$ is connected, it must lie in some $\mathcal{V}_{v'}$, and thus coincide with it.

Let $\psi := \mathfrak{B}(\phi)$ be defined as follows: $\psi(v)$ is the unique v' such that ϕ sends \mathcal{V}_v to $\mathcal{V}_{v'}$, and $\psi(e)$ is the unique e' such that ϕ sends \mathcal{E}_e to $\mathcal{E}_{e'}$.

Observe that ψ is bijective since applying the above construction to ϕ^{-1} yields an inverse map to ψ , and it remains to verify that ψ is actually a morphism of graphs, i.e. ψ preserves adjacency.

Recall that $h^+(e) \in \mathcal{V}_u$ if and only if u is the origin of e . Fix an edge $e \in EC$ with origin $u \in VC$, such that $h^+(e) \in \mathcal{V}_u$. Let $u' := \psi(u)$ and $e' := \psi(e)$. Then

$$h^+(e') = \phi(h^+(e)) \in \phi(\mathcal{V}_u) = \mathcal{V}_{u'},$$

and since $h^+(e') \in \mathcal{V}_{u'}$, u' is the origin of e' . Replacing $+$ by $-$ and “origin” by “terminus”, we conclude that ψ is indeed a morphism of graphs.

\mathfrak{B} is a homomorphism. Let $\phi, \phi' \in \text{Aut}(\mathcal{C}^*)$ be two automorphisms, and suppose that ϕ sends \mathcal{V}_v to $\mathcal{V}_{v'}$, while ϕ' sends $\mathcal{V}_{v'}$ to $\mathcal{V}_{v''}$. Then $\phi'\phi$ sends \mathcal{V}_v to $\mathcal{V}_{v''}$, and $\mathfrak{B}(\phi'\phi)(v) = v'' = \mathfrak{B}(\phi')(\mathfrak{B}(\phi)(v))$. An analogous statement holds for edges, and thus \mathfrak{B} is a morphism of groups.

Injectivity of \mathfrak{B} . Let ϕ be an element of $\text{Aut}\mathcal{C}^*$. If $e, e' \in EC$ are two edges, then $\phi(\mathcal{E}_e) = \mathcal{E}_{e'}$ if and only if $\mathfrak{B}(\phi)(e) = e'$. If $\mathfrak{B}(\phi) = \text{id}$, then we have that $\phi(\mathcal{E}_e) = \mathcal{E}_e$. However, edge graphs have trivial automorphism group, and thus $\phi|_{\mathcal{E}_e} = \text{id}$. Since the graph \mathcal{C}^* is folded, we obtain that ϕ is the identity on \mathcal{C}^* .

Surjectivity of \mathfrak{B} . Let $\psi \in \text{Aut}\mathcal{C}$. Our goal is to find some $\phi \in \text{Aut}\mathcal{C}^*$ such that $\mathfrak{B}(\phi) = \psi$. Define

$$\begin{aligned} \phi|_{\mathcal{E}_e} &:= \iota_{\psi(e)} \circ (\iota_e|_{\mathcal{E}_e})^{-1}, & \text{for every edge } e \in EC, \\ \phi|_{\mathcal{V}_v} &:= \iota_{\psi(v)} \circ (\iota_v|_{\mathcal{V}_v})^{-1}, & \text{for every vertex } v \in VC. \end{aligned}$$

By construction, we see already that, assuming ϕ to be a well-defined automorphism, its image under \mathfrak{B} is ψ , and it remains to verify that the piecewise definitions above actually agree. The only vertices lying in two different subgraphs (one vertex- and one edge-subgraph) are the boundary vertices. Fix an edge e with terminus v and label s , such that $\iota_e(h_s^-) = h^-(e) = \iota_v(\chi(e, -))$. Then

$$\begin{aligned} \phi|_{\mathcal{E}_e}(h^-(e)) &= \iota_{\psi(e)} \iota_e^{-1} \iota_e(h_s^-) = \iota_{\psi(e)}(h_s^-) = h^-(\psi(e)), \\ \phi|_{\mathcal{V}_v}(h^-(e)) &= \iota_{\psi(v)} \iota_v^{-1} \iota_v(\chi(e, -)) = \iota_{\psi(v)}(\chi(e, -)) = h^-(\psi(e)) \end{aligned}$$

and $\phi|_{\mathcal{E}_e}(h^-(e)) = \phi|_{\mathcal{V}_v}(h^-(e))$, as desired. Replacing $-$ by $+$ and “terminus” by “origin”, we conclude that $\phi|_{\mathcal{E}_e}(h^+(e)) = \phi|_{\mathcal{V}_v}(h^+(e))$, as well.

Finally, observe that ϕ is an automorphism since the construction above applied to ψ^{-1} , the inverse of ψ , yields an inverse of ϕ . □

5. ADDING MORE FACTORS

In this section we shall concentrate mostly on properties of telescopic groups, which do not have direct applications to the combinatorial results from Section 7, and rather stay on the purely group-theoretic side of our study.

PROPOSITION 5.1. *Let G_1 be a freely telescopic group and let G_2 be a group having a finite index subgroup H such that $N_{G_2}(H) = H$. Then $G_1 * G_2$ is freely telescopic.*

Note that, in particular, the free product of two freely telescopic groups is freely telescopic. The proof follows the same kind of argument as Lemmas 3.5–3.6.

Proof. Let Γ be a finite group. Since G_1 is freely telescopic, there exists a finite-index free subgroup H_1 satisfying $N_{G_1}(H_1)/H_1 \cong \Gamma$. By assumption, there exists a finite-index free subgroup $H_2 < G_2$ such that $N_{G_2}(H_2)/H_2$ is trivial.

Consider the inclusions morphisms $\phi_i : G_i \rightarrow G_1 \times G_2$, for $i = 1, 2$, and the morphism $\phi : G_1 * G_2 \rightarrow G_1 \times G_2$ induced by the universal mapping property of the free product. The latter is easily seen to be surjective.

Observe that $H_1 \times H_2$ is a finite index subgroup of $G_1 \times G_2$, and that the normaliser $N_{G_1 \times G_2}(H_1 \times H_2)/H_1 \times H_2$ is isomorphic to $N_{G_1}(H_1)/H_1 \times N_{G_2}(H_2)/H_2$, while the latter is simply $N_{G_1}(H_1)/H_1$. Then, let us consider the preimage \tilde{H} of $H_1 \times H_2$ under the morphism ϕ defined above. Since ϕ is an epimorphism, by Lemma 3.4, we have immediately that $G_1 * G_2$ is telescopic. It remains to show that $\tilde{H} = \phi^{-1}(H_1 \times H_2)$ is a free subgroup of $G_1 * G_2$.

By Kurosh's Subgroup Theorem, \tilde{H} can be written as $\tilde{H} = F(X) * (*_i u_i K_{1,i} u_i^{-1}) * (*_j w_j K_{2,j} w_j^{-1})$ for a number of subgroups $K_{1,i}$ of G_1 , $K_{2,j}$ of G_2 , a subset X of $G_1 * G_2$ and words u_i, w_j in $G_1 * G_2$. If \tilde{H} were not free, then some of the $K_{1,i}, K_{2,j}$ would not be free either. Assume, without loss of generality, that $K = K_{1,i}$ is not free. Then, by construction, $u_1 K u_1^{-1}$ is a non-free subgroup of \tilde{H} and thus

$$\phi(u_1 K u_1^{-1}) = \phi(u_1) \phi(K) \phi(u_1)^{-1} = \phi(u_1) \phi_1(K) \phi(u_1)^{-1}$$

is a non-free subgroup of H_1 , by using the injectivity of ϕ_1 . Therefore H_1 cannot be free, and the proposition follows. \square

6. AN ASYMPTOTIC ESTIMATE

THEOREM 6.1. *Let T be a finite free product of cyclic groups, different from $\mathbb{Z}_2 * \mathbb{Z}_2$. Then for any finite group Γ , there exist constants $A > 1$, $B > 0$ and $M \in \mathbb{N}$ such that for all $d \geq M$ the set $F(T, \Gamma, d) = \{\text{free subgroups } H < T \text{ of index } \leq d \text{ with } N_T(H)/H \cong \Gamma, \text{ up to conjugacy}\}$ has cardinality $\geq A^{Bd \log d}$.*

The proof of the above theorem is conceptually simple. Fix the aforementioned groups T, Γ and index d , and let $N = N(T, \Gamma, d)$ be an integer whose dependence on T, Γ, d will be clarified later. We shall verify that sufficiently many non-isomorphic graphs \mathcal{H}_σ on N vertices can be built, with respect to an additional parameter σ , introduced below. Then the previously used edge-links in the construction of the Schreier graph for $H < T$ (based on a Cayley graph of Γ) will be combined with one of many possible choices of \mathcal{H}_σ . Once we show that there are $\geq A^{Bd \log d}$ non-isomorphic instances of \mathcal{H}_σ , the result follows. On the other hand, the cardinality of $F(T, \Gamma, d)$ is $\leq C^{Dd \log d}$, for some constants $C > 1, D > 0$, with $d \geq M$ by [2]. Thus, the *growth type* of the cardinality of $F(T, \Gamma, d)$ is d^d , and is independent of T and Γ .

One natural condition on the graphs \mathcal{H}_σ is that they do not contain "roots" that have property P from Section 3.2.3, i.e. vertices v for which

- $v \cdot r = v \cdot c$, if $T = \mathbb{Z}_p * \mathbb{Z}_q$ (with $p \geq q, p \geq 3, q \geq 2$), or
- $v \cdot r g b g r g r = v$ and $v \cdot b = v \cdot g$, if $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$.

If an \mathcal{H}_σ contained such a root, then the combination of an edge-link with \mathcal{H}_σ would not have a *unique* root any more.

The following lemma provides all the necessary details about constructing \mathcal{H}_σ 's, which is the first step towards the proof of Theorem 6.1.

LEMMA 6.2. *Let T be $\mathbb{Z}_p * \mathbb{Z}_q$, with $p \geq q, p \geq 3, q \geq 2$, resp. $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$. Then there exist constants $k, K > 0$, such that for any N a multiple of pq , resp. N a multiple of 8, we have at least KN^{kN} non-isomorphic non-degenerate $(T, \{r, c\})$ -graphs, resp. $(T, \{r, g, b\})$ -graphs, with two split vertices.*

Proof. First, we treat the case $T = \mathbb{Z}_p * \mathbb{Z}_q$, with $p \geq q, p \geq 3$ and $q \geq 2$. Let N be a multiple of pq , and let us consider N vertices ordered as v_0, \dots, v_{N-1} . Similar to Section 3.2.3, we start by drawing red p -cycles of the form

$$v_{kp} \rightarrow v_{kp+1} \rightarrow \dots \rightarrow v_{kp+p-1} \rightarrow v_{kp},$$

for each $k = 0, \dots, N/p - 1$. Then we draw a cyan edge $v_0 \rightarrow v_1$, which creates a “double edge” from v_0 to v_1 , and continue by adding as few cyan q -cycles as possible in order to obtain a connected graph, without creating any more double edges. Let \mathcal{H}_\emptyset denote the resulting graph.

One sees that since the initial graph is a disjoint union of red p -cycles, then $\leq 2N/p$ vertices will have to be joined by cyan edges in order to get a connected graph, and then at most q other vertices will be joined in order to close up the last cyan q -cycle. Observe that at least $N - 2N/p - q \geq N/3 - q \geq N/4$ (for N large enough) vertices will remain “free” in \mathcal{H}_\emptyset . Let $F \geq N/4$ be the exact number of remaining free vertices, and let $D := \lfloor F \rfloor_q$, where $\lfloor x \rfloor_q$ denotes the greatest multiple of q that is smaller than x , for any natural numbers q and x . Then, for N large enough, we also have $D \geq N/4$.

Now, let us consider any ordering σ on the free vertices in \mathcal{H}_\emptyset . By drawing cyan edges between the consecutive vertices in σ , and closing up the cycles when necessary, any such choice of σ yields a non-degenerate graph \mathcal{H}_σ .

First of all, observe that after fixing a basepoint v_0 in \mathcal{H}_\emptyset , the graphs \mathcal{H}_σ and \mathcal{H}_τ , for any orderings σ and τ , are basepoint-isomorphic if and only if $\mathcal{H}_\sigma = \mathcal{H}_\tau$. Indeed, since $\mathcal{H}_\emptyset \leq \mathcal{H}_\sigma$, any isomorphism $\mathcal{H}_\sigma \rightarrow \mathcal{H}_\tau$ restricts to an embedding $\mathcal{H}_\emptyset \hookrightarrow \mathcal{H}_\tau$, only one of which exists by foldedness: namely, the identity map.

Our next step is estimating the number of non-isomorphic graphs \mathcal{H}_σ , as above, without a basepoint. To this end, observe that two orderings σ and τ of the remaining free vertices yield the same graph if and only if τ can be obtained from σ by a permutation of the respective “ q -blocks” (each consisting of the vertices in a cyan cycle created in accordance with σ or τ), as well as by cyclic permutations within each “block”. Indeed, permuting the blocks, together with cyclic permutations inside each block, obviously does not change the resulting graph. Also, if two orderings result in the same graph, then the graph isomorphism implies that the orderings are equal up to the aforementioned permutations. Thus, the number of non-basepoint-isomorphic extensions of \mathcal{H}_\emptyset by using orderings is

$$\geq \frac{D!}{(D/q)! q^{D/q}}.$$

However, since the resulting graphs will be used later on in the constructions of “edge-links”, we do not want them to contain any double edges. This means that no vertices v with property $P : v \cdot r = v \cdot c$, except for $v_0 \rightarrow v_1$ (joined by a double edge by construction) are allowed, since this would contradict the uniqueness of roots (in the sense of having property P from Section 3.2.3).

We will therefore consider only the orderings that will not produce such edges. Therefore, we start with D free vertices and at each step choose any of the remaining vertices, with the restriction that no double edge be created. Observe that there are three possible cases: the to-be-chosen vertex can be the first of a new cyan cycle, it can lie in the middle of such a cycle or be its last one with respect to the chosen ordering. In the first case, no restriction is imposed since choosing the first vertex of a cycle does not create edges at all. In the second case, one edge is created, and thus one vertex is “illegal” in the sense that choosing it would result in a double edge. In the third case, two vertices are “illegal”.

We have to determine if any initial choice of vertices in cyan cycles can be extended to a legal ordering as above. Such an extension may fail under certain circumstances: e.g. if two vertices remain to be ordered, then no legal choice may be possible. In order to circumvent this problem, we shall only consider orderings of $D - q$ vertices, since any legal prefix of length $D - q$ can be extended to a legal ordering of length D (if $q = 2$, then one has to consider orderings of the first $D - 4$ vertices instead).

With the above points in mind, one readily sees that the number of legal orderings is at least $\geq (D - 2)(D - 3) \dots (q - 1) = (D - 2)! / (q - 2)!$. Once again, taking into account the fact that orderings yield non-basepoint isomorphic graphs if and only if they are obtained one from another by permutations of the blocks and by cyclic permutations within those blocks, we conclude the existence of

$$(*) \qquad \geq \frac{1}{(q - 2)!} \cdot \frac{(D - 2)!}{(D/q)! q^{D/q}}$$

legal choices of σ defining non-isomorphic graphs \mathcal{H}_σ without basepoint.

Now, take any \mathcal{H}_σ , split its vertex v_0 , and call the resulting graph \mathcal{H}'_σ , while v_r, v_c will be its dangling vertices. We claim that:

- (1) each \mathcal{H}'_σ is connected;
- (2) if σ and τ are distinct (up to the aforementioned permutations), then \mathcal{H}'_σ and \mathcal{H}'_τ are non-isomorphic without basepoint;
- (3) \mathcal{H}'_σ does not have double edges.

The first property holds since we split at the *unique* double edge of \mathcal{H}_σ . The second property follows from the fact any isomorphism between \mathcal{H}'_σ and \mathcal{H}'_τ must fix the dangling vertices. Then, were the split graphs \mathcal{H}'_σ and \mathcal{H}'_τ isomorphic, then \mathcal{H}_σ and \mathcal{H}_τ would be so. Finally, since we always split the vertex v_0 of \mathcal{H}_σ at the end of its unique double edge, none of the latter remains, and the third property is verified.

Thus we conclude the existence of as many as (*) non-isomorphic graphs \mathcal{H}'_σ , without a root and having two dangling vertices. By invoking Stirling's formula and using the fact that $D - 2 \geq \frac{D}{q}$ for D big enough, the desired asymptotic estimate follows.

The argument for $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ is similar and not spelled out in detail here. First, we take N a big enough multiple of 8 and then assemble a chain of alternating red-green edges connecting N vertices. Finally, we estimate the number of different graphs obtained by adding blue edges in an analogous way to the above argument. \square

Proof of Theorem 6.1. Observe that it suffices to verify the statement for $T = \mathbb{Z}_p * \mathbb{Z}_q$, with $p \geq q$, $p \geq 3$, $q \geq 2$, and $T = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$, only. Indeed, for any surjection $f : G \rightarrow K$ and any subgroups $K_1, K_2 \leq K$, their pre-images $f^{-1}(K_1)$ and $f^{-1}(K_2)$ are conjugate if and only if K_1 and K_2 are so. Therefore, the inductive constructions of Section 3.3 preserve the size of the set in the statement of the theorem.

Assume that T is either $\mathbb{Z}_p * \mathbb{Z}_q$, as above, or $\mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ and choose N as in Lemma 6.2. Then there are at least KN^{kN} non-isomorphic graphs \mathcal{H}'_σ with two dangling vertices.

Let us fix such a graph \mathcal{H}_σ . Then consider the edge-link \mathcal{L}_e constructed in Section 3.2.3, and glue \mathcal{H}'_σ to \mathcal{L}_e via the dangling vertices in order to obtain $\mathcal{L}_e + \mathcal{H}'_\sigma$. The graph $\mathcal{L}_e + \mathcal{H}'_\sigma$ still has a unique root, which is the only condition for the arguments of Section 3.2 to hold. Following those constructions once again, for a given finite group Γ , results in a finite-index subgroup $H_\sigma < T$ with $N_T(H_\sigma)/H_\sigma \cong \Gamma$.

For any two choices of graphs \mathcal{H}'_σ and \mathcal{H}'_τ as above, the resulting subgroups H_σ and H_τ are not conjugate. Indeed if H_σ and H_τ were conjugate, then $\text{Sch}_{T,X}(H_\sigma)$ and $\text{Sch}_{T,X}(H_\tau)$ would be isomorphic (with base-points not necessarily matched), so that the roots of the former Schreier graph would be sent to the roots of the latter one, and the instances of \mathcal{H}'_σ would be mapped isomorphically to the instances of \mathcal{H}'_τ , which would imply that σ and τ coincide up to the above mentioned permutations. Thus, there are $\geq KN^{kN}$ non-conjugate free subgroups H_σ of T with $N_T(H_\sigma)/H_\sigma \cong \Gamma$ having same index.

Finally, let us estimate the index of H . For any fixed Γ with a given generating set S , let e_Γ and v_Γ be the following numbers:

- $e_\Gamma :=$ the number of edge-links appearing in the original construction of Section 3.2.3 = $|\Gamma| \cdot |S| \cdot (|S| - 1)/2$;
- $v_\Gamma :=$ the number of vertices appearing in the original construction.

Now, with the modified construction involving \mathcal{H}'_σ , we obtain that the index of the corresponding subgroup H_σ is

$$d = v_\Gamma + e_\Gamma N,$$

which is the number of vertices in the resulting Schreier graph. Thus, for such a d and N sufficiently large, there are

$$\geq K \left(\frac{d - v_\Gamma}{e_\Gamma} \right)^{k \frac{d - v_\Gamma}{e_\Gamma}}$$

conjugacy classes of index d subgroups $H < T$ such that $N_T(H)/H \cong \Gamma$.

In order to conclude for an arbitrary d , as in the statement, just consider index $\lfloor d \rfloor_{v_\Gamma + e_\Gamma N}$ subgroups. Then the above construction of \mathcal{H}_σ 's yields the desired result. \square

7. SYMMETRIES OF MAPS, PAVINGS AND CONSTELLATIONS

As mentioned in Section 1, an oriented map M on n darts can be thought of as index n free subgroup H_M of $\Delta^+ = \mathbb{Z}_2 * \mathbb{Z}$. Moreover, the group of orientation-preserving automorphisms of M is $\text{Aut } M \cong N(H_M)/H_M$, where $N(H_M)$ is the normaliser of H_M in Δ^+ , cf. [3] and [13, Theorem 3.8].

An analogous statement holds for several other classes of combinatorial objects, such as

- oriented hypermaps, with $\Delta^+ = \mathbb{Z} * \mathbb{Z}$ [16, 17];
- oriented (p, q) -hypermaps, with $\Delta^+ = \mathbb{Z}_p * \mathbb{Z}_q$ ($p \geq 3, q \geq 2$) [5];
- oriented pavings (or three-dimensional maps), with $\Delta^+ = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$ [4],
- length $k \geq 3$ constellations (in the sense of [15, Definition 1.1.1]), with $\Delta^+ = \underbrace{\mathbb{Z} * \dots * \mathbb{Z}}_{k-1}$.

The following theorem generalises the respective results of [7, 8, 11, 19] to the case of (p, q) -hypermaps.

THEOREM 7.1. *For any finite group Γ , and n sufficiently large, there exist $\sim n^n$ non-isomorphic oriented (p, q) -hypermaps H on n darts with $\text{Aut } H \cong \Gamma$.*

Here by the symbol $f(x) \sim x^x$, for a function $f(x) : \mathbb{N} \rightarrow \mathbb{R}$, we mean its *rate of growth*, i.e. that there exist positive constants A_1, B_1, A_2, B_2 such that $A_1 x^{B_1 x} \leq f(x) \leq A_2 x^{B_2 x}$, for x sufficiently large.

The family of (p, q) -hypermaps naturally comprises the cases of maps ($p = \infty, q = 2$) and hypermaps ($p = q = \infty$). Some other interesting classes of maps, such as triangulations ($p = 3, q = 2$) and quadrangulations ($p = 4, q = 2$) of surfaces, or their bi-coloured triangulations ($p = q = 3$) also satisfy the above theorem.

As well, we can now easily estimate the number of three-dimensional pavings, cf. [1, 4], with a given automorphism group.

THEOREM 7.2. *For any finite group Γ , and n sufficiently large, there exist $\sim n^n$ non-isomorphic oriented pavings P on n darts with $\text{Aut } P \cong \Gamma$.*

An analogous result holds for constellations as defined in [15, Definition 1.1.1].

THEOREM 7.3. *For any finite group Γ , any $k \geq 3$, and n sufficiently large, there exist $\sim n^n$ non-isomorphic length k constellations C on n darts with $\text{Aut } C \cong \Gamma$.*

Since each length $k \geq 3$ constellation defines a branched covering of the sphere \mathbb{S}^2 with k branch points, we can reformulate the above theorem in a more geometric language.

THEOREM 7.4. *For any finite group Γ , any $k \geq 3$, and n sufficiently large, there exist $\sim n^n$ non-isomorphic degree n branched coverings of \mathbb{S}^2 with k branch points and deck transformation group Γ .*

All the above theorems are fairly obvious corollaries of the results in Section 6. Indeed, the lower bound on the number of non-isomorphic hypermaps, resp. pavings, on n darts with given automorphism group follows from Theorem 6.1, and the upper bound of $(n!)^k$, with an appropriate fixed $k \geq 2$, is trivial. Then, invoking Stirling’s formula provides the rate of growth $\sim n^n$.

The statements of Theorems 7.1–7.4 should be contrasted with [9, Lemma 1], which implies that the corresponding combinatorial objects “mostly” have only trivial automorphism groups, i.e. are *asymmetric*. However, those with a given non-trivial automorphism group are still numerous, though not as abundant as asymmetric ones.

Apparently, this technique can be applied to many naturally arising classes of oriented maps, hypermaps, and pavings: exactly those describable as free subgroups of a certain “universal group” Δ^+ , which is a finite free product of cyclic groups. For any such class we obtain that any finite group is realisable as automorphism group by infinitely many of its members (more precisely, super-exponentially many depending on the number of darts). It is worth mentioning that not all families of maps admit such a wide variety of symmetries, e.g. the maps with underlying graph a tree [10].

Acknowledgements. The authors would like to thank Derek Holt (University of Warwick, UK) for answering an author’s question on math.stackexchange⁽²⁾, Laura Ciobanu (Heriot–Watt University, UK), Adrian Tanasa (Université Bordeaux, France) and Gareth A. Jones (University of Southampton, UK) for numerous fruitful discussions, and the anonymous referee for helpful remarks and suggestions.

REFERENCES

- [1] Didier Arquès and Patrice Koch, *Pavages tridimensionnels*, Bigre + Globule **61–62** (1989), 5–15.
- [2] Béla Bollobás, *The asymptotic number of unlabelled regular graphs*, J. Lond. Math. Soc. **26** (1982), 201–206.
- [3] Antonio Breda d’Azevedo, Alexander Mednykh, and Roman Nedela, *Enumeration of maps regardless of genus: geometric approach*, Discrete Math. **310** (2010), no. 6–7, 1184–1203.
- [4] Laura Ciobanu and Alexander Kolpakov, *Three-dimensional maps and subgroup growth*, <https://arxiv.org/abs/1712.01418>, 2017.
- [5] ———, *Free subgroups of free products and combinatorial hypermaps*, Discrete Math. **342** (2019), no. 5, 1415–1433.
- [6] Robert Cori, *Un code pour les graphes planaires et ses applications.*, Astérisque, vol. 27, Société Mathématique de France (SMF), Paris, 1975.
- [7] Robert Cori and Antonio Machì, *Construction of maps with prescribed automorphism group*, Theor. Comput. Sci. **21** (1982), no. 1, 91–98.
- [8] ———, *Maps, hypermaps and their automorphisms: a survey, I, II, III*, Expo. Math. **10** (1992), 403–427, 429–447, 449–467.
- [9] Michael Drmota and Roman Nedela, *Asymptotic enumeration of reversible maps regardless of genus*, Ars Math. Contemp. **5** (2012), 77–97.

⁽²⁾<https://math.stackexchange.com/questions/2767720/preimage-of-product-of-free-subgroups-in-free-product-is-free>

- [10] Valeriĭ Zalmanovich Feĭnberg, *Automorphism groups of trees*, Dokl. Akad. Nauk BSSR **13** (1969), 1065–1067.
- [11] Roberto Frucht, *Herstellung von Graphen mit vorgegebener abstrakter Gruppe*, Compos. Math. **6** (1938), 239–250.
- [12] Gareth A. Jones, *Realisation of groups as automorphism groups in permutational categories*, <https://arxiv.org/abs/1807.00547>, 2018.
- [13] Gareth A. Jones and David Singerman, *Theory of maps on orientable surfaces*, Proc. Lond. Math. Soc. **S3-37** (1978), no. 2, 273–307.
- [14] Ilya Kapovich and Alexei Myasnikov, *Stallings foldings and subgroups of free groups*, J. Algebra **248** (2002), no. 2, 608–668.
- [15] Sergei K. Lando and Alexander K. Zvonkin, *Graphs on surfaces and their applications*, Encycl. Math. Sci., vol. 141, Springer, 2004.
- [16] Alexander Mednykh and Roman Nedela, *Enumeration of unrooted maps of a given genus*, J. Comb. Theory, Ser. B **96** (2006), no. 5, 706–729.
- [17] ———, *Enumeration of unrooted hypermaps of a given genus*, Discrete Math. **310** (2010), no. 3, 518–526.
- [18] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003.
- [19] Josef Širáň and Martin Škoviera, *Orientable and non-orientable maps with given automorphism groups*, Australas. J. Comb. **7** (1993), 47–53.
- [20] Jean-Claude Spehner, *Merging in maps and in pavings*, Theor. Comput. Sci. **86** (1991), no. 2, 205–232.
- [21] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565.
- [22] W. Wilson Stothers, *Free subgroups of the free product of cyclic groups*, Math. Comput. **32** (1978), no. 144, 1274–1280.
- [23] William T. Tutte, *What is a map?*, in New Directions in Graph Theory (F. Harary, ed.), Academic Press, New York, 1973, pp. 309–325.
- [24] Cora Welsch, *Subgroup graph methods for presentations of finitely generated groups and the contractibility of associated simplicial complexes*, J. Algebra **486** (2017), 119–156.

RÉMI BOTTINELLI, Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2000 Neuchâtel, Switzerland
E-mail : remi.bottinelli@unine.ch

LAURA GRAVE DE PERALTA, Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2000 Neuchâtel, Switzerland
E-mail : laura.grave@unine.ch

ALEXANDER KOLPAKOV, Institut de Mathématiques, Université de Neuchâtel, Rue Emile-Argand 11, CH-2000 Neuchâtel, Switzerland
E-mail : kolpakov.alexander@gmail.com