

Manuscript version: Author's Accepted Manuscript

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

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/141499>

How to cite:

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

Copyright and reuse:

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

© 2020 Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



Publisher's statement:

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

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

Polynomial-time proofs that groups are hyperbolic

Derek Holt, Stephen Linton, Max Neunhöffer, Richard Parker,
Markus Pfeiffer and Colva M. Roney-Dougal*

* = corresponding author

July 22, 2020

Abstract

It is undecidable in general whether a given finitely presented group is word hyperbolic. We use the concept of *pregroups*, introduced by Stallings in [22], to define a new class of van Kampen diagrams, which represent groups as quotients of *virtually* free groups. We then present a polynomial-time procedure that analyses these diagrams, and either returns an explicit linear Dehn function for the presentation, or returns **fail**, together with its reasons for failure. Furthermore, if our procedure succeeds we are often able to produce in polynomial time a word problem solver for the presentation that runs in linear time. Our algorithms have been implemented, and when successful they are many orders of magnitude faster than KBMAG, the only comparable publicly available software.

Keywords: Hyperbolic groups; word problem; van Kampen diagrams; curvature.

1 Introduction

The Dehn function of a finitely presented group is linearly bounded if and only if the group is hyperbolic. We describe a new, polynomial-time procedure for proving that a group defined by a finite presentation is hyperbolic, by establishing such a linear upper bound on its Dehn function. Our procedure returns a positive answer significantly faster than other methods, and in particular always terminates in low degree polynomial time, although sometimes it will terminate with **fail** even when the input group is hyperbolic. Our approach has the added advantage that it can sometimes be carried out by hand, which can enable one to prove the hyperbolicity of infinite families of groups.

A finitely generated group is (word) hyperbolic if its Cayley graph is negatively curved as a geometric metric space; that is, if its geodesic triangles are uniformly slim. There are several good sources, such as [1], for an introduction to and development of the basic properties of hyperbolic groups. Another useful reference for the specific properties that we need in this paper is [14, Chapter 6]. In particular, hyperbolic groups are finitely presentable, and they admit a Dehn algorithm. Furthermore, for groups that are defined by a finite presentation, this last condition implies that the group is hyperbolic.

We use [16, Chapter V] as reference for the theory of van Kampen diagrams over group presentations, and its application to groups defined by presentations that satisfy various small cancellation hypotheses. The arguments used in the proofs of these results can be formulated in terms of the assignment of curvature to the vertices, edges and faces of reduced van Kampen diagrams. The idea is to show that, under appropriate conditions, the curvature in those parts of the diagrams that are not close to the boundary is non-positive. For example, one specific conclusion of [16, Theorem 4.4] is that groups with presentations that satisfy $C'(\lambda)$ for $\lambda \leq 1/6$, or $T(4)$ together with $C'(\lambda)$ for $\lambda \leq 1/4$,

have Dehn algorithms. It is not assumed in these results that the group presentations in question are finite, but we shall be working only with finite presentations in this paper, in which case these conditions imply that the group is hyperbolic.

The algorithmic methods developed in this paper involve the assignment of curvature to van Kampen diagrams in the manner described above, such that the total curvature of every diagram is 1. A serious limitation of methods that rely on small cancellation conditions is that they are unlikely to be satisfied in the presence of short defining relators such as powers x^n , where n is small and x is a generator. However, such relators are present in many of the most interesting group presentations. Our methods use the theory of *pregroups*, developed by Stallings in [22]. This theory enable us to remove short relators, replacing them with certain other relators of length three (the pregroup relators), which we then ignore when considering generalisations of small cancellation.

Our general aim is to assign the curvature in such a way that vertices and edges have zero curvature, faces labelled by pregroup relators have non-positive curvature, and other faces that are not close to the boundary of the diagram have curvature that is bounded above by some constant $-\varepsilon < 0$. Our principal theoretical result is Theorem 5.9, which states roughly that, if we can assign curvature in this manner to all reduced diagrams over the presentation, then the group has a Dehn function that is bounded above by a linear function that we can specify explicitly in terms of ε and various basic parameters of the defining presentation. So the group is hyperbolic.

Although we cannot expect such methods to work for all hyperbolic groups, we can explore a variety of methods of assigning curvature, which we call *curvature distribution schemes*. In this paper we restrict attention to a single such scheme, which we call **RSym**. In Theorem 6.13, we apply Theorem 5.9 to calculate an explicit linear bound on the Dehn function that is satisfied in the event that **RSym** succeeds in assigning curvature in the required fashion. As a simple remedy in examples in which **RSym** does not succeed and some, but not all, interior faces of a diagram end up with zero or positive curvature, we could try to transfer some of the negative curvature from those faces that already have it to those that do not. This process is hard to implement in a computer algorithm, but it can often be done by hand, which significantly increases the applicability of the methods.

The principal algorithmic challenge is to prove that **RSym** succeeds on a sufficiently large set of reduced diagrams over the input presentation. Our main algorithm, **RSymVerify**, which is described in Section 7, attempts to achieve this. It is technically complicated and involves a detailed study of the possible neighbourhoods of interior faces in diagrams over the presentation.

For hyperbolic groups, the word problem is solvable in linear time by a Dehn algorithm. However, the current best results require as a preprocessing step the computation of the set S of all words of length up to 8δ that are trivial in the group, where geodesic triangles in the Cayley graph are δ -thin. Given a linear bound λn on the Dehn function $D(n)$ of a finitely-presented group G , it is therefore theoretically possible to use brute force to test all such words for triviality in G , and hence to construct the set S . However, this requires time and space that are exponential in both δ and λ , and so is completely impractical. We instead devise an additional polynomial-time test, which, if satisfied, enables the polynomial-time construction of a linear time word problem solver. This additional test is also the basis of future joint work by the sixth author, which will give a polynomial-time construction of a quadratic time solver for the conjugacy problem, the second of Dehn's classic problems.

Here is a breakdown of the contents of the paper. In Section 2, we summarise the required properties of pregroups, and define a new kind of presentation, called a *pregroup presentation*, for a group G . It was shown by Rimlinger in [20] that a finitely generated group H is virtually free if and only if H is the universal group $U(P)$ of a finite pregroup P : see Theorem 2.14. Pregroup presentations enable us to view the group G as a quotient

of a virtually free group $U(P)$, rather than just as a quotient of a free group, and hence to ignore any failures of small cancellation on the defining relators of $U(P)$.

In Section 3 we define *coloured* van Kampen diagrams over these new pregroup presentations, where the relators of the virtually free group $U(P)$ (which we collect in a set V_P) are coloured red, and the additional relators (which we collect in a set \mathcal{R}) are green. We show in Proposition 3.17 that, given any coloured van Kampen diagram Γ satisfying a certain technical condition, there exists a coloured van Kampen diagram Γ' , with the same boundary word as Γ , whose area is bounded by an explicit linear function of the number of green faces of Γ . Hence, to prove that a group is hyperbolic it suffices to prove a linear upper bound on the number of green faces appearing in any reduced coloured diagram of boundary length n .

In Section 4 we show that if we replace our presentation by a certain related presentation, then we can assume without loss of generality that each vertex of a coloured diagram is incident with at least *two* green faces. This property will be critical to our later curvature analysis, and is automatically satisfied by diagrams over free groups, since for them all faces are green.

Section 5 is devoted to the definition and general discussion of curvature distribution schemes. These provide an overall schema for the design of many possible methods for proving that a group given by a finite pregroup presentation is hyperbolic: since a pregroup presentation is a generalisation of a standard presentation, these methods apply to all finite presentations. As mentioned earlier, in Theorem 5.9 we characterise how these schema can produce explicit bounds on the Dehn function.

In Section 6 we present the **RSym** curvature distribution scheme mentioned earlier. For reasons of space and ease of comprehension, we restrict attention to this scheme in this paper, but our approach can be used to define many others. Theorem 6.13 gives an explicit bound on the Dehn function of the presentation when **RSym** succeeds on all coloured van Kampen diagrams of minimal coloured area.

In Section 7 we prove (see Theorems 7.20 and 7.22) that under some mild and easily testable assumptions on the set \mathcal{R} of green relators, one can test whether **RSym** succeeds on *all* of the (infinitely many) coloured van Kampen diagrams of minimal area. This test is carried out by our procedure **RSymVerify** (Procedure 7.19), which runs in time $O(|X|^5 + r^3|X|^4|\mathcal{R}|^2)$, where X is the set of generators, and r is the length of the longest green relator. Our assumptions hold, for example, for all groups given as quotients of free products of free and finite groups. We also prove that without these assumptions, one can test whether **RSym** succeeds on all minimal diagrams in time polynomial in $|X|$ and $r|\mathcal{R}|$: our procedure to do this is called **RSymIntVerify** (Procedure 7.30).

In Section 8 we go on to consider the word problem. Whilst a successful run of **RSymVerify** or **RSymIntVerify** proves an explicit linear bound on the Dehn function, it is rarely practical to construct a set of Dehn rewrites. We present a low degree polynomial-time method to construct a word problem solver: see Theorem 8.6. The construction of the solver succeeds in many but not all examples in which **RSym** succeeds, and the solver itself runs in linear time: see Theorem 8.9 and Proposition 8.12.

In Section 9 we consider a variety of examples of finite group presentations, and show how **RSym** can be used by hand to prove that the groups are hyperbolic. In particular, we prove that **RSym** succeeds on groups satisfying any of a wide variety of small cancellation conditions, we use **RSym** to analyse two infinite families of presentations, and we discuss a range of possible future applications of **RSym** to problems concerning the hyperbolicity of finitely-presented groups.

Our procedures have been released as part of both the **GAP** [6] and **MAGMA** [2] computer algebra systems, and in Section 10 we present runtimes on a variety of examples, including some with very large numbers of generators and relations. Almost none of these examples could have been analysed using previously existing methods, due to the size of the presentations.

Since we have introduced many new terms and much new notation, we conclude with an Appendix containing lists of all new terms, notation, and procedures.

As far as we know, the only other publicly available software that can prove hyperbolicity of a group defined by an arbitrary finite presentation is the first author's KBMAG package [13] for computing automatic structures. Hyperbolicity is verified by proving that geodesic bigons in the Cayley graph are uniformly thin, as described in [12, Section 5]. It was proved by Papasoglu in [19] that this property implies hyperbolicity, but it does not provide a useful bound on the Dehn function. An algorithm for computing the "thinness" constant for geodesic triangles in the Cayley graph of a hyperbolic group is described in [5], but this is of limited applicability in practice, on account of its high memory requirements. Even on the simplest examples, the KBMAG programs involve far too many computational steps for them to be carried out by hand, and they can only be applied to individual presentations. The automatic structure does however provide a fast method (at worst quadratic time) of reducing words to normal form and hence solving the word problem in the group.

Shortly before submitting this paper, we became aware of a paper by Lysenok [17], which explores similar concepts of redistributing curvature to prove hyperbolicity to those presented in Section 5 of this paper. His main theorem is similar to our Theorem 5.9, but the ideas are less fully developed.

2 Pregroup presentations

In this section we introduce pregroups, establish some of their elementary properties, and show that any quotient of a virtually free group by finitely many relators can be defined by a finite pregroup presentation. Pregroups were first defined by Stallings in [22].

Definition 2.1. A *pregroup* is a set P , with a distinguished element 1, equipped with a partial multiplication $(x, y) \rightarrow xy$ which is defined for $(x, y) \in D(P) \subseteq P \times P$, and with an involution $\sigma : x \rightarrow x^\sigma$, satisfying the following axioms, for all $x, y, z, t \in P$:

- (P1) $(1, x), (x, 1) \in D(P)$ and $1x = x1 = x$;
- (P2) $(x, x^\sigma), (x^\sigma, x) \in D(P)$ and $xx^\sigma = x^\sigma x = 1$;
- (P3) if $(x, y) \in D(P)$ then $(y^\sigma, x^\sigma) \in D(P)$ and $(xy)^\sigma = y^\sigma x^\sigma$;
- (P4) if $(x, y), (y, z) \in D(P)$ then $(xy, z) \in D(P)$ if and only if $(x, yz) \in D(P)$, in which case $(xy)z = x(yz)$;
- (P5) if $(x, y), (y, z), (z, t) \in D(P)$ then at least one of $(xy, z), (yz, t) \in D(P)$.

Since we will often be working with words over P , if we wish to emphasise that two (or more) consecutive letters, say x and y , of a word w are to be multiplied we shall write $[xy]$.

Note that (P2) implies that $1^\sigma = 1$, and that (P1), (P2) and (P4) imply that inverses are unique: if $xy = 1$ then $y = x^\sigma$. It was shown in [9] that (P3) follows from (P1), (P2) and (P4), but we include it to keep our numbering consistent with the literature.

Definition 2.2. Let P be a pregroup. We define $X = P \setminus \{1\}$, and let σ be the involution on X . We write X^σ to denote X , equipped with this involution, but will sometimes omit the σ , when the meaning is clear. We shall write $F(X^\sigma)$ to denote the group defined by the presentation $\langle X \mid xx^\sigma : x \in X \rangle$. If σ has cycle structure $1^k 2^l$ on X , then $F(X^\sigma)$ is the free product of k copies of C_2 and l copies of \mathbb{Z} .

Let V_P be the set of all length three relators over X of the form $\{xy[xy]^\sigma : x, y \in X, (x, y) \in D(P), x \neq y^\sigma\}$. The *universal group* $U(P)$ of P is the group given by

$$\langle X \mid \{xx^\sigma : x \in X\} \cup V_P \rangle = F(X^\sigma) / \langle\langle V_P \rangle\rangle,$$

where $\langle\langle V_P \rangle\rangle$ denotes the normal closure of V_P in $F(X^\sigma)$.

Since this presentation of $U(P)$ is on a set of monoid generators that is closed under inversion, we can and shall write the elements of $U(P)$ as words over X , and use x^σ rather than x^{-1} to denote the inverse of x . More generally, if $w = x_1 \dots x_n \in F(X^\sigma)$, then $w^{-1} =_{F(X^\sigma)} x_n^\sigma x_{n-1}^\sigma \dots x_1^\sigma$.

Stallings in [22] defines $U(P)$ as the universal group of P in a categorical sense: every morphism from P to a group G factors through $U(P)$.

Remark 2.3. It is an easy exercise using the pregroup axioms to show that if $(x, y) \in D(P)$ and $z = [xy]$ then $(y, z^\sigma), (z^\sigma, x) \in D(P)$, so that if $xyz^\sigma \in V_P$ then all products of cyclic pairs of letters are defined in P .

Example 2.4. A pregroup P such that $U(P) = F(X^\sigma)$ is free of rank n can be made by letting X have $2n$ elements, defining σ to be fixed-point-free on X , and letting the only products be $xx^\sigma = 1$, $1x = x1 = x$, and $1 \cdot 1 = 1$, for all $x \in X$.

Example 2.5. A pregroup P such that $U(P)$ is the free product of finite groups G and H can be made as follows. We let P have elements the disjoint union of $\{1\}$, $G \setminus \{1\}$ and $H \setminus \{1\}$. We define σ to be the inversion map on both G and H , and to fix 1. We let $D(P) = (G \times G) \cup (H \times H)$, and define all products as in the parent groups.

More generally, if the finite groups G and H intersect in a subgroup I , and again $P = G \cup H$ with inversion and $D(P)$ defined as before, then $U(P)$ is the amalgamated free product $G *_I H$.

Definition 2.6. We define a word $w \in X^*$ to be σ -reduced if w contains no consecutive pairs xx^σ of letters: this is a slight generalisation of free reduction. We define *cyclically* σ -reduced similarly.

The word $w = x_1 \dots x_n \in X^*$ is P -reduced if either $n \leq 1$, or $n > 1$ and no pair (x_i, x_{i+1}) lies in $D(P)$. The word w is *cyclically* P -reduced if either (i) $n \leq 1$; or (ii) w is P -reduced, $n > 1$, and $(x_n, x_1) \notin D(P)$.

Stallings defines a relation \approx on the set of P -reduced words in X^* as follows.

Definition 2.7. Let $v = x_1 \dots x_n \in X^*$ be P -reduced and let $w = y_1 \dots y_m$ be any word in X^* . Then we write $v \approx w$ if $n = m$ and there exist $s_0 = 1, s_1, \dots, s_{n-1}, s_n = 1 \in P$ such that $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i) \in D(P)$ for all i , and $y_i = [s_{i-1}^\sigma x_i s_i]$. We say that w is an *interleave* of v . In the case when $s_i \neq 1$ for a single value of i , we call the transformation from v to w a *single rewrite*.

Notice that if $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i) \in D(P)$ then it follows from (P4) that $(s_{i-1}^\sigma, [x_i s_i]) \in D(P)$, and that $y_i = (s_{i-1}^\sigma x_i) s_i = s_{i-1}^\sigma (x_i s_i)$.

Example 2.8. Let $G = \langle a \rangle$ and $H = \langle b \rangle$ be cyclic of order 6 with $a^2 = b^2$ and $I := G \cap H = \langle a^2 \rangle$, and let $P = G \cup H = \{1, a_1, a_3, a_5, b_1, b_3, b_5, i_2, i_4\}$ as in Example 2.5 above, with the interpretation $a_j = a^i$, $b_j = b^i$ and $i_j = a^j = b^j \in I$. Let $v = a_1 b_1 a_3$. Then, by choosing $s_1 = i_2$ and $s_2 = i_4$, we obtain the interleave $w = a_3 b_3 a_5$.

Theorem 2.9 ([22, 3.A.2.7, 3.A.2.11, 3.A.4.5 & 3.A.4.6]). *Let P be a pregroup, let $X = P \setminus \{1\}$, and let $v, w \in U(P)$, with v a P -reduced word. Then*

- (i) *if $v \approx w$ then w is P -reduced;*
- (ii) *interleaving is an equivalence relation on the set of P -reduced words over X ;*
- (iii) *each element $g \in U(P)$ can be represented by a P -reduced word in X^* ;*
- (iv) *if w is P -reduced, then v and w represent the same element of $U(P)$ if and only if $u \approx v$; in particular, P embeds into $U(P)$.*

Corollary 2.10. *Let P be a finite pregroup. Then the word problem in $U(P)$ is soluble in linear time.*

Proof. The only P -reduced word representing $1_{U(P)}$ is the empty word, so we can solve the word problem in $U(P)$ by reducing words using the products in $D(P)$. This process is a Dehn algorithm, which by [4] requires time linear in the length of the input word. \square

Definition 2.11. We now define a new type of presentation, which we shall call a *pregroup presentation*. Let P be a pregroup, let $X = P \setminus \{1\}$, let σ be the involution giving inverses in X , and let $\mathcal{R} \subset X^*$ be a set of cyclically P -reduced words. We write

$$\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$$

to define a group presentation $\mathcal{P} = \langle X \mid \{xx^\sigma : x \in X\} \cup V_P \cup \mathcal{R} \rangle$ on the set X of monoid generators.

Example 2.12. We construct two pregroup presentations for $G = \langle x, y \mid x^\ell, y^m, (xy)^n \rangle$ (with $\ell, m \geq 2$). For the first, let P be the pregroup with universal group $C_\ell * C_m$, as in Example 2.5, so that

$$P = \{1, x = x_1, x_2, \dots, x_{\ell-1}, y = y_1, y_2, \dots, y_{m-1}\}$$

with each $x_i =_P x^i$ and $y_i =_P y^i$. Then σ fixes 1, maps each x_i to $x_{\ell-i}$ and maps each y_j to y_{m-j} . The set V_P consists of all triples $x_i x_j x_{i+j}^\sigma$ and $y_i y_j y_{i+j}^\sigma$, where $+$ represents modular arithmetic (and no subscript is equal to zero). Let $\mathcal{R} = \{(xy)^n\}$. Then $\mathcal{P} = \langle (P \setminus \{1\})^\sigma \mid V_P \mid \mathcal{R} \rangle$ is a pregroup presentation for G .

For the second, assume for convenience that $\ell, m \geq 3$, and let $Q = \{1, x, X, y, Y\}$ be a pregroup with 1 as the identity, σ exchanging x with X and y with Y , and no other products defined. Notice that $U(Q)$ is a free group of rank 2. Let $\mathcal{S} = \{x^\ell, y^m, (xy)^n\}$. Then $\mathcal{Q} = \langle \{x, X, y, Y\}^\sigma \mid \emptyset \mid \mathcal{S} \rangle$ is also a pregroup presentation for G .

Remark 2.13. We shall assume throughout the rest of the paper that there are no relators of the form x^2 for $x \in X$ in \mathcal{R} and that, instead, we have chosen a pregroup P such that $x^\sigma = x$. This can always be achieved by, for example, choosing P such that $U(P) = F(X^\sigma)$. Notice also that $\mathcal{R} \cap V_P = \emptyset$, since each element of \mathcal{R} is cyclically P -reduced.

We finish this section by considering the applicability of these presentations.

Theorem 2.14 ([20, Corollary to Theorem B]). *A finitely generated group G is virtually free if and only if G is the universal group of a finite pregroup.*

The class of virtually free groups includes amalgamated free products of finite groups, and HNN extensions with finite base groups, which is the source of many of the pregroups that are useful in the algorithmic applications to proving hyperbolicity that are described in this paper. More generally, a group is virtually free if and only if it is the fundamental group of a finite graph of groups with finite vertex groups [21, Proposition 11].

For the remainder of the paper, we shall be working with groups given by finite pregroup presentations. The following immediate corollary shows that this includes all quotients of virtually free groups by finitely many additional relators.

Corollary 2.15. *Let a group G have pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$, as in Definition 2.11. Then $G \cong U(P)/\langle\langle \mathcal{R} \rangle\rangle$, where $\langle\langle \mathcal{R} \rangle\rangle$ denotes the normal closure of \mathcal{R} in $U(P)$. Furthermore, any group that is a quotient of a virtually free group by finitely many additional relators has a finite pregroup presentation.*

3 Diagrams over pregroups

In this section, we introduce coloured van Kampen diagrams, which are a natural generalisation of van Kampen diagrams to pregroup presentations. After completing the introductory material, our main result is Proposition 3.17, which shows that if a word of length n can be written as a product of conjugates of k relators from \mathcal{R}^\pm over $U(P)$, then it can be written as a product of conjugates of $\lambda k + n$ relators from $\mathcal{R}^\pm \cup V_P$ over $F(X^\sigma)$, where λ depends only on the maximum length r of the relators in \mathcal{R} , and \mathcal{R}^\pm denotes $\mathcal{R} \cup \{R^{-1} : R \in \mathcal{R}\}$.

This section contains many new definitions, and we remind the reader that the Appendix contains a list of all new terms and notation.

In general we follow standard terminology for van Kampen diagrams, as given in [16, Chapter 5, §1] for example. For clarity, we record some definitions that will be useful in what follows.

Definition 3.1. We shall orient each face clockwise. We shall count all incidences with multiplicities, for example a vertex may be incident more than once with the same face.

In the present article, we shall require our diagrams to be simply connected. There is therefore a unique *external face*, and its label is the *external word*. All other faces are *internal*. If an element $x \in X$ is self-inverse in P , then x has order 2 in $U(P)$ and we will identify x with x^σ , so that an edge may have label x on both sides.

We will refer to a nontrivial path of maximal length that is common to two adjacent faces of a diagram as a *consolidated edge*.

We denote the boundary of a face f or a diagram Γ by $\partial(f)$ and $\partial(\Gamma)$, respectively. We consider $\partial(f)$ (and $\partial(\Gamma)$) to contain both vertices and edges, but abuse notation and write $|\partial(f)|$ for the number of edges. An internal face f of a diagram Γ is a *boundary face* if $|\partial(f) \cap \partial(\Gamma)| \geq 1$. A vertex or edge of Γ is a *boundary vertex* or edge if it is contained in $\partial(\Gamma)$.

Definition 3.2. A *coloured van Kampen diagram* over the pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ is a van Kampen diagram with edge labels from X^σ , and face labels from $V_P \cup \mathcal{R}^\pm$, in which the faces labelled by an element of V_P are coloured red, and the faces labelled by an element of \mathcal{R}^\pm , together with the external face, are coloured green. We shall often refer to a coloured van Kampen diagram as a *coloured diagram*.

For v a vertex in a coloured diagram Γ , we shall write $\delta(v, \Gamma)$ for the degree of v , $\delta_G(v, \Gamma)$ for *green degree* of v : the number of green faces incident with v in Γ , and $\delta_R(v, \Gamma)$ for the number of red faces incident with v in Γ .

Notice from our relator set V_P that all red faces are triangles: we shall often refer to them as such. If a word w is a product of conjugates of exactly k relators from \mathcal{R}^\pm in $U(P)$, then there exists a coloured diagram for w with exactly k internal green faces. The proof of this is essentially the same as for standard van Kampen diagrams; see [16, Chapter V, Theorem 1.1], for example.

Definition 3.3. The area of a coloured diagram Γ is its total number of internal faces, both red and green, and is denoted $\text{Area}(\Gamma)$.

However, when comparing areas of diagrams, it is convenient to count green faces first and then red triangles.

Definition 3.4. Let Γ be a coloured diagram. The *coloured area* $\text{CArea}(\Gamma)$ of Γ is an ordered pair $(a, b) \in \mathbb{N} \times \mathbb{N}$, where a is the number of internal green faces of Γ and b is the number of red triangles. Let Δ be a coloured diagram with $\text{CArea}(\Delta) = (c, d)$. We say that $\text{CArea}(\Gamma) \leq \text{CArea}(\Delta)$ if $a < c$ or if $a = c$ and $b \leq d$. A diagram has *minimal coloured area* for a word w if its coloured area is minimal over all diagrams with boundary word w .

Definition 3.5. Let Γ be a coloured van Kampen diagram. A *subdiagram* of Γ is a subset of the edges, vertices and internal faces of Γ which, together with a new external face coloured green, form a coloured diagram in their own right.

In particular, we do not allow annular subdiagrams.

Definition 3.6. A coloured diagram is *semi- σ -reduced* if no two distinct adjacent faces are labelled by w_1w_2 and $w_2^{-1}w_1^{-1}$ for some relator $w_1w_2 \in V_P \cup \mathcal{R}^\pm$ and have a common consolidated edge labelled by w_1 and w_1^{-1} . It is *σ -reduced* if the same also holds for a single face adjacent to itself.

Our definition of a σ -reduced coloured diagram corresponds to the usual definition of a reduced diagram; see [16, p241], for example. Unfortunately, the proof in [16, Chapter V, Lemma 2.1] that a diagram of minimal area is reduced breaks down in our situation for faces adjacent to themselves, because two cyclic conjugates of a P -reduced word can be mutually inverse in $F(X^\sigma)$: we shall eventually get around this problem, and the first step in this direction is Proposition 3.9.

First, however, we make a natural generalisation of σ -reduction.

Definition 3.7. A coloured diagram is *semi- P -reduced* if no two distinct adjacent green faces are labelled by w_1w_2 and $w_3^{-1}w_1^{-1}$ and have a common consolidated edge labelled by w_1 and w_1^{-1} , where w_2 and w_3 are equal in $U(P)$ (which, by Theorem 2.9 is equivalent to $w_2 \approx w_3$).

Notice in particular that a semi- P -reduced diagram is semi- σ -reduced.

Proposition 3.8. *Let Γ be a coloured diagram with boundary word w . Then there exists a semi- P -reduced coloured diagram Δ with boundary word w such that $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$. If Γ is not already semi- P -reduced then this inequality is strict. Notice in particular that the diagram Δ is semi- σ -reduced.*

Proof. If a coloured diagram Γ fails to be semi- P -reduced, then Γ contains two adjacent green faces labelled by w_1w_2 and $w_3^{-1}w_1^{-1}$, as in Definition 3.7.

Since $w_2 =_{U(P)} w_3$, we can remove the consolidated edge labelled w_1 , identify any consecutive edges with inverse labels, and fill in the resulting region, with label the cyclically σ -reduced word resulting from $w_2w_3^{-1}$, by a number of red triangles, yielding a diagram Γ_1 . It is possible that $\text{Area}(\Gamma_1) > \text{Area}(\Gamma)$, but $\text{CArea}(\Gamma_1) < \text{CArea}(\Gamma)$, since Γ_1 has two fewer green faces than Γ . The process terminates at a semi- P -reduced diagram Δ . \square

The following result will enable us to restrict our attention to σ -reduced diagrams later in the paper.

Proposition 3.9. *Let G have pregroup presentation \mathcal{P} . Suppose that there exists a coloured diagram Γ over \mathcal{P} that contains a face f that is adjacent to itself in such a way that Γ is not σ -reduced. Then f is green.*

Furthermore, there exists $t \in X$ such that

- (i) $t^\sigma = t$ and $t =_G 1$;
- (ii) *all coloured diagrams Δ of minimal coloured area with boundary label t are σ -reduced and semi- P -reduced, and satisfy $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$.*

Proof. The only way that a red triangle could be adjacent to itself is if an element of X is trivial in $U(P)$, contradicting Theorem 2.9(iv). Hence f is green.

Let f be adjacent to itself via a consolidated edge labeled w_1 , so that reading $\partial(f)$ from the side of the edge labelled w_1 gives w_1w_2 , and from the other side gives $w_1^{-1}w_2^{-1}$. Then w_2 contains a subword w_1^{-1} , so there are words v_1 and v_2 with $w_2 =_{F(X^\sigma)} v_1w_1^{-1}v_2$ and $w =_{F(X^\sigma)} w_1w_2 =_{F(X^\sigma)} w_1v_1w_1^{-1}v_2$.

We have assumed that w_2 is the inverse of the cyclic subword of w that starts just after w_1^{-1} and finishes just before it, namely $v_2 w_1 v_1$. So $w_2 =_{F(X^\sigma)} v_1^{-1} w_1^{-1} v_2^{-1}$ and hence, since there is no cancellation in these products, $v_1 =_{F(X^\sigma)} v_1^{-1}$ and $v_2 =_{F(X^\sigma)} v_2^{-1}$. Both v_1 and v_2 are σ -reduced, so each is an $F(X^\sigma)$ -conjugate of a self-inverse element of X^σ .

The face f in Γ encloses regions with boundary labels v_1 and v_2 , so the corresponding involutions are trivial in G . In fact these regions are subdiagrams Γ_1 and Γ_2 of Γ , and at least one of them, say Γ_1 , does not contain f , and hence has coloured area less than that of Γ . By identifying inverse pairs of edges on $\partial(\Gamma_1)$ we obtain a diagram with the same coloured area as Γ_1 , and with boundary label a self-inverse element of X^σ .

Now let Δ be a diagram with smallest possible coloured area such that its boundary label is some $t \in X$ with $t = t^\sigma$. Then Δ is semi- P -reduced by Proposition 3.8, and it must be σ -reduced, since otherwise we could repeat the above argument and obtain such a diagram with smaller coloured area. Furthermore, $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma_1) < \text{CArea}(\Gamma)$. \square

Definition 3.10. A *plane graph* is a planar graph embedded in the plane, so that the faces are determined. We denote the sets of all edges, vertices, and internal faces of a coloured van Kampen diagram or plane graph Γ by $E(\Gamma)$, $V(\Gamma)$ and $F(\Gamma)$, respectively, and if the faces of Γ are coloured then we write $F(\Gamma) = F_G(\Gamma) \cup F_R(\Gamma)$, where $F_G(\Gamma)$ and $F_R(\Gamma)$ denote the sets of internal green and red faces of Γ . For an edge e of Γ , we define $\delta_G(e, \Gamma)$ to be the number of green faces incident with e . So $\delta_G(e, \Gamma) = 0, 1$ or 2 .

Recall from Definition 3.1 that we count incidences of faces with edges and vertices with multiplicity: we do the same in the corresponding plane graph. Recall from Definition 3.2 that the external face of any coloured diagram Γ is green, and so contributes 1 to the value of $\delta_G(e, \Gamma)$ for each edge on its boundary. Recall also Definition 3.2 of $\delta_G(v, \Gamma)$.

In the proof of the next result, we use the well-known Euler formula $|V(\Gamma)| + |F(\Gamma)| - |E(\Gamma)| = 1$ (we have 1 rather than 2 because $F(\Gamma)$ does not include the external face).

Proposition 3.11. *Let Γ be a simply-connected plane graph, where all faces have been coloured red or green and the unique external face is green. Assume that the boundary of every red face has length 3. Then, with the notation of Definition 3.10, we have*

$$\sum_{e \in E(\Gamma)} \delta_G(e, \Gamma) = |F_R(\Gamma)| + 2 \left(1 - |F_G(\Gamma)| + \sum_{v \in V(\Gamma)} (\delta_G(v, \Gamma) - 1) \right).$$

Proof. Let $E := |E(\Gamma)|$, $V := |V(\Gamma)|$, $F := |F(\Gamma)|$, $F_R := |F_R(\Gamma)|$ and $F_G := |F_G(\Gamma)|$. The proof is by induction on F_R , so suppose first that $F_R = 0$. Then $F_G = F$ and $\sum_{v \in V(\Gamma)} (\delta_G(v, \Gamma) - 1) = 2E - V$, so the right hand side of the above equation is $2(1 - F + 2E - V)$, which by Euler's formula is equal to $2E = \sum_{e \in E} \delta_G(e, \Gamma)$.

For the inductive step, it is enough to prove that the equation remains true when we change the colour of an internal triangle from green to red, since any two-face-coloured graph can be made by first colouring all faces one colour and then changing the colour of some faces. The triangle has three sides, so when it changes colour, $\sum_{e \in E} \delta_G(e, \Gamma)$ decreases by 3, the value of F_R increases by 1 and F_G decreases by 1. The triangle is incident with three vertices, so $\sum_{v \in V} (\delta_G(v, \Gamma) - 1)$ decreases by 3. Thus, the formula still holds. \square

Lemma 3.12. *Let Γ be a coloured diagram with cyclically σ -reduced boundary word. Then Γ contains no vertices of degree 1.*

Proof. We assumed in Definition 2.11 that all boundary labels of internal green faces are cyclically P -reduced, and hence in particular are cyclically σ -reduced. Furthermore, all boundary labels of red triangles are σ -reduced by the definition of V_P . We have now assumed that the boundary word of Γ is cyclically σ -reduced. If v is a vertex in Γ of degree 1, then the boundary of the face containing v has label containing a subword xx^σ , where x is the label of the unique edge incident with v . This is a contradiction. \square

A number of the results that we prove later require the assumption that certain elements of X are not trivial in G , as in Proposition 3.9. A potential problem with this assumption is that, on the one hand the triviality of generators is known to be undecidable in general in finitely presented groups, but on the other hand our algorithms need to be able to test whether it holds on the presentations that we want to test for hyperbolicity.

Fortunately, as we shall see in Theorem 6.12, there is a way of avoiding these difficulties. Although we cannot assume in our proofs that our diagrams contain no loops, we can typically assume that, on consideration of a minimal counterexample to whatever we are trying to prove, if there is a such a loop in the diagram Γ under consideration, then Γ has the smallest possible coloured area among all such diagrams. The following definition formalises this assumption.

Definition 3.13. A letter $x \in X$ such that $x^\sigma = x$ or x is a letter of a relator in V_P is called a V^σ -letter. A coloured diagram Γ over \mathcal{P} is *loop-minimal* if every coloured diagram Δ over \mathcal{P} that contains a loop labelled by some V^σ -letter $x \in X$ satisfies $\text{CArea}(\Delta) \geq \text{CArea}(\Gamma)$.

Example 3.14. 1. If no element of X is trivial in G , then all coloured diagrams over \mathcal{P} are loop-minimal.
2. If $U(P)$ is free, and constructed as in Example 2.4, then there are no V^σ -letters, and so all diagrams over \mathcal{P} are loop-minimal.
3. Assume that at least one V^σ -letter is trivial in G , and let Δ be a diagram of smallest coloured area with boundary word a single V^σ -letter. Then the set of loop-minimal coloured diagrams is the set of diagrams with coloured area less than or equal to $\text{CArea}(\Delta)$.

Our algorithms to prove hyperbolicity will be able to verify for certain presentations that no V^σ -letter is trivial in G , so that all diagrams over these presentations are loop-minimal: see Theorem 6.12. This approach will succeed for almost all of the examples in Section 9, and for the remainder we shall demonstrate other methods to prove that no element of X is trivial in G .

In the remainder of this section, we shall work towards a proof of Proposition 3.17: the total area of a loop-minimal coloured diagram is bounded above by a linear function of the boundary length and the number of green faces. First we shall prove two results which allow us to assume that every vertex v in a coloured diagram Γ satisfies $\delta_G(v, \Gamma) \geq 1$.

Lemma 3.15. *Let Γ be a coloured diagram over \mathcal{P} with boundary word w . Assume that Γ contains a vertex v with three consecutive adjacent red triangles, and that none of the edges in any red triangle incident with v is a loop based at v .*

Then there exists a diagram Δ over \mathcal{P} with boundary word w , with the same green faces as Γ , satisfying $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$, in which v is incident with at least one fewer red triangle than it is in Γ , and in which none of the edges of any of the red triangles incident with v is a loop based at v .

Proof. Let f_1, f_2 and f_3 be the three consecutive red triangles, with edge labels as in the left hand picture of Figure 1. We assume that the edges labelled $\{a, c, e\}$ are pairwise distinct, and that the edges $\{c, e, g\}$ are pairwise distinct, but we allow the possibility that the edge labelled a is equal to the edge labelled g (so that v has degree 3). Some of the




Figure 1: Reducing the degree of v

other vertices and edges in Figure 1 may not be distinct, but notice that the assumption on loops based at v implies that none of v_1, v_2, v_3 and v_4 coincide with v , and hence that f_1, f_2 and f_3 are pairwise distinct.

Then $a^\sigma =_P bc^\sigma$, $d =_P c^\sigma e$ and $g =_P ef$. Hence by Axiom (P5) at least one of the pairs $([bc^\sigma], e), ([c^\sigma e], f)$ is in $D(P)$. Suppose that $([bc^\sigma], e) \in D(P)$ (the other case is similar), and let $x = [bc^\sigma e]$. Then

$$x =_P a^\sigma e =_P [bc^\sigma]e =_P b[c^\sigma e] =_P bd.$$

If $x =_P 1$ then $e = a$ and $d = b^\sigma$, so Γ is not semi- σ -reduced, and v has degree at least four (since $a^\sigma e$ is not a subword of the label of any red triangle). We may delete faces f_1 and f_2 , identifying v_1 with v_3 , the directed edge labelled b with the one labelled d^σ , and the edge labelled a with the one labelled e , to produce a diagram Δ . Then $\delta_R(v, \Delta) = \delta_R(v, \Gamma) - 2$, and Δ has no loops based at v (since the only amalgamated vertices are v_1 and v_3 , which do not coincide with v).

Otherwise, $x \neq 1$, and so $axe^\sigma, x^\sigma bd \in V_P$, and there is a diagram Δ in which the triangles f_1 and f_2 have been replaced by triangles f'_1 and f'_2 respectively, with labels axe^σ and $x^\sigma bd$ (as in the right hand picture of Figure 1). We have deleted an edge $\{v, v_2\}$, and added an edge $\{v_1, v_3\}$. Since v_1 and v_3 are not equal to v , the number of red and green faces of Δ is identical to that of Γ , but $\delta(v, \Delta) < \delta(v, \Gamma)$. Since the only added edge is not incident with v , the condition on loops based at v still holds. \square

Theorem 3.16. *Let Γ be a loop-minimal coloured diagram over \mathcal{P} with cyclically σ -reduced boundary word w . Then w is the boundary word of a loop-minimal coloured diagram Δ with $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$, the same green faces as Γ , and every vertex v of Δ satisfies $\delta_G(v, \Delta) \geq 1$. Furthermore, if Γ contains a vertex v with $\delta_G(v, \Gamma) = 0$, then $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$.*

Proof. Suppose that v is a vertex of Γ such that $\delta_G(v, \Gamma) = 0$. We shall show that it is possible to modify Γ to produce a diagram Δ with the same boundary word and green faces, but fewer red triangles. Since $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$, the new diagram Δ must be loop-minimal. A vertex of degree 1 does not exist in Γ by Lemma 3.12.

We prove the result first under the assumption that the three vertices of any red triangle are distinct. Suppose that $\delta(v, \Gamma) = 2$. Let the outgoing edge labels be a and b . Then one of the incident triangles has label $a^\sigma b[b^\sigma a]$ and the other has label $b^\sigma a[a^\sigma b]$, and so Γ is not semi- σ -reduced. Thus there is a diagram Δ with boundary word w in which v and both triangles have been removed from Γ , and in which the edges labelled $a^\sigma b$ and $b^\sigma a$ have been identified. Then Δ has the same green faces as Γ , but two fewer red triangles.

Suppose $\delta(v, \Gamma) \geq 3$. Since $\delta_G(v, \Gamma) = 0$, the vertex $v \notin \partial(\Gamma)$, so the loop-minimality of Γ ensures that there are no loops based at v . Hence by Lemma 3.15, we can replace Γ by a diagram Γ_1 with boundary word w and the same green faces as Γ , such that $\text{CArea}(\Gamma_1) \leq \text{CArea}(\Gamma)$ and $\delta_R(v, \Gamma_1) \leq \delta_R(v, \Gamma) - 1$, and in which there are no loops




Figure 2: Triangles with two or three coincident vertices

based at v . By repeating this process, we eventually reduce to a diagram Γ_k in which $\delta(v, \Gamma_k) = 2$, and the coloured area can then be reduced, yielding Δ . This completes the proof in the case that no red triangle of Γ meets itself at one or more vertices.

Now we allow for the possibility that not all vertices of a red triangle are distinct. Figure 2 shows the possible configurations of a red triangle in Γ in which two or more of the vertices coincide. The red triangle is labelled T and other regions of the diagram, which may contain additional vertices, edges, and faces, are labelled Θ , Θ_1 or Θ_2 .

In the third of these configurations, two consecutive letters of a word in V_P σ -cancel, which contradicts the definition of V_P . In the second and the fourth, there is an internal loop labelled by a single letter from V_P , contradicting our assumption of loop-minimality.

It therefore remains only to consider the case where the whole of Γ has the structure of the first diagram in Figure 2. Here, the sub-diagram labelled Θ has no loops labelled by V^σ -letters, and so Θ can have no red triangles with coincident vertices. The boundary word of Θ is σ -reduced, by definition of V_P , and so, by the arguments above, we may assume that all vertices of Θ have green degree at least 1. But we have to consider the possibility that the only green face of Θ that is incident with the vertex labelled v in Γ is the external face of Θ , which is coloured red as a face of Γ . We shall now show that this case does not occur.

If $\delta(v, \Gamma) = 2$, then the face of Θ incident with v is a red triangle, and so Θ contains a loop labelled by a V^σ -letter, contradicting our assumption of loop-minimality. If $\delta(v, \Gamma) \geq 3$, there are three consecutive red triangles incident with v . Since Γ has no loops based at v , we can repeatedly apply Lemma 3.15 to v to reduce its degree, and hence reduce the coloured area, reaching a diagram Γ' in which $\delta(v, \Gamma') = 2$. This yields the same contradiction to our assumption of loop-minimality as at the beginning of this paragraph. \square

Proposition 3.17. *Let Γ be a loop-minimal coloured van Kampen diagram with cyclically P -reduced boundary word w of length n and let $r = \max\{|R| : R \in \mathcal{R}\}$. Then there exists a loop-minimal coloured diagram Δ with boundary word w , such that $\text{Area}(\Delta) \leq (3 + r)|F_G(\Delta)| + n - 2$. If every vertex v of Γ satisfies $\delta_G(v, \Gamma) \geq 1$ then Δ can be taken to be Γ .*

Proof. By Theorem 3.16, there exists a loop-minimal coloured diagram Δ with boundary word w and the same green faces as Γ , in which every vertex v satisfies $\delta_G(v, \Delta) - 1 \geq 0$. Hence for Δ , Proposition 3.11 gives the inequality

$$|F_R(\Delta)| \leq \sum_{e \in E(\Delta)} \delta_G(e, \Delta) - 2 + 2|F_G(\Delta)|.$$

Each contribution to $\delta_G(e, \Delta)$ comes either from the boundary of the external face (of length n), or from an internal green face (of which there are $|F_G(\Delta)| = |F_G(\Gamma)|$, each of boundary length at most r), so we deduce that

$$|F_R(\Delta)| \leq n + r|F_G(\Delta)| - 2 + 2|F_G(\Delta)|,$$

and $\text{Area}(\Delta) = |F_R(\Delta)| + |F_G(\Delta)| \leq (3+r)|F_G(\Delta)| + n - 2$, as claimed. \square

4 Interleaving the green relators

In Subsection 4.1 we shall generalise interleaving (Definition 2.7) to cyclic interleaving, and show that this gives an equivalence relation on cyclically P -reduced words. Then in Subsection 4.2 we shall prove our main result in this section, Proposition 4.10. This shows that if we replace \mathcal{R} by the (finite) set $\mathcal{I}(\mathcal{R})$ of all cyclic interleaves of elements of \mathcal{R} , then a cyclically P -reduced word w is equal to 1 in $G = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ if and only if some cyclic interleave of w is the boundary of a coloured diagram over $\langle X^\sigma \mid V_P \mid \mathcal{I}(\mathcal{R}) \rangle$ in which each vertex is incident with at least two green faces. Finally, in Subsection 4.3 we shall study the regions of coloured diagrams that are composed entirely of red triangles.

We remind the reader that all newly defined terms and notation are listed in the Appendix.

4.1 Cyclic interleaving

Recall Definition 2.7 of interleaving, and that by Theorem 2.9 this yields an equivalence relation on P -reduced words, with one equivalence class for each element of $U(P)$.

Lemma 4.1. *Let $v = x_1x_2 \cdots x_n \in X^*$ be cyclically P -reduced. If $v \approx w$, then w is cyclically P -reduced.*

Proof. If $n \leq 1$ then the result is trivial, so assume that $n > 1$, and let $s_0 = 1, s_1, \dots, s_n = 1$ be the elements of P from Definition 2.7.

By the transitivity of \approx , it suffices to consider a single rewrite, so assume that s_i is nontrivial for a single i . The result is immediate unless $i = 1$ or $i = n - 1$, so assume without loss of generality that $s_1 \neq 1$.

Suppose first that $n \geq 3$. Then $x_nx_1x_2 \cdots x_{n-1}$ is P -reduced by assumption, and $x_nx_1x_2 \cdots x_{n-1} \approx x_n[x_1s_1][s_1^\sigma x_2] \cdots x_{n-1}$. By Theorem 2.9(i), the word $x_n[x_1s_1][s_1^\sigma x_2] \cdots x_{n-1}$ is P -reduced, so $(x_n, [x_1s_1]) \notin D(P)$, which proves the result.

Otherwise, $n = 2$ and $w = [x_1s_1][s_1^\sigma x_2]$. We need to show that $([s_1^\sigma x_2], [x_1s_1]) \notin D(P)$, so assume the contrary. We apply Axiom (P5) with

$$(x, y, z, t) = (s_1, s_1^\sigma x_2, x_1s_1, s_1^\sigma)$$

and conclude that at least one of $(x_2, x_1s_1), (s_1^\sigma x_2, x_1) \in D(P)$. Suppose that $(x_2, x_1s_1) \in D(P)$ (the other case is similar). Applying (P5) again with $(x, y, z, t) = (x_2, x_1s_1, s_1^\sigma, x_2)$ gives $(x_1, x_2) \in D(P)$ or $(x_2, x_1) \in D(P)$, contradicting v being cyclically P -reduced. \square

We now define a coarser relation than \approx , by allowing the elements s_0 and s_n from Definition 2.7 to be nontrivial but equal.

Definition 4.2. Let $v = x_1 \cdots x_n \in X^*$ be cyclically P -reduced, and let $w = y_1 \cdots y_m \in X^*$. Then we write $v \approx^c w$ if $m = n$ and either $n \leq 1$ and $v = w$, or $n > 1$ and there exist $s_0, s_1, \dots, s_{n-1}, s_n = s_0 \in P$ such that $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i) \in D(P)$ for $1 \leq i \leq n$ and $y_i = [s_{i-1}^\sigma x_i s_i]$. We say that w is a *cyclic interleave* of v .

Example 4.3. Consider the pregroup P from Example 2.8, and let $v = a_1b_1a_3b_5 \in U(P)$. Notice that since $(b_5, a_1) \notin D(P)$ the word v is cyclically P -reduced.

Then one cyclic interleave w of v can be made by setting $s_0 = s_1 = s_4 = i_2$ and $s_2 = s_3 = i_4$. With this choice of interleaving elements we find that $w = a_1b_3a_3b_3$.

Theorem 4.4. *Let $v = x_1 \cdots x_n \in X^*$ be cyclically P -reduced. If $w = y_1 \cdots y_n \approx^c v$, then w is cyclically P -reduced. Furthermore, \approx^c is an equivalence relation on the set of all cyclically P -reduced words.*

Proof. Since \approx^c is the identity relation on words of length at most 1, we may assume without loss of generality that $n > 1$.

We can move from v to w by a sequence of single rewrites. By Lemma 4.1, a single rewrite with $i \neq n$ replaces v by another cyclically P -reduced word and, by applying the lemma to a cyclic permutation of v , we see that the same applies when $i = n$. So w is cyclically P -reduced.

To show that \approx^c is an equivalence relation, it is sufficient to prove that if $w \approx^c u$, where $u = z_1 \cdots z_n$ is the result of applying a single rewrite to w , then $v \approx^c u$.

Suppose first that this rewrite consists of replacing (y_n, y_1) by $([y_n t], [t^\sigma y_1])$ for some $t \in P$, so that $u = [t^\sigma y_1] y_2 \cdots y_{n-1} [y_n t]$. Let $s_0, s_1, \dots, s_{n-1}, s_n = s_0 \in P$ be as in Definition 4.2, cyclically interleaving v to w . Then $y_n = [s_{n-1}^\sigma x_n s_n]$ and $y_1 = [s_n^\sigma x_1 s_1]$. Since w is cyclically P -reduced, $y_n y_1$ is P -reduced, and hence $(s_n, t) \in D(P)$ by [22, 3.A.2.6]. So, putting $s'_i = s_i$ for $1 \leq i < n$ and $s'_0 = s'_n = [s_n t]$, we have $z_i = s_{i-1}'^\sigma x_i s'_i$ for $0 \leq i \leq n$ and $v \approx^c u$ as claimed.

The argument in the case when the single rewrite from w to u consists of the replacement of (y_i, y_{i+1}) with $1 \leq i < n$ is similar. \square

Definition 4.5. Let $w \in X^*$ be cyclically P -reduced. We denote by $\mathcal{I}(w)$, called the *cyclic interleave class* of w , the set

$$\mathcal{I}(w) = \{v \in X^* : v \approx^c w\}.$$

We record the following easy lemma for later use.

Lemma 4.6. *Let $v = x_1 \cdots x_n \in X^*$ be cyclically P -reduced with $n > 1$, and suppose that $w \in \mathcal{I}(v)$. Then we can obtain w from v by applying a sequence of at most n single rewrites.*

Proof. By definition of the cyclic interleave class of v , there are elements y_1, \dots, y_n and $s_0, s_1, \dots, s_n = s_0$ of X with $w = y_1 \cdots y_n$ and $y_i = [s_{i-1}^\sigma x_i s_i]$ for $1 \leq i \leq n$. By Theorem 4.4, the relation \approx^c is transitive, so we can introduce the s_i individually (and in any order) as single rewrites, where each such rewrite consists of replacing the current word $y'_1 y'_2 \cdots y'_n$ by either $y'_1 \cdots y'_i y''_{i+1} \cdots y'_n$ with $y''_i = y'_i s_i$ and $y''_{i+1} = s_i^\sigma y'_{i+1}$ for some i with $1 \leq i < n$, or by $y'_1 y'_2 \cdots y'_{n-1} y''_n$ with $y''_1 = s_0^\sigma y_1$ and $y''_n = y'_n s_0$. \square

Definition 4.7. We write $\mathcal{I}(\mathcal{R})$ for $\cup_{R \in \mathcal{R}} \mathcal{I}(R)$, and let $\mathcal{I}(\mathcal{P})$ be the pregroup presentation

$$\langle X^\sigma \mid V_P \mid \mathcal{I}(\mathcal{R}) \rangle.$$

Note in particular that, since X and \mathcal{R} are finite, so is $\mathcal{I}(\mathcal{R})$. Faces of diagrams over $\mathcal{I}(\mathcal{P})$ are coloured red and green, just as for coloured diagrams over \mathcal{P} .

Theorem 4.8. *The normal closure of $\langle V_P \cup \mathcal{R} \rangle$ in $F(X^\sigma)$ is equal to the normal closure of $\langle V_P \cup \mathcal{I}(\mathcal{R}) \rangle$ in $F(X^\sigma)$. Hence, G is defined by \mathcal{P} if and only if G is defined by $\mathcal{I}(\mathcal{P})$.*

Proof. One containment is clear, since $\mathcal{R} \subset \mathcal{I}(\mathcal{R})$. For the other, by Lemma 4.6 each element of $\mathcal{I}(\mathcal{R})$ is made by applying a finite number of single rewrites to each $R = x_1 \cdots x_n \in \mathcal{R}$, and so is conjugate to a word that is equal in $U(P)$ to an element of \mathcal{R} . The final statement follows from Corollary 2.15. \square

As a result of the above theorem, we shall move between working with a presentation \mathcal{P} and a presentation $\mathcal{I}(\mathcal{P})$ without further comment.

4.2 Green-rich diagrams

Recall Definitions 3.1 and 3.2 and in particular our conventions on counting incidence, and on the boundary of diagrams. The following definition will be used throughout the rest of the paper, as it is key to our methods of proving hyperbolicity.

Definition 4.9. A coloured diagram Γ in which each vertex v satisfies $\delta_G(v, \Gamma) \geq 2$ is *green-rich*.

The following is our main result in this section.

Proposition 4.10. *Let Γ be a loop-minimal coloured diagram over $\mathcal{I}(\mathcal{P})$ with cyclically σ -reduced boundary word w .*

- (i) *Assume first that w is cyclically P -reduced and that, if $|w| = 1$, then the unique boundary face of Γ is green. Then some $w' \in \mathcal{I}(w)$ is the boundary word of a green-rich loop-minimal coloured diagram Δ over $\mathcal{I}(\mathcal{P})$ with $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$.*
- (ii) *If, instead, $|w| > 1$ and w is not cyclically P -reduced, then w is the boundary word of a loop-minimal coloured diagram Δ over $\mathcal{I}(\mathcal{P})$ with $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$, in which all non-boundary vertices v satisfy $\delta_G(v, \Delta) \geq 2$.*

In both cases, if Γ is not green-rich, then $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$.

Proof. The assumptions of loop-minimality, and that if $|w| = 1$ then the unique boundary face of Γ is green, together imply that Γ does not have a unique boundary face that is red. The loop-minimality then implies that the vertices of each red triangle of Γ are distinct, and that if $|w| = 1$, with v_0 the unique boundary vertex, then $\delta_G(v_0, \Gamma) \geq 3$.

By Theorem 3.16 we may assume that every vertex v of Γ satisfies $\delta_G(v, \Gamma) \geq 1$. Assume that $\delta_G(v, \Gamma) = 1$, and that if $v \in \partial(\Gamma)$ then w is cyclically P -reduced. Let ab be the length two subword of the boundary label of the unique green face f that is incident with v , so that v is between a and b . We shall construct a coloured diagram Δ with boundary word $w' \in \mathcal{I}(w)$, in which v no longer exists, and such that $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$. The loop-minimality of Γ implies that Δ is loop-minimal.

If v is incident with a unique red triangle T , then $\delta(v, \Gamma) = 2$, which implies that $b^\sigma a^\sigma$ is a subword of the boundary label of T , and so $(b^\sigma, a^\sigma) \in D(P)$. Elements of \mathcal{R} (and by Theorem 4.4 also of $\mathcal{I}(\mathcal{R})$) are cyclically P -reduced, so $v \in \partial(\Gamma)$. But in this case w was assumed to be cyclically P -reduced, a contradiction.

If v is incident with exactly two red triangles, then they must be distinct and share an edge. Let the third edge incident with v be labelled c , so that the triangles have labels with subwords $c^\sigma a^\sigma$ and $b^\sigma c$. Then (a, c) and (c^σ, b) are also in $D(P)$, and $([ac], [c^\sigma b]) \notin D(P)$ by Theorem 4.4. A single rewrite can therefore be applied to the label of f , replacing ab with $[ac][c^\sigma b]$. This has the effect of replacing f by a face labelled with a cyclic interleave of the label of f , removing the vertex v and its two incident red triangles from Γ , and leaving the number of green faces unchanged.

Assume finally that $\delta_R(v, \Gamma) \geq 3$. Since v is not a boundary vertex of a diagram with boundary length 1, the loop-minimality of Γ means that there are no loops based at v . Hence we can repeatedly apply Lemma 3.15 to make a diagram Γ' in which $\delta_R(v, \Gamma') = 2$, and then delete v and the final two red triangles incident with it, as in the previous paragraph. \square

In what follows, we shall often assume that we work with green-rich van Kampen diagrams. But, since we have not been able to eliminate the existence of non-green-rich loop-minimal diagrams with boundary length 1 and the unique boundary face a red triangle, we need to take that possibility into account when developing algorithms, as we shall do in Theorem 6.12.

4.3 Red blobs

We now turn our attention to the regions of coloured diagrams that are comprised entirely of red triangles.

Definition 4.11. A *red blob* in a coloured diagram Γ is a nonempty subset B of the set of closed red triangles of Γ , with the property that any nonempty proper subset C of B has at least one edge in common with $B \setminus C$. Equivalently, the induced subgraph \bar{B} of the dual graph $\bar{\Gamma}$ of Γ on those vertices that correspond to the triangles in B is connected.

A red blob is *simply connected* if its interior is homeomorphic to a disc: its boundary may pass more than once through a vertex.

Notice that a simply-connected red blob corresponds to a van Kampen diagram over $U(P)$, and so in particular has boundary length greater than one, by Theorem 2.9. (However, not all van Kampen diagrams over $U(P)$ correspond to red blobs, since the interior of such diagrams may be disconnected).

Lemma 4.12. *Let B be a red blob in a coloured diagram Γ , with boundary length l and area t . Then $l \leq t + 2$, and $l \leq t$ if B is not simply connected. Furthermore, if B is simply connected, and every vertex of B lies on $\partial(B)$ (which holds in particular when Γ is green-rich), then $l = t + 2$.*

Proof. Let \bar{B} be the induced subgraph of $\bar{\Gamma}$ that corresponds to the triangles in B , as in Definition 4.11. A vertex of degree 1 in \bar{B} corresponds to a triangle in B that has two edges on $\partial(B)$. Deleting this triangle reduces both the number of boundary edges and the number of triangles by 1, and \bar{B} remains connected. We repeatedly remove degree 1 vertices from \bar{B} until none remain. At that stage, either a single vertex remains, in which case $l = t + 2$, or the remaining vertices all have degree at least two. Such vertices correspond to triangles with at most one edge on $\partial(B)$, so the number of triangles is at least the boundary length, and $l \leq t$.

If B is not simply connected, then \bar{B} contains a circuit, and so the second of the above two situations arises, and $l \leq t$.

Now assume that all vertices of B lie on $\partial(B)$, and B is simply connected. We shall show that \bar{B} is a tree, from which the final claim follows. By way of contradiction, let C be a circuit in \bar{B} . It is a standard result from graph theory (see, for example, [24, Corollary 4.15]) that the corresponding edges in B form a cutset in B . Hence the circuit must enclose at least one vertex of B . We have assumed that all vertices of B lie on $\partial(B)$, so this contradicts the fact that B is simply connected. \square

Proposition 4.13. *Let Γ be a loop-minimal coloured diagram with cyclically σ -reduced boundary word w . Then there exists a loop-minimal coloured diagram Δ with boundary word w , with $\text{CArea}(\Delta) \leq \text{CArea}(\Gamma)$, and in which the (cyclic) boundary word of each simply connected red blob has no proper subword equal to 1 in $U(P)$. If Γ is green-rich then Δ is green-rich. Furthermore, if Γ does not have the required property already, then $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$.*

Proof. Let B be a simply connected red blob in Γ with boundary word $w = x_1 \dots x_n$. First note that by Theorem 3.16, we may assume that all vertices of Γ have green degree at least one. Hence all vertices of B lie on $\partial(B)$, and $\text{Area}(B) = n - 2$ by Lemma 4.12.

We first consider σ -reduction, so assume (without loss of generality) that $w = x_1 x_1^\sigma w_1$. Notice that $w =_{U(P)} 1$, so $w_1 =_{U(P)} 1$. Hence we can identify the vertices at the beginning of the edge labelled x_1 and the end of the edge labelled x_1^σ , and replace B by a coloured sub-diagram Θ consisting of a red blob B_1 with boundary word w_1 , with a single edge added to the boundary. The blob B_1 is simply connected with all vertices on the boundary,

and $|\partial(B_1)| = |w_1| = n - 2$, so $\text{Area}(B_1) = n - 4$ by Lemma 4.12. The diagram Δ in which B has been replaced by Θ satisfies $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$, so Δ is loop-minimal. Replacing B by Θ cannot decrease the green degrees of vertices, so if Γ is green-rich then Δ is green-rich.

Now we consider subwords of w of length greater than 2. Assume that w has a factorisation $w = w_1 w_2$ such that w_1 and w_2 have lengths $b_1, b_2 \geq 3$, and $w_1 =_{U(P)} 1 =_{U(P)} w_2$. We can produce a new diagram Δ in which the two vertices where w_1 and w_2 start and end are identified, and B has been replaced by two red blobs B_1 and B_2 with boundary words w_1 and w_2 , of area $b_1 - 2$ and $b_2 - 2$, respectively. From $b_1 + b_2 = n$, we see that $\text{Area}(B_1) + \text{Area}(B_2) = n - 4$, so $\text{CArea}(\Delta) < \text{CArea}(\Gamma)$, and hence Δ is loop-minimal. As before, if Γ was green-rich then Δ is still green-rich. \square

Definition 4.14. We say that $a \in X$ *intermults with* $b \in X$ if $b \neq a^\sigma$ and either $(a, b) \in D(P)$ or there exists $x \in X$ such that $(a, x), (x^\sigma, b) \in D(P)$. We also say that (a, b) is an *intermult pair*.

Example 4.15. If P is the natural pregroup for a free product (with no amalgamation), constructed as in Example 2.5, then the intermult pairs are precisely the non-inverse pairs of non-identity elements contained within a free factor.

If P is instead the natural pregroup for a free product with non-trivial amalgamation, then the intermult pairs are all non-inverse pairs of non-identity elements.

In the following lemma, we do not assume that the red blob is simply connected: it may therefore have more than one boundary word. The following lemma will be used in the algorithmic part of this paper, to reduce the number of possible boundary words of red blobs.

Lemma 4.16. If $g, a \in X$ and ga is a subword of a boundary word of a red blob B with σ -reduced boundary words, then g intermults with a .

Proof. First notice that the assumption that B has σ -reduced boundary words shows that $g \neq a^\sigma$.

Let v be the vertex between g and a on the boundary of B . Reading clockwise around v from a as far as g , let the labels of the outgoing edges be $a = a_1, a_2, \dots, a_k, g^\sigma = a_{k+1}$ (the outgoing label is g^σ not g). Notice that the edges labelled a_2, \dots, a_k are all interior to B , since ga is a subword of a boundary word. Notice also that

$$(a_{i+1}^\sigma, a_i) \in D(P) \quad \text{for } 1 \leq i \leq k. \quad (1)$$

We shall show by induction that (1) implies that $a_{k+1} = g$ intermults with $a_1 = a$.

If $k = 1$ then $(a_2^\sigma, a_1) = (g, a) \in D(P)$. If $k = 2$ then $(a_2^\sigma, a), (g, a_2) \in D(P)$, so (g, a) is an intermult pair. Assume that $k \geq 3$, and for $1 \leq i \leq 3$ let $b_i = [a_{i+1}^\sigma a_i]^\sigma$, so that the boundary label of each triangle is $a_{i+1}^\sigma a_i b_i$. Then applying Axiom (P5) to $(b_1, a_2^\sigma), (a_2^\sigma, a_3), (a_3, b_3) \in D(P)$ shows that at least one of $([b_1 a_2^\sigma], a_3), (a_2^\sigma, a_3 b_3) \in D(P)$. Since $b_1 a_2^\sigma = a_1^\sigma$ and $a_3 b_3 = a_4$, at least one of $(a_3^\sigma, a_1), (a_4^\sigma, a_2) \in D(P)$, so the result follows by induction. \square

5 Curvature distribution schemes

In this section, we introduce the concept of curvature distribution schemes, and prove that they can be used to show that groups given by a pregroup presentation satisfy an explicit linear isoperimetric inequality, and hence are hyperbolic.

Definition 5.1. Let Γ be a coloured van Kampen diagram with vertex set $V(\Gamma)$, edge set $E(\Gamma)$, set of red triangles $F_R(\Gamma)$ and set of internal green faces $F_G(\Gamma)$. Let $F(\Gamma) = F_R(\Gamma) \cup F_G(\Gamma)$. A *curvature distribution* is a function $\rho_\Gamma : V(\Gamma) \cup E(\Gamma) \cup F(\Gamma) \rightarrow \mathbb{R}$ such that

$$\sum_{x \in V(\Gamma) \cup E(\Gamma) \cup F(\Gamma)} \rho_\Gamma(x) = 1.$$

Definition 5.2. Let \mathcal{K} be a set of coloured diagrams over $\mathcal{I}(\mathcal{P})$. A *curvature distribution scheme* on \mathcal{K} is a map $\Psi : \mathcal{K} \rightarrow \{\rho_\Gamma : \Gamma \in \mathcal{K}\}$, that associates a curvature distribution to every diagram in \mathcal{K} .

Example 5.3. For any coloured diagram Γ we can define a curvature distribution by setting $\rho_\Gamma(v) := +1$ for each vertex v , setting $\rho_\Gamma(e) := -1$ for each edge e , and setting $\rho_\Gamma(f) := +1$ for each internal face f . Euler's formula ensures that the total sum of all curvature values is $+1$. Since this defines a curvature distribution for every diagram, it gives rise to a curvature distribution scheme on \mathcal{K} , where \mathcal{K} is any set of coloured diagrams over $\mathcal{I}(\mathcal{P})$, for any pregroup presentation \mathcal{P} .

Definition 5.4. Let Γ be a plane graph, and let $\bar{\Gamma}$ be its dual. Let f_1 and f_2 be faces of Γ , corresponding to vertices v_1 and v_2 of $\bar{\Gamma}$. The *dual distance* in Γ from f_1 to f_2 is the distance in $\bar{\Gamma}$ from v_1 to v_2 .

Definition 5.5. The *pregroup Dehn function* $\text{PD}(n) : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{Z}$ of a pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ is defined as follows. For each σ -reduced word $w \in X^*$ with $w =_G 1$, let $A(w)$ be the smallest area of a coloured diagram over \mathcal{P} with boundary label w . Then $\text{PD}(n) := \max\{A(w) : w \in X^*, w =_G 1, |w| \leq n\}$.

$\text{PD}(n)$ may differ from the standard Dehn function $\text{D}(n)$ of a corresponding group presentation, because faces of standard van Kampen diagrams labelled by relators x^2 , corresponding to generators $x \in X$ with $x = x^\sigma$, are not counted. To bound $\text{D}(n)$ in terms of $\text{PD}(n)$, we need to fix a corresponding group presentation.

Definition 5.6. Let $\mathcal{P} = \langle X \mid V_P \mid \mathcal{R} \rangle$. Let Y be a minimal subset of X such that $X = Y \cup Y^\sigma$ (so Y generates \mathcal{P} as a group). Let subsets V'_P and \mathcal{T} of $F(Y)$ be constructed from V_P and \mathcal{R} , respectively, by replacing all symbols $x \in X \setminus Y$ by $(x^\sigma)^{-1}$. Then the *standard group presentation corresponding to \mathcal{P}* is $\mathcal{P}_G := \langle Y \mid \{x^2 : x \in Y, x = x^\sigma\} \cup V'_P \cup \mathcal{T} \rangle$.

Example 5.7. Let $\mathcal{P} = \langle x, y, z \mid y^3, z^3 \mid (xz)^7, (xyxz)^4 \rangle$ with $x^\sigma = x$ and $y^\sigma = z$. Then choosing $Y = \{x, y\}$ gives $\mathcal{P}_G = \langle x, y \mid x^2, y^3, y^{-3}, (xy^{-1})^7, (xyxy^{-1})^4 \rangle$ (where we could of course omit the redundant relator y^{-3}).

The following bound is not at all tight, but suffices to show that if $\text{PD}(n)$ is linear then so is $\text{D}(n)$.

Lemma 5.8. Let $\text{PD}(n)$ and $\text{D}(n)$ be the pregroup and standard Dehn functions of \mathcal{P} and \mathcal{P}_G , respectively. Let r_I be the maximum number of involutory generators appearing in any $R \in V_P \cup \mathcal{R}$. Then $\text{D}(n) \leq r_I \text{PD}(n) + n/2$.

Proof. To change a coloured diagram into a standard van Kampen diagram, we first replace each edge label from $X \setminus Y$ by the inverse of the corresponding element of Y . This produces a diagram that is almost a van Kampen diagram, except that involutions $x \in P$ may appear on both sides of an edge. But we can rectify that as follows. For $R \in V_P \cup \mathcal{R}$, let R_I denote the number of involutory generators occurring in R (with multiplicity). Then, in a diagram Γ over $F(X^\sigma)$, a face previously labelled by $R^{\pm 1}$ needs

to have R_I digons (with boundary labels of the form x^2) added to its boundary to correct the edge labels. There are at most $\text{PD}(n)$ faces in Γ , and at most n boundary edges. Each boundary edge is incident either with an internal face, or twice with the external face, so the result follows. \square

The following theorem is one of the key results in this paper. It appears technical, but the insight behind it is straightforward. We shall show how curvature distribution schemes give sufficient conditions for the area of a van Kampen diagram Γ over \mathcal{P} to be bounded by a multiple of the boundary length n : this is our generalisation of small cancellation.

We first ensure that all of the positive curvature is associated with the green faces of Γ , and that there is a fixed upper bound m on the curvature of any face. Then we ensure that the green faces at dual distance greater than d from the boundary in fact have curvature bounded above by $-\varepsilon$, for some fixed d and fixed $\varepsilon > 0$. Notice that the number of green faces at dual distance at most d from $\partial(\Gamma)$ is bounded by a function of n and the length of the longest relator, and this function is linear in n . Hence, since the total curvature must sum to 1, the total number of green faces is bounded by a linear multiple of n . If the number of red faces is bounded linearly in terms of the number of green faces (for example, if Γ satisfies the conditions of Proposition 3.17), then the desired proof of linear area follows.

Theorem 5.9. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation for a group G , and let r be the maximum length of a relator in \mathcal{R} . Let \mathcal{K} be a set of coloured van Kampen diagrams over $\mathcal{I}(\mathcal{P})$, and let $\Psi : \mathcal{K} \rightarrow \{\rho_\Gamma : \Gamma \in \mathcal{K}\}$ be a curvature distribution scheme.*

Assume that there exist constants $\varepsilon \in \mathbb{R}_{>0}$, $\lambda, \mu, m \in \mathbb{R}_{\geq 0}$ and $d \in \mathbb{Z}_{>0}$ such that the following conditions hold, for all $\Gamma \in \mathcal{K}$:

- (a) $\rho_\Gamma(x) \leq 0$ for all $x \in V(\Gamma) \cup E(\Gamma) \cup F_R(\Gamma)$,
- (b) $\rho_\Gamma(f) \leq -\varepsilon$ for all faces $f \in F_G(\Gamma)$ that are dual distance at least $d+1$ from the external face,
- (c) if $\text{Area}(\Gamma) > 1$, then $\rho_\Gamma(f) \leq m$ for all faces $f \in F_G(\Gamma)$ that are dual distance at most d from the external face,
- (d) $\text{Area}(\Gamma) \leq \lambda|F_G(\Gamma)| + \mu|\partial(\Gamma)|$.

Then each $\Gamma \in \mathcal{K}$ with boundary length n and area greater than 1 satisfies

$$\text{Area}(\Gamma) \leq f(n) = \lambda \left(n \frac{(r-1)^d - 1}{r-2} \left(1 + \frac{m}{\varepsilon} \right) - \frac{1}{\varepsilon} \right) + \mu n. \quad (2)$$

Assume now that, in addition, the following holds:

- (e) *if $w \in X^*$ is cyclically P -reduced, and satisfies $w =_G 1$, then there exists a diagram $\Gamma \in \mathcal{K}$ with boundary word some $w' \in \mathcal{I}(w)$.*

Then the group G is hyperbolic. In particular, if $\mathcal{I}(w) = \{w\}$ for all $w \in X^$, then the pregroup Dehn function of \mathcal{P} is bounded above by $\max\{f(n), 1\}$.*

Proof. We show first that Equation (2) holds. Let $\Gamma \in \mathcal{K}$ have boundary length n , and let $F := F(\Gamma)$ and $F_G := F_G(\Gamma)$. If $\text{Area}(\Gamma) = 1$, then Equation (2) does not apply, so assume that $\text{Area}(\Gamma) > 1$.

Let I be the set of green faces that are dual distance at least $d+1$ from the external face (the set I may be empty). From Condition (a) we deduce that $\sum_{f \in F_G} \rho_\Gamma(f) = \sum_{f \in I} \rho_\Gamma(f) + \sum_{f \in F_G \setminus I} \rho_\Gamma(f) \geq 1$. Combinatorial considerations show that

$$|F_G \setminus I| \leq n + n(r-1) + n(r-1)^2 + \cdots + n(r-1)^{d-1} = n \frac{(r-1)^d - 1}{r-2}.$$

Condition (b) yields $\sum_{f \in I} \rho_\Gamma(f) \leq -\varepsilon|I|$ and then applying Condition (c) (since $\text{Area}(\Gamma) > 1$), we deduce that

$$\varepsilon|I| \leq -\sum_{f \in I} \rho_\Gamma(f) \leq \sum_{f \in F_G \setminus I} \rho_\Gamma(f) - 1 \leq mn \frac{(r-1)^d - 1}{r-2} - 1.$$

From this we get $|I| \leq \frac{1}{\varepsilon} \left(mn \frac{(r-1)^d - 1}{r-2} - 1 \right)$, and so by Condition (d) we see that

$$\begin{aligned} |F| &\leq \lambda(|I| + |F_G \setminus I|) + \mu n \\ &\leq \lambda \left(\frac{1}{\varepsilon} \left(mn \frac{(r-1)^d - 1}{r-2} - 1 \right) + n \frac{(r-1)^d - 1}{r-2} \right) + \mu n \end{aligned}$$

and Equation (2) follows.

Now assume that Condition (e) also holds. Each single rewrite of the boundary word of a diagram Γ adds two red triangles to the diagram, as in the proof of Proposition 4.10. It therefore follows from Lemma 4.6 that, if there is a diagram Γ with boundary label a word w of length n then, for any $w' \in \mathcal{I}(w)$, there is a coloured diagram of area at most $\text{Area}(\Gamma) + 2n$ with boundary label w' . So there is a linear upper bound on the pregroup Dehn function, and hence by Lemma 5.8 the Dehn function, of G . The remaining assertions now follow. \square

In the curvature scheme that we shall study in the remainder of this paper, we shall generally set $d = 1$, and we shall prove that $m = 1/2$.

Corollary 5.10. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation for a group G , such that each $x \in X$ is nontrivial in G . Let r be the maximum length of a relator in \mathcal{R} , let \mathcal{K} contain all diagrams over \mathcal{P} of minimal coloured area for each cyclically P -reduced word w that is trivial in G , and let Ψ be a curvature distribution scheme on \mathcal{K} .*

If there exists $\varepsilon > 0$ such that Conditions (a), (b) and (c) of Theorem 5.9 hold, with $m = 1/2$ and $d = 1$, then G is hyperbolic, and the pregroup Dehn function of \mathcal{P} is bounded above by

$$n \left(4 + r + \frac{3+r}{2\varepsilon} \right) - \frac{3+r}{\varepsilon}.$$

Proof. Notice that since each $x \in X$ is nontrivial in G , all diagrams are loop-minimal. By Theorem 3.16, each diagram Γ of minimal coloured area for its boundary word satisfies $\delta_G(v, \Gamma) \geq 1$ for each vertex v . Hence by Proposition 3.17 each diagram of minimal coloured area for its boundary word satisfies Condition (d) of Theorem 5.9, with $\lambda = 3+r$ and $\mu = 1$. Substituting for λ, μ, m and d into (2) yields

$$f(n) = (3+r) \left(n \left(1 + \frac{1}{2\varepsilon} \right) - \frac{1}{\varepsilon} \right) + n.$$

This gives an upper bound on the area of all diagrams whose area is greater than 1. The assumption that each $x \in X$ is nontrivial in G implies that any diagram of area 1 has $n \geq 2$, and one may check that if $n \geq 2$ then $f(n) \geq 1$. Hence $f(n)$ bounds the pregroup Dehn function of G . \square

In general, it is not practical to let \mathcal{K} consist *only* of diagrams of minimal coloured area for their boundary word, as membership of \mathcal{K} cannot easily be tested. We shall however define in the next section a useful set of diagrams with an easily-testable membership condition. We shall also deal with the condition in the above corollary that each generator is nontrivial in the group G .

The remainder of this paper presents and analyses one curvature distribution scheme, chosen because it can be tested in time that is bounded by a low-degree polynomial function of $|X|$, $|\mathcal{R}|$ and r , and because it verifies that V^σ -letters are nontrivial in G . There are, of course, infinitely many possible such schemes, and we leave as an open problem the development of others that are also computationally or theoretically useful.

6 The RSym scheme

In this section we describe a curvature distribution scheme that treats each vertex and each edge of each diagram symmetrically, and so is called the RSym scheme. We first specify the set \mathcal{D} of diagrams on which RSym operates.

We remind the reader that all definitions and notation are recorded in the Appendix.

Definition 6.1. Let \mathcal{P} be a pregroup presentation. Then \mathcal{D} denotes the set of all coloured diagrams Γ over $\mathcal{I}(\mathcal{P})$ with the following properties:

1. the boundary word of Γ is cyclically P -reduced (see Definition 2.6);
2. Γ is σ -reduced and semi- P -reduced (see Definitions 3.6 and 3.7);
3. Γ is green-rich (see Definition 4.9);
4. no proper subword of the boundary word of a simply connected red blob in Γ is equal to 1 in $U(P)$.

Recall Definitions 3.1 and 3.2 for our conventions on coloured diagrams.

Definition 6.2. In a coloured diagram, we shall consider each edge to be composed of two coloured *half-edges*, oppositely oriented. Each half-edge is associated with the face on that side, and inherits its colour and orientation from that face.

The following algorithm, **ComputeRSym**, takes as input a diagram $\Gamma \in \mathcal{D}$, and returns a curvature distribution $\kappa_\Gamma : \Gamma \rightarrow \mathbb{R}$. The algorithm assigns and alters curvature on the vertices, edges and faces of Γ in several successive steps: the external face has curvature 0 throughout. In the algorithm description, when we say (for example) that a half-edge e gives curvature c to vertex v , we mean that the curvature of e is reduced by c , and that of v is increased by c . When we say that a vertex v distributes its curvature equally among green faces f_1, \dots, f_k , we mean that, if $k > 0$, then the current curvature c of v is replaced by 0, and c/k is added to the curvature of each of f_1, \dots, f_k .

Algorithm 6.3. **ComputeRSym**(Γ):

- Step 1 Initially, each vertex, red triangle, and internal green face of Γ has curvature $+1$, and each half-edge has curvature $-1/2$.
- Step 2 Each green half-edge gives curvature $-1/2$ to its end vertex, and each red half-edge gives curvature $-1/2$ to its triangle.
- Step 3 Each vertex distributes its curvature equally amongst its incident internal green faces, counting incidences with multiplicity.
- Step 4 Each red blob B such that $\partial(B) \not\subseteq \partial(\Gamma)$ sums the curvatures of its red triangles, to get the *blob curvature* $\beta(B)$. A red blob with $b := |\partial(B) \setminus \partial(\Gamma)| > 0$ then gives curvature $\beta(B)/b$ across each such edge to the (internal) green face on the other side.
- Step 5 Return the function $\kappa_\Gamma : V(\Gamma) \cup E(\Gamma) \cup F(\Gamma) \rightarrow \mathbb{R}$, where $\kappa_\Gamma(x)$ is the current curvature of x .

Definition 6.4. We define **RSym** to be the map from \mathcal{D} to $\{\kappa_\Gamma(x) : x \in \mathcal{D}\}$ evaluated by **ComputeRSym**, so that $\kappa_\Gamma = \text{RSym}(\Gamma)$. We denote the curvature given by a vertex v to a

face f in Step 3 of **ComputerSym** by $\chi(v, f, \Gamma)$, noting that if f is incident more than once with v then f will receive a proper multiple of $\chi(v, f, \Gamma)$ of curvature from v . Similarly, we denote the curvature given by a blob B to a face f in Step 4 of **ComputerSym** by $\chi(B, f, \Gamma)$. We shall omit the Γ from $\chi(v, f, \Gamma)$ and $\chi(B, f, \Gamma)$ when the meaning is clear.

Since the curvature in Step 1 of **ComputerSym** is precisely the curvature distribution from Example 5.3, and curvature is neither created nor destroyed by the algorithm, the following is immediate:

Proposition 6.5. *RSym is a curvature distribution scheme on \mathcal{D} .*

Recall Definition 5.4 of *dual distance*.

Definition 6.6. Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation, and let $\varepsilon > 0$ be a constant. We say that **RSym** *succeeds with constant ε* on a diagram $\Gamma \in \mathcal{D}$ if $\kappa_\Gamma(f) \leq -\varepsilon$ for all internal non-boundary green faces of Γ .

More generally if, for some $d \geq 1$, we can bound $\kappa_\Gamma(f) \leq -\varepsilon$ for all green faces of Γ that are at dual distance at least $d + 1$ from the external face, then we say that **RSym** *succeeds with constant ε at level d* . (So the default level is $d = 1$.)

We say that **RSym** *succeeds on \mathcal{P} with constant ε* (at level d) if this is true for every $\Gamma \in \mathcal{D}$, and **RSym** *succeeds on \mathcal{P}* (at level d) if there exists an $\varepsilon > 0$ for which **RSym** succeeds.

Our goal in the rest of this section is to show that, if **RSym** succeeds on a pregroup presentation \mathcal{P} , then the group presented by \mathcal{P} is hyperbolic. Before we can do that, we need to study the behaviour of **RSym**, and then prove two technical lemmas which will allow us to deal with our frequent assumption of loop-minimality in earlier sections.

We first show, amongst other things, that for each $\Gamma \in \mathcal{D}$ the curvature distribution $\kappa_\Gamma = \text{RSym}(\Gamma)$ satisfies Condition (a) of Theorem 5.9 for vertices.

Lemma 6.7. *Let v be a vertex of a diagram $\Gamma \in \mathcal{D}$, incident with $v_G := \delta_G(v, \Gamma)$ green faces, of which x are the external face. If $x \neq v_G$ then let f be a non-external green face incident with v . Then*

- (i) $\kappa_\Gamma(v) \leq 0$, and $\kappa_\Gamma(v) = 0$ if $x \neq v_G$;
- (ii) if $x \neq v_G$ then $\chi(v, f, \Gamma) = \frac{2-v_G}{2(v_G-x)}$;
- (iii) if $v_G > 2$ and $x \neq v_G$ then $\chi(v, f, \Gamma) \leq -1/6$.

Proof. The vertex v begins with curvature $+1$, and $v_G \geq 2$ since Γ is green-rich. Thus v has at least two incoming green half-edges, so receives at most -1 in curvature in Step 2 of **ComputerSym**(Γ). Thus Part (i) holds, and Part (ii) is now clear. For Part (iii), notice that if $v_G > 2$ and $x \neq v_G$ then the maximum value of $\frac{2-v_G}{2(v_G-x)}$ is attained when $x = 0$ and $v_G = 3$. \square

We now show that, for each Γ in \mathcal{D} , the curvature distribution κ_Γ satisfies Condition (a) of Theorem 5.9 for red faces. Recall our conventions in Definition 3.1 on boundaries of faces.

Lemma 6.8. *Let B be a red blob composed of t triangles in a diagram $\Gamma \in \mathcal{D}$. Then $\kappa_\Gamma(T) \leq 0$ for each triangle T of B .*

Let $d := |\partial(B) \cap \partial(\Gamma)|$. Then

$$\chi(B, f, \Gamma) = \frac{-t}{2|\partial(B) \setminus \partial(\Gamma)|} \leq \frac{-t}{2(t-d)+4} \leq -\frac{1}{6}.$$




Figure 3: Isolated boundary vertex

Proof. After Step 2 of `ComputerSym`(Γ), the curvature of each triangle T is $-1/2$, so $\kappa_\Gamma(T) \leq 0$, as required. Hence $\chi(B, f, \Gamma) = -t/(2|\partial(B) \setminus \partial(\Gamma)|)$. By Lemma 4.12, $|\partial(B)| \leq t + 2$, so

$$\frac{-t}{2|\partial(B) \setminus \partial(\Gamma)|} \leq \frac{-t}{2(t-d)+4} \leq \frac{-t}{2t+4} \leq -\frac{1}{6}.$$

□

Recall Definition 3.1 of a consolidated edge. It follows from the fact that all diagrams $\Gamma \in \mathcal{D}$ are green-rich that if a consolidated edge of Γ has length greater than 1, then both of the incident faces are green. We now show that for all diagrams $\Gamma \in \mathcal{D}$, the curvature distribution κ_Γ satisfies Condition (c) of Theorem 5.9, with $m = 1/2$ and $d = 1$. The second part of the next lemma will be used when we attack the word problem, in Section 8.

Lemma 6.9. *Let $\Gamma \in \mathcal{D}$ have area greater than 1, and let f be a boundary green face of Γ . Then $\kappa_\Gamma(f) \leq 1/2$.*

Furthermore, if $\kappa_\Gamma(f) > 0$ then the consolidated edges and vertices in $\overline{\partial(f) \setminus \partial(\Gamma)}$ form a single path p , and at most three of the vertices in p lie on $\partial(\Gamma)$. If there are three such vertices, let v be the middle one (as in Figure 3). Then $\delta_G(v, \Gamma) \geq 4$, and f is incident with no red blobs at v .

Proof. First assume that $\partial(f) \setminus \partial(\Gamma)$ contains no edge. Since $\text{Area}(\Gamma) > 1$, each vertex v_0 that lies on both $\partial(f)$ and an edge of $\partial(\Gamma) \setminus \partial(f)$ satisfies $\delta_G(v_0, \Gamma) \geq 3$, and is incident at least twice with the external face. By Lemma 6.7, each such vertex v_0 therefore satisfies $\chi(v_0, f, \Gamma) \leq -1/2$, so $\kappa_\Gamma(f) \leq 1/2$ and both claims follow.

Assume instead that $\partial(f) \setminus \partial(\Gamma)$ contains an edge. Let β denote a maximal sequence of incident vertices and (consolidated) edges on $\partial(f)$, such that each edge of β is internal in Γ , and let v_1 and v_2 be the vertices at the beginning and end of β . If $\delta_G(v_i, \Gamma) \geq 3$ then we can apply Lemma 6.7 with $x \geq 1$ and $v_G \geq 3$ to deduce that $\chi(v_i, f, \Gamma) \leq -1/4$, and so $\kappa_\Gamma(f) \leq 1/2$. So assume that $\delta_G(v_i, \Gamma) = 2$ for at least one $i \in \{1, 2\}$. Then f is adjacent to a red blob B_i at v_i , and $|\partial(B_i) \cap \partial(\Gamma)| \geq 1$. By Lemma 6.8, $\chi(B_i, f, \Gamma) \leq -1/4$. It follows that $\kappa_\Gamma(f) \leq 1/2$ unless $\delta_G(v_1, \Gamma) = \delta_G(v_2, \Gamma) = 2$ and $B_1 = B_2$. But in this case, $|\partial(B_1) \cap \partial(\Gamma)| \geq 2$, and so $\chi(B_1, f, \Gamma) \leq -1/2$, by Lemma 6.8.

Suppose now that $\kappa_\Gamma(f) > 0$. Then, by the previous paragraph, there must be exactly one such maximal sequence β on $\partial(f)$. Suppose that β contains a vertex $v \neq v_1, v_2$ that lies on $\partial(\Gamma)$. If $\delta_G(v, \Gamma) = 2$, then there are two red blobs adjacent to f at v (either or both of which may be equal to B_1 or B_2), giving additional combined curvature at most $-1/2$ to f , and so $\kappa_\Gamma(f) \leq 0$, contrary to assumption. If $\delta_G(v, \Gamma) = 3$, then f is adjacent

to at least one red blob B_3 at v , and the combined additional curvature that B_3 and v give to f is at most $-1/2$, giving $\kappa_\Gamma(f) \leq 0$ again. Hence $\delta_G(v, \Gamma) \geq 4$, and so v gives at most $-1/3$ of curvature to f , and there can be at most one such v . \square

In particular, we have now shown that, if **RSym** succeeds on a diagram $\Gamma \in \mathcal{D}$, then κ_Γ satisfies Conditions (a), (b) and (c) (with $m = 1/2$) of Theorem 5.9. Since all diagrams in \mathcal{D} are green-rich, if no V^σ -letter is trivial in G it follows immediately from Proposition 3.17 that κ_Γ satisfies Condition (d).

We shall show next that the set \mathcal{D} satisfies Condition (e) of Theorem 5.9, provided that no V^σ -letter is trivial in G .

Proposition 6.10. *Let \mathcal{P} be a pregroup presentation for a group G , and let w be a cyclically P -reduced word that is equal to 1 in G . Assume that no V^σ -letter is trivial in G . Then there exists $w' \in \mathcal{I}(w)$ that is the boundary word of a coloured diagram $\Gamma \in \mathcal{D}$.*

Proof. Let Γ be a coloured diagram of minimal coloured area, amongst all coloured diagrams of words $w' \in \mathcal{I}(w)$, and let w' be the boundary word of Γ . We shall show that $\Gamma \in \mathcal{D}$.

Since $w' \approx^c w$, it follows from Theorem 4.4 that w' is cyclically P -reduced. The assumption that Γ has minimal coloured area implies that Γ is semi- P -reduced, by Proposition 3.8.

We have assumed that no V^σ -letter is trivial in G , so Γ is loop-minimal, and if $|\partial(\Gamma)| = 1$ then the unique boundary face is green. No letter $x \in X$ such that $x = x^\sigma$ is trivial in G , so by Proposition 3.9 the diagram Γ is σ -reduced. It now follows from Proposition 4.10(i) and the minimality of $\text{CArea}(\Gamma)$ that Γ is green-rich.

Finally, the loop-minimality of Γ and the minimality of $\text{CArea}(\Gamma)$ imply that the boundary word of each simply connected red blob has no proper subwords equal to 1 in $U(P)$, by Proposition 4.13. Hence $\Gamma \in \mathcal{D}$. \square

It remains to deal with the assumption that no V^σ -letter is trivial in G . To do so, we first prove that, subject to some easily-checkable conditions on the set \mathcal{R} of relators, if **RSym** succeeds then there are no diagrams in \mathcal{D} of boundary length two.

Lemma 6.11. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation. Suppose that no $R \in \mathcal{R}$ has length 1 or 2 and that no two distinct cyclic conjugates of relators $R, S \in \mathcal{I}(\mathcal{R})^\pm$ have a common prefix consisting of all but one letter of R or S . Let Γ be a diagram in \mathcal{D} with boundary length 2. Then **RSym** does not succeed on Γ .*

Proof. Suppose that **RSym** succeeds on Γ . Since each $R \in \mathcal{R}$ has length at least three, $\text{Area}(\Gamma) > 1$. Each boundary face f of Γ satisfies $\kappa_\Gamma(f) \leq 1/2$ if f is green, by Lemma 6.9, and $\kappa_\Gamma(f) \leq 0$ if f is red. Hence, since all of the positive curvature of κ_Γ lies on the boundary faces and sums to at least 1, the diagram Γ has exactly two boundary faces, f_1 and f_2 say, both green, and $\kappa_\Gamma(f_1) = \kappa_\Gamma(f_2) = 1/2$. Now any other green face f of Γ would satisfy $\kappa_\Gamma(f) < 0$, so no such face exists.

The two vertices v_1 and v_2 on $\partial(\Gamma)$ both satisfy $\delta_G(v_i, \Gamma) \geq 3$, so by Lemma 6.7 in Step 3 of **ComputeRSym** they each give curvature at most $-1/4$ to each of f_1 and f_2 . On the other hand, the assumptions on common prefixes of relators imply that $\text{Area}(\Gamma) > 2$. So f_1 and f_2 are both adjacent to red blobs, and by Lemma 6.8 they receive a negative amount of curvature from these blobs in Step 4 of **ComputeRSym**, giving a contradiction. \square

We now give a condition under which the success of **RSym** shows that no V^σ -letter is trivial in G .




Figure 4: Boundary vertex with red degree 4

Theorem 6.12. *Let \mathcal{P} be a pregroup presentation for a group G . Assume that no $R \in \mathcal{R}$ has length 1 or 2 and that no two distinct cyclic conjugates of relators $R, S \in \mathcal{I}(\mathcal{R})^\pm$ have a common prefix consisting of all but one letter of R or S . If \mathbf{RSym} succeeds on \mathcal{P} at level 1, then no V^σ -letter is trivial in G .*

Proof. Suppose that some V^σ -letter x is equal to 1 in G , and let Γ be a coloured diagram over $\mathcal{I}(\mathcal{P})$ with boundary word x , and with smallest possible coloured area for diagrams with boundary word a single V^σ -letter. We do *not* assume that $\Gamma \in \mathcal{D}$. We shall show that Γ does not exist.

Proposition 3.8 shows that Γ is semi- P -reduced. The diagram Γ is loop-minimal by definition, and Γ is σ -reduced by Proposition 3.9. Our assumption that no $R \in \mathcal{R}$ has length 1 implies that $\text{Area}(\Gamma) > 1$. Let f be the unique boundary face of Γ .

Suppose that Γ is green-rich. Then, since we chose Γ to have minimal coloured area, Proposition 4.13 implies that $\Gamma \in \mathcal{D}$. Hence \mathbf{RSym} succeeds on Γ for some $\varepsilon > 0$, and in particular Lemma 6.9 implies that $\kappa_\Gamma(f) \leq 1/2$, which contradicts the total curvature of Γ being 1.

Hence Γ is not green-rich. By Proposition 4.10 (i), this can only occur when f is a red triangle. The fact that Γ is loop-minimal implies that Γ looks like the left hand picture in Figure 2. The boundary label of the subdiagram labelled Θ in Figure 2 is cyclically σ -reduced, since the label of f is in V_P , so by Proposition 4.10 (ii), with the possible exception of v and the other vertex in Figure 2, which we shall call u , all vertices w in Θ satisfy $\delta_G(w, \Theta) = \delta_G(w, \Gamma) \geq 2$. We shall show that $\delta_G(u, \Gamma) \geq 2$ and $\delta_G(v, \Gamma) \geq 2$, and hence that Γ is in fact green-rich, a contradiction.

Consider first the vertex labelled v in Figure 2. Theorem 3.16 shows that $\delta_G(v, \Gamma) \geq 1$, so assume, by way of contradiction, that $\delta_G(v, \Gamma) = 1$, and let f_1 be the green face incident with v . If $\delta_R(v, \Gamma) = 1$, then the boundary label z of f_1 is not cyclically P -reduced, a contradiction. If $\delta_R(v, \Gamma) = 2$, then both incident red faces are adjacent to f_1 . Then f_1 can be replaced by a green face with boundary label an interleave of z , yielding a diagram with boundary word x but with smaller coloured area, a contradiction. If $\delta_R(v, \Gamma) \geq 3$ then, since the loop-minimality of Γ implies that there are no loops labelled by elements of V_P based at v , we can apply Lemma 3.15 to reduce $\delta_R(v, \Gamma)$ to two, and then we reach a contradiction as before. Hence $\delta_G(v, \Gamma) \geq 2$.

Assume next that $\delta_G(u, \Gamma) = 1$ (so the only green face incident with u is the external face), and notice that $\delta_R(u, \Gamma) \geq 3$, since u is incident twice with the red face f , and with at least one other red face. If $\delta_R(u, \Gamma) = 3$, then Γ contains a loop at v labelled by a letter from V_P , contradicting the minimality of Γ . If $\delta_R(u, \Gamma) = 4$, then Θ consists of two red triangles that meet at v , and enclose a subdiagram Δ of boundary length 2, as in Figure 4. Let the boundary label of Δ be ab , and notice that $b \neq a^\sigma$, as otherwise there would exist a diagram proving that $x =_{U(P)} 1$, contradicting Theorem 2.9. If (a, b) or (b, a) is in $D(P)$, then there would exist a diagram consisting of Δ surrounded by a single red triangle, which would have boundary a single letter from V_P but have coloured area less than that of Γ , contradicting the minimality of Γ . Hence Δ is a σ -reduced, semi- P -reduced, green-rich loop-minimal coloured diagram with cyclically P -reduced boundary word, and the

minimality of Γ means that Proposition 4.13 shows that $\Delta \in \mathcal{D}$. However, we showed in Lemma 6.11 that **RSym** fails on all diagrams in \mathcal{D} of boundary length 2, contradicting our assumption that **RSym** succeeds. Hence $\delta_R(u, \Gamma) \geq 5$, and so $\delta_R(u, \Theta) \geq 3$. Applying Lemma 3.15 to Θ (in which there are no loops based at u with labels from V_P), allows one to reduce $\delta_R(u, \Theta)$ down to three, eventually yielding a contradiction as in the previous paragraph. Hence $\delta_G(u, \Gamma) \geq 2$, and so Γ is in fact green-rich, a contradiction. \square

Finally, we are able to prove that, subject to the same conditions on \mathcal{R} as in Theorem 6.12, if **RSym** succeeds on a pregroup presentation \mathcal{P} for a group G , then G is hyperbolic. Recall Definition 5.5 of a pregroup Dehn function.

Theorem 6.13. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation of a group G , and let r be the maximum length of a relator in \mathcal{R} .*

Assume that no $R \in \mathcal{R}$ has length 1 or 2 and that no two distinct cyclic conjugates of relators $R, S \in \mathcal{I}(\mathcal{R})^\pm$ have a common prefix consisting of all but one letter of R or S .

- (i) *Suppose that **RSym** succeeds on the presentation \mathcal{P} (at level 1) for some $\varepsilon > 0$. Then the pregroup Dehn function of $\mathcal{I}(\mathcal{P})$ is bounded above by*

$$f(n) = n \left(6 + r + \frac{3+r}{2\varepsilon} \right) - \frac{3+r}{\varepsilon}.$$

- (ii) *If V_P is empty, and **RSym** succeeds on \mathcal{P} , then the pregroup Dehn function of $\mathcal{I}(\mathcal{P})$ is bounded above by $n(\frac{1}{2\varepsilon} + 1) - \frac{1}{\varepsilon}$.*
- (iii) *If V_P is nonempty, $\mathcal{I}(w) = w$ for all cyclically P -reduced words w , and **RSym** succeeds on \mathcal{P} , then the pregroup Dehn function of \mathcal{P} is bounded above by $f(n) - 2n$, where $f(n)$ is as in Part (i).*
- (iv) *If no V^σ -letter is trivial in G , and **RSym** succeeds at level d on \mathcal{P} then the pregroup Dehn function of $\mathcal{I}(\mathcal{P})$ is bounded above by*

$$n \left((3+r) \frac{(r-1)^d - 1}{r-2} \left(1 + \frac{1}{\varepsilon} \right) + 3 \right) - \frac{3+r}{\varepsilon}.$$

*In particular, if **RSym** succeeds at level 1, or if no V^σ -letter is trivial in G and **RSym** succeeds at level d , then G is hyperbolic.*

Proof. We first prove (i), by showing that **RSym** satisfies all conditions of Theorem 5.9. By Proposition 6.5, **RSym** is a curvature distribution scheme on \mathcal{D} . Let $\Gamma \in \mathcal{D}$, and let $\kappa_\Gamma = \mathbf{RSym}(\Gamma)$. The fact that κ_Γ is non-positive on vertices and red triangles follows from Lemmas 6.7 and 6.8, and it is clear from Step 2 of **ComputeRSym** that $\kappa_\Gamma(e) = 0$ for each edge e . Hence κ_Γ satisfies Condition (a). Condition (b) is satisfied with $d = 1$ by our assumption that **RSym** succeeds (at level 1). By Lemma 6.9, κ_Γ satisfies Condition (c) with $m = 1/2$. By Theorem 6.12, no V^σ -letter is trivial in G . Hence all coloured diagrams over $\mathcal{I}(\mathcal{P})$ are loop-minimal and green-rich, so by Proposition 3.17 all diagrams in \mathcal{D} satisfy Condition (d) with $\lambda = 3 + r$ and $\mu = 1$.

It now follows from Theorem 5.9 that if $\Gamma \in \mathcal{D}$ has boundary length n and area greater than 1 then, as in Corollary 5.10,

$$\text{Area}(\Gamma) \leq n \left(4 + r + \frac{3+r}{2\varepsilon} \right) - \frac{3+r}{\varepsilon}.$$

A diagram of area 1 has boundary length $n \geq 3$, since no $R \in \mathcal{R}$ has length less than 3. For $n \geq 3$ the above bound evaluates to at least 1, so in fact this area bound applies to all $\Gamma \in \mathcal{D}$.

We showed in Proposition 6.10 that if $w =_G 1$ then there exists an $w' \in \mathcal{I}(w)$ that is the boundary of a coloured diagram $\Gamma \in \mathcal{D}$. Hence it follows from the definition of $\mathcal{I}(w)$ and Lemma 4.6 that there is a coloured diagram Γ' with boundary word w and area at most $\text{Area}(\Gamma) + 2n$, which gives the bound in the theorem statement.

(ii) Since there are no red triangles, κ_Γ satisfies Condition (d) of Theorem 5.9 with $\lambda = 1$ and $\mu = 0$, so the formula in Part (i) simplifies as given (and is valid for diagrams consisting of a single face).

(iii) We keep λ and μ as in Part (i), but take $w = w'$ in the final paragraph of the proof.

(iv) The assumption that no V^σ -letter is trivial in G means that all diagrams are loop-minimal. Hence as in Part (i), by Proposition 3.17 we can set $\lambda = 3 + r$ and $\mu = 1$ in Theorem 5.9. We can then apply Theorem 5.9 with $m = 1$ and the specified value of d to diagrams $\Gamma \in \mathcal{D}$. Proposition 6.10 applies, since all diagrams are loop-minimal, and so we can complete the argument as in the case $d = 1$. \square

7 A polynomial-time RSym tester

In this section we describe a pair of polynomial-time procedures, $\text{RSymVerify}(\mathcal{P}, \varepsilon)$ and $\text{RSymIntVerify}(\mathcal{P}, \varepsilon)$ (see Procedures 7.19 and 7.30) that attempt to verify that **RSym** succeeds on a given presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ with a given value of ε . They return either **true** or **fail**, together with some additional data in the event of **fail**.

If **true** is returned, then **RSym** is guaranteed to succeed on \mathcal{P} with constant ε . If **fail** is returned, then it does not necessarily mean that **RSym** does not succeed on \mathcal{P} . The additional returned data describes one or more configurations that could arise in a diagram in \mathcal{D} over $\mathcal{I}(\mathcal{P})$ on which **RSym** might fail, but such a diagram may not exist. The user can either attempt to show that such a diagram does not exist, or try again with a smaller value of ε .

For convenience of exposition, **RSymVerify** works under the assumption that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$. This is a commonly-occurring special case – for example all quotients of free products of free and finite groups can be presented this way – and is the case that is currently implemented (see Section 10). **RSymIntVerify** is the generalisation of **RSymVerify** to the case where $\mathcal{I}(\mathcal{R}) \neq \mathcal{R}$, and will be presented in Subsection 7.8.

Since we have had to introduce many new data structures and auxiliary functions to produce a polynomial-time procedure, we shall start by giving an outline of our approach. We remind the reader that our newly-defined terms, our notation, and our procedures are all listed in the Appendix.

Recall from Definition 6.6 that to test whether **RSym** succeeds on \mathcal{P} means to test whether **RSym** succeeds on all of the coloured diagrams Γ in \mathcal{D} . There are infinitely many such diagrams, but fortunately there are only finitely many elements in \mathcal{R} . Hence **RSymVerify** will consider each possible relator $R \in \mathcal{R}^\pm$ as the label of a non-boundary green face f of such a diagram, and attempt to show that $\kappa_\Gamma(f) \leq \varepsilon$.

After describing some pre-processing, which is not part of **RSymVerify**, we begin the real work in Subsection 7.2, where we introduce some data structures that will be used to work efficiently with our relators as cyclic words, and to record information about the possible faces edge-incident with such a green face f labelled by R .

The curvature $\kappa_\Gamma(f)$ is equal to $1 + \sum_{v \in \partial(f)} \chi(v, f, \Gamma) + \sum_B \chi(B, f, \Gamma)$, where the second sum is over those red blobs B which share an edge with f . We will therefore bound $\kappa_\Gamma(f)$ by bounding $\chi(v, f, \Gamma)$ and $\chi(B, f, \Gamma)$, over all diagrams $\Gamma \in \mathcal{D}$ containing a face f labelled by R . There are infinitely many possible vertices and red blobs in diagrams in \mathcal{D} , but we shall take a pragmatic approach of determining partial data about those vertices and blobs which could give f close to zero curvature in Steps 3 and 4

of `ComputerSym`, and otherwise use an upper bound of $-1/3$ for vertices and $-5/14$ for red blobs. This analysis results in the creation of two functions, called `Vertex` and `Blob` (Algorithms 7.7 and 7.12), which take as input information about f and its adjacent faces, and return these curvature bounds.

The next complication arises from our desire for a low degree polynomial time algorithm. This means that we cannot run, say, a backtrack search which tries all possible ways of fitting neighbouring faces around f to bound $\kappa_\Gamma(f)$, as this would lead to a complexity with the length r of the longest relator in the exponent. To keep the complexity where we want it, we use a combinatorial result: Lemma 7.21. This tells us that, provided we associate to each vertex on f the length of the consolidated edge before it, and require the cumulative curvature coming from red blobs and vertices to be less than the proportion of $|R|$ that their edges take up, we can avoid doing a backtrack search. To make this approach work, we need a careful analysis of edge lengths and their corresponding curvature values, and this is the content of Subsection 7.5.

In Subsection 7.6, we shall finally be able to present `RSymVerify`, then in Subsection 7.7 we shall prove Theorem 7.22, which states that `RSymVerify` runs in time $O(|X|^5 + r^3|X||\mathcal{R}|^2)$, where r is the length of the longest relator.

Finally, in Subsection 7.8, we shall describe the modifications that must be made if $\mathcal{I}(\mathcal{R}) \neq \mathcal{R}$, and present Procedure 7.30 (`RSymIntVerify`), which is the generalisation of `RSymVerify` to this case. The outline of the procedure and the majority of the subroutines barely change, but we must work with slightly more complex data structures to handle the potentially exponential size set $\mathcal{I}(\mathcal{R})$, whilst still terminating in polynomial time.

The majority of the subroutines used by `RSymVerify` will be useful for testing other curvature distribution schemes, not just `RSym`.

7.1 Preprocessing

Before running `RSymVerify`, we assume that some preprocessing has been done to the presentation, to ensure that the assumptions of Theorem 6.13 hold, and to improve the likelihood that `RSym` succeeds. The first two steps of preprocessing are done *before* a pregroup P is chosen, when we just have a group presentation $\langle X \mid \mathcal{R} \rangle$.

Preprocessing Step 1: Eliminate any relators of the form x or xy with $x, y \in X$, and $x \neq y$, by eliminating generators. Delete any relators of the form x^2 with $x \in X$, and require that $x^\sigma = x$ in the pregroup P .

Preprocessing Step 2: Look for pairs $R_1, R_2 \in \mathcal{R}$ for which there are distinct cyclic conjugates S_1, S_2 of $R_1^{\pm 1}, R_2^{\pm 1}$ that have a common prefix of length greater than half of $|S_1|$. That is, $S_1 = ww_1$ and $S_2 = ww_2$ with $|w| > |w_1|$. If $R_1 \neq R_2$, then replace R_2 by the shorter relator $w_1^{-1}w_2$. Notice that it is not possible to have $R_1 = R_2$ and $S_1 \neq S_2$ with $|w_1| = |w_2| = 1$.

We may need to carry out these two steps repeatedly, but at the end of them all relators have length at least 3, and no two distinct cyclic conjugates of relators $R, S \in \mathcal{R}^\pm$ have a common prefix consisting of all but one letter of R or S . See, for example, [11, Section 5.3.3] for discussion of how to do this efficiently. Note that we are only carrying out Steps 1 and 3 of the simplification in [11]: we are not attempting to eliminate generators using relators of length greater than two.

Preprocessing Step 3: Define a pregroup P on the remaining generators (adding additional generators if necessary to close P), and ensure that all elements of \mathcal{R} are cyclically P -reduced. Ideally, as many remaining relators of length 3 as possible should become elements of V_P rather than \mathcal{R} .

After this, all hypotheses in Theorem 6.13 are satisfied. Recall that we assume for now that after Preprocessing Step 3, $\mathcal{I}(\mathcal{R}) = \mathcal{R}$: see Subsection 7.8 for further preprocessing that is required if this assumption does not hold.

7.2 Steps, locations and places

RSymVerify needs to check that, for every diagram $\Gamma \in \mathcal{D}$, every non-boundary internal green face $f \in \Gamma$ receives at most $-1 - \varepsilon$ of curvature from its incident vertices and red blobs in Steps 3 and 4 of **ComputerSym**. Each such face f has boundary label some $R \in \mathcal{R}^\pm = \mathcal{I}(\mathcal{R}^\pm)$, and $\partial(f)$ is split up into the consolidated edges (see Definition 3.1) that f shares with its adjacent faces in Γ . In this subsection, we describe how we represent this decomposition into consolidated edges.

The idea is to consider each relator $R \in \mathcal{R}$ in turn, and to use the relators in V_P and \mathcal{R}^\pm to determine the possible decompositions into consolidated edges e_i of the boundary of a non-boundary face f with label R . We do not need to consider \mathcal{R}^{-1} , as the situation for each R^{-1} will be equivalent to that for R . Each such decomposition corresponds to an expression of some cyclic conjugate R' of R as $w_1 w_2 \cdots w_k$, where w_i is the label of e_i .

We attach a colour $C_i \in \{\mathbf{G}, \mathbf{R}\}$ to each w_i in the decomposition, which is the colour of the adjacent region (green face or red blob): there could be more than one decomposition $R' = w_1 w_2 \cdots w_k$ with the same w_i but with different colours C_i . We have assumed that $\Gamma \in \mathcal{D}$, so Γ is green-rich. Hence if $C_i = \mathbf{R}$ then $|w_i| = 1$. For reasons that will become clear shortly, we do not allow $C_1 = \mathbf{R}$ and $C_k = \mathbf{G}$; in that situation, we shall instead consider the decomposition $w_2 \cdots w_k w_1$.

Definition 7.1. If $C_i = \mathbf{G}$ then let ϵ_i be the maximum value of $\chi(v_i, f, \Gamma)$, considered over all possible diagrams $\Gamma \in \mathcal{D}$ in which w_i labels a maximal green consolidated edge on f . If $C_i = \mathbf{R}$ then let ϵ_i be the maximum value of $\chi(v_i, f, \Gamma) + \chi(B, f, \Gamma)$, considered over all possible diagrams $\Gamma \in \mathcal{D}$ on which w_i labels a red consolidated edge of f , and all possible incident red blobs B at w_i .

If we find a decomposition with $\epsilon_1 + \cdots + \epsilon_k > -1 - \varepsilon$, then **RSymVerify** returns **fail** and gives details of the decomposition.

We shall see later, in Lemma 7.17, that $\epsilon_i \leq -1/6$ when $C_i = \mathbf{R}$ and when $C_i = C_{i+1} = \mathbf{G}$, but when $C_i = \mathbf{G}$ and $C_{i+1} = \mathbf{R}$, we can have $\epsilon_i = 0$. In our main algorithm it is convenient to have an upper bound of $-1/6$ on all curvature contributions and, for this reason, we combine subwords $w_i w_{i+1}$ with $C_i = \mathbf{G}$ and $C_{i+1} = \mathbf{R}$ into a single unit, which we call a *step*, and use the curvature contributions from these steps, rather from each individual w_i .

Definition 7.2. For a given coloured decomposition $R' = w_1 w_2 \cdots w_k$ as above, a *step* consists either of a single subword w_i , or of two consecutive subwords $w_i w_{i+1}$, determined as follows, where subscripts should be interpreted cyclically.

- (i) If $C_i = \mathbf{G}$ and $C_{i+1} = \mathbf{R}$, then $w_i w_{i+1}$ is a step.
- (ii) If neither $w_{i-1} w_i$ nor $w_i w_{i+1}$ is a step by Condition (i), then w_i is a step.

We have disallowed the combination, $C_1 = \mathbf{R}$, $C_k = \mathbf{G}$, so this cannot give rise to a step $w_k w_1$. Note that the steps are the same for any cyclic permutation of the decomposition that does not violate the $C_1 = \mathbf{R}$, $C_k = \mathbf{G}$ condition.

Let ϵ_i be as in Definition 7.1. We define the *stepwise curvature* χ of a step to be $\chi = \epsilon_i$ when the step is w_i and $\chi = \epsilon_i + \epsilon_{i+1}$ when it is $w_i w_{i+1}$. The *length* of a step is the number of letters of R' that it comprises.

To carry out the required estimates of upper bounds on the curvature $\epsilon_1 + \cdots + \epsilon_k$, we need to study the possible diagrams Γ in which the decomposition $w_1 \cdots w_k$ of R under consideration can occur. We can choose the amount of detail in which we analyse possible neighbourhoods of the face labelled R in such diagrams. More detail may lead to better estimates, but will take longer to compute. In **RSymVerify**, we generally limit our consideration to the faces of $\Gamma \in \mathcal{D}$ that have at least one edge in common with the face labelled R .



Figure 5: (a) Instantiation of place on face f ; (b) Instantiation of partial vertex on face f

To devise efficient algorithms for carrying out this analysis and for storing the information in a useful form, we need to devise some suitable data structures. The definitions of *locations* and *places* that follow may seem somewhat arbitrary on a first reading but, after some experimentation, they have turned to be efficacious for the purpose in hand.

Definition 7.3. Let $R \in \mathcal{R}^\pm$, and fix a word $w = x_1x_2 \cdots x_{|w|}$ such that $R = w^k$ with k maximal amongst such expressions for R . A *location* on R is an ordered triple (i, a, b) , denoted $R(i, a, b)$, where $i \in \{1, \dots, |w|\}$, $a = x_{i-1}$ (or $x_{|w|}$ if $i = 1$), and $b = x_i$.

For example, if $R = abab = (ab)^2$ then the locations are $R(1, b, a)$ and $R(2, a, b)$. We shall present a method to find such a w and k in the proof of Theorem 7.22. Note that writing R as a proper power, where possible, is not essential to the running of our algorithm, but merely reduces the number of locations. It therefore helps with hand calculations, as in Section 9, and also with the running time of our implementations, but if the reader wishes to think of every pair of letters on every relator as being a different location, no great harm will be done.

In the case of places, we needed to distinguish between places that could conceivably occur, and those that really do occur in some diagram $\Gamma \in \mathcal{D}$: recall that all such diagrams are σ -reduced.

Definition 7.4. A *potential place* \mathbf{P} is a triple $(R(i, a, b), c, C)$, where $R(i, a, b)$ is a location, $c \in X$, and $C \in \{\mathbf{G}, \mathbf{R}\}$. A potential place is a *place* if it is *instantiable*, in the following sense. (See Figure 5 (a).)

- (i) There exists a σ -reduced diagram Γ (see Definition 3.6) with a face f labelled R , a face f_2 meeting f at b , and a vertex between a and b on $\partial(f)$ of degree at least three;
- (ii) the half-edge on f_2 after b^σ is labelled c ;
- (ii) if $C = \mathbf{G}$ then f_2 is green, and if $C = \mathbf{R}$ then f_2 is a red blob.

We say that \mathbf{P} is *green* if $C = \mathbf{G}$ and *red* otherwise.

Notice that if $C = \mathbf{G}$ then, by the fact that Γ is σ -reduced, there exists a location $R'(j, b^\sigma, c)$ such that the label of R' beginning at b^σ is not equal in $F(X^\sigma)$ to the inverse of the label of R that ends at b . If $C = \mathbf{R}$, then b^σ must intermult with c , as in Definition 4.14.

RSymVerify computes an array of all intermult pairs, and then finds all locations of relators $R \in \mathcal{R}^\pm$. For each location $R(i, a, b)$ with $R \in \mathcal{R}$, it must find all instantiable places. To do so, it considers each letter $c \in X$, and each $C \in \{\mathbf{R}, \mathbf{G}\}$. For $C = \mathbf{R}$ it checks that b^σ intermults with c . For $C = \mathbf{G}$ it checks that there exists a location $R'(j, b^\sigma, c)$ such that a diagram of area two with faces labelled by R and R' , sharing the edge labelled b , is σ -reduced.

7.3 Vertex data and the Vertex function

Fix a relator $R \in \mathcal{R}$ and let f be a non-boundary face labelled R in a diagram in \mathcal{D} . Whilst decomposing R into steps, we find an upper bound on $\sum_{v \in \partial(f)} \chi(v, f, \Gamma)$. Although

Table 1: Vertex curvature $\chi(v, f, \Gamma)$

$\delta_G(v, \Gamma)$	$v \notin \partial(\Gamma)$	$v \in \partial(\Gamma)$
2	0	0
3	$-1/6$	$-1/4$
4	$-1/4$	$-1/3$
5	$-3/10$	$(-3/8)$
6	$-1/3$	$(-2/5)$
≥ 7	$(\leq -5/14)$	$(\leq -5/12)$

there are infinitely many possible vertices v that can arise in such a diagram, almost all give f not much more than $-1/2$ of curvature. In this subsection we show how to produce a list of descriptions of the finitely many possible vertices v on $\partial(f)$ such that $\chi(v, f, \Gamma) \geq -1/3$. For reasons of efficiency, we do not completely describe each such v , but only store information on the possibilities for three consecutive incident faces, reading anticlockwise, together with an upper bound on $\chi(v, f, \Gamma)$, over all $\Gamma \in \mathcal{D}$.

From now on, we shall often think of the triangles within a red blob as having been merged, and treat the blob as having no internal structure other than its area. As a consequence, we will never consider a vertex to have more than one consecutive incident red face. The following lemma therefore bounds the total degree of the vertices that we shall classify, as well as their green degree.

Lemma 7.5. *Let v be a vertex in a diagram $\Gamma \in \mathcal{D}$, and let f be an internal green face incident with v . If v is incident more than once with the external face, then $\chi(v, f, \Gamma) \leq -1/2$. Otherwise, the curvature $\chi(v, f, \Gamma)$ is as in Table 1.*

Proof. Since $\Gamma \in \mathcal{D}$, the diagram Γ is green-rich, so $\delta_G(v, \Gamma) \geq 2$. Let x be the number of times that v is incident with the external face. Then $\delta_G(v, \Gamma) > x$ because f is incident with v . By Lemma 6.7, v gives curvature $\chi = \chi(v, f, \Gamma) = (2 - \delta_G(v, \Gamma)) / (2(\delta_G(v, \Gamma) - x))$ to f . The result follows by computing χ for $\delta_G(v, \Gamma) \leq 7$ and $0 \leq x < \delta_G(v, \Gamma)$, and noticing that χ decreases as $\delta_G(v, \Gamma)$ and x increase. \square

Let $\mathbf{P} = (R(i, a, b), c, C)$ be a place, and let v be a vertex on $\partial(f)$ between the edges labelled a and b in a diagram $\Gamma \in \mathcal{D}$ that instantiates \mathbf{P} , in which f is non-boundary. Let f_1 and f_2 be the faces incident with the edges labelled a and b , respectively, as in Figure 5 (b), so that f_2 is as in Definition 7.4. We describe how to encode this situation, and efficiently bound $\chi(v, f, \Gamma)$.

Suppose that v is internal and has degree k . Let the outgoing edges from v , reading anticlockwise, be labelled $x_1 = a^\sigma$, $x_2 = b$, $x_3 = c$, x_4, \dots, x_k , and let the face adjacent to v with edges labelled x_i^σ , x_{i+1} have colour C_i for $1 \leq i \leq k$ (with $x_{k+1} = x_1$). In particular, $C_1 = G$ and $C_2 = C$. To help us compute the required upper bound on $\chi(v, f, \Gamma)$, we record these edge labels and face colours in a certain directed graph \mathcal{G} . To avoid confusion with the vertices and edges of coloured diagrams, we shall always refer to the vertices and edges of \mathcal{G} as \mathcal{G} -vertices and \mathcal{G} -edges.

The location $R(i, a, b)$ of v on f is encoded by the \mathcal{G} -vertex (a^σ, b, G) , and there exist \mathcal{G} -vertices labelled $(x_i^\sigma, x_{i+1}, C_i)$ for each i . There are \mathcal{G} -edges from \mathcal{G} -vertex $(x_{i-1}^\sigma, x_i, C_{i-1})$ to \mathcal{G} -vertex $(x_i^\sigma, x_{i+1}, C_i)$ for each i , so the vertex v of Γ is represented by a circuit of length k in \mathcal{G} starting at the \mathcal{G} -vertex (a^σ, b, G) . We assign the \mathcal{G} -edge from $(x_{i-1}^\sigma, x_i, C_{i-1})$ to $(x_i^\sigma, x_{i+1}, C_i)$ the weight 1 if $C_{i-1} = G$ and 0 if $C_{i-1} = R$. Then $\delta_G(v, \Gamma)$ is equal to the total weight of the circuit in \mathcal{G} , and so we can bound $\chi(v, f, \Gamma)$ by bounding the weight of circuits in \mathcal{G} that start at (a^σ, b, G) .

We shall now present the formal definition of the vertex graph \mathcal{G} .

Definition 7.6. The *vertex graph* \mathcal{G} of \mathcal{P} has \mathcal{G} -vertices of the form (a, b, C) with $a, b \in X$ and $C \in \{G, R\}$. There is a *green \mathcal{G} -vertex* (a, b, G) if and only if there exists a location $R(i, a, b)$. There is a *red \mathcal{G} -vertex* (a, b, R) if and only if (a, b) is an intermult pair (see Lemma 4.16).

There is a (directed) \mathcal{G} -edge from (a, b, G) to (b^σ, c, G) if there exist locations $R(i, a, b)$ and $R'(j, b^\sigma, c)$ such that the one-face or two-face diagram in which faces labelled R and R' share this edge labelled b is σ -reduced. There is a \mathcal{G} -edge from each (a, b, G) to each (b^σ, c, R) . There is a \mathcal{G} -edge from each (a, b, R) to each (b^σ, c, G) . There are no \mathcal{G} -edges between red \mathcal{G} -vertices, since we do not allow red blobs to share edges with other red blobs. The \mathcal{G} -edges have weight 1 if their source is green and weight 0 if it is red.

After computing the list of all places, the next step in **RSymVerify** is to construct \mathcal{G} . We also store a list of the locations that correspond to each green \mathcal{G} -vertex. For each \mathcal{G} -vertex ν , and for each path ν_1, ν, ν_2 in \mathcal{G} , we let $w(\nu_2, \nu_1)$ denote the smallest weight of a path in \mathcal{G} from ν_2 to ν_1 with at least one \mathcal{G} -edge. (So in the case $\nu_1 = \nu_2$ this cannot be 0.) If there is no \mathcal{G} -path from ν_2 to ν_1 , then we take its weight to be infinite. We can use the Johnson-Dijkstra algorithm [15] to find and store the weights of all of these paths.

The **Vertex** function takes as input a triple (ν_1, ν, ν_2) of \mathcal{G} -vertices for which ν is green and there is a (directed) path ν_1, ν, ν_2 in \mathcal{G} . The existence of such a directed path means that ν_1, ν and ν_2 represent subwords of boundary labels, and colours, of adjacent faces f_1, f and f_2 around a vertex v in a coloured diagram, as in Figure 5(b). The **Vertex** function returns an upper bound on $\chi(v, f, \Gamma)$.

Algorithm 7.7. **Vertex** (ν_1, ν, ν_2) : Require that (ν_1, ν) and (ν, ν_2) are \mathcal{G} -edges.

1. If ν_1 and ν_2 are both green, then return $-1/6$, $-1/4$, $-3/10$, or $-1/3$ when $w(\nu_2, \nu_1)$ is respectively 1, 2, 3, or greater than 3.
2. If ν_1 is green and ν_2 is red, then return 0, $-1/6$, or $-1/4$ when $w(\nu_2, \nu_1)$ is respectively 0, 1, or greater than 1.
3. If ν_1 is red and ν_2 is green, then return 0, $-1/6$, or $-1/4$ when $w(\nu_2, \nu_1)$ is respectively 1, 2, or greater than 2.
4. If ν_1 and ν_2 are both red, then return 0.

Lemma 7.8. Let f_1, f, f_2 be three consecutive faces around a vertex v in a diagram $\Gamma \in \mathcal{D}$, in locations corresponding to \mathcal{G} -vertices ν_1, ν, ν_2 . Then $\chi(v, f, \Gamma) \leq \mathbf{Vertex}(\nu_1, \nu, \nu_2)$.

Proof. It is clear from Definition 7.6 that there are \mathcal{G} -vertices ν_1, ν, ν_2 , as required, and \mathcal{G} -edges from ν_1 to ν and from ν to ν_2 . Let $\chi = \chi(v, f, \Gamma)$, and let $\beta = \mathbf{Vertex}(\nu_1, \nu, \nu_2)$.

Assume first that ν_1 and ν_2 are both green, so that $\delta_G(v, \Gamma) \geq 3$, and $\delta_G(v, \Gamma) \geq 4$ if v is boundary. If $\chi > -1/3$ then, by Lemma 7.5, v is not boundary and $\delta_G(v, \Gamma) \leq 5$, so $w(\nu_2, \nu_1) \leq 3$. If $\beta = -1/3$, then $w(\nu_2, \nu_1) \geq 4$. Hence the shortest \mathcal{G} -path from ν_2 to ν_1 passes through at least three additional green \mathcal{G} -vertices (and possibly some red ones), and so either v is boundary or $\delta_G(v, \Gamma) \geq 6$. Hence by Lemma 7.5, $\chi \leq -1/3 = \beta$. If $\beta = -3/10$, then $w(\nu_2, \nu_1) = 3$, so either v is boundary or $\delta_G(v, \Gamma) \geq 5$. Therefore by Lemma 7.5, $\chi \leq -3/10 = \beta$. If $\beta = -1/4$, then $w(\nu_2, \nu_1) = 2$, so $\delta_G(v, \Gamma) \geq 4$, and hence by Lemma 7.5, $\chi \leq -1/4 = \beta$. Similarly, if $\beta = -1/6$ then $w(\nu_2, \nu_1) = 1$, so $\delta_G(v, \Gamma) \geq 3$, and $\chi \leq -1/6 = \beta$.

Next assume that ν_1 is green and ν_2 is red, so that $\delta_G(v, \Gamma) \geq 2$, and $\delta_G(v, \Gamma) \geq 3$ if v is boundary. If $\beta = -1/4$ then $w(\nu_2, \nu_1) \geq 2$, so the shortest \mathcal{G} path from ν_2 to ν_1 passes through at least two additional green vertices. Hence either v is boundary, or $\delta_G(v, \Gamma) \geq 4$, and so by Lemma 7.5 $\chi \leq -1/4 = \beta$. If $\beta = -1/6$ then $w(\nu_2, \nu_1) = 1$, so either v is boundary or $\delta_G(v, \Gamma) \geq 3$, and hence $\chi \leq -1/6 = \beta$.

Table 2: Bounds on simply connected red blob curvature

$ \partial(B) $	$ \partial(B) \cap \partial(\Gamma) $	$\chi(B, f, \Gamma)$
3	0	$-1/6$
3	1	$-1/4$
4	0	$-1/4$
4	1	$-1/3$
5	0	$-3/10$
6	0	$-1/3$

The case ν_1 red and ν_2 green is similar, except that the weights $w(\nu_2, \nu_1)$ are increased by one, since every edge leaving ν_2 has weight one.

Finally, if ν_1 and ν_2 are both red, then $\beta = 0$, which by Lemma 7.5 is an upper bound on χ . \square

Remark 7.9. When testing **RSym** at level 1, each non-boundary face $f \in \Gamma$ may be at dual distance two from the external face. Since we are not recording whether the next edges on f_1 and f_2 (the ones labelled c and d in Figure 5(b)) are boundary edges, we allow the **Vertex** function to consider them as boundary. See Remark 7.14.

Similarly, as a consequence of the possible presence of the external face at dual distance two, the bracketed values of $\chi(v, f, \Gamma)$ in Table 1 are not used by **RSymVerify**, and we impose a lower bound of $-1/3$ on the return values of the **Vertex** function. However if, for example, we are testing **RSym** at level 2 and $V_P = \emptyset$ (so there are no red triangles in any diagram, and $U(P)$ is a free product of copies of \mathbb{Z} and C_2), then all boundary vertices v of a face f to be tested satisfy $\delta_G(v, \Gamma) \geq 6$, and so we can make use of these smaller curvature values.

7.4 Red blob data and the Blob function

Similarly to vertices, there are infinitely many possible red blobs in diagrams in \mathcal{D} . The methods described in this subsection collect information about possible red blobs B such that there exists a diagram $\Gamma \in \mathcal{D}$ and a green face f of Γ , with $\chi(B, f, \Gamma) > -5/14$.

Lemma 7.10. *Let B be a red blob in a diagram $\Gamma \in \mathcal{D}$, let f be an internal green face adjacent to B at an edge e . If B is not simply connected or if $|\partial(B) \cap \partial(\Gamma)| \geq 2$ then $\chi(B, f, \Gamma) \leq -1/2$. If $\chi(B, f, \Gamma) > -5/14$, then $|\partial(B)|$, $|\partial(B) \cap \partial(\Gamma)|$, and $\chi(B, f, \Gamma)$ are as in Table 2.*

Proof. Let B have boundary length l and area t . If B is not simply connected, then by Lemma 4.12, $l \leq t$, so $\chi(B, f, \Gamma) \leq -1/2$. Hence B is simply connected and Γ is green-rich, and so $t = l - 2$, by Lemma 4.12. The result now follows easily from Lemma 6.8. \square

Recall from Lemma 4.16 and Definition 6.1 that two consecutive letters of the boundary word of a red blob in a diagram in \mathcal{D} must intermult. We create a function **Blob**(a, b, c), which takes as input $(a, b, c) \in X^3$ such that (a, b) and (b, c) intermult, and returns an upper bound on

$$\{\chi(B, f, \Gamma) : \Gamma \in \mathcal{D} \text{ contains both a red blob } B \text{ with } abc \text{ a subword of its boundary label, and a green face } f \text{ that is incident with } B \text{ at } b\}.$$

Definition 7.11. We call $x \in X$ an \mathcal{R} -letter if x occurs in an element of $\mathcal{I}(\mathcal{R}^\pm)$.

Notice that non- \mathcal{R} -letters can only appear on the boundary of a diagram, and that our present assumption that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ implies that an \mathcal{R} -letter occurs in an element of \mathcal{R}^\pm .

It is straightforward (see the proof of Theorem 7.22 for details) to compute a list \mathcal{B} of all cyclic words $w \in X^*$ that satisfy all of the following conditions.

1. The word w is equal to 1 in $U(P)$.
2. $3 \leq |w| \leq 6$.
3. Each consecutive pair of letters in w intermult.
4. No proper nonempty subword of w is equal to 1 in $U(P)$.
5. w contains at most one non- \mathcal{R} -letter, and none if $|w| > 4$.

When **RSymVerify** bounds $\chi(B, f, \Gamma)$, it will have specified three consecutive letters a, b, c on $\partial(B)$.

Algorithm 7.12. $\text{Blob}(a, b, c)$: Require that $(a, b), (b, c)$ intermult.

1. If abc is a cyclic subword of a word in $w \in \mathcal{B}$ then return the maximal curvature from Lemma 7.10 over all such words w , with $|\partial(B) \cap \partial(\Gamma)|$ assumed to be nonzero if and only if w contains a non- \mathcal{R} -letter.
2. Otherwise, if at most one of a and c is not an \mathcal{R} -letter then return $-5/14$.
3. Otherwise, return $-1/2$.

Lemma 7.13. *Let B be a red blob with subword abc of its boundary word, in a diagram $\Gamma \in \mathcal{D}$, and let f be the green face incident with B at b . Then $\chi(B, f, \Gamma) \leq \text{Blob}(a, b, c)$.*

Proof. By Lemma 4.16 both (a, b) and (b, c) intermult, so $\text{Blob}(a, b, c)$ is defined.

If B is not simply connected, then from Lemma 7.10 we see that $\chi(B, f, \Gamma) \leq -1/2 \leq \text{Blob}(a, b, c)$, so assume that B is simply connected, and let w be the boundary word of B . If w contains at least two \mathcal{R} -letters, then $\partial(B) \cap \partial(\Gamma) \geq 2$, and so $\chi(B, f, \Gamma) \leq -1/2 \leq \text{Blob}(a, b, c)$, by Lemma 7.10, so assume that w contains at most one non- \mathcal{R} -letter, and in particular that at most one of a and c is not an \mathcal{R} -letter.

Let l be the length of w . If $l \geq 7$ then $\chi(B, f, \Gamma) \leq -5/14 \leq \text{Blob}(a, b, c)$, so assume that $l \leq 6$. If $l \in \{5, 6\}$ and w contains an \mathcal{R} -letter, then $\chi(B, f, \Gamma) \leq -5/14 \leq \text{Blob}(a, b, c)$, so assume not.

Then: w is equal to 1 in $U(P)$, since B is simply-connected; each consecutive pair of letters of w intermult, by Lemma 4.16; and no proper empty subword of w is equal to 1 in $U(P)$, since $\Gamma \in \mathcal{D}$ (see Definition 6.1). Hence $w \in \mathcal{B}$, and so $\chi(B, f, \Gamma) \leq \text{Blob}(a, b, c)$, as required. \square

Remark 7.14. In Remark 7.9 we observed that the **Vertex** function often assumes that the face f is dual distance two from the external face. At present this data is not being used by the **Blob** function, which may be bounding curvature as if any corresponding blobs have no boundary edges. Some curvature is potentially being missed. We plan to rectify this in future versions of **RSymVerify**, by modifying our definition of places to also record whether the edge labelled by the extra letter is on the boundary, and hence enabling the **Vertex** and **Blob** functions to use this extra data.




Figure 6: One-step reachable places

7.5 One-step reachable places and the OneStep lists

Recall Definition 7.2 of *step* and *step curvature*. In this subsection, we describe how to find the steps, and use the **Vertex** and **Blob** functions to bound the corresponding step curvature. For each place \mathbf{P} on each relator R , we shall create a list $\mathbf{OneStep}(\mathbf{P})$ of those places \mathbf{Q} on R that can be reached from \mathbf{P} in a single step, together with the largest possible associated step curvature χ .

If $R = w^k$ is proper power, then each place occurs k times on R . So a place \mathbf{Q} could occur several times in $\mathbf{OneStep}(\mathbf{P})$, corresponding to different positions on R relative to \mathbf{P} . To differentiate them, we store the number of letters of R between \mathbf{P} and \mathbf{Q} for each item on the list.

Definition 7.15. Let \mathbf{P} be a place with location $R(i, a, b)$. A place \mathbf{Q} is *one-step reachable* at distance l from \mathbf{P} , where $1 \leq l < |R|$, if the following hold:

- (i) \mathbf{Q} has location $R(j, s, t)$ for some $s, t \in X$, where $j = i + l$ (interpreted cyclically).
- (ii) If \mathbf{P} is red, then $l = 1$ (and so $s = b$).
- (iii) If \mathbf{P} is green, then exactly one of the following occurs:
 - (a) there exists a green face f' instantiating \mathbf{P} , and a consolidated edge between f and f' of length l from the location of \mathbf{P} to that of \mathbf{Q} , and \mathbf{Q} is green;
 - (b) there exists an *intermediate place* \mathbf{P}' whose location is $R(j - 1, u, s)$ and whose colour is red, there is a green face f' instantiating \mathbf{P} such that there is a consolidated edge between f and f' of length $l - 1$ between the locations of \mathbf{P} and \mathbf{P}' , and there is red edge between \mathbf{P}' and \mathbf{Q} .

For each place $\mathbf{P} = (R(i, a, b), c, C)$ on R , we compute the list $\mathbf{OneStep}(\mathbf{P})$ as follows. In the description below, by *including* an item (\mathbf{Q}, l, χ) in $\mathbf{OneStep}(\mathbf{P})$, we mean append it to the list if there is no entry of the form (\mathbf{Q}, l, χ') already or, if there is such an entry with $\chi > \chi'$, then replace that entry with (\mathbf{Q}, l, χ) . (If there is such an entry with $\chi \leq \chi'$, we do nothing).

Algorithm 7.16. $\mathbf{ComputeOneStep}(\mathbf{P} = (R(i, a, b), c, C))$:

Step 1 Initialise $\mathbf{OneStep}(\mathbf{P})$ as an empty list.

Step 2 **Case $C = R$.** For each place $\mathbf{Q} = (R(i+1, b, d), x, C')$, and for each $y \in X$ such that y intermults with b^σ , proceed as follows. (See Figure 6 (a).)

Let $\chi_1 := \text{Blob}(y, b^\sigma, c)$, and let $\chi_2 := \text{Vertex}((y, b^\sigma, R), (b, d, G), (d^\sigma, x, C'))$.

Include $(\mathbf{Q}, 1, \chi_1 + \chi_2)$ in $\text{OneStep}(\mathbf{P})$.

Case $C = G$. For each location $R'(k, b^\sigma, c)$ instantiating \mathbf{P} , proceed as follows.

For each place $\mathbf{P}' = (R(j, d, e), x, C')$ on R that can be reached from \mathbf{P} by a single (not necessarily maximal) consolidated edge α between R and R' , let $\nu_1 := (y, d^\sigma, G)$ be the green \mathcal{G} -vertex corresponding to the location on R' at the end of α , and let $l := \ell(\alpha)$. For each out-neighbour $\nu_2 := (e^\sigma, x, C')$ of the \mathcal{G} -vertex $\nu := (d, e, G)$, compute $\chi' := \text{Vertex}(\nu_1, \nu, \nu_2)$.

- (i) If \mathbf{P}' is green then include (\mathbf{P}', l, χ') in $\text{OneStep}(\mathbf{P})$. (See Figure 6 (b).)
- (ii) If \mathbf{P}' is red then \mathbf{P}' is the intermediate place of the step. Find all places \mathbf{Q} that are one letter further along R than \mathbf{P}' . (See Figure 6 (c).) Just as in Case R , compute the combined maximum curvature χ'' returned by the red blob between \mathbf{P}' and \mathbf{Q} and the vertex at \mathbf{Q} . Include $(\mathbf{Q}, l+1, \chi' + \chi'')$ in $\text{OneStep}(\mathbf{P})$.

Lemma 7.17. *Let $\mathbf{P} = (R(i, a, b), c, C)$ and $\mathbf{Q} = (R(j, d, e), c', C')$ be places on the same relator R . Then the following are equivalent.*

- (i) *The place \mathbf{Q} is one-step reachable from \mathbf{P} at distance l .*
- (ii) *There exists a coloured decomposition of a cyclic conjugate R' of R such that a subword w_k or $w_k w_{k+1}$ of R' between a location of \mathbf{P} and a location of \mathbf{Q} is a step of length l , the face adjacent to f at w_k has colour C , and edge after w_k^{-1} labelled c , and the face adjacent to f at the letter after the end of the step has colour C' and next letter c' .*
- (iii) *There exists χ such that $(\mathbf{Q}, l, \chi) \in \text{OneStep}(\mathbf{P})$.*

Furthermore, if $(\mathbf{Q}, l, \chi) \in \text{OneStep}(\mathbf{P})$ then χ is an upper bound for the step curvature, and $\chi \leq -1/6$.

Proof. It follows from the definitions that (i) and (ii) are equivalent: the only thing for the reader to check is that enough conditions have been placed in (ii) to uniquely identify the place \mathbf{Q} at distance l from \mathbf{P} .

It is also clear that the OneStep algorithm finds all one-step reachable places, and does not find any places that are not one-step reachable, so the equivalence of (i) and (iii) follows.

The step curvature is the sum of the curvature given to f by at most two vertices and at most one blob. Hence the fact that χ is an upper bound on the step curvature is immediate from the fact that the Vertex and Blob functions return an upper bound on $\chi(v, f, \Gamma)$ and $\chi(B, f, \Gamma)$ (see Lemmas 7.8 and 7.13).

When $C = R$, or when $C = G$ and $\mathbf{P}' \neq \mathbf{Q}$ so that \mathbf{P}' is red, the claim that $\chi \leq -1/6$ follows from $\text{Blob}(y, b, c) \leq \chi(B, f, \Gamma) \leq -1/6$ for all y, b, c, B . Otherwise $C = G$ and $\mathbf{P}' = \mathbf{Q}$ is green, and the vertex v at any instantiation of \mathbf{P}' in a diagram $\Gamma \in \mathcal{D}$ has three consecutive incident green faces. These are encoded by three green \mathcal{G} -vertices ν_1, ν and ν_2 which form a directed \mathcal{G} -path, and so $\chi(v, f, \Gamma) \leq \text{Vertex}(\nu_1, \nu, \nu_2) \leq -1/6$. \square

7.6 The main RSymVerify procedure

The user chooses a value of $\varepsilon > 0$ to test, and we must check whether the steps computed in the lists $\text{OneStep}(\mathbf{P})$ could be combined around a relator to leave it with more than $-\varepsilon$

of curvature. In this subsection, we shall present the main procedure $\text{RSymVerify}(\mathcal{P}, \varepsilon)$ which carries out these checks and prove that it works.

After computing the data and functions from the previous subsections, $\text{RSymVerify}(\mathcal{P}, \varepsilon)$ runs a sub-procedure, $\text{RSymVerifyAtPlace}(\mathbf{P}_s, \varepsilon)$, at each start place \mathbf{P}_s on each relator $R \in \mathcal{R}$ in turn. If every call to RSymVerifyAtPlace returns **true**, then RSymVerify returns **true**, but if any fail, then it aborts and returns **fail**.

$\text{RSymVerifyAtPlace}(\mathbf{P}_s, \varepsilon)$ creates a list L , whose entries are quadruples (\mathbf{Q}, l, k, ψ) . The first three components represent a place \mathbf{Q} at distance l from \mathbf{P}_s along R that can be reached from \mathbf{P}_s in k steps. The final component ψ is equal to $((1 + \varepsilon)l/|R|) + \chi$, where χ is the largest possible total curvature arising from these k steps. If $\psi \leq 0$ then we are on track for a final curvature of most $-\varepsilon$, whereas if $\psi > 0$ then we are not. We shall show in the proof of Theorem 7.20 that there is no need to keep a record of situations in which $\psi < 0$. In other words, we may assume that if the test fails then $\psi \geq 0$ after each step in the failing decomposition.

By Lemma 7.17, the largest possible step curvature is $-1/6$, so the cumulative curvature after $\lceil 6(1 + \varepsilon) \rceil$ steps is at most $-\varepsilon$, and we only need consider $\lceil 6(1 + \varepsilon) \rceil - 1$ steps from \mathbf{P}_s . There can also be at most r steps, where r is the length of the longest relator in \mathcal{R} , so we define

$$\zeta := \min(\lceil 6(1 + \varepsilon) \rceil - 1, r),$$

and use ζ as an upper bound on the number of steps. Notice that if $\varepsilon < 1$ then $\zeta \leq 11$: the default value of ε in our implementations (see Section 10) is $1/10$.

Similarly to Algorithm 7.16, by *including* an entry $(\mathbf{Q}, l + l', i, \phi')$ in a list L , we mean appending it to L if there is no entry $(\mathbf{Q}, l + l', j, \phi'')$ in L or, if there is such an entry with $\phi' > \phi''$, then replacing it by $(\mathbf{Q}, l + l', i, \phi')$.

Procedure 7.18. $\text{RSymVerifyAtPlace}(\mathbf{P}_s = (R(i, a, b), c, C), \varepsilon)$:

Step 1 Initialise $L := [(\mathbf{P}_s, 0, 0, 0)]$.

Step 2 For $i := 1$ to ζ do:

For each $(\mathbf{P}, l, k, \psi) \in L$ with $k = i - 1$, and for each $(\mathbf{Q}, l', \chi) \in \text{OneStep}(\mathbf{P})$ with $l + l' \leq |R|$, do:

- (i) Let $\psi' := \psi + \chi + (1 + \varepsilon)l'/|R|$.
- (ii) If $\psi' < 0$, or if $l + l' = |R|$ and $\mathbf{Q} \neq \mathbf{P}_s$, then do nothing;
- (iii) else if $\psi' > 0$, $\mathbf{Q} = \mathbf{P}_s$ and $l + l' = |R|$, then return **fail** and L ;
- (iv) else include $(\mathbf{Q}, l + l', i, \psi')$ in L .

Step 3 Return **true**.

We now make a couple of remarks on Procedure 7.18. First, notice that if on the i th iteration of Step 2 there are no places \mathbf{Q} that can be reached from any \mathbf{P} with non-positive curvature, then RSymVerifyAtPlace stops early. As we shall see in the proof of Theorem 7.20, Lemma 7.21 implies that, if there is a decomposition of a cyclic conjugate of R that leads to failure of RSymVerifyAtPlace , then there is a start place from which each intermediate place can be reached with non-negative curvature.

Another observation is that, in Step 2 (iv), a list entry $(\mathbf{Q}, l + l', k, \psi'')$ can be replaced by $(\mathbf{Q}, l + l', i, \psi')$ with $i > k$. If there is a failing decomposition of R' involving the original entry, then there must be a (possibly longer) failing decomposition involving the new entry.

We are now able to summarise the overall procedure, RSymVerify .

Procedure 7.19. $\text{RSymVerify}(\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle, \varepsilon)$:

Step 1 For all $a, b \in X$, use V_P and σ to test whether a and b intermult: Subsection 7.2. Store the list of intermult pairs.

- Step 2 Express each relator $R \in \mathcal{R}^\pm$ as a power w^k and find all locations on w : Subsection 7.2. Store the list of locations.
- Step 3 For each location $R(i, a, b)$, and for each choice of extra letter $x \in X$ and colour $C \in \{\mathbf{G}, \mathbf{R}\}$, test whether the potential place $(R(i, a, b), x, C)$ is instantiable: Subsection 7.2. Store a list of all places, and for each green place store the list of locations $R'(k, b^\sigma, c)$ which instantiate it.
- Step 4 Use the list of places and the list of intermult pairs to compute the vertex graph \mathcal{G} . Store \mathcal{G} and the lists of locations corresponding to each green \mathcal{G} -vertex: Subsection 7.3.
- Step 5 For each pair of \mathcal{G} -vertices ν_2 and ν_1 , find the minimal weight of a non-trivial \mathcal{G} -path from ν_2 to ν_1 . Hence create the **Vertex** function: Algorithm 7.7.
- Step 6 Identify \mathcal{R} -letters and compute the list \mathcal{B} , to create the **Blob** function: Algorithm 7.12.
- Step 7 For each relator $R \in \mathcal{R}$, and each place \mathbf{P} on R , use the **Vertex** and **Blob** functions to run **ComputeOneStep**(\mathbf{P}): Algorithm 7.16. Store the list **OneStep**(\mathbf{P}), for each such place \mathbf{P} .
- Step 8 For each relator $R \in \mathcal{R}$, and each place \mathbf{P}_s on R , run **RSymVerifyAtPlace**($\mathbf{P}_s, \varepsilon$). If **RSymVerifyAtPlace**($\mathbf{P}_s, \varepsilon$) ever returns **fail** and a list L , then return **fail** and L , otherwise do nothing.
- Step 9 Return **true**.

Theorem 7.20. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation, such that $\mathcal{I}(R) = R$ for all $R \in \mathcal{R}$, and let $\varepsilon > 0$. If **RSymVerify**(\mathcal{P}, ε) returns **true** then **RSym** succeeds on \mathcal{P} with constant ε .*

Before proving Theorem 7.20 we prove a useful combinatorial lemma.

Lemma 7.21. *Let $\ell \in \mathbb{Z}_{>0}$, let $L := \{1, 2, \dots, \ell\}$, and let $a_1, a_2, \dots, a_\ell \in \mathbb{R}$. For $m \in \mathbb{Z}$, denote by \overline{m} the element of L with $m \equiv \overline{m} \pmod{\ell}$. If $S := \sum_{m \in L} a_m \geq 0$ then there exists $j \in L$ such that for all $i \in \mathbb{Z}_{>0}$ the partial sum*

$$s_{j,i} := \sum_{m=0}^{i-1} a_{\overline{j+m}} \geq 0.$$

Proof. If $S = s_{1,\ell}$ is the minimum of $\{s_{1,i} \mid i \in L\}$, then all partial sums starting at a_1 are positive, so we can set $j = 1$. Otherwise, choose $j \in L \setminus \{1\}$ such that $s_{1,j-1} \leq s_{1,i}$ for all $i \in L$.

Notice that $s_{a,b} + s_{a+b,c} = s_{a,b+c}$ for all $a, b, c \in \mathbb{Z}_{>0}$. Thus $s_{j,i} = s_{1,j+i-1} - s_{1,j-1}$ for all $i \in \mathbb{Z}_{>0}$. However, $s_{1,j+i-1} = s_{1,\overline{j+i-1}} + kS$ for some $k \in \mathbb{Z}_{\geq 0}$. So in any case, $s_{j,i} = kS + s_{1,\overline{j+i-1}} - s_{1,j-1} \geq 0$ by the choice of j . \square

PROOF OF THEOREM 7.20 Suppose that **RSym** does not succeed on \mathcal{P} with the constant ε . Then there exists a diagram $\Gamma \in \mathcal{D}$ over \mathcal{P} , and a face $f \in \Gamma$, such that f is green, has no boundary edges, and satisfies $\kappa_\Gamma(f) > -\varepsilon$, where $\kappa_\Gamma = \mathbf{RSym}(\Gamma)$. We shall show that **RSymVerifyAtPlace**($\mathbf{P}_s, \varepsilon$) returns **fail** for some \mathbf{P}_s on f .

Let the label on $\partial(f)$ be $R \in \mathcal{R}^\pm$. If $R \notin \mathcal{R}$, then the corresponding face f' in the diagram where all faces have labels the inverses of the labels of those in Γ will also satisfy $\kappa_\Gamma(f') > -\varepsilon$, so assume without loss of generality that $R \in \mathcal{R}$.

As discussed in Subsection 7.2, for some cyclic conjugate R' of R , we have $R' = w_1 w_2 \cdots w_k$, where each w_i labels a consolidated edge e_i in Γ , and each w_i has an associated colour $C_i \in \{\mathbf{G}, \mathbf{R}\}$ describing the colour of the other face incident with e_i in Γ . Recall that we do not allow the combination $C_1 = \mathbf{R}$, $C_k = \mathbf{G}$. From this we derive a

decomposition $R' = v_1 v_2 \cdots v_\ell$, where each v_i labels a step and is equal either to a single w_j or to some $w_j w_{j+1}$ with w_j green and w_{j+1} red.

Let the step curvature given to f by the step corresponding to v_i be χ_i , let l_i be the number of letters in v_i , let $\lambda_i = (1 + \varepsilon)l_i/|R|$, and let $a_i = \chi_i + \lambda_i$. Then $\kappa_\Gamma(f) = 1 + \sum_{i=1}^\ell \chi_i > -\varepsilon$, so $\sum_{i=1}^\ell a_i > 0$.

By Lemma 7.21 there exists an i such that the partial sums

$$a_i, a_i + a_{i+1}, \dots, a_i + \cdots + a_\ell + a_1 + \cdots + a_{i-1}$$

are all non-negative. Replace R' if necessary by its cyclic conjugate $R'' := v_i v_{i+1} \cdots v_{i-1}$, and observe that the steps induced by the corresponding decomposition into consolidated edges are the same subwords v_i as before.

Let \mathbf{P}_s be the place on R at the beginning of R'' (which is instantiable because $\Gamma \in \mathcal{D}$). We showed in Lemma 7.17 that each step v_i corresponds to a pair \mathbf{P}, \mathbf{Q} of places on R , and that there exist l and χ such that $(\mathbf{Q}, l, \chi) \in \text{OneStep}(\mathbf{P})$, with $\chi \geq \chi_i$. $\text{RSymVerifyAtPlace}(\mathbf{P}_s, \varepsilon)$ uses χ in place of χ_i , so the partial sums calculated for each place are greater than or equal to the actual curvature sums, and in particular are all non-negative. Thus $\text{RSymVerifyAtPlace}(\mathbf{P}_s, \varepsilon)$ returns **fail**. \square

7.7 Complexity of RSymVerify

We now show that RSymVerify runs in time polynomial in $|X|$, $|\mathcal{R}|$ and r , where $r := \max\{|R| : R \in \mathcal{R}\}$ is the length of the longest relator.

Recall that we assume that *before* running RSymVerify , the presentation has been simplified and the pregroup has been defined: see Subsection 7.1. The presentation simplification process involves comparing subwords of cyclic conjugates of the relators, and any simplification reduces the total length of the presentation, so it is clear that this can be carried out in polynomial time.

Theorem 7.22. *$\text{RSymVerify}(\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle, \varepsilon)$ runs in time $O(|X|^5 + r^3|X|^4|\mathcal{R}|^2)$.*

Proof. We shall work through Steps 1 to 8 of Procedure 7.19, bounding the time complexity of each step. We are not attempting to find the optimal bounds, simply to show that the process runs in low-degree polynomial time. We assume that products and inverses in the pregroup can be computed in constant time.

In Step 1, we compute an $X \times X$ boolean array describing the set of all intermult pairs $(a, b) \in X^2$. For each $a \in X$, and for each $b \in X \setminus \{a^\sigma\}$, we must check whether $(a, b) \in D(P)$ and, if not, whether there exists an $x \in X$ such that $(a, x) \in D(P)$ and $(x^\sigma, b) \in D(P)$. This can be done in time $O(|X|^3)$.

In Step 2, for each $R \in \mathcal{R}$ we first find w that maximises the value of k for which $R = w^k$. For $2 \leq l \leq |R|/2$, we let w be the length l prefix of R , and test whether $w^{|R|/l} = R$, in total time $O(r^2|\mathcal{R}|)$. There are at most $2r|\mathcal{R}|$ locations defined by \mathcal{R}^\pm . When compiling the list of locations, we record which pairs of locations are mutually inverse, in the sense that they describe corresponding positions in inverse pairs of relators.

In Step 3, we find the $O(r|X||\mathcal{R}|)$ places $\mathbf{P} = (R(i, a, b), c, C)$ with $R \in \mathcal{R}^\pm$. To do so, for each of the $O(r|\mathcal{R}|)$ locations $R(i, a, b)$ we first find all c such that (b^σ, c) is an intermult pair, and hence all instantiable red places $\mathbf{P} = (R(i, a, b), c, R)$, in time $O(|X|)$. Then, for each location $R'(j, b^\sigma, c)$ there is a green place $(R(i, a, b), c, G)$ if and only if the locations $R(i, a, b)$ and $R'(j, b^\sigma, c)$ are not mutual inverses. We computed the inverse pairs of locations in Step 2. So Step 3 requires time $O(r|\mathcal{R}|(|X| + r|\mathcal{R}|)) = O(r^2|X||\mathcal{R}|^2)$.

In Step 4, we compute the vertex graph \mathcal{G} . It has at most $2|X|^2$ \mathcal{G} -vertices. We can find these, and also list the locations corresponding to each green \mathcal{G} -vertex, in time $O(r|\mathcal{R}| + |X|^2)$. Let $\nu = (a, b, G)$ be a green \mathcal{G} -vertex. There is a \mathcal{G} -edge from ν to the \mathcal{G} -vertex $\nu_1 = (c, d, G)$ if and only if (i) $c = b^\sigma$; and (ii) if $d = a^\sigma$ then there is more than

one location corresponding to ν_1 . (For Condition (ii), note that there is at least one such location, since there is one coming from the corresponding position of the inverse R' of each relator R associated with ν . But if there was only one, then the only diagram in which faces labelled R and R' shared the edge labelled b would fail to be σ -reduced. So Condition (ii) ensures that there exists a σ -reduced diagram instantiating this \mathcal{G} -edge.) These two conditions can be tested in constant time for each \mathcal{G} -vertex ν_1 . There is a \mathcal{G} -edge from ν to each \mathcal{G} -vertex (b^σ, c, R) and one from each \mathcal{G} -vertex (c, a^σ, R) to ν . We can define all of these edges from and to ν in time $O(|X|)$. There are no edges between red \mathcal{G} -vertices. So each \mathcal{G} -vertex has \mathcal{G} -degree $O(|X|)$, and we can find the $O(|X|^3)$ \mathcal{G} -edges and assign their weights in time $O(|X|^4)$. The time complexity of Step 4 is $O(r|\mathcal{R}| + |X|^4)$.

In Step 5, we compute the values of $\text{Vertex}(\nu_1, \nu, \nu_2)$. We begin by using the Johnson–Dijkstra algorithm [15] to find the smallest weights of paths between all pairs of \mathcal{G} -vertices in the vertex graph. This algorithm runs in time $O(|V|^2 \log |V| + |V||E|)$, on a graph with $|V|$ vertices and $|E|$ edges, so $O(|X|^5)$ in our case, and returns a matrix of path weights. We then consider each of the $O(|X|^2)$ green \mathcal{G} -vertices ν in turn, and for each of the $O(|X|^2)$ directed \mathcal{G} -paths ν_1, ν, ν_2 , we record the appropriate curvature value for Algorithm 7.7. The total time complexity of Step 5 is $O(|X|^5)$.

In Step 6, we compute the values of the **Blob** function. We first identify the set of \mathcal{R} -letters, in time $O(r|\mathcal{R}|)$, and store this information as a boolean array. We next use the intermult table to construct all of the $O(|X|^5)$ words of length l between 3 and 5 such that each cyclically consecutive pair of letters intermult, and that include at most one non- \mathcal{R} -letter when they have length 3 or 4, and none otherwise. We then discard all such words w that have length 3 or 4 and are not equal to 1 in $U(P)$, or have length 5 and are not equal in $U(P)$ to some $a \in P$. Finally, if $l = 5$, and $w =_{U(P)} a \neq 1$, then we check that a^σ is an \mathcal{R} -letter, and that all cyclic subwords of wa^σ of length 2 or 3 are not equal to 1 in $U(P)$. The checks on each word take constant time, so the time complexity of Step 6 is $O(|X|^5)$.

In Step 7, we compute $\text{OneStep}(\mathbf{P})$, for each of the $O(r|X||\mathcal{R}|)$ places \mathbf{P} in turn. If \mathbf{P} is red then there are $O(|X|)$ 1-step reachable places \mathbf{Q} and, for each of them we make $O(|X|)$ calls to both **Vertex** and **Blob** to find the maximum step curvature. If \mathbf{P} is green, then there are $O(r|X|)$ 1-step reachable places (\mathbf{Q}, l) . To find them, we look up all $O(r|\mathcal{R}|)$ locations for the second face f_1 that instantiates \mathbf{P} , and for each of them we find the length l_1 of the maximal consolidated edge between f and f_1 in time $O(r)$. For each such f_1 , there are $O(l|X|) = O(r|X|)$ possibilities for the place \mathbf{P}' at the end of the consolidated edge. If \mathbf{P}' is green, then $\mathbf{P}' = \mathbf{Q}$, and we include its step curvature with a single call to **Vertex**. If \mathbf{P}' is red, then we need a further $O(|X|^2)$ calls to **Vertex** and **Blob** to find all possible triples (\mathbf{Q}, l, χ) at the end of the step. So the time complexity of Step 7 is $O(r^3|X|^4|\mathcal{R}|^2)$.

Step 8 runs **RSymVerifyAtPlace** at each of the $O(r|X||\mathcal{R}|)$ places. The length of the list L constructed by **RSymVerifyAtPlace** is $O(r|X|)$. Each item on L is considered at most $\zeta \leq r$ times, so the time complexity of Step 8 is $O(r^3|X|^2|\mathcal{R}|)$. \square

7.8 RSym with interleaving

In the previous subsections, we described a procedure, $\text{RSymVerify}(\mathcal{P}, \varepsilon)$, that checks whether **RSym** succeeds on a pregroup presentation \mathcal{P} under the assumption that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$. We now describe the modifications that we make when this assumption does not hold, defining a more general procedure $\text{RSymIntVerify}(\mathcal{P}, \varepsilon)$. It follows the same overall steps as **RSymVerify**, and the reader may wish to refer to Procedure 7.19 for these steps. After presenting the key ideas, we describe **RSymIntVerify** at the end of this subsection: see Procedure 7.30.

We remind the reader that all terms, notation and procedures are listed in the Appendix.

The set $\mathcal{I}(\mathcal{R})$ is potentially of exponential size, since it might be possible to interleave between each pair of letters of each $R \in \mathcal{R}$ (for an example of this, let R be any cyclically P -reduced word in $U(P)$, where P is the pregroup for a free product with amalgamation given in Example 2.8). Despite this, we shall see that **RSymIntVerify** runs in polynomial time.

The overall strategy of **RSymIntVerify** is the same as that of **RSymVerify**: we consider each relator $R \in \mathcal{R}$ in turn, and look for ways to decompose each cyclic conjugate R' of each element of $\mathcal{I}(R)$ into words $w_1 w_2 \cdots w_k$ that maximise the curvature received by R' in Steps 3 and 4 of **ComputeRSym**. Each w_i is an interleave of $[s_{i-1}^\sigma w'_{i1}] w'_{i2} \cdots [w'_{il_i} s_i]$ for some $s_{i-1}, s_i \in P$, where $w'_i = w'_{i1} \cdots w'_{il_i}$ is the corresponding subword of the corresponding cyclic conjugate of R . It follows from Lemma 4.6 that all elements of $\mathcal{I}(R)$ have a description of this form.

In the preprocessing stage (see Subsection 7.1), we first carry out Preprocessing Steps 1 to 3. Once the pregroup has been chosen, if $\mathcal{I}(\mathcal{R}) \neq \mathcal{R}$ then in Preprocessing Step 4 we find all $R_1, R_2 \in \mathcal{R}$ for which there exist distinct cyclic conjugates S_1, S_2 in $\mathcal{I}(R_1^\pm), \mathcal{I}(R_2^\pm)$ with $S_1 = w w_1, S_2 = w w_2$ and $|w| > |w_1|$. Each such common prefix w is equal in $U(P)$ to $s^\sigma w' t$ for some $s, t \in P$, where w' has length $|w|$ and is a prefix of a cyclic conjugate of R_1 or R_1^{-1} . We can solve the word problem in $U(P)$ in linear time by Corollary 2.10, so we can find all such S_1, S_2 in polynomial time.

As in preprocessing Step 2, if we find such a pair and $R_1 \neq R_2$, then we replace R_2 by $w_1^{-1} w_2$. It is conceivable that $R_1 = R_2, S_1 \neq S_2$ and $|w_1| = |w_2| = 1$. In that case we use the implied length two relator $w_1^{-1} w_2$ to adjust the pregroup. So again the simplification process ensures that the hypothesis, and hence also the conclusion, of Theorem 6.13 holds.

Definition 7.23. For each $(a, b) \in X \times X$ the *interleave set*, denoted $\mathcal{I}(a, b)$, consists of all $s \in P$ such that $(a, s), (s^\sigma, b) \in D(P)$. We explicitly permit $s = 1$, so that no interleave set is empty.

Example 7.24. Let P_1 be the pregroup for a free product $G * H$, as in the first part of Example 2.5. If $g, h \in G$ then $\mathcal{I}(g, h) = G$, whilst if $g \in G$ and $h \in H$ then $\mathcal{I}(g, h) = 1$.

Let P_2 be the pregroup for a free product with amalgamation, as in the second part of Example 2.5. Then $\mathcal{I}(a, b) = P_2$, for all $a, b \in X$.

Definition 7.25. Let $R \in \mathcal{R}^\pm$. A *decorated location* on R is a 4-tuple (i, a, b, s) , denoted $R(i, a, b, s)$, where $R(i, a, b)$ is a location, and $s \in \mathcal{I}(a, b)$. Let the subword of the cyclic word R containing the location $R(i, a, b)$ be *dabe* with $d, e \in X$. The *pre-interleave set* $\text{Pre}(R(i))$ of $R(i, a, b, s)$ is $\mathcal{I}(d, a)$, and the *post-interleave set* $\text{Post}(R(i))$ is $\mathcal{I}(b, e)$.

We now generalise Definition 7.4 to cover non-trivial interleaves. Recall that $[a_1 a_2 \cdots a_n]$, with $a_i \in P$, denotes the element of P that is the product in $U(P)$ of the a_i .

Definition 7.26. A *potential decorated place* \mathbf{P} is a triple $(R(i, a, b, s), c, C)$, where $R(i, a, b, s)$ is a decorated location, $c \in X$, and $C \in \{\mathbf{G}, \mathbf{R}\}$. A potential decorated place is a *decorated place* if it is instantiable in the following sense:

1. There exists a σ -reduced and semi- P -reduced diagram Γ over $\mathcal{I}(\mathcal{P})$ with a face f labelled by an element of $\mathcal{I}(R)$, elements $t \in \text{Pre}(R(i))$ and $u \in \text{Post}(R(i))$, a face f_2 meeting f at an edge labelled $[s^\sigma b u] \in X$, and the vertex between the edges of f labelled $[t^\sigma a s] \in X$ and $[s^\sigma b u]$ has degree at least three;
2. the half-edge on f_2 after $[s^\sigma b u]^\sigma$ is labelled c ;
3. if $C = \mathbf{G}$ then f_2 is green, and if $C = \mathbf{R}$, then f_2 is a red blob.

Notice in particular that we require $([s^\sigma b], u), ([t^\sigma a], s) \in D(P)$.

We determine whether a potential decorated place $\mathbf{P} = (R(i, a, b, s), c, C)$ is instantiable as follows. If $C = \mathbf{R}$, then \mathbf{P} is a decorated place if and only if there exists $u \in \text{Post}(R(i))$ such that $[s^\sigma bu]^\sigma$ is an element of X that intermults with c . When $C = \mathbf{G}$, for each decorated location $R'(j, d, e, v)$, we first check whether there exists $x \in \text{Post}(R'(j))$ such that $[v^\sigma ex] =_P c$. Then, for each $u \in \text{Post}(R(i))$ we check whether there exists $y \in \text{Pre}(R'(j))$ such that $[y^\sigma dv] =_P [s^\sigma bu]^\sigma$. Finally, we check whether the resulting diagram is σ -reduced and semi- P -reduced. If there exist such $R'(j, d, e, v)$, u , x and y , then \mathbf{P} is a green decorated place.

Definition 7.27. The *decorated vertex graph* \mathcal{V} of $\mathcal{I}(\mathcal{P})$ has two sets of vertices. There is a *green \mathcal{V} -vertex* (a, b, s, \mathbf{G}) if and only if there exists a decorated location $R(i, a, b, s)$. There is a *red \mathcal{V} -vertex* (a, b, \mathbf{R}) for each intermult pair (a, b) .

There is a \mathcal{V} -edge from (a, b, s, \mathbf{G}) to (d, e, v, \mathbf{G}) if there exist decorated locations $R(i, a, b, s)$ and $R'(j, d, e, v)$ such that each one-face or two-face diagram with faces labelled by elements of $\mathcal{I}(R)$ and $\mathcal{I}(R')$, sharing an edge at these locations labelled $[s^\sigma bu] \in X$ on the R -side and $[y^\sigma dv] = [s^\sigma bu]^\sigma$ on the R' -side, is σ -reduced and semi- P -reduced, for some $u \in \text{Post}(R(i))$ and $y \in \text{Pre}(R'(j))$.

There is a \mathcal{V} -edge from (a, b, s, \mathbf{G}) to each $([s^\sigma bu]^\sigma, c, \mathbf{R})$, where $u \in \text{Post}(R(i))$ for some $R(i, a, b, s)$ and $[s^\sigma bu] \in X$. There is a \mathcal{V} -edge from (a, b, \mathbf{R}) to (c, d, t, \mathbf{G}) if and only if there exists a decorated location $R(j, c, d, t)$, and a $u \in \text{Pre}(R(j))$, such that $[u^\sigma ct] = b^\sigma$.

The \mathcal{V} -edges have weight 1 if their source is green, and weight 0 if it is red.

The most significant difference between **RSymVerify** and **RSymIntVerify** is in finding the one-step reachable decorated places (see Algorithm 7.16). To simplify the exposition, we shall break this task into two parts: finding the edges between green faces, and finding the steps. It would be quicker to carry out these tasks concurrently.

The following algorithm, **FindEdges**, takes as input $R \in \mathcal{R}$ and $S \in \mathcal{R}^\pm$, and returns a list $L_{R,S}$ of all possible consolidated edges e between faces f and f_1 with labels in $\mathcal{I}(R)$ and $\mathcal{I}(S)$. The list $L_{R,S}$ consists of a 5-tuple for each (not necessarily maximal) consolidated edge e : the decorated locations of R at the beginning and end of e in f , the two corresponding decorated locations of S , and the length of e .

Algorithm 7.28. **FindEdges**(R, S):

Step 1 Initialise $L_{R,S} := []$.

Step 2 For each pair of decorated locations $R(i, u_0, u_1, s)$ and $S(j, v_1, v_0, t)$, with corresponding cyclic conjugates $u_0 u_1 \dots u_{n-1}$ of R and $v_0^\sigma v_1^\sigma \dots v_{m-1}^\sigma$ of S^{-1} :

- (a) Test, using \mathcal{V} , whether these could be the beginning of a consolidated edge e .
- (b) If so, then consider each possible consolidated edge length $l = 1, 2, \dots, r$. For each $s_l \in \mathcal{I}(u_l, u_{l+1})$ and each $t_l \in \mathcal{I}(v_{l+1}, v_l)$, if

$$[s^\sigma u_1] u_2 \dots u_{l-1} [u_l s_l] =_{U(P)} ([t_l^\sigma v_l] v_{l-1} \dots v_2 [v_1 t_l])^{-1}$$

then add $(R(i, u_0, u_1, s), R(i+l, u_l, u_{l+1}, s_l), S(j, v_1, v_0, t), S(j-l, v_{l+1}, v_l, t_l), l)$ to $L_{R,S}$. If not, then do nothing.

Step 3 Return $L_{R,S}$.

Lemma 7.29. Let $R \in \mathcal{R}$ and $S \in \mathcal{R}^\pm$. Then **FindEdges**(R, S) returns all of the consolidated edges between cyclic conjugates of elements of $\mathcal{I}(R)$ and $\mathcal{I}(S)$. Furthermore, **FindEdges**(R, S) runs in time $O(r^4 |X|^4)$.

Proof. To simplify the notation, we will write $R = u_1 u_2 \cdots u_n$ and $S = v_m v_{m-1} \cdots v_1$. Assume that

$$\begin{aligned} R' &= [s_0^\sigma u_1 s_1] [s_1^\sigma u_2 s_2] \cdots [s_{n-1}^\sigma u_n s_0] \in \mathcal{I}(R), \\ S' &= [t_0^\sigma v_m t_{m-1}] [t_{m-1}^\sigma v_{m-1} t_{m-2}] \cdots [t_1^\sigma v_1 t_0] \in \mathcal{I}(S) \end{aligned}$$

have a consolidated edge e between them.

Since **FindEdges** considers each possible pair of starting decorated locations, we may assume without loss of generality that e is labelled by the subwords

$$[s_0^\sigma u_1 s_1] [s_1^\sigma u_2 s_2] \cdots [s_{l-1}^\sigma u_l s_l] \quad \text{and} \quad [t_l^\sigma v_l t_{l-1}] [t_{l-1}^\sigma v_{l-1} t_{l-2}] \cdots [t_1^\sigma v_1 t_0]$$

of R' and S' , respectively. Hence

$$[s_0^\sigma u_1 s_1] [s_1^\sigma u_2 s_2] \cdots [s_{l-1}^\sigma u_l s_l] =_{F(X^\sigma)} [t_0^\sigma v_1 t_1] [t_1^\sigma v_2 t_2] \cdots [t_{l-1}^\sigma v_l t_l],$$

and these words are P -reduced. So

$$\begin{aligned} [s_0^\sigma u_1] u_2 \cdots u_{l-1} [u_l s_l] &=_{U(P)} [s_0^\sigma u_1 s_1] [s_1^\sigma u_2 s_2] \cdots [s_{l-1}^\sigma u_l s_l] \\ &=_{U(P)} [t_0^\sigma v_1 t_1] [t_1^\sigma v_2 t_2] \cdots [t_{l-1}^\sigma v_l t_l] \\ &=_{U(P)} [t_0^\sigma v_1] v_2 \cdots v_{l-1} [v_l t_l] \\ &=_{U(P)} ([t_l^\sigma v_l] v_{l-1} \cdots v_2 [v_1 t_0])^{-1}. \end{aligned}$$

So e will be found by **FindEdges**(R, S).

For the complexity claims, notice that there are $O(r^2 |X|^2)$ decorated locations in Step 2 of **FindEdges**, that $l \leq r$, that for each l we consider $O(|X|^2)$ pairs (s_l, t_l) of interleaving elements, and that

$$[s_0^\sigma u_1] u_2 \cdots u_{l-1} [u_l s_l] [t_l^\sigma v_l] v_{l-1} \cdots v_2 [v_1 t_0] =_{U(P)} 1$$

can be tested in time $O(l) = O(r)$, by Corollary 2.10. \square

For each decorated place $\mathbf{P} = (R(i, a, b, s), c, C)$, we compute a list **OneStep**(\mathbf{P}) of decorated places that are 1-step reachable from \mathbf{P} , together with an upper bound on the corresponding step curvature, as follows.

If $C = \mathbf{R}$ then for each decorated place $\mathbf{Q} = (R(i+1, b, d, t), x, C')$, for each $y \in X$ such that y intermults with $[t^\sigma b^\sigma s]$, and for each \mathcal{V} -vertex ν_2 of colour C' with a \mathcal{V} -edge from (b, d, t, \mathbf{G}) to ν_2 , we let $\chi_1 = \mathbf{Blob}(y, [t^\sigma b^\sigma s], c)$ and $\chi_2 = \mathbf{Vertex}((y, [t^\sigma b^\sigma s], \mathbf{R}), (b, d, t, \mathbf{G}), \nu_2)$. We include $(\mathbf{Q}, 1, \chi_1 + \chi_2)$ in **OneStep**(\mathbf{P}).

If $C = \mathbf{G}$ then we use those 5-tuples in the list $L_{R,S}$ with first entry $R(i, a, b, s)$ to locate the possible decorated places $\mathbf{P}' = (R(j, d, e, s_l), c, C')$ that can be reached from \mathbf{P} by a single consolidated edge. For each such 5-tuple, the fifth data item specifies the length of this edge, the second identifies the location of \mathbf{P}' , and the fourth identifies the \mathcal{V} -vertex ν_1 , in the notation of Case $C = \mathbf{G}$ in Algorithm 7.16. Furthermore, the component c of \mathbf{P} must be equal to $[t^\sigma v_0 u]$ for some $u \in P$, where $R_1(j, v_1, v_0, t)$ is the second entry of the 5-tuple. Otherwise, Case \mathbf{G} is as in Algorithm 7.16.

Here is an overall summary of **RSymIntVerify**.

Procedure 7.30. **RSymIntVerify** ($\mathcal{P} = \langle X \mid V_P \mid \mathcal{R} \rangle, \varepsilon$):

- Step 1 For all $a, b \in X$, use V_P to test whether $(a, b) \in D(P)$, and store the result. If $(a, b) \notin D(P)$ then compute and store the interleave set $\mathcal{I}(a, b)$.
- Step 2 Express each relator $r \in \mathcal{R}$ as a power w^k and find all decorated locations on w . Store the pre- and post-interleave sets of each decorated location.
- Step 3 For each decorated location, find all corresponding decorated places.
- Step 4 Compute the decorated vertex graph, \mathcal{V} .
- Step 5 Use \mathcal{V} to create the **Vertex** function.

- Step 6 The set of \mathcal{R} -letters is all elements of X of the form $[s^\sigma bu]$, where $R(i, a, b, s)$ is a decorated location and $u \in \text{Post}(R(i))$. Create the **Blob** function.
- Step 7 For each $R, S \in \mathcal{R}$, use **FindEdges** to compute the list $L_{R,S}$ of possible consolidated edges between faces with labels in $\mathcal{I}(R)$ and $\mathcal{I}(S)$. Hence, for each decorated place $\mathbf{P} = (R(i, a, b, s), c, C)$, construct the list **OneStep**(\mathbf{P}).
- Step 8 From each decorated start place \mathbf{P}_s , run **RSymVerifyAtPlace**($\mathbf{P}_s, \varepsilon$) (using decorated places rather than places). If it returns **fail** and a list L , then return **fail** and L , otherwise do nothing.
- Step 9 Return **true**.

It is clear that **RSymIntVerify** runs in polynomial time, although with a higher time complexity than that of **RSymVerify**.

8 RSym and the word problem

Suppose that **RSym** succeeds on a presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ for a group G . We shall show in this section that this leads to a linear time algorithm for solving the word problem in G , which can be made into a practical algorithm in many examples. We remind the reader that all newly defined terms, notation and procedures are listed in the Appendix.

One approach to solving the word problem is to use the following result.

Proposition 8.1. *Let G be defined by the pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$, and let \mathcal{P}_G be the standard group presentation of G . Let r be the length of the longest relator in $V_P \cup \mathcal{R}$. Suppose that, for some constant λ , the Dehn function $D(n)$ of \mathcal{P}_G satisfies $D(n) \leq \lambda n$ for all $n \geq 0$. Then any word w over X with $w =_G 1$ has a subword of length at most $384\lambda r(r-1) + 64$ that is not geodesic.*

Proof. It is shown in the proof of [14, Theorem 6.5.3] that G is hyperbolic and that all geodesic triangles in its Cayley graph of G are δ -slim with $\delta \leq 96\lambda r^2 + 4$. In fact the result of [14, Lemma 6.5.1] can easily be improved from $\text{area}(\Delta) \geq mn/l^2$ to $\text{area}(\Delta) \geq 4mn/l(l-1)$, which results in the improved bound $\delta \leq 24\lambda r(r-1) + 4$. It is proved in [14, Theorem 6.1.3] that all geodesic triangles in the Cayley graph are 4δ -thin, and then [14, Theorem 6.4.1] implies that any word w with $w =_G 1$ must contain a non-geodesic word of length at most 16δ , which is at most $384\lambda r(r-1) + 64$. \square

We proved in Theorem 6.13 that if **RSym** succeeds with constant ε , then the pregroup Dehn function $\text{PD}(n) \leq \lambda_0 n$, where the constant λ_0 depends only on r and ε . It follows from this and Lemma 5.8 that $D(n) \leq \lambda n$, with $\lambda = r\lambda_0 + 1/2$. So we can apply Proposition 8.1, and compute $\gamma = 384\lambda r(r-1) + 64$ explicitly, which is typically a moderately large but not a huge number. We can therefore solve the word problem in linear time using a Dehn algorithm, provided that we can solve it for words of length at most γ , which can in principal be accomplished in constant time using a brute force algorithm which tests all products of conjugates of up to $D(n)$ relators. This is clearly impractical. In practice, one possibility is to use **KBAG** for this purpose, but that is only possible if **KBAG** can compute the automatic structure, which may not be feasible, particularly for examples with large numbers of generators.

An alternative approach to solving the word problem, which succeeds in many examples, is to use the success of **RSym** directly to produce a linear-time word problem solver **RSymSolve**, and the main purpose of this section is to describe how to do that. It is not guaranteed that such a solver can be constructed, even when **RSym** succeeds, but the attempted construction of **RSymSolve** takes low-degree polynomial time.

The remainder of the section is structured as follows: first, in Definition 8.2 we define an extra condition that **RSym** may satisfy, which guarantees that **RSymSolve** works. Then

we present a procedure called **VerifySolver**, which tests for this extra condition. Finally we describe the algorithm **RSymSolve**, which solves the word problem, and prove its correctness and complexity. It is similar in spirit to a Dehn algorithm, with complications arising from the fact that we work over $\mathcal{I}(\mathcal{P})$ whilst only storing rewrites arising from $V_P \cup \mathcal{R}$, and that we need to work with P -reductions rather than just free reductions.

Definition 8.2. *RSym verifies a solver for $\mathcal{I}(\mathcal{P})$ if, for any green boundary face f in any $\Gamma \in \mathcal{D}$ with $\kappa_\Gamma(f) > 0$, the removal of f shortens $\partial(\Gamma)$.*

We assume in this section that **RSym** has succeeded. In fact, it is possible to use **RSym** to solve the word problem under somewhat weaker hypotheses: see Remark 8.11.

The procedure **VerifySolver** seeks to check that **RSym** verifies a solver for $\mathcal{I}(\mathcal{P})$. It is very similar to the main **RSym** tester, except that some of the places are on $\partial(\Gamma)$. We shall assume that **RSymVerify**(\mathcal{P}, ε) has returned **true** for some $\varepsilon > 0$, and that the data computed have been stored.

We describe **VerifySolver** only for the case where $\mathcal{I}(\mathcal{R}) = \mathcal{R}$: the modifications necessary when $\mathcal{I}(\mathcal{R}) \neq \mathcal{R}$ are straightforward. We shall describe the word problem solver **RSymSolve** in the general case, when we do not assume that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$: as we shall explain in Remark 8.10 below, if **VerifySolver** succeeds then we can use a standard Dehn algorithm when $\mathcal{I}(\mathcal{R}) = \mathcal{R}$.

Procedure 8.3. **VerifySolver**($\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$):

Step 1 For each $R \in \mathcal{R}$ do

(a) For each place \mathbf{P}_s on R do

(i) If **VerifySolverAtPlace**(\mathbf{P}_s) returns **fail** then return **fail**.

Step 2 Return **true**.

For a start place $\mathbf{P}_s = (R(i, a, b), c, C)$ on a face f labelled by R , **VerifySolverAtPlace** works along all possible sequences of internal edges of f starting at \mathbf{P}_s , bounding the resulting curvature of f . We seek to show that if $\kappa_\Gamma(f) > 0$, then these internal edges take up less than half of $\partial(f)$.

The vertex where **VerifySolverAtPlace** terminates is on $\partial(\Gamma)$, so need not be a place in the sense of Definition 7.4, as it need not be instantiable for any choice of extra letter $c \in X$. We now rectify this.

Definition 8.4. A *terminal place* is a triple $(R(i, a, b), \text{terminal}, G)$ where $R(i, a, b)$ is a location. A terminal place is green, and has no extra letter. For use in this section, Definition 7.15 should be modified, to say that no place is 1-step reachable *from* a terminal place.

For the rest of this section, a *place* may be terminal, unless specified otherwise. Before running **VerifySolver** we re-run **ComputeOneStep**(\mathbf{P}) (Algorithm 7.16), to also find the one-step reachable terminal places from each non-terminal place \mathbf{P} . The curvature values χ of the triples (\mathbf{Q}, l, χ) in **OneStep**(\mathbf{P}), when \mathbf{Q} is terminal, are as follows (we shall justify them in the proof of Theorem 8.6).

$C = R$. We set χ to be the maximum curvature given to a face labelled R by a boundary red blob at \mathbf{P} (calculated using Lemma 7.10 with $|\delta(B) \cap \delta(\Gamma)| \geq 1$).

$C = G$. If $\mathbf{P}' = \mathbf{Q}$ we set $\chi = -1/4$. If $\mathbf{P}' \neq \mathbf{Q}$ then we let χ be the maximum sum of the curvature given to a face labelled R by the vertex at \mathbf{P}' and a boundary red blob at \mathbf{P}' .

For a non-terminal place \mathbf{P}_s on R , **VerifySolverAtPlace**(\mathbf{P}_s) makes a list of 4-tuples (\mathbf{Q}, l, t, ψ) . The first three components represent a place \mathbf{Q} at distance l from \mathbf{P}_s along R that can be reached from \mathbf{P}_s in t steps, where if \mathbf{P}_s is green then the first “step” consists of only the boundary vertex. The final component ψ is $1 + \chi$, where χ is the largest

possible curvature given to R by these t steps: unlike in Subsection 7.6 we do not adjust for the length of the step. The algorithm **VerifySolverAtPlace**(\mathbf{P}_s) returns **fail** if there is a sequence of neighbouring green faces and red blobs, starting at \mathbf{P}_s and ending at a terminal place, that occupies at least half of $\partial(f)$ and from which f receives more than -1 of curvature. It returns **true** otherwise.

In the description below, the meaning of *including* an entry in the list L is the same as in Procedure 7.18.

Procedure 8.5. **VerifySolverAtPlace**($\mathbf{P}_s = (R(i, a, b), c, C)$):

- Step 1 Let $n := |R|$, and initialise $L := []$.
- Step 2 If $C = \mathbf{G}$ then include $(\mathbf{P}_s, 0, 1, 3/4)$ in L .
- Step 3 If $C = \mathbf{R}$ then, for each non-terminal place \mathbf{P}_1 at distance 1 from \mathbf{P}_s , calculate the maximum value of $\chi = \chi(B, f, \Gamma) + \chi(v, f, \Gamma)$, where B is boundary red blob between \mathbf{P}_s and \mathbf{P}_1 , and v is the vertex at \mathbf{P}_1 . Include $(\mathbf{P}_1, 1, 1, 1 + \chi)$ in L .
- Step 4 For $i := 1$ to 3, for each $(\mathbf{P}, l, i, \psi) \in L$, and for each $(\mathbf{Q}, l', \chi) \in \mathbf{OneStep}(\mathbf{P})$ do:
 - (a) If $l + l' < n/2$ and \mathbf{Q} is not terminal then let $\psi' := \psi + \chi$;
if $\psi' > 0$ then include $(\mathbf{Q}, l + l', i + 1, \psi')$ in L .
 - (b) If $l + l' \geq n/2$, \mathbf{Q} is terminal, and $\psi + \chi > 0$, then return **fail** and L .
- Step 5 Return **true**.

Note that we do nothing in Step 4 if neither of the specified conditions hold.

Theorem 8.6. *Assume that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$. If **VerifySolver** returns **true**, then **RSym** verifies a solver for \mathcal{P} . The procedure **VerifySolver** runs in polynomial time.*

Proof. Let f be a boundary green face of a diagram $\Gamma \in \mathcal{D}$ such that $\kappa_\Gamma(f) > 0$, and assume that at most half of the edges of f are contained in $\partial(\Gamma)$. Let the label of f be $R \in \mathcal{R}$. We show that **VerifySolverAtPlace** returns **fail** for at least one start place \mathbf{P}_s on R .

By Lemma 7.5, a boundary vertex v with $\delta_G(v, \Gamma) \geq 3$ satisfies $\chi(v, f, \Gamma) \leq -1/4$, so Step 2 of **VerifySolverAtPlace**, and the value of χ in **OneStep**(\mathbf{P}) when \mathbf{P} is green and $\mathbf{Q} = \mathbf{P}'$ is terminal, correctly bound $\chi(v, f, \Gamma)$ when the first or last edge of $\partial(f) \setminus \partial(\Gamma)$ is incident with a green internal face.

Step 3 of **VerifySolverAtPlace**, and the remaining cases of **OneStep**(\mathbf{P}) when \mathbf{Q} is terminal, bound $\chi(B, f, \Gamma)$ as if the B is on the boundary. To see why this is correct, first notice that the label of any non-boundary red blob is permissable as a label of a boundary red blob. Let v be the vertex at \mathbf{P}_s , and let B be a red blob at \mathbf{P}_s with boundary length l and area t (the case where the red blob is at the end of $\partial(f) \setminus \partial(\Gamma)$ is equivalent). By Lemma 4.12, $\chi(B, f, \Gamma)$ is maximised by assuming that B is simply-connected, in which case $l = t + 2$ since Γ is green-rich. If B has a boundary edge at \mathbf{P}_s , then by Lemmas 6.7 and 6.8, $\chi(v, f, \Gamma) = 0$ and $\chi(B, f, \Gamma) \leq \frac{-t}{2(t+1)}$. If B has no boundary edge at \mathbf{P}_s , then $\chi(v, f, \Gamma) \leq -1/4$ and $\chi(B, f, \Gamma) \leq \frac{-t}{2(t+2)}$. For all t

$$\frac{-t}{2(t+1)} > \frac{-t}{2(t+2)} - \frac{1}{4}$$

so $\chi(v, f, \Gamma) + \chi(B, f, \Gamma) \leq \frac{-t}{2(t+1)}$. Hence Step 3 of **VerifySolverAtPlace** and **OneStep**(\mathbf{P}) correctly bound the curvature received by R when the first or last edge of $\partial(f) \setminus \partial(\Gamma)$ is incident with with a red blob. We showed in Lemma 7.17 that the **OneStep** lists correctly bound all other step curvatures.

We showed in Lemma 6.9 that $\partial(f) \cap \partial(\Gamma)$ consists of at most one consolidated edge together with at most one isolated vertex v , and that the internal edges of $\partial(f)$ form a path.

Therefore each place on R at a vertex along these internal edges (except for the last one) has a face of Γ instantiating it, and so is not terminal. Therefore **VerifySolverAtPlace** is correct to require in Steps 3 and 4(a) that all places are non-terminal.

It remains only to show that $\partial(f) \setminus \partial(\Gamma)$ consists of at most three steps along R , after the initial one (which may not formally be a “step”). Assume first that $f \cap \partial(\Gamma)$ contains an isolated vertex v . Then we can decompose $\partial(f)$ into $v_1, \beta_1, v, \beta_2, v_2, \beta_3$, where v_1, v_2, v are vertices on $\partial(\Gamma)$, and each β_i is a sequence of edges and vertices such that $\beta_3 \subset \partial(\Gamma)$, and β_1 and β_2 are internal to Γ (see Figure 3). As we have just seen, the curvature given to f by the vertices and blobs at the beginning of β_1 and the end of β_2 sums to at most $-1/2$. We showed in Lemma 6.9 that f is incident with no red blobs at v , and that $\delta_G(v, \Gamma) \geq 4$, so $\chi(v, f, \Gamma) \leq -1/3$.

Since $\kappa_\Gamma(f)$ is assumed to be positive, by Lemmas 7.5 and 7.10 the face f is adjacent to only these two green internal faces (and possibly some red blobs at v_1 and v_2), and so **VerifySolverAtPlace** will return **fail** when $i \leq 2$, with \mathbf{P}_s a place at v_1 .

Next assume that $f \cap \partial(\Gamma)$ consists of a single consolidated edge e . Write the boundary of f as e, v_1, β, v_2 , where β is a sequence of edges and vertices internal to Γ . The vertices v_1 and v_2 , together with any incident red blobs, give at most $-1/2$ of curvature to f . Excluding the red blobs that might be incident with v_1 and v_2 , by Lemmas 7.5 and 7.10 the path β can be incident with at most two red blobs, or at most two vertices of green degree greater than two, or exactly one of each. Hence f is adjacent to at most three internal green faces, and **VerifySolver** will return **fail**.

The complexity claims follow as in the proof of Theorem 7.22. \square

We now show how to solve the word problem, provided that **RSym** verifies a solver. First we show that any word $w = x_1 \dots x_n$ can be cyclically P -reduced in linear time.

Proposition 8.7. *Let $w = x_1 \dots x_n \in X^*$. On input w and $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$, a cyclically P -reduced word w' that is conjugate in $U(P)$ to w can be found in time $O(|w| - |w'|) = O(n)$.*

Proof. The word problem over $U(P)$ can be solved in linear time by Corollary 2.10, so without loss of generality we may assume that $x_1 \dots x_n$ is P -reduced.

It remains to consider cyclic P -reduction. Assume we have two pointers, **start**, initially pointing at x_1 and **end**, initially pointing at x_n . We check whether $(x_n, x_1) \in D(P)$. If $[x_n x_1] =_P 1$, we move **start** to x_2 and **end** to x_{n-1} . If $[x_n x_1] =_P a \neq 1$, we replace x_1 by a , and move **end** to x_{n-1} . We continue this process until the letters s and t pointed to by **end** and **start** satisfy $(s, t) \notin D(P)$. We then test $(t, u) \in D(P)$, where u is the letter after the one to which **start** points. If not, we are done. If $tu =_P 1$ then we move **start** forwards by two letters. If $tu =_P a \neq 1$ then we replace u by a and move **start** forward by one letter. We continue moving **start** and **end** towards the middle of w until no further reductions are possible. We then return the word reading from **start** to **end**.

Let w' be the resulting word, with $m = |w| - |w'|$. Then each pointer moves $O(m)$ times, and $O(m)$ products in P are calculated. \square

We compute a list \mathcal{L} , whose entries are pairs of words $(u, v) = (u_1 \dots u_k, v_1 \dots v_l) \in X^* \times X^*$, where $[s^\sigma u_1]u_2 \dots u_{k-1}[u_k t]([s^\sigma v_1]v_2 \dots v_{l-1}[v_l t])^{-1}$ is a cyclic conjugate of some $R \in \mathcal{R}^\pm$ for some $s, t \in P$, and $k = \lceil (|R| + 1)/2 \rceil$.

In the light of Proposition 8.7 we may assume that the input to **RSymSolve** is a cyclically P -reduced word $w = x_1 \dots x_n \in X^*$. Let r be the length of the longest relator in \mathcal{R} . In the description below, we interpret all subscripts cyclically, so that $x_{n+1} = x_1$.

Algorithm 8.8. **RSymSolve**($w = x_1 \dots x_n$):

Step 1 Store w as a doubly-linked list: each letter has a pointer to the letter before it, and the letter after it.

- Step 2 Set $\alpha := 1$.
- Step 3 For $\alpha \leq i \leq n$, search for $m \in \{1, \dots, \lceil (r+1)/2 \rceil\}$, $a \in \mathcal{I}(x_{i-1}, x_i)$, $b \in \mathcal{I}(x_{i+m-1}, x_{i+m})$ and $(u, v) \in \mathcal{L}$ such that $[a^\sigma x_i]x_{i+1} \dots x_{i+m-2}[x_{i+m-1}b] =_{U(P)} u$.
- (a) Let $i, m, a, b, v := v_1 \dots v_l$ be the first such found, if any.
 - (b) If none such exist then $w \neq_G 1$. Terminate and return **false**.
- Step 4 Replace x_{i-1} by $[x_{i-1}a]$ and replace x_{i+m} by $[b^\sigma x_{i+m}]$. Put a pointer **CutStart** to the new x_{i-1} and a pointer **CutEnd** to the new x_{i+m} . Store v as a doubly-linked list, with pointers **NewStart** to v_1 and **NewEnd** to v_l .
- Step 5 P -reduce at the beginning of v : If $[x_{i-1}v_1] =_P s \neq 1$, then replace v_1 by s , and move **CutStart** to x_{i-2} . If $[x_{i-1}v_1] =_P 1$, then move **NewStart** and **CutStart** to v_2 and x_{i-2} .
- Step 6 Repeat Step 5 until one of the following: no further reductions are found; **CutStart** should be moved back past x_1 ; **NewStart** should be moved forward past v_l .
- Step 7 Provided that there is at least one letter left in v , perform Steps 5 and 6 (with the appropriate pointers) to P -reduce at the end of v .
- Step 8 Update the links in the list describing w so that whatever remains of v is now inserted into the correct place in $x_1 \dots x_n$, yielding a word w_1 .
- Step 9 Cyclically P -reduce w_1 , as in the proof of Proposition 8.7, yielding a word w_2 . If w_2 is empty, then terminate and return **true**.
- Step 10 Let j be the position in w_2 to which **CutStart** points, and let $\alpha := \max\{1, j - \lceil (r+1)/2 \rceil + 1\}$. Replace n by $|w_2|$, and go to Step 3 with w_2 in place of w .

(When returning to Step 3, we start the search for the next rewrite at x_α , as earlier untouched letters will still not be eligible for rewriting.)

Theorem 8.9. *Let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a pregroup presentation for a group G , such that **RSym** succeeds on \mathcal{P} . If **VerifySolver** succeeds on $\mathcal{I}(\mathcal{P})$ then for all $n \in \mathbb{N}$, and for all $x_1 \dots x_n \in X^*$, the algorithm **RSymSolve**($x_1 \dots x_n$) correctly tests whether $x_1 \dots x_n =_G 1$ in time $O(n)$.*

Proof. Let $w = x_1 \dots x_n \in X^*$. By Proposition 8.7 in $O(n)$ we can replace w by a word $w_1 = y_1 \dots y_k$ that is cyclically P -reduced. By Theorem 6.13, since **RSym** succeeds on \mathcal{P} , for some $w_2 \in \mathcal{I}(w_1)$ there exists a diagram $\Gamma \in \mathcal{D}$ with boundary word w_2 . By Lemma 6.9, either Γ consists of a single face f_1 , which must be green and have curvature $+1$, or there are at least two boundary faces f_1 and f_2 of Γ with positive curvature. In this second case, since **VerifySolver** succeeds, the faces f_1 and f_2 each have more than half of their boundary length as a continuous subword of the boundary of Γ .

Hence a rewrite applies to at least one subword $z_i \dots z_{i+m-1}$ of w_2 . Such a subword is equal in $U(P)$ to $[a^\sigma y_i] \dots [y_{i+m-1}b]$ for some $a \in \mathcal{I}(y_{i-1}, y_i)$ and $b \in \mathcal{I}(y_{i+m-1}, y_m)$, and so will be found by **RSymSolve**.

On input a cyclically P -reduced word w of length n , **RSymSolve** runs $O(n)$ tests of equality in $U(P)$ of words of the form $a^\sigma w' b$, where $w' = t_1 \dots t_m$ is a subword of w with $m \leq (r+1)/2$ such that $(a^\sigma, t_1) \in D(P)$ and $(t_m, b) \in D(P)$, with the first entry of each pair in \mathcal{L} . It also tests for $O(n)$ (cyclic) P -reductions. Every time a letter of w (or of a substring v for replacement into w) is changed, it is due to a shortening of $|w|$, so at most $O(n)$ letter replacements occur. \square

Remark 8.10. If **VerifySolver** succeeds and $\mathcal{I}(\mathcal{R}) = \mathcal{R}$, then we can replace **RSymSolve** by a standard Dehn algorithm using the length reducing rewrite rules derived from $V_P \cup \mathcal{R}$. The presence of the rules from V_P ensures that we reduce our input word w to a P -reduced

word. There is no need to carry out any additional cyclic reduction or cyclic P -reduction on w .

This is because cyclic (P -)reduction might delete some letters at the beginning and end of w , and might insert a single new letter at the beginning of w . If the resulting cyclically P -reduced word w' was equal to the identity in G , then a diagram $\Gamma \in \mathcal{D}$ for w' would either consist of a single green face, or have at least two green regions with more than half of their length on the boundary. In the former case, all but one letter of w' is a subword of the original word w . In the latter case, the label of the intersection of at least one of these two regions with the boundary of Γ would be a subword of w . So this subword would be reduced in length by the application of one of the rewrite rules.

Although a standard Dehn algorithm is no faster than **RSymSolve** in terms of complexity (both are linear in the length of w) it has the advantage that it can be implemented efficiently using a two stack model, as described for example in [4].

An improvement to **VerifySolver** is sometimes possible: for example, when P is the standard pregroup for a free product of finite and free groups.

Remark 8.11. Consider the situation where if (a, b) is an intermult pair, then $(a, b) \in D(P)$. Note that this implies in particular that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$. In this case we design an upgrade to **VerifySolver**, which we call **VerifySolverTrivInt**.

Let Γ be a diagram in \mathcal{D} and let f be a boundary face of Γ with $\kappa_\Gamma(f) > 0$, with k boundary edges and l internal edges. By Lemma 6.9, the edges in $\partial(f) \setminus \partial(\Gamma)$ form a path e_1, \dots, e_l , say.

If exactly one of e_1 or e_l is incident to a red blob whose next edge is on the boundary, then deleting f leaves a red blob B with two edges appearing consecutively on $\partial(\Gamma)$. By Lemma 4.16 the labels a and b of these two edges intermult. Hence $(a, b) \in D(P)$, so the new boundary word can be P -reduced, deleting at least one red triangle. Hence, deleting f followed by P -reduction shortens $\partial(\Gamma)$ by at least $(k + 1) - l$ edges. If both e_1 and e_l are incident with boundary red blobs, then deleting f , followed by P -reduction, shortens $\partial(\Gamma)$ by at least $(k + 2) - l$ edges.

Hence one can produce a more powerful algorithm, **VerifySolverTrivInt**, by modifying **VerifySolverAtPlace** and the **OneStep** values for terminal places. The list L from **VerifySolverAtPlace** should now contain two types of entries, those that record places on paths that start at a boundary red blob, and those that do not.

In Step 4(a) of **VerifySolverAtPlace**, we replace $l + l' < n/2$ by $l + l' < (n + 2)/2$.

In Step 4(b) of **VerifySolverAtPlace**, if there is a boundary red blob at either \mathbf{P}_s or just before \mathbf{Q} , then failure is *only* reported if $l + l' \geq (n + 1)/2$. If \mathbf{P}_s and the edge before \mathbf{Q} are *both* incident with boundary red blobs, then failure is only reported if $l + l' \geq (n + 2)/2$.

When using **VerifySolverTrivInt**, appropriate additions need to be made to the list \mathcal{L} of rewrites for **RSymSolve** to adjoin the letters which can appear on such boundary red blobs. We shall refer to the enhanced version as **RSymSolveTrivInt**.

As in Remark 8.10, in many situations (we omit the details) we can use a standard Dehn algorithm in place of **RSymSolveTrivInt**, with the modified list \mathcal{L} .

We can use the success of **VerifySolver** or **VerifySolverTrivInt** to lower our bound on the Dehn function of G . Recall Definition 5.5 of the pregroup Dehn function.

Proposition 8.12. *Let $\text{PD}(n)$ be the pregroup Dehn function of \mathcal{P} . If both **RSym** and **VerifySolver** succeed, then $\text{PD}(n) \leq n$. If **RSym** and **VerifySolverTrivInt** succeed, then $\text{PD}(n) \leq 3n$.*

Proof. The first claim is clear, since **RSymSolve** is a variation of a Dehn algorithm. For the second, notice that the removal of at most three faces from each diagram results in a shortening of the boundary word. \square

Remark 8.13. Assume that we know that all V^σ -letters are nontrivial in G . Then with a little effort one may show that we can deduce the same bounds on the pregroup Dehn function of G from the success of `RSymSolve` and `RSymSolveTrivInt`, even when `RSym` fails, provided that `RSym` is able to show that all green faces at dual distance at least two from the external face have *non-positive* curvature. The assumption that no V^σ -letters are trivial in G means that all diagrams are loop-minimal, which permits us to use Proposition 3.17, and to deduce that if $w =_G 1$ then there is a diagram $\Gamma \in \mathcal{D}$ with boundary word w . We excluded this case from our earlier analysis as it gave no immediate upper bound on the Dehn function, but `RSymSolve` and `RSymSolveTrivInt` together provide such a bound.

9 Applications of `RSym`

In this section we shall first show that `RSym` generalises several small cancellation conditions. We then show how `RSym` can be verified by hand to prove the hyperbolicity of various infinite families of presentations. This is an advantage over the algorithm based on the theory of automatic groups that is used by the `KB MAG` package, which can only handle individual groups and is not susceptible to hand-calculation. Finally, we discuss possible further applications of `RSym`.

We remind the reader that all new terms, notation and procedures are listed in the Appendix.

9.1 Small cancellation conditions

As a first example of the applicability of `RSym`, we consider various small cancellation conditions, thereby recovering the result proved in [7, Corollary 3.3]. Furthermore, in many cases `RSymSolve` solves the word problem.

Theorem 9.1. *Let $\mathcal{Q} = \langle Y \mid \mathcal{R} \rangle$ be a group presentation for a group G , satisfying $C(p) - T(q)$ for some $(p, q) \in \{(7, 3), (5, 4), (4, 5)\}$. Then `RSym` succeeds on \mathcal{Q} with $\varepsilon = -1/6$, $-1/4$ and $-1/5$, respectively. If $(p, q) = (3, 7)$, then `RSym` succeeds at level 2 with $\varepsilon = -1/14$. In all of these cases, G is hyperbolic.*

Proof. We let $X = Y \dot{\cup} \{y^\sigma : y \in Y\}$ and set $V_P = \emptyset$, just as in Example 2.4. Then let $\mathcal{P} = \langle X \mid \emptyset \mid \mathcal{R} \rangle$, so each coloured diagram over \mathcal{P} is also a diagram over \mathcal{Q} . Let Γ be a reduced coloured diagram; notice that each face of Γ is green, so $\Gamma \in \mathcal{D}$. By [16, Chapter V, Lemma 2.2], the fact that \mathcal{Q} satisfies $C(p) - T(q)$ means that all non-boundary faces of Γ have at least p edges (and hence at least p vertices), and all non-boundary vertices v satisfy $\delta(v, \Gamma) = \delta_G(v, \Gamma) \leq q$.

First assume that $(p, q) \in \{(7, 3), (5, 4), (4, 5)\}$, and let f be a non-boundary face of Γ . A vertex v of f that is not on $\partial(\Gamma)$ satisfies

$$\chi(v, f, \Gamma) = \frac{2 - \delta(v, \Gamma)}{2\delta(v, \Gamma)} = \frac{1}{\delta(v, \Gamma)} - \frac{1}{2} \leq \frac{1}{q} - \frac{1}{2},$$

by Lemma 6.7(ii). Since f is a non-boundary face, a vertex v of f on $\partial(\Gamma)$ has degree at least 4 and so $\chi(v, f, \Gamma) \leq -1/3$ by Lemma 7.5. The claim that `RSym` succeeds, and the stated values of ε , now follow from the fact that f has at least p incident vertices of degree at least q . Since `RSym` succeeds at level 1, it follows from Theorem 6.13 that G is hyperbolic.

So suppose instead that $(p, q) = (3, 7)$, so that each non-boundary face has at least 3 incident edges, each of degree at least 7. If f has only three incident vertices, all boundary vertices of degree exactly four, then $\kappa_\Gamma(f) = 0$, so `RSym` fails at level 1. We therefore

apply **RSym** at level 2 to Γ , and let f be an internal face at dual distance at least 2 from $\partial(\Gamma)$. Then $\delta(v, \Gamma) \geq 6$ for all vertices v of f on $\partial(\Gamma)$, so Lemma 7.5 tells us that $\chi(v, f, \Gamma) \leq -2/5$. Since the non-boundary vertices v of f satisfy $\chi(v, f, \Gamma) \leq -5/14$ in this case, and $-2/5 < -5/14$, we conclude that **RSym** succeeds at level 2 with $\varepsilon = -1/14$ as claimed. Since $V_P = \emptyset$, there are no V^σ -letters, so G is hyperbolic by Theorem 6.13. \square

With metric small cancellation conditions, **RSymSolve** solves the word problem.

Theorem 9.2. *Let \mathcal{Q} be a group presentation satisfying $C'(1/6)$ or $C'(1/4) - T(4)$. Then both **RSym** and **VerifySolver** succeed on \mathcal{Q} .*

Proof. The success of **RSym** follows from Theorem 9.1. **VerifySolver** considers up to three steps from each place on each relator, with the vertices at each end giving curvature at most $-1/4$, and the intermediate vertices giving curvature at most $1/q - 1/2$, where $q \in \{3, 4\}$. Hence for $C'(1/6)$ a boundary face f with $\kappa_\Gamma(f) > 0$ has at most three internal consolidated edges, and for $C'(1/4) - T(4)$ it has at most two. By Lemma 6.9 these internal consolidated edges are contiguous, and comprise less than half of the length of the relator. \square

Our second example considers the generalisation of small cancellation to amalgamated free products, as described in [16, Chapter V §11]. Let X_1, \dots, X_m be finite groups with proper subgroups $A_i \leq X_i$, let $A = A_1$, and let $\psi_i : A \rightarrow A_i$ be isomorphisms. Let $F = \langle *X_i : A = \psi_i(A_i) \rangle$ be the free product of the X_i , amalgamated over the A_i .

A *normal form* for $g \in F \setminus \{1\}$ is any expression $y_1 y_2 \cdots y_n$ such that $g =_F y_1 y_2 \cdots y_n$, each $y_i \in X_j$ for some j , successive y_i come from different X_j , and no y_i is in A unless $n = 1$. The length n , and the factors in which the y_i lie, are uniquely determined by g . An element $g \in F \setminus \{1\}$ with normal form $y_1 \cdots y_n$ is *cyclically reduced* if $n = 1$ or y_1 and y_n are in different factors, and *weakly cyclically reduced* if $n = 1$ or $y_n y_1 \notin A$. A product of normal forms $y_1 \cdots y_n x_1 \cdots x_m$ is *semi-reduced* if $y_n x_1 \notin A$ and neither normal form is a single element of A .

Let \mathcal{R} be a set of weakly cyclically reduced elements of $F \setminus A$. The *symmetrised* set $\widehat{\mathcal{R}}$ consists of all normal forms of all weakly cyclically reduced F -conjugates of elements of \mathcal{R}^\pm . A normal form $b \in F \setminus A$ is a *piece* if there exist distinct $R_1, R_2 \in \widehat{\mathcal{R}}$ such that $R_1 =_F b c_1$ and $R_2 =_F b c_2$, where c_1 and c_2 are normal forms, and the products $b c_1$ and $b c_2$ are both semi-reduced.

Notice that if $R_1 = x_1 \dots x_n$ and $R_2 = x_1 \dots x_k y_{k+1} \dots y_m$ are normal forms of elements of \mathcal{R} with $m, n \geq k + 2$, and x_{k+1} and y_{k+1} are in the same free factor X_i , then $x_1 \dots x_{k+1}$ is a piece. To see this, observe that we can write $y_{k+1} = x_{k+1} z$ for some $z \in X_i$. If $z \notin A$ then $x_1 \dots x_{k+1} \cdot z y_{k+2} \dots y_m$ is a semi-reduced product of two normal forms. If $z \in A$ then let $z_{k+2} = z y_{k+2} \notin A$, and notice that the product $x_1 \dots x_{k+1} \cdot z_{k+2} y_{k+3} \dots y_m$ is a semi-reduced product of two normal forms.

Definition 9.3. A symmetrised set $\widehat{\mathcal{R}}$ satisfies $\mathcal{C}'_{FA}(\lambda)$, where $\lambda \in \mathbb{R}_{>0}$, if

- (i) $|R| > 1/\lambda$ for all $R \in \widehat{\mathcal{R}}$;
- (ii) if $R \in \widehat{\mathcal{R}}$ is equal in F to a semi-reduced product bc , where b is a piece, and c is a normal form, then $|b| < \lambda|R|$.

We have not attempted to optimise the value of ε in the following result.

Theorem 9.4. *Let X_1, \dots, X_m be finite groups, let $F = \langle *X_i : A = \psi_i(A_i) \rangle$ be a free product with amalgamation, and let \mathcal{R} be a finite set of cyclically reduced elements of F such that $\widehat{\mathcal{R}}$ satisfies $\mathcal{C}'_{FA}(1/6)$.*

*Let $P = X_1 \cup \bigcup_{i>1} (X_i \setminus A_i)$, with products defined within each X_i but not across factors. Then P is a pregroup and **RSym** succeeds on the presentation $\langle (P \setminus 1)^\sigma \mid V_P \mid \mathcal{R} \rangle$ with $\varepsilon = 1/(2r)$.*

Proof. The fact that P is a pregroup and $U(P) = F$ is established in [22, 3.A.5.2]. The set $\mathcal{I}(\mathcal{R})$ consists of all normal forms of cyclic conjugates of elements of \mathcal{R} , and so is contained in $\widehat{\mathcal{R}}$. The set $\widehat{\mathcal{R}}$ also contains elements of F that are weakly cyclically reduced but not cyclically P -reduced, but these are not labels of faces in van Kampen diagrams (coloured or otherwise).

We shall use the approach of **RSymVerifyAtPlace** (Procedure 7.18) and consider each possible decomposition of a face f in a diagram $\Gamma \in \mathcal{D}$ that is labelled by $R \in \mathcal{I}(\mathcal{R})$ into steps. We shall consider the cumulative curvature of f , namely the curvature value ψ stored as the last entry of the 4-tuples in the list L created by **RSymVerifyAtPlace**. We shall show that this is negative after at most two steps, and hence $\kappa_\Gamma(f) \leq -1/2r$. Let the first step have length l , let χ_1 be the curvature of the first step, and let $n = |R|$. Then the cumulative curvature ψ_1 after the first step is

$$\psi_1 = \chi_1 + (1 + \varepsilon) \frac{l}{n} = \chi_1 + \frac{2r+1}{2r} \cdot \frac{l}{n}.$$

Each element of $\widehat{\mathcal{R}}$ has length at least 7, by the assumption that $\widehat{\mathcal{R}}$ satisfies $\mathcal{C}'_{FA}(1/6)$. Therefore $r \geq n \geq 7$, and in particular $(2r+1)/(2r) \leq 15/14$. By Lemma 7.17, the step curvature $\chi_1 \leq -1/6$, so if $l = 1$ then $\psi_1 \leq -2/147 < 0$, and we are done. Hence we may assume that $l > 1$, and in particular that the first edge is green. Let $S \in \mathcal{I}(\mathcal{R})$ be such that f and a face f_1 labelled by S are incident with a consolidated edge e labelled $x_1 x_2 \cdots x_k$. Notice that $n = 6k + m$ for some $m > 0$, by Condition $\mathcal{C}'_{FA}(1/6)$.

If the step consists just of the edge e , so that we are in Case 3(a) of Definition 7.15, then $l = k$. Hence

$$\psi_1 \leq -\frac{1}{6} + \frac{2r+1}{2r} \cdot \frac{k}{6k+m},$$

and a short calculation shows that this is always negative.

So assume that the step consists of e , ending at a vertex v_1 , and then a red edge (labelled x_{k+1}) between f and a blob B_1 , then a vertex v_2 . Hence $l = k + 1$, and a short calculation shows that $\frac{2r+1}{2r} \cdot \frac{k+1}{6k+m} < 3/10$, so we may assume that $\chi_1 > -3/10$. Notice that $\chi_1 = \chi(v_1, f, \Gamma) + \chi(B_1, f, \Gamma) + \chi(v_2, f, \Gamma)$, and so in particular $\delta_G(v_1, \Gamma) = 2$ (so f and f_1 are both adjacent to B_1), the length $|\partial(B_1)| \leq 4$, and $\delta_G(v_2, \Gamma) = 2$. In particular, $\chi_1 = \chi(B_1, f, \Gamma)$.

Let x_{k+1} lie in the free factor X_i . If the corresponding letter of S also lies in X_i , then R contains a piece of length $k + 1$. Hence $\frac{k+1}{6k+m} < 1/6$, so $m \geq 7$, and a short calculation shows that $\frac{2r+1}{2r} \cdot \frac{k+1}{6k+m} \leq \frac{2r+1}{2r} \cdot \frac{k+1}{6k+7} < 1/6$, and so ψ_1 is negative. Hence we may assume that the corresponding letter of S does not lie in X_i , so B_1 contains two triangles, and $\chi_1 = \chi(B_1, f, \Gamma) = -1/4$.

We are therefore done if $\frac{2r+1}{2r} \cdot \frac{k+1}{6k+m} < 1/4$, so assume otherwise. Since $\frac{2r+1}{2r} \leq 15/14$, a short calculation shows that $k = 1$ and $m \leq 2$, so n is 7 or 8. The assumption that $\widehat{\mathcal{R}}$ satisfies $\mathcal{C}'_{FA}(1/6)$ now implies that there are no pieces on R of length 2.

Let f_2 be the green face incident with B_1 and v_2 . Then the edge shared by f_2 and B_1 has label from X_i . Since one edge label of f_2 at v_2 is from X_i and no pieces have length 2, the faces f_2 and f cannot be edge-incident after v_2 . Since $\delta_G(v_2) = 2$, it follows that the next step is red, corresponding to a blob B_2 . The face f_2 is edge-incident with B_2 , so B_2 cannot be a single red triangle. Hence $|\partial(B_2)| \geq 4$. Thus the curvature χ_2 of this second step is at most $-1/4$, and so $\chi_1 + \chi_2 \leq -1/2$. However, the sum of the two step lengths is 3, and since $(15/14) \cdot (3/7) < 1/2$, the cumulative curvature after two steps is negative. \square

9.2 Families of presentations

We shall now consider some infinite families of presentations where we can show by hand that **RSym** succeeds. To help make our descriptions clear and concise, we shall not work through every step of **RSymVerify**, but just extract the parts that we need.

For our first family of examples we consider the triangle groups. The following result is well known, but it illustrates how we can use **RSym** to provide a straightforward proof.

Proposition 9.5. *Let $G = \langle x, y \mid x^\ell, y^m, (xy)^n \rangle$ with $2 \leq \ell \leq m \leq n$ and $1/\ell + 1/m + 1/n < 1$. Then **RSym** and **VerifySolverTrivInt** both succeed on a pregroup presentation of G , and so G is hyperbolic.*

Proof. Let \mathcal{P} be the first pregroup presentation for G from Example 2.12. Then $\mathcal{R}^\pm = \{R_1 := (xy)^n, R_2 := (x^{-1}y^{-1})^n =_{U(P)} (x_{\ell-1}y_{m-1})^n\}$. Let Γ be a σ -reduced coloured diagram over \mathcal{P} .

Suppose first that $\ell \geq 3$ and hence $n \geq 4$ and $|R_1| = |R_2| \geq 8$. Then it is not possible for two internal green faces of Γ to share an edge, and so there are no instantiable green places. Hence all steps are red and have length 1. Now each step curvature is at most $-1/6$ by Lemma 7.17, so each non-boundary face f of Γ satisfies $\kappa_\Gamma(f) \leq 1 - 8 \cdot \frac{1}{6} = -1/3$, and so **RSym** succeeds with $\varepsilon = 1/3$.

We now show that **VerifySolverTrivInt** succeeds when $\ell \geq 3$, and hence that **RSymSolveTrivInt** solves the word problem. Let f be a boundary face of Γ with $\kappa_\Gamma(f) > 0$. We must show that at least half of $\partial(f)$ is on $\partial(\Gamma)$. If f has at most one internal edge (which is red, so has length 1), then at least $2n - 1$ edges of f are boundary, so assume that f has at least two internal edges. Then the ends of at least two of these internal edges must intersect $\partial(\Gamma)$ non-trivially and, by Lemma 7.10, the steps corresponding to those edges have step curvature at most $-1/4$. Hence $\kappa_\Gamma(f) > 0$ implies that f has at most four internal edges, which must be contiguous by Lemma 6.9. Hence if $n > 4$ then more than half of $\partial(f)$ is on $\partial(\Gamma)$, and so **VerifySolver** succeeds. If $n = 4$, and f satisfies $\kappa_\Gamma(f) > 0$ and has four contiguous internal edges, then the first or last such edge is incident with a boundary red blob. If (a, b) is an intermult pair then $(a, b) \in D(P)$, so **VerifySolverTrivInt** succeeds.

So suppose for the remainder of the proof that $\ell = 2$, and so $m \geq 3$. Then on R_1 there is a single instantiable (non-terminal) green place $\mathbf{P}_1 = (R_1(1, y, x), y, \mathbf{G})$. Furthermore, the consolidated edges between two adjacent green faces in any diagram $\Gamma \in \mathcal{D}$ have length 1. So each step in a decomposition of R_1 has length at most 2.

Assume that $n \geq 7$, and notice in particular this holds when $m = 3$. Then $\kappa_\Gamma(R_1) \leq -1/6$, and so **RSym** succeeds with $\varepsilon = 1/6$. For **RSymSolveTrivInt** notice that the longest possible sequence of edges between a boundary face f with positive curvature and the interior of a diagram is length 7 (with label $y(xy)^3$) resulting in $\kappa_\Gamma(f) = -1/6$. So **VerifySolver** fails when $n = 7$, but this sequence of edges has a boundary red blob at each end, so **VerifySolverTrivInt** succeeds by Remark 8.11. Note that **VerifySolver** succeeds when $n \geq 8$. This completes the proof when $m = 3$.

So suppose that $m \geq 4$ and hence $n \geq 5$. We shall show that the stepwise curvature of steps of length 2 is at most $-1/4$, and hence that **RSym** succeeds with $\varepsilon = 1/4$. A step of length 2 in a green face f of Γ consists of two consolidated edges e_1 and e_2 of length 1, labelled x and y , and the adjacent faces are green and red, respectively. Let v be the vertex between e_1 and e_2 , and B the blob incident with f at e_2 . Then $\chi(B, f, \Gamma) \leq -1/6$, by Lemma 7.10, so by Lemma 7.5 if $\delta_G(v, \Gamma) \geq 3$ then the step curvature is at most $-1/3$.

Otherwise, $\delta(v, \Gamma) = 3$ and $\delta_G(v, \Gamma) = 2$. Then B has two successive edges both labelled $y^{-1} = y_{m-1}$. If $\text{Area}(B) > 1$, then $\chi(B, f, \Gamma) \leq -1/4$. Otherwise B is a triangle. But then the third edge of $\partial(B)$ is labelled y_2 , which is not equal to y or y^{-1} since $m \geq 4$. Hence B is a boundary red blob, and so again $\chi(B, f, \Gamma) \leq -1/4$.

Finally, we consider the word problem for $m \geq 4$. Similarly to the previous cases, the longest possible sequence of edges between a boundary face with positive curvature and the interior is length 5, with label $y(xy)^2$, and so `VerifySolverTrivInt` succeeds. \square

The next result is a more complicated application and is, as far we know, new. (Note that there are general results due to Gromov and others that adding a suitably high power of a non-torsion element to a presentation of a hyperbolic group yields another hyperbolic group.)

For this proof, we shall modify `ComputeRSym`, by adding an extra step, which is currently not implemented and is used only for hand calculations. After Step 4 of `ComputeRSym`, we insert the following.

Step 4+ (Optional) Each green face with curvature less than $-\varepsilon$, for some user-determined ε , gives some of its curvature to any adjacent non-boundary green faces f for which f has curvature greater than $-\varepsilon$.

Definition 9.6. We shall refer to the curvature distribution calculated by this modification of `RSym` as `RSym+`.

Theorem 9.7. Let $G = \langle x, y \mid x^2, y^3, (xy)^m, (xyxy^{-1})^n \rangle$. Then there exists a pregroup presentation \mathcal{P} of G such that the following hold.

- (i) If $m \geq 13$ and $n \geq 7$, then `RSym` and `VerifySolverTrivInt` succeed on \mathcal{P} .
- (ii) If $m \geq 7$ and $n \geq 19$, or if $m \geq 25$ and $n \geq 4$, then `RSym+` succeeds on \mathcal{P} at level 2.

Furthermore, G is hyperbolic in all of these situations.

Proof. We let $P = \{1, x, y, Y\}$, with products $x = x^\sigma$, $Y = y^\sigma$, $y^2 = Y$ and $Y^2 = y$. So the only red blobs in diagrams Γ in \mathcal{D} are single red triangles with boundary yyy or YYY , and we can take $\mathcal{R}^\pm = \{R_1 := (xy)^m, R_2 := (xY)^m, R_3 := (xyxY)^n\}$.

We start by listing the possible labels of consolidated edges between two green faces in any diagram $\Gamma \in \mathcal{D}$, but we omit those in which one of the two faces is labelled R_2 , since these correspond in an obvious way to those with a face labelled R_1 . In each case, we have specified the words labelling the two faces, with the labels of the consolidated edge positioned at the beginning of each of the two words.

1. x between (faces labelled) $(xy)^m$ and $(xy)^m$;
2. x between $(xy)^m$ and $(xyxY)^n$;
3. x between $(xy)^m$ and $(xYxy)^n$;
4. xy between $(xy)^m$ and $(Yxyx)^n$;
5. xyx between $(xy)^m$ and $(xYxy)^n$;
6. yx between $(yx)^m$ and $(xYxy)^n$;
7. y between $(yx)^m$ and $(Yxyx)^n$;
8. x between $(xyxY)^n$ and $(xYxy)^n$.

Since consolidated edges have length at most 3, each step has length at most 4 and so, if $m \geq 13$ and $n \geq 7$, then there are least 7 steps in any decomposition of a relator and hence `RSym` succeeds with $\varepsilon = 1/6$, by Lemma 7.17. The longest possible sequence of edges between a boundary face with positive curvature and the interior of Γ is comprised of three steps together with a red triangle at the beginning; and hence has length at most 13. Any such sequence of edges has a red triangle at each end of it, so if $m \geq 13$ and $n \geq 7$ then `VerifySolverTrivInt` succeeds. This proves Part (i) of the theorem.

For Part (ii), assume that $m \geq 7$. Let f_1 be a non-boundary face with boundary label $R_1 = (xy)^m$, in a diagram $\Gamma \in \mathcal{D}$. We consider the possible decompositions of R_1 into

steps. A step of length k on R_1 constitutes a proportion $k/(2m) \leq k/14$ of $|R_1|$, and the step curvature is at most $-1/6$. So, when $k \leq 2$, the step curvature is less than its length requires on average for $\kappa_\Gamma(f_1) \leq 0$. In fact, the step curvature is less than required by a factor of at least $7/6$.

Consider a consolidated edge labelled xy , as in item 4 of the list above. The edges of all red triangles are labelled y or Y , and the letter following xy in R_1 is x , so such an edge must be followed by another green consolidated edge. So this step consists of the consolidated edge xy only and hence has length 2. Hence the step curvature is less than required by a factor of at least $7/6$. This deals with items 1, 2, 3, 4 and 7 of the list above.

Consider next a consolidated edge e labelled xyx , as in item 5 of the list above, and let v be the vertex at the end of e . There are two places that could come at v , namely $\mathbf{P}_1 = (R_1(2, x, y), Y, R)$ and $\mathbf{P}_2 = (R_1(2, x, y), x, G)$. For \mathbf{P}_2 , the step consists of e , and it is easily checked that $\delta_G(v, \Gamma) \geq 4$ and hence $\chi(v, f_1, \Gamma) \leq -1/4$. Since $1/4 > 3/14$, such a step gives less than its proportionate contribution to f_1 , by a factor of at least $7/6$. The same is true for \mathbf{P}_1 , except when $\delta(v, \Gamma) = 3$ and $\delta_G(v, \Gamma) = 2$.

Similar considerations apply to consolidated edges labelled yx , as in item 6 of the list above, so there are just two types of steps that give more than their proportionate contribution to f_1 , namely those consisting of a consolidated edge labelled xyx or yx together with a red edge, with the property that the vertex in the middle of the step has total degree 3. These steps have lengths 4 and 3, respectively, and have curvature $-1/6$.

Let us call these consolidated edges labelled xyx or yx in these steps *bad consolidated edges*. Then the other face f_2 incident with a bad consolidated edge is labelled $R_3 = (xyxY)^n$ and, since a bad consolidated edge on R_3 is immediately preceded by a red edge labelled y , there can be at most n bad consolidated edges on $\partial(f_2)$. (Note that the bad consolidated edges of f_2 could also be incident with faces labelled $R_2 = (xY)^m$, but the same restrictions apply.)

Now suppose that $n \geq 19$. Then, since the steps have length at most 4 and the step curvature is at most $-1/6$, a non-boundary face f_2 labelled R_3 satisfies $\kappa_\Gamma(f_2) \leq 1 - n/6$. For such faces, we can now apply Step 4+ of the algorithm to compute \mathbf{RSym}^+ , as follows. Fix some small $\varepsilon > 0$. Then f_2 donates curvature $-1/6 + (1 + \varepsilon)/n$ across each of its bad consolidated edges. Since there at most n of these, it still has curvature at most $-\varepsilon$ after making these donations.

A face f_1 labelled R_1 (corresponding considerations apply to faces labelled R_2) that is at dual distance at least 3 from $\partial(\Gamma)$ receives at most $-1/6 + (1 + \varepsilon)/19$ curvature across each bad consolidated edge in Step 4+ of the algorithm to compute \mathbf{RSym}^+ . If f_1 has d bad consolidated edges, then $d \leq m/2$, so the curvature of f_1 before and after Step 4+ is at most

$$1 - \left(\frac{2m - 4d}{2m} \cdot \frac{7}{6} \right) - \frac{d}{6} \quad \text{and} \quad 1 - \left(\frac{2m - 4d}{2m} \cdot \frac{7}{6} \right) - \frac{d}{3} + \frac{d(1 + \varepsilon)}{19}.$$

It can be checked that this is negative for all $m \geq 7$ and $d \leq m/2$, with ε close to 0.

The proof in the case $n \geq 4$ and $m \geq 25$ is analogous, and is omitted.

To deduce hyperbolicity of G , we apply Theorem 6.13(i) when \mathbf{RSym} succeeds at level 1. To apply Theorem 6.13(iii) in the cases when \mathbf{RSym} succeeds only at level 2, we need to know that neither x nor y is trivial in G . We could verify that in each of the individual cases by describing a finite homomorphic image of G in which the images of x and y are nontrivial, but it is quicker just to observe that if either x or y were trivial in G then G would be finite of order at most 3 and hence hyperbolic. \square

It seems possible that by working a little harder we could slightly improve the above result to include more pairs (m, n) , but we have not done this, because we can use the KBMAG package to test hyperbolicity in individual cases. By doing that, and combining it with the result of Theorem 9.7, we obtain

Theorem 9.8. *Let $G = \langle x, y \mid x^2, y^3, (xy)^m, (xyxy^{-1})^n \rangle$. Then G is infinite hyperbolic whenever any of the following conditions hold.*

- $m = 7$ and $n \geq 13$;
- $m = 8$ and $n \geq 8$;
- $m = 9$ and $n \geq 7$;
- $m = 10$ and $n \geq 6$;
- $m \geq 11$ and $n \geq 5$;
- $m \geq 15$ and $n \geq 4$.

In the cases $(m, n) = (7, 12), (8, 7), (9, 6), (10, 5)$ and $(14, 4)$, the group G is automatic and infinite. In each of these examples, by using straightforward searches through the elements of G of bounded length, we were able to find a pair g, h of commuting elements that project onto a free abelian group of rank 2 in an abelian quotient of a suitably chosen subgroup of finite index in G . So these groups all contain free abelian subgroups of rank 2, and hence they are not hyperbolic.

These groups are finite for some smaller values of m and n . The final case to be settled was $(m, n) = (13, 4)$, which was proved finite by coset enumeration. In the cases $(m, n) = (7, 10), (7, 11), (8, 6)$ and $(12, 4)$, G has been proved to be infinite, and we conjecture that it is not automatic, and hence also not hyperbolic, but we are unable to prove this. See [8] for details and further references on the finiteness question.

9.3 Further examples

In an interesting combination of **RSym** with some theoretical arguments involving curvature, Chalk proves in [3] that the Fibonacci groups $F(2, n)$ are hyperbolic for odd $n \geq 11$. The hyperbolicity of $F(2, n)$ for even $n \geq 8$ was proved earlier in [10], and that of $F(2, 9)$ has been proved computationally, using KBMAG. Furthermore, $F(2, n)$ is finite for $n = 1, 2, 3, 4, 5, 7$ and virtually free abelian of rank 3 (and so not hyperbolic) for $n = 6$. So this completes the proof that $F(2, n)$ is hyperbolic if and only if $n \neq 6$. Additionally, in as yet unpublished work, Chalk has used some concepts from **RSym** to prove that the Heineken group $\langle x, y, z \mid [x, [x, y]] = z, [y, [y, z]] = x, [z, [z, x]] = y \rangle$ is hyperbolic.

Whilst in this paper we have only proved that **RSym** naturally generalises the classical small cancellation conditions $C(p)$, $C'(1/p)$ and $T(q)$ over free groups, and free products with amalgamation, we are confident that **RSym** naturally generalises a wide variety of other small cancellation conditions. For example, in [16] there is a form of small cancellation for groups constructed as HNN extensions, which we have not analysed only because the construction of a pregroup describing an HNN extension is a little technical. Metric small cancellation has been defined over graphs of groups, and used to prove hyperbolicity of large families of groups [18]. Other conditions for small cancellation over free groups have been introduced by many authors: for example, Condition $V(6)$ in [23]. We think it is likely that **RSym** generalises most, or even all, of these, although it would be some work to check all of the details.

More speculatively, we believe that it is possible that more powerful curvature distribution schemes than **RSym** could be used to tackle a wide range of problems regarding the hyperbolicity of finitely-presented groups. **RSym**, even with the modification we have called **RSym**⁺, is rarely useful for 1-relator groups, but curvature distribution schemes that allowed the same relator to be treated differently in different contexts might well be useful. Similarly, curvature distribution schemes that permitted the curvature to be moved (bounded distances) across diagrams could be useful for the Restricted Burnside Problem. We chose to present **RSym** in this paper because it can be tested in low-degree polynomial time, but if one is willing to accept a higher degree polynomial cost, or perhaps a cost

with exponent the length of the longest relator, then schemes could be devised which would prove the hyperbolicity of much wider classes of finite presentations.

10 Implementation

We have implemented `RSymVerify`, for the case where $\mathcal{I}(R) = R$ for all $R \in \mathcal{R}$, in the computer algebra systems GAP and MAGMA, as `IsHyperbolic`. It is in the released version of MAGMA, and in the deposited GAP package `Walrus`. The two implementations are moderately different in their details, so we have used each of them as a test of correctness of the other. We have provided methods to produce a pregroup whose universal group is a given free product of free and finite groups, as in Examples 2.4 and 2.5. The user is then able to add any additional relators. We have also implemented `VerifySolverTrivInt` and `RSymSolveTrivInt` in MAGMA.

In this section we describe some run times, using the MAGMA version. The experiments were run on a MacBook Pro laptop with a 3.1GHz processor, and all set $\varepsilon = 1/10$. We have not compared timings with the KBMAG package, as with the exception of the very smallest presentations we found that KBMAG did not appear to terminate.

We first ran `IsHyperbolic` on presentations of the form $\langle x, y \mid x^2, y^m, (xy)^n \rangle$, constructed as a quotient of the free product $C_2 * C_m$, for $3 \leq m \leq 6$ and $n \in \{5, 10, 15\}$. As expected, it succeeded for all $(m, n) \neq (3, 5)$. The time taken was not noticeably dependent on m or n and was less than 0.01 seconds for each trial.

We then tested presentations of the form $\langle x, y \mid x^2, y^3, (xy)^m, [x, y]^n \rangle$, again constructed as a quotient of $C_2 * C_3$, for $10 \leq m \leq 20$ and $6 \leq n \leq 15$. `IsHyperbolic` failed for $m \leq 12$ or $n = 6$, and otherwise succeeded on all trials. Again, the time taken was not noticeably dependent on m or n and was less than 0.01 seconds for each trial.

We have also run experiments with randomly chosen relators, and the results appear in Table 3. For each, we take the average time for 20 sets of random relators with the given parameters. After each run time we give the number of times `IsHyperbolic` successfully proved that the group was hyperbolic, with $\varepsilon = 1/10$.

For random quotients of free groups we choose random, freely cyclically reduced words of the given length as additional relators. For random quotients of free products of two groups we choose random nontrivial group elements alternating between the two factors. For random quotients of three finite groups, we choose a factor at random (other than the previous factor) and then a random nontrivial element from that factor. For free products with a nontrivial free factor we allow the free factor to be chosen twice in a row, but not then to choose the inverse of the previously-chosen letter.

Let F be a free group of rank n , and consider the quotient of F by r random, freely cyclically reduced relators of length 3. There are $2m(4m^2 - 6m + 3) \sim (2m)^3$ such words of length 3 over $\{a_1^{\pm 1}, \dots, a_m^{\pm 1}\}$, so define the *density* $d \in (0, 1)$ of the presentation by $r = (2m)^{3d}$. Żuk showed in [25] that if $d < 1/2$ then the probability that \mathcal{P} defines an infinite hyperbolic group tends to 1 as $m \rightarrow \infty$, whilst if $d > 1/2$ then the probability that \mathcal{P} defines the trivial group tends to 1 as $m \rightarrow \infty$. These asymptotic results tell us what to expect when we choose r random cyclically reduced relators of length 3 in the cases when r/n is either very small or very large, and it seemed interesting to study the case when $n \rightarrow \infty$ with r/n constant. We used our MAGMA implementation of `IsHyperbolic` to investigate this situation experimentally, and also attempted to analyse it theoretically.

Provided that we enforce our condition that there are no pieces of length 2 in the presentation, the most common cause of failure of `RSym` for moderate values of r/n is the possible existence of an internal vertex of degree 3 in a van Kampen diagram. A simple calculation, of which we omit the details, shows that the expected number of triples $\{a, b, c\}$ of distinct elements of X which could label the edges incident with such a vertex tends to $\lambda := 9(r/n)^3/2$ as $n \rightarrow \infty$. Assuming that the number of such vertices forms a

Table 3: Run times averaged over 20 randomly-chosen examples

A free group of rank 2 with m random relators of length n			
$m = 2$	$n = 20$	30	40
		0.02 (0)	0.04 (20) 0.07 (20)
$m = 3$	$n = 25$	35	45
		0.05 (0)	0.10 (20) 0.17 (20)
$m = 10$	$n = 40$	50	60
		1.42 (11)	2.21 (20) 3.26 (20)
$m = 40$	$n = 52$	62	72
		47.83 (12)	70.68 (20) 103.00 (20)
A free group of rank 10 with m random relators of length n			
$m = 10$	$n = 8$	20	30
		0.13 (8)	1.02 (20) 2.33 (20)
$m = 20$	$n = 10$	20	30
		1.02 (3)	3.77 (20) 6.97 (20)
$m = 30$	$n = 13$	20	30
		4.00 (19)	7.01 (20) 12.31 (20)
$m = 50$	$n = 15$	25	35
		8.82 (18)	20.02 (20) 35.90 (20)
A free group of rank 100 with m random relators of length n			
$m = 30$	$n = 4$	10	20
		0.09 (14)	0.91 (20) 7.11 (20)
$m = 50$	$n = 4$	10	20
		0.33 (6)	4.23 (20) 41.78 (20)
$m = 70$	$n = 5$	10	50
		1.49 (18)	13.51 (20) 132.21 (20)
$C_2 * C_3$ with m random relators of length n			
$m = 1$	$n = 96$	120	160
		0.39 (1)	0.59 (19) 1.02 (20)
$m = 2$	$n = 120$	160	200
		2.33 (3)	4.32 (19) 6.74 (20)
$m = 5$	$n = 200$	240	280
		40.60 (20)	62.55 (20) 83.64 (20)
$C_3 * C_3 * C_3$ with m random relators of length n			
$m = 1$	$n = 12$	24	36
		0.01 (8)	0.03 (19) 0.06 (20)
$m = 2$	$n = 20$	30	40
		0.06 (5)	0.12 (20) 0.19 (20)
$m = 5$	$n = 25$	55	75
		0.41 (1)	1.82 (20) 3.43 (20)
$C_3 * A_5 * F_3$ with m random relators of length n			
$m = 2$	$n = 5$	10	20
		1.74 (4)	7.60 (19) 38.36 (20)
$m = 3$	$n = 12$	20	30
		26.58 (19)	98.67 (20) 302.00 (20)
$m = 5$	$n = 15$	25	35
		184.92 (19)	638.24 (20) 1575.63 (20)

Poisson distribution, this would imply that the probability of there being no such triples would tend to $\exp(-\lambda)$. This estimate agrees surprisingly well with our experiments with **RSym**. When $r/n = 1/2$, for example, we have $\exp(-\lambda) \simeq 0.570$ and, the proportion of successes over 1000 runs of our implementation with $n = 100, 500$ and 1000 , were $0.510, 0.577$, and 0.569 .

If $d > 1/3$, then the probability that two relators share a subword of length 2, and hence that our “preprocessing step” simplifies the presentation, tends to 1, and renders the presentation non-random. It is therefore unclear to us how to complete the analysis.

11 Appendix: Glossary and list of notation

List of mathematical terms

Term	See	Term	See
area	3.3	location	7.3
blob curvature	6.3	loop-minimal	3.13
boundary face, edge, vertex	3.1	one-step reachable	7.15
coloured area	3.4	P -reduced, semi- P -reduced	3.7
coloured (van Kampen) diagram	3.2	place, potential place,	7.4
consolidated edge	3.1	plane graph	3.10
curvature distribution	5.1	post-interleave set	7.25
curvature distribution scheme	5.2	pregroup	2.1
cyclic interleave	4.2	pregroup presentation	2.11
cyclic interleave class	4.5	pregroup Dehn function	5.5
(cyclically) σ -reduced	2.6	pre-interleave set	7.25
(cyclically) P -reduced	2.6	\mathcal{R} -letter	7.11
decorated location	7.25	red blob	4.11
decorated place	7.26	red degree	3.2
decorated vertex graph	7.27	red place	7.4
dual distance	5.4	σ -reduced, semi- σ -reduced	3.6
external face	3.1	simply connected red blob	4.11
external word	3.1	single rewrite	2.7
\mathcal{G} -vertex	7.6	standard group presentation	5.6
green degree	3.2	step	7.2
green place	7.4	stepwise curvature	7.2
green-rich	4.9	subdiagram	3.5
half-edge	6.2	succeeds with constant ε	6.6
intermult, intermult pair	4.14	succeeds at level d	6.6
internal face	3.1	terminal place	8.4
interleave	2.7	universal group	2.2
interleave set	7.23	V^σ -letter	3.13
intermediate place	7.15	verifies a solver	8.2
length of step	7.2	vertex graph	7.6

List of notation

Symbol	See	Symbol	See
$\text{Area}(\Gamma)$	3.3	$\text{RSym}+$	9.6
$\beta(B)$	6.3	$U(P)$	2.2
$\chi(v, f, \Gamma), \chi(B, f, \Gamma)$	6.4	\mathcal{V}	7.27
$\text{CArea}(\Gamma)$	3.4	V_P	2.2
∂	3.1	X^σ	2.2
\mathcal{D}	6.1	\approx	2.7
$D(P)$	2.1	\approx^c	4.2
$D(n)$	5.5	$[ab]$	2.1
$\delta_G(v, \Gamma), \delta_R(v, \Gamma)$	3.2		
$\delta_G(e, \Gamma)$	3.10		
ϵ_i	7.1		
$F(X^\sigma)$	2.2		
\mathcal{G}	7.6		
$\mathcal{I}(\mathcal{P}), \mathcal{I}(\mathcal{R})$	4.7		
$\mathcal{I}(w)$	4.5		
$\mathcal{I}(a, b)$	7.23		
κ_Γ	6.3		
$\text{OneStep}(\mathbf{P})$	7.16		
$\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$	2.11		
\mathcal{P}_G	5.6		
$\text{PD}(n)$	5.5		
$\text{Post}(R(i))$	7.25		
$\text{Pre}(R(i))$	7.25		
$R(i, a, b)$	7.3		
RSym	6.4		

Procedures

Name	See
Blob	7.12
FindEdges	7.28
ComputeOneStep	7.16
ComputeRSym	6.3
RSymIntVerify	7.30
RSymSolve	8.8
RSymSolveTrivInt	8.11
RSymVerify	7.19
RSymVerifyAtPlace	7.18
VerifySolver	8.3
VerifySolverTrivInt	8.11
VerifySolverAtPlace	8.5
Vertex	7.7

Acknowledgements This work was supported by EPSRC grant number EP/I03582X/1. We would like to thank P.E. Holmes for many useful conversations about the fundamental ideas underpinning this project. We would also like to thank Simon Jurina for a very careful reading of a final draft of this article.

Addresses Derek Holt, Mathematics Institute, University of Warwick, Coventry, CB4 7AL, UK. D.F.Holt@warwick.ac.uk
Steve Linton and Markus Pfeiffer, School of Computer Science, University of St Andrews, Fife KY16 9SX, UK. steve.linton@st-andrews.ac.uk, markus.pfeiffer@st-andrews.ac.uk
Max Neunhöffer, Im Bendchen 35a, 50169 Kerpen, Germany. max09hoeffer.de
Richard Parker, 70 York St, Cambridge CB1 2PY, UK. richpark7920@gmail.com
Colva M. Roney-Dougal, Mathematical Institute, University of St Andrews, Fife KY16 9SS, UK. colva.roney-dougal@st-andrews.ac.uk

References

- [1] J. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro and H. Short, *Notes on word-hyperbolic groups*, in E. Ghys, A. Haefliger and A. Verjovsky, eds., *Proceedings of the Conference Group Theory from a Geometric Viewpoint* held in I.C.T.P., Trieste, March 1990, World Scientific, Singapore, 1991.
- [2] W. Bosma, J. Cannon and C. Playoust. *The Magma algebra system. I. The user language*, J. Symbolic Comput. **24** (1997), 235–265.
- [3] C. Chalk. *Fibonacci groups, $F(2, n)$, are hyperbolic for n odd and $n \geq 11$* . arXiv:2005.10653.

- [4] B. Domanski and M. Anshell. *The complexity of Dehn's algorithm for word problems in groups*. *J. Algebra* **6** (1985), 543–549.
- [5] D. B. A. Epstein and D. F. Holt. Computation in word-hyperbolic groups. *Internat. J. Algebra Comput.* **11** (2001), 467–487.
- [6] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.10.1*; 2019, <https://www.gap-system.org>.
- [7] S. M. Gersten and H. Short. Small cancellation theory and automatic groups. *Invent. Math.* **102**(2) (1990), 305–334.
- [8] G. Havas and D. F. Holt. On Coxeter's families of group presentations. *J. Algebra* **324**(5) (2010), 1076–1082.
- [9] A. H. M. Hoare. Pregroups and length functions. *Math. Proc. Cambridge Philos. Soc.* **104**(1) (1988), 21–30.
- [10] H. Helling, A. C. Kim, and J. L. Mennicke. A geometric study of Fibonacci groups. *Journal of Lie Theory* **8**(1) 1998, 1–23.
- [11] Derek F. Holt, Bettina Eick, and Eamonn A. O'Brien. *Handbook of computational group theory*. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, 2005.
- [12] D. F. Holt, The Warwick automatic groups software. *Geometrical and Computational Perspectives on Infinite Groups*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 25, ed. Gilbert Baumslag et. al., 1996, 69–82.
- [13] D. F. Holt, “KB MAG - Knuth–Bendix in monoids and automatic groups”, software package (1995), available from <http://homepages.warwick.ac.uk/~mareg/download/kbmag2/>
- [14] D. F. Holt, S. Rees and C. E. Röver, *Groups, languages and automata*, London Mathematical Society Student Texts **88**, Cambridge University Press, 2017.
- [15] D. B. Johnson. Efficient algorithms for shortest paths in sparse networks. *J. ACM* **24**(1) (1977), 1–13.
- [16] R. C. Lyndon & P. E. Schupp. *Combinatorial group theory*. Springer-Verlag, 1977.
- [17] I. G. Lysenok. An approach to the study of finitely presented groups based on the notion of discrete curvature. *Mathematical Notes* **103**(4) (2018), 568–575.
- [18] A. Martin. Complexes of groups and geometric small cancellation over graphs of groups. *Bull. Soc. Math. France* **145** (2017), 193–223.
- [19] P. Papasoglu. Strongly geodesically automatic groups are hyperbolic. *Invent. Math.* **121**(2) (1995), 323–334.
- [20] F. Rimlinger. *Pregroups and Bass-Serre theory*. Mem. Amer. Math. Soc. **65**(361) (1987).
- [21] J.-P. Serre. *Trees*. 2nd Edition, Springer, 2003.
- [22] J. R. Stallings. *Group theory and three-dimensional manifolds*. Yale University Press, New Haven, Conn., 1971.
- [23] U. Weiss. On biautomaticity of non-homogenous small-cancellation groups. *Internat. J. Algebra Comput.* **17**(4) (2007), 797–820.
- [24] Robin J. Wilson. *Introduction to graph theory*. Prentice Hall. 5th Edition, 2010.
- [25] A. Żuk. Property (T) and Kazhdan constants for discrete groups. *Geom. Funct. Anal.* **13** (2003), 643–670.