On on a conjecture of Karrass and Solitar

Warren Dicks and Benjamin Steinberg

December 15, 2015

1 Graphs

1.1 Definitions. A *graph* Γ consists of a set V of *vertices*, a set E of *edges*, an *initial incidence function* $\iota : E \to V$, and a *terminal incidence function* $\tau : E \to V$. Sometimes we shall treat Γ as the set $V \mathbin{\dot\cup} E$, where $\dot\cup$ denotes disjoint union. Other times, we shall treat Γ as a one-dimensional CW-complex in which each edge e corresponds to the $]0,1[$ part of a copy of $[0,1]$ that is given with an attaching map that carries 0 to ιe and 1 to τe.

In the free group $\langle E \mid \;\; \rangle$, set $E^{\pm 1} := E \cup E^{-1}$. For each $e \in E$, set $\iota(e^{-1}) := \tau(e)$ and $\tau(e^{-1}) := \iota(e)$. For each $v \in V$, set $\operatorname{link}_\Gamma(v) := \{e \in E^{\pm 1} : \iota e = v\}$ and $\deg_\Gamma(v) := |\operatorname{link}_\Gamma(v)|$, called the *degree* of v in Γ. By a Γ-*path* we mean a sequence γ of the form

$$(v_0, e_1, v_1, e_2, v_2, \ldots, v_{n-1}, e_n, v_n)$$

where $n \geqslant 0$ and, for each $i \in \{1, 2, \ldots, n\}$, $e_i \in E^{\pm 1}$, $v_{i-1} = \iota e_i$ and $v_i = \tau e_i$; we sometimes find it helpful to depict γ as $v_0 \xrightarrow{e_1} v_1 \xrightarrow{e_2} v_2 \cdots v_{n-1} \xrightarrow{e_n} v_n$. We call n the *length* of γ. For each $i \in \{1, 2, \ldots, n-1\}$, we say that (e_i, v_i, e_{i+1}) is a *turn in* γ, and that the turn is a *backtracking* if $e_{i+1} = e_i^{-1}$. We say γ is *closed* if $v_0 = v_n$, γ is *reduced* if it has no backtrackings, and γ is a *circle-path* if it is closed, reduced, and satisfies $|\{v_1, v_2, \ldots, v_n\}| = n \geqslant 1$.

Let $E^{\pm 1 *}$ denote the free monoid on $E^{\pm 1}$, endowed with the natural inversion map. We call $\prod_{i=1}^{n} e_i \in E^{\pm 1 *}$ the Γ-*label* of γ, and let $\operatorname{Labels}(\Gamma) \subseteq E^{\pm 1 *}$ denote the set of Γ-labels of Γ-paths. All the Γ-paths of length 0 have Γ-label 1, while each Γ-path of positive length can be reconstructed from its Γ-label. The subset $\operatorname{Labels}(\Gamma)$ of $E^{\pm 1 *}$ is closed under the inversion operation, but need not be closed under the multiplication operation. We shall be performing only multiplications that do produce another element of $\operatorname{Labels}(\Gamma)$.

A graph is *connected* if any two vertices lie in some path. Every graph is the disjoint union of its maximal nonempty, connected subgraphs, called its *components*. A *circle-graph* is a finite, nonempty, connected graph in which each vertex has degree two. A *tree* is a nonempty, connected graph with no circle-subgraphs. A *line-segment-graph* is a finite tree in which exactly two vertices have degree one.

2 Arborizing Steinberg's argument

The case of the following where $C = \{1\}$ is due to Steinberg.

2.1 Theorem. *Let $G := A *_C B$ with C finite and $A \neq C \neq B$, and let W denote the subsemigroup of G generated by $(B - C)(A - C)$. If H is a finitely generated, infinite-index subgroup of G, then there exists some $w \in W$ such that $H \cap \bigcup_{g \in G} (wW \cup \{w\})^g = \emptyset$; hence, each element of H is non-conjugate to each cyclically reduced word that begins with w or w^{-1}.*

Proof. Let T denote the graph with vertex-set $(G/A) \mathbin{\dot\cup} (G/B)$ and edge-set G/C, with each edge gC of T having initial vertex gA and terminal vertex gB. By Bass-Serre theory, T is a G-tree. For any vertices v_1 and v_2 of T, $T[v_1, v_2]$ denotes the intersection of all subtrees of T that contain $\{v_1, v_2\}$; if $v_1 \neq v_2$, then $T[v_1, v_2]$ is a line-segment-subgraph.

Let S be a finite generating set for H. Set $T_0 := \bigcup_{s \in S \cup \{1\}} T[A, sA]$ and $H{\cdot}T_0 := \bigcup_{h \in H} h{\cdot}T_0$. Then $H{\cdot}T_0$ is an H-subgraph of T, and T_0 is a finite subtree of $H{\cdot}T_0$ such that, for each $s \in S$, $s{\cdot}T_0 \cap T_0 \neq \emptyset$. Hence, the component of $H{\cdot}T_0$ which contains T_0 is stabilized by S, and hence by H; it follows that $H{\cdot}T_0$

is an H-subtree of T. Also, $H{\cdot}T_0$ is H-finite, that is, the graph $H\backslash(H{\cdot}T_0)$ is finite. Fix a $\subseteq$-minimal H-subtree T_H of $H{\cdot}T_0$.

Consider any $w \in W$. Set $(T_H : T[A, wB]) := \{q \in G : T[qA, qwB] \subseteq T_H\}$. Write $w = \prod_{i=1}^{n}(b_i a_i)$ where $n \geqslant 1$ and, for each $i \in \{1, 2, \ldots, n\}$, $b_i \in B - C$ and $a_i \in A - C$. Let e_0 denote the edge $1C$ of T. Then the reduced T-path from A to wB is

$$A \xrightarrow{e_0} B \xrightarrow{b_1\cdot e_0^{-1}} b_1a_1A \xrightarrow{b_1a_1\cdot e_0} b_1a_1B \cdots \xrightarrow{\prod_{i=1}^{n-1}(b_ia_i)b_n\cdot e_0^{-1}} \prod_{i=1}^{n}(b_ia_i){\cdot}A \xrightarrow{\prod_{i=1}^{n}(b_ia_i)\cdot e_0} \prod_{i=1}^{n}(b_ia_i){\cdot}B.$$

Here, $e_0 \in T[A, wB]$, and, for each $q \in (T_H : T[A, wB])$, $q{\cdot}e_0 \in T_H$. Notice that $(T_H : T[A, wB])$ is a left H-subset of G and we have a map

$$H\backslash(T_H : T[A, wB]) \to H\backslash T_H, \qquad Hq \mapsto Hq{\cdot}e_0.$$

Since C is finite, this map is finite-to-one. Since $H\backslash T_H$ is finite, $H\backslash(T_H : T[A, wB])$ is finite.

Fix a $w \in W$ that $\subseteq$-minimizes the H-finite H-subset $(T_H : T[A, wB])$ of G. Fix a $w' \in wW \cup \{w\}$. It suffices to show that $w' \notin \bigcup_{g\in G} H^g$.

Suppose this fails. Fix a $q \in G$ such that $w' \in H^q$. Then $q^{-1}{\cdot}T_H$ is a $\langle w'\rangle$-subtree of T. Now $w' \in W$ and $T[A, w'B] \supseteq T[A, wB]$. Set $T_{\langle w'\rangle} := \bigcup_{n\in\mathbb{Z}} w'^n{\cdot}T[A, w'B]$. It can be seen that $T_{\langle w'\rangle}$ is a $\subseteq$-minimal $\langle w'\rangle$-subtree of T, on which w' acts by translation. If $T_{\langle w'\rangle} \cap q^{-1}{\cdot}T_H = \emptyset$, then there exist unique vertices v_1 of $T_{\langle w'\rangle}$ and v_2 of $q^{-1}{\cdot}T_H$ that $\subseteq$-minimize $T[v_1, v_2]$, and then w' fixes v_1 and v_2, which contradicts w' acting by translation on $T_{\langle w'\rangle}$. Hence, $T_{\langle w'\rangle} \cap q^{-1}{\cdot}T_H \neq \emptyset$. Now $T_{\langle w'\rangle} \cap q^{-1}{\cdot}T_H$ is a $\langle w'\rangle$-subtree of $T_{\langle w'\rangle}$, and hence is all of $T_{\langle w'\rangle}$. Thus, $T[A, wB] \subseteq q^{-1}{\cdot}T_H$. Hence, $T[qA, qwB] \subseteq T_H$. Since H has infinite index in G and C is finite, $H\backslash G/C$ is infinite. Hence, $G/C \nsubseteq T_H$. Thus, $G/B \nsubseteq T_H$. Fix a $g \in G$ such that $gB \in T - T_H$. Let v denote that vertex of T_H which $\subseteq$-minimizes $T[gB, v]$. Since T_H is the $\subseteq$-smallest subtree of T that contains $H{\cdot}v$, we see that $T_H \subseteq \bigcup_{h\in H} T_H[qA, h{\cdot}v]$. Fix an $h \in H$ such that $qwB \in T_H[qA, h{\cdot}v]$. Then $h^{-1}qwB \in T_H[h^{-1}qA, v] \subsetneq T[h^{-1}qA, gB]$. Hence, $wB \in T[A, q^{-1}hgB]$. It follows that $q^{-1}hgB = w''B$ for some $w'' \in W$. Now we have $T[A, wB] \subseteq T[A, w''B]$, and, hence, $(T_H : T[A, wB]) \supseteq (T_H : T[A, w''B])$. Equality holds by the $\subseteq$-minimality of $(T_H : T[A, wB])$, and, hence, $q \in (T_H : T[A, w''B])$. Thus, $qw''B \in T_H$. Hence, $h^{-1}qw''B \in T_H$. Since $w''B = q^{-1}hgB$, this says that $gB \in T_H$, which is a contradiction. $\square$

2.2 Remarks. In the statement of Theorem 2.1 above, the hypothesis that H is finitely generated can be weakened to the hypothesis that H is generated by some set S' of the form

(2.2.1) $\{h_1, h_2, \ldots, h_n\} \cup \big(H \cap (A^{g_1} \cup A^{g_2} \cup \cdots \cup A^{g_k} \cup B^{g_{k+1}} \cup B^{g_{k+2}} \cup \cdots \cup B^{g_m})\big), n \geqslant 0, m \geqslant k \geqslant 0,$

where $\{h_1, h_2, \ldots, h_n\} \subseteq H$ and $\{g_1, g_2, \ldots, g_k, g_{k+1}, \ldots, g_m\} \subseteq G$, since the above proof of Theorem 2.1 remains valid if S is replaced with S' and T_0 is replaced with

$$T_0' := \{A\} \cup (\textstyle\bigcup_{i=1}^{n} T[A, h_iA]\,) \cup (\bigcup_{j=1}^{k} T[A, g_j^{-1}A]\,) \cup (\bigcup_{j=k+1}^{m} T[A, g_j^{-1}B]\,).$$

Collins and Turner (1994) assigned to a subgroup of a free product of groups a certain cardinal that they called its Kurosh rank, and they pointed out that, by a result of Baer and Levi (1936), this cardinal is uniquely determined by the subgroup and the given set of free-product factors. In the $C = \{1\}$ case of Theorem 2.1, it can be shown by Bass-Serre theory that a subgroup H of $A{*}B$ has finite Kurosh rank if and only if H has a generating set as in (2.2.1).

3 Small-cancellation theory for amalgamated free products

In this section, we shall prove the following.

3.1 Theorem. *Let $G := A *_C B$ with $A \neq C \neq B$, and let $g_0 := \prod_{i=1}^{n}(b_ia_i) \in G$ where $n \geqslant 1$ and, for each $i \in \{1, 2, \ldots, n\}$, $a_i \in A - C$ and $b_i \in B - C$. Suppose that, for all cyclic permutations g_1, g_2 of the word g_0, all $\epsilon_1, \epsilon_2 \in \{-1, 1\}$, and all $c_1, c_2 \in C$,*

(3.1.1) *if $c_1g_1^{\epsilon_1} = g_2^{\epsilon_2}c_2$, then $c_1g_1^{6\epsilon_1} = g_2^{6\epsilon_2}c_1$.*

Let p_0 be a nontrivial element of the normal closure of g_0^6 in G. Then some G-conjugate of p_0 equals g_0^6, or g_0^{-6}, or a cyclically reduced product $\prod_{k=1}^{\ell}(w_k'w_k'')$ of nontrivial reduced words such that each w_k' is a subword of g_0^6 or g_0^{-6} of length $6n{-}1$, $8n{-}1$, or $10n{-}1$, and $\sum_{j=1}^{3} j{\cdot}\ell_j \geqslant 6$, where ℓ_j denotes the number of k such that w_k' has length $(j+2)2n{-}1$. In particular, p_0 is conjugate to some cyclically reduced word that begins with g_0^2 or g_0^{-2}.

3.2 Remarks. Small-cancellation theory was initiated by Dehn, Greendlinger, and Lyndon for free groups, by Lyndon for free products, and by Schupp for amalgamated free products. As our proof of Theorem 3.1 amounts to one of the fundamental arguments of this theory, we take the opportunity to present an introductory exposition of it, by writing our proof carefully in a relatively unsophisticated language, with no attempt made to draw more than the one conclusion.

The $\mathrm{C}'(\frac{1}{6})$ analogue of Theorem 3.1 for free groups is due to Greendlinger, in an extension of work of Dehn.

The case of Theorem 3.1 where $C = \{1\}$ is due to Duncan and Howie; here, (3.1.1) clearly holds.

Juhász studies small-cancellation theory for amalgamated free products in the situation where the following three conditions hold:

- C is malnormal in G;
- g_0 is not a C-proper power in the sense of Juhász; and
- $Ch^{-1}C \neq ChC$ for each letter h occurring in g_0.[1]

Together these imply that

(3.2.1) C is malnormal in G, and the set of cyclic permutations of $g_0^{\pm 1}$ meets $4n$ C-bi-cosets,

which in turn implies (3.1.1), for here $g_1 = g_2$, $\epsilon_1 = \epsilon_2$, and $c_1 = c_2 = 1$.

If C is malnormal in G and nontrivial, then both $(B - C)(A - C)$ and $(A - C)(B - C)$ contain at least two C-bi-cosets. (If $bca = c_1bac_2$, then $bC = c_1bC$ and $Ca = Cac_2$, then $c_1 = 1 = c_2$ by malnormality, and then $c = 1$.) In the situation of interest to us, we may arrange for n to be a prime number greater than 4, $Cb_na_nC \neq Cb_{n-1}a_{n-1}C$, and $Ca_{n-3}b_{n-2}C \neq Ca_{n-4}b_{n-3}C$. Then the set of cyclic permutations of g_0 meets $2n$ C-bi-cosets, and the set of cyclic permutations of g_0^{-1} meets $2n$ C-bi-cosets. To ensure (3.2.1) one could further assume that, for each $i \in \{1, 2, \ldots, n\}$, $Cb_ia_iC \neq Cb_i^{-1}a_{i-1}^{-1}C$ where $a_0 := a_n$, but that is somewhat unsatisfactory; it does hold if $Cb_i^{-1}C \neq Cb_iC$.

3.3 Definitions. Set $\widehat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$, the Riemann sphere.

Let Γ be a finite, nonempty, connected graph engraved on $\mathbb{C}$ such that Γ is the one-skeleton of a CW-structure for $\widehat{\mathbb{C}}$ and the edges of Γ are smooth curves. Thus, each component of $\widehat{\mathbb{C}} - \Gamma$ is a topological open disk, called a *face for* Γ. Let V, E, and F, denote the sets of vertices of Γ, edges of Γ, and faces for Γ, respectively.

Consider any $v \in V$. Some small circle Δ_v in $\mathbb{C}$ centered at v has an induced clockwise circle-graph structure in which the Δ_v-vertices and Δ_v-edges are the intersections of Δ_v with the Γ-edges and Γ-faces incident to v, with multiplicities. We identify the vertex-set of Δ_v with $\mathrm{link}_\Gamma(v)$. Any edge of Δ_v is a subset of some $f \in F$, and its initial vertex e and terminal vertex e' lie in $\mathrm{link}_\Gamma(v)$; here, we call (e^{-1}, v, e') a *turn for* f.

If E is nonempty and f is a face for Γ, then the set of all the turns for f may be concatenated, yielding a sequence of the form $(e_1, v_1, e_2, v_2, \cdots, e_n, v_n, e_1)$ with $n \geqslant 1$. The closed Γ-path $(v_n, e_1, v_1, e_2, v_2, \cdots, e_n, v_n)$ is called a *boundary path* for f. The set of boundary paths for f forms a single orbit under cyclic shifting. In particular, f uniquely determines a conjugacy class, denoted $[f]$, in $\langle E \mid \quad \rangle$ such that $\prod_{i=1}^{n} e_i \in [f]$. By the *boundary* of f we mean the subgraph of Γ with vertex-set $\{v_1, v_2, \ldots, v_n\}$ and edge-set $\{e_1, e_2, \ldots, e_n\}$. The boundary of f is a circle-graph if and only if one/every boundary path for f is a circle-path.

For any circle-subgraph Δ of Γ, the two faces for Δ induce a partitioning of F into two proper subsets, and each edge of Δ lies in the boundary of two different faces for Γ.

It follows from Lemma 2 of van Kampen (1933) that, for any enumeration $(f_1, f_2, \ldots, f_m)$ of the set F,

$$[f_1]^{-1} \ \subseteq\ [f_2] \cdots [f_m] \ \subseteq\ \langle E \mid \quad \rangle; \tag{3.3.1}$$

to see (3.3.1), one finds some enumeration $(f_1, f_2, \ldots, f_m)$ of F and some expression of 1 as an element of $[f_1][f_2] \cdots [f_m]$, by recursively eliminating one edge from one circle-subgraph of Γ until Γ becomes a tree. Since $[f_1][f_2] = [f_2][f_1]$, the choice of enumeration can then be made arbitrary.

We now start the proof of Theorem 3.1.

[1] Rather than this, Juhász requires that G have no elements of order two, which seems to be insufficient.

3.4 Notation. For each $X \in \{A, B, C\}$, let $\widetilde{X}$ be a set given with a bijective maps of sets $X \to \widetilde{X}$, $x \mapsto \widetilde{x}$. Set $S := \widetilde{A} \mathbin{\dot\cup} \widetilde{B} \mathbin{\dot\cup} \{\widetilde{e}\}$ where $\widetilde{e}$ is a new symbol. Let Ω denote the graph with two vertices $\iota\widetilde{e}$ and $\tau\widetilde{e}$, called the *A-vertex* and the *B-vertex*, respectively, and edge set S, where the incidence functions of Ω may be depicted as $\iota\widetilde{e} \xrightarrow{a} \iota\widetilde{e}$ with $a \in \widetilde{A}$, $\tau\widetilde{e} \xrightarrow{b} \tau\widetilde{e}$ with $b \in \widetilde{B}$, and $\iota\widetilde{e} \xrightarrow{\widetilde{e}} \tau\widetilde{e}$; these are called the *$A$-edges*, the *$B$-edges*, and the *$\widetilde{e}$-edge* of Ω, respectively. By sending $\widetilde{e}$ to 1, we obtain a natural map of sets $S \to G$, and, hence, a morphism of monoids-with-involution $S^{\pm 1*} \to G$, and, hence, a map $\pi : \mathrm{Labels}(\Omega) \to G$. The natural maps $\widetilde{C}^{\,\pm 1*} \to \widetilde{A}^{\,\pm 1*}$ and $\widetilde{C}^{\,\pm 1*} \to \widetilde{B}^{\,\pm 1*}$ will be denoted by $c \mapsto c_A$ and $c \mapsto c_B$, respectively. The Ω-label of a closed Ω-path is said to be *G-reduced* if it cyclically contains no subword having any of the following forms: aa' with a, $a' \in \widetilde{A}^{\,\pm 1}$; bb' with b, $b' \in \widetilde{B}^{\,\pm 1}$; $\widetilde{e}^{\,-1}\widetilde{e}$; $\widetilde{e}\,\widetilde{e}^{\,-1}$; and $\widetilde{e}^{\,-1}c_A\widetilde{e}$ and $\widetilde{e}c_B\widetilde{e}^{\,-1}$ with $c \in \widetilde{C}^{\,\pm 1}$.

By an *Ω-graph* we mean a graph Γ given with a graph morphism $\Gamma \to \Omega$. We then speak of the *A-vertices*, the *B-vertices*, the *A-edges*, the *B-edges*, and the *$\widetilde{e}$-edges* of Γ. We have natural maps of sets $\mathrm{Labels}(\Gamma) \to \mathrm{Labels}(\Omega) \to G$, and will speak of the *Ω-label* and the *G-label* for any Γ-path. We shall sometimes use the notation $v \xrightarrow{g}\!\!\to w$ to refer to a Γ-path with G-label g, and similarly for Ω-labels and Γ-labels. For Γ-paths of length one, we may write $v \xrightarrow{g} w$. A closed Γ-path is said to be *G-reduced* if its Ω-label is G-reduced. A Γ-path is said to be *G-trivial* if its G-label is 1, and *G-nontrivial* otherwise.

In the remainder of this section, we shall be considering a specific finite, connected Ω-graph Γ engraved on $\mathbb{C}$. In this situation, for any face f for Γ, we say that f is *G-reduced* if one/each boundary path for f is G-reduced, and similar terminology applies with 'G-reduced' replaced with 'G-trivial' or 'G-nontrivial'. For $g \in G$, we say that f is a *g-face* if there exists some boundary path for f which has g as its G-label. We say that f is an *A-polygon* if all the vertices in its boundary are A-vertices, and therefore all the edges in its boundary are A-edges. Similar terminology applies with B in place of A. A circle-subgraph Δ of Γ will be said to be *G-reduced*, *G-trivial*, or *G-nontrivial* if one/each of the boundary paths of one/each face for Δ is G-reduced, G-trivial, or G-nontrivial, respectively.

3.5 Notation. Set $r_0 := g_0^6 \in G$. We may write $p_0 = \prod_{j=1}^{m}(h_j r_0^{\epsilon_j} h_j^{-1})$ where $m \geqslant 0$ and, for each $j \in \{1, 2, \ldots, m\}$, $h_j \in G$ and $\epsilon_j \in \{-1, 1\}$. We take m to be smallest possible. Since $p_0 \neq 1$, $m \geqslant 1$.

We now engrave an Ol'shanskii-van Kampen graph Γ on $\mathbb{C}$ that encodes the given data.

In $\mathrm{Labels}(\Omega)$, set $\widehat{g_0} := \prod_{i=1}^{n}(\widetilde{e}\,\widetilde{b_i}\,\widetilde{e}^{\,-1}\widetilde{a_i})$ and $\widehat{r_0} := \widehat{g_0}^{\,6}$. Recall the map $\pi : \mathrm{Labels}(\Omega) \to G$. For each $j \in \{1, \ldots, m\}$, firstly, choose a suitable $d_j \geqslant 0$ and some $\widehat{h_j} \in \widetilde{A}(\widetilde{e}\,\widetilde{B}\,\widetilde{e}^{\,-1}\widetilde{A})^{d_j} \subseteq \mathrm{Labels}(\Omega)$ such that $\pi(\widehat{h_j}) = h_j$, and, secondly, endow each of the line segments $[\frac{3j-3}{3m}, \frac{3j-2}{3m}]$, $[\frac{3j-2}{3m}, \frac{3j-1}{3m}]$, and $[\frac{3j-1}{3m}, \frac{3j}{3m}]$ with the structure of a line-segment Ω-graph such that the resulting path from left to right has the Ω-label $\widehat{h_j}$, $\widehat{r_0}^{\,\epsilon_j}$, and $\widehat{h_j}^{\,-1}$, respectively. The overall result is an Ω-graph engraved on $[0, 1]$. Let Γ_0 denote this graph.

It follows from our construction that no vertex of Γ_0 is incident to two $\widetilde{e}$-edges of Γ_0. As Γ_0 will contain all the vertices and all the $\widetilde{e}$-edges of the Ω-graph Γ on $\mathbb{C}$ that we are building, no vertex of Γ will be incident to two $\widetilde{e}$-edges of Γ.

For each $j \in \{1, 2, \ldots, m\}$, we now add to Γ_0 an A-edge with Ω-label $\widetilde{1_A}$, initial vertex $\frac{3j-1}{3m}$, and terminal vertex $\frac{3j-2}{3m}$, engraved as a semicircle in the lower half-plane, that is, the component of $\mathbb{C} - \mathbb{R}$ containing $-\mathbf{i}$. Recall that the Γ_0-path from left to right for the graph engraved on $[\frac{3j-2}{3m}, \frac{3j-1}{3m}]$ has Ω-label $\widehat{r_0}^{\,\epsilon_j}$. The overall result is an Ω-graph engraved on $\mathbb{C}$. Let Γ_1 denote this graph.

It can be seen that the number of faces for Γ_1 is $m+1$; for each $j \in \{1, 2, \ldots, m\}$, we have an $r_0^{\,-\epsilon_j}$-face whose boundary path at $\frac{3j-2}{3m}$ has Ω-label $(\widehat{r_0}^{\,\epsilon_j}\widetilde{1_A})^{-1}$, and we have a p_0-face whose boundary path at 0 has Ω-label

$$\widehat{h_1}\;\widehat{r_0}^{\,\epsilon_1}\;\widehat{h_1}^{\,-1}\;\widehat{h_2}\;\widehat{r_0}^{\,\epsilon_2}\;\widehat{h_2}^{\,-1}\;\cdots\;\widehat{h_m}\;\widehat{r_0}^{\,\epsilon_m}\;\widehat{h_m}^{\,-1}\;\widehat{h_m}\;\widetilde{1_A}\;\widehat{h_m}^{\,-1}\;\widehat{h_{m-1}}\;\widetilde{1_A}\;\widehat{h_{m-1}}^{\,-1}\;\cdots\;\widehat{h_1}\;\widetilde{1_A}\widehat{h_1}^{\,-1}.$$

We shall speak of *the* p_0-face, *the* $r_0^{\pm 1}$-faces, and, soon, *the* 1-faces. The boundary and the boundary paths for the p_0-face will be called the *p_0-boundary* and the *p_0-boundary paths*, respectively, and the edges of the p_0-boundary will be called the *p_0-edges*. Analogous terminology applies with $r_0^{\pm 1}$ in place of p_0. At the moment all edges are p_0-edges, but that will soon change. The $r_0^{\pm 1}$-boundaries are circle-graphs, and that will be maintained. The p_0-boundary is not a circle-graph, but soon it will be.

The next step is to engrave a larger Ω-graph Γ_2 on $\mathbb{C}$ by cutting corners off the $r_0^{\pm 1}$-faces and the p_0-face until they are all G-reduced, as follows.

If $v_1 \xrightarrow{\widetilde{a}^{\epsilon}} v_2 \xrightarrow{\widetilde{a'}^{\epsilon'}} v_3$ occurs in a p_0-boundary path, with a, $a' \in A$, we add a semicircle with path

$v_1 \xrightarrow{\widetilde{a^{\epsilon}a'^{\epsilon'}}} v_3$ across the p_0-face, creating a new p_0-face by cutting off a new A-edge and a new 1-face that is an A-trigon.

There is an analogous operation with B in place of A.

If some $v_1 \xrightarrow{\widetilde{e}^{-1}} v_2 \xrightarrow{c_A} v_3 \xrightarrow{\widetilde{e}} v_4$ or $v_1 \xrightarrow{\widetilde{e}^{-1}} v_2 \xrightarrow{(c_A)^{-1}} v_3 \xrightarrow{\widetilde{e}} v_4$ occurs in a p_0-boundary path, with $c \in \widetilde{C}$, we add a semicircle with path $v_1 \xrightarrow{c_B} v_4$ or $v_4 \xrightarrow{c_B} v_1$, respectively, across the p_0-face. Similarly, if $v_1 \xrightarrow{\widetilde{e}} v_2 \xrightarrow{c_B} v_3 \xrightarrow{\widetilde{e}^{-1}} v_4$ or $v_1 \xrightarrow{\widetilde{e}} v_2 \xrightarrow{(c_B)^{-1}} v_3 \xrightarrow{\widetilde{e}^{-1}} v_4$ occurs, we add $v_1 \xrightarrow{c_A} v_4$ or $v_4 \xrightarrow{c_A} v_1$, respectively, across the p_0-face. Each such corner cut off consists of an A- or B-edge and a 1-face called an *$\widetilde{e}$-tetragon*.

In the lower half-plane, we successively cut corners off the p_0-face until the resulting graph has a semicircle with path $1 \xrightarrow{\widetilde{1_A}} 0$. Thenceforth, the p_0-boundary will be a circle-graph.

Now, in the upper half-plane, we successively cut corners off the p_0-face until it is G-reduced. If the p_0-face is cut down to a digon, and then a monogon, the edges in these two steps cannot be represented as semicircles and other curves will be required. However, these steps arise if and only if the p_0-face is an A- or B-polygon, that is, $p_0 \in A \cup B$, and we shall see eventually that this does not happen, which implies some cases of the Freiheitssätze of Lyndon, Schupp, Collins and Perraud, Howie, Juhász, and others.

Turning now to the $r_0^{\pm1}$-faces, we cut off one A-trigon from each $r_0^{\pm1}$-face, and then the resulting $r_0^{\pm1}$-faces are G-reduced. This completes the construction of Γ_2.

The faces for Γ_2 are the p_0-face, the m $r_0^{\pm1}$-faces, and the 1-faces, each of which is an A- or B-trigon, -digon, or -monogon, or an $\widetilde{e}$-tetragon, and one trigon may have two vertices. We remark that ∞ lies in an A-trigon. The boundaries are all circle-subgraphs, except for a hypothetical two-vertex trigon.

The graph Γ we are constructing will be a subgraph of Γ_2. Consider any circle-subgraph Δ of Γ_2. By the *exterior* of Δ, we mean the face for Δ that contains the p_0-face for Γ_2; the other face for Δ is called the *interior* of Δ. All of this terminology will have similar meanings at all stages of the construction of Γ. We recursively search for a G-trivial A- or B-circle-subgraph Δ of (the current) Γ_2 whose interior is not a face for Γ_2, and, on finding such a Δ we eliminate from Γ_2 all the edges and vertices that lie in the interior of Δ, thus fusing some vertices, edges, and faces into a new 1-face for the next Γ_2. The new 1-face is then an A- or B-polygon. When all G-trivial A- or B-circle-subgraphs of Γ_2 are boundaries of 1-faces, we have completed the construction of Γ.

3.6 Remarks. It can be seen that the G-trivial A- or B-circle-subgraphs of Γ have pairwise disjoint edge-sets.

The faces for Γ are the p_0-face, any surviving $r_0^{\pm1}$-faces, and the 1-faces. Each 1-face is an A- or B-polygon or an $\widetilde{e}$-tetragon. It follows from (3.3.1) and the minimality of m that the number of $r_0^{\pm1}$-faces is m; that is, they all survive.

All faces have boundaries that are circle-graphs, except for a hypothetical two-vertex trigon.

For the remainder of proof of Theorem 3.1, we fix the following.

3.7 Notation. Let Δ be an arbitrary circle-subgraph of Γ. Let Υ be the subgraph of Γ obtained by deleting all those vertices and edges of Γ that lie in the exterior of Δ. It is easy to see that Υ is nonempty and connected. Now $\widehat{\mathbb{C}}$ is a CW-complex with the vertices and edges of Υ, and the faces for Υ.

Let $p_0' \in G$ be the G-label of some boundary path of the exterior of Δ. We shall call the exterior of Δ the *p_0'-face* for Υ, and call Δ the *p_0'-boundary.* The faces for Υ are then the p_0'-face and various 1-faces and $r_0^{\pm1}$-faces. Let m' denote the number of $r_0^{\pm1}$-faces for Υ. By (3.3.1), p_0' is a product of m' conjugates of $r_0^{\pm1}$; also by (3.3.1), p_0 is a product of a conjugate of p_0' and $m - m'$ conjugates of $r_0^{\pm1}$. Hence, m' is minimal, since m is.

3.8 Notation. Let X denote the CW-complex obtained from the CW-complex $\widehat{\mathbb{C}}$ by the following sequence of operations: all the A-edges and A-faces for Υ are collapsed to A-vertices, with the exception that if Δ is a G-trivial A-circle-graph, then the p_0'-face is left untouched even though its boundary is being collapsed to an A-vertex; then similarly, with B in place of A; at this point all edges are $\widetilde{e}$-edges joining A-vertices to B-vertices and all $\widetilde{e}$-tetragons have been collapsed to $\widetilde{e}$-digons, and now each $\widetilde{e}$-digon is to be collapsed to a single $\widetilde{e}$-edge. Notice that if Δ is a G-nontrivial A- or B-circle-graph, then the p_0'-face is to be collapsed to a vertex; this situation was already exceptional for allowing a two-vertex trigon, and eventually it will be an excluded case.

Now X has a CW-structure with vertices, edges, and faces. The one-skeleton $X^{(1)}$ is a graph on X. There are $m'+1$ or m' faces; these are m' $r_0^{\pm 1}$-faces, together with a p'_0-face if it has not been collapsed. All the X-faces have new attaching maps involving only their old $\widetilde{e}$-edges.

We now present the viewpoint that the cells of X are subsets of $\widehat{\mathbb{C}}$.

A *megavertex* M *for* Υ has as its one-skeleton $M^{(1)}$ a component of that graph which is obtained from Υ by deleting all the $\widetilde{e}$-edges, and to $M^{(1)}$ are added those faces for Υ whose boundaries lie in $M^{(1)}$, with the exception that if Δ is a G-trivial A- or B-circle-graph, the p'_0-face is *not* included in M.

Each $\widetilde{e}$-edge of Υ is incident to exactly two faces for Υ, each of which is an $\widetilde{e}$-tetragon, an $r_0^{\pm 1}$-face, or the p'_0-face, since all other faces are A- or B-polygons. By an $\widetilde{e}$*-band for* Υ we mean a component of that topological subspace of $\widehat{\mathbb{C}}$ which is the union of all the $\widetilde{e}$-tetragons for Υ and all the $\widetilde{e}$-edges of Υ. An $\widetilde{e}$-band may be an open annulus, with boundary paths having Ω-labels c_A and c_B for some $c \in \widetilde{C}^{\pm 1 *}$. The second possibility is that an $\widetilde{e}$-band may be an $\widetilde{e}$-edge that appears with its inverse in a p'_0-boundary path. The third and final possibility is that an $\widetilde{e}$-band may be an open disk plus two $\widetilde{e}$-edges, with a boundary path of the form $v_0 \xrightarrow{e_1} v_1 \xrightarrow{c_B} v_2 \xrightarrow{e_2^{-1}} v_3 \xrightarrow{c_A^{-1}} v_0$ for some $c \in \widetilde{C}^{\pm 1*}$, where, for $i = 1, 2$, e_i is an edge with an $\widetilde{e}$-label, and $e_i^{(-1)^i}$ appears in an $r_0^{\pm 1}$- or p'_0-boundary path. Since no vertex is incident to two $\widetilde{e}$-edges, the C-labelled subgraphs associated with annular, resp. disk, $\widetilde{e}$-bands are circle-subgraphs, resp. line-segment-subgraphs.

We may then view the X-vertices as the megavertices for Υ, the X-edges as the $\widetilde{e}$-bands for Υ, and the X-faces as the m' $r_0^{\pm 1}$-faces for Υ and, if it is not in a megavertex, the p'_0-face for Υ.

3.9 Notation. Create from X a CW-complex Y with a different one-skeleton by recursively making a degree-one vertex and its incident edge into interior points of the incident face until there are no degree-one vertices left, and then recursively making a degree-two vertex that is incident to two different edges into an interior point of a single edge, until either $Y^{(1)}$ has no degree-two vertices or $Y^{(1)}$ is a one-edge circle graph. By the Y*-length* of a Y-edge, we mean the number of X-edges it contains.

Let V, E, and F, denote the sets of Y-vertices, Y-edges, and Y-faces, respectively. Then F consists of m' $r_0^{\pm 1}$-faces and at most one p'_0-face. For each $f \in F$, the Y-boundary of f will be denoted ∂f.

Let F_r denote the set of m' $r_0^{\pm 1}$-faces in Y, and let E_r denote the set of Y-edges that are in Y-boundaries of elements of F_r. Set $F_p := F - F_r$, and let E_p denote the set of Y-edges that are in Y-boundaries of elements of F_p.

Set $E_{rp} := E_r \cap E_p$ and $E_{rr} := E_r - E_p$. Let F_{rr} denote the set of $r_0^{\pm 1}$-faces whose Y-boundaries have all edges in E_{rr}. Let $F_{rp} := F_r - F_{rr}$, the set of $r_0^{\pm 1}$-faces whose Y-boundaries have at least one Y-edge in E_p. For each $f_r \in F_r$, set $\deg_{rr}(f_r) := |E_{rr} \cap \partial f_r|$ and $\deg_{rp}(f_r) := |E_{rp} \cap \partial f_r|$. We define $\deg_{rr}(F_{rr}) := \sum\limits_{f_r \in F_{rr}} \deg_{rr}(f_r)$, and so on.

3.10 Lemma. *With the foregoing notation, the following hold.*

(*i*) *For all* f, $f' \in F_r$, *if* $f \neq f'$, *then each* Y*-edge of* $\partial f \cap \partial f'$ *has* Y*-length at most* $2n$.

(*ii*) $|V| - |E| + |F| = 2$.

(*iii*) *If* $|F_r| = 0$, *then* $p'_0 = 1$. *If* $|F_r| = 1$, *then* p'_0 *is a* G*-conjugate of* r_0 *or* r_0^{-1}.

(*iv*) *If* $|F_r| \geqslant 2$, *then* $6|F| - 12 \geqslant 2|E|$.

(*v*) *For each* $f \in F_r$, ∂f *is a circle-subgraph of* $Y^{(1)}$.

(*vi*) $6|F_{rr}| \leqslant \deg_{rr}(F_{rr}) \leqslant \deg_{rr}(F_r) = 2|E_{rr}|$.

(*vii*) *Every closed path in a megavertex for* Υ *is* G*-trivial.*

Proof. We argue by induction on the number of faces for Υ that lie in the interior of Δ.

If there is only one face in the interior of Δ, then that face is an A- or B-polygon, an $\widetilde{e}$-tetragon, or an $r_0^{\pm 1}$-face. An A- or B-polygon gets collapsed to a vertex, while the complementary p'_0-face, as a special case, remains an open disk, and then Y is a sphere and $Y^{(1)}$ is a vertex. An $\widetilde{e}$-tetragon gets reduced to a vertex, in stages, while the complementary p'_0-face remains an open disk, and then Y is a sphere and $Y^{(1)}$ is a vertex. For an $r_0^{\pm 1}$-face, the A- and B-edges get collapsed to vertices, and then Y is a sphere and $Y^{(1)}$ is a circle-graph with one edge of Y-length $12n$. In all these cases, (*i*)–(*vii*) hold.

Thus, we may assume that there are at least two faces for Υ in the interior of Δ, and that, for any circle-subgraph Δ' of Υ, if $\Delta' \neq \Delta$, then the analogues of (*i*)–(*vii*) with Δ' in place of Δ all hold.

Two $\widetilde{e}$-edges that appear in an Υ-boundary path of an $r_0^{\pm 1}$-face with only one intervening A- or B-edge cannot lie in some $\widetilde{e}$-band, by the induction hypothesis applied to (*vii*) and the fact that no letter of r_0 lies in C. Hence, the X-boundary path of each $r_0^{\pm 1}$-face has no backtrackings. Hence, each Y-boundary path of each $r_0^{\pm 1}$-face has Y-length $12n$.

(i). Suppose that (i) fails. In terms of Υ, this entails that we have all of the following data:

- $\gamma = (v_0, e_0, v_1, x_1, v_2, e_1, v_3, x_2, v_4, e_2, v_5, x_3, \ldots, x_{2n-1}, v_{4n-2}, e_{2n-1}, v_{4n-1}, x_{2n}, v_{4n}, e_{2n}, v_{4n+1})$ is an Υ-path which is the beginning of an Υ-boundary path of f. Each v_i is a vertex, each e_i has an $\widetilde{e}^{\,\pm 1}$-label, and each x_i has an $(\widetilde{A} \cup \widetilde{B})^{\pm 1}$-label. Let q_i denote the G-label of x_i, and set $g_1 := \prod_{i=1}^{2n} q_i$, which is the G-label of γ. Then g_1^6 is the G-label of the Υ-boundary path of f that begins with γ, and g_1 is a cyclic permutation of g_0 or g_0^{-1}.
- $\gamma' = (v'_0, e'_0, v'_1, x'_1, v'_2, e'_1, v'_3, x'_2, v'_4, e'_2, v'_5, x'_3, \ldots, x'_{2n-1}, v'_{4n-2}, e'_{2n-1}, v'_{4n-1}, x'_{2n}, v'_{4n}, e'_{2n}, v'_{4n+1})$ is an Υ-path which is the beginning of the inverse of an Υ-boundary path of f'. Each v'_i is a vertex, each e'_i has an $\widetilde{e}^{\,\pm 1}$-label, and each x'_i has an $(\widetilde{A} \cup \widetilde{B})^{\pm 1}$-label. Let q'_i denote the G-label of x'_i, and set $g'_1 := \prod_{i=1}^{2n} q'_i$, which is the G-label of γ'. Then $g_1'^{-6}$ is the G-label of the Υ-boundary path of f' that ends with γ'^{-1}, and g'_1 is a cyclic permutation of g_0 or g_0^{-1}.
- For each $i \in \{0, 1, 2, \ldots, 2n\}$, $v'_{2i} \xrightarrow{\gamma_{2i}} v_{2i} \xrightarrow{e_i} v_{2i+1} \xrightarrow{\gamma_{2i+1}} v_{2i+1} \xrightarrow{e_i'^{-1}} v'_{2i}$ is an Υ-boundary path of some $\widetilde{e}$-band f_i. (By the construction of Γ, $e_i \neq e'_i$.) For each $j \in \{0, 1, 2, \ldots, 4n+1\}$, let $c_j \in C$ denote the G-label of γ_j. Then $c_{2i} = c_{2i+1}$.

For each $i \in \{1, 2, \ldots, 2n\}$, the closed megavertex-path
$$v'_{2i-1} \xrightarrow{\gamma_{2i-1}} v_{2i-1} \xrightarrow{x_i} v_{2i} \xrightarrow{\gamma_{2i}^{-1}} v'_{2i} \xrightarrow{x_i'^{-1}} v'_{2i-1}$$
lies in the interior of Δ, and, hence, is G-trivial, by the induction hypothesis applied to (vii); thus, $c_{2i-1} q_i = q'_i c_{2i} = q'_i c_{2i+1}$ in G. Now
$$c_1 g_1 = c_1 \prod_{i=1}^{2n} q_i = (\prod_{i=1}^{2n} q'_i) c_{4n+1} = g'_1 c_{4n+1}.$$
By (3.1.1), $c_1 g_1^6 = g_1'^6 c_1 = g_1'^6 c_0$. We then eliminate all the $\widetilde{e}$-edges that lie in the band f_0. This fuses $f \cup f_0 \cup f'$ into a single face, which is a 1-face since one of its boundary paths has G-label $c_1 g_1^6 c_0^{-1} g_1'^{-6}$. In summary, we have a new graph engraved on $\widehat{\mathbb{C}}$, with $m' - 2$ $r_0^{\pm 1}$-faces, one p'_0-face, and various 1-faces. By (3.3.1), p'_0 is then a product of $m' - 2$ conjugates of $r_0^{\pm 1}$. This contradicts the minimality of m'. Hence, (i) holds.

(ii). It suffices to prove the analogous equality with X in place of Y, and for X it suffices to show that each megavertex for Υ has Euler characteristic 1, and that each $\widetilde{e}$-band for Υ has Euler characteristic -1.

Consider any A-megavertex M for Υ.

Consider the case where $\Delta \subseteq M$. Then Δ is an A-circle-subgraph of Γ. If Δ were G-trivial, then there would be only one face for Υ in the interior of Δ, by the construction of Γ, and we are assuming this is not the case. Hence Δ is G-nontrivial, and then the p'_0-face also lies in M, by the construction of X. If some edge of Δ lies in a G-trivial A-circle-subgraph, we could find a G-nontrivial A-circle-subgraph that is different from Δ. This contradicts the induction hypothesis applied to (vii). Hence, in all cases, no edge in M lies in the boundary of two faces of M.

Thus we could eliminate one edge from the boundary of each face, as well as the face, eventually transforming M into a connected graph, without changing the Euler characteristic. If this graph were not a tree, then it would contain an A-circle-subgraph Δ' which is not the boundary of a face of M. If Δ' is G-trivial, then, by the construction of Γ, Δ' is the boundary of an A-face, which is a face of M. If Δ' is G-nontrivial, then by the induction hypothesis applied to (vii), $\Delta' = \Delta$, and by the construction of X, the p'_0-face of X is a face of M. These are contradictions. Hence the graph that would remain would be a tree, and M must have Euler characteristic 1.

Consider any $\widetilde{e}$-band N for Υ.

Consider first the case where N is an annulus. At least one boundary component Δ' of N is not Δ. By the induction hypothesis applied to (vii), the C-circle-subgraph Δ' is G-trivial. By the definition of $\widetilde{e}$-tetragons, the other boundary component Δ'' of N is also G-trivial. By the construction of Γ, both Δ' and Δ'' are boundaries of C-polygons in the interior of Δ, and N does not exist.

It now remains to consider the case where N is not an annulus. Here, the number of edges in N is one more than the number of faces in N, and there are no vertices in N. Thus, the Euler characteristic of N is -1.

Since $\widehat{\mathbb{C}}$ has Euler characteristic 2, so too do X and Y.

(iii) follows from (3.3.1).

(iv). If $|F_r| \geqslant 2$, then, by (i), $Y^{(1)}$ is not a one-edge circle-graph. By the definition of Y, every Y-vertex has degree at least three. Thus, $3|V| \leqslant \deg(V) = 2|E|$. By ($ii$),

$$6|F| - 12 = 6|E| - 6|V| \geqslant 6|E| - 4|E| = 2|E|.$$

(v). Suppose that (v) fails. Then there exists some $r_0^{\pm 1}$-face f and some megavertex M for Υ that contains two A- or B-edges of the Υ-boundary of f. It follows that there exist two vertices v_1, v_2 in Υ, and three line-segment-subgraphs L, ∂_-, and ∂_+ in Υ all with endpoints v_1 and v_2, such that the following hold: $\partial_- \cap \partial_+ = \partial_- \cap L = \partial_+ \cap L = \{v_1, v_2\}$; $\partial_- \cup \partial_+$ is the Υ-boundary of f; $L \subseteq M$; and f lies in the interior of the circle-subgraph $\partial_+ \cup L$ and in the exterior of the circle-subgraph $\partial_- \cup L$.

One of the G-labels of the circle-subgraph $\partial_- \cup L$ is the product of an odd-length subword of a cyclic permutation of $r_0^{\pm 1}$ and an element of $A \cup B$. Hence, $\partial_- \cup L$ is G-nontrivial. By (3.3.1), $|F_r| \geqslant 2$. By (iv), $6|F| - 12 \geqslant 2|E|$.

It follows from the induction hypothesis applied to (v) that $\Delta = \partial_+ \cup L$ and that there is no other incidence of a megavertex for Υ containing two edges of the Υ-boundary of an $r_0^{\pm 1}$-face for Υ. Although (v) fails, no Y-edge and its inverse both occur in the Y-boundary path of an $r_0^{\pm 1}$-face of Y. It follows from (i) that $6|F_{rr}| \leqslant \deg_{rr}(F_{rr}) \leqslant 2|E_{rr}|$.

In Y, L is collapsed to a vertex, and f becomes a pinched annulus that engulfs the interior of $\partial_- \cup L$. Since ∂_+ contains at least one $\widetilde{e}$-edge, the p_0'-face does not get collapsed in X and Y. It follows that $|F_p| = 1$, $F_{rp} = \{f\}$, $E_r = E$, and E_{rp} consists of a single Y-edge obtained as a quotient of ∂_+. Hence, $|E_{rr}| = |E| - 1$ and $|F_{rr}| = |F| - 2$.

Now $2|E| \leqslant 6|F| - 12 = 6|F_{rr}| \leqslant 2|E_{rr}| = 2|E| - 2$, which is a contradiction. Hence, (v) holds.

(vi) follows from (i) and (v).

(vii). Suppose that (vii) fails. Any closed path contains either a backtracking or a circle-subpath. It follows that there exists a G-nontrivial A- or B-circle-subgraph Δ' in Υ. By the induction hypothesis applied to (vii), $\Delta' = \Delta$. Thus, $p_0' \in (A \cup B) - \{1\}$. By (iii), $|F_r| \geqslant 2$. By (iv), $6|F| - 12 \geqslant 2|E|$.

In the construction of X and Y, the p_0'-face and its Υ-boundary are collapsed to a vertex. Thus, in Y, $F_p = E_p = F_{rp} = \emptyset$, $F = F_{rr}$, and $E = E_{rr}$. By (vi), $6|F| \leqslant 2|E|$, which contradicts $6|F| - 12 \geqslant 2|E|$. Hence, (vii) holds. □

3.11 Notation. By Lemma 3.10(vii), Δ is not a G-nontrivial A- or B-circle-subgraph, and, hence, Y has a p_0'-face. The p_0'-face of Y will be denoted f_p.

For each $j \in \{1, 2, 3\}$, set $F(j) := \{f_r \in F_{rp} \mid \deg_{rp}(f_r) = 1 \text{ and } \deg_{rr}(f_r) = 4 - j\}$.

3.12 Corollary. *If $|F_r| \geqslant 2$ and ∂f_p is a circle-subgraph of $Y^{(1)}$, then $\sum_{j=1}^{3} j \cdot |F(j)| \geqslant 6$.*

Proof. By Lemma 3.10(vii), $|F_r| = |F| - 1$. Since ∂f_p is a circle-graph, $E_r = E$. Since $|F_r| \geqslant 2$, $6|F| - 2|E| \geqslant 12$, by Lemma 3.10(iv). Hence, $6|F_r| - 2|E_r| = 6|F| - 6 - 2|E| \geqslant 6$.

By Lemma 3.10(vi), $6|F_{rr}| \leqslant \deg_{rr}(F_{rr}) \leqslant \deg_{rr}(F_r) = 2|E_{rr}|$. Hence,

$$\deg_{rr}(F_{rp}) = \deg_{rr}(F_r) - \deg_{rr}(F_{rr}) \leqslant 2|E_{rr}| - 6|F_{rr}|.$$

Clearly, $\deg_{rp}(F_{rp}) = |E_{rp}|$. Now,

$$\begin{aligned} \sum_{f_r \in F_{rp}} \big(6 - 2\deg_{rp}(f_r) - \deg_{rr}(f_r)\big) &= 6|F_{rp}| - 2\deg_{rp}(F_{rp}) - \deg_{rr}(F_{rp}) \\ &\geqslant 6|F_{rp}| - 2|E_{rp}| - (2|E_{rr}| - 6|F_{rr}|) = 6|F_r| - 2|E_r| \geqslant 6. \end{aligned}$$

Consider any $f_r \in F_{rp}$. Since ∂f_p is a circle-graph, and ∂f_r is a circle-graph by Lemma 3.10(v), it can be seen that in ∂f_r the E_{rp}-edges are separated from each other by intervening E_{rr}-edges. Hence, either $\deg_{rp}(f_r) = 1$ or $\deg_{rr}(f_r) \geqslant \deg_{rp}(f_r)$ and $6 - 2\deg_{rp}(f_r) - \deg_{rr}(f_r) \leqslant 0$. The result then follows. □

Proof of Theorem 3.1. We now take Δ to be the p_0-boundary in Γ. Then $\Upsilon = \Gamma$, and we take $p_0' = p_0$. Recall that Δ is G-reduced, by the construction of Γ. It follows from Lemma 3.10(vii) that the X-boundary of f_p has no backtrackings; hence, the Y-length of a Y-boundary path of f_p equals the length of our cyclically reduced expression for p_0.

We argue by induction on m. If $m = 1$, then p_0 is a G-conjugate of r_0 or r_0^{-1}, as desired. Thus, we may assume that $m \geqslant 2$, and that the corresponding result holds for smaller m. There are two cases.

If ∂f_p is not a circle-subgraph of $Y^{(1)}$, then some megavertex M for Γ contains two edges of Δ, say M is an A-megavertex. By (3.3.1), there then exist $x_1, x_2, y_1, y_2 \in A$ and $p_1, p_2 \in G$ such that the following hold: $p_1x_1p_2x_2$ is a reduced expression of a cyclic permutation of our cyclically reduced expression for p_0; p_1y_1 is a product of n_1 conjugates of $r_0^{\pm 1}$; p_2y_2 is a product of n_2 conjugates of $r_0^{\pm 1}$; and $n_1 + n_2 = m$. It can be seen that p_1y_1 and p_2y_2 are nontrivial. By the induction hypothesis, there exist G-conjugates of p_1y_1 and p_2y_2 that have factorizations of a desired form, with distinguished factors that are subwords of $r_0^{\pm 1}$. As it is cyclically reduced, p_1x_1 then has a factorization of a desired form that employs all but one of the distinguished factors of the G-conjugate of p_1y_1, and in the case where the G-conjugate of p_1y_1 is $r_0^{\pm 1}$, p_1x_1 has a distinguished factor of length $10n{-}1$. If ℓ'_1, ℓ'_2, ℓ'_3 denote the numbers of the distinguished factors of p_1x_1 of lengths $6n{-}1$, $8n{-}1$, and $10n{-}1$, respectively, then $\ell'_1 + 2\ell'_2 + 3\ell'_3 \geqslant 3$. The analogous result holds for p_2x_2. Now $p_1x_1p_2x_2$ has been written in the desired form.

It remains to consider the case where ∂f_p is a circle-subgraph of $Y^{(1)}$. We have $|F_r| = m \geqslant 2$. Consider any $j \in \{1,2,3\}$ and any $f_r \in F(j)$. By Lemma 3.10(i), the E_{rr}-part of ∂f_r has Y-length at most $(4-j)(2n)$. Thus, the E_{rp}-edge of ∂f_r has Y-length at least $12n - \big((4-j)(2n)\big)$, that is, $(j+2)2n$. Hence, the G-label of the E_{rp}-edge of ∂f_r has length at least $(j+2)2n-1$, since then we are counting megavertices. This G-label is a cyclic subword of $r_0^{\pm 1}$, which then contains a subword of $r_0^{\pm 1}$ of length $(j+2)2n-1$, since this is less than $10n$. As in the proof of Lemma 3.10(i), up to left and right multiplication by elements of C, this subword of $r_0^{\pm 1}$ equals a subword of a cyclic permutation of our cyclically reduced expression for p_0, since the E_{rp}-edge of ∂f_r is an edge of ∂f_p. Now Corollary 3.12 gives a desired factorization of a cyclically reduced G-conjugate of p_0, since the Y-vertices of ∂f_p correspond to megavertices for Γ that give nontrivial letters, which can absorb elements of C. $\square$