Proceedings of the Edinburgh Mathematical Society (2000) **43**, 295–307 ©

EQUATIONS WITH TORSION-FREE COEFFICIENTS

ANDREW CLIFFORD[1] AND RICHARD Z. GOLDSTEIN[2]

[1] *Department of Mathematics and Statistics, The College of New Jersey, PO Box 7718, Ewing, NJ 08628-0718, USA*

[2] *Department of Mathematics and Statistics, State University of New York at Albany, Albany, NY 12222, USA*

(Received 6 April 1998)

Abstract In this paper we generalize techniques used by Klyachko and the authors to prove some tessellation results about S^2. These results are applied to prove the solvability of certain equations with torsion-free coefficients.

Keywords: equations over groups; relative pictures

AMS 1991 *Mathematics subject classification:* Primary 20E05

1. Introduction

Since Howie [**7**] introduced them, relative diagrams have been used by many authors to discuss equations over groups. It seems that most reasonably general results about tessellations of the two-sphere can be translated into results about group theory. Work following this scheme can be found in [**3, 4, 6, 8**].

We consider the dual notion of relative diagrams, relative pictures. In short, relative pictures are directed graphs embedded in S^2 whose corners are labelled. We require that each vertex has a neighbourhood that looks like one of two star graphs, which we call patterns. First, we will develop some machinery to help us analyse these graphs. Next, we use this machinery to get a handle on what types of regions must occur in these graphs. We then apply these tessellation results to prove the following group theoretic results.

Theorem 1.1. *Let*

$$e = \prod_{i=1}^{N} a_i t^{m_i}$$

be an equation over the group G so that

(1) *if e has a subword of the form ta_it^{-1} or $t^{-1}a_it$, then a_i has infinite order;*

(2) *if e has two distinct subwords ta_it^{-1} and ta_jt^{-1}, then no positive word in $\{a_i, a_j^{-1}\}$ represents the identity of G; and*

https://doi.org/10.1017/S0013091500020939 Published online by Cambridge University Press

(3) *if e has two distinct subwords $t^{-1}a_it$ and $t^{-1}a_jt$, then no positive word in $\{a_i, a_j^{-1}\}$ represents the identity of G.*

Then e is solvable over G.

Theorem 1.2. *Let G be a group and let e be an equation of the form:*

$$e = \prod_{j=1}^{p}\left[tk_{(j,1)}t^{-1}k_{(j,2)}tk_{(j,3)}t^{-1}k_{(j,4)}\dots tk_{(j,2n_j-1)}t^{-1}k_{(j,2n_j)}\left(\prod_{r=1}^{m_j} ty_{(j,r)}\right)\right],$$

where

(1) *$n_j \geqslant 0$ and $m_j \geqslant 1$ for $1 \leqslant j \leqslant p$;*

(2) *$k_{(j,i)}$ has infinite order for all i and j;*

(3) *if $j \neq j'$, then the identity of G cannot be represented as a positive word in $\{k_{(j,1)}, k_{(j',1)}^{-1}\}$ or as a positive word in $\{k_{(j,2n_j)}, k_{(j',2n_{j'})}^{-1}\}$.*

Then e is solvable over G.

2. Patterns and graphs

Let Γ be a connected directed graph embedded in the two-sphere S^2. A *region* of Γ is a component of $S^2 - \Gamma$. A region R of Γ will be called *consistent* if the edges in the boundary of R are consistently directed. (This is slightly different from the definition found in [**3**].) Let e be an edge of Γ. A *germ* of e is an equivalence class of 'small' subintervals of e, each of which share the same vertex of e. So each edge has exactly two germs.

Let v be a vertex of Γ. A *corner* of v is a component of $D - \Gamma$, where D is some suitably small neighbourhood of v. If $\deg(v) = d$, then v has d corners, each of which is between two adjacent germs of v. If the two germs adjacent to the corner k of v are directed away from v, then k will be called a *source* corner; if they are directed toward v, then k will be called a *sink* corner. Otherwise, k will be a *neutral* corner.

A *pattern* P is a finite directed tree embedded in S^2 with a specific vertex v, called the *centre*, that is adjacent to every other vertex. (P is a directed star graph.) We will refer to corners of v as corners of P. We will call an edge of P *significant* if it is adjacent to a neutral corner.

If P is a pattern, we obtain the inverse pattern $\bar{P}$ by reversing the orientation on each edge of P and reflecting in S^2. Let E and $\bar{E}$ be the sets of edges of P and $\bar{P}$, respectively. For each edge σ of E, the corresponding edge of $\bar{P}$ will be called $\bar{\sigma}$. Similarly, if k is a corner of P, then the corresponding corner of $\bar{P}$ will be called $\bar{k}$.

Let P be a pattern. A finite, connected directed graph Γ embedded in S^2 is called a *P-graph* if each vertex v of Γ looks like the centre of P (in which case we call v a positive vertex) or of $\bar{P}$ (whence v is a negative vertex). More precisely, for each v of Γ, there is an orientation- and direction-preserving isomorphism of embedded graphs,

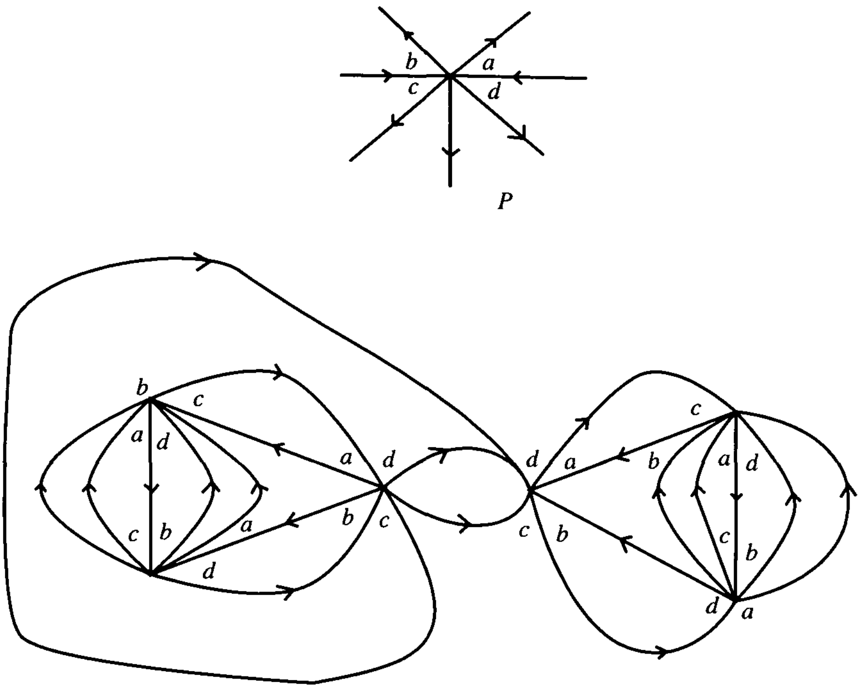

Figure 1. A pattern P and a P-graph.

ϕ_v, from a neighbourhood of v onto a neighbourhood of the centre of P or of $\bar{P}$. The map ϕ_v defines a correspondence between the germs of Γ at v and the edges of P or $\bar{P}$, and a correspondence between the corners of v and the corners of P or $\bar{P}$. We use these correspondences to label the germs of edges of Γ with the edges of P and $\bar{P}$, and the corners of Γ with the corners of P and $\bar{P}$.

Let R be a consistent region of the P-graph Γ. We will say that R is a *type 1* region if there is some corner k of P so that every corner of R at a positive vertex is labelled k and every corner of R at a negative vertex is labelled $\bar{k}$. Similarly, we will say that R is a *type 2* region if there is some corner k of P and some corner $\bar{l}$ of $\bar{P}$ so that every corner of R at a positive vertex is labelled k and every corner of R at a negative vertex is labelled $\bar{l}$.

We say that the pattern P is *type 1* if every P-graph has at least two type 1 regions; we say that P is *type 2* if every P-graph has at least two type 2 regions. In [**3**], we defined the class of type K patterns. The main result of that paper is that every type K pattern is type 1. In Figure 1, we give an example of a pattern P and a P-graph, Γ. (Here, we labelled some of the corners of P to depict the correspondence between corners of Γ and corners of P and $\bar{P}$.) Evidently, Γ has no type 1 regions. So, we cannot expect that every pattern is type 1.

In this paper, we will prove that every pattern whose centre is not a source or a sink is type 2. To this end, will use the following characterization of type 2 regions.

https://doi.org/10.1017/S0013091500020939 Published online by Cambridge University Press

Lemma 2.1. *Let R be a consistent region of the P-graph Γ and let σ and τ be edges of P so that as one travels around the boundary of R in some direction, one leaves a positive vertex on the germ labelled σ and a negative vertex on the germ labelled $\bar{\tau}$. Then R is a type 2 region.*

Proof. Assume we are given R, σ and τ, as in the statement of the lemma. We need to show that R is a type 2 region. Without loss of generality, assume that σ and τ are realized, as in the statement of the lemma, by travelling clockwise around R.

Since R is a consistent region, the corner that is immediately clockwise of the edge σ on P is neutral. Call this corner k. We see that every corner of R at a positive vertex is labelled k. Similarly, let $\bar{l}$ be the corner immediately clockwise of $\bar{\tau}$ on $\bar{P}$. Then every corner of R at a negative vertex is labelled $\bar{l}$. So R is a type 2 region. □

We see that we will want to find regions that are bounded by a certain type of simple closed path on Γ. Let σ and τ be edges of P. A $(\sigma, \bar{\tau})$-path on Γ is a simple closed path C, so that as one travels along it in some direction, one leaves positive vertices on the germ labelled σ and negative vertices on the germ labelled $\bar{\tau}$. In what follows, $(\sigma, \bar{\tau})$-paths will usually be consistently oriented. This may or may not coincide with the direction C is travelled to make it a $(\sigma, \bar{\tau})$-path. If this latter direction is clockwise (respectively, counterclockwise) with respect to the disc D bounded by C, we say the pair (C, D) is clockwise (respectively, counterclockwise).

If one of σ and τ is directed toward the centre of P and the other is directed away from the centre of P, we call the pair $(\sigma, \bar{\tau})$ *forcing*. The reason for this terminology will become evident in the following lemma, which we call the Forcing Lemma.

Lemma 2.2 (Forcing Lemma). *Let Γ be a P-graph. Assume that C is a simple closed path on Γ and that D is one of the discs that it bounds. Assume that at each positive vertex on C, the germ labelled σ lies in the closure of D, and that at each negative vertex of C the germ labelled $\bar{\tau}$ lies in the closure of D. Furthermore, assume that there is at least one vertex v on C so that the appropriate germ at v lies in the interior of D. Then, if $(\sigma, \bar{\tau})$ is forcing, there is a $(\sigma, \bar{\tau})$-path that bounds a disc D' that is strictly contained in D.*

Proof. Each positive vertex has exactly one germ labelled σ and each negative vertex has exactly one germ labelled $\bar{\tau}$. Furthermore, the assumptions on σ and τ imply that if an edge of Γ has one germ labelled σ or $\bar{\tau}$, the other germ is not labelled with either of these two labels. We can find the $(\sigma, \bar{\tau})$-path by leaving the vertex v on the appropriate germ. At the next vertex we reach, the germ labelled either σ or $\bar{\tau}$ is not the germ on which we arrived. So we may continue our path by leaving that vertex on the appropriate germ. Since Γ is finite, this process will eventually lead us to a vertex that we have already visited. This completes a closed path, which is a $(\sigma, \bar{\tau})$-path.

The assumptions on the germs of vertices of C assure us that whenever we leave from a vertex on the boundary of D, we do not leave D. Since we entered the interior of D, the resulting $(\sigma, \bar{\tau})$-path bounds a disc D' that is strictly contained in D. □

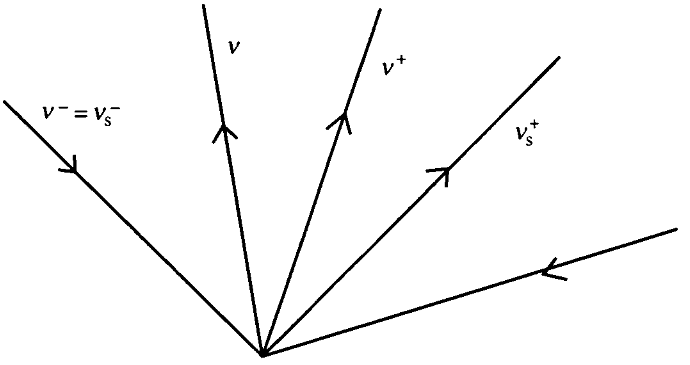

Figure 2. Significant edges.

Note that this lemma remains valid if we assume that the path C is trivial. In this case, the lemma just asserts that there is some $(\sigma, \bar{\tau})$-path on Γ.

3. First tessellation result

We shall call a pattern *stable* if its centre is either a source or a sink. In this paper, we will assume that all patterns are not stable. So, every pattern will have at least two neutral corners and at least two significant edges.

Let P be such a pattern. If ν is an edge of P, we will refer to the edges next to it as ν^- and ν^+, where ν^- is immediately counterclockwise of ν and ν^+ is immediately clockwise of ν. Moreover, let ν_s^- be the first significant edge one reaches by going counterclockwise from ν around the centre of P. Similarly, let ν_s^+ be the first significant edge one reaches by going clockwise from ν (see Figure 2).

Note that on $\bar{P}$, $\bar{\nu}_s^-$ (respectively, $\bar{\nu}_s^+$) is the first significant edge clockwise (respectively, counterclockwise) of $\bar{\nu}$. The following observation is stated without proof.

Lemma 3.1. *Let Γ be a P-graph. Let C be a simple closed path that is consistently directed bounding a disc D to the right of the direction of C.*

Assume v is a positive vertex on C so that C involves the germs labelled μ directed toward v and ν directed away from v. Then, if μ_s^- is directed toward v, it lies in the interior of D, otherwise it lies in the closure of D. Similarly, if ν_s^+ is directed away from v it lies interior to v, otherwise it lies in the closure of D.

Assume v is a negative vertex on C so that C involves the germs labelled $\bar{\mu}$ directed toward v and $\bar{\nu}$ directed away from v. Then, if $\bar{\mu}_s^+$ is directed toward v, it lies in the interior of D, otherwise it lies in the closure of D. Similarly, if $\bar{\nu}_s^-$ is directed away from v it lies interior to v, otherwise it lies in the closure of D.

Theorem 3.2. *Let P be a non-stable pattern. Then P is type 2.*

Proof. Let P be a pattern and let Γ be a P-graph. Let E be the set of edges of P.

Consider the set A of forcing pairs $(\sigma, \bar{\tau}) \in E \times \bar{E}$. Since P is not stable, A is not empty. We define an *A-path* to be a pair (C, D), where C is a $(\sigma, \bar{\tau})$-path for some $(\sigma, \bar{\tau}) \in A$ and D is a topological disc it bounds.

https://doi.org/10.1017/S0013091500020939 Published online by Cambridge University Press

We define a partial order $<$ on the set of A-paths by $(C, D) < (C', D')$ if D is strictly contained in D'. By the Forcing Lemma, we know there is at least one $(\sigma, \bar{\tau})$-path C on Γ for some $(\sigma, \bar{\tau}) \in A$. Now, C bounds two discs D and D'. The A-paths (C, D) and (C, D') are not comparable under $<$. It follows that Γ must have at least two minimal A-paths.

From Lemma 2.1, it suffices to show that if (C, D) is a minimal A-path, then D is a region of Γ. We do this contrapositively. Moreover, we need only prove this in the case where (C, D) is clockwise as the counterclockwise case follows from this by a switch of ambient orientation.

Let (C, D) be an A-path with C a $(\sigma, \bar{\tau})$-path for some $(\sigma, \bar{\tau}) \in A$ and D not a region. We may assume that at each positive vertex v of C, the corner in D adjacent to the germ labelled σ is neutral. If not, then the germ labelled σ^+ is directed in the same direction as σ, so $(\sigma^+, \bar{\tau})$ is a forcing pair. From Lemma 3.1, the germ at each positive vertex labelled σ^+ lies interior to D. From the Forcing Lemma, there is a $(\sigma^+, \bar{\tau})$-path C' bounding a disc D' with D' strictly contained in D. So (C, D) is not minimal.

Similarly, we may assume that at each negative vertex v of C, the corner in D adjacent to the germ labelled $\bar{\tau}$ is neutral. It follows that $(\sigma^+, \bar{\tau}^-)$ is a forcing pair. Furthermore, at each positive vertex of C, the germ labelled σ^+ lies in the closure of D. Similarly, at each negative vertex of C, the germ labelled $\bar{\tau}^-$ is in the closure of D. Since D is not a region, one of these germs must lie in the interior of D. It follows from the Forcing Lemma that (C, D) is not minimal.

This ends the proof. $\square$

4. Representing patterns with words

In [**3**], we defined the class of type K patterns. These were non-stable patterns that had no sink corners, so that the collection of source corners were contiguous as one circled the centre. For the next two sections, we wish to generalize this to patterns that we call *pre-stable*. A pre-stable pattern is a non-stable pattern with no sink corners. Such a pattern will be called *non-singular* if it has at least one source corner. In [**3**], we showed that each type K pattern was a type 1 pattern. Figure 1 shows a pre-stable pattern that is not type 1. From the previous section, we know that every non-stable pattern is type 2. This includes pre-stable patterns. However, there is more we can say. In the next section, we will show that if P is non-singular and pre-stable and Γ is a P graph, then there are two regions R_1 and R_2 of Γ so that each R_ι is either type 1 or type 2, where the corner labels of R_ι come out of a restricted set of neutral corners of P.

In order to discuss the patterns that we are interested in, we will set up a correspondence between patterns and elements of the free monoid $F[t, t^{-1}]$. We associate to the word $e \in F[t, t^{-1}]$ the pattern P_e having centre v. The pattern P_e has an edge directed away from v for each occurrence of t in e and an edge directed toward v for each occurrence of t^{-1} in e, so that if one circles around v in a counterclockwise direction reading t for every edge leaving v and t^{-1} for every edge entering v, one will read a cyclic conjugate of the word e. This correspondence is well defined up to cyclic conjugacy of elements of

$F[t, t^{-1}]$ and ambient isotopies of patterns. Obtain the formal inverse, e^{-1}, from e, by writing e backwards and exchanging every positive power of t for a negative power and vice versa. It follows that $P_{e^{-1}} = \bar{P}_e$.

In fact, we will be interested in keeping track of the neutral corners of P_e. So, we will adapt the above correspondence to obtain for each pattern P a word

$$e_P = \prod_{i=1}^{N} (k_i) t^{m_i}$$

in the free group $F[X \cup \{t\}]$. Here, $X = \{k_i\}$ is a set of corners of P that contains all neutral corners. (It should be noted that P_{e_P} is the pattern dual to the relative cell corresponding to e as described in [**1**].) For example, if $e = at^2bt^{-1}ct^3dt^{-1}$, then P_e is shown in Figure 1.

It should be clear that a pattern P is pre-stable exactly when the corresponding word e_P, when considered cyclically, does not have two consecutive occurrences of t^{-1}. It is also non-singular if it has consecutive occurrences of t.

We will now describe the class of non-singular pre-stable patterns by describing their associated words.

An *alternating piece* is a word of the form

$$\alpha_j = tk_{(j,1)}t^{-1}k_{(j,2)}tk_{(j,3)}t^{-1}k_{(j,4)} \dots tk_{(j,2n_j-1)}t^{-1}k_{(j,2n_j)}t,$$

where $n_j \geqslant 1$. Then, the class of words in question are of the form

$$e = \prod_{j=1}^{p} \alpha_j t^{m_j},$$

where α_j is an alternating piece and $m_j \geqslant 0$ for each j. Let P_e be the corresponding pattern with the set of neutral corners given by $X = \{k_{(j,i)} \mid 1 \leqslant j \leqslant p;\ 1 \leqslant i \leqslant 2n_j\}$. We realize that the set of corners of P_e is being used to define P_e. This is only for notational convenience. Since the exponent sum of e is positive and the occurrences of t^{-1} in e are isolated by positive occurrences of t, P_e is non-singular and pre-stable. Furthermore, every non-singular pre-stable pattern or its inverse can be obtained in this manner.

5. Second tessellation result

Theorem 5.1. *Let P be a non-singular pre-stable pattern with associated word*

$$e_P = \prod_{j=1}^{p} \alpha_j t^{m_j}$$

as described above. Let Γ be a P graph. Then there are two consistent regions R_1 and R_2 of Γ so that for each $\iota = 1, 2$ there exist i and j satisfying one of the following conditions:

(1) *each corner of R_ι is positive and labelled by the corner $k_{(j,i)}$ of P; or*

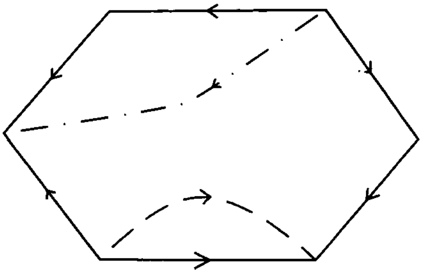

Figure 3.

(2) *each corner of* R_ι *is negative and labelled by the corner* $\overline{k}_{(j,i)}$ *of* $\bar{P}$*; or*

(3) *each positive corner of* R_ι *is labelled* $k_{(j,1)}$ *and each negative corner is labelled* $\overline{k}_{(i,1)}$*; or*

(4) *each positive corner of* R_ι *is labelled* $k_{(j,2n_j)}$ *and each negative corner is labelled* $\overline{k}_{(i,2n_i)}$.

Proof. Let P and e_P be as described in the statement of the theorem, and let E be the set of edges of P.

Let Γ be a P-graph. As described previously, the correspondence of vertices of Γ with the patterns P_e and $\bar{P}_e$ induces a labelling of the germs of Γ by elements of $E \cup \bar{E}$; and a labelling of the neutral corners of Γ by elements of $X \cup \bar{X}$.

As in [**3**], we need to add a collection of dotted edges to Γ. Let R be a region that is not consistent. We pair the corners of R that are not neutral, so that a dotted edge runs from each source corner to the sink corner to which it has been paired. We do this in such a way as to keep these added dotted edges from intersecting. Call the resulting graph $\hat{\Gamma}$ (see Figure 3).

Every source corner of Γ is at a positive vertex; every sink corner of Γ is at a negative vertex. Every source corner of Γ received the initial germ of a dotted edge, and each sink corner received the terminal germ of a dotted edge. So, each added dotted edge runs from a positive vertex to a negative vertex.

If we create $\hat{P}$ by adding a dotted edge to each source corner of P_e and obtain $\bar{\hat{P}}$ from $\hat{P}$ as before, then $\hat{\Gamma}$ is a $\hat{P}$-graph. To complete the story, if we let the variable s represent the new dotted edges of $\hat{P}$ in the same way that t represented the old solid edges of P_e, then

$$e_{\hat{P}} = \prod_{j=1}^{p} \alpha_j s(ts)^{m_j}.$$

We will call this word $\hat{e}$.

We extend the definitions used previously to deal with graphs in which some of the edges are dotted.

Let σ and τ be edges of $\hat{P}$. We will call the pair $(\sigma, \bar{\tau})$ *forcing* if either one of σ and τ is directed toward the centre of P and the other is directed away from the centre; or one of these edges is solid and the other is dotted. A $(\sigma, \bar{\tau})$-path will be defined as previously.

Let $(\sigma, \bar{\tau})$ be a forcing pair. As before, if an edge of Γ has one germ labelled either σ or $\bar{\tau}$, then the other germ is not labelled either σ or $\bar{\tau}$. Because of this, the Forcing Lemma holds for $\hat{\Gamma}$.

As above, we wish to consider a family of paths on $\hat{\Gamma}$. To this end, we will give specific names to some of the edges of $\hat{P}$.

Consider the subgraph of $\hat{P}$ corresponding to the subword $w_j = \alpha_j s$ of $\hat{e}$. For each $j = 1, \ldots, p$, let $r_j = 2n_j + 1$ be the number of occurrences of t and t^{-1} in w_j. These occurrences correspond to r_j edges of $\hat{P}$, which alternate direction (relative to the central vertex v). We refer to the edge corresponding to the ith occurrence of t in α_j as the ordered pair $[j, i]$. Finally, we refer to the dotted edge of $\hat{P}$ corresponding to the sole occurrence of s in w_j as s_j.

We now have names for every significant solid edge of P. Also, there are p dotted edges, which have names. The dotted edge s_j is next to the solid edge $[j, r_j]$ for each $1 \leqslant j \leqslant p$. Furthermore, the edge $[j, i]$ is directed toward v if i is even and away from v if i is odd.

Consider the following six sets of forcing pairs:

(i) $A_1 = \{([a, 2], \overline{[b, 1]}) \mid 1 \leqslant a,\ b \leqslant p\}$;

(ii) $A_2 = \{([a, 1], \overline{[b, 2]}) \mid 1 \leqslant a,\ b \leqslant p\}$;

(iii) $A_3 = \{([a, r_a], \overline{[b, r_b - 1]}) \mid 1 \leqslant a,\ b \leqslant p\}$;

(iv) $A_4 = \{([a, r_a - 1], \overline{[b, r_b]}) \mid 1 \leqslant k,\ l \leqslant p\}$;

(v) $A_5 = \{([a, u], \bar{s}_b) \mid 1 \leqslant a, b \leqslant p;\ 2 \leqslant u \leqslant r_a - 1\}$;

(vi) $A_6 = \{(s_a, \overline{[b, u]}) \mid 1 \leqslant a, b \leqslant p;\ 2 \leqslant u \leqslant r_b - 1\}$;

and let $A = A_1 \cup A_2 \cup A_3 \cup A_4 \cup A_5 \cup A_6$.

As before, an ordered pair (C, D) is an A-path if C is a $(\sigma, \bar{\tau})$-path for some $(\sigma, \bar{\tau}) \in A$ and C bounds D; and $(C, D) < (C', D')$ if D' properly contains D.

We note that if (C, D) is an A-path and D is a region, then C does not contain any dotted edges. This is because every dotted edge connects a positive vertex to a negative vertex and if $(\sigma, \bar{\tau}) \in A$ and one of σ or $\bar{\tau}$ is dotted then the other is not adjacent (on $\hat{P}$ or $\bar{\hat{P}}$) to a dotted edge. Let us run through the possibilities of labels on the corners of D.

First, assume (C, D) is counterclockwise with C a $(\sigma, \bar{\tau})$-path.

If $(\sigma, \bar{\tau}) \in A_1$, then there exists some j so that each corner of D is labelled $k_{(j,2)}$.

If $(\sigma, \bar{\tau}) \in A_2$, then there exist some i and j so that each corner of D is labelled $k_{(j,1)}$ or $\overline{k_{(i,1)}}$.

If $(\sigma, \bar{\tau}) \in A_3$, then there exist some i and j so that each corner of D is labelled $k_{(j,2n_j)}$ or $\overline{k_{(i,2n_i)}}$.

If $(\sigma, \bar{\tau}) \in A_4$, then there exists some j so that each corner of D is labelled $\overline{k_{(j,2n_j-1)}}$.

If $(\sigma, \bar{\tau}) \in A_5$, then there exist some i and j so that each corner of D is labelled $k_{(j,i)}$.

If $(\sigma, \bar{\tau}) \in A_6$, then there exist some i and j so that each corner of D is labelled $\overline{k_{(j,i)}}$.

Next, assume that (C, D) is clockwise with C a $(\sigma, \bar{\tau})$-path.

If $(\sigma, \bar{\tau}) \in A_1$, then there exist some i and j so that each corner of D is labelled $k_{(j,1)}$ or $\overline{k_{(i,1)}}$.

If $(\sigma, \bar{\tau}) \in A_2$, then there exists some j so that each corner of D is labelled $\overline{k_{(j,2)}}$.

If $(\sigma, \bar{\tau}) \in A_3$, then there exists some j so that each corner of D is labelled $k_{(j,2n_j-1)}$.

If $(\sigma, \bar{\tau}) \in A_4$, then there exist some i and j so that each corner of D is labelled $k_{(j,2n_j)}$ or $\overline{k_{(i,2n_i)}}$.

If $(\sigma, \bar{\tau}) \in A_5$, then there exist some i and j so that each corner of D is labelled $k_{(j,i)}$.

If $(\sigma, \bar{\tau}) \in A_6$, then there exist some i and j so that each corner of D is labelled $\overline{k_{(j,i)}}$.

In short, we see that if (C, D) is an A-path with D a region, then the labels on the corners of D satisfy the conclusions of the theorem for $D = R_\iota$. So, it suffices to show that if (C, D) is an A-path and D is not a region, then (C, D) is not minimal.

Let (C, D) be an A-path in which D is not a region. So C is a $(\sigma, \bar{\tau})$-path for some $(\sigma, \bar{\tau}) \in A$. We will prove the case where C is counterclockwise with respect to D. The clockwise case is proved similarly. We do this in cases as follows.

Case 5.2. $(\sigma, \bar{\tau}) = ([a, 2], \overline{[b, 1]}) \in A_1$ for some a and b. Here, as we traverse C, we are travelling against the directions of the arrows on C. So each germ we arrive on is directed away from the corresponding vertex. Consider the edge $\overline{[b, 1]}$ on the pattern $\hat{\bar{P}}$. If we travel around the centre in a counterclockwise direction, we pass the dotted edge $\overline{s_{b-1}}$ before we reach a solid edge that is directed away from the centre. (Here $b-1$ is taken modulo p.)

If there is a negative vertex v on C, the germ at v that is labelled $\overline{s_{b-1}}$ enters the interior of D. Now, the forcing pair $([a, 2], \overline{s_{b-1}}) \in A_5 \subset A$. It follows from the Forcing Lemma that (C, D) is not minimal.

If there is no negative vertex on C, then at each positive vertex, the germ labelled $[a, 3]$ is in the closure of D, and at least one of these germs lies in the interior of D. If $r_a = 3$, the Forcing Lemma assures us that there is a $([a, 3], \overline{[a, r_a - 1]})$-path in D. If $r_a \neq 3$, the Forcing Lemma assures us that there is a $([a, 3], \overline{s_a})$-path in D. In either case, (C, D) is not minimal.

Case 5.3. $(\sigma, \bar{\tau}) = ([a, 1], \overline{[b, 2]}) \in A_2$. Here, at each positive vertex the germ labelled $[a, 2]$ lies in the closure of D, and at each negative vertex the germ labelled $\overline{[b, 1]}$ lies in the closure of D. Since D is not a region, at least one of these germs enters the interior of D. Now, $([a, 2], \overline{[b, 1]}) \in A_1$. So, (C, D) is not minimal.

https://doi.org/10.1017/S0013091500020939 Published online by Cambridge University Press

Case 5.4. $(\sigma, \bar{\tau}) = ([a, r_a], \overline{[b, r_b - 1]}) \in A_1$ for some a and b. As in Case 5.2, as we traverse C we are travelling against the direction of the arrows on C. Consider the edge $[a, r_a]$ on the pattern $\hat{P}$. If we travel around the centre in a counterclockwise direction, we pass the dotted edge s_a before we reach a solid edge that is directed toward the centre.

If there is a positive vertex v on C, the germ at v, which is labelled s_a, enters the interior of D. Now, the forcing pair $(s_a, \overline{[b, r_b - 1]}) \in A_6 \subset A$. It follows from the Forcing Lemma that (C, D) is not minimal.

If there is no positive vertex on C, then at each negative vertex, the germ labelled $\overline{[b, r_b - 2]}$ is in the closure of D and at least one of these germs lies in the interior of D. If $r_b - 2 = 1$, the Forcing Lemma assures us that there is a $([b, 2], \overline{[b, r_b - 2]})$-path in D. If $r_b - 2 \neq 1$, the Forcing Lemma assures us that there is a $(s_b, \overline{[b, r_b - 2]})$-path in D. In either case (C, D) is not minimal.

Case 5.5. $(\sigma, \bar{\tau}) = ([a, r_a - 1], \overline{[b, r_b]}) \in A_4$. Here, at each positive vertex the germ labelled $[a, r_a]$ lies in the closure of D, and at each negative vertex the germ labelled $\overline{[b, r_b - 1]}$ lies in the closure of D. Since D is not a region, at least one of these germs enters the interior of D. Now, $([a, 2], \overline{[b, 1]}) \in A_3$. So, (C, D) is not minimal.

Case 5.6. $(\sigma, \bar{\tau}) = ([a, u], \overline{s_b}) \in A_5$ for some $2 \leqslant u \leqslant r_a - 1$. Here, at each positive vertex, the germ labelled $[a, u + 1]$ is in the closure of D, and at each negative vertex, the germ labelled $[b, r_b]$ lies in the closure of D.

If $u+1 \leqslant r_a - 1$, then the forcing pair $([a, u + 1], \overline{s_b}) \in A_5$, and there is a $([a, u + 1], \overline{s_b})$-path that bounds a disc strictly contained in D. So (C, D) is not minimal.

If $u+1 = r_a$, then $u = r_a - 1$ and $([a, u], \overline{[b, r_b]}) \in A_4$, and there is a $([a, u], \overline{[b, r_b]})$-path that bounds a disc strictly contained in D. So (C, D) is not minimal.

Case 5.7. $(\sigma, \bar{\tau}) = (s_a, \overline{[b, u]}) \in A_6$ for some $2 \leqslant u \leqslant r_b - 1$. Here, at each negative vertex, the germ labelled $\overline{[b, u - 1]}$ is in the closure of D. If $u - 1 \geqslant 2$, then the forcing pair $(s_a, \overline{[b, u - 1]}) \in A_6$, and there is a $(s_a, \overline{[b, u - 1]})$-path that bounds a disc strictly contained in D. So (C, D) is not minimal.

If $u - 1 = 1$, then $u = 2$. So at each negative vertex, the germ labelled $\overline{[b, u]}$ is directed away from the negative vertex. This means that one reaches positive vertices on germs directed toward them. So, as in Cases 5.2 and 5.4, the germ at a positive vertex of C labelled $[a + 1, 1]$ lies in the interior of D. (Here $a + 1$ is taken modulo p.) Here $([a + 1, 1], \overline{[b, u]}) \in A_2$ and there is a $([a + 1, 1], \overline{[b, u]})$-path that bounds a disc strictly contained in D. So, (C, D) is not minimal.

There are no other cases.

This ends the proof. □

6. Equations over groups

In this final section, we prove Theorems 1.1 and 1.2 stated in § 1.

Let

$$e = \prod_{i=1}^{N} a_i t^{m_i}$$

https://doi.org/10.1017/S0013091500020939 Published online by Cambridge University Press

be an equation in the variable t, whose coefficient set $a_1, \ldots, a_N$ is a subset of the group G. We say that e is solvable over G if the inclusion induced homomorphism $G \to \langle G, t \mid e \rangle$ is injective. See [**2**] and [**5**] for descriptions of equations over groups.

Proof of Theorem 1.1. Let e be an equation over the group G, as described in the statement of Theorem 1.1. We may assume that e has non-negative exponent sum. If every occurrence of t is positive, then the result follows from work of Levin [**9**]. So, we will assume e has positive and negative occurrences of t, whence P_e is not stable.

Assume that e is not solvable over G. It is a consequence of the work of Howie that there exists a P_e graph Γ with a specified region Δ_0, so that the word in G obtained by reading around any region $\Delta \neq \Delta_0$ is a relation in G, while the word obtained by reading around Δ_0 is not a relation of G. Such a graph is called a relative picture over $\langle G, t \mid e \rangle$ and is discussed more fully in [**1**]. We may assume that Γ is reduced in the sense of Seradski [**10**]. In particular, this means that there is no edge of Γ connecting a positive and negative vertex with the germ at the negative vertex labelled with the inverse label of the germ at the positive vertex.

From Theorem 1.1, P_e is type 2. So, Γ has two regions of type 2. At least one of these, say Δ is different from Δ_0. Let C be the consistently directed boundary of Δ. So, there are i and j such that the word ω obtained by reading around Δ involves only a_i (read from corners at positive vertices) and $\bar{a}_j$ (read from corners at negative vertices). If C is oriented clockwise around Δ, there are corresponding subwords ta_it^{-1} and ta_jt^{-1} of e; if C is oriented counterclockwise around Δ, there are subwords $t^{-1}a_it$ and $t^{-1}a_jt$ of e.

If $i \neq j$, then there is no word ω involving only a_i and $\bar{a}_j$ that is a relation in G. Therefore, i must equal j and there must be both positive and negative vertices on C. In this case, Γ is not reduced.

This contradiction proves the theorem. □

In [**2**], Brodskii and Howie define the sign index of an equation

$$e = \prod_{i=1}^{N} a_i t^{m_i}$$

to be the number of changes of sign in the cyclic sequence $(m_1, m_2, \ldots, m_N)$. The sign index of any equation is even. We note that if the sign index of an equation is greater than two, then the first condition of Theorem 1.1 is redundant in the presence of the others.

In [**9**], Levin proved that if the sign index of the equation

$$e = \prod_{i=1}^{N} a_i t^{m_i}$$

is zero, then e is solvable over any group. In [**11**], Stallings proved that if the sign index of e is two and a_i and a_j are the coefficients where these two sign changes take place (that is to say, $m_{i-1}m_i < 0$ and $m_{j-1}m_j < 0$), then e is solvable over any group in which a_i

and a_j have infinite order. Theorem 1.1 can be viewed as an extension of these previous results.

Proof of Theorem 1.2. Let e be as in the statement of Theorem 1.2. As explained above, if e is not solvable over G, then there is a reduced P_e graph Γ with a region Δ so that Δ satisfies the conclusion of Theorem 1.2 and so that the word obtained by reading around Δ is a relation of G. However, the assumptions on e make this impossible. $\square$

We point out that these results can be obtained using standard small cancellation theory arguments in the case when $p \geqslant 6$. In fact, the referee has a small cancellation proof for the case $p = 5$.

Acknowledgements. The authors thank the referee for many insightful comments as to both the form and content of this article.

References

1. W. A. Bogley and S. J. Pride, Aspherical relative presentations, *Proc. Edinb. Math. Soc.* **35** (1992), 1–39.
2. S. D. Brodskii and J. Howie, One relator products of torsion-free groups, *Glasgow Math. J.* **35** (1993), 99–104.
3. A. Clifford and R. Z. Goldstein, Tesselations of S^2 and equations over torsion-free groups, *Proc. Edinb. Math. Soc.* **38** (1995), 485–493.
4. M. Edjvet, Equations over groups and a theorem of Higman, Neumann, and Neumann, *Proc. Lond. Math. Soc.* **62** (1991), 563–589.
5. M. Edjvet and A. Juhasz, On equations over groups, *Int. J. Alg. Comp.* **4** (1994), 451–468.
6. S. M. Gersten, Reducible diagrams and equations over groups, in *Essays in group theory* (ed. S. M. Gersten) (Mathematical Sciences Research Institute Publications, New York, 1987).
7. J. Howie, The solution of length three equations over groups, *Proc. Edinb. Math. Soc.* **26** (1983), 89–96.
8. A. A. Klyachko, A funny property of sphere and equations over groups, *Commun. Alg.* **21** (1993), 2555–2575.
9. F. Levin, Solutions of equations over groups, *Bull. Am. Math. Soc.* **68** (1962), 603–604.
10. A. J. Sieradski, A coloring test for asphericity, *Q. J. Math. Oxford* **34** (1983), 97–106.
11. J. R. Stallings, A graph-theoretic lemma and group embeddings, in *Combinatorial group theory and topology* (ed. S. M. Gersten and J. R. Stallings), Annals of Mathematical Studies, vol. 111 (1987), pp. 145–155.