SOLVING DECISION PROBLEMS IN FINITELY PRESENTED GROUPS VIA GENERALIZED SMALL CANCELLATION THEORY

Šimon Jurina

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

2023

Identifiers to use to cite or link to this thesis:

DOI: https://doi.org/10.17630/sta/429
http://hdl.handle.net/10023/27549

Solving decision problems in finitely presented groups via generalized small cancellation theory

Šimon Jurina

University of St Andrews

This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy (PhD)
at the University of St Andrews

August 2022

Abstract

This thesis studies two decision problems for finitely presented groups. Using a standard RAM model of computation, in which the basic arithmetical operations on integers are assumed to take constant time, in Part I we develop an algorithm $\mathbf{IsConjugate}$, which on input a (finite) presentation defining a hyperbolic group G, correctly decides whether $w_1 \in X^*$ and $w_2 \in X^*$ are conjugate in G, and if so, then for each $i \in \{1, 2\}$, returns a cyclically reduced word r_i that is conjugate in G to w_i, and an $x \in X^*$ such that $r_2 =_G x^{-1} r_1 x$ (hence if w_1 and w_2 are already cyclically reduced, then it returns an $x \in X^*$ such that $w_2 =_G x^{-1} w_1 x$). Moreover, $\mathbf{IsConjugate}$ can be constructed in polynomial-time in the input presentation $\langle X \mid \mathcal{R} \rangle$, and $\mathbf{IsConjugate}$ runs in time $O((|w_1| + |w_2|) \cdot \min\{|w_1|, |w_2|\})$.

$\mathbf{IsConjugate}$ has been implemented in the `MAGMA` software, and our experiments show that the run times agree with the worst-case time complexities. Thus, $\mathbf{IsConjugate}$ is the most efficient general practically implementable conjugacy problem solver for hyperbolic groups.

It is undecidable in general whether a given finitely presented group is hyperbolic. In Part II of this thesis, we present a polynomial-time procedure $\mathbf{VerifyHypVertex}$ which on input a finite presentation for a group G, returns `true` only if G is hyperbolic. $\mathbf{VerifyHypVertex}$ generalizes the methods from [34], and in particular succeeds on all presentations on which the implementation from [34] succeeds, and many additional presentations as well. The algorithms have been implemented in `MAGMA`, and the experiments show that they return a positive answer on many examples on which other comparable publicly available methods fail, such as KBMAG.

Acknowledgements

General acknowledgements

I would like to express my deepest gratitude to the following people:

To my supervisors, Professor Peter Cameron and Professor Colva Roney-Dougal, for your guidance throughout the PhD process, for incredibly helpful mathematical discussions, and for teaching me how to write mathematics. Without your help I doubt I would have come this far.

To my examiners, Professor James David Mitchell and Professor Sarah Rees, for reading my thesis thoroughly, and for organizing and conducting the PhD viva.

To my mother Jarmila, my father Lubomír, and my sister Alena, whose unconditional love and relentless support have helped me to overcome the worst moments, and inspire me to become a better person every day.

And to my friends Horacio Guerra, Carlos Pérez and Yunchen Xiao, for having lots of laughs together, for sharing your experience doing a PhD, and for always encouraging me to work hard on my thesis.

Funding

This work was supported by the Engineering and Physical Sciences Research Council (grant number EP/R513337/1) and the University of St Andrews (School of Mathematics and Statistics).

Research data access statement

The supplementary code for this thesis is available at https://github.com/simonjurina/PhDThesisCode and at https://doi.org/10.17630/0f7a7c71-26f5-4346-9bf1-3da33e354721.

Candidate's declaration

I, Šimon Jurina, hereby certify that this thesis, which is approximately 70 000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a candidate for the degree of Ph.D in September 2018.

Date: 09/03/2023 Signature of candidate:

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Ph.D in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date: 13/3/2023 Signature of supervisor:

Permission for publication

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis.

I, Šimon Jurina, have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

PRINTED COPY

No embargo on print copy

ELECTRONIC COPY

No embargo on electronic copy

Date: 13/3/2023

Signature of candidate: Signature of supervisor:

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Šimon Jurina, understand that by declaring that I have original research data or digital outputs, I should make every effort in meeting the University's and research funders' requirements on the deposit and sharing of research data or research digital outputs.

Date: 09/03/2023 Signature of candidate:

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

No embargo on underpinning research data or digital outputs.

Date: 09/03/2023 Signature of candidate:

Date: 13/3/2023 Signature of supervisor:

Contents

Chapter 1

Introduction

In this thesis we study two problems for finitely presented groups. In the description of these problems, an *algorithm* is a computational method consisting of finitely many steps that always returns a definite answer, while a *procedure* is a computational method that might consist of infinitely many steps, and is not guaranteed to return a definite answer. Furthermore, a finitely presented group is hyperbolic if and only if its Dehn function is linearly bounded, and if X is a finite set, then X^* denotes the set of all words over X.

1. Given a finite presentation $\langle X \mid \mathcal{R} \rangle$ defining a hyperbolic group G, develop an algorithm that decides in finite time, whether or not two words $w_1 \in X^*$ and $w_2 \in X^*$ are conjugate in G.

2. Given a finite presentation $\langle X \mid \mathcal{R} \rangle$ of a group G, produce a procedure that returns in finite time a positive answer only if G is hyperbolic, else it returns `fail` – meaning that G might be hyperbolic, but the procedure is unable to prove it.

In both Problems 1-2, we further want our algorithms to be at worst polynomial-time (in particular in Problem 1 quadratic time), and to be practically implementable. We emphasize that we assume the RAM model of computation, in which the basic arithmetical operations on integers can be performed in constant time. This is a common assumption when analysing complexity of algorithms. For example, in [2] it is used as default. The assumption is reasonable as long as we are not working with integers outside of range that can be represented with a given number of digits, so on a typical machine, integers with absolute value larger than 2^{31}. For example in Problem 1, it is a reasonable assumption for input words not longer than about 2^{31}.

There are at least two types of *decision problems* for groups. Type one: given a (finite) presentation of a group G, does G have some property P? Type two: given a (finite) group presentation $\mathcal{P}$ and a finite list of words w_i in the generators of $\mathcal{P}$, does the list have a property P? We say that a decision problem is *decidable* if and only if there exists an algorithm to solve it.

The following three decision problem proposed by Max Dehn in 1911 (see [15]) turned out to be fundamental in group theory.

Dehn 1 Given a group G with a finite presentation $\langle X \mid \mathcal{R} \rangle$ and a word w, is w equal to the identity in G (the *Word Problem*)?

Dehn 2 Given a group G with a finite presentation $\langle X \mid \mathcal{R} \rangle$, and two word w_1 and w_2, are w_1 and w_2 conjugate in G (the *Conjugacy Problem*)?

Dehn 3 Given two finite presentations $G_1 = \langle X_1 \mid \mathcal{R}_1 \rangle$ and $G_2 = \langle X_2 \mid \mathcal{R}_2 \rangle$, are G_1 and G_2 isomorphic (the *Isomorphism Problem*)?

In 1912 (see [16]) he developed an algorithm that solves both the word and the conjugacy problem for fundamental groups of closed orientable two-dimensional manifolds of genus greater or equal to two. In 1950s and 1960s several authors extended Dehn's algorithm to solve the word and the conjugacy problem for group presentations satisfying various small cancellation conditions (see [30, 41, 50]). Both the word problem and the conjugacy problem, are however, undecidable in general. Novikov in 1955 (see [46]) proved that there exists a finitely presented group for which the word problem is undecidable. It follows immediately that the conjugacy problem is also undecidable, since we cannot decide whether a given element is conjugate to the identity. In 1958, Boone independently provided a different proof of the undecibility of the word problem (see [5]).

The undecidability of the word and the conjugacy problem is an instance of a more general and surprising phenomenon that most 'reasonable' properties of finitely presented groups are undecidable.

Definition 1.0.1. A property P is a *Markov property* of finitely presented groups if P satisfies the following 3 conditions.

1. P is preserved by isomorphism.
2. There exists a finitely presented group with property P.
3. There exists a finitely presented group that does not embed into any finitely presented group with property P.

The Adian-Rabin theorem states that Markov properties are undecidable:

Theorem 1.0.2. (The Adian-Rabin Theorem) *Let P be a Markov property of finitely presented groups. Then there does not exist an algorithm that, given a finite presentation $\langle X \mid \mathcal{R} \rangle$, decides whether or not the group defined by $\langle X \mid \mathcal{R} \rangle$ has property P.*

The Adian-Rabin theorem was first proved by Adian in 1955 (see [1]), and independently, by Rabin in 1958 (see [48]). Examples of Markov properties are: being trivial, being finite, being abelian, being finitely presented with solvable word problem, being hyperbolic.

Thus, it is undecidable in general whether a given finitely presented group is hyperbolic. That is why our goal in Problem 2 is to develop a procedure that returns a positive answer on as many group presentations as possible that define hyperbolic groups, but to not develop an algorithm, since that is impossible in general. Also, note that as being a trivial group is a Markov property, the isomorphism problem is also undecidable in general. However, many important decision problems are decidable for hyperbolic groups, even for the more general class of relatively hyperbolic groups. In his seminal article (see [31]) Gromov showed that the word and the conjugacy problem are decidable for hyperbolic groups. In more recent papers (see [13, 14]) Dahmani and Guirardel proved that the isomorphism problem for hyperbolic groups is also decidable. Thus, all Dehn's important decision problems have been positively resolved for hyperbolic groups.

Let G be a hyperbolic group relative to a collection of its subgroups $H_1, \ldots, H_m$. Farb in [25] proved that the word problem is solvable for G provided that it is solvable for each of the subgroups $H_1, \ldots, H_m$. More generally, in the influential monograph on relatively hyperbolic groups (see [47]), Osin showed that the *membership problem* is solvable for each H_i (the membership problem for a subgroup K of a group H is to decide whether a given element $h \in H$ satisfies $h \in K$). At around the same time, Bumagin (see [9]) showed that the conjugacy problem for G is solvable if it is solvable for each H_i.

Compressed decision problems for hyperbolic and relatively hyperbolic groups have also been studied recently. The classical decision problems for groups such as the word and the conjugacy problem are often challenging as huge intermediate words might arise during the computation. These words are sometimes highly compressible, and one can try to compute with these condensed representatives instead of the words themselves. Many authors successfully developed theories of the compression techniques in group theory in recent years (see [18, 21, 35, 36, 40, 44, 45]). In particular, Holt et al. in [35] showed that the *compressed word problem* for hyperbolic groups is solvable in polynomial-time. Finally, in [36] Holt and Rees showed that the compressed word problem in a group that is hyperbolic relative to a collection of free abelian subgroups is also solvable in polynomial time.

One of the main ideas in our solutions to Problems 1-2 is representing groups via finite pregroup presentations (first defined in [34], see Definition 2.3.14). The concept of pregroups was introduced by Stallings in [53]. The work of Rimlinger (see [49]) provides an extension of the theory that enables us to view a group G, given by a finite pregroup presentation, as a quotient of a virtually free group, and not just as a quotient of a free group. We can then ignore any failure of small cancellation on a certain subset of relators of G, leading to a generalisation of small cancellation theory (see for example [42, Chapter 5]).

Part I of this thesis studies Problem 1. The first solution to the conjugacy problem for hyperbolic groups was published by Gromov in the aforementioned article [31], the second by Gersten and Short [27] in the more general context of biautomatic groups. Both of these algorithms run in exponential time in the length of the input words. In [7], Bridson and Haefliger developed a polynomial-time conjugacy problem solver (as it is given its complexity is $O(n^3)$, where n is the length of the input words, but it can be improved to $O(n^2)$). Finally, Epstein and Holt describe a linear-time solution in [24] (the second author together with Buckley provided a linear time solution for finite lists of group elements in [8]).

However, even though these algorithms are great accomplishments, only the solution of

Gersten and Short has been implemented (by Wakefield, see [55, Chapters 5 & 6]). In [43], Marshall developed and implemented a more efficient algorithm, but it provides a solution only for elements with infinite order and despite the fact it runs fast on typical examples, it is not theoretically a linear-time algorithm. Moreover, both of these algorithms as well as the algorithm of Epstein and Holt assume that we can precompute an automatic structure of the input group, but the time complexity of the currently best known algorithm KBMAG (see [33]) for finding these automatic structures is not bounded in the size of the input.

We give a new method for solving the problem, which gives a quadratic time solution, and as far as we know it is currently the most efficient general practically implementable conjugacy problem solver for hyperbolic groups. (See Definition 2.6.14 of a valid pregroup presentation – we shall see that groups defined by valid pregroup presentations are hyperbolic, and Definition 3.1.15 of a cyclically $\mathcal{P}$-reduced word).

Theorem 1.0.3. *Let G be a group defined by a valid pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \cup \mathcal{R}\rangle$, and let $r := \max\{|R| : R \in \mathcal{R}\}$ be the length of the longest relator in $\mathcal{R}$. Then it is possible to construct in time $O(r^4|\mathcal{R}^2||X|^9 + r^2|\mathcal{R}||X|^{11})$ an algorithm* $\mathbf{IsConjugate}$*, which correctly decides whether $w_1 \in X^*$ and $w_2 \in X^*$ are conjugate in G and, if so, then for each $i \in \{1, 2\}$, returns a cyclically $\mathcal{P}$-reduced word r_i that is conjugate in G to w_i, and an $x \in X^*$ such that $r_2 =_G x^{-1}r_1x$ (hence if w_1 and w_2 are already cyclically $\mathcal{P}$-reduced, then it returns an $x \in X^*$ such that $w_2 =_G x^{-1}w_1x$). Moreover,* $\mathbf{IsConjugate}$ *runs in time $O((|w_1| + |w_2|) \cdot \min\{|w_1|, |w_2|\})$.*

We shall prove Theorem 1.0.3 in Section 6.4 by using new theory developed in Chapters 3-6. Furthermore, we implemented $\mathbf{IsConjugate}$ in the MAGMA software (see [6]), and the reported run times (see Chapter 7) agree with Theorem 1.0.3.

In Part II of this thesis, we describe a new polynomial-time procedure $\mathbf{VerifyHypVertex}$ that seeks to find a linear bound on the Dehn function of a finitely presented group. $\mathbf{VerifyHyp}$ $\mathbf{Vertex}$ is the most general polynomial-time procedure for proving hyperbolicity of finitely presented groups, and our experiments show that it returns a positive answer on all presentations on which the $\mathbf{RSymVerify}$ procedure developed by Holt et al. in [34] succeeds. Additionally we show that there are many presentations on which it succeeds but either $\mathbf{RSym}$ $\mathbf{Verify}$ or (other comparable publicly available method) KBMAG fails. Also, the iterative nature of $\mathbf{VerifyHypVertex}$ allows the user to choose how much work the algorithm will do.

We build on the theory of [34], and work with a new type of van Kampen diagrams (defined over pregroup presentations). The general idea used in our algorithms is to assign curvature to vertices, edges and faces of a van Kampen diagram Γ in such a way that the overall curvature of Γ sums to 1; vertices, edges and the external face of Γ have curvature 0; faces of Γ labelled by a pre-determined subset of the relators have also curvature 0; and faces of Γ labelled by other relators, that are sufficiently far from the boundary of Γ have curvature smaller than $-\varepsilon$ for some $\varepsilon \in \mathbb{R}_{>0}$. If we can achieve this for a suitable set of van Kampen diagrams, then we obtain a linear bound on the Dehn function, thus proving that the input presentation defines a hyperbolic group.

The following theorem is the main result of Part II (see Definition 2.3.20 for the meaning of $\mathcal{I}(\mathcal{R})$).

Theorem 1.0.4. *Let G be a group given by a finite pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ such that $\mathcal{I}(R) = R$ for all $R \in \mathcal{R}$, and let $h \geq 1$.*

If $\mathbf{VerifyHypVertex}$ *returns* `true` *on input $\mathcal{P}$ on iteration i for some $i \leq h$, then G is hyperbolic, and an explicit bound on the Dehn function of G can be calculated. Moreover, the time complexity of* $\mathbf{VerifyHypVertex}$ *is $O(r^9|\mathcal{R}|^9|X|^9)$, where $r := \max\{|R| : R \in \mathcal{R}\}$.*

Theorem 1.0.4 will be proved in Section 11.3 by applying new results presented in Chapters 9-11. $\mathbf{VerifyHypVertex}$ has been implemented in `MAGMA`, and our experiments (see Chapter 12) confirm that there are many presentations on which it succeeds but either $\mathbf{RSymVerify}$ or KBMAG fails.

Chapter 2

Preliminaries and notation

In this chapter we give definitions and notation that we shall use throughout this thesis. We split the chapter into 6 sections: Section 2.1 recalls basic concepts from the theory of metric spaces, Section 2.2 gives two equivalent definitions of hyperbolicity of a finitely generated group, Section 2.3 presents the concept of pregroups and pregroup presentations, Section 2.4 presents a new result that the universal group (see Definition 2.3.4) of a finite pregroup satisfying a certain technical condition is isomorphic to a free product of finitely many factors, with each factor a finite or an infinite cyclic group, Section 2.5 presents definitions of coloured diagrams over pregroup presentations, and Section 2.6 presents the concept of curvature distributions schemes.

2.1 Metric spaces

In this section we present the elementary theory of metric spaces that we shall use in this thesis. We took the definitions mostly from [11].

We start by recalling the definition of a metric space.

Definition 2.1.1. [11, Definition 1.1.1 & Example 1.1.2] A *metric space* is a pair (X, d) where X is a set, and $d : X \times X \to [0, \infty)$ is a function, called a *metric*, that satisfies the following conditions, for all $x, y, z \in X$:

(a) $d(x, y) = d(y, x)$;

(b) $d(x, y) = 0$ if and only if $x = y$;

(c) $d(x, z) \leq d(x, y) + d(y, z)$.

If $Y \subseteq X$, then the pair (Y, d') with $d' : Y \times Y \to [0, \infty)$ such that $d'(x, y) = d(x, y)$ for all $x, y \in Y$ is a metric space, called a *sub-space* of (X, d).

We shall work extensively in this thesis with objects embedded in Euclidean spaces:

Example 2.1.2. The *n-dimensional Euclidean space* $\mathbb{R}^n$ is equipped with the *Euclidean metric* d defined as follows, for all $x = (p_i)_{i=1}^n, y = (q_i)_{i=1}^n \in \mathbb{R}^n$:

$$d(x, y) = \sqrt{\sum_{i=1}^{n} (p_i - q_i)^2}.$$

When $n = 1$, the subset $(0, 1) \subseteq \mathbb{R}$ is called the *open unit interval*.

Definition 2.1.3. [11, Definition 1.1.6] Let (X, d) be a metric space. The *open ball of radius r centred at a point $x \in X$*, denoted by $\mathbf{B}_X(x, r)$, is defined as

$$\mathbf{B}_X(x, r) = \{y \in X \mid d(x, y) < r\}.$$

A subset $Y \subseteq X$ is *open* if for every $y \in Y$ there exists an $r > 0$ such that $\mathbf{B}_X(y, r) \subseteq Y$. A subset $Y \subseteq X$ is *closed* if its complement $X \setminus Y$ is open.

Example 2.1.4. The ball $\mathbf{B}_{\mathbb{R}^2}(0, 1)$ is called the *open unit disc*.

Definition 2.1.5. [11, Definition 1.1.12] Let Y be a subset of a metric space (X, d). The *interior* of Y, denoted by Y°, is the set

$$Y^\circ = \bigcup \{G : G \text{ is open and } G \subseteq Y\}.$$

The *closure* of Y, denoted by $\overline{Y}$, is the set

$$\overline{Y} = \bigcap \{F : F \text{ is closed and } F \supseteq Y\}.$$

Finally, the *boundary* of Y, denoted by $\partial(Y)$, is the set

$$\partial(Y) = \overline{Y} \cap \overline{X \setminus Y}$$

.

Definition 2.1.6. [11, [Definitions 1.3.1 & 1.3.11] Let (X, d_1) and (Y, d_2) be metric spaces. A function $f : X \to Y$ is said to be *continuous at a point* $x \in X$ if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that if $d(x, a) < \delta$, then $d(f(x), f(a)) < \varepsilon$. The function f is *continuous* if it is continuous at every point of X.

A function $f : X \to Y$ is said to be a *homeomorphism* if f is bijective and both f and f^{-1} are continuous. Two metric spaces are *homeomorphic* if there exists a homeomorphism from one to the other.

We shall use the following terminology (based on [42, Chapter 5]) in Part 1 of this thesis when working with coloured diagrams defined in Section 2.5.

Definition 2.1.7. A *vertex* is a point of $\mathbb{R}^2$. An *edge* is a bounded subset of $\mathbb{R}^2$ homeomorphic to the open unit interval. A *region* is a bounded subset of $\mathbb{R}^2$ homeomorphic to the open unit disc.

A *path* is a sequence $p = v_0, e_1, v_1, \ldots, v_{n-1}, e_n, v_n$ of vertices v_i and edges e_i such that for $1 \leq i \leq n$, we have $\partial(e_i) = \{v_{i-1}\} \cup \{v_i\}$. A path can be empty, or consist of a single vertex. The vertices v_0 and v_n are the *endpoints* of p. If $v_0 = v_n$, then we say that p is *closed*, or sometimes, that p is a *cycle* (we emphasize that a cycle may have a repeating vertex that is not its endpoint). The *inverse path of* p, denoted by p^{-1}, is the path $p^{-1} = v_n, e_n, \ldots, e_1, v_0$. We write $|p|$ for the *length* of p, which is the number of edges of p (note that $|p| = 0$ if p contains no edge). To simplify notation, if a path $p = v_0, e_1, v_1, \ldots, v_{n-1}, e_n, v_n$ satisfies $|p| \geq 1$, then we often just write $p = e_1 e_2 \ldots e_n$. Then a sub-path q of p is either a path $e_i \ldots e_j$ for $1 \leq i \leq j \leq n$, or a single vertex of p. A path $v_0, e_1, v_1, \ldots, v_{n-1}, e_n, v_n$ is *simple* if it does not have repeating vertices other than the endpoints, i.e. $v_i = v_j$ and $i < j$ implies that $i = 0$ and $j = n$. If we say that a closed path p is of the form $p_1 p_2 \ldots p_n$, then p is a sequence of simple sub-paths p_i.

Lemma 2.1.8. *Suppose that $p = e_1 \ldots e_n$ is a path. Then $(1, \ldots, n)$ contains a non-empty sub-sequence $(i_1, \ldots, i_m)$ such that $q = e_{i_1} \ldots e_{i_m}$ is a simple path with the same endpoints as p.*

Proof. The proof is by induction on n. If $n = 1$, then p is simple, so take $q = p$. Assume that the lemma holds for $n-1$. If p is simple, then we can take $q = p$, so assume that p is not simple. Then there exists a closed sub-path $r = e_i \ldots e_j \neq p$ of p. Consider a path p' obtained by deleting the whole of r except its endpoints from p, i.e. the path $p' = e_1 \ldots e_{i-1} \ldots e_{j+1} \ldots e_n$. By induction $(1, \ldots, i-1, \ldots, j+1, \ldots, n)$ contains a non-empty sub-sequence $(i_1, \ldots, i_m)$ such that $q = e_{i_1} \ldots e_{i_m}$ is a simple path with the same endpoints as p'. But note that p' has the same endpoints as p, and $(i_1, \ldots, i_m)$ is a sub-sequence of $(1, \ldots, n)$, so we are done. ■

In [11, Definition 1.5.1] connected metric spaces are defined. We extend the definition and include familiar definitions of simply-connected and annular subsets of $\mathbb{R}^2$.

Definition 2.1.9. A metric space (X, d) is *connected* if there are no subsets of X that are both simultaneously open and closed other than X and $\emptyset$. If $Y \subseteq X$, then Y is connected if (Y, d) is connected. If Y is connected and is properly contained in no other connected subset, then we say that Y is a *component*. If Y is not connected, then we say that Y is *disconnected*.

A connected subset $X \subseteq \mathbb{R}^2$ is called

1. *simply-connected* if $\mathbb{R}^2 \setminus X$ is connected;
2. *annular* if $\mathbb{R}^2 \setminus X$ is comprised of two components.

We conclude this section with the definition (taken from [37]) of a geodesic metric space and the concept of a graph metric defined on a Cayley graph of a group (taken from [3, Chapter 0]) that we shall use in Section 2.2 to define a hyperbolic group.

Definition 2.1.10. A metric space (X, d) is called *geodesic* if for all $x, y \in X$ there exists a distance preserving bijection $f : [0, k] \to X$ such that $f(0) = x$ and $f(k) = y$, where $k = d(x, y)$. The image of f is called a *geodesic path* from x to y, written as $[xy]$.

Example 2.1.11. Let G be a group generated by a finite set X, and let $\Gamma = \Gamma(G, X)$ be the Cayley graph of G. Define the *graph metric* d on Γ by assigning length 1 to each edge, and defining the distance between two points to be the length of the shortest path between them. Then as G is finitely generated and Γ is connected, (Γ, d) is a geodesic metric space.

2.2 Hyperbolic groups

The notion of hyperbolic groups was introduced and developed by Mikhail Gromov in 1987 (see [31]). There are a number of good sources for an introduction into theory of hyperbolic groups. Beside [31] classic texts include for example [12] and [28]. For the development of basic properties of hyperbolic groups, we recommend [3]. Another good reference is [37, Chapter 6], which presents some useful results for our work in this thesis.

Let G be a group generated by a finite set X, and let $\Gamma = \Gamma(G, X)$ be the Cayley graph of G. Recall from Example 2.1.11 that Γ equipped with the graph metric d defines a geodesic metric space. So we let a *geodesic triangle* xyz in Γ consist of three points x, y, z with geodesic paths $[xy]$, $[yz]$, $[zx]$. The hyperbolicity of Γ (and hence of G) can be defined in terms of 'slimness' of geodesic triangles.

Definition 2.2.1. We say that G (and Γ) is δ-*hyperbolic* if each geodesic triangle xyz of Γ is δ-*slim*: there exists a $\delta > 0$ such that for any point p on one of the sides of xyz, there is a point q in the union of the other two sides of xyz with $d(p, q) < \delta$.

Crucially, in [31] Gromov showed that hyperbolicity is independent of the choice of generators.

Theorem 2.2.2. (Gromov) *If G is hyperbolic and Γ is a Cayley graph of G, then Γ is δ-hyperbolic for some $\delta > 0$.*

Moreover, it was proved by Rips that every hyperbolic group is finitely presented, see [12, Theorem 2.3, Chapter 5].

There are several equivalent definitions of hyperbolicity of a finitely generated group. The one that we shall use is defined with respect to the linearity of the Dehn function.

Definition 2.2.3. [37, Chapter 3] Let G be a group defined by a finite presentation $\mathcal{Q} = \langle X \mid \mathcal{R} \rangle$, and let $F = F(X)$ be the free group on X. A word w over X satisfies $w =_G 1$ if and only if there exist $R_i \in \mathcal{R}^{\pm 1}$ and $u_i \in F(X)$, $1 \leq i \leq k$, such that

$$w =_F u_1^{-1} R_1 u_1 \cdots u_k^{-1} R_k u_k.$$

We call the expression on the right-hand side a *factorisation* of w of length k over $\mathcal{Q}$ (so a factorisation of w of length k is a product of k conjugates of relators in $\mathcal{R}^{\pm 1}$). The *area of w with respect to $\mathcal{Q}$*, written as $\mathbf{Area}(w, \mathcal{Q})$, is the length of the shortest factorisation of w over

$\mathcal{Q}$. The *Dehn function* of $\mathcal{Q}$ is defined as follows:

$$\begin{aligned} f &: \mathbb{N} \to \mathbb{N} \\ n &\mapsto \max\{\mathbf{Area}(w, \mathcal{Q}) : |w| \leq n\}. \end{aligned}$$

It is a standard result that the linearity of the Dehn function is independent of the choice of the presentation for a given group G. It has also been shown (for example, see [37, Theorem 6.5.3]) that there are no groups with sub-quadratic Dehn functions that are greater than linear, i.e. of the form $n \mapsto An^\alpha + B$ with $A, B > 0$ and $\alpha \in (1, 2)$.

The key result is that linearity of the Dehn function implies hyperbolicity.

Theorem 2.2.4. *[37,* Theorem 6.6.1*] For a finitely generated group G with a Cayley graph $\Gamma = \Gamma(G, X)$ the following are equivalent, and hence each provides a definition of hyperbolicity.*

1. *There exists a $\delta > 0$ such that each geodesic triangle in Γ is δ-slim.*
2. *G has a linear Dehn function.*

2.3 Pregroups and their presentations

A significant downside of methods that are applicable only under small cancellation conditions is that they are often not satisfied when short relators are present, e.g. x^n where n is small and x is a generator. But many of the most important group presentations contain such relators. The theory of pregroups will enable us to replace short relators with other relators of length three (the pregroup relators), which we will then ignore in our methods generalizing small cancellation theory.

In this section we shall present definition of pregroups, and then collect some elementary statements about them and explain how one can present any quotient of a virtually free group by finitely many additional relators with a *finite pregroup presentation*. The section is based entirely on [34], where pregroup presentations were first defined.

Definition 2.3.1. [34, Definition 2.1] A *pregroup* is a set P with a distinguished element $1 \in P$ and equipped with a partial multiplication $(x, y) \to [xy]$, which is defined for $(x, y) \in D(P) \subseteq P \times P$, and an involution $\sigma : P \to P, x \mapsto x^\sigma$, satisfying the following axioms, for all $x, y, z, t \in P$:

$(P1)$ $(x, 1), (1, x) \in D(P)$ 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 $([xy], z), ([yz], t) \in D(P)$.

Lemma 2.3.2. *We have* $1^\sigma = 1$*, and if* $[xy] = 1$*, then* $y = x^\sigma$*.*

Proof. By Axioms P1-P2 we have $1 = [11^\sigma] = 1^\sigma$, so the first statement holds.

Suppose that $[xy] = 1$. Then $(x, y) \in D(P)$, by Axiom P2, $(y, y^\sigma) \in D(P)$ and $[yy^\sigma] = 1$, and by Axiom P1, $(x, 1), (1, y^\sigma) \in D(P)$, so $(x, [yy^\sigma]), ([xy], y^\sigma) \in D(P)$. Therefore, by Axioms P1, P2 and P4 we have

$$x = [x1] = [x[yy^\sigma]] = [[xy]y^\sigma] = [1y^\sigma] = y^\sigma,$$

as required. ■

Example 2.3.3. The following multiplication table defines a pregroup $P = \{1, A, B, C, D, E, F, G, H, I\}$ (a blank space means that the product is undefined).

Table 2.1: Pregroup multiplication table

1	A	B	C	D	E	F	G	H	I
A	G	D	B	C			1		
B	C	1	A	G			D		
C	D	G	1	A			B		
D	B	A	G	1			C		
E					H			1	
F									1
G	1	C	D	B			A		
H					1			E	
I						1			

Definition 2.3.4. [34, Definition 2.1] Let P be a pregroup. We denote by X^σ the set $X = P \setminus \{1\}$ equipped with the involution σ, but will sometimes omit the σ when the meaning is clear. We shall write $F(X^\sigma)$ to denote the group defined by the monoid presentation $\langle X \mid xx^\sigma : x \in X\rangle$.

Let V_P be the set of all length three words 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 $F(X^\sigma)$.

The fact that the presentation of $U(P)$ is over an inverse-closed set of monoid generators allows us to write the elements of $U(P)$ as words over X. Also, we shall often write x^σ to mean the inverse of x in $F(X^\sigma)$ rather than x^{-1}. More generally, for $w = x_1x_2\ldots x_n \in X^*$ (recall that X^* denotes the set of all words over X), we shall write $w^\sigma = x_n^\sigma x_{n-1}^\sigma \ldots x_1^\sigma \in X^*$. Then if $w \in F(X^\sigma)$, then $w^{-1} =_{F(X^\sigma)} w^\sigma$.

Remark 2.3.5. 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 $\mathbb{Z}$.

Also, if $xy[xy]^\sigma \in V_P$, then $xy =_{U(P)} [xy]$, hence $x =_{U(P)} [xy]y^\sigma$, so $([xy], y^\sigma) \in D(P)$, and $[xy]y^\sigma x^\sigma \in V_P$. In other words, if $R \in V_P$, then $R^{-1} \in V_P$.

Example 2.3.6. [34, Example 2.4] We can construct a pregroup P such that $U(P) = F(X^\sigma)$ is a free group of rank n by letting X be a set with $|X| = 2n$, letting $x^\sigma \neq x$ for all $x \in X$, and letting the only products be $xx^\sigma = 1$, $1x = x1 = x$, and $1 \cdot 1 = 1$, for all $x \in X$.

Example 2.3.7. [34, Example 2.5] A construction of a pregroup P for which $U(P)$ is the free product of finite groups G and H is: we let P be equal to the disjoint union of $\{1\}$, $G \setminus \{1\}$ and $H \setminus \{1\}$; define $1^\sigma = 1$ and for $g \in G$, $h \in H$, define g^σ to be the inverse of g in G and h^σ to be the inverse of h in H; and finally let $D(P) = (G \times G) \cup (H \times H)$, and define all products as in G and H.

Definition 2.3.8. [34, Definition 3.13] We say that $x \in X$ is a *V^σ-letter* if and only if $x^\sigma = x$ or x is a letter of a relator in V_P.

The following definition extends the notion of a freely (cyclically) reduced word in a free group.

Definition 2.3.9. [34, Definition 2.6] Let $w = x_1 x_2 \ldots x_n \in X^*$. If w contains no sub-words xx^σ with $x \in X$, then we say that w is *σ-reduced*. We say that w is *cyclically σ-reduced* if w is *σ-reduced* and not of the form $x^\sigma w' x$ for some $x \in X$ and $w' \in X^*$.

More generally, the word w 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 $(x_n, x_1) \notin D(P)$.

There is an equivalence relation $\approx$ defined (by Stallings, see [53]) on the set of P-reduced words:

Definition 2.3.10. [34, Definition 2.7] Let $v = x_1 x_2 \ldots x_n \in X^*$ be a P-reduced word and let $w = y_1 y_2 \ldots y_m$ be any word over X. We write $v \approx w$ if $n = m$ and there exist $1 = s_0, s_1, s_2, \ldots, s_{n-1}, s_n = 1 \in P$ 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 where $s_i \neq 1$ for a single value of i, we call a transformation from v to w a *single rewrite*.

By Axiom P4 (see Definition 2.3.1) it follows that if $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i]) \in D(P)$, then $(s_{i-1}^\sigma, [x_i s_i]) \in D(P)$, and that $y_i = [s_{i-1}^\sigma x_i] s_i = s_{i-1}^\sigma [x_i s_i]$.

Theorem 2.3.11. *[53, 3.A.2.7, 3.A.2.11, 3.A.4.5 & 3.A.4.6] Let P be a pregroup and let $X = P \setminus \{1\}$. Let $u, v \in U(P)$ such that u is P-reduced. Then*

(i) if $u \approx v$ then v 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 as a P-reduced element over X^*;

(iv) if v is P-reduced, then u and v represent the same element of $U(P)$ if and only if $u \approx w$. In particular, P embeds into $U(P)$.

Corollary 2.3.12. *[34, Corollary 2.10] Let P be a finite pregroup. Then the word problem in $U(P)$ can be solved in linear time.*

Proof. Let $w \in X^*$. By Theorem 2.3.11 $w =_{U(P)} 1$ if and only if the P-reduced form of w is the empty word. Hence if $w =_{U(P)} 1$; w has length $|w| \geq 1$ and w is σ-reduced, then w cannot be P-reduced, so a length-reducing rewrite derived from V_P applies to w. Thus, $U(P)$ is a finitely generated group (as P is finite) and Dehn's algorithm solves the word problem in $U(P)$. ■

Definition 2.3.13. [34, Definition 4.14] Let $a, b \in X$. We say that (a, b) is an *intermult pair* and that a *intermults* with b if $b \neq a^\sigma$, and either $(a, b) \in D(P)$ or there is $x \in X$ such that $(a, x), (x^\sigma, b) \in D(P)$.

Note that by the definition of V_P, if $ab[ab]^\sigma \in V_P$, then (a, b) is an intermult pair.

We now present the definition of a pregroup presentation.

Definition 2.3.14. [34, Definition 2.11] Let P be a pregroup, let $X = P \setminus \{1\}$, let σ be the involution that gives inverses in X, and let $\mathcal{R}$ be a set of cyclically P-reduced words over X. We define the pregroup presentation to be the 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, and write $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$.

Observe that $\mathcal{R} \cap V_P = \emptyset$ since each word in $\mathcal{R}$ is cyclically P-reduced.

Assumption 2.3.15. We assume throughout this thesis that there is no $x \in X$ such that $x^2 \in \mathcal{R}$, that no $R \in \mathcal{R}$ has length $|R| \in \{1, 2\}$, and that 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.

Before running our algorithms, we shall always assume that the preprocessing from [34, Section 7.1] has been done to the input pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$. This process ensures that $\mathcal{P}$ satisfies Assumption 2.3.15.

Theorem 2.3.16. *[49, Corollary to Theorem B] A finitely generated group is virtually free if and only if it is a universal group of a finite pregroup.*

It is well-known that all amalgamated free products of finite groups and HNN extensions with finite base groups are virtually free – these classes provide many useful examples for the algorithmic solutions in this thesis. More generally, Serre in [51, Proposition 11] classifies virtually free groups as fundamental groups of finite graphs of groups with finite vertex groups.

We shall be working with groups given by finite pregroup presentations. The following corollary to Theorem 2.3.16 states that these are precisely the quotients of virtually free groups by finitely many additional relators.

Corollary 2.3.17. *[34, Corollary 2.15] Let G have a pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$. 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 $U(P)$. Moreover, any group that is a quotient of a virtually free group by finitely many additional relators has a finite pregroup presentation.*

In [34, Section 4.1] a coarser relation than interleaving (see Definition 2.3.10) was introduced on the set of cyclically P-reduced words, where the condition $s_0 = 1 = s_n$ was replaced by $s_0 = s_n$.

Definition 2.3.18. Let $w = x_1x_2 \dots x_n \in X^*$ be cyclically P-reduced, and let $v = y_1y_2 \dots y_n$ be any word over X. We say that v is a *cyclic interleave* of w, and write $w \approx^c v$ if either $n \le 1$ and $w = v$, or if $n > 1$ and there exist $s_0, s_1, s_2, \dots, s_{n-1}, s_n = s_0 \in P$ such that $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i]) \in D(P)$ for $1 \le i \le n$, and $y_i = [s_{i-1}^\sigma x_i s_i]$.

Theorem 2.3.19. *[34, Theorem 4.4] Let $w \in X^*$ be cyclically P-reduced. If $v \in X^*$ satisfies $v \approx^c w$, then v is cyclically P-reduced. Moreover, $\approx^c$ is an equivalence relation on the set of all cyclically P-reduced words.*

Definition 2.3.20. [34, Definition 4.5] Let $w \in X^*$ be cyclically P-reduced. The *cyclic interleave class* of w is the set $\mathcal{I}(w)$ defined as

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

We further write $\mathcal{I}(\mathcal{R})$ for $\cup_{R\in\mathcal{R}}\mathcal{I}(R)$.

Lemma 2.3.21. *Let $w = x_1 \cdots x_n \in X^*$ be P-reduced and $n \ge 2$. If either*

(a) w' is an interleave of w, or

(b) w is cyclically P-reduced and $w' \in \mathcal{I}(w)$,

then we can obtain w' from w by applying a sequence of at most n single rewrites.

Proof. The proof is the same as the proof of [34, Lemma 4.6], but we include it for completeness. If Assumptions (b) holds, then the lemma is precisely [34, Lemma 4.6]. So suppose that Assumption (a) holds. By Definition 2.3.10 there exist sequences $1 = s_0, s_1, s_2, \dots, s_{n-1}, s_n = 1 \in P$ and $y_1, y_2, \dots, y_n \in X$ such that $w' = y_1 \dots y_n$, $(s_{i-1}^\sigma, x_i), (x_i, s_i), ([s_{i-1}^\sigma x_i], s_i]) \in D(P)$ for all i, and $y_i = [s_{i-1}^\sigma x_i s_i]$. By Theorem 2.3.11 $\approx$ is an equivalence relation on the set of P-reduced words. Therefore, we can construct a sequence $w' = w_m, w_{m-1}, \dots, w_0 = w$ of P-reduced words such that w_{i+1} can be obtained from w_i by replacing a (cyclic sub-word) $y'_jy'_{j+1}$ of w_i with a length 2 word $[y'_js_j][s_j^\sigma y'_{j+1}]$ with $s_j \ne 1$. There are at most $n-1$ letters s_j with $s_j \ne_P 1$, so we are done. ∎

The following terminology was introduced in [23].

Definition 2.3.22. A group is *plain* if it is isomorphic to a free product of finitely many factors, with each factor a finite group or an infinite cyclic group.

We shall consider groups given by finite pregroup presentations $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$ such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$. This is a common case – combining the constructions from Examples 2.3.6-2.3.7, we deduce that all quotients of plain groups by finitely many additional relators can be defined in this way. E.g. the pregroup P from Example 2.3.3 satisfies $U(P) \cong S_3 \times C_3 \times \mathbb{Z}$, so we can define $\mathcal{P}$ with $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ by adding additional relators to the presentation defining $U(P)$.

Definition 2.3.23. Let P be a pregroup, and let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$ be a pregroup presentation with $X = P \setminus \{1\}$. We say that P and $\mathcal{P}$ satisfy *trivial-interleaving* if whenever $(a, b) \in P \times P$ is an intermult pair, then $(a, b) \in D(P)$. In particular, we have $\mathcal{I}(\mathcal{R}) = \mathcal{R}$.

2.4 Finite pregroups satisfying trivial-interleaving

In this section we shall give an example of a finite pregroup satisfying trivial-interleaving that has a different form than the standard constructions for obtaining free groups (see Example 2.3.6) and free products of finite groups (see Example 2.3.7), and prove that if P is a finite pregroup satisfying trivial-interleaving, then $U(P)$ is a plain group. All uncited results in this section are new.

Example 2.4.1. Consider a pregroup $P = \{1, x, y, z, t, s, m, X, Y, Z, T, S, M\}$ given by Table 2.2, where a blank space means that the product is undefined. Then P satisfies trivial-interleaving (we used the `GAP` package `Walrus` to produce the pregroup multiplication table and check that P satisfies trivial-interleaving, see [26]). Since $(X, Y) \in D(P)$; $[XY] = T$; and $(\alpha, \alpha) \notin D(P)$ for all $\alpha \in X$, we deduce that P has different form than the constructions from Examples 2.3.6-2.3.7. In the next proposition we show that $U(P)$ is in fact a free group.

Proposition 2.4.2. *Let P be the pregroup from Example 2.4.1. Then $U(P)$ is a free group.*

Proof. It was shown by Stallings in [54] that every torsion-free virtually free group is free. Hence by Theorem 2.3.16 it suffices to show that $U(P)$ is torsion-free. Let $w \in X^*$ with $w \neq_{U(P)} 1$. Since $w \neq_{U(P)} 1$, there exists a cyclically P-reduced $U(P)$-conjugate w' of w with $w' \neq_{U(P)} 1$. It suffices to show that w' has infinite order.

Write $w' = w_1 \dots w_n$ with $w_i \in X^\sigma$ for all $1 \leq i \leq n$. Suppose first that $n = 1$. Then from Table 2.2 we have $(w', w') \notin D(P)$. Hence for every $k \geq 1$, the word $w'^k = w_1^k$ is P-reduced. If $k = 1$, then by Theorem 2.3.11 we have $w'^k \neq_{U(P)} 1$. Otherwise, w'^k has length $|w'^k| \geq 2$, so by Definition 2.3.10 we have $w'^k \not\approx 1$, hence again by Theorem 2.3.11 $w'^k \neq_{U(P)} 1$. Thus, w' has infinite order.

Suppose instead that $n \geq 2$. Then as by Definition 2.3.9 we have $(w_n, w_1) \notin D(P)$ and $(w_i, w_{i+1}) \notin D(P)$ for $1 \leq i \leq n-1$, it follows that for every $k \geq 1$, w'^k is P-reduced. So by Definition 2.3.10 $w'^k \not\approx 1$, and by Theorem 2.3.11 $w'^k \neq_{U(P)} 1$. The proposition follows. ∎

Table 2.2: Pregroup multiplication table

1	X	Y	Z	T	S	M	x	y	z	t	s	m
X		T			M		1					
Y			S					1		x		
Z									1		y	t
T			M					X		1		
S									Y		1	x
M									T		X	1
x	1			Y		S						
y		1			Z		t					
z			1					s		m		
t	y			1		Z						
s		z			1		m					
m	s			z		1						

For the rest of this section our aim is to show that if P is a finite pregroup satisfying trivial-interleaving, then $U(P)$ is a plain group. The next two definitions were taken from [23].

Definition 2.4.3. A *rewriting system* is a pair (ζ, T), where ζ is a finite alphabet of symbols, and $T \subseteq \zeta^* \times \zeta^*$ is a set of rewriting rules.

Let $w \in \zeta^*$. If $l, r \in T$ and l is a sub-word of w, then we can replace l by r in w to obtain a word v. Then $w \overset{*}{\to} v$ means that v can be obtained from w by applying a finite sequence of rewrites. If w can be transformed into v by a sequence consisting of a single rewrite, then we write $w \to v$. If we further define, for all $w, v \in \zeta^*$: $w \overset{*}{\to} w$, and $w \overset{*}{\to} v$ if and only if $v \overset{*}{\to} w$, then $\overset{*}{\to}$ becomes an equivalence relation on ζ^*.

Now if $v_1, w_1, v_2, w_2 \in \zeta^*$ are such that $[v_1]_{\overset{*}{\to}} = [w_1]_{\overset{*}{\to}}$ and $[v_2]_{\overset{*}{\to}} = [w_2]_{\overset{*}{\to}}$, then $v_1v_2 \overset{*}{\to} w_1v_2 \overset{*}{\to} w_1w_2$, hence $[v_1v_2]_{\overset{*}{\to}} = [w_1w_2]_{\overset{*}{\to}}$. Thus, the operation of concatenation of representatives gives us a well-defined product on the set of equivalence classes, and it turns this set into a monoid: we will restrict our focus on rewriting systems that define universal groups of finite pregroups.

Definition 2.4.4. [23, Definitions 2.1 & 2.2]. A rewriting system (ζ, T) is called

1. *finite* if both ζ and T are finite;

2. *confluent* if, for all $u, v, w \in \zeta^*$ such that $w \overset{*}{\to} u$ and $w \overset{*}{\to} v$, there exists q such that $u \overset{*}{\to} q$ and $v \overset{*}{\to} q$;

3. *strongly-confluent* if, for all $u, v, w \in \zeta^*$ such that $w \overset{*}{\to} u$ and $w \overset{*}{\to} v$, there exists q such that $u \to q$ and $v \to q$;

4. *terminating* if every rewriting sequence terminates in a finite number of steps;

5. *convergent* if it is both confluent and terminating;

6. *length-reducing* if $|r| < |l|$ for every $(l, r) \in T$;

7. *monadic* if $|r| \leq 1$ for every $(l, r) \in T$.

Let P be a pregroup. The next definition gives us two rewriting systems for obtaining $U(P)$.

Definition 2.4.5. [17, Definition 3.1]

1. The rewriting system $S_1 \subseteq P^* \times P^*$ is defined as follows:

$$\begin{aligned} 1 &\to \eta \quad (= \text{the empty word}) \\ ab &\to [ab] \quad (\text{if } (a,b) \in D(P)) \\ ab &\to [ax][x^\sigma b] \quad (\text{if } (a,x), (x^\sigma, b) \in D(P)). \end{aligned}$$

2. Let $X = P \setminus \{1\}$. The rewriting system $S(P) \subseteq X^* \times X^*$ is defined as follows:

$$\begin{aligned} ab &\to \eta \quad (\text{if } (a,b) \in D(P) \text{ and } [ab] = 1) \\ ab &\to [ab] \quad (\text{if } (a,b) \in D(P) \text{ and } [ab] \neq 1) \\ ab &\to [ax][x^\sigma b] \quad (\text{if } (a,x), (x^\sigma, b) \in D(P) \text{ and } (a,b) \notin D(P)). \end{aligned}$$

It was shown in [17] that $U(P)$ is given by

$$X^*/\{(l,r) \mid (l,r) \in S_1\}$$

and

$$X^*/\{(l,r) \mid (l,r) \in S(P)\}.$$

In the proof of [19, Theorem 8.4] the following is shown.

Lemma 2.4.6. *The system S_1 is strongly-confluent.*

Subsequently, in [19, Remark 8.6] it was noted that $S(P)$ is confluent.

Proposition 2.4.7. *Let P be a finite pregroup satisfying trivial-interleaving. Then $S(P)$ is a finite, convergent, length-reducing, monadic rewriting system.*

Proof. Since P is finite, so is $S(P)$. Observe that the third rewriting rule never applies since P satisfies trivial-interleaving, hence $S(P)$ is both length-reducing and monadic. A finite length-reducing rewriting system is terminating, so as $S(P)$ is confluent, $S(P)$ is convergent. ■

Furthermore, in [23, Theorem 5.3] Eisenberg and Piggott provided a positive answer to Gilman's conjecture (see [29]).

Theorem 2.4.8. (Gilman's conjecture) *Let G be a group. Then G admits a presentation by a finite, convergent, length-reducing, monadic rewriting system if and only if G is a plain group.*

Combining Proposition 2.4.7 with Theorem 2.4.8 we obtain our desired result.

Theorem 2.4.9. *Let P be a finite pregroup satisfying trivial-interleaving. Then $U(P)$ is a plain group.*

2.5 Coloured diagrams over pregroups

In this section we present a natural generalisation of van Kampen diagrams for finite pregroup presentations, coloured van Kampen diagrams, that were first defined in [34]. We also give the definition of coloured annular diagrams, which will be extensively studied in Part I of this thesis. We use the standard terminology for maps and diagrams given in [42, Chapter 5], and this section presents definitions from [42] that we shall use throughout our work. Otherwise, most definitions in this section are from [34].

Recall Definition 2.1.5 of the closure of a subset of a metric space, Definition 2.1.7 of a vertex, an edge, a region and a path, and Definition 2.1.9 of connected, simply-connected, and annular subsets of $\mathbb{R}^2$.

Definition 2.5.1. A *map* M is a finite collection of vertices, edges, and regions which are pairwise disjoint and satisfy the following conditions.

(i) If e is an edge of M, then there are (not necessarily distinct) vertices a and b in M such that $\overline{e} = e \cup \{a\} \cup \{b\}$.

(ii) The boundary, $\partial(R)$, of each region R of M is connected and there is a set of edges $e_1, \ldots, e_n$ in M such that $\partial(R) = \overline{e_1} \cup \ldots \cup \overline{e_n}$.

We shall call the regions of M *internal faces* of M, and the components of $\mathbb{R}^2 \setminus M$ *external faces* of M. We shall also use M to denote the set-theoretic union of its vertices, edges, and internal faces. Then the boundary of M is denoted as $\partial(M)$. If e is an edge with $\overline{e} = e \cup \{a\} \cup \{b\}$, the vertices a and b are called the *endpoints* of e. An edge with equal endpoints is a *loop*.

We shall always assume that M is connected. Note also that the boundaries of faces of M are paths.

Definition 2.5.2. Let M be a map with subsets A and B consisting of vertices, edges and internal faces of M. We say that A is *incident* with B if $\overline{A} \cap \overline{B} \neq \emptyset$. We say that A is *edge-incident* with B if $\overline{A} \cap \overline{B}$ contains an edge of M.

Suppose that x is a vertex or an edge of M. If $x \in \overline{A} \cap \overline{B}$, then we say that A *shares* x with B, and that x is *common* to A and B.

Definition 2.5.3. Let M be an annular map, let O be the unbounded component of $\mathbb{R}^2 \setminus M$, and let I be the bounded component of $\mathbb{R}^2 \setminus M$. We call $\partial(M) \cap \partial(O)$ the *outer boundary* of M, which will be denoted by ω, and $\partial(M) \cap \partial(I)$ the *inner boundary* of M, which will be denoted by τ.

Definition 2.5.4. The maps will be oriented. Let M be a map. If M is annular, then we orient the external face with boundary ω counter-clockwise, and all other faces (including the external face with boundary τ) clockwise. Otherwise, all (including external) faces of M are oriented clockwise. Each edge of M is composed of two *half-edges*. Each half-edge is associated with the face on one of the sides distinct from the one associated with the other half-edge, and inherits its orientation from that face. If F is a face of M with the given orientation, any cycle α of minimal length which includes all the edges of $\partial(F)$ is a *boundary cycle* of F.

We now present the definition of diagrams and coloured diagrams, which extends [34, Definition 3.2] and defines (coloured) diagrams over pregroup presentations that are not necessarily simply-connected. For the rest of this section fix a finite pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.20), and let G be defined by $\mathcal{P}$.

Definition 2.5.5. A *diagram over* $\mathcal{P}$ is a map Γ and a function ϕ assigning to each half-edge of Γ a label from X^σ. If F is a face of Γ with the given orientation, then a label of F is a word resulted from concatenating the labels of half-edges of some boundary cycle of F oriented by F. The labels of the external faces of Γ are called the *boundary words* of Γ.

A *coloured diagram over* $\mathcal{P}$ is a diagram Γ over $\mathcal{P}$ in which labels of internal faces are from $V_P \cup \mathcal{R}^{\pm 1}$. The internal faces of Γ labelled by relators from V_P are coloured red, and the ones labelled by relators from $\mathcal{R}^{\pm 1}$, together with each external face, are coloured green. Each half-edge inherits beside its orientation also its colour from that face. The red faces of Γ are called *red triangles*.

A *coloured van Kampen diagram* is a coloured diagram that is simply-connected, and a *coloured annular diagram* is simply an annular coloured diagram.

The next definition summarizes terminology for coloured diagrams that we shall use throughout this thesis. It is based on [34, Definition 3.1].

Definition 2.5.6. Let Γ be a coloured diagram. If an element $x \in X$ satisfies $x^\sigma = x$, then x has order 2 in $U(P)$, and we will identify x with x^σ, so than an edge may have label x on both sides. If $F, F' \subseteq \Gamma$ are faces, we write $|\partial(F) \cap \partial(F')|$ for the number of edges in $\partial(F) \cap \partial(F')$ (note that $|\partial(F) \cap \partial(F')| = 0$ if $\partial(F) \cap \partial(F')$ contains no edge). Similarly, we write $|\partial(F) \cap \partial(\Gamma)|$ for the number of edges in $\partial(F) \cap \partial(\Gamma)$. A *consolidated edge* between (not necessarily internal) faces F and F' is a non-empty path of maximal length that is a sub-path of both $\partial(F)$ and $\partial(F')$. We write

$$\mathbf{Area}(\Gamma) := \text{ the number of internal faces of } \Gamma.$$

Vertices and edges contained in $\partial(\Gamma)$ are called *boundary vertices* and *boundary edges*. An internal face F of Γ is a *boundary face* if and only if $\partial(F) \cap \partial(\Gamma)$ contains an edge. Vertices, edges and internal faces which are not boundary are called *interior*.

All incidences are counted with multiplicities, so for example a vertex v can be incident more than once with the same face F (note that F is incident $n \geq 1$ times with v if and only

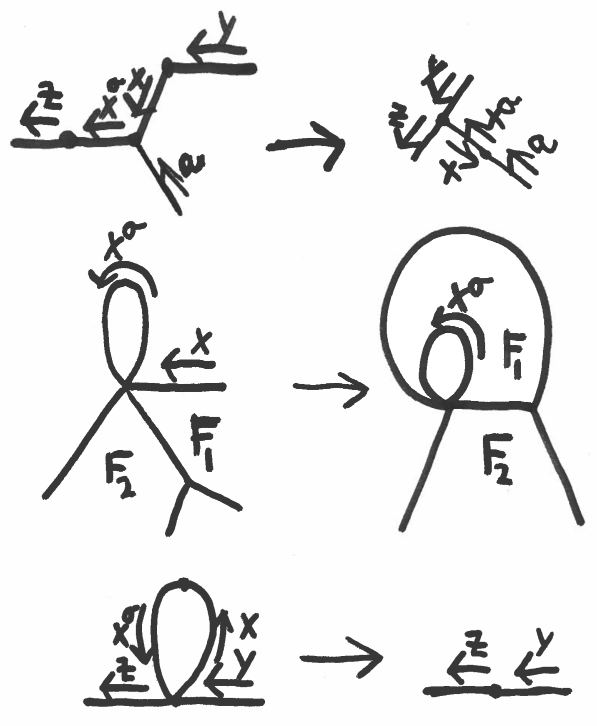

Figure 2.1: σ-reducion of w', see the proof of Lemma 2.5.7.

if $\partial(F)$ passes through v n times). We denote by $\delta(v,\Gamma)$ the *degree* of v in Γ, $\delta_G(v,\Gamma)$ for the *green degree* of v in Γ: the number of green faces of Γ incident with v, $\delta_G^I(v,\Gamma)$ for the *internal green degree* of v in Γ: the number of internal green faces of Γ incident with v, and $\delta_R(v,\Gamma)$ for the *red degree* of v in Γ: the number of red faces of Γ incident with v. When it is clear which coloured diagram Γ is considered, we omit the Γ.

Finally, Γ is *green-rich* if every vertex $v \in \Gamma$ satisfies $\delta_G(v,\Gamma) \geq 2$.

Lemma 2.5.7. *Let $w = x_1 \ldots x_n \in X^*$. Then both of the following statements hold.*

1. *Suppose that w is a boundary word of a simply-connected diagram Γ over $\mathcal{P}$. Then the following holds.*
   - *(i) There exists a simply-connected diagram over $\mathcal{P}$ with a σ-reduced boundary word equal to w in $F(X^\sigma)$.*
   - *(ii) There exists a simply-connected diagram over $\mathcal{P}$ with a P-reduced boundary word equal to w in $U(P)$.*
2. *Suppose that either*
   - *(a) w is P-reduced and there exists a simply-connected diagram Γ over $\mathcal{P}$ with boundary word w' that is an interleave of w; or*
   - *(b) w is cyclically P-reduced and there exists a simply-connected diagram Γ over $\mathcal{P}$ with boundary word $w' \in \mathcal{I}(w)$.*

Then w is a boundary word of a simply-connected diagram over $\mathcal{P}$ of area at most $\mathbf{Area}(\Gamma) + 2n$.

Proof. Part 1. Let Γ' be a simply-connected diagram over $\mathcal{P}$, and let w' be the label of some boundary cycle of Γ' with endpoint v'. Suppose first that w' contains a sub-word xx^σ. Then as in the case of free reduction (see [42, Chapter 5, Section 1]), we can fold the edges (or delete the sub-diagram of Γ' enclosed by them, see Figure 2.1) with labels x and x^σ to produce a simply-connected diagram with boundary cycle beginning at v' and labelled by a word equal to w' in $F(X^\sigma)$, but in which the sub-word xx^σ was deleted from w.

Suppose instead that w' contains two consecutive letters x, y with $x \neq y^\sigma$ and $(x, y) \in D(P)$, and let $e, f \subseteq \partial(\Gamma')$ be the edges labelled x and y respectively. Then we attach a red triangle T to Γ' (labelled by $xy[xy]^\sigma$) at the path ef. This produces a simply-connected diagram with boundary cycle beginning at v' and labelled by a word equal to w' in $U(P)$, but in which the sub-word xy was replaced by $[xy]$.

As by Corollary 2.3.12 we can solve the word problem in $U(P)$ in linear time, by the previous two paragraphs we can obtain both diagrams in finitely many steps. Hence Part 1 follows.

Part 2. We prove the lemma under the Assumption (a) as the proof under the Assumption (b) is very similar. If $w = w'$, then Γ satisfies the lemma. So suppose that $w' \neq w$. Then by Definition 2.3.10 $|w| \geq 2$. Hence by Lemma 2.3.21 we can obtain w from w' by applying a sequence of at most n single rewrites. Therefore, there is a sequence $w = w_m, w_{m-1}, \ldots, w_0 = w' = y_1 \ldots y_n$ (with $1 \leq m \leq n$) of P-reduced words such that for each j, w_{j+1} can be obtained from w_j by replacing a (cyclic sub-word) $y_i' y_{i+1}'$ of w_j with a length 2 word $[y_i' s_i][s_i^\sigma y_{i+1}']$ with $s_i \neq 1$. Hence we can obtain a simply-connected diagram Γ' with boundary word w by constructing a sequence $\Gamma' = \Gamma_m, \Gamma_{m-1}, \ldots, \Gamma_0 = \Gamma$ of simply-connected diagrams with boundary words w_i, where Γ_{j+1} is obtained from Γ_j by attaching two triangles T_i and T_i' at edges labelled by y_i' and y_{i+1}' respectively to Γ_j, such that T_i and T_i' share an edge labelled s_i, and have labels $y_i' s_i [y_i' s_i]^\sigma$ and $s_i^\sigma y_{i+1}' [s_i^\sigma y_{i+1}']^\sigma$ respectively. In particular, $\mathbf{Area}(\Gamma') \leq \mathbf{Area}(\Gamma) + 2n$, so we are done. ∎

The following theorem is a standard result which shows that coloured van Kampen diagrams are a great tool for studying the word problem and for showing hyperbolicity. In its proof, we use the ideas from the proofs of [42, Chapter 5, Theorem 1.1 & Lemma 1.2].

Theorem 2.5.8. (van Kampen's lemma) *Let $w \in X^*$. Then all of the following statements hold.*

1. *Suppose that w is P-reduced and $w =_G 1$. Then there exists a coloured van Kampen diagram Γ over $\mathcal{P}$ with boundary word w.*

2. *Suppose that w is σ-reduced and $w =_{U(P)} 1$. Then there there exists a coloured van Kampen diagram Γ over $\mathcal{P}$ with boundary word w and with no green faces.*

3. *Suppose that w is a boundary word of a coloured van Kampen diagram Γ over $\mathcal{P}$ that contains m internal faces $F_1, \ldots, F_m$. Then there exist labels R_i of F_i and $u_i \in F(X^\sigma)$, $1 \leq i \leq m$, such that*

$$w =_{F(X^\sigma)} u_1^{-1} R_1 u_1 \cdots u_m^{-1} R_m u_m.$$

Hence $w =_G 1$, and if $F_1, \ldots, F_m$ are all red, then $w =_{U(P)} 1$.

Proof. Part 1. If $w =_{U(P)} 1$, then by Theorem 2.3.11 w is the empty word and we take Γ to consist of a single vertex. Otherwise, by Corollary 2.3.17 there exist $R_i \in \mathcal{R}^{\pm 1}$ and $u_i \in U(P)$, $1 \leq i \leq k$, such that

$$w =_{U(P)} u_1^{-1} R_1 u_1 \cdots u_k^{-1} R_k u_k.$$

We construct Γ as follows. First, let v be a vertex. Then attach paths p_i to v with endpoints v and v_i such that $|p_i| = |u_i|$, and such that $(p_1, p_2, \ldots, p_k)$ is a cyclic tuple when reading counter-clockwise around v. Now for each vertex v_i, attach an internal green face F_i with label R_i^{-1}, where R_i^{-1} is the label of the boundary cycle of F_i with endpoint v_i. Call the resulting map Γ'. Next label the half-edges of p_i so that the boundary cycle of Γ' with endpoint v has label $\Pi_{i=1}^k (u_i')^{-1} R_i u_i'$, where each u_i' is a P-reduced word representing u_i. Then by Lemma 2.5.7 there exists a simply-connected diagram over $\mathcal{P}$ with a P-reduced boundary word w' with $w' =_{U(P)} \Pi_{i=1}^k (u_i')^{-1} R_i u_i' =_{U(P)} w$, and such that all its internal faces are labelled by elements from $V_P \cup \mathcal{R}^{\pm 1}$. By Theorem 2.3.11 w' is an interleave of w, so applying Lemma 2.5.7 again shows that Γ exists.

Part 2. If $w =_{F(X^\sigma)} 1$, then w is the empty word and we take Γ to consist of a single vertex. Otherwise, there exist $R_i \in V_P$ and $u_i \in F(X^\sigma)$, $1 \leq i \leq k$, such that $w =_{F(X^\sigma)} \Pi_{i=1}^k u_i^{-1} R_i u_i$. Hence we proceed as in Part 1 and construct a simply-connected diagram Γ' with boundary word $\Pi_{i=1}^k u_i^{-1} R_i u_i$ and with internal faces $F_1, F_2, \ldots, F_k$ labelled by $R_1^{-1}, \ldots, R_k^{-1}$ respectively. By Remark 2.3.5 we have $R_i^{-1} \in V_P$, so by Lemma 2.5.7 there exists a simply-connected diagram over $\mathcal{P}$ with a σ-reduced boundary word w' equal to $\Pi_{i=1}^k u_i^{-1} R_i u_i$ in $F(X^\sigma)$ and with all internal faces labelled by relators from V_P. As w is σ-reduced, we have $w = w'$, so Part 2 follows.

Part 3. The proof is by induction on m. Base case $m = 0$. There is nothing to prove, but Γ is a tree, so $w =_{F(X^\sigma)} 1$, and hence $w =_G 1$. So suppose that $m \geq 1$, and let l be the boundary cycle of Γ with label w. Now there exists an internal face F with an edge e on $\partial(\Gamma)$. Let x be the letter of w that labels e, write $w = s_1 x s_2$ for some $s_1, s_2 \in X^*$, and let $x^\sigma s$ be a label of F. Delete e from Γ to obtain a coloured diagram Γ'. Then Γ' is simply-connected, with $m-1$ internal faces, and there is a boundary cycle of Γ' with the same endpoints as l, and with label $s_1 s s_2$. By induction we can list the internal faces of Γ' as $F_1, \ldots, F_{m-1}$ to obtain

$$s_1 s s_2 =_{F(X^\sigma)} u_1^{-1} R_1 u_1 \cdots u_{m-1}^{-1} R_{m-1} u_{m-1},$$

where R_i is a label of F_i and $u_i \in F(X^\sigma)$ for $1 \leq i \leq m-1$. Now note that $w = s_1 x s_2 =_{F(X^\sigma)} (s_1 s s_2)(s_2^{-1} s^{-1} x s_2)$, and $s^{-1} x$ is the inverse in $F(X^\sigma)$ of the label of F.

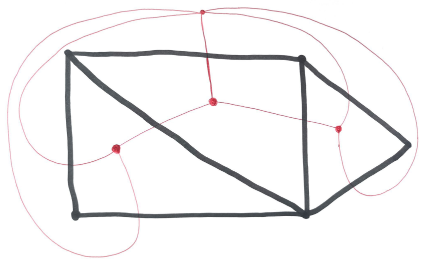

Figure 2.2: A coloured diagram Γ (in black) and its dual Γ^* (in red), see Definition 2.5.12.

By Remark 2.3.5 $s^{-1}x$ is a cyclic conjugate of some $R \in V_P \cup \mathcal{R}^{\pm 1}$. Therefore, there exists $u \in X^*$ such that $s^{-1}x =_{F(X^\sigma)} u^{-1}Ru$, and hence $s_2^{-1}s^{-1}xs_2 =_{F(X^\sigma)} s_2^{-1}u^{-1}Rus_2$. Thus, Part 2 follows by taking $u_m =_{F(X^\sigma)} us_2$ and $R_m = R$. ∎

We extend [34, Definition 3.4] and define a coloured area for coloured diagrams.

Definition 2.5.9. Let Γ be a coloured diagram. We define the coloured area of Γ, denoted as $\mathbf{CArea}(\Gamma)$, to be an ordered pair (a, b), where a is the number of internal green faces of Γ and b is the number of red triangles. Suppose that Δ is a coloured diagram with $\mathbf{CArea}(\Delta) = (c, d)$. We say that $\mathbf{CArea}(\Gamma) \leq \mathbf{CArea}(\Delta)$ if $a < c$ (in which case $\mathbf{CArea}(\Gamma) < \mathbf{CArea}(\Delta)$) or if $a = c$ and $b \leq d$ (if $b < d$ then we say $\mathbf{CArea}(\Gamma) < \mathbf{CArea}(\Delta)$).

We next define sub-diagrams ([34, Definition 3.5]) and islands ([42, Chapter 5, page 257]).

Definition 2.5.10. Let Γ be a coloured diagram. A *sub-diagram* of Γ is a subset of the edges, vertices and internal faces of Γ which, together with new external faces coloured green, form a coloured diagram in its own right.

An *island* of a coloured annular diagram Γ_A is a sub-diagram of Γ_A bounded by a closed path of the form $\omega_1\tau_1$, where $\omega_1 \subseteq \omega$; $\tau_1 \subseteq \tau$; and $|\omega_1|, |\tau_1| \geq 1$. The endpoints of ω_1 are called the *endpoints of* E. A *bridge* is an edge in $\omega \cap \tau$. We say that Γ_A is *island-free* if Γ_A contains no islands.

Remark 2.5.11. We shall encounter cases where Γ_A is a single island.

The next definition considers a special kind of sub-diagrams, consisting entirely of red triangles, that as we shall see have a great impact on the overall structure of diagrams that will study.

Definition 2.5.12. [34, Definition 4.11] A *red blob* in a coloured diagram Γ is a non-empty subset B of the set of closures of red triangles of Γ, with the property that any non-empty proper subset C of B has at least one edge in common with $B \setminus C$. Equivalently, the induced sub-graph B^* of the dual graph Γ^* of Γ (see Figure 2.2 for an example of a dual graph) on those vertices that correspond to the triangles in B is connected. In particular, B is a sub-diagram of Γ. To simplify our statements, if $\mathbf{Area}(B) = 1$, then we often say that B is a red triangle, but we always mean that B is the closure of the triangle contained in it.

Semi-σ-reduction

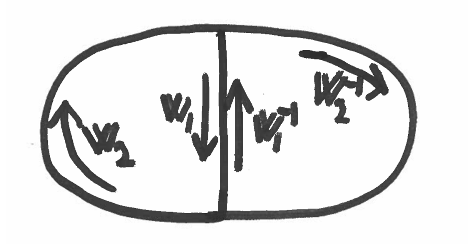

Figure 2.3: Two faces making the ambient diagram not semi-σ-reduced, see Definition 2.5.15.

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

Recall Definition 2.3.13 of an intermult pair, and Definition 2.3.23 of trivial-interleaving. Part 1 of the next result ([34, Lemma 4.16]) is stated only for coloured simply-connected diagrams, but the proof does not assume it.

Lemma 2.5.13. *Let Γ be a coloured diagram over $\mathcal{P}$ with a red blob B. Suppose that $a, b \in X$ and ab is a sub-word of a boundary word of B. Then both of the following statements hold.*

1. *If $b \neq a^\sigma$, then a intermults with b.*

2. *If, in addition, $\mathcal{P}$ satisfies trivial-interleaving, then $(a, b) \in D(P)$.*

Proof. Part 1. This is ([34, Lemma 4.16]).

Part 2. If $b = a^\sigma$, then by Axiom P2 we have $(a, b) \in D(P)$. Otherwise, by Part 1, (a, b) is an intermult pair, so as $\mathcal{P}$ satisfies trivial-interleaving, we have $(a, b) \in D(P)$. ∎

Definition 2.5.14. Let Γ be a coloured diagram. The *1-skeleton* of Γ is a graph Γ^1 defined as

$$\Gamma^1 := \bigcup_{F \text{ is a face of } \Gamma} \partial(F).$$

Definition 2.5.15. [34, Definition 3.6] Let Γ be a coloured diagram over $\mathcal{P}$. We say that Γ is *semi-σ-reduced* if no two distinct incident faces are labelled by $w_1 w_2$ and $w_2^{-1} w_1^{-1}$ for some relator $w_1 w_2 \in V_P \cup \mathcal{R}^\pm$ and have a common consolidated edge labelled by w_1 and w_1^{-1} (see Figure 2.3). It is *σ-reduced* if the same holds for a single face edge-incident with itself (see Figure 2.4).

A natural generalization of semi-σ-reduction is the following definition.

Definition 2.5.16. [34, Definition 3.7] We say that a coloured diagram Γ over $\mathcal{P}$ is *semi-P-reduced* if no two distinct incident green faces are labelled by $w_1 w_2$ and $w_3^{-1} w_1^{-1}$ and have a common consolidated edge labelled by w_1 and w_1^{-1}, where $w_2 =_{U(P)} w_3$ (see Figure 2.5).

σ-reduction

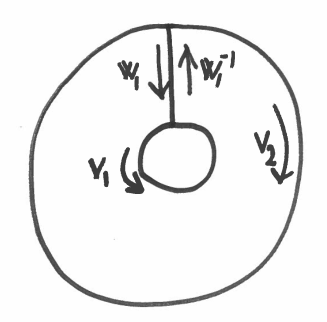

Figure 2.4: A face making the ambient diagram not σ-reduced. The words w_1, v_1, v_2 satisfy $v_1 w_1^{-1} v_2 =_{F(X^\sigma)} (v_2 w_1 v_1)^{-1} = v_1^{-1} w_1^{-1} v_2^{-1}$, see Definition 2.5.15.

Semi-P-reduction

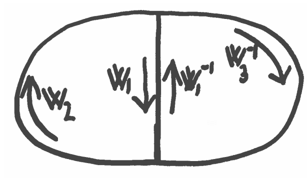

Figure 2.5: Two green faces making the ambient diagram not semi-P-reduced. The words w_2 and w_3 satisfy $w_2 =_{U(P)} w_3$, see Definition 2.5.16.

We conclude this section by presenting an important set $\mathcal{D}$ of coloured van Kampen diagrams, proven fruitful for showing hyperbolicity of groups defined by finite pregroup presentations. These will later inspire us to define new sets of annular diagrams with very similar properties that enable us to solve the conjugacy problem in quadratic time.

Definition 2.5.17. [34, Definition 6.1] We define $\mathcal{D}$ to be the set of all coloured van Kampen diagrams Γ over $\mathcal{P}$ with the following properties.

1. The boundary word of Γ is cyclically P-reduced (see Definition 2.3.9).
2. Γ is σ-reduced and semi-P-reduced (see Definitions 2.5.15-2.5.16).
3. Γ is green-rich (see Definition 2.5.6).
4. No proper sub-word of the (cyclic) boundary word of a simply-connected red blob of Γ is equal to 1 in $U(P)$.

2.6 Curvature distribution schemes

Similarly as in [34, Sections 5 & 6], we shall now explain ways of assigning curvature to coloured annular and simply-connected diagrams. Throughout this section, let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$ be a finite pregroup presentation such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.20), and let G be the group defined by $\mathcal{P}$.

The next two definitions extend [34, Definitions 5.1 & 5.2] and define curvature distributions and curvature distributions schemes on coloured diagrams that are not necessarily simply-connected.

Definition 2.6.1. Let Γ be a coloured diagram over $\mathcal{P}$ with vertex set $V(\Gamma)$, edge set $E(\Gamma)$, and set $F(\Gamma)$ of internal faces. A *curvature distribution on* Γ is a function $\rho_\Gamma : V(\Gamma) \cup E(\Gamma) \cup F(\Gamma) \to \mathbb{R}$ such that

$$\rho_\Gamma(\Gamma) = \begin{cases} \sum\limits_{x \in V(\Gamma)\cup E(\Gamma)\cup F(\Gamma)} \rho_\Gamma(x) = 0 & \text{if } \Gamma \text{ is annular,} \\ \sum\limits_{x \in V(\Gamma)\cup E(\Gamma)\cup F(\Gamma)} \rho_\Gamma(x) = 1 & \text{if } \Gamma \text{ is simply-connected.} \end{cases}$$

Definition 2.6.2. Let $\mathcal{K}$ be a non-empty set of coloured annular and simply-connected diagrams over $\mathcal{P}$. A *curvature distribution scheme* on $\mathcal{K}$ is a map $\mathcal{Y} : \mathcal{K} \to \{\rho_\Gamma : \Gamma \in \mathcal{K}\}$ that associates a curvature distribution to every diagram in $\mathcal{K}$.

Example 2.6.3. For an annular or simply-connected coloured diagram Γ, define a function ρ_Γ by setting $\rho(v) := +1$ for each vertex $v \in V(\Gamma)$, $\rho(e) := -1$ for each edge $e \in E(\Gamma)$ and $\rho(F) := 1$ for each internal face F of Γ. Euler's formula then tells us that if Γ is simply-connected, then $\rho_\Gamma(\Gamma) = 1$, and if Γ is annular, then $\rho_\Gamma(\Gamma) = 0$. Thus, ρ_Γ is a curvature distribution. It follows that the map $\Psi : \mathcal{K} \to \{\rho_\Gamma : \Gamma \in \mathcal{K}\}$ is a curvature distribution scheme for any non-empty set $\mathcal{K}$ of coloured annular or simply-connected diagrams over $\mathcal{P}$.

We shall now describe a key curvature distribution scheme, $\mathbf{RSym}$ (see [34, Definition 6.4]), that is computed by the algorithm $\mathbf{ComputeRSym}$ (see [34, Algorithm 6.3]). $\mathbf{ComputeRSym}$ as given in [34] operates on diagrams in $\mathcal{D}$ (see Definition 2.5.17). However, we need more flexibility for our work, hence we allow $\mathbf{ComputeRSym}$ to take as input any coloured annular or simply-connected diagram Γ over $\mathcal{P}$. Apart from that, the algorithm remains unchanged. In particular, $\mathbf{ComputeRSym}$ returns (see Proposition 2.6.6) a curvature distribution $\kappa_\Gamma : \Gamma \to \mathbb{R}$.

$\mathbf{ComputeRSym}$ assigns and alters curvature on the vertices, edges and faces of Γ in several successive steps, where the external face has curvature 0 throughout. As in the description of [34, Algorithm 6.3], when we say (for example) that a half-edge e *gives* curvature c to a 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 $F_1, \ldots, F_k$, 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 $F_1, \ldots, F_k$.

Algorithm 2.6.4. $\mathbf{ComputeRSym}(\Gamma)$:
// Γ: a coloured simply-connected or annular diagram.

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 edge of $\partial(B) \setminus \partial(\Gamma)$ to the (internal) green face on the other side.

Step 5 Return the function $\kappa_\Gamma : V(\Gamma) \cup E(\Gamma) \cup F(\Gamma) \to \mathbb{R}$, where $\kappa_\Gamma(x)$ is the current curvature of x.

Definition 2.6.5. Let $\mathcal{K}$ be a non-empty set of coloured annular and simply-connected diagrams over $\mathcal{P}$. We define $\mathbf{RSym}$ to be the map from $\mathcal{K}$ to $\{\kappa_\Gamma(x) : \Gamma \in \mathcal{K}\}$ evaluated by $\mathbf{ComputeRSym}$. We shall omit the Γ from $\kappa_\Gamma(x)$ and write just $\kappa(x)$ when the meaning is clear.

If B is a red blob of an annular or a simply-connected diagram Γ, then we define

$$\kappa_\Gamma(B) = \sum_{T\colon \text{a red triangle of } B} \kappa_\Gamma(T).$$

Proposition 2.6.6. *Let $\mathcal{K}$ be a non-empty set of coloured annular and simply-connected diagrams over $\mathcal{P}$. Then* $\mathbf{RSym}$ *is a curvature distribution scheme on $\mathcal{K}$.*

Proof. Note that the curvature in Step 1 is precisely the curvature distribution from Example 2.6.3. Since curvature is neither created nor destroyed by $\mathbf{ComputeRSym}$, the proposition follows. ■

Definition 2.6.7. Let Γ be a coloured annular or simply-connected diagram with a green face F. We say that a vertex v *is curvature incident with* F if v is incident with F, and that a red blob B *is curvature incident with* F if B is edge-incident with F.

Definition 2.6.8. Let x be a vertex or a red blob of an annular or simply-connected diagram Γ, and let $F \subseteq \Gamma$ be a green face. We let

1. $\zeta(x, \Gamma)$ be the total curvature that x gives to internal green faces of Γ in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma)$;
2. $\chi(x, \Gamma)$ be the curvature that x gives to a single internal green face of Γ across each curvature incidence in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma)$;
3. $\chi(x, F, \Gamma)$ be the total curvature that x gives to F in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma)$.

The following two lemmas are important for the work in this thesis. They were proved in [34] under the assumption that the vertex v is a vertex of a diagram in $\mathcal{D}$, but the proofs used only the fact that as the diagrams in $\mathcal{D}$ are green-rich, v has green degree at least 2. Hence they can be stated as follows.

Lemma 2.6.9. *[34,* Lemma 6.7*] Let v be a vertex in a coloured annular or simply-connected diagram Γ. Assume that v is incident k times with external faces, and that $\delta_G(v, \Gamma) \geq 2$. Then*

Table 2.3: Vertex curvature $\chi(v, \Gamma)$.

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

(i) $\kappa_\Gamma(v) \leq 0$ *and* $\kappa_\Gamma(v) = 0$ *if* $k \neq \delta_G(v, \Gamma)$*;*

(ii) If $\delta_G(v, \Gamma) \neq k$*, then* $\chi(v, \Gamma) = \frac{2-\delta_G(v,\Gamma)}{2\cdot(\delta_G(v,\Gamma)-k)}$ *(hence if* $k = 2$*, then* $\chi(v, \Gamma) = -1/2$*);*

(iii) If $\delta_G(v, \Gamma) > 2$ *and* $\delta_G(v, \Gamma) \neq k$ *then* $\chi(v, \Gamma) \leq -1/6$.

Lemma 2.6.10. *[34,* Lemma 7.5*] Let v be a vertex in a coloured annular or simply-connected diagram Γ incident with an internal green face, and such that* $\delta_G(v, \Gamma) \geq 2$.

If v is incident k times with external faces for some $k > 1$, then $\chi(v, \Gamma) \leq -1/2$. Otherwise, the curvature $\chi(v, \Gamma)$ is as in Table 2.3.

Definition 2.6.11. [34, Definition 6.6] We say that $\mathbf{RSym}$ *succeeds with a constant* $\varepsilon > 0$ on a diagram $\Gamma \in \mathcal{D}$ if $\kappa_\Gamma(F) \leq -\varepsilon$ for all interior green faces F of Γ.

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

In [34, Section 7] the authors describe a polynomial-time procedure $\mathbf{RSymVerify}(\mathcal{P}, \varepsilon)$ ([34, Procedure 7.19]) with input $\mathcal{P}$ (the procedure assumes that the preprocessing from [34, Section 7.1] has been done to $\mathcal{P}$, so that it satisfies Assumption 2.3.15) and a constant $\varepsilon > 0$, such that if $\mathbf{RSymVerify}(\mathcal{P}, \varepsilon)$ returns `true`, then $\mathbf{RSym}$ succeeds on $\mathcal{P}$ with ε (see [34, Theorem 7.16]). Hence by [34, Theorem 6.13] G is then hyperbolic, and an explicit linear bound on the Dehn function of G can be calculated (see Definition 2.2.3 and Theorem 2.2.4).

The success of $\mathbf{RSym}$ on $\mathcal{P}$ also implies the following useful result, which is stated only for V^σ-letters, but the proof does not assume it.

Proposition 2.6.12. *[34,* Theorem 6.12*] Assume that $\mathcal{P}$ satisfies Assumption 2.3.15, and that* $\mathbf{RSym}$ *succeeds on $\mathcal{P}$. Then no $x \in X^\sigma$ is trivial in G.*

In [34, Section 8] the following condition on $\mathcal{P}$ is introduced, which enables one to solve the word problem in G in linear time.

Definition 2.6.13. $\mathbf{RSym}$ *verifies a solver* for $\mathcal{P}$ if, for any boundary green face F in any $\Gamma \in \mathcal{D}$ with $\kappa_\Gamma(F) > 0$, the removal of F shortens $\partial(\Gamma)$.

In [34, Section 8] a polynomial-time procedure $\mathbf{VerifySolver}$ (see [34, Procedure 8.3 & Procedure 8.5]) is described, such that if $\mathbf{VerifySolver}(\mathcal{P})$ returns `true`, then $\mathbf{RSym}$ verifies a solver for $\mathcal{P}$ (see [34, Theorem 8.4]). Subsequently, an algorithm $\mathbf{RSymSolve}$ is given (see [34, Algorithm 8.8]), which gives a linear-time solution to the word problem in G if $\mathbf{RSym}$ succeeds on $\mathcal{P}$ and $\mathbf{VerifySolver}(\mathcal{P})$ returns `true` (see [34, Theorem 8.6]). $\mathbf{RSymSolve}$ is a highly technical algorithm, but we shall use it only when the input pregroup presentation satisfies trivial-interleaving (see Definition 2.3.23): in Chapter 6 for development of $\mathbf{IsConjugate}$. Section 6.1 gives the complete description of this simplified version.

We conclude this section with several useful definitions for working with pregroup presentations.

Definition 2.6.14. We say that $\mathcal{P}$ is *sound* if $\mathcal{P}$ satisfies trivial-interleaving; $\mathbf{RSym}$ succeeds on $\mathcal{P}$; and $\mathbf{VerifySolver}(\mathcal{P})$ returns `true`, and that $\mathcal{P}$ is *proper* if no $R \in \mathcal{R}$ is conjugate in $F(X^\sigma)$ to R^{-1}.

We say that $\mathcal{P}$ is *valid* if $\mathcal{P}$ is sound and proper.

Remark 2.6.15. If $\mathcal{P}$ is sound, then G is hyperbolic, and $\mathbf{RSymSolve}$ solves the word problem in G.

Moreover, by [34, Remark 8.10] we can solve the word problem in G by the standard Dehn algorithm using the length reducing rewrite rules derived from $V_P \cup \mathcal{R}$, i.e. $\mathcal{P}$ is a Dehn presentation.

Part I

Conjugacy problem in hyperbolic groups

Brief Outline

Part I of this thesis is structured as follows. In Chapter 3, we define two subsets of coloured annular diagrams (called $\mathcal{T}$ and $\mathcal{S}$) that we shall use for development of $\mathbf{IsConjugate}$.

Definition 2.6.16. Let Γ be a coloured diagram with dual Γ^*, and let $f_1 \subseteq \Gamma$ and $f_2 \subseteq \Gamma$ be faces corresponding to vertices v_1 and v_2 of Γ^*. The *dual distance* from f_1 to f_2 in Γ is the distance from v_1 to v_2 in Γ^*.

In Chapters 4-5 we study the structure of diagrams $\Gamma_A \in \mathcal{T}$, and prove that for any edge e that lies on a boundary ρ of Γ_A, there is an internal face F with $\overline{e} \cap \partial(F) \neq \emptyset$ that is either at dual distance (when treating each red blob of Γ_A as a single face) at most three from the external face with boundary $\rho' \neq \rho$, or $\partial(F) \cap \rho' \neq \emptyset$ (this will enable us to make $\mathbf{IsConjugate}$ quadratic).

In Chapter 6 we describe $\mathbf{IsConjugate}$, and prove Theorem 1.0.3 stated in Chapter 1, which is the main result of Part I. In Chapter 7 we present experiments with our implementation, and show that the reported run times agree with Theorem 1.0.3. Finally, Chapter 8 includes suggestions for improvements and generalizations of $\mathbf{IsConjugate}$.

Chapter 3

Conjugacy diagrams over pregroups

In this chapter we give a procedure for obtaining the subsets $\mathcal{T}$ and $\mathcal{S}$ of coloured annular diagrams. Throughout this whole chapter, let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a finite pregroup presentation for a group G such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.23).

3.1 Preliminaries

In this section we introduce the set $\mathcal{T}$, and determine some elementary properties of coloured annular diagrams. Throughout the whole section let $w_1, w_2 \in X^*$.

In the proof of the next theorem we use the following concept (recall Definition 2.5.14 of the 1-skeleton of a coloured diagram).

Definition 3.1.1. Let Γ_A be a coloured annular diagram with a path $p \in \Gamma_A^1$ intersecting both boundaries of Γ_A. A process of *cutting* Γ_A *open along* p results in creating a coloured simply-connected diagram Γ (see Figure 3.1), where two disjoint copies $p_1, p_2 \subseteq \Gamma$ of p are created, and where each point of $\Gamma_A \setminus p$ is mapped to precisely one point of Γ.

Let Γ_A be an annular diagram, with the external face O with boundary ω. Recall from Definition 2.5.4 that O is oriented counter-clockwise, that all other faces of Γ_A are oriented clockwise, and that all faces of a simply-connected diagram are oriented clockwise.

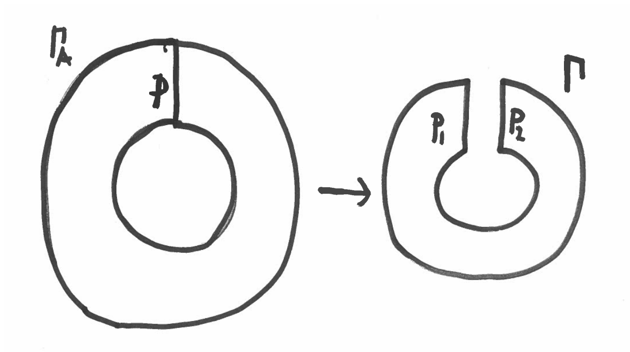

Figure 3.1: Cutting Γ_A open along p, see Definition 3.1.1.

Theorem 3.1.2. *Both of the following statements hold.*

1. *Suppose that w_1 and w_2 are cyclically P-reduced, non-trivial in G, and G-conjugate but not $U(P)$-conjugate. Then w_1 and w_2 are boundary words of some coloured annular diagram over $\mathcal{P}$.*

2. *Suppose that w_1 and w_2 are boundary words of some coloured annular diagram Γ_A over $\mathcal{P}$. Then w_1 and w_2 are G-conjugate; and if w_1 and w_2 are not $U(P)$-conjugate, then $\mathbf{CArea}(\Gamma_A) \geq (1, 0)$.*

In particular, if w_1 and w_2 are cyclically P-reduced, non-trivial in G, and not $U(P)$-conjugate, then w_1 and w_2 are G-conjugate if and only if they are boundary words of some coloured annular diagram over $\mathcal{P}$.

Proof. Part 1. We use the ideas from the proof of [42, Chapter 5, Lemma 5.2]. By the assumptions on w_1 and w_2 and Corollary 2.3.17, there exist $\alpha \in U(P)$, and relators $R_i \in \mathcal{R}^{\pm 1}$ and $u_i \in U(P)$, $1 \leq i \leq k$, such that

$$\alpha^{-1} w_2 \alpha w_1^{-1} =_{U(P)} u_1^{-1} R_1 u_1 \cdots u_k^{-1} R_k u_k,$$

hence $w_1 =_{U(P)} (\Pi_{i=1}^k u_i^{-1} R_i u_i)^{-1} \alpha^{-1} w_2 \alpha$. We proceed as in the proof of Theorem 2.5.8 and construct a balloon diagram Γ' with base point v such that reading counter-clockwise around v, we obtain a cyclic tuple $(p_0, p_1, \ldots, p_k)$ of paths (starting at v) with labels $(\alpha')^{-1}, (u_1')^{-1}$, $\ldots, (u_k')^{-1}$ (where α' is a P-reduced word representing α and each u_i' is a P-reduced word representing u_i) and endpoints v and v_i respectively, and a cyclic tuple of internal faces $(F_0, F_1$, $\ldots, F_k)$, where each F_i is attached to p_i at v_i, F_0 is labelled by w_2, and each $F_i \neq F_0$ is green and labelled by R_i^{-1}. By Lemma 2.5.7 there exists a simply-connected diagram Γ' with P-reduced boundary word w_1' (when reading counter-clokwise around $\partial(\Gamma')$) such that $w_1' =_{U(P)} (\Pi_{i=1}^k (u_i')^{-1} R_i u_i')^{-1} (\alpha')^{-1} w_2 \alpha' =_{U(P)} w_1$, and in which all internal faces other than F_0 are labelled by elements from $V_P \cup \mathcal{R}^{\pm 1}$. By Theorem 2.3.11 w_1' is an interleave of w_1, hence applying Lemma 2.5.7 again shows there exists a simply-connected diagram Γ with boundary word w_1.

If in the construction of Γ the face F_0 was deleted, then by Theorem 2.5.8 $w_1 =_G 1$, a contradiction. Hence $F_0 \subseteq \Gamma$. Now delete F_0 from Γ. The obtained diagram is a coloured annular diagram with the outer boundary labelled by w_1, and the inner boundary labelled by w_2, as required.

Part 2. By relabelling, if necessary, we can without loss of generality assume that w_1 labels the outer boundary of Γ_A. Since Γ_A is connected there exists a simple path p with label some $\alpha \in X^*$ such that cutting Γ_A open along p gives us a coloured van Kampen diagram Γ with boundary word $W = \alpha w_2 \alpha^\sigma w_1^\sigma$. By Theorem 2.5.8 we have $W =_G 1$; and if Γ has no green faces, then $W =_{U(P)} 1$, so we are done.

The last statement follows directly from Parts 1-2. ∎

Definition 3.1.3. Let Γ_A be a coloured annular diagram over $\mathcal{P}$, with boundary words w_1 and w_2. We say that Γ_A has *minimal coloured area* if for every coloured annular diagram Δ_A over $\mathcal{P}$ with boundary words w_1 and w_2 we have $\mathbf{CArea}(\Gamma_A) \leq \mathbf{CArea}(\Delta_A)$.

Similarly, we say that a coloured simply-connected diagram Γ over $\mathcal{P}$ with boundary word w has *minimal coloured area* if $\mathbf{CArea}(\Gamma)$ is minimal among all simply-connected diagrams over $\mathcal{P}$ with boundary word w.

Definition 3.1.4. A *minimal coloured conjugacy diagram* for w_1 and w_2 is a coloured annular diagram with boundary words w_1 and w_2 of minimal coloured area.

Proposition 3.1.5. *If w_1 and w_2 are cyclically P-reduced, non-trivial in G and not $U(P)$-conjugate, then w_1 and w_2 are G-conjugate if and only if there exists a minimal coloured conjugacy diagram for w_1 and w_2.*

Proof. Suppose that w_1 and w_2 are G-conjugate, and let S be the set of all coloured annular diagrams with boundary words w_1 and w_2. By Theorem 3.1.2 S is non-empty, so let $\Gamma_A \in S$, and let $\mathbf{CArea}(\Gamma_A) = (a, b)$. Now using Definition 2.5.9 define an equivalence relation $\sim$ on S as follows, for all $\Gamma'_A, \Delta_A \in S$: $\Gamma'_A \sim \Delta_A$ if and only if $\mathbf{CArea}(\Gamma'_A) = \mathbf{CArea}(\Delta_A)$. Let P be the set of all equivalence classes, and for $0 \leq k \leq a$, let P_k be the set of all equivalence classes containing diagrams with exactly k internal green faces. Note that the coloured area gives us a total order on P, and each set P_k contains a minimal element. Hence as there are only finitely many such elements, one of them is a minimal element of P. Thus, there exists a minimal coloured conjugacy diagram for w_1 and w_2.

Suppose instead that there exists a minimal coloured conjugacy diagram for w_1 and w_2. Then by Theorem 3.1.2 w_1 and w_2 are G-conjugate. ∎

Recall that a loop is an edge with equal endpoints, and Definition 2.5.3 that ω and τ denote the outer and the inner boundary of an annular diagram respectively.

Definition 3.1.6. Let Γ_A be a coloured annular diagram. We say that Γ_A is *loop-minimal* if Γ_A satisfies the following condition.

(*) If l is a loop in Γ_A labelled by a V^σ-letter (see Definition 2.3.8), then $\bar{l} \in \{\omega, \tau\}$.

A *layer* of Γ_A is an annular sub-diagram of Γ_A that does not contain interior loops. A *boundary layer* Γ of Γ_A is a layer of Γ_A such that $\rho \subseteq \partial(\Gamma)$ for some $\rho \in \{\omega, \tau\}$.

The next result will be frequently used in our work (recall Definition 2.6.14 of a valid pregroup presentation).

Lemma 3.1.7. *Let Γ_A be a coloured annular diagram, defined over a valid pregroup presentation. Suppose that Γ_A contains a loop l, and let C be the bounded component of $\mathbb{R}^2 \setminus \bar{l}$. Then $\overline{C} \cap \Gamma_A$ is an annular sub-diagram of Γ_A. Hence Γ_A is a face-disjoint union of layers $\bigcup_{j=1}^{n} \Gamma_j$, and there is i such that $\bar{l}$ is one of the boundaries of Γ_i.*

Proof. Suppose for a contradiction that $\overline{C} \cap \Gamma_A$ is not annular. Then $\overline{C} \cap \Gamma_A$ is simply-connected. Hence the label of l is a single letter trivial in G, which by Proposition 2.6.12 contradicts our assumption that Γ_A is defined over a valid pregroup presentation. ∎

We now define $\mathcal{T}$.

Definition 3.1.8. Assume that $\mathcal{P}$ is valid. We define $\mathcal{T}$ to be the set of all coloured annular diagrams Γ_A over $\mathcal{P}$ that satisfy all of the following axioms.

(T_1) The boundaries of Γ_A are simple closed paths, and no internal green face has more than $1/2$ of its length on a single boundary of Γ_A.

(T_2) Γ_A is σ-reduced and semi-P-reduced (see Definitions 2.5.15 and 2.5.16).

(T_3) Γ_A is green-rich (see Definition 2.5.6).

(T_4) No proper sub-word of the (cyclic) boundary word of a simply-connected red blob of Γ_A is equal to 1 in $U(P)$.

(T_5) Γ_A is loop-minimal (see Definition 3.1.6).

(T_6) Each internal green face F of Γ_A contains a boundary edge and $\kappa_{\Gamma_A}(F) = 0$.

Recall Definition 2.6.8 that $\chi(x, \Gamma_A)$ is the curvature that x gives to a single internal green face of Γ_A across each curvature incidence.

Lemma 3.1.9. *Let $\Gamma_A \in \mathcal{T}$ have a vertex v incident with an internal green face. If $v \in \omega \cap \tau$, then $\chi(v, \Gamma_A) = -1/2$, else $\chi(v, \Gamma_A) > -1/2$.*

Proof. By Axiom T_1, the boundaries of Γ_A are simple closed paths. Hence if $v \in \omega \cap \tau$, then v is incident exactly twice with external faces of Γ_A, else v is incident exactly once with them. So the lemma follows from Lemma 2.6.9. ∎

Recall Definition 2.5.10 of an island and a bridge.

Lemma 3.1.10. *Let $\Gamma_A \in \mathcal{T}$ contain an island. Then Γ_A is a union of islands and bridges. Suppose further that Γ_A contains at least two islands, and let Δ_A be a diagram resulted from deleting some island of Γ_A and identifying its endpoints. Then $\Delta_A \in \mathcal{T}$, and Δ_A is a union of islands and bridges.*

Proof. By Axiom T_1, the boundaries of Γ_A are simple closed paths, hence Γ_A is a union of islands and bridges.

To prove the second statement of the lemma, first note that Δ_A is also a union of islands and bridges. Hence Δ_A is annular, the boundaries ω_1 and τ_1 of Δ_A are simple closed paths, and each internal green face and each red blob of Δ_A is contained in some island of Γ_A. Therefore, no internal green face has more than $1/2$ of its length on a single boundary of Δ_A, so Δ_A satisfies Axiom T_1; Δ_A is σ-reduced and semi-P-reduced (Axiom T_2); each vertex $v \notin \omega_1 \cap \tau_1$ satisfies $\delta_G(v, \Delta_A) = \delta_G(v, \Gamma_A) \geq 2$ and each vertex $v \in \omega_1 \cap \tau_1$ also satisfies $\delta_G(v, \Delta_A) \geq 2$, hence Δ_A satisfies Axiom T_3; Δ_A satisfies Axiom T_4; and Δ_A is loop-minimal. Thus, Δ_A satisfies Axioms T_1-T_5. Moreover, as Γ_A satisfies Axiom T_6, each internal green face F of Δ_A contains a boundary edge.

So it remains to show that $\kappa_{\Delta_A}(F) = 0$. Assume that B is a red blob of Δ_A edge-incident with F. Then $\mathbf{Area}(B)$ and the number of edge-incidences with F remained unchanged, and $|\partial(B) \setminus \partial(\Delta_A)| = |\partial(B) \setminus \partial(\Gamma_A)|$, so by the description of **ComputeRSym** (see Algorithm 2.6.4) we have $\chi(B, F, \Delta_A) = \chi(B, F, \Gamma_A)$. Suppose that v is a vertex of Δ_A incident with F. If $v \in \omega_1 \cap \tau_1$, then since ω_1 and τ_1 are simple closed paths, v is incident exactly twice with external faces of Δ_A, hence by Lemma 2.6.9 $\chi(v, \Delta_A) = -1/2$. Suppose further that v resulted from identifying two endpoints v_1 and v_2 of some islands of Γ_A. Since $v_1, v_2 \in \omega \cap \tau$, we have $\chi(v_1, \Gamma_A) = -1/2 = \chi(v_2, \Gamma_A)$. Moreover, the sum of the numbers of incidences of v_1 and v_2 with F in Γ_A is equal to the number of incidences of v with F in Δ_A. Hence $\chi(v, F, \Delta_A) = \chi(v_1, F, \Gamma_A) + \chi(v_2, F, \Gamma_A)$. Otherwise, we have $v \in \omega \cap \tau$, so $\chi(v, \Gamma_A) = -1/2$, and the number of incidences with F remained unchanged, hence $\chi(v, F, \Delta_A) = \chi(v, F, \Gamma_A)$. If $v \notin \omega_1 \cap \tau_1$, then the green degree of v, the internal green degree of v (see Definition 2.5.6), and the number of incidences with F remained all unchanged, so $\chi(v, F, \Delta_A) = \chi(v, F, \Gamma_A)$. Thus, as by Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$, it follows that $\kappa_{\Delta_A}(F) = 0$, hence $\Delta_A \in \mathcal{T}$. ■

Using the techniques from [34], we now prove several auxiliary results that will help us to show that minimal conjugacy diagrams share many properties with diagrams in $\mathcal{T}$.

Lemma 3.1.11. *Let Γ_A be a coloured annular diagram with cyclically σ-reduced boundary words. Then there are no vertices of degree 1 in Γ_A.*

Proof. The proof is essentially the same as the proof of [34, Lemma 3.12]. If v is a vertex of Γ_A with degree 1, then xx^σ is a sub-word of some label w of the face containing v, where x labels the unique edge incident with v. This is a contradiction since w is cyclically σ-reduced. ■

In the proof of the next lemma we use the ideas from the proof of [34, Proposition 3.8].

Lemma 3.1.12. *Let Γ_A be a minimal coloured conjugacy diagram. Then Γ_A is semi-P-reduced (hence also semi-σ-reduced).*

Proof. Suppose not. Since Γ_A is not semi-P-reduced, by Definition 2.5.16 there are internal green faces labelled by $w_1 w_2$ and $w_3^{-1} w_1^{-1}$ that have a common consolidated edge labelled by w_1 and w_1^{-1}, where $w_2 w_3^{-1} =_{U(P)} 1$. Therefore, we can delete the consolidated edge labelled w_1 from Γ_A and identify consecutive edges with inverse labels. The resulted region has a cyclically σ-reduced boundary word equal to 1 in $U(P)$, so by Theorem 2.5.8 we can fill in it with red triangles. The obtained diagram has the same boundary words as Γ_A, but strictly smaller coloured area, a contradiction. ■

Definition 3.1.13. Let Γ be a coloured annular/simply-connected diagram over $\mathcal{P}$. A *smaller sibling of* Γ is a coloured annular/simply-connected diagram Δ over $\mathcal{P}$ with the same boundary words/word as Γ, with the same green faces as Γ, and satisfying $\mathbf{CArea}(\Delta) \leq \mathbf{CArea}(\Gamma)$.

The next result is stated only for coloured simply-connected diagrams, but its proof does not assume it.

Lemma 3.1.14. *[34,* Lemma 3.15*] Let Γ_A be a coloured annular or simply-connected diagram. Assume that Γ contains a vertex v with three consecutive edge-incident 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 smaller sibling of Γ 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.*

Recall Definition 2.3.9 of a (cyclically) σ/P-reduced word.

Definition 3.1.15. A word $w \in X^*$ is *$\mathcal{R}$-reduced* if w does not contain a sub-word s such that there exists a cyclic conjugate R of some $R' \in \mathcal{R}^{\pm 1}$ that can be written as $R = usv$ for some $u, v \in X^*$ and $|s| > |R|/2$. We define *cyclically $\mathcal{R}$-reduced* similarly.

A word w is (*cyclically*) *$\mathcal{P}$-reduced* if w is both (*cyclically*) *P-reduced* and (*cyclically*) *$\mathcal{R}$-reduced*.

Lemma 3.1.16. *Let Γ_A be a coloured annular diagram with a cyclically $\mathcal{P}$-reduced boundary word w that labels $\rho \in \{\omega, \tau\}$. Then no internal green face of Γ_A has more than $1/2$ of its length on ρ.*

Proof. This is immediate from Definition 3.1.15 since w is cyclically R-reduced. ∎

3.2 Loop-free minimal conjugacy diagrams

This section studies minimal coloured conjugacy diagrams (see Definition 3.1.4) that contain no loops labelled by V^σ-letters (see Definition 2.3.8).

Definition 3.2.1. Let Γ_A be a coloured annular diagram over $\mathcal{P}$. We say that Γ_A is *loop-free* if Γ_A contains no loops labelled by V^σ-letters.

The following theorem is the main result of this section (recall Definition 2.6.14 of a valid pregroup presentation and Definition 3.1.8 of the set $\mathcal{T}$).

Theorem 3.2.2. *Assume that $\mathcal{P}$ is valid, and let Γ_A be a minimal coloured conjugacy diagram with cyclically $\mathcal{P}$-reduced boundary words. If Γ_A is loop-free, then $\Gamma_A \in \mathcal{T}$.*

Remark 3.2.3. Let $w_1, w_2 \in X^*$ be cyclically $\mathcal{P}$-reduced and G-conjugate. Assume that $\mathcal{P}$ is valid, and that no cyclically $\mathcal{P}$-reduced $w \in X^*$ with $|w| \geq 2$ is G-conjugate to a V^σ-letter. Then by Theorem 3.2.2, unless $|w_1| = 1 = |w_2|$, all minimal coloured conjugacy diagrams Γ_A for w_1 and w_2 are guaranteed to satisfy $\Gamma_A \in \mathcal{T}$. Hence to solve the conjugacy problem for G, it suffices to solve three problems.

1. Finding conjugacy classes of single letters.

2. Testing if any cyclically $\mathcal{P}$-reduced word w with $|w| \geq 2$ is G-conjugate to a V^σ-letter.
3. For cyclically $\mathcal{P}$-reduced words w_1' and w_2', testing if there exists a diagram $\Gamma_A \in \mathcal{T}$ with boundary words w_1' and w_2'.

Recall that a consolidated edge between faces F and F' is a non-empty path of maximal length that is a sub-path of both $\partial(F)$ and $\partial(F')$. In the statement of the next lemma we treat the red blob B as a single face.

Lemma 3.2.4. *Let Γ be a green-rich coloured diagram over $\mathcal{P}$ that contains a red blob B. Then all consolidated edges between B and any green face of Γ_A have length at most one.*

Proof. Assume for a contradiction that there is a consolidated edge l with $|l| \geq 2$ between B and some green face of Γ. Then l contains a vertex v with $\delta_G(v) = 1$, a contradiction as Γ is green-rich. ∎

Recall the notation $\delta_G(v, \Gamma)$ from Definition 2.5.6. The next proof is inspired by the proof of [34, Theorem 3.16].

Lemma 3.2.5. *Assume that $\mathcal{P}$ is valid, and let Γ_A be a minimal conjugacy diagram over $\mathcal{P}$ with cyclically $\mathcal{P}$-reduced boundary words. Assume that Γ_A contains a vertex v that is not incident with any loop labelled by a V^σ-letter. Then $\delta_G(v, \Gamma_A) \geq 2$.*

Hence if Γ_A is loop-free, then is is green-rich.

Proof. We first show that $\delta_G(v, \Gamma_A) \geq 1$. Assume for a contradiction that $\delta_G(v, \Gamma_A) = 0$. By Lemma 3.1.11 we have $\delta(v, \Gamma_A) \geq 2$. Suppose first that $\delta(v, \Gamma_A) = 2$. Then v is incident with two red triangles edge-incident by two edges meeting at v, and labelled $ab[ab]^\sigma$ and $b^\sigma a^\sigma [ab]$ respectively for some $a, b \in X$. Hence Γ_A is not semi-σ-reduced, contradicting Lemma 3.1.12. So suppose that $\delta(v, \Gamma_A) \geq 3$. Then by Lemma 3.1.14 there exists a smaller sibling of Γ_A in which v is incident with at least one fewer red triangle than it is in Γ_A, and in which none of the edges of any of the red triangles incident with v is a loop based at v. By repeating this process we obtain a smaller sibling Δ_A of Γ_A in which $\delta(v, \Delta_A) = 2$, and in which none of the edges of any of the red triangles incident with v is a loop based at v. So Δ_A is a minimal conjugacy diagram, and as before Δ_A is not semi-σ-reduced, contradicting Lemma 3.1.12.

It remains to show that $\delta_G(v, \Gamma_A) \neq 1$. Suppose otherwise. Let F be the unique green face incident with v, and let e and f be the edges of F with $v \in \overline{e} \cap \overline{f}$, and with labels a and b respectively. Since there is no loop labelled by a V^σ-letter based at v, from $\delta(v, \Gamma_A) \geq 2$ we have $e \neq f$. Hence ef is a sub-path of $\partial(F)$, so ab is a sub-word of some label S of F, and e and f are edges of some (not necessarily distinct) red faces F_1 and F_2 of Γ_A respectively. If $F_1 = F_2$, then $(a, b) \in D(P)$, so S is not cyclically P-reduced, a contradiction. Hence suppose that $F_1 \neq F_2$. Then since F is the only green face incident with v and F is incident once with v, we deduce that ef is a sub-path of a consolidated edge between F and the red blob B with $v \in B$, so by Lemma 2.5.13 $(a, b) \in D(P)$, a contradiction.

The last statement follows from the first. ∎

The next result is again stated only for coloured simply-connected diagrams, but the proof does not assume it.

Lemma 3.2.6. *[34,* Lemma 4.12*] Let B be a red blob in a coloured diagram Γ over $\mathcal{P}$ with boundary length l and area t. Then $l \le t+2$, and $l \le t$ if B is not simply-connected. Moreover, if B is simply-connected, and every vertex of B lies on $\partial(B)$ (which holds in particular if all vertices of Γ have green degree at least 1), then $l = t+2$.*

Using Lemma 3.2.6 we obtain:

Lemma 3.2.7. *Let Γ be a coloured annular or simply-connected diagram over $\mathcal{P}$ of minimal coloured area and in which all vertices have green degree at least 1. Then no proper sub-word of the (cyclic) boundary word of a simply-connected red blob of Γ is equal to 1 in $U(P)$.*

Hence the same holds for a loop-free minimal conjugacy diagram defined over a valid pregroup presentation $\mathcal{P}'$ that has cyclically $\mathcal{P}'$-reduced boundary words.

Proof. We use the ideas from the proof of [34, Proposition 4.13]. Assume that there is a simply-connected red blob B of Γ that does not have the stated properties. Since all vertices of Γ have green degree at least 1, we deduce that all vertices of B lie on $\partial(B)$. Suppose first that B has a boundary word w such that $w = w'xx^\sigma$. Since B is simply-connected, we have $w =_{U(P)} 1$, so $w' =_{U(P)} 1$. Hence we can identify the vertex at beginning of the edge with label x with the vertex at the end of the edge labelled x^σ, and replace B with a red blob B' with boundary label w' and with a single edge added to the boundary. By Lemma 3.2.6 $\mathbf{Area}(B') = |\partial(B')| - 2 < |\partial(B)| - 2 = \mathbf{Area}(B)$. Hence the diagram with the red blob B' is a smaller sibling of Γ, and with a strictly smaller coloured area, a contradiction.

Now suppose that B has a boundary word w such that $w = w_1w_2$ with $|w_1|, |w_2| > 2$ and $w_1 =_{U(P)} 1 =_{U(P)} w_2$. Then we can identify the vertices at the beginning and the end of w_1, and replace B by two blobs B_1 and B_2 with boundary words w_1 and w_2, and with

$$\mathbf{Area}(B_1) + \mathbf{Area}(B_2) = |\partial(B_1)| - 2 + |\partial(B_2)| - 2 = |\partial(B)| - 4 = \mathbf{Area}(B) - 2.$$

So we again obtained a smaller sibling of Γ with a strictly smaller coloured area.

Now note that the final statement follows from the first statement and Lemma 3.2.5. ∎

In the next result we derive the formula for calculating $\chi(B, \Gamma_A)$ (see Definition 2.6.8) of a simply-connected red blob B.

Lemma 3.2.8. *Let B be a red blob composed of t triangles in a green-rich coloured annular diagram Γ_A. If B is not simply-connected, then $\chi(B, \Gamma_A) \le -1/2$.*

Let $d = |\partial(B) \cap \partial(\Gamma_A)|$. If B is simply-connected then

$$\chi(B, \Gamma_A) = \frac{-t}{2(t-d)+4}. \tag{3.1}$$

In particular, for small values of t, the curvature values are as in Table 3.1.

Proof. First note that from the description of **ComputeRSym** (see Algorithm 2.6.4) it follows that

$$\chi(B, \Gamma_A) = \frac{-t}{2|\partial(B) \setminus \partial(\Gamma_A)|}. \tag{3.2}$$

Let $l = |\partial(B)|$. By Lemma 3.2.6 $l \leq t$ if B is not simply-connected, hence by (3.2) $\chi(B, \Gamma_A) \leq -1/2$.

Suppose that B is simply-connected. Since Γ_A is green-rich, by Lemma 3.2.6 $l = t + 2$. Thus, from (3.2) we deduce that (3.1) holds for $\chi(B, \Gamma_A)$. ∎

Table 3.1: Red blob curvature

$\mathbf{Area}(B)$	$\partial(B) \cap \partial(\Gamma_A)$	$\chi(B, \Gamma_A)$
≥ 1	≥ 2	$\leq -1/2$
1	0	$-1/6$
1	1	$-1/4$
2	0	$-1/4$
2	1	$-1/3$
3	0	$-3/10$
3	1	$-3/8$
4	0	$-1/3$
4	1	$-2/5$
5	0	$-5/14$
5	1	$-5/12$

Proposition 3.2.9. *Assume that $\mathcal{P}$ is valid. Let Γ_A be a green-rich coloured annular diagram over $\mathcal{P}$ that has a cyclically $\mathcal{P}$-reduced boundary word w that labels $\rho \in \{\omega, \tau\}$. Then ρ is a simple closed path. In particular, if E is the external face with boundary ρ, then no boundary vertex v of Γ_A is incident more than once with E.*

Proof. Assume for a contradiction that ρ is not a simple closed path. Then either ρ contains a vertex with degree 1, or it contains a closed path l such that $l \neq \rho$. By green-richness of Γ_A we deduce that the latter holds. Now l encloses a coloured simply-connected sub-diagram of Γ_A with boundary word u, such that u is a contiguous sub-word of a cyclic conjugate of w. By Remark 2.6.15 $\mathcal{P}$ is a Dehn presentation. Hence as w is cyclically P-reduced, u is not $\mathcal{R}$-reduced, so w is not cyclically $\mathcal{P}$-reduced, a contradiction. ∎

Recall Definition 2.5.10 of an island.

Lemma 3.2.10. *Let E be an island of a green-rich annular diagram Γ_A with endpoints v_1 and v_2, and bounded by the closed path $\omega_1\tau_1$. Then* $\mathbf{CArea}(E) \geq (1, 0)$.

Hence if Γ_A is a green-rich annular diagram whose boundaries are simple closed paths, then $\kappa_{\Gamma_A}(T) = 0$ for each red triangle T.

Proof. Suppose that E has no green faces. Then E is a simply-connected red blob. Since Γ_A is green-rich, by Lemma 3.2.4 all consolidated edges between E and any external face of Γ_A

have length at most 1. Hence as by Definition 2.5.10 $\omega_1 \neq \tau_1$, we have $|\partial(E)| = 2$, so by Lemma 3.2.6 $\mathbf{Area}(E) = 0$, a contradiction.

For the final statement, suppose that $B \subseteq \Gamma_A$ is a red blob that is not edge-incident with any internal green face of Γ_A. Then as Γ_A is green-rich and its boundaries are simple closed paths, B is an island of Γ_A, contradicting the lemma. So $\kappa_{\Gamma_A}(T) = 0$ for each triangle T. ■

Conjugacy in $U(P)$ was studied in [17], where the following was proved.

Theorem 3.2.11. *Let $w_1 = x_1x_2 \ldots x_n, w_2 \in X^*$ be cyclically P-reduced and $U(P)$-conjugate. If $n = 1$, then there exists $c \in P$ such that $w_2 = [c^\sigma wc]$, where $(c^\sigma w), (w, c), ([c^\sigma w], c) \in D(P)$.*

Otherwise, for some i, we have

$$w_2 = [c^\sigma x_i]x_{i+1} \ldots x_n x_1 \ldots [x_{i-1}c] \in U(P),$$

where $c, [c^\sigma x_i], [x_{i-1}c] \in P$.

Proof. The case $n = 1$ is [17, Corollary 4.5]. The rest is [17, Theorem 4.6]. ■

For the case $n \geq 2$, we can improve this result when $\mathcal{P}$ satisfies trivial-interleaving (see Definition 2.3.23):

Corollary 3.2.12. *Assume that $\mathcal{P}$ satisfies trivial-interleaving, and that $w_1 = x_1x_2 \ldots x_n, w_2 \in X^*$ are cyclically P-reduced and $U(P)$-conjugate. If $|w_1| \geq 2$, then for some i, we have*

$$w_2 = x_i x_{i+1} \ldots x_n x_1 \ldots x_{i-1}.$$

Proof. By Theorem 3.2.11 we have $w_2 \in \mathcal{I}(x_i \ldots x_n x_1 \ldots x_{i-1})$ (see Definition 2.3.20) for some i. Hence as $\mathcal{P}$ satisfies trivial-interleaving, the corollary follows. ■

Proposition 3.2.13. *Assume that $\mathcal{P}$ is valid. Then all coloured semi-σ-reduced annular diagrams over $\mathcal{P}$ are σ-reduced. In addition, we can test whether $\mathcal{P}$ is proper in time $O(r|\mathcal{R}|)$, where $r := \max\{|R| : R \in \mathcal{R}\}$.*

Proof. Let Γ_A be a coloured annular diagram over $\mathcal{P}$ and suppose that f is a face edge-incident with itself in such a way that Γ_A is not σ-reduced. First note that the only way that a red triangle could be edge-incident with itself is when some element of X is trivial in $U(P)$, which contradicts Theorem 2.3.11. Hence f is green. Now f is edge-incident with itself by a consolidated edge l with label w_1, say. By our assumption reading $\partial(f)$ from the side of l gives w_1w_2, and from the other side gives $w =_{F(X^\sigma)} w_1^{-1}w_2^{-1}$ (see Figure 2.4). Hence w is a cyclic conjugate of w_1w_2 and of $w_2^{-1}w_1^{-1} =_{F(X^\sigma)} (w_1w_2)^{-1}$. So w_1w_2 is an $F(X^\sigma)$-conjugate of its inverse in $F(X^\sigma)$, and hence R contradicts our assumption that $\mathcal{P}$ is proper (since $\mathcal{P}$ is valid).

Now if $R \in \mathcal{R}$ and R^{-1} are conjugate in $F(X^\sigma)$, then they are cyclic conjugates as they are cyclically σ-reduced. So R is a contiguous sub-word of $(R^{-1})^2$, and by [2, Section 9.1] we can test this in time $O(r)$ by the Knuth-Morris-Pratt (KMP) string-searching algorithm. The final complexity claim follows from this. ■

Lemma 3.2.14. *Let Γ_A be a green-rich coloured annular diagram of area greater than 1 that contains a boundary green face F. If $\kappa_{\Gamma_A}(F) > 0$ then the following statements hold.*

1. *The consolidated edges and vertices in $\overline{\partial(F) \setminus \partial(\Gamma_A)}$ form a single path p, and at most three of the vertices in p lie on $\partial(\Gamma_A)$. If there are three such vertices, let v be the middle one. Then $\delta_G(v, \Gamma_A) \geq 4$, and F is incident with no red blobs at v.*

2. *F does not have an edge on both boundaries of Γ_A.*

Proof. Since Γ_A is green-rich, Lemmas 2.6.9 and 3.2.8 hold, which are analogues of [34, Lemma 6.7 & 6.8] that are applicable to coloured annular diagrams. Hence the proof of [34, Lemma 6.9] shows Part 1.

To prove Part 2, assume for a contradiction that F contains an edge on both boundaries of Γ_A. Let e and f be consolidated edges of F with $|e|, |f| \geq 1$ and $e \subseteq \omega$, $f \subseteq \tau$. By Part 1 we have $\partial(F) \cap \partial(\Gamma_A) = e \cup f$, and $\partial(e) \cap \partial(f)$ contains a vertex v with $v \in \omega \cap \tau$. By Lemma 2.6.10 we have $\chi(v, F, \Gamma_A) \leq -1/2$.

Let v and v_1 be the endpoints of e, and let v and v_2 be the endpoints of f. Let $x \in \{v_i\}_{i=1}^2$. If $\delta_G(x) \geq 3$, then by Lemma 2.6.10 $\chi(x, F, \Gamma_A) \leq -1/4$. If $\delta_G(x) = 2$, then there is a red blob B edge-incident with F at x, and with an edge $g \subseteq \partial(\Gamma_A)$ such that $x \in \overline{g}$. Hence by Lemma 3.2.8 we have $\chi(B, F, \Gamma_A) \leq -1/4$. We deduce that for each $1 \leq i \leq 2$, there is x_i such that either $x_i = v_i$ and $\chi(x_i, F, \Gamma_A) \leq -1/4$, or x_i is a red blob with edges $g \subseteq \partial(F)$, $h \subseteq \partial(\Gamma_A)$ such that $v_i \in \overline{g} \cap \overline{h}$, and $\chi(x_i, F, \Gamma_A) \leq -1/4$. Hence if $x_1 \neq x_2$ and $x_1 \neq v \neq x_2$, then

$$\kappa_{\Gamma_A}(F) \leq 1 + 2 \cdot (-1/4) - 1/2 = 0,$$

a contradiction. Otherwise, at least one of the following cases holds.

(a) $x_i = v$ for some $i \in \{1, 2\}$ and v is incident at least twice with F.

(b) $x_1 = x_2$ and x_1 is a vertex with $x_1 \in \omega \cap \tau$.

(c) $x_1 = x_2$ and x_1 is a red blob with edges on both boundaries of Γ_A.

Assume first that Case (a) holds. Then $\chi(v, F, \Gamma_A) \leq -1$, so $\kappa_{\Gamma_A}(F) \leq 0$, a contradiction. Assume that Case (b) holds instead. Then by Lemma 2.6.10 $\chi(x_1, F, \Gamma_A) \leq -1/2$, so $\kappa_{\Gamma_A}(F) \leq 1 + 2 \cdot (-1/2) = 0$, a contradiction. Finally, assume that Case (c) holds. Then by Lemma 3.2.8 we have $\chi(x_1, F, \Gamma_A) \leq -1/2$, so again $\kappa_{\Gamma_A}(F) \leq 0$. ∎

Proof of Theorem 3.2.2. Let w_1 and w_2 be the boundary words of Γ_A. By Lemma 3.2.5 Γ_A satisfies Axiom T_3. Hence as w_1 and w_2 are cyclically $\mathcal{P}$-reduced, by Lemma 3.1.16 and Proposition 3.2.9 Γ_A satisfies Axiom T_1. Furthermore, by Lemma 3.1.12 and Proposition 3.2.13 Γ_A satisfies Axiom T_2; by Lemma 3.2.7 Γ_A satisfies Axioms T_4; and as Γ_A is loop-free, Γ_A satisfies Axiom T_5.

So it remains to show that Γ_A satisfies Axiom T_6. Hence we can assume that Γ_A contains at least one green face, as otherwise there is nothing to prove. Since $\mathcal{P}$ is valid, there exists $\varepsilon > 0$

such that $\mathbf{RSym}(\mathcal{P}, \varepsilon)$ succeeds. Now if there is an interior green face F with $\kappa_{\Gamma_A}(F) > -\epsilon$, then there is a decomposition of the label of F into steps representing red blobs edge-incident with F and vertices incident with F that makes $\mathbf{RSymVerify}(\mathcal{P}, \varepsilon)$ fail, a contradiction. Hence $\kappa_{\Gamma_A}(F) \leq -\epsilon$.

We have $\kappa_{\Gamma_A}(e) = 0$ for each edge $e \in \Gamma_A$. By Axiom T_1, ω and τ are simple closed paths, hence each vertex $v \in \Gamma_A$ is incident at most twice with external faces. By Lemma 2.6.9 we therefore have $\kappa_{\Gamma_A}(v) = 0$. Moreover, by Axioms T_1 and T_3, Γ_A satisfies assumptions of Lemma 3.2.10, hence by Lemma 3.2.10 $\kappa_{\Gamma_A}(T) = 0$ for each red triangle T. Thus, as $\kappa(\Gamma_A) = 0$ and all non-zero curvature is on the internal green faces of Γ_A, if Γ_A contains precisely one green face F, then $\kappa_{\Gamma_A}(F) = 0$. So from above F contains a boundary edge. Hence we can assume that Γ_A contains at least two green faces.

Suppose that F is a boundary green face labelled by R and $\kappa_{\Gamma_A}(F) > 0$. By Lemma 3.2.14 either $\partial(F) \cap \partial(\Gamma_A) = l$, where l is a single consolidated edge with $l \subseteq \rho \in \{\omega, \tau\}$, or $\partial(F) \cap \partial(\Gamma_A) = l \cup \{v\}$, where l has properties as before and v is a vertex. Hence by Axiom T_1, F does not have more than $1/2$ of its length on $\partial(\Gamma_A)$. Therefore, we can find a decomposition of R into steps representing red blobs edge-incident with F and vertices incident with F that makes $\mathbf{VerifySolver}(\mathcal{P})$ return `fail`, a contradiction since $\mathcal{P}$ is valid. Thus, as $\kappa(\Gamma_A) = 0$ and all interior green faces F have $\kappa_{\Gamma_A}(F) \leq -\epsilon$, we conclude that Γ_A has no interior green faces, and all boundary green faces F satisfy $\kappa_{\Gamma_A}(F) = 0$. ∎

3.3 Conjugacy diagrams containing loops labelled by V^σ-letters

This subsection analyses the case where a minimal coloured conjugacy diagram is not loop-free (see Definitions 3.1.4 and 3.2.1). Recall Definition 3.1.6 of a (boundary) layer. Throughout this section, assume that $\mathcal{P}$ is valid, and let $w_1, w_2 \in X^*$ be cyclically $\mathcal{P}$-reduced and G-conjugate.

Definition 3.3.1. Let Γ be a coloured annular diagram over $\mathcal{P}$. We say that Γ is *decomposable* if Γ is a face-disjoint union of non-empty annular sub-diagrams $\Gamma_1 \cup \Delta_1$ or $\Gamma_1 \cup \Delta_1 \cup \Delta_2$, where $\mathbf{CArea}(\Delta_1) = (0, 3) = \mathbf{CArea}(\Delta_2)$, and the length 1 boundary of Δ_i is a boundary of Γ.

We call Δ_i a *boundary red blob of* Γ, and Γ_1 the *core of* Γ.

Definition 3.3.2. We define $\mathcal{S}$ to be the set of all coloured annular diagrams Γ_A over $\mathcal{P}$ each of which is a face-disjoint union of layers Γ with area at least 1 and such that: for non-boundary layers Γ, both boundaries of Γ have length 1, and one of the following 3 statements holds for Γ.

1. The boundary words of Γ are single letters G-conjugate by some single letter.

2. $\Gamma \in \mathcal{T}$.

3. Γ is decomposable, where

    (i) the core Γ_1 of Γ satisfies $\Gamma_1 \in \mathcal{T}$;

(ii) the label of the length 2 boundary of each boundary red blob B of Γ is non-trivial in P, is not P-reduced, and is not equal to the other boundary word of B in P-reduced form.

For boundary layers Γ the same holds, but at least one of the boundaries of Γ is a boundary of Γ_A, so may have length greater than 1.

Our aim is to show the following theorem (recall Definition 3.1.13 of a smaller sibling).

Theorem 3.3.3. *Assume that w_1 and w_2 are non-trivial in G and are not $U(P)$-conjugate, and that Γ_A^m is a minimal conjugacy diagram for w_1 and w_2. Then there exists a smaller sibling Γ_A of Γ_A^m (hence Γ_A is a minimal conjugacy diagram for w_1 and w_2) such that $\Gamma_A \in \mathcal{S}$. Moreover, if Γ_A^m is a single layer, then Γ_A is a single layer.*

Assume that w_1 and w_2 are non-trivial in G and are not $U(P)$-conjugate. Since w_1 and w_2 are cyclically $\mathcal{P}$-reduced and G-conjugate, by Proposition 3.1.5 there exists a minimal conjugacy diagram Γ_A^m for w_1 and w_2. Assume that Γ_A^m contains a loop l. Then by Lemma 3.1.7 Γ_A^m is a face-disjoint union of annular sub-diagrams Δ_1 and Δ_2, where Δ_1 is bounded by ω and $\bar{l}$, and Δ_2 is bounded by $\bar{l}$ and τ. Note that Δ_1 and Δ_2 can have area 0.

The next lemma considers the degrees of the endpoint v of l. Recall that $\delta_G(v', \Gamma)$ (and $\delta_R(v', \Gamma)$) is the number of green (and red) faces incident with a vertex v' in a coloured diagram Γ, and that $\delta_G^I(v', \Gamma)$ is the number of internal green faces incident with v' in Γ.

Lemma 3.3.4. *For $i \in \{1, 2\}$, one of the following statements holds.*

1. *We have $\delta_G(v, \Delta_i) \geq 2$.*

2. *We have $\delta_G(v, \Delta_i) = 1$ and $\delta_R(v, \Delta_i) \geq 3$.*

Proof. Since v is a boundary vertex of Δ_i, we have $\delta_G(v, \Delta_i) \geq 1$. Assume that $\delta_G(v, \Delta_i) = 1$. Then $\delta_G^I(v, \Delta_i) = 0$ and $\delta_R(v, \Delta_i) \geq 1$. In fact, the unique triangle of Δ_i containing l is incident more than once with v, hence $\delta_R(v, \Delta_i) \geq 2$. Assume that $\delta_R(v, \Delta_i) = 2$. Then as by Lemma 3.1.11 each vertex $v' \in \Gamma_A^m$ has $\delta(v', \Gamma_A^m) \geq 2$, it follows that v is incident with both external faces of Δ_i, contradicting $\delta_G(v, \Delta_i) = 1$, so $\delta_R(v, \Delta_i) \geq 3$. ∎

Theorem 3.3.5. *Assume that w_1 and w_2 are non-trivial in G and are not $U(P)$-conjugate, and that Γ_A^m contains a loop l. Then Γ_A^m is a face-disjoint union of n layers for some n, and there exists a smaller sibling Γ_A of Γ_A^m such that Γ_A is a face-disjoint union of layers $\bigcup_{j=1}^{n} \Gamma_j$ with area at least 1. Furthermore, for some i, the closure $\bar{l}$ is one of the boundaries of Γ_i, and Γ_i satisfies one of the following 3 statements.*

1. *The boundary words of Γ_i are single letters that are G-conjugate by some single letter.*

2. *Γ_i satisfies Axioms $T_1 - T_5$.*

3. *Γ_i is decomposable, where*

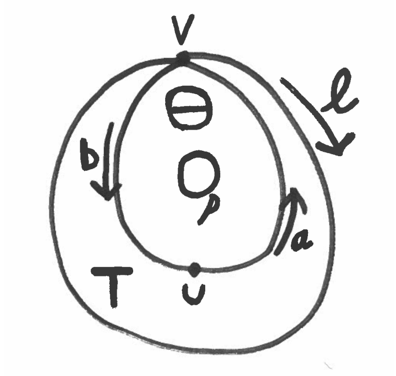

Figure 3.2: The sub-diagram Θ of Δ_i, see the proof of Theorem 3.3.5.

(i) the core of Γ_i satisfies Axioms $T_2 - T_5$;

(ii) the label of the length 2 boundary of each boundary red blob B of Γ_i is non-trivial in P, is not P-reduced, and is not equal to the other boundary word of B in P-reduced form.

Proof. By Lemma 3.1.7 Γ_A^m is a face-disjoint union of layers $\bigcup_{j=1}^n \Delta_j$, and for some i, the closure $\bar{l}$ is one of the boundaries of Δ_i. Since w_1 and w_2 are not $U(P)$-conjugate, by Theorem 3.1.2 $\mathbf{CArea}(\Gamma_A^m) \geq (1,0)$. Hence for all $1 \leq j \leq n$ we can assume $\mathbf{Area}(\Delta_j) \geq 1$. Without loss of generality assume that Δ_i is contained in the annular sub-diagram with boundaries $\bar{l}$ and ω. Let v be the endpoint of l, let $\partial(\Delta_i) = \{l, \rho\}$, and let t and w_3 be the labels of l and ρ respectively when oriented by external faces of Δ_i. Assume that t and w_3 are single letters G-conjugate by some single letter. Then $\Gamma_A = \Gamma_A^m$ satisfies the theorem. So suppose throughout the rest of the proof that t and w_3 are not single letters G-conjugate by any single letter.

Assume first that Case 1 of Proposition 3.3.4 holds for v, and that either $|\rho| \geq 2$, or ρ is a loop and the endpoint v_1 of ρ also satisfies $\delta_G(v_1, \Delta_i) \geq 2$. By Assumption 2.3.15 all $R \in \mathcal{R}$ satisfy $|R| \geq 3$, hence the boundary words of Δ_i are cyclically $\mathcal{P}$-reduced. Since by Lemma 3.1.12 and Proposition 3.2.13 Γ_A^m satisfies Axiom T_2, so does Δ_i. The fact that Γ_A^m is a minimal conjugacy diagram implies that the same holds for Δ_i. Hence by Lemma 3.2.5 all vertices $v' \in \Delta_i$ that are not incident with any loop labelled by a V^σ-letter satisfy $\delta_G(v', \Delta_i) \geq 2$. Therefore, Δ_i satisfies Axiom T_3, so by Lemma 3.2.7 Δ_i satisfies Axiom T_4; and as the boundary words of Δ_i are cyclically $\mathcal{P}$-reduced, by Lemma 3.1.16 and Proposition 3.2.9 Δ_i satisfies Axiom T_1. Thus, as Δ_i is loop-minimal (see Definition 3.1.6), Δ_i satisfies Axioms $T_1 - T_5$.

Next assume that v_1 satisfies Case 2 of Proposition 3.3.4. We first show that there exists a smaller sibling Γ_A of Γ_A^m with layers $\Gamma_1, \ldots, \Gamma_n$ as in the statement of the theorem, and such that $\partial(\Gamma_i) = \{\bar{l}, \rho\}$, Γ_i is decomposable, and the label of the length 2 boundary of each boundary red blob B of Γ_i is non-trivial in P, is not P-reduced, and is not equal to the other boundary word of B in P-reduced form.

There is a unique (red) face T of Δ_i with $l \subseteq \partial(T)$. Hence we can let Θ be the annular sub-diagram of Δ_i bounded by ρ and the edges e_1 and e_2 of T with $e_1 \neq l \neq e_2$ (see Figure 3.2). Now e_1 and e_2 meet at v, and a vertex u, say. Let a and b be labels of e_1 and e_2 respectively when oriented by T.

We first show that there is no other loop l_2 in Δ_i incident with v. Suppose for a contradiction that such a loop exists. Then as $\delta_G^I(v, \Delta_i) = 0$, Theorem 2.3.11 implies that the labels of l and l_2 are equal in P. So we can identify l and l_2 and delete the sub-diagram bounded by l and l_2 from Γ_A^m. This yields a coloured annular diagram Δ_A with boundary words w_1 and w_2 and with $\mathbf{CArea}(\Delta_A) < \mathbf{CArea}(\Gamma_A^m)$, contradicting Definition 3.1.4. Hence no such l_2 exists in Δ_i.

Next we show that without loss of generality $\delta_R(v, \Delta_i) = 4$. Suppose first that $\delta_R(v, \Delta_i) = 3$. Then as there is no loop $l_2 \neq l$ in Δ_i incident with v, there is a loop l' based at u with label t'. Hence t and t' are $U(P)$-conjugate by b. Since Δ_i contains no interior loops, we have $l' = \rho$, contradicting our assumption on t and w_3 from the first paragraph. Now suppose that $\delta_R(v, \Delta_i) \geq 5$. Then $\delta_R(v, \Theta) \geq 3$. Hence as there are no loops labelled by V^σ-letters based at v in Θ, we can repeatedly use Lemma 3.1.14 to reduce $\delta_R(v, \Delta_i)$ to four.

Hence as $\delta_R(v, \Delta_i) = 4$, there are two red triangles $T_1 \subseteq \Theta$ and $T_2 \subseteq \Theta$ distinct from T and incident with both v and u. Let Δ be the annular sub-diagram of Θ such that $\Theta = \Delta \cup \overline{T_1} \cup \overline{T_2}$. It follows that $\partial(\Delta) = \{\rho, e_3 e_4\}$, where $e_{2i+1} \subseteq \partial(T_i)$. Let c and d be inverses of labels of e_3 and e_4 when oriented by T_1 and T_2, and let x be the label of the common edge of T_1 and T_2 when oriented by T_1 (see Figure 3.3). Now $bcdb^\sigma =_{U(P)} t =_{U(P)} xdcx^\sigma$. Hence if $c = d^\sigma$ or $d = c^\sigma$, then $t =_{U(P)} 1$, contradicting Theorem 2.3.11. So $c \neq d^\sigma$ and $d \neq c^\sigma$. In addition, by Lemma 2.5.13 we have $(c, d), (d, c) \in D(P)$ since $\mathcal{P}$ satisfies trivial-interleaving and both cd and dc are boundary words of the red blob $\overline{T} \cup \overline{T_1} \cup \overline{T_2}$. Suppose that $[cd] =_P t$. Then we can delete the triangles T, T_1, T_2 from Γ_A^m, and add a new triangle with label $d^\sigma c^\sigma t$ to Γ_A^m as in Figure 3.4. The obtained annular diagram has boundary words w_1 and w_2 and strictly smaller coloured area than Γ_A^m, contradicting Definition 3.1.4. Hence $[cd] \neq_P t$, and similarly $[dc] \neq_P t$.

Suppose first that $|\rho| \geq 2$. Then Δ_i is decomposable, $\overline{T} \cup \overline{T_1} \cup \overline{T_2}$ is the boundary red blob of Δ_i, and Δ is the core of Δ_i. Furthermore, by the previous paragraph the label of the length 2 boundary of $\overline{T} \cup \overline{T_1} \cup \overline{T_2}$ is non-trivial in P, is not P-reduced, and is not equal to t.

Suppose instead that ρ is a loop. Let w be the vertex incident with e_3 and e_4 distinct from u. We show that there is no loop based at u or w. If there are loops based at both u and w, then one of these loops contradicts Lemma 3.1.7. If there is precisely one loop l' with label t' based at u or w, then t' is G-conjugate to t by b or x. Hence as Δ_i contains no interior loops, we have $\rho = l'$, contradicting our assumption on t and w_3 from the first paragraph. Hence the endpoint $v_1 \in \rho$ satisfies $u \neq v_1 \neq w$, so applying Proposition 3.3.4 and the arguments above shows that either $\delta_G(v_1, \Delta_i) \geq 2$, or without loss of generality we can assume $\delta_G(v_1, \Delta_i) = 1$ and $\delta_R(v_1, \Delta_i) = 4$. Thus, Δ_i is decomposable, and the label of the length 2 boundary of each boundary red blob B of Δ_i is non-trivial in P, and is not P-reduced, and is not equal to the other boundary word of B in P-reduced form. Hence we showed that there exists a diagram Γ_A with properties as in the third paragraph.

Case $\delta_R(v, \Delta_i) = 4$**.**

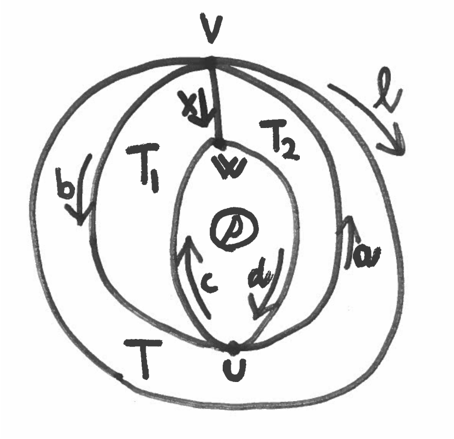

Figure 3.3: A boundary red blob comprised of three triangles T, T_1, T_2, see the proof of Theorem 3.3.5.

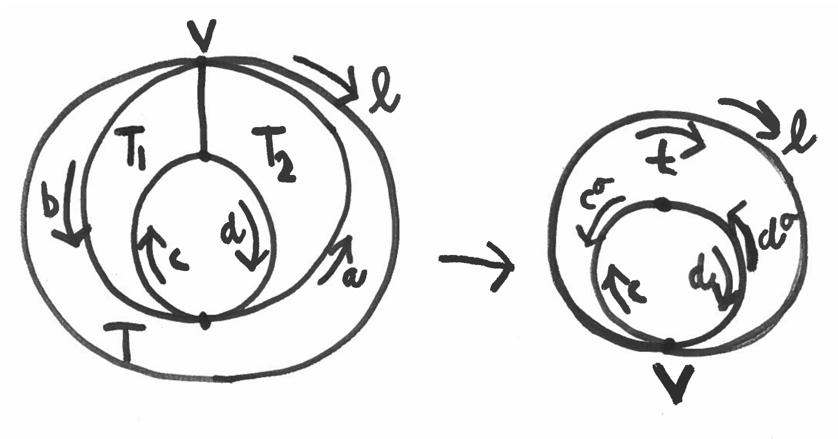

Figure 3.4: Reducing the coloured area of Γ_A^m, see the proof of Theorem 3.3.5.

Let Γ be the core of Γ_i, and note that the boundary words of Γ_i are cyclically $\mathcal{P}$-reduced. Since $\mathbf{CArea}(\Gamma_A) = \mathbf{CArea}(\Gamma_A^m)$, Γ_i and Γ are minimal conjugacy diagrams. Therefore, by Lemma 3.1.12 Γ is semi-P-reduced, hence by Proposition 3.2.13 Γ satisfies Axiom T_2. By Lemma 3.2.5 all vertices $v' \in \Gamma_i$ that are not incident with any loop labelled by a V^σ-letter satisfy $\delta_G(v', \Gamma_i) \geq 2$. Therefore, Γ satisfies Axiom T_3, so by Lemma 3.2.7 Γ satisfies Axiom T_4. Thus, as Γ is loop-minimal, Γ satisfies Axioms T_2-T_5.

Finally, assume that $\delta_G(v, \Delta_i) \geq 2$, that ρ is a loop, and that the endpoint v_1 of ρ satisfies $\delta_G(v_1, \Delta_i) = 1$. Then applying arguments from the previous case shows that there exists a smaller sibling of Γ_A^m with layers $\Gamma_1, \ldots, \Gamma_n$ as in the statement of the theorem, such that Γ_i is decomposable with precisely one boundary red blob B, and B and the core of Γ_i satisfy Statement 3. The theorem follows. ∎

Proof of Theorem 3.3.3. Since w_1 and w_2 are not $U(P)$-conjugate, by Theorem 3.1.2

$\mathbf{CArea}(\Gamma_A^m) \geq (1,0)$. So if $\mathbf{CArea}(\Gamma_A^m)$ contains no loops, then Γ_A^m is a single layer with area at least 1, and by Theorem 3.2.2 we have $\Gamma_A^m \in \mathcal{T}$, so $\Gamma_A^m \in \mathcal{S}$. Hence we can assume that Γ_A^m contains a loop. Then applying Theorem 3.3.5 repeatedly shows that there exists a smaller sibling Γ_A of Γ_A^m (hence Γ_A is a minimal conjugacy diagram for w_1 and w_2) such that Γ_A is a face-disjoint union of n layers, where each layer satisfies one of the Statements 1-3 of Theorem 3.3.5. In particular, if Γ_A^m is a single layer, then Γ_A is a single layer.

So it remains to prove that if a layer $\Gamma \subseteq \Gamma_A$ satisfies Statement 2 of Theorem 3.3.5, then Γ satisfies Axiom T_6; and if Γ satisfies Statement 3 of Theorem 3.3.5, then the core Γ_1 of Γ satisfies Axioms T_1 and T_6. Suppose first that Γ satisfies Statement 2 of Theorem 3.3.5. Then as $\mathcal{P}$ is valid, we can use the same arguments as in the proof of Theorem 3.2.2 to deduce that Γ satisfies Axiom T_6.

Assume instead that Γ satisfies Statement 3 of Theorem 3.3.5. Since we can assume that Γ does not satisfy Statement 1 of Theorem 3.3.5, by Theorem 3.2.11 it follows that $\mathbf{CArea}(\Gamma) \geq (1,0)$. Let ρ be a length 2 boundary of some boundary red blob of Γ, and let $F \subseteq \Gamma_1$ be a green face. Suppose that ρ is a sub-path of $\partial(F)$. Then as Γ_1 is green-rich, we contradict Lemma 3.2.4. Hence if $\rho \subseteq \partial(F)$, then the label R of F satisfies $|R| \geq 4$. Else F contains at most one edge of ρ. Therefore, since by Assumption 2.3.15 we have $|R| \geq 3$, it follows that F does not have more than $1/2$ of its length on ρ. Now the boundaries of each boundary red blob of Γ are simple closed paths (see Figure 3.3), so in particular, ρ is a simple closed path. Also, as the boundary words of Γ_A are cyclically $\mathcal{P}$-reduced, the label of every boundary ρ' of Γ_1 that is not a boundary of any boundary red blob of Γ is cyclically $\mathcal{P}$-reduced. Hence by Lemma 3.1.16 no internal green of Γ_1 has more than $1/2$ of its length on ρ'; and as Γ_1 is green-rich, by Proposition 3.2.9 ρ' is a simple closed path. It follows that Γ_1 satisfies Axiom T_1. Thus, applying the arguments of the proof of Theorem 3.2.2 again shows that Γ_1 satisfies Axiom T_6. ∎

Chapter 4

Foundational theory for diagrams in $\mathcal{T}$

4.1 Introduction and statement of the Three Face Theorem

We study diagrams in the set $\mathcal{T}$ from Definition 3.1.8. Recall Definition 2.5.3 that ω and τ denote the outer and the inner boundary of an annular diagram respectively, and Definition 2.6.16 of dual distance.

Theorem 1. (Three Face Theorem) *Let $\Gamma_A \in \mathcal{T}$. Treat each red blob of Γ_A as a single face. Then for all edges $e \subseteq \omega$, at least one of the following statements holds.*

1. *There is a vertex $v \in \tau$ incident with e.*

2. *There is an internal face F with $\overline{e} \cap \partial(F) \neq \emptyset$ such that either $\partial(F) \cap \tau \neq \emptyset$, or F is at dual distance at most three from the external face with boundary τ. Furthermore, either $e \subseteq \partial(F)$; or e is red, F is green and F has an edge $e_1 \subseteq \omega$ incident with e.*

We shall prove Theorem 1 in Section 5.6. Theorem 1 will enable us to make our conjugacy problem solver quadratic.

Since our analysis is extensive, let us define a subset $\mathcal{U} \subseteq \mathcal{T}$ of diagrams satisfying two conditions that will help us to split the proofs into simpler cases.

Definition 4.1.1. Let Γ_A be a coloured annular diagram with a red blob B. We say that B is *bad* if B is annular and $\Gamma_A \setminus B^\circ$ decomposes as a disjoint union of two annular diagrams. We say that B is *good* if B is not bad.

Recall Definition 2.6.7 of curvature incidence.

Definition 4.1.2. Let $\mathcal{U}$ be the set of all diagrams $\Gamma_A \in \mathcal{T}$ satisfying the following conditions.

1. Every vertex and every red blob of Γ_A is curvature incident with each internal green face of Γ_A at most once.

2. There are no bad red blobs in Γ_A.

We always assume that $\Gamma_A \in \mathcal{T}$ contains at least one internal green face unless stated otherwise.

4.2 Structural properties of red blobs and internal green faces

In this section we collect some foundational results about internal green faces and red blobs of diagrams in the set $\mathcal{T}$. We start by proving that certain pathological cases do not arise.

Recall that if we say that a closed path p is of the form $p_1p_2\dots p_n$, then p is a sequence of simple sub-paths p_i.

Definition 4.2.1. Let Γ_A be a coloured annular diagram with a non-empty sequence $\mathcal{F} = (D_1, D_2, \dots, D_n)$ consisting of internal green faces and red blobs such that $D_i \neq D_j$ for $i \neq j$, and with an external face I with boundary τ.

A *sub-diagram of* Γ_A *bounded by* $\mathcal{F}$ is equal to $\overline{C \setminus I}$ for some bounded component C of $\mathbb{R}^2 \setminus \cup_{D\in\mathcal{F}}\overline{D}$ that is bounded by a closed path of the form $p_1 \dots p_n$ with $p_i \subseteq D_i$ and $|p_i| \geq 1$, and (possibly) τ.

It is worth noting that by Definition 4.2.1 a sub-diagram K of Γ_A bounded by $\mathcal{F}$ has $K \subseteq \Gamma_A$, hence K is a sub-diagram of Γ_A in the sense of Definition 2.5.10.

Lemma 4.2.2. *Let $\Gamma_A \in \mathcal{T}$, and let $\mathcal{F}$ be a non-empty sequence consisting of at most one internal green face and at most one red blob of Γ_A. If K is a sub-diagram of Γ_A bounded by $\mathcal{F}$, then K contains an edge of $\partial(\Gamma_A)$.*

Proof. Suppose for a contradiction that there is a sub-diagram K of Γ_A bounded by $\mathcal{F}$ that does not contain a boundary edge. Let $p = \partial(K)$. Suppose first that $\mathcal{F}$ consists of a single internal green face F. As $|p| \geq 1$, it follows that p is comprised of at least one sub-path of $\partial(F)$ that contains an edge of F. Since K does contain any boundary edge, we have that K does not contain τ, so K is simply-connected. Now if K contains an internal green face F_1, then F_1 contradicts Axiom T_6. Hence K is a simply-connected red blob.

Suppose first that p is a sub-path of $\partial(F)$. Then Lemma 3.2.4 implies $|p| = 1$, contradicting Theorem 2.3.11 since P embeds into $U(P)$.

Next suppose that p is comprised of two sub-paths p_1 and p_2 of $\partial(F)$ with $|p_1|, |p_2| \geq 1$. By applying Lemma 3.2.4 again we have $|p_1| = 1 = |p_2|$. But by Lemma 3.2.6 we then have $\mathbf{Area}(K) = 0$, a contradiction.

Finally, suppose that for some $n \geq 3$: $p = p_1, p_2, \dots, p_n$, where each p_i is a sub-path of $\partial(F)$ with $|p_i| \geq 1$. Since F is homeomorphic to a disc, $\mathbb{R}^2 \setminus \overline{F}$ contains besides K° at least two bounded components C such that $\partial(C)$ is a simple closed path in Γ_A^1 (see Definition 2.5.14) with $|\partial(C)| \geq 1$, and is a sub-path of $\partial(F)$. Now the closure K' of one of these two components does not contain τ, so K' has no edge of $\partial(\Gamma_A)$. Thus, as $\partial(K')$ is a sub-path of $\partial(F)$, we have a contradiction as before.

Now suppose that $\mathcal{F}$ contains a blob B. Since K contains no edge of $\partial(\Gamma_A)$, it follows from $|p \cap \partial(B)| \geq 1$ that K contains an internal green face F. But F contradicts Axiom T_6. ∎

Lemma 4.2.3. *Let x be an internal green face or a simply-connected red blob of $\Gamma_A \in \mathcal{T}$. Then all of the following statements hold.*

Figure 4.1: Possible cases when $\partial(x)$ passes more than once through some vertex, and x is either a simply-connected red blob B or an internal green face F of $\Gamma_A \in \mathcal{T}$, see Lemma 4.2.3.

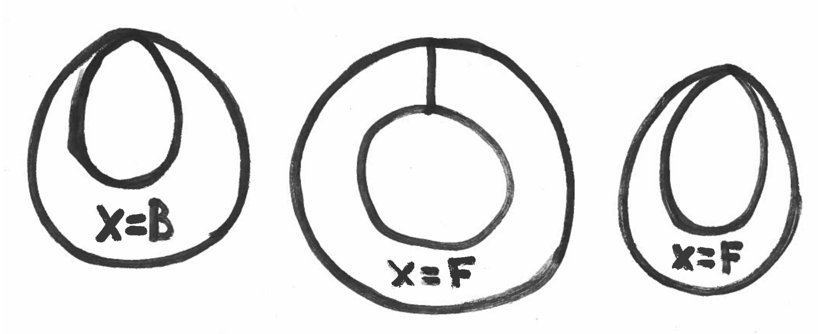

1. *No vertex is incident more than twice with x.*
2. *Suppose that x is a red blob B. Then $\partial(B)$ passes more than once through at most one vertex (see Figure 4.1).*
3. *Suppose that x is a green face F. Then there are at most two vertices of degree at least three incident more than once with F, and if there are two such vertices $v \neq w$, then F is edge-incident with itself by a consolidated edge e with $\partial(e) = \{v\} \cup \{w\}$ (see Figure 4.1). Hence if $\partial(F)$ passes through some vertex more than once, then $\partial(\Gamma_A) \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths.*

Proof. For Part 1, suppose for a contradiction that v exists. Then as x° is homeomorphic to a disc, there are at least two bounded components C_1 and C_2 of $\mathbb{R}^2 \setminus \overline{x}$ such that $\partial(C_i)$ is a simple closed path with $|\partial(C_i)| \geq 1$. Now one of these components C_i has $\tau \not\subseteq \overline{C_i}$, so $\overline{C_i}$ contradicts Lemma 4.2.2.

For Part 2, suppose for a contradiction that $\partial(B)$ passes more than once through at least two vertices. Since B° is homeomorphic to a disc, $\mathbb{R}^2 \setminus B$ contains at least two bounded components with properties as in the previous paragraph, a contradiction.

For Part 3, assume for a contradiction that such an x exists. Then by Part 1 there are either at least three vertices of degree at least three incident twice with F, or there are precisely two such vertices and F is not edge-incident with itself by any consolidated edge. Hence as F is homeomorphic to a disc, $\mathbb{R}^2 \setminus \overline{F}$ contains at least two bounded components with properties as in the first paragraph, a contradiction.

To prove the final statement, assume that $\partial(F)$ passes more than once through some vertex, but that $\partial(\Gamma_A) \setminus (\overline{F})^\circ$ does not decompose as stated. By Axiom T_1, ω and τ are simple closed paths, hence by the previous paragraph $\mathbb{R}^2 \setminus \overline{F}$ contains again a bounded component that contradicts Lemma 4.2.2. ∎

Lemma 4.2.4. *Let B be an annular red blob. Then $\partial(B)$ is comprised of the outer boundary ω_1 and the inner boundary τ_1, which are simple closed paths with $\omega_1 \cap \tau_1 = \emptyset$.*

Hence $\Gamma_A \in \mathcal{T}$ contains at most one bad red blob, and if Γ_A contains such a blob, then Γ_A is island-free.

Proof. By Definition 2.5.12 B^* is connected and $\mathbb{R}^2 \setminus B^\circ$ is comprised of two components, so the first statement follows.

Suppose that $\Gamma_A \in \mathcal{T}$ contains a bad red blob B. Then by Definition 4.1.1 $\Gamma_A \setminus B^\circ$ decomposes as a disjoint union of two annular diagrams. Hence by Lemma 3.1.10 Γ_A is island-free. Since by Axiom T_6 each internal green face of Γ_A contains a boundary edge, it follows that B is the only bad red blob of Γ_A. ∎

Lemma 4.2.5. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B. Then*

1. *B is good if and only if B is simply-connected;*
2. *if B is simply-connected and $\partial(B)$ passes through some vertex more than once, then $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths.*

Proof. For Part 1, note that the reverse implication follows directly from Definition 4.1.1. So assume that B is good, and suppose first that B is annular. Then by Definition 2.5.12 $\mathbb{R}^2 \setminus B^\circ$ is comprised of two components; and as B is good, by Definition 4.1.1 $\Gamma_A \setminus B^\circ$ does not decompose as a disjoint union of two annular diagrams. Hence there is a bounded component C of $\mathbb{R}^2 \setminus B$ such that $\overline{C}$ contains no edge of $\partial(\Gamma_A)$, and $\partial(C)$ is a simple closed path with $|\partial(C)| \geq 1$, contradicting Lemma 4.2.2.

Now assume that B is not annular. Since B is not simply-connected, $\mathbb{R}^2 \setminus B^\circ$ contains at least 3 components, hence there is at least one bounded component of $\mathbb{R}^2 \setminus B$ with properties as in the previous paragraph, a contradiction.

For Part 2, suppose for a contradiction that $\Gamma_A \setminus B^\circ$ does not decompose as stated. By Axiom T_1, ω and τ are simple closed paths, hence by Lemma 4.2.3 there is a bounded component of $\mathbb{R}^2 \setminus B$ with properties as in the first paragraph, a contradiction. ∎

Lemma 4.2.6. *Let $\Gamma_A \in \mathcal{T}$ contain an internal green face F and a red blob B such that $\partial(F) \cap \partial(B)$ contains a path $p \in \Gamma_A^1$ with $|p| > 1$. Then $\partial(F)$ or $\partial(B)$ passes more than once through each vertex of p common to two edges of p.*

Proof. Suppose for a contradiction that p contains a vertex v common to two edges of p such that $\partial(F)$ and $\partial(B)$ does not pass through v more than once. Then there is a consolidated edge between F and B of length at least two. But this contradicts Lemma 3.2.4. ∎

Lemma 4.2.7. *Let $\Gamma_A \in \mathcal{T}$ contain an internal green face F and a simply-connected red blob B edge-incident n times with F for some $n \geq 2$. Then $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at least $n-1$ bounded components.*

Proof. By Lemma 4.2.3 $\partial(B)$ does not pass more than twice through any vertex of Γ_A, and there is at most one vertex v such that $\partial(B)$ passes through v twice.

Suppose there is such a v. By Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Now precisely one of the diagrams in this decomposition contains F, call it Γ. Assume that $\partial(F)$ does not pass more than once through any vertex. Then $\partial(F)$ is a simple closed path. Hence since v is the only vertex such that $\partial(B)$ passes through v twice, by Lemma 4.2.6 $\partial(F) \cap \partial(B)$ contains at most one path $p \in \Gamma_A^1$ with $|p| > 1$; and if such a p exists, then p cannot be closed and $|p| = 2$. Hence as $\partial(F) \cap \partial(B)$ contains n edges, it follows that $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at least $n - 1$ bounded components.

Now suppose that $\partial(F)$ passes more than once through some vertex. By Part 3 of Lemma 4.2.3 $\Gamma_A \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of two annular diagrams, so $\mathbb{R}^2 \setminus \overline{F}$ contains a bounded component C; and $\partial(F)$ passes more than once through at most one vertex of $\partial(F) \cap \partial(B)$. Hence $\partial(F) \cap \partial(B)$ contains at most two vertices through which $\partial(F)$ or $\partial(B)$ passes more than once, so by Lemma 4.2.6 $\partial(F) \cap \partial(B)$ contains at most two paths $p \in \Gamma_A^1$ with $|p| > 1$; and if such a p exists, then $|p| \leq 3$, $|p| = 2$ if p is closed, and if $|p| = 3$ or if p is closed, then p is the only path of $\partial(F) \cap \partial(B)$ with $|p| > 1$. Therefore, as $F \subseteq \Gamma$, we deduce that $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at least $n - 2$ bounded components distinct from C, so the lemma follows.

Now the case where $\partial(F)$ passes more than once through some vertex and $\partial(B)$ is a simple closed path is similar to the case analysed in the second paragraph, so it remains to consider the case where $\partial(F)$ and $\partial(B)$ are both simple closed paths. By Lemma 4.2.6 all paths $p \in \Gamma_A^1$ with $p \subseteq \partial(F) \cap \partial(B)$ satisfy $|p| \leq 1$, so the lemma follows as $\partial(F) \cap \partial(B)$ contains n edges. ■

The following proposition provides a useful restriction.

Proposition 4.2.8. *Let $\Gamma_A \in \mathcal{T}$ contain an internal green face F and a red blob $B \in \mathcal{S}_F$. Then $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at most one bounded component (and if such a component exists, then it contains the external face of Γ_A with boundary τ); B is edge-incident at most twice with F; and if B is simply-connected and edge-incident twice with F, then Γ_A is island-free and $\Gamma_A \setminus (\overline{F} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths.*

In particular, for every internal green face F, every element of $\mathcal{S}_F$ is curvature incident at most twice with F.

Proof. Suppose first that $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains more than one bounded component, or precisely one bounded component, which does not contain the external face E with boundary τ. Then the closure of at least one of these components contradicts Lemma 4.2.2. Hence the first statement of the lemma holds.

Next assume that B is edge-incident more than twice with F. By Lemma 4.2.5 B is simply-connected. So by Lemma 4.2.7 $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at least two bounded components, contradicting the first statement of the lemma. Therefore, B is edge-incident at most twice with F, and by Lemma 4.2.3 the last statement of the lemma holds.

Finally, suppose that B is simply-connected and edge-incident twice with F. Let e and f be the distinct edges of $\partial(B) \cap \partial(F)$. By Lemma 4.2.7 and the first statement of the lemma $\mathbb{R}^2 \setminus$

$(\overline{F} \cup B)$ contains precisely one bounded component, which contains E. Hence $\Gamma_A \setminus (\overline{F} \cup B)^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2, where $\partial(\Gamma_1) = \{\omega, \tau_1\}$ and $\partial(\Gamma_2) = \{\omega_1, \tau\}$. By Axiom T_1, ω and τ are simple closed paths. We show that ω_1 and τ_1 are also simple closed paths. Let $\rho \in \{\omega_1, \tau_1\}$, and suppose first that $\rho \subseteq \partial(x)$ for some $x \in \{F, B\}$ such that $\partial(x)$ passes through some vertex more than once. By Lemmas 4.2.3 and 4.2.5 $\Gamma_A \setminus (\overline{x})^\circ$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths. Hence ρ is a simple closed path. Otherwise, since $e, f \subseteq \partial(B) \cap \partial(F)$, we deduce that ρ is a simple closed path of the form p_1p_2, where $p_1 \subseteq \partial(F)$ and $|p_1| \geq 1$, and $p_2 \subseteq \partial(B)$ and $|p_2| \geq 1$.

It remains to show that Γ_A is island-free. Suppose not. By Lemma 3.1.10 F is contained in some island E of Γ_A (see Definition 2.5.10). Let v_1 and v_2 be the endpoints of E. Since $\omega \subseteq \partial(\Gamma_1)$ and $\tau \subseteq \partial(\Gamma_2)$, we have $v_1 = v_2$, and either $v_1 \in \partial(F)$ or $v_1 \in \partial(B)$. But then as $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at most one bounded component, ef is a path and there is a vertex of ef common to e and f that contradicts Lemma 4.2.6. ∎

4.3 Curvature neighbourhoods

The next two definitions introduce sets containing both vertices and red blobs. These may seem unintuitive, but these are the objects that give non-zero curvature to its curvature incident internal green faces in Steps 3 and 4 of $\mathbf{ComputeRSym}$ (see Algorithm 2.6.4), so are crucial.

Definition 4.3.1. For an internal green face F of an annular diagram Γ_A, we define the *curvature neighbourhood of* F to be

$$\begin{aligned}\mathcal{S}_F :=\{&\text{Red blobs and vertices giving } F \text{ non-zero curvature in Steps 3 and 4}\\ &\text{of } \mathbf{ComputeRSym}(\Gamma_A).\}\end{aligned}$$

Lemma 4.3.2. *Let F be an internal green face of a green-rich coloured annular diagram Γ_A. Then the curvature neighbourhood of F is the set of all red blobs edge-incident with F and vertices v of F with $\delta_G(v) \geq 3$.*

Hence for all vertices v of F with $\delta(v, \Gamma_A) \geq 3$, either $v \in \mathcal{S}_F$, or v is incident with a red blob in $\mathcal{S}_F$.

Proof. For the first statement it suffices to show that all vertices v of F with $\delta_G(v) = 2$ give F curvature 0 in Step 3 of $\mathbf{ComputeRSym}$, which follows from Lemma 2.6.10. For the second statement, since Γ_A is green-rich, either $\delta_G(v) \geq 3$, so $v \in \mathcal{S}_F$, or $\delta_G(v) = 2$ and $\delta_R(v) \geq 1$, hence v is incident with a red blob in $\mathcal{S}_F$. ∎

Definition 4.3.3. For an internal green face F of an annular diagram Γ_A, we define the *boundary curvature neighbourhood of* F to be the set $\mathcal{B}_F$ consisting of all vertices v and all red blobs B of Γ_A that satisfy the following conditions.

1. v is incident with F; v is boundary; and $\delta_G(v) \geq 3$.

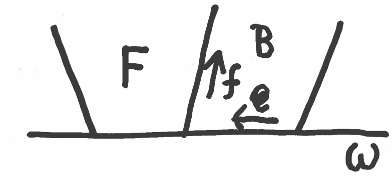

Figure 4.2: A red blob B with $B \in \mathcal{B}_F$, see Definition 4.3.3.

2. There are edges e and f such that e and f are consecutive on $\partial(B)$; $e \subseteq \partial(\Gamma_A)$; and $f \subseteq \partial(F)$ (see Figure 4.2).

From Lemma 4.3.2 it follows that $\mathcal{B}_F \subseteq \mathcal{S}_F$. Recall Definition 2.5.14 of the 1-skeleton of a coloured diagram.

Lemma 4.3.4. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Then all of the following statements hold.*

1. *All boundary vertices v of F of degree at least 3 are either in $\mathcal{B}_F$, or there is a red blob in $\mathcal{B}_F$ edge-incident with F at v; and if v lies on precisely one boundary ρ of Γ_A, then any red blob in $\mathcal{B}_F$ edge-incident with F at v has an edge on ρ.*

2. *Each element of $\mathcal{B}_F$ gives F curvature of at most $-1/4$ across each curvature incidence in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma_A)$.*

3. *$|\mathcal{B}_F| \geq 1$, and if $|\mathcal{B}_F| = 1$, then the element in $\mathcal{B}_F$ is curvature incident exactly twice with F.*

4. *The elements of $\mathcal{B}_F$ collectively give F curvature of at most $-1/2$ in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma_A)$. In particular, the elements of $\mathcal{S}_F \setminus \mathcal{B}_F$ collectively give F curvature $-1/2 \leq \chi \leq 0$ in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma_A)$.*

Proof. To prove Part 1, let v be a boundary vertex of F with $\delta(v, \Gamma_A) \geq 3$. If $\delta_G(v, \Gamma_A) \geq 3$, then $v \in \mathcal{B}_F$. Otherwise, $\delta_G(v, \Gamma_A) = 2$, so $\delta_R(v, \Gamma_A) \geq 1$ and there is a red blob in $\mathcal{B}_F$ edge-incident with F at v. The second statement follows directly from Definition 4.3.3.

Part 2 follows from Lemmas 2.6.10 and 3.2.8 as all elements of $\mathcal{B}_F$ are curvature incident with an external face and F.

By Axiom T_6, F has a boundary consolidated edge e with $|e| \geq 1$. Let v and w be the endpoints of e, and let $\rho \in \{\omega, \tau\}$ be such that $e \subseteq \rho$. By Part 1, each $t \in \{v, w\}$ is either in $\mathcal{B}_F$, or is incident with a red blob in $\mathcal{B}_F$. Hence $|\mathcal{B}_F| \geq 1$. Suppose that $|\mathcal{B}_F| = 1$. If $\mathcal{B}_F$ contains a vertex, then $v = w$, and by Lemma 4.2.3 v is incident exactly twice with F. If $\mathcal{B}_F$ contains a red blob, then there is a red blob $B \in \mathcal{B}_F$ incident with v and w; $\delta_G(v) = 2 = \delta_G(w)$; and $\partial(F) \cap \rho = e$. We show that B is edge-incident more than once with F. Suppose not. Then F is contained in a sub-diagram K of Γ_A bounded by a closed path ee', where e' is the common edge of F and B. In particular, ee' is a sub-path of $\partial(F)$. Hence if K contains τ or if F is not contained in the bounded component of $\mathbb{R}^2 \setminus ee'$, then F is not simply-connected,

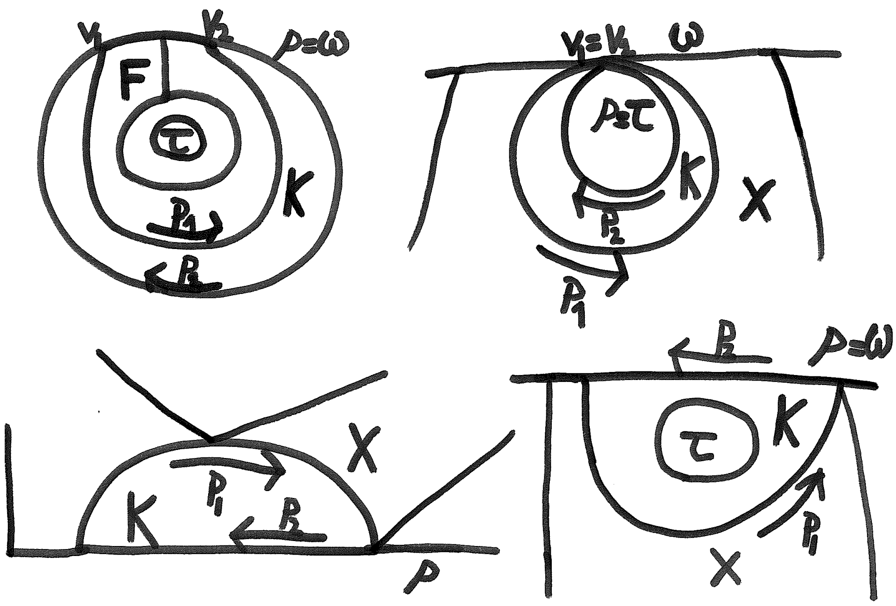

Figure 4.3: Configurations of a sub-diagram K of Γ_A well-bounded by x and ρ. The notation $x = F$ used in the top left picture of the figure means that x is an internal green face F of Γ_A, see Definition 4.4.2.

a contradiction. Therefore, K is simply-connected and in fact $K = \overline{F}$, so $\partial(F) = ee'$. By Axiom T_1, the label R of F does not have more than $1/2$ of its length on ρ. Thus, $|R| = 2$, which contradicts Assumption 2.3.15 that no $R \in \mathcal{R}$ has $|R| \in \{1, 2\}$. Hence by Proposition 4.2.8 B is edge-incident exactly twice with F, so Part 3 follows.

By Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$, so Part 4 is immediate from Parts 2-3. ■

4.4 Tricky boundaries of red blobs and internal green faces

Let x be a red blob or an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. In this section we introduce concepts fruitful for studying cases when $\partial(x) \cap \rho$ is complicated.

Recall that a consolidated edge between two faces $F_1, F_2 \subseteq \Gamma_A$ is a non-empty path of maximal length that is a sub-path of both $\partial(F_1)$ and $\partial(F_2)$, and that a path can consist of a single vertex, in which case has length zero.

Definition 4.4.1. Let x be a red blob or an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. We say that x is *tricky for* Γ_A *with respect to* ρ if $\partial(x) \cap \rho \neq \emptyset$ and $\partial(x) \cap \rho$ is not a single consolidated edge.

Recall Definition 2.5.10 of a sub-diagram in a coloured diagram, and Definition 2.5.14 of the 1-skeleton of a coloured diagram. Recall also that if we say that a closed path p is of the form $p_1 p_2 \ldots p_n$, then p is a sequence of simple sub-paths p_i.

We shall work extensively with the concepts presented in the next two definitions.

Definition 4.4.2. Let x be a red blob or an internal green face of an annular diagram Γ_A, and let $\rho \in \{\omega, \tau\}$. A *sub-diagram of* Γ_A *well-bounded by* x *and* ρ is a sub-diagram K of Γ_A with $K \subseteq \Gamma_A \setminus (\overline{x})^\circ$ that is bounded by a closed path $p \in \Gamma_A^1$ of the form p_1p_2, where $p_1 \subseteq \partial(x)$ and $|p_1| \geq 1$; $p_2 \subseteq \rho$; and K° is contained in some bounded component of $\mathbb{R}^2 \setminus p$ (see Figure 4.3). The endpoints of p_1 are called the *corners* of K.

We say that K is *trivial* if $\mathbf{CArea}(K) < (1, 0)$, and that K is *well-connected* if either K is simply-connected, or is annular and consisting of a single island with boundaries p_1 and p_2 (so that K is not of a form as shown in the bottom right picture of Figure 4.3).

Definition 4.4.3. Let x be a red blob or an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. A sub-diagram K of Γ_A well-bounded by x and ρ is *well-contained* if K is non-trivial, K is well-connected, and $\partial(x)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho$.

We shall show in Sections 4.5-4.6 that if x is a tricky green face or a tricky red blob for Γ_A with respect to ρ, then there exists a well-contained sub-diagram of Γ_A given by x and ρ.

Let Γ be a coloured annular or simply-connected diagram. Recall Definition 2.6.8 that $\chi(x, \Gamma)$ is the curvature that x gives to a single internal green face across each curvature incidence, and that $\zeta(x, \Gamma)$ is the total curvature that x gives to internal green faces. Recall also that by Algorithm 2.6.4, if $e \subseteq \Gamma$ is an edge, then $\kappa_\Gamma(e) = 0$; and if $F \subseteq \Gamma$ is a green face, then $\kappa_\Gamma(F) = 1 + \sum_{v \in \partial(F)} \chi(v, F, \Gamma) + \sum_B \chi(B, F, \Gamma)$, where the last sum is over all red blobs edge-incident with F.

Definition 4.4.4. Let K be a sub-diagram of a coloured annular or simply-connected diagram Γ. We let K^Γ be the subset of Γ that is equal to K but excluding any external faces.

If K^Γ contains an internal green face, then we define

$$\kappa(K^\Gamma) = \sum_{F\colon \text{a green face of } K^\Gamma} \kappa_\Gamma(F).$$

Finally, if $x \in \Gamma$ is a red blob or a vertex, then we let $\zeta(x, K^\Gamma)$ be the total curvature that x gives to internal green faces of K in Steps 3 and 4 of $\mathbf{ComputeRSym}(\Gamma)$.

Note that $\kappa(K^\Gamma) = N + \sum_{x \in \Gamma} \zeta(x, K^\Gamma)$, where N is the number of internal green faces of K. Furthermore, by Definition 4.4.3 if K is a well-contained sub-diagram of $\Gamma_A \in \mathcal{T}$, then $\kappa(K)$ and $\kappa(K^{\Gamma_A})$ are defined. We now present several useful results that analyse well-contained sub-diagrams given by red blobs and ρ.

Throughout the rest of this section we shall use the following notation.

Notation 4.4.5. Let K be a well-contained sub-diagram of Γ_A given by a red blob B and some $\rho \in \{\omega, \tau\}$. Let p_1 and p_2 be as in Definition 4.4.2, and let v_1 and v_2 be the corners of K. Write $p_1 = e_1e_2 \ldots e_n$, where $v_1 \in \overline{e}_1$ and $v_2 \in \overline{e}_n$. Further, let $w_i = \overline{e}_i \cap \overline{e}_{i+1}$ for $1 \leq i \leq n-1$, and let D_i be the green face of K with $e_i \subseteq \partial(D_i)$.

Lemma 4.4.6. *Let* $\Gamma_A \in \mathcal{T}$ *contain a red blob* B *and a well-contained sub-diagram* K *given by* B *and* ρ. *Let* v_1 *and* v_2 *be the corners of* K, *and let* $n = |p_1|$.

Suppose that $\partial(D) \cap \rho$ is a single consolidated edge for all green faces $D \subseteq K$. Then

$$1/2 + \sum_{i=1}^{2} \chi(v_i, \Gamma_A) = -n \cdot (\chi(B, \Gamma_A) + 1/2). \tag{4.1}$$

Proof. First note that by Axiom T_6, $\kappa_{\Gamma_A}(F) = 0$ for all green faces $F \subseteq K^{\Gamma_A}$ and all internal green face of Γ_A contain a boundary edge, hence $\kappa(K^{\Gamma_A}) = 0$; and as K is well-connected, all internal green faces of K have an edge on p_2. We will prove the lemma by deriving a closed-form expression of $\zeta(x, K^{\Gamma_A})$ for each each $x \in \Gamma_A$ with $\zeta(x, K^{\Gamma_A}) \neq 0$ and $x \notin K$; by finding for each $x \in K$, an $m \in \mathbb{Z}$ such that $\zeta(x, K) - \zeta(x, K^{\Gamma_A}) = m$; and by showing that if $x \subseteq K$ is a red blob or a vertex, then $\kappa_K(x) = 0$: from which we have $\kappa(K) = N + \sum_{x \in K} \zeta(x, K)$, where N is the number of internal green faces of K, and so

$$\begin{aligned}\kappa(K) - \kappa(K^{\Gamma_A}) &= N + \sum_{x \in K} \zeta(x, K) - N - \sum_{x \in \Gamma_A} \zeta(x, K^{\Gamma_A}) \\ &= \sum_{x \in K} \zeta(x, K) - \sum_{x \in \Gamma_A} \zeta(x, K^{\Gamma_A}).\end{aligned}$$

Suppose that $B' \subseteq \Gamma_A$ is a red blob with $\zeta(B', K^{\Gamma_A}) \neq 0$. Then either $B' = B$, or $B' \subseteq K$; and if $B' = B$, then B' gives each D_i curvature $\chi(B, \Gamma_A)$ across e_i, hence $\chi(B', K^{\Gamma_A}) = n \cdot \chi(B, \Gamma_A)$. Now suppose that $B' \subseteq K$ is a red blob. Since $N \geq 1$, B' is edge-incident with some internal green face of K, so $\kappa_K(B) = 0$. Also, all internal green faces of Γ_A edge-incident with B' are contained in K, hence $\zeta(B', K) - \zeta(B', K^{\Gamma_A}) = 0$.

If $v \in \Gamma_A$ is a vertex with $\zeta(v, K^{\Gamma_A}) \neq 0$, then $v \in K$. So for each vertex $v \in K$, let us find $m \in \mathbb{Z}$ such that $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = m$, and show $\kappa_K(v) = 0$. First suppose that $v \in K \setminus p_1$. Then all internal green faces of Γ_A incident with v are contained in K and $\delta_G(v, K) = \delta_G(v, \Gamma_A) \geq 2$, hence by Lemma 2.6.9 $\kappa_K(v) = 0$, and $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = 0$. Next suppose that $v \in \{w_i\}_{i=1}^{n-1}$. Since $\partial(B)$ does not pass more than once through any vertex of $p_1 \setminus p_2$, we deduce that p_1 is a sub-path of $\partial(B)$, so all green faces of Γ_A incident with v are contained in K. Now $v \in \partial(K) \setminus \partial(\Gamma_A)$, hence $\delta_G(v, K) = \delta_G(v, \Gamma_A) + 1 \geq 3$, so $\kappa_K(v) = 0$ and $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = -1/2$.

Finally, suppose that $v \in \{v_1, v_2\}$. First note that as K is well-connected, $v_1 \neq v_2$ if and only if K is simply-connected, and if $v_1 = v = v_2$, then K is annular and v is incident precisely twice with external faces of K. Now as $\partial(D) \cap \rho$ is a single consolidated edge for all green faces $D \subseteq K$ and D has an edge on p_2, we have that $\{D_1, D_n\}$ is the set of all internal green face of K incident with v_1 or v_2. Moreover, D_1 and D_n are incident precisely once with v_1 and v_2 respectively unless $v_1 = v_2$ and $D_1 = D_n$, in which case D_1 is incident precisely twice with v_1; and if $v_1 = v_2$ and $D_1 \neq D_n$, then v_1 is incident with both D_1 and D_n. Hence $\delta_G(v, K) \geq 2$ and $\kappa_K(v) = 0$. Moreover, if $v_1 \neq v_2$ then $\delta_G(v, K) = 2$, so $\chi(v, K) = 0 = \zeta(v, K)$; and $\zeta(v, K^{\Gamma_A}) = \chi(v, \Gamma_A)$. Otherwise, $\zeta(v, K^{\Gamma_A}) = 2 \cdot \chi(v, \Gamma_A) = \sum_{i=1}^{2} \chi(v_i, \Gamma_A)$ and $\zeta(v, K) = 2 \cdot \chi(v, K)$; and as v is incident precisely twice with external faces of K and v is incident with D_1, by Lemma 2.6.9 we have $\chi(v, K) = -1/2$, so $\zeta(v, K) = -1$.

Suppose that K is simply-connected, so $v_1 \neq v_2$. By Proposition 2.6.6 we have $\kappa(K) = 1$. Hence by using the observations from the previous 4 paragraphs we deduce that

$$\begin{aligned} 1 = \kappa(K) - \kappa(K^{\Gamma_A}) &= -(n-1)\cdot 1/2 - n\cdot\chi(B,\Gamma_A) - \sum_{i=1}^{2}\zeta(v_i, K^{\Gamma_A}) \\ &= -n\cdot(\chi(B,\Gamma_A)+1/2) + 1/2 - \sum_{i=1}^{2}\chi(v_i,\Gamma_A). \end{aligned}$$

Hence $1/2 + \sum_{i=1}^{2}\chi(v_i,\Gamma_A) = -n\cdot(\chi(B,\Gamma_A)+1/2)$, and the result follows.

Now suppose that K is not simply-connected. So $v_1 = v_2$ and K is annular, and by Proposition 2.6.6 we have $\kappa(K) = 0$. Hence by the first 4 paragraphs we have

$$\begin{aligned} 0 = \kappa(K) - \kappa(K^{\Gamma_A}) &= -(n-1)\cdot 1/2 + \zeta(v_1,K) - n\cdot\chi(B,\Gamma_A) - \zeta(v_1,K^{\Gamma_A}) \\ &= -n\cdot(\chi(B,\Gamma_A)+1/2) + 1/2 + \zeta(v_1,K) - \zeta(v_1,K^{\Gamma_A}) \\ &= -n\cdot(\chi(B,\Gamma_A)+1/2) + 1/2 - 1 - \sum_{i=1}^{2}\chi(v_i,\Gamma_A) \\ &= -n\cdot(\chi(B,\Gamma_A)+1/2) - 1/2 - \sum_{i=1}^{2}\chi(v_i,\Gamma_A). \end{aligned}$$

So $1/2 + \sum_{i=1}^{2}\chi(v_i,\Gamma_A) = -n\cdot(\chi(B,\Gamma_A)+1/2)$. ■

Lemma 4.4.7. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B and a well-contained sub-diagram K given by B and ρ, and assume that $\partial(D)\cap\rho$ is a single consolidated edge for all green faces $D \subseteq K$. Then $|p_1| \geq 2$.*

Proof. Suppose for a contradiction that $|p_1| = 1$. By assumption, $\partial(D_1)\cap\rho$ is a single consolidated edge, so $\partial(D_1) = e \cup p_1$, where e is the boundary consolidated edge of D_1. By Axiom T_1, the label R of D_1 does not have more than $1/2$ of its length on ρ. So $|R| = 2$, which contradicts Assumption 2.3.15 that no $R \in \mathcal{R}$ has $|R| \in \{1,2\}$. ■

Lemma 4.4.8. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B and a well-contained sub-diagram K given by B and ρ. Let v_1 and v_2 be the corners of K.*

Assume that $v_1, v_2 \notin \omega\cap\tau$, and that $\partial(D)\cap\rho$ is a single consolidated edge for all green faces $D \subseteq K$. Then for $i \in \{1,2\}$: if $\delta_G(v_i,\Gamma_A) \geq 3$, then $\delta_G(v_{3-i},\Gamma_A) \geq 3$.

Proof. By symmetry, it suffices to deduce a contradiction when $\delta_G(v_1,\Gamma_A) \geq 3$ and $\delta_G(v_2,\Gamma_A) = 2$. Since $\delta_G(v_2,\Gamma_A) = 2$, we have $\chi(v_2,\Gamma_A) = 0$. Hence by Lemma 4.4.6

$$1/2 + \chi(v_1,\Gamma_A) = -n\cdot(\chi(B,\Gamma_A)+1/2).$$

By Lemma 3.1.9 we have $\chi(v_1,\Gamma_A) > -1/2$, hence $\chi(B,\Gamma_A) < -1/2$. By Axiom T_6 all internal green faces F of Γ_A have $\kappa_{\Gamma_A}(F) = 0$, hence B is edge-incident at most once with each such F.

Note that by assumption $v_1 \neq v_2$, so as $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of K and by Axiom T_6, D has an edge on p_2, it follows that D_1 is the only internal green face of K incident with v_1. By Lemma 4.4.7, $n \geq 2$, so there is a vertex v_3 incident with e_1 such that $v_3 \neq v_1$, and $v_3 \in p_2 \setminus p_1$ implies that v_3 is interior. Hence $\delta_G(v_1, \Gamma_A) \geq 3$ implies $B \in \mathcal{S}_{D_1} \setminus \mathcal{B}_{D_1}$. Let $e \subseteq p_2$ be the boundary consolidated edge of D_1 with $|e| \geq 1$. As $v_1 \neq v_2$, there is a vertex v_4 incident with e such that $v_4 \neq v_1$. By Lemma 4.3.4, D_1 receives curvature of at most $-1/4$ from v_4, or a red blob $B_4 \in \mathcal{B}_{D_1}$ incident with v_4. Combining this curvature with $\chi(v_1, D_1, \Gamma_A) \leq -1/4$ and $\chi(B, D_1, \Gamma_A) < -1/2$ we have $\kappa_{\Gamma_A}(D_1) < 0$, a contradiction. ∎

Lemma 4.4.9. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B and a well-contained sub-diagram K given by B and ρ. Let v_1 and v_2 be the corners of K.*

Assume that $v_1, v_2 \notin \omega \cap \tau$, and that $\partial(D) \cap \rho$ is a single consolidated edge for all green faces $D \subseteq K$. If $\delta_G(v_1, \Gamma_A) = 2 = \delta_G(v_2, \Gamma_A)$, then $|p_1| = 2$ and $\chi(B, \Gamma_A) = -3/4$.

Proof. Since $\delta_G(v_i, \Gamma_A) = 2$ for $i \in \{1, 2\}$, we have $\chi(v_i, \Gamma_A) = 0$. Hence Lemma 4.4.6 shows that

$$1/2 = -n \cdot (\chi(B, \Gamma_A) + 1/2), \tag{4.2}$$

so $\chi(B, \Gamma_A) < -1/2$. Therefore, by Axiom T_6, B is edge-incident with each internal green face of Γ_A at most once.

By Lemma 4.4.7 we have $n \geq 2$. Suppose that $n \geq 3$. Since $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of K and by Axiom T_6, D has an edge on p_2, we deduce that $\{D_1, D_n\}$ is the set of all internal green face of K incident with v_1 or v_2. Hence there is $1 < i < n$ such that $v_1, v_2 \notin \partial(D_i)$. This means $B \notin \mathcal{B}_{D_i}$. But $\chi(B, D_i, \Gamma_A) < -1/2$, contradicting Part 4 of Lemma 4.3.4.

Hence $n = 2$. Then from (4.2) we have $\chi(B, \Gamma_A) = -3/4$. ∎

4.5 Intersection of boundaries of internal green faces with $\partial(\Gamma_A)$

Let F be an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. Axiom T_6 motivates us to study $\partial(F) \cap \rho$. The main result (see Theorem 4.5.13) in this section gives us the complete description of $\partial(F) \cap \rho$.

Let us first prove two useful results when Γ_A contains an island (see Definition 2.5.10) that we shall use in the proof of Theorem 4.5.13. Recall Definition 4.1.2 of the set $\mathcal{U}$, and Definitions 4.3.1 and 4.3.3 of the (boundary) curvature neighbourhood of an internal green face.

Proposition 4.5.1. *Let $\Gamma_A \in \mathcal{T} \setminus \mathcal{U}$ contain an island E, with endpoints v_1 and v_2. Then $\Gamma_A = E$, and E contains a green face incident with v_1 and v_2.*

Proof. By Lemma 4.2.4 all red blobs of Γ_A are good (see Definition 4.1.1), so by Lemma 4.2.5 they are all simply-connected. So as $\Gamma_A \in \mathcal{T} \setminus \mathcal{U}$ and by Lemma 3.1.10 Γ_A is a union of islands and bridges, we can without loss of generality assume that there is an internal green face $F \subseteq E$ that is curvature incident more than once with some element of $\mathcal{S}_F$.

By Proposition 4.2.8 this element is a vertex v, and by Lemma 4.2.3 $\Gamma_A \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2. Since each Γ_i contains one of the boundaries of Γ_A, we have $v_1 = v = v_2$, so $\Gamma_A = E$. ■

Lemma 4.5.2. *Let $\Gamma_A \in \mathcal{T}$. Assume that Γ_A contains an island E with endpoints v_1 and v_2, bounded by the closed path $\omega_1\tau_1$, and with an internal green face F_1 such that $v_1, v_2 \in \partial(F_1)$. Then one of the following statements holds.*

1. *$E = \overline{F_1}$, hence $\partial(F_1) \cap \omega = \omega_1$ and $\partial(F_1) \cap \tau = \tau_1$.*

2. *There is an internal green face $F_2 \neq F_1$ such that $E = \overline{F_1} \cup \overline{F_2}$, and the following statements hold.*

   (i) $\partial(F_1) \cap \partial(F_2)$ is a single consolidated edge e with $\partial(e) = \{v_1\} \cup \{v_2\}$.

   (ii) For $i \in \{1, 2\}$ we have $\partial(F_i) \cap \partial(E) \in \{\omega_1, \tau_1\}$ and $\partial(F_i) = (\partial(F_i) \cap \partial(E)) \cup e$.

Proof. By Lemma 2.6.9 we have $\chi(v_1, \Gamma_A) = -1/2 = \chi(v_2, \Gamma_A)$. So as $v_1, v_2 \in \partial(F_1)$ and by Axiom T_6 we have $\kappa_{\Gamma_A}(F_1) = 0$, it follows that $\mathcal{S}_{F_1} = \{v_1, v_2\}$. Hence either $E = \overline{F_1}$ and Case 1 holds, or there is an internal green face F_2 such that $\partial(F_1) \cap \partial(F_2)$ is a single consolidated edge e with $\partial(e) = \{v_1\} \cup \{v_2\}$. Similarly as for F_1 we deduce that $\mathcal{S}_{F_2} = \{v_1, v_2\}$. Hence as by Axiom T_6, F_2 contains a boundary edge, we deduce that $E = \overline{F_1} \cup \overline{F_2}$ and Case 2 holds. ■

Recall Definition 4.4.3 of a well-contained sub-diagram given by an internal green face and $\rho \in \{\omega, \tau\}$.

Lemma 4.5.3. *Let F be a tricky green face for Γ_A with respect to ρ. Then there exists a well-contained sub-diagram of Γ_A given by F and ρ.*

Proof. We first show that there is a sub-diagram K of Γ_A well-bounded by F and ρ (see Definition 4.4.2) that is well-connected, and such that $\partial(F)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho$.

Assume first that $\partial(F)$ does not pass more than once through any vertex. Then $\partial(F)$ is a simple closed path; and since F is tricky, $\partial(F) \cap \rho \neq \emptyset$ and $\partial(F) \cap \rho$ is not a single consolidated edge, so $\partial(F) \cap \rho$ is disconnected. Hence there is a simply-connected sub-diagram that is well-bounded by F and ρ.

Suppose instead that $\partial(F)$ passes more than once through some vertex v. By Lemma 4.2.3 $\partial(F)$ passes through v twice, and $\Gamma_A \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2 with $\rho \subseteq \Gamma_1$, where Γ_1 contains precisely one vertex w incident more than once with F. Since $\partial(F) \cap \rho \neq \emptyset$ and $\partial(F) \cap \rho$ is not a single consolidated edge, there is a path $p \in \partial(\Gamma_1) \cap \partial(F)$ with $|p| \geq 1$ intersecting ρ only at its endpoints, hence by Lemma 2.1.8 there is a simple path $p_1 \subseteq \partial(\Gamma_1) \cap \partial(F)$ with $|p_1| \geq 1$ intersecting ρ only at its

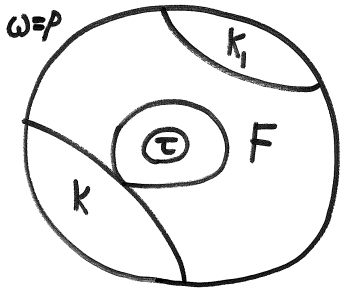

Figure 4.4: A case where $\partial(F)$ passes more than once some vertex of $\partial(K) \setminus \rho$, see the proof of Lemma 4.5.3.

endpoints. Now there is a sub-path p_2 of ρ with the same endpoints as p_1, and the sub-diagram K bounded by the closed path p_1p_2 satisfies $K \subseteq \Gamma_1$ (so in particular K° is contained in some bounded component of $\mathbb{R}^2 \setminus p_1p_2$). Hence K is well-bounded by F and ρ, and K is well-connected. Suppose that $w \in \partial(K) \setminus \rho$. Then as w is the only vertex of Γ_1 through which $\partial(F)$ passes more than once and $\partial(F) \cap \rho$ is not a single consolidated edge, there is a simple path $p' \in \partial(\Gamma_1) \cap \partial(F)$ with $|p'| \geq 1$ distinct from p_1 intersecting ρ only at its endpoints. It follows that there is a sub-diagram $K_1 \subseteq \Gamma_1$ well-bounded by F and ρ with $\partial(K_1) \setminus \rho = p'$ (see Figure 4.4), and such that $\partial(F)$ does not pass more than once through any vertex of $\partial(K_1) \setminus \rho$. Since $K_1 \subseteq \Gamma_1$, we deduce that K_1 is well-connected. Hence we showed that there is a sub-diagram K of Γ_A well-bounded by F and ρ that is well-connected, and such that $\partial(F)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho$.

It remains to show that K is non-trivial. Suppose not. Then as K is well-connected, K is a simply-connected red blob. As $\partial(F)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho$, from Lemma 3.2.4 we have $|\partial(K)| \leq 2$. Hence by Lemma 3.2.6 $\mathbf{Area}(K) \leq 0$, a contradiction. ■

Using Lemma 4.5.3 we can now present the following definition.

Definition 4.5.4. Let F be a tricky green face for Γ_A with respect to ρ, and let X be such that if Γ_A is island-free, then $X = \Gamma_A$, else X is the island containing F. By Lemma 4.5.3 there is a well-contained sub-diagram of Γ_A given by F and ρ. Then F is called a *minimal tricky green face for* Γ_A *with respect to* ρ if there is a well-contained sub-diagram K of Γ_A given by F and ρ of minimal area among all well-contained sub-diagrams given by tricky green faces $F' \subseteq X$ (for Γ_A with respect to ρ) and ρ.

K is called a *minimal well-contained sub-diagram given by* F *and* ρ.

Before proving Theorem 4.5.13, let us prove several useful results that consider minimal well-contained sub-diagrams given by tricky green faces and ρ. These sub-diagrams play a central role in the proof of Theorem 4.5.13.

Lemma 4.5.5. *Let F be a minimal tricky green face for Γ_A with respect to ρ, and let K be a minimal well-contained sub-diagram given by F and ρ. Then $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of K.*

Proof. Suppose that there is an internal green face D of K such that $\partial(D) \cap \rho$ is not a single consolidated edge. By Axiom T_6 we have $\partial(D) \cap \rho \neq \emptyset$, hence D is tricky for Γ_A with respect to ρ. Therefore, by Lemma 4.5.3 there is a well-contained sub-diagram L given by D and ρ. Since D is contained in K, we have $L \subseteq K$. But $D \not\subseteq L$, hence $\mathbf{Area}(L) < \mathbf{Area}(K)$, a contradiction. ■

Definition 4.5.6. Let K be a sub-diagram of an annular diagram Γ_A well-bounded by an internal green face F and ρ. A *curvature corner of* K is an $x \in K$ satisfying one of the following two statements.

1. x is a corner of K incident with some internal green face of K.
2. x is a red blob in $\mathcal{B}_F \cap K$ edge-incident with F at some corner of K.

Lemma 4.5.7. *Let $\Gamma_A \in \mathcal{T}$ contain a well-contained sub-diagram K given by an internal green face F and ρ, and let v_1 and v_2 be the corners of K. Then both of the following two statements hold.*

1. If $v \in \{v_1, v_2\}$ is a curvature corner of K, then $v \in \mathcal{B}_F \cap K$.

2. The set S of curvature corners of K is one of the following:

(i) $S = \{v_1, v_2\}$;

(ii) $S = \{v, B\}$, where $v \in \{v_1, v_2\}$, and B is a red blob;

(iii) $S = \{B_1, B_2\}$, where B_1 and B_2 are both red blobs.

Proof. Since K is well-contained, by Definition 4.4.3 K is well-connected and $\mathbf{CArea}(K) \geq (1, 0)$, hence by Axiom T_6 we have $|\partial(K) \cap \rho| \geq 1$.

Part 1. By Definition 4.5.6 v is incident with some internal green face of K, so as v is incident with F and $v \in \rho$, we have $\delta_G(v) \geq 3$, hence $v \in \mathcal{B}_F \cap K$.

Part 2. Suppose first that $v_1 \neq v_2$, and let $v' \in \{v_1, v_2\}$. If $v' \notin S$, then no internal green face of K is incident with v'. So as K is well-connected, there is a red blob $B \in \mathcal{B}_F \cap K$ edge-incident with F at v', and there is precisely one such blob. Otherwise, some internal green face of K is incident with v', and hence no blob $B \subseteq \Gamma_A$ is edge-incident with F at v' and satisfies $B \in \mathcal{B}_F \cap K$. So Part 2 follows.

Suppose instead that $v_1 = v_2$. If $v_1 \notin S$, then since K is well-connected, there is a red blob $B \in \mathcal{B}_F \cap K$ edge-incident with F at v_1, and there are at most two such blobs. Otherwise, there is at most one blob in S. ■

Throughout the rest of this section if we say that the intersection $\partial(B) \cap \partial(F)$ of a red blob B and a green face F is a single edge, we mean that $\partial(B) \cap \partial(F)$ is a single consolidated edge with $|\partial(B) \cap \partial(F)| = 1$.

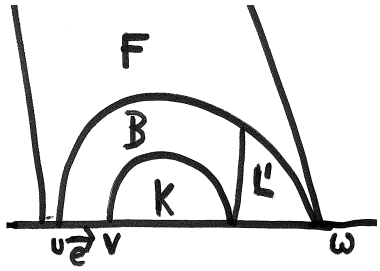

Figure 4.5: Depiction of the sub-diagram K well-bounded by B and ρ, where $L = B \cup K \cup L'$ and $\rho = \omega$, see the proof of Lemma 4.5.9.

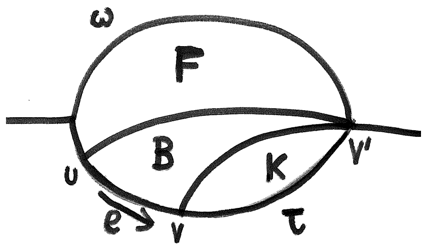

Figure 4.6: Depiction of the sub-diagram K well-bounded by B and ρ, where $L = B \cup K$, $\rho = \tau$ and $v' \in \omega \cap \tau$, see the proof of Lemma 4.5.9.

Lemma 4.5.8. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B contained in a sub-diagram K of Γ_A and edge-incident with an internal green face $F \in \Gamma_A \setminus K$ such that $\partial(F) \cap \partial(K)$ is a sub-path of $\partial(F)$. Suppose that either*

(a) K is simply-connected, or

(b) is well-bounded by F and $\rho \in \{\omega, \tau\}$ and well-connected.

Then $\partial(B) \cap \partial(F)$ is a single edge.

Proof. Suppose that there is a bounded component C of $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ such that $\partial(C)$ is of the form $p_1 p_2$, where $p_1 \subseteq \partial(B)$ and $|p_1| \geq 1$, and $p_2 \subseteq \partial(F)$. As K is either simply-connected, or is well-bounded by F and ρ and well-connected, we deduce that $\overline{C}$ contains no edge of $\partial(\Gamma_A)$, hence $\overline{C}$ contradicts Lemma 4.2.2. So no such component C exists.

Therefore, if $\partial(B) \cap \partial(F)$ is not a single edge, then $\partial(F) \cap \partial(B)$ is a single path p with $|p| \geq 2$; and as $\partial(F) \cap \partial(K)$ is a sub-path of $\partial(F)$, by Lemma 4.2.6 $\partial(B)$ passes more than once through some vertex of p common to two edges of p. In particular, by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams. But this is impossible since $B \subseteq K$ and K satisfies at least one of the Assumptions (a)-(b). ■

Lemma 4.5.9. *Let F be a minimal tricky green face for Γ_A with respect to ρ, let L be a minimal well-contained sub-diagram of Γ_A given by F and ρ, and let x be a curvature corner of L.*

If x is a red blob B, then B is simply-connected, $\partial(B) \cap \partial(F)$ is a single edge, $\partial(B)$ contains no edge on the boundary ρ' of Γ_A distinct from ρ, and $\partial(B) \cap \rho$ is a single edge.

Proof. Suppose that $x = B$. By Definition 4.4.3 L is well-bounded by F and ρ and well-connected, and $\partial(F)$ does not pass more than one through any vertex of $\partial(L) \setminus \rho$, so $\partial(F) \cap \partial(L)$ is a sub-path of $\partial(F)$. Hence by Lemma 4.5.8 $\partial(B) \cap \partial(F)$ is a single edge, and as $B \subseteq L$, we deduce that B is good: so by Lemma 4.2.5 B is simply-connected; and $\partial(B)$ contains no edge on ρ'.

It remains to prove that $\partial(B) \cap \rho$ is a single edge. Suppose not. Let u be a corner of L with $u \in B$, and let e be an edge of $\partial(B) \cap \rho$ with $u \in \overline{e}$. By Definition 4.4.3 we have $\mathbf{CArea}(L) \geq (1, 0)$, hence as L is well-connected, by Axiom T_6, $\partial(L) \cap \rho$ contains a green edge. So we can let v be the endpoint of e distinct from u. By Lemma 3.2.4 all consolidated edges between B and the external faces of Γ_A have length at most one, hence there is a path $p \in \Gamma_A^1 \cap \partial(B)$ with $|p| \geq 1$, intersecting ρ only at its endpoints $v, v' \in \rho$. By Lemma 2.1.8 there is a simple path $p_1 \subseteq \Gamma_A^1 \cap \partial(B)$ with $|p_1| \geq 1$ intersecting ρ only at v and v'. Hence there is a sub-diagram K well-bounded by B and ρ, and with corners v and v' (see Figure 4.5).

We show that K is well-contained. Since L is well-connected, it follows that K is simply-connected. Hence as $|\partial(B) \cap \partial(K)| \geq 1$, we have $\mathbf{CArea}(K) \geq (1, 0)$. If $\partial(B)$ passes more than once through some vertex w, then by Lemma 4.2.3 w is unique, and by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Hence as $B \subseteq L$, we deduce that w is a corner of L, so $w \in \rho$. Thus, $\partial(B)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho$, so K is well-contained.

As K is simply-connected, Axiom T_6 and $\mathbf{CArea}(K) \geq (1, 0)$ imply that $|\partial(K) \cap \rho| \geq 1$. So $v \notin \omega \cap \tau$ since $v \in \overline{e}$. As by Lemma 4.5.5 $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of L, the same holds for K. In particular, since by Axiom T_6 they all have an edge on $\partial(K) \cap \rho$, we have $\delta_G(v) = 2$. Suppose first that $v' \in \omega \cap \tau$. Then $\partial(L) \cap \rho = \overline{e} \cup (\partial(K) \cap \rho)$. Hence v' is a corner of L, and e is the only boundary edge of B (see Figure 4.6). Therefore, by Lemma 3.2.8 we have $\chi(B, \Gamma_A) > -1/2$. Further, by Lemmas 2.6.10 and 3.1.9 we have $\chi(v, \Gamma_A) = 0$ and $\chi(v', \Gamma_A) = -1/2$, so applying Lemma 4.4.6 shows that

$$0 = 1/2 + \chi(v', \Gamma_A) = -|\partial(B) \cap \partial(K)| \cdot (\chi(B, \Gamma_A) + 1/2) < 0,$$

a contradiction.

Suppose instead that $v' \notin \omega \cap \tau$. Then K satisfies assumptions of Lemmas 4.4.8-4.4.9. Since $\delta_G(v) = 2$, by Lemma 4.4.8 $\delta_G(v') = 2$, hence Lemma 4.4.9 gives $\chi(B, \Gamma_A) = -3/4$. Now as $\delta_G(v) = 2 = \delta_G(v')$ and $\partial(F) \cap \rho$ is not a single consolidated edge, by Part 1 of Lemma 4.3.4 there is $y \notin K$ with $y \in \mathcal{B}_F$. By Part 2 of Lemma 4.3.4 we have $\chi(y, F, \Gamma_A) \leq -1/4$. Since $\kappa_{\Gamma_A}(F) = 0$ and $\chi(B, F, \Gamma_A) = -3/4$, we deduce that $\mathcal{S}_F = \{B, y\} = \mathcal{B}_F$, that B is edge-incident once with F, and that y is either a red triangle edge-incident once with F that contains one boundary edge, or a vertex with $\delta_G(y) = 3$ incident once with F and such that

$y \notin \omega \cap \tau$. In particular, B is the only element of $\mathcal{S}_F$ contained in K.

Suppose first that y is a red triangle. Then all vertices of F have green degree 2, hence as y is edge-incident once with F and $|\partial(y) \cap \partial(\Gamma_A)| = 1$, we deduce that F is contained in a sub-diagram K_1 of Γ_A well-bounded by y and ρ such that $\partial(K_1) = \overline{e'}p'$, where e' is the common edge of F and y, p' is a sub-path of ρ, and $\rho \setminus K_1$ is the boundary edge b of y. But then $e'b$ is a sub-path of $\partial(y)$, so y is not simply-connected, a contradiction as y is a red triangle.

Assume instead that y is a vertex. Then $\mathcal{S}_F$ contains no red blob outside K, hence as y is incident once with F, it follows that $\mathcal{S}_F$ contains at least two distinct vertices, contradicting $\mathcal{S}_F = \{B, y\}$. ■

The following configuration appears multiple times in our proofs.

Definition 4.5.10. Let F be an internal green face of a coloured annular diagram Γ_A, and let $\rho \in \{\omega, \tau\}$. Assume that for each $i \in \{1, 2\}$, there is x_i such that x_i is either a simply-connected red blob edge-incident with F containing at most two boundary edges, or a boundary vertex incident with F and with $x_i \in \rho$.

Then F, x_1, x_2 are called a *neighbourhood of* ρ if $\Gamma_A \setminus (\overline{F} \cup x_1 \cup x_2)^\circ$ contains a sub-diagram K of Γ_A that satisfies the following 3 conditions.

1. $\mathbf{CArea}(K) \geq (1, 0)$.

2. $\partial(K)$ is a closed path in Γ_A^1 of the form $p_1 p_3 p_2 p_4$ satisfying the following conditions.

   (i) p_4 is a sub-path of ρ.

   (ii) If x_i is a vertex, then $p_i = x_i$ and x_i is incident with some internal green face of K, else p_i is a sub-path of $\partial(x_i)$ with $|p_i| \geq 1$ and with the following properties:

   (a) for the vertex $v \in p_i \cap p_4$: the external face with boundary ρ is the only green face of Γ_A incident with v that is not contained in K, and is incident precisely once with v;

   (b) For all vertices $v \in p_i^\circ$: all green faces of Γ_A incident with v are contained in K.

   (iii) p_3 is a sub-path of $\partial(F)$ such that if x_1 and x_2 are both vertices, then $|p_3| \geq 1$, and for all vertices $v \in p_3 \setminus p_4$ the only green face of Γ_A incident with v that is not contained in K is F, which is incident precisely once with v.

3. If K is not simply-connected then x_1 and x_2 are both vertices, and K is annular and consisting of a single island with boundaries p_3 and p_4.

We call K *a sub-diagram of* Γ_A *bounded by* F, x_1, x_2 *and* ρ.

By Definition 4.5.10 it follows that $\kappa(K)$ and $\kappa(K^{\Gamma_A})$ are defined (see Definition 4.4.4). Even though Definition 4.5.10 may seem complicated, the following lemma provides a useful restriction.

Lemma 4.5.11. *Let $\Gamma_A \in \mathcal{T}$ contain a neighbourhood F, x_1, x_2 of $\rho \in \{\omega, \tau\}$, and let K be a sub-diagram of Γ_A bounded by F, x_1, x_2 and ρ. Assume that $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of K. Then each of x_1 and x_2 is either a vertex of $\omega \cap \tau$, or is a red blob with two boundary edges.*

Proof. By Axiom T_6 all green faces $F \subseteq K^{\Gamma_A}$ satisfy $\kappa_{\Gamma_A}(F) = 0$ and all internal green faces of Γ_A contain a boundary edge, hence $\kappa(K^{\Gamma_A}) = 0$; and Condition 3 of Definition 4.5.10 implies that all internal green faces of K have an edge on p_4, so as $\mathbf{CArea}(K) \geq (1, 0)$, we have $|p_4| \geq 1$. The strategy of the proof is deriving a closed-form expression of $\zeta(x, K^{\Gamma_A})$ (see Definition 4.4.4) for each $x \in \Gamma_A$ with $\zeta(x, K^{\Gamma_A}) \neq 0$ and $x \notin K$; finding for each $x \in K$, an $m \in \mathbb{Z}$ such that $\zeta(x, K) - \zeta(x, K^{\Gamma_A}) \leq m$; and showing that if $x \subseteq K$ is a red blob or a vertex, then $\kappa_K(x) = 0$, from which it follows that $\kappa(K) = N + \sum_{x \in K} \zeta(x, K)$, where N is the number of internal green faces of K, and hence

$$\begin{aligned}\kappa(K) - \kappa(K^{\Gamma_A}) &= N + \sum_{x \in K} \zeta(x, K) - N - \sum_{x \in \Gamma_A} \zeta(x, K^{\Gamma_A}) \\ &= \sum_{x \in K} \zeta(x, K) - \sum_{x \in \Gamma_A} \zeta(x, K^{\Gamma_A}).\end{aligned}$$

Using this we show that $\chi(x_i, \Gamma_A) \leq -1/2$ for $1 \leq i \leq 2$, and then apply Lemmas 3.1.9 and 3.2.8 to deduce the lemma.

Suppose that $B \subseteq \Gamma_A$ is a red blob with $\zeta(B, K^{\Gamma_A}) \neq 0$. Then either $B = x_i$ for some $i \in \{1, 2\}$, or $B \subseteq K$; and if $B = x_i$, then B gives curvature $\chi(x_i, \Gamma_A)$ across each edge of p_i to precisely one internal green face of K, so $\chi(B, K^{\Gamma_A}) = |p_i| \cdot \chi(x_i, \Gamma_A)$. Now suppose that $B \subseteq K$ is a red blob. Since $N \geq 1$, B is edge-incident with some internal green face of K, so $\kappa_K(B) = 0$, and note that $\zeta(B, K) = \zeta(B, \Gamma_A)$. If B has an edge on p_3, then by Condition (iii) of Definition 4.5.10, F is the only internal green face of Γ_A edge-incident with B that is not contained in K. Now B gives F some of its negative curvature in Γ_A, hence $\zeta(B, K) - \zeta(B, K^{\Gamma_A}) < 0$. Otherwise, as $\zeta(B, K) = \zeta(B, \Gamma_A)$, we have $\zeta(B, K) - \zeta(B, K^{\Gamma_A}) = 0$. Thus, $\zeta(B, K) - \zeta(B, K^{\Gamma_A}) \leq 0$.

If $v \in \Gamma_A$ is a vertex with $\zeta(v, K^{\Gamma_A}) \neq 0$, then $v \in K$. So for each vertex $v \in K$, let us find $m \in \mathbb{Z}$ such that $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) \leq m$, and show $\kappa_K(v) = 0$. First suppose that $v \in K \setminus (p_1 \cup p_2 \cup p_3)$. Then $\delta_G(v, K) = \delta_G(v, \Gamma_A) \geq 2$, and all internal green faces of Γ_A incident with v are contained in K. Hence by Lemma 2.6.9 we have $\kappa_K(v) = 0$, and $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = 0$. Next suppose that $v \in p_3 \setminus p_4$. Then $v \in \partial(K) \setminus \partial(\Gamma_A)$, so from Condition (iii) of Definition 4.5.10 it follows that $\delta_G(v, K) = \delta_G(v, \Gamma_A) \geq 2$ and v is incident with some internal green face of K, hence $\kappa_K(v) = 0$ and $\zeta(v, K) - \zeta(v, \Gamma_A) = 0$. Furthermore, if $\delta_G(v, \Gamma_A) \geq 3$, then v gives F some of its negative curvature in Γ_A, so $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) < 0$, and hence $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) \leq 0$.

Next suppose that x_i is a red blob for some $i \in \{1, 2\}$, and suppose that $v \in p_i \cap p_4$. Since $v \in \partial(K)$, from Definition 4.5.10 (ii) we have $\delta_G(v, K) = \delta_G(v, \Gamma_A) \geq 2$ and v is incident with some internal green face of K, hence $\kappa_K(v) = 0$ and $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = 0$. Now suppose that $v \in p_i^\circ$, and note that there are $|p_i| - 1$ vertices lying on p_i°. We have $v \in \partial(K)$, hence applying Definition 4.5.10 (ii) shows that $\delta_G(v, K) = \delta_G(v, \Gamma_A) + 1 \geq 3$, so $\kappa_K(v) = 0$ and $\zeta(v, K) - \zeta(v, K^{\Gamma_A}) = -1/2$.

Finally, suppose that x_i is a vertex for some $i \in \{1, 2\}$, and that $v = x_i$. First note that by Condition 3 of Definition 4.5.10, $x_1 = v = x_2$ if and only if K is not simply-connected; and if K is not simply-connected, then it is annular, and v is incident precisely twice with external faces of K. Now by Condition (ii) of Definition 4.5.10 v is incident with some internal green face of K, so $\delta_G(v, K) \geq 2$, and hence $\kappa_K(v) = 0$. By assumption, $\partial(D) \cap \rho$ is a single consolidated edge for all internal green faces D of K. Therefore, as they all have an edge on p_4, if $x_1 \neq x_2$, then there is precisely one internal green face D_i in K incident with x_i, and D_i is incident precisely once with x_i. Otherwise, x_1 is incident with at most two internal green faces $D_1, D_2 \subseteq K$; and x_1 is incident precisely once with both D_1 and D_2 unless $D_1 = D_2$, in which case x_1 is incident precisely twice with D_1. Hence if $x_1 \neq x_2$, then $\zeta(v, K^{\Gamma_A}) = \chi(v, \Gamma_A)$ and $\zeta(v, K) = \chi(v, K)$; and $\delta_G(v, K) = 2$, so $\chi(v, K) = 0 = \zeta(v, K)$. Otherwise, $\zeta(v, K^{\Gamma_A}) = 2 \cdot \chi(v, \Gamma_A)$ and $\zeta(v, K) = 2 \cdot \chi(v, K)$.

We are ready to prove the lemma. Suppose first that x_1 and x_2 are both vertices, and assume further that K is simply-connected. By Proposition 2.6.6 we have $\kappa(K) = 1$, and by the previous paragraph we have $x_1 \neq x_2$ and $\chi(x_i, K) = 0 = \zeta(x_i, K)$. Hence by the above observations we have

$$1 = \kappa(K) - \kappa(K^{\Gamma_A}) \leq -\sum_{i=1}^{2} \zeta(x_i, K^{\Gamma_A}) = -\sum_{i=1}^{2} \chi(x_i, \Gamma_A).$$

So $\sum_{i=1}^{2} \chi(x_i, \Gamma_A) \leq -1$. Thus, by Lemma 3.1.9 $x_1, x_2 \in \omega \cap \tau$.

Now assume that K is not simply-connected, so by the fifth paragraph $x_1 = x_2$; K is annular: so by Proposition 2.6.6 $\kappa(K) = 0$; and x_1 is incident precisely twice with external faces of K: hence as x_1 is incident with D_1, by Lemma 2.6.9 $\chi(x_1, K) = -1/2$, and so $\zeta(x_1, K) = -1$. It follows that

$$0 = \kappa(K) - \kappa(K^{\Gamma_A}) \leq \zeta(x_1, K) - \zeta(x_1, K^{\Gamma_A}) = -1 - (2 \cdot \chi(x_1, \Gamma_A)).$$

So $2 \cdot \chi(x_1, \Gamma_A) \leq -1$, which implies $\chi(x_1, \Gamma_A) \leq -1/2$, and the lemma follows from Lemma 3.1.9.

Next suppose that x_1 is a red blob, and x_2 is a vertex. Then by the fifth paragraph K is simply-connected, so $\kappa(K) = 1$. Hence letting $n = |p_1|$ and using the observations from the first 5 paragraphs we have

$$\begin{aligned}
1 = \kappa(K) - \kappa(K^{\Gamma_A}) &\leq -(n-1) \cdot 1/2 + \zeta(x_2, K) - n \cdot \chi(x_1, \Gamma_A) - \zeta(x_2, K^{\Gamma_A}) \\
&= -n \cdot (\chi(x_1, \Gamma_A) + 1/2) + 1/2 + \zeta(x_2, K) - \zeta(x_2, K^{\Gamma_A}) \\
&= -n \cdot (\chi(x_1, \Gamma_A) + 1/2) + 1/2 + 0 - \chi(x_2, \Gamma_A) \\
&= -n \cdot (\chi(x_1, \Gamma_A) + 1/2) + (1/2 - \chi(x_2, \Gamma_A)).
\end{aligned}$$

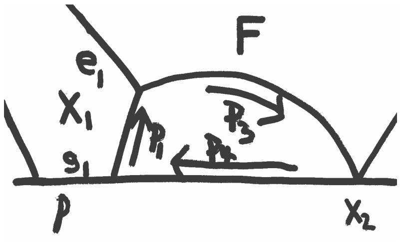

Figure 4.7: A case where x_1 is a red blob, and x_2 is a vertex, see Lemma 4.5.12.

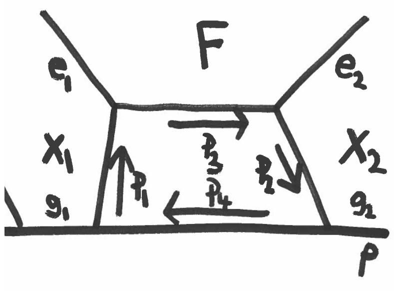

Figure 4.8: A case where x_1 and x_2 are both red blobs, see Lemma 4.5.12.

By Lemma 3.1.9 we have $\chi(x_2, \Gamma_A) \geq -1/2$, hence $-n \cdot (\chi(x_1, \Gamma_A) + 1/2) \geq 0$, which implies $\chi(x_1, \Gamma_A) \leq -1/2$. Since x_1 contains at most two boundary edges, by Lemma 3.2.8 $\chi(x_1, \Gamma_A) \geq -1/2$, so $\chi(x_1, \Gamma_A) = -1/2$, and therefore $\chi(x_2, \Gamma_A) = -1/2$. Hence x_1 contains two boundary edges, and $x_2 \in \omega \cap \tau$.

Finally, suppose that x_1 and x_2 are both red blobs. By the fifth paragraph K is simply-connected, so $\kappa(K) = 1$. Let $n = |p_1|$, $m = |p_2|$. By the first 5 paragraphs we have

$$\begin{aligned}
1 &= \kappa(K) - \kappa(K^{\Gamma_A}) \\
&\leq -(n-1) \cdot 1/2 - (m-1) \cdot 1/2 - n \cdot \chi(x_1, \Gamma_A) - m \cdot \chi(x_2, \Gamma_A) \\
&= 1 - n \cdot (1/2 + \chi(x_1, \Gamma_A)) - m \cdot (1/2 + \chi(x_2, \Gamma_A)).
\end{aligned}$$

Hence $0 \leq -n \cdot (1/2 + \chi(x_1, \Gamma_A)) - m \cdot (1/2 + \chi(x_2, \Gamma_A))$. Since by Lemma 3.2.8 $\chi(x_i, \Gamma_A) \geq -1/2$ for each $i \in \{1, 2\}$, we have $\chi(x_i, \Gamma_A) = -1/2$, and the lemma follows. ∎

In the statement of the next lemma we allow tracing boundaries of faces in both directions.

Lemma 4.5.12. *Let $\Gamma_A \in \mathcal{T}$. Assume that the following statements hold.*

1. *Γ_A contains an internal green face F, and for each $i \in \{1, 2\}$, it contains x_i such that x_i is either a simply-connected red blob containing at most two boundary edges, or a boundary vertex incident with F and with $x_i \in \rho \in \{\omega, \tau\}$.*

2. *There is a sub-diagram K that satisfies the following 3 conditions.*

   (i) $\mathbf{CArea}(K) \geq (1, 0)$.

(ii) *K is bounded by a closed $p \in \Gamma_A^1$ of the form $p_1p_3p_2p_4$ satisfying the following conditions (see Figures 4.7 and 4.8).*

(a) *p_4 is a sub-path of ρ.*

(b) *If x_i is a vertex, then $p_i = x_i$ and x_i is incident with some internal green face of K, else p_i is a sub-path of $\partial(x_i)$ with $|p_i| \geq 1$.*

(c) *p_3 is a sub-path of $\partial(F)$ such that if at most one of x_1 or x_2 is a red blob, then $|p_3| \geq 1$; and if $|p_3| < 1$, then F is incident precisely once with p_3.*

(d) *If x_i is a red blob for some $i \in \{1, 2\}$, then $\partial(x_i) \cap \rho$ contains an edge g_i such that g_ip_i is a sub-path of $\partial(x_i)$, and $\partial(F) \cap \partial(x_i)$ contains an edge e_i such that p_ie_i is a sub-path of $\partial(x_i)$, and e_ip_3 is a sub-path of $\partial(F)$. Moreover, the common endpoint of g_i and p_i (which lies on $p_i \cap p_4$) does not lie on $\omega \cap \tau$.*

(iii) *If K is not simply-connected then x_1 and x_2 are both vertices, and K is annular and consisting of a single island with boundaries p_3 and p_4.*

Then F, x_1, x_2 are a neighbourhood of ρ, and K is a sub-diagram bounded by F, x_1, x_2 and ρ.

Proof. Suppose first that x_1 and x_2 are both vertices. As p_3 is a sub-path of $\partial(F)$, for all vertices $v \in p_3 \setminus p_4$: F is the only green face of Γ_A incident with v that is not contained in K, and is incident precisely once with v. Hence the lemma follows.

Next assume that x_i is a red blob for exactly one $i \in \{1, 2\}$. As the vertex $v \in p_i \cap p_4$ satisfies $v \notin \omega \cap \tau$, and g_ip_i is a sub-path of $\partial(x_i)$, it follows that the external face E with boundary ρ is the only green face incident with v that is not contained in K. Furthermore, since by Axiom T_1, ω and τ are simple closed paths, E is incident precisely once with v. Note also that as p_i is a sub-path of $\partial(x_i)$, for all vertices $v \in p_i^\circ$: all green faces incident with v are contained in K. Similarly, as p_3 and e_ip_3 are sub-paths of $\partial(F)$, and p_ie_i is a sub-path of $\partial(x_i)$, it follows that for all vertices $v \in p_3 \setminus p_4$: the only green face of Γ_A incident with v that is not contained in K is F, which is incident precisely once with v. So the lemma follows again.

Finally, assume that x_i is a red blob for each $i \in \{1, 2\}$. Then similarly as in the previous case we deduce that p_1 and p_2 satisfy Condition (ii) of Definition 4.5.10. Now since p_3 is a sub-path of $\partial(F)$, and for each $i \in \{1, 2\}$, p_ie_i is a sub-path of $\partial(x_i)$, and e_ip_3 is a sub-path of $\partial(F)$, we deduce that for all vertices $v \in p_3$: F is the only green face incident with v and not contained in K, and is incident precisely once with v. Hence we are done. ∎

We can now present the main result of this section.

Theorem 4.5.13. *Let D be an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. Then one of the following statements holds.*

1. *$\partial(D) \cap \rho = \emptyset$.*

2. *$\partial(D) \cap \rho$ is a single consolidated edge.*

Figure 4.9: The sub-diagram K_1 bounded by $r_1r_3r_4$, see the proof of Theorem 4.5.13.

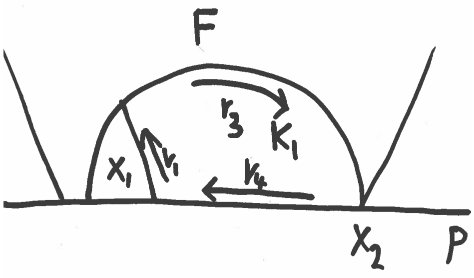

3. *$\partial(D)\cap\partial(\Gamma_A)$ is a single consolidated edge with endpoints v_1 and v_2 with $v_1, v_2 \in \omega\cap\tau$ and $\partial(D)\cap\rho = \{v_1, v_2\}$: hence D is contained in an island E with $\mathbf{CArea}(E) = (2, 0)$ that satisfies Lemma 4.5.2.*

Proof. Suppose that there is a green face $D \subseteq \Gamma_A$ such that $\partial(D) \cap \rho \neq \emptyset$ and $\partial(D) \cap \rho$ is not a single consolidated edge. Then D is tricky for Γ_A with respect to ρ (see Definition 4.4.1). By Lemma 4.5.3 there exists a well-contained sub-diagram of Γ_A given by D and ρ (see Definition 4.4.3). Hence we can let F be a minimal tricky green face for Γ_A with respect to ρ (see Definition 4.5.4), and let K be a minimal well-contained sub-diagram given by F and ρ. We show that Statement 3 of the theorem holds for F, and from that we deduce it also holds for D.

Let $p_1 = \partial(K) \cap \partial(F)$ and $p_2 = \partial(K) \cap \rho$. Since $\partial(F)$ does not pass more than once through any vertex of $p_1 \setminus p_2$, we have that p_1 is a sub-path of $\partial(F)$. By Lemma 4.5.5 $\partial(D') \cap \rho$ is a single consolidated edge for all internal green faces D' of K, and by Lemma 4.5.7 we can let x_1 and x_2 be (not necessarily distinct) curvature corners of K. Then by Definition 4.5.6 for each $i \in \{1, 2\}$, if x_i is a vertex, then x_i is incident with some internal green face of K.

Suppose first that x_1 is a red blob, and x_2 is a vertex. Let

$$K_1 := \overline{K \setminus x_1}.$$

Since K is well-connected, it follows that K_1 is simply-connected; and $\mathbf{CArea}(K) \geq (1, 0)$ implies that $\mathbf{CArea}(K_1) \geq (1, 0)$. By Lemma 4.5.9 x_1 is simply-connected, and $\partial(x_1) \cap \partial(F)$ and $\partial(x_1) \cap \rho$ are single edges e_1 and g_1 respectively. Hence as K_1 is simply-connected, $\mathbf{CArea}(K_1) \geq (1, 0)$, and by Axiom T_6 each internal green face of Γ_A contains a boundary edge, we deduce that K_1 is bounded by a closed path $p \in \Gamma^1_A$ of the form $r_1r_3r_4$, where r_1 is a sub-path of $\partial(x_1)$ with $|r_1| \geq 1$, r_3 is a sub-path of p_1 with $|r_3| \geq 1$, and r_4 is a sub-path of p_2 with $|r_4| \geq 1$ (see Figure 4.9). Moreover, we have that g_1r_1 and r_1e_1 are sub-paths of $\partial(x_1)$, e_1r_3 is a sub-path of $\partial(F)$, and the common endpoint of g_1 and r_1 does not lie on $\omega \cap \tau$, since $g_1 \subseteq \rho$ and $|r_4| \geq 1$. Hence as by the second paragraph x_2 is incident with some internal green face of K (with is contained in K_1), by Lemma 4.5.12 F, x_1, x_2 are a neighbourhood of ρ, and K_1 is a sub-diagram of Γ_A bounded by F, x_1, x_2 and ρ. Lemma

Figure 4.10: The sub-diagram K_1 bounded by $r_1r_3r_2r_4$, see the proof of Theorem 4.5.13.

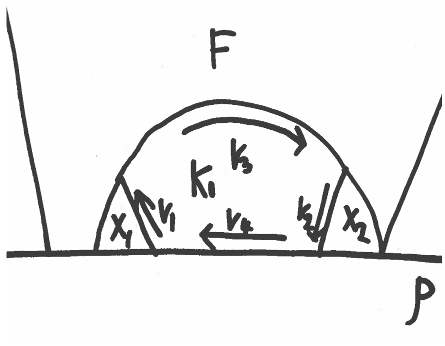

4.5.11 then implies that x_1 contains two boundary edges, contradicting Lemma 4.5.9.

Assume next that x_1 and x_2 are both red blobs. Let

$$K_1 := \overline{K \setminus (x_1 \cup x_2)}.$$

Then similarly as in the previous case, K_1 is simply-connected, $\mathbf{CArea}(K_1) \geq (1, 0)$, and for each $i \in \{1, 2\}$: x_i is simply-connected, and $\partial(x_i) \cap \partial(F)$ and $\partial(x_i) \cap \rho$ are single edges e_i and g_i respectively. Therefore, as K_1 is simply-connected, $\mathbf{CArea}(K_1) \geq (1, 0)$, and each internal green face of Γ_A contains a boundary edge, it follows that K_1 is bounded by a closed path $p \in \Gamma_A^1$ of the form $r_1r_3r_2r_4$, where for each $i \in \{1, 2\}$, r_i is a sub-path of $\partial(x_i)$ with $|r_i| \geq 1$; r_3 is a sub-path of p_1 such that if $|r_3| < 1$, then F is incident precisely once with r_3; r_4 is a sub-path of p_2 with $|r_4| \geq 1$ (see Figure 4.10); and the edges e_i, g_i satisfy Condition (d) of Lemma 4.5.12 (replace p_i by r_i in the statements). Therefore, by Lemma 4.5.11 F, x_1, x_2 are a neighbourhood of ρ, and K_1 is a sub-diagram of Γ_A bounded by F, x_1, x_2 and ρ. So by Lemma 4.5.12 x_1 contains two boundary edges, a contradiction.

Hence x_1 and x_2 are both vertices. Then since K is well-contained, p_1 is a sub-path of $\partial(F)$ and by the second paragraph x_1 and x_2 are incident with some internal green face of K, by Lemma 4.5.12 F, x_1, x_2 are a neighbourhood of ρ, and K is a sub-diagram of Γ_A bounded by F, x_1, x_2 and ρ. Hence by Lemma 4.5.11 we have $x_1, x_2 \in \omega \cap \tau$, and so by Lemma 4.5.2 Statement 3 of the theorem holds for F.

Hence D is contained in some island E of Γ_A. By above the minimal tricky green face F' for Γ_A with respect to ρ contained in E satisfies Statement 3 of the theorem, hence $D = F'$, as required. ∎

Theorem 4.5.13 is a powerful result. We finish this section with four results whose proofs rely on it.

Corollary 4.5.14. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Then for any $\rho \in \{\omega, \tau\}$ and any red blob B edge-incident once with F, F is not contained in a sub-diagram K of Γ_A well-bounded by B and ρ, where $\partial(K) \cap \partial(B)$ is the closure of the common edge of F and B.*

Proof. Suppose for a contradiction that such a K exists. Since $\partial(K) \cap \partial(B)$ is the closure of the common edge of F and B, we have $\partial(F) \cap \rho \neq \emptyset$, hence by Theorem 4.5.13 $\partial(F) \cap \rho$ is a single consolidated edge. So $\partial(F) \cap \rho = \partial(K) \cap \rho$, and $\partial(K) \cap \rho$ is the consolidated edge between F and the external face with boundary ρ. Hence $\partial(K)$ is a sub-path of $\partial(F)$. So if K is not well-connected, then F is not simply-connected, a contradiction. Therefore, we can use the same argument as in the proof of Part 3 of Lemma 4.3.4 to deduce that the label R of F satisfies $|R| = 2$, which contradicts Assumption 2.3.15 that no $R \in \mathcal{R}$ satisfies $|R| \in \{1, 2\}$. ∎

Corollary 4.5.15. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Assume that F contains a consolidated edge $e \subseteq \rho \in \{\omega, \tau\}$ with $|e| \geq 1$ and with endpoints v and w. Then there are $x, y \in \mathcal{B}_F$ that lie on or are incident with e such that x and y collectively give F curvature of at most $-1/2$. Moreover, if $z \in \{x, y\}$ is a red blob incident with $t \in \{v, w\}$ such that $t \notin \omega \cap \tau$, then z has an edge on ρ.*

Proof. If $v = w$, then v is incident twice with F, so by Lemma 2.6.10 we have $\chi(v, F, \Gamma_A) \leq -1/2$. Hence assume that $v \neq w$. By Part 1 of Lemma 4.3.4 there are $x, y \in \mathcal{B}_F$ that lie on or are incident with e, and x and y satisfy the second statement of the corollary. Suppose that $x \neq y$. Then by Part 2 of Lemma 4.3.4 we have $\chi(x, F, \Gamma_A) + \chi(y, F, \Gamma_A) \leq -1/4 - 1/4 = -1/2$.

So suppose that $x = y$. Then as $v \neq w$, it follows that x is a red blob B. By Lemma 2.6.10 we can further assume that $\delta_G(v) = 2 = \delta_G(w)$. If $|\partial(B) \cap \partial(F)| > 1$, then by Part 2 of Lemma 4.3.4 we have $\chi(B, F, \Gamma_A) \leq -1/2$. So suppose that B is edge-incident once with F. If $|\partial(B) \cap \partial(\Gamma_A)| \geq 2$ or if B is not simply-connected, then the lemma follows from Lemma 3.2.8. Hence assume that $|\partial(B) \cap \partial(\Gamma_A)| = 1$ and that B is simply-connected.

By Theorem 4.5.13 we have $\partial(F) \cap \rho = e$. Hence since $\delta_G(v) = 2 = \delta_G(w)$, we have $\rho = ef$, where f is the boundary edge of B. So as $|\partial(B) \cap \partial(F)| = 1$, Corollary 4.5.14 implies that B is contained in a simply-connected diagram K of Γ_A well-bounded by F and ρ and such that $\partial(K) = fe' = \partial(B)$, where e' is the common edge of F and B. But then Lemma 3.2.6 implies $\mathbf{Area}(B) = 0$, a contradiction. ∎

Lemma 4.5.16. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. Assume that there is a vertex $v \in \mathcal{B}_F$ incident at least twice with F, and v is incident with a consolidated edge e of F with $|e| \geq 1$ and $e \subseteq \rho$. Then $\rho = e$.*

Proof. Since v is incident at least twice with F, we have $\Gamma_A \in \mathcal{T} \setminus \mathcal{U}$ (see Definition 4.1.2), so if F is contained in an island E of Γ_A, then by Proposition 4.5.1 $E = \Gamma_A$ and some green face contained in E is incident with both of its endpoints, hence by Lemma 4.5.2 $\rho = e$. So assume that Γ_A is island-free. By Theorem 4.5.13 we have $\partial(F) \cap \rho = e$. Now there can be at most two boundary edges of F that lie on ρ and are incident with v, and if there are two such edges or if e is a loop, then the lemma holds. Hence assume for a contradiction that there is precisely one boundary edge of F incident with v and that e is not a loop.

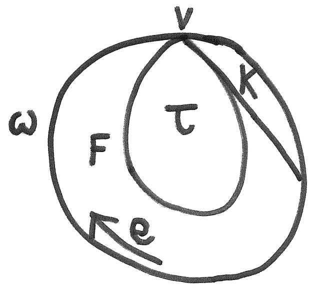

Figure 4.11: Configuration of a green face incident twice with a boundary vertex, see the proof of Lemma 4.5.16.

By Lemma 4.2.3 $\partial(F)$ passes through v twice, and $\Gamma_A \setminus (\overline{F})^{\circ}$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths. Let Γ be the one with $\rho \subseteq \Gamma$. Since the boundaries of Γ are simple and $e \subseteq \rho$, Γ is a union of islands and bridges (see Definition 2.5.10). Hence by the assumption from the previous paragraph Γ contains a simply-connected sub-diagram K of Γ_A well-bounded by F and ρ (see Figure 4.11 with $\rho = \omega$). But then $\partial(F) \cap \rho \neq e$, a contradiction. ■

Corollary 4.5.17. *Let v be a boundary vertex of $\Gamma_A \in \mathcal{T}$ such that $v \notin \omega \cap \tau$. If $\delta_G(v) \geq 4$, then there is an internal green face F incident with v and with an edge on the opposite boundary from that on which v lies, and v is not incident with any boundary edge of F.*

Proof. Without loss of generality assume that $v \in \omega$. Now there are at most two internal green faces of Γ_A that contain an edge of ω incident with v, and by Lemma 4.5.16 if such face F is incident at least twice with v, then F is the only internal green face that contains an edge of ω incident with v. By Lemma 4.2.3 no vertex is incident more than twice with an internal green face, hence as $\delta_G(v) \geq 4$, there is an internal green face F incident with v and such that v is not incident with any boundary edge of F.

By Axiom T_6, F contains a boundary consolidated edge f with $|f| \geq 1$. If $f \subseteq \tau$, then the lemma holds. So assume that $f \subseteq \omega$. Then $\partial(F) \cap \omega$ is not a single consolidated edge, contradicting Theorem 4.5.13. ■

4.6 More on red blobs

In this section we continue to study red blobs of diagrams in $\mathcal{T}$. In Section 4.6.1 we shall describe simply-connected red blobs that contain at least two boundary edges, in Section 4.6.2 we shall characterize the structure of diagrams containing bad red blobs (see Definition 4.1.1), and in Section 4.6.3 we shall show that even when a diagram in $\mathcal{T}$ contains a red blob with a complicated structure, its boundary words are conjugate by a word of length at most $2r + 1$, where r is the length of the longest green relator.

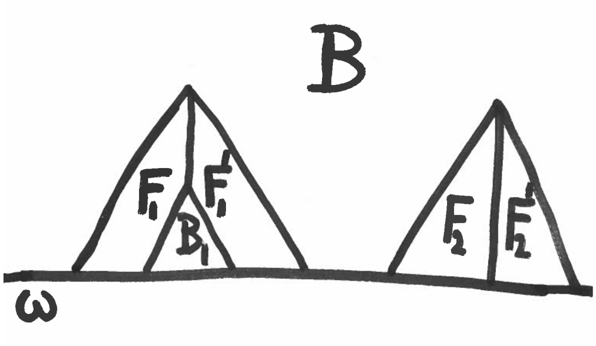

Figure 4.12: The structure of components in $\Gamma_A \setminus B$, see Theorem 4.6.2.

4.6.1 Simply-connected red blobs

Definition 4.6.1. Let B be a simply-connected red blob of an annular diagram Γ_A. Then B is called *complicated* if B contains at least two edges on ω or on τ.

The main result of this subsection is the following theorem. Recall Definition 2.5.10 of an island of a coloured annular diagram, and that an endpoint of an island is a vertex on $\omega \cap \tau$. Recall also Definitions 4.3.1 and 4.3.3 of the (boundary) curvature neighbourhood of an internal green face.

Theorem 4.6.2. *Let B be a complicated red blob of $\Gamma_A \in \mathcal{T}$. Then there exists an island E with $B \subseteq E$, and $E \setminus B$ is a disjoint union of components $T_1, \ldots, T_k$ for some k with the following properties (see Figure 4.12). For all i:*

1. *T_i is homeomorphic to a disc and $\partial(T_i)$ intersects at most one of the boundaries of Γ_A;*
2. *T_i contains precisely two internal green faces F_i and F_i' with $\partial(F_i) \cap \partial(F_i') \neq \emptyset$ and $B \in \mathcal{B}_{F_i} \cap \mathcal{B}_{F_i'}$;*
3. *Either $\overline{T} = \overline{F_i} \cup \overline{F_i'}$ or $\overline{T} = \overline{F_i} \cup \overline{F_i'} \cup \overline{B_i}$, where in the second case B_i is a red triangle with $B_i \in \mathcal{B}_{F_i} \cap \mathcal{B}_{F_i'}$.*

In particular, no internal green face of E is incident with any endpoint of E, the red blob B has an edge on both boundaries of Γ_A, and in E, no $\rho \in \{\omega, \tau\}$ contains two consecutive red edges.

Recall Definition 2.6.8 that $\chi(B, \Gamma)$ is the curvature that a red blob B gives to a single internal green face across each edge-incidence, and that if we say that a closed path p is of the form $p_1 p_2 \ldots p_n$, then p is a sequence of simple sub-paths p_i. Recall also Definition 4.4.3 of a well-contained sub-diagram given by a red blob and $\rho \in \{\omega, \tau\}$.

Lemma 4.6.3. *Let $\Gamma_A \in \mathcal{T}$ contain a red blob B and a well-contained sub-diagram K given by B and ρ. Suppose that K contains a corner v that lies on an edge of $\partial(B) \cap \rho$, and that $\partial(B)$ does not pass more than through v. Then $\delta_G(v) = 2$ and $\chi(B, \Gamma_A) \leq -1/2$.*

Proof. Let $p = \partial(K) \cap \rho$ and $n = |\partial(K) \cap \partial(F)|$. Since K contains an internal green face and K is well-connected (see Definition 4.4.2), by Axiom T_6 we have $|p| \geq 1$. Now B has an edge on ρ at v, and so does K, hence $v \notin \omega \cap \tau$, and the other corner w of K is distinct from v. Furthermore, by Axiom T_6 and Theorem 4.5.13 every green face $F \subseteq K$ has an edge on p and $\partial(F) \cap \rho$ is a single consolidated edge. Hence as $\partial(B)$ does not pass more than once through v, we have $\delta_G(v) = 2$, and therefore by Lemma 2.6.10 we have $\chi(v, \Gamma_A) = 0$. By Lemma 4.4.6 we have

$$1/2 + \chi(v, \Gamma_A) + \chi(w, \Gamma_A) = 1/2 + \chi(w, \Gamma_A) = -n \cdot (\chi(B, \Gamma_A) + 1/2).$$

From Lemma 3.1.9 we have $\chi(w, \Gamma_A) \geq -1/2$. Hence

$$0 \leq 1/2 + \chi(w, \Gamma_A) = -n \cdot (\chi(B, \Gamma_A) + 1/2),$$

and the lemma follows. ∎

Lemma 4.6.4. *Let B be a simply-connected red blob of $\Gamma_A \in \mathcal{T}$. Suppose that B is tricky for Γ_A with respect to $\rho \in \{\omega, \tau\}$, and that B contains an edge on ρ. Then the following statements hold.*

1. *There is a sub-diagram of Γ_A well-bounded by B and ρ that is well-connected, and with a corner lying on an edge of $\partial(B) \cap \rho$.*

2. *If Γ_A contains an island and B contains at least two edges on ρ, then there is a sub-diagram satisfying Part 1 with no corner on $\omega \cap \tau$.*

Proof. Part 1. First note that by Lemma 3.2.4 all consolidated edges between B and the external faces of Γ_A have length at most 1, and by definition of a tricky red blob $\partial(B) \cap \rho$ is not a single consolidated edge. Hence as B has an edge on ρ, it follows that there is a path $p \in \Gamma_A^1 \cap \partial(B)$ (see Definition 2.5.14) with $|p| \geq 1$ intersecting ρ only at its endpoints, where one of the endpoints of p is a vertex v lying on an edge of $\partial(B) \cap \rho$. So by Lemma 2.1.8 there is a simple path $p_1 \in \Gamma_A^1 \cap \partial(B)$ with $|p_1| \geq 1$ intersecting ρ only at its endpoints, and with $v \in p_1$. Now there is a sub-path p_2 of ρ with the same endpoints as p_1, and the sub-diagram K bounded by the closed path p_1p_2 is well-bounded by B and ρ.

If K is well-connected, then K satisfies Part 1, so suppose not. Then as K is well-bounded by B and ρ, it follows that $\tau \subseteq K$ and $\rho = \omega$. Therefore, as B contains an edge on ρ and $\partial(B) \cap \rho$ is not a single consolidated edge, there is a path $p' \in \Gamma_A^1 \cap \partial(B)$ distinct from p that intersects ρ only at its endpoints, and one of the endpoints of p' is a vertex v' lying on an edge of $\partial(B) \cap \rho$. Hence similarly as in the previous paragraph we deduce that there is a sub-diagram $K_1 \neq K$ well-bounded by B and ρ (see Figure 4.13); with $v' \in K_1$; and as $K_1 \neq K$, it follows that K_1 is well-connected. Hence Part 1 follows.

Part 2. By Lemma 3.1.10 B is contained in an island of Γ_A. As all consolidated edges between B and the external faces of Γ_A have length at most 1, there is a path $p \in \Gamma_A^1 \cap \partial(B)$

with $|p| \geq 1$ intersecting ρ only at its endpoints v and w, where $w \notin \omega \cap \tau$, and v lies on an edge of $\partial(B) \cap \rho$, so $v \notin \omega \cap \tau$. By Lemma 2.1.8 there is a simple path $p_1 \in \Gamma_A^1 \cap \partial(B)$ with $|p_1| \geq 1$ intersecting ρ only at v and w, so we can use similar arguments as in Part 1 to deduce that there is a sub-diagram K well-bounded by B and ρ and with corners v and w. As K is contained in an island, K is well-connected, so we are done. ∎

The next lemma gives precise value of $\chi(B, \Gamma_A)$ when B is complicated.

Lemma 4.6.5. *Let B be a simply-connected red blob of $\Gamma_A \in \mathcal{T}$. Then both of the following statements hold.*

1. *Suppose that B is tricky for Γ_A with respect to $\rho \in \{\omega, \tau\}$, and that B contains an edge on ρ. Then $\chi(B, \Gamma_A) \leq -1/2$, and if B contains at least two edges on ρ, then $\chi(B, \Gamma_A) = -3/4$.*

2. *If B' is a complicated red blob, then $\chi(B', \Gamma_A) = -3/4$.*

Proof. Part 1. By Lemma 4.6.4 there is a sub-diagram K of Γ_A well-bounded by B and ρ that is well-connected, with corners v and w, where v lies on an edge e of $\partial(B) \cap \rho$, and if B contains at least two edges on ρ, then without loss of generality we can assume that $v, w \notin \omega \cap \tau$. Note that $|\partial(K) \cap \partial(B)| \geq 1$ implies $\mathbf{CArea}(K) \geq (1, 0)$, so by Axiom T_6, $|\partial(K) \cap \rho| \geq 1$; and by Axiom T_6 and Theorem 4.5.13, $\partial(F) \cap \rho$ is a single consolidated edge for all green faces $F \subseteq K$. Suppose that $\partial(B)$ does not pass more than once through any vertex of $(\partial(K) \setminus \rho) \cup \{v\}$. Then K is well-contained, so by Lemma 4.6.3 $\delta_G(v) = 2$ and $\chi(B, \Gamma_A) \leq -1/2$. Suppose further that B contains at least two edges on ρ, so that $v, w \notin \omega \cap \tau$. Then from $\delta_G(v) = 2$ and Lemma 4.4.8 we have $\delta_G(w) = 2$, and so Lemma 4.4.9 gives $\chi(B, \Gamma_A) = -3/4$. Hence Part 1 holds for B.

Suppose instead that $\partial(B)$ passes more than once through some vertex $u \in (\partial(K) \setminus \rho) \cup \{v\}$. By Lemma 4.2.3 u is the only vertex through which $\partial(B)$ passes more than once, and by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2 (with $\rho \subseteq \Gamma_1$, say) whose boundaries are simple closed paths. In particular, Γ_1 contains all sub-diagrams well-bounded by B and ρ.

Suppose first that $u \neq v$. By definition of a tricky red blob $\partial(B) \cap \rho$ is not a single consolidated edge, hence as B contains an edge on ρ, there is a path $p \in \partial(\Gamma_1) \cap \partial(B)$ with $|p| \geq 1$ and $K \setminus \rho \nsubseteq p$, and intersecting ρ only at its endpoints v' and w', where v' lies on an edge of $\partial(B) \cap \rho$. Hence by Lemma 2.1.8 there is a simple path $p_1 \in \partial(\Gamma_1) \cap \partial(B)$ with $|p_1| \geq 1$ and $p_1 \neq K \setminus \rho$, and intersecting ρ only at v' and w'. Now note that there is a sub-path p_2 of ρ with endpoints v' and w', so the sub-diagram $K_1 \subseteq \Gamma_1$ bounded by $p_1 p_2$ is well-bounded by B and ρ (see Figure 4.14). Since $K_1 \subseteq \Gamma_1$ and B is simply-connected, K_1 is well-connected, and $\partial(K_1)$ does not intersect both ω and τ, so $v', w' \notin \omega \cap \tau$. Also, as $p_1 \neq K \setminus \rho$ and u is the only vertex through which $\partial(B)$ passes more than once, the simply-connectedness of B implies that $\partial(B)$ does not pass more than once through any vertex of $\partial(B) \cap \partial(K_1)$. Now $|\partial(K_1) \cap \partial(B)| \geq 1$ implies $\mathbf{CArea}(K_1) \geq (1, 0)$, hence K_1 is well-contained; and Axiom T_6 and Theorem 4.5.13 imply that $\partial(F) \cap \rho$ is a single consolidated

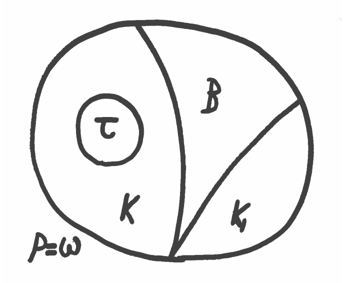

Figure 4.13: A case where K is not well-connected, see the proof of Lemma 4.6.4.

edge for all green faces $F \subseteq K_1$. By Lemma 4.6.3 we have $\delta_G(v') = 2$. Therefore, by Lemma 4.4.8 $\delta_G(w') = 2$, so by Lemma 4.4.9 $\chi(B, \Gamma_A) = -3/4$. Hence Part 1 holds for B.

Suppose instead that $u = v$. Then $\delta_G(v) \geq 3$; and since $|\partial(K) \cap \rho| \geq 1$ and $e \subseteq \partial(B) \cap \rho$, we have $w \neq v$. Therefore, as v is the only vertex through which $\partial(B)$ passes more than once, $\partial(B)$ does not pass more than once through any vertex of $\partial(K) \setminus \rho \cup \{w\}$, so K is well-contained.

Suppose first that $\partial(B) \cap \rho = \overline{e}$. It suffices to show $\chi(B, \Gamma_A) \leq -1/2$. We have $w \in \overline{e}$, so as $\partial(B)$ does not pass more than one through w, by Lemma 4.6.3 $\delta_G(w) = 2$ and $\chi(B, \Gamma_A) \leq -1/2$. Hence B satisfies Part 1.

Finally, suppose that $\partial(B) \cap \rho \neq \overline{e}$. As $e \subseteq \rho$ and the boundaries of Γ_1 are simple closed paths, Γ_1 is a union of islands and bridges. Therefore, there is a simply-connected sub-diagram $K_1 \subseteq \Gamma_1 \setminus K^\circ$ that is well-bounded by B and ρ and with corners v' and w' such that $v' \in \overline{e}$. Since B is simply-connected and v is the only vertex through which $\partial(B)$ passes more than once, K_1 does not intersect both ω and τ: so $v', w' \notin \omega \cap \tau$; and $\partial(B)$ does not pass more than once through any vertex of $\partial(B) \cap \partial(K_1)$. Now note that from $|\partial(K_1) \cap \partial(B)| \geq 1$ we have $\mathbf{CArea}(K_1) \geq (1, 0)$: so K_1 is well-contained, by Axiom T_6 and Theorem 4.5.13 $\partial(F) \cap \rho$ is a single consolidated edge for all green faces $F \subseteq K_1$, and by Lemma 4.6.3 we have $\delta_G(v') = 2$. Hence by Lemmas 4.4.8-4.4.9 we have $\chi(B, \Gamma_A) = -3/4$. Thus, Part 1 follows.

Part 2. Suppose that B' is a complicated red blob, with at least two edges on $\rho \in \{\omega, \tau\}$. As by Lemma 3.2.4 all consolidated edges between B' and the external faces of Γ_A have length at most one, B' is tricky for Γ_A with respected to ρ, so $\chi(B', \Gamma_A) = -3/4$. ∎

The next few results consider blobs that are not assumed to be simply-connected.

Lemma 4.6.6. *Suppose that B is a red blob of $\Gamma_A \in \mathcal{T}$ with $\chi(B, \Gamma_A) = -3/4$ that is edge-incident with an internal green face F. Then*

1. *$\mathcal{S}_F = \{B, x\} = \mathcal{B}_F$ has size two, B is edge-incident once with F, and x is either a red triangle edge-incident once with F that contains one edge of $\partial(\Gamma_A)$, or a boundary vertex incident once with F and such that $\delta_G(x) = 3$ and $x \notin \omega \cap \tau$;*

2. *$\partial(F) \cap \partial(\Gamma_A)$ is a single consolidated edge e with $|e| \geq 1$, and with endpoints t_1 and*

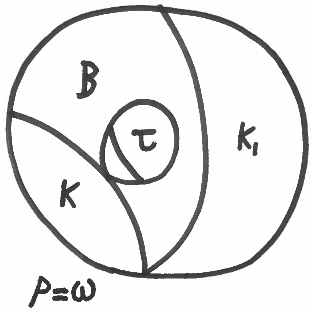

Figure 4.14: A case where $\partial(B)$ passes more than once through some vertex of $\partial(K) \setminus \rho$, see the proof of Lemma 4.6.5.

t_2, where t_1 lies on the common edge f of F and B, $\delta_G(t_1) = 2$ and $\delta_G(t_2) \leq 3$, and if $t_2 \in \partial(B)$, then $\delta_G(t_2) = 3$;

3. *the other endpoint t_3 of f satisfies $t_3 \notin \partial(\Gamma_A)$ and $\delta_G(t_3) = 2$.*

Proof. By Axiom T_6, F has a boundary consolidated edge e with $|e| \geq 1$. Assume without a loss of generality that $e \subseteq \omega$, and let $\partial(e) = \{t_1, t_2\}$. By Theorem 4.5.13 we have $\partial(F) \cap \omega = e$. Since $\chi(B, F, \Gamma_A) = -3/4$ and $\kappa_{\Gamma_A}(F) = 0$, we have that B is edge-incident once with F. Also, by Part 2 of Lemma 4.3.4 every element of $\mathcal{B}_F$ curvature incident more than once with F gives F curvature of at most $-1/2$. Hence no such element exists, so by Part 3 of Lemma 4.3.4 we have $|\mathcal{B}_F| > 1$. Therefore, since $\chi(B, F, \Gamma_A) = -3/4$, by Lemmas 2.6.10 and 3.2.8 Part 1 follows.

To prove Part 2, we first show that t_1 lies on f, that $\delta_G(t_1) = 2$ and $\delta_G(t_2) \leq 3$, and if $t_2 \in \partial(B)$, then $\delta_G(t_2) = 3$. By Corollary 4.5.15 there are $\alpha, \beta \in \mathcal{B}_F$ that are equal to or are incident with t_1 and t_2 such that α and β collectively give F curvature of at most $-1/2$. Suppose first that no $t \in \{t_1, t_2\}$ lies on f. Then $B \notin \{\alpha, \beta\}$, hence by Part 1 we have $\alpha = x = \beta$, and x gives F curvature $-1/4$, a contradiction.

Assume next that $t_1, t_2 \in \partial(B)$ and $\delta_G(t_1) = 2 = \delta_G(t_2)$. Then t_1 and t_2 are endpoints of f. Furthermore, as t_1, t_2, B are all curvature incident once with F, we have $t_1 \neq t_2$, and from $\partial(F) \cap \omega = e$ we deduce that ef is a sub-path of $\partial(F)$. Hence since F is simply-connected, we deduce that F is contained in a sub-diagram bounded by the closed path ef, which is in particular well-bounded by B and ω (see Definition 4.4.2), contradicting Corollary 4.5.14. Hence if $t_1, t_2 \in \partial(B)$, then either $\delta_G(t_1) \geq 3$ or $\delta_G(t_2) \geq 3$, so by Part 1, either $\delta_G(t_1) = 3$ or $\delta_G(t_2) = 3$.

By the second paragraph there exists $i \in \{1, 2\}$ such that $t_i \in \overline{f}$. If t_i is unique, then as $t_1 \neq t_2$, by Part 1 of Lemma 4.3.4 there is $y \in \mathcal{B}_F \setminus \{B, t_i\}$ equal to or incident with $t_j \in \{t_1, t_2\} \setminus \{t_i\}$. So by Part 1 we have $\delta_G(t_i) = 2$ and $\delta_G(t_j) \leq 3$. If $t_1, t_2 \in \partial(B)$, then by the previous paragraph there exists $i \in \{1, 2\}$ such that $\delta_G(t_i) = 3$, and therefore by Part 1 the $t_j \in \{t_1, t_2\} \setminus \{t_i\}$ satisfies $\delta_G(t_j) = 2$, so $t_j \in \overline{f}$. Hence by relabelling, if necessary, we can assume that t_1 lies on f and $\delta_G(t_1) = 2$, so $\delta_G(t_2) \leq 3$, and if $t_2 \in \partial(B)$, then $\delta_G(t_2) = 3$.

It remains to show $\partial(F)\cap\partial(\Gamma_A) = e$ and prove Part 3. If all endpoints of f are in $\{t_1, t_2\}$, then as t_1, t_2, B are all curvature incident once with F and $\partial(F)\cap\omega = e$, we deduce that ef is a sub-path of $\partial(F)$, and so we get a contradiction as in the third paragraph. So let t_3 be the endpoint of f such that $t_3 \notin \{t_1, t_2\}$.

Suppose for a contradiction that $t_3 \in \partial(\Gamma_A)$. If $t_3 \in \omega$, then as $t_3 \in \overline{f}$ and $\partial(F)\cap\omega = e$, we have $t_3 \in \{t_1, t_2\}$, a contradiction. Hence $t_3 \in \tau$. Let $h \subseteq \tau$ be the boundary consolidated edge of F with $t_3 \in h$. As by Part 1, B is edge-incident once with F, by Part 1 of Lemma 4.3.4 there exists $y \in \mathcal{B}_F \setminus \{B\}$ curvature incident with F at some endpoint of h (note that if $|h| = 0$, then the only endpoint of h is t_3), and as by Part 1 we have $\partial(F)\cap(\omega\cap\tau) = \emptyset$, either $y \in \tau\setminus\omega$, or y is a red blob with an edge on τ. In particular, by Part 1 we have $\delta_G(t_2) = 2$, so $t_2 \notin \partial(B)$ and $\delta_R(t_2) \geq 1$. Therefore, there is a red blob $B_1 \neq B$ edge-incident with F at t_2, and with an edge on ω. By Part 1, B_1 is a red triangle with precisely one boundary edge, hence $B_1 \neq y$, and so $|\mathcal{S}_F| \geq 3$, a contradiction. Hence t_3 is interior. So by Part 1 we have $t_3 \notin \mathcal{S}_F$, and therefore $\delta_G(t_3) = 2$.

Finally, suppose for a contradiction that $\partial(F)\cap\tau \neq \emptyset$. Then as $\partial(F)\cap(\omega\cap\tau) = \emptyset$, by Part 1 of Lemma 4.3.4 either $\mathcal{S}_F$ contains a vertex on $\tau\setminus\omega$, or a red blob edge-incident with F at a vertex on $\tau\setminus\omega$, and with an edge on τ. Hence as t_3 is interior, there exists $y \in \mathcal{S}_F \setminus \{B, t_2\}$. Hence by Part 1 we have $\delta_G(t_2) = 2$, so $t_2 \notin \partial(B)$ and $\delta_R(t_2) \geq 1$, and as in the previous paragraph we deduce that there is a red triangle edge-incident with F at t_2, and with precisely one boundary edge, on ω. Hence again $|\mathcal{S}_F| \geq 3$, a contradiction. So as $\partial(F)\cap\tau = \emptyset$, we have $\partial(F)\cap\partial(\Gamma_A) = e$, as required. ∎

Lemma 4.6.7. *Suppose that B is a red blob of $\Gamma_A \in \mathcal{T}$ with $\chi(B,\Gamma_A) = -3/4$ that is edge-incident with an internal green face F. Then F is contained in a component T of $\Gamma_A \setminus B$ that satisfies Statements 1-3 of Theorem 4.6.2.*

Proof. Since B and F satisfy assumptions of Lemma 4.6.6, we can let t_1, t_2, t_3, e, f and x be as in Lemma 4.6.6. Without loss of generality assume that $e \subseteq \omega$. As t_3 is interior and $\delta_G(t_3) = 2$, there is precisely one internal green face F' incident with F, and F' is incident with t_3 and edge-incident with B. Since B is edge-incident with F', Lemma 4.6.6 holds for F'.

Assume first that $\delta_G(t_2) = 2$. Then x is a red blob B_1. As $\delta_G(t_3) = 2$, F and F' share a consolidated edge p_2 with endpoints t_3 and a vertex t_4 with $t_4 \in \partial(B_1)$. Since $\mathcal{S}_F = \{B, B_1\}$, we have $\delta_G(t_4) = 2$, so B_1 is edge-incident with F'. By Part 1 of Lemma 4.6.6 we therefore have $\mathcal{B}_{F'} = \{B, B_1\}$, and B and B_1 are edge-incident once with F'. So $\partial(F')\cap\omega \neq \emptyset$. Let t_5 be the vertex of F' lying on $\partial(B_1)\cap\omega$, and let $p_3 \subseteq \omega$ be the consolidated edge of F' with $t_5 \in p_3$. By Part 2 of Lemma 4.6.6 we have $\partial(F')\cap\partial(\Gamma_A) = p_3$ and $|p_3| \geq 1$.

Since $\mathcal{S}_{F'} = \{B_1, B\} = \mathcal{B}_{F'}$, and B and B_1 are edge-incident once with F', it follows that p_3 is incident with a boundary edge of B. Let p_4 and p_5 be the common edges of F' and B and of F' and B_1 respectively. Since $\mathcal{S}_{F'} = \{B_1, B\} = \mathcal{B}_{F'}$, all vertices $v \in \partial(F')$ satisfy $\delta_G(v) = 2$. Hence $\partial(F') = p_2p_4p_3p_5$. Let p_6 be the common edge of F and B_1. From $\mathcal{S}_F = \{B, B_1\} = \mathcal{B}_F$ we deduce that that all vertices $v \in \partial(F)$ satisfy $\delta_G(v) = 2$. Hence

$\partial(F) = efp_2p_6$. Let p_7 be the boundary edge of B_1. It follows that F, F', B_1° are contained in a component T of $\Gamma_A \setminus B$ bounded by the closed path $efp_4p_3p_7$. Hence $\overline{T} = \overline{F} \cup \overline{F'} \cup \overline{B_1}$. In particular, T is homeomorphic to a disc and $\partial(T) \cap \tau = \emptyset$, so the lemma follows.

Assume instead that $\delta_G(t_2) = 3$. Then $x = t_2$. As t_3 is interior and $\delta_G(t_3) = 2$, from $\mathcal{B}_F = \{B, t_2\} = \mathcal{S}_F$ it follows that there is a consolidated edge p_2 common to F and F' incident with t_2 and t_3. Hence by Part 1 of Lemma 4.6.6 we have $\mathcal{B}_{F'} = \{B, t_2\} = \mathcal{S}_{F'}$, B is edge-incident once with F', and t_2 is incident once with F'. In particular, F' has a consolidated edge p_3 on ω with $t_2 \in p_3$. By Part 2 of Lemma 4.6.6 we have $\partial(F') \cap \partial(\Gamma_A) = p_3$ and $|p_3| \geq 1$, and from $\mathcal{B}_{F'} = \{B, t_2\} = \mathcal{S}_{F'}$ we have that p_3 is incident with a boundary edge of B.

By applying $\mathcal{S}_F = \mathcal{B}_F = \{B, t_2\} = \mathcal{B}_{F'} = \mathcal{S}_{F'}$ once again we deduce that all vertices $v \in (\partial(F) \cup \partial(F')) \setminus \{t_2\}$ satisfy $\delta_G(v) = 2$. So $\partial(F) = efp_2$ and $\partial(F') = p_2p_4p_3$, where p_4 is the common edge of F' and B. Therefore, F and F' are contained in a component T of $\Gamma_A \setminus B$ bounded by the closed path efp_4p_3. So $\overline{T} = \overline{F} \cup \overline{F'}$, and hence T is homeomorphic to a disc and $\partial(T) \cap \tau = \emptyset$. ■

Proof of Theorem 4.6.2. By Lemma 4.6.5 we have $\chi(B, \Gamma_A) = -3/4$ since B is complicated. Assume that some $\rho \in \{\omega, \tau\}$ is equal to $\overline{e}$ for some edge $e \subseteq B$. Let v be the endpoint of e. If there is an internal green face F incident with v and edge-incident with B, then by Axiom T_6, F contains an edge on $\rho' \in \{\omega, \tau\}$ with $\rho' \neq \rho$, so $\partial(F) \cap \omega \neq \emptyset \neq \partial(F) \cap \tau$, contradicting Lemma 4.6.6. Hence Γ_A contains an island.

Recall that we assume $\mathbf{CArea}(\Gamma_A) \geq (1, 0)$, and by Lemma 3.2.10 B is not an island of Γ_A. By Lemma 4.6.7 each internal green face of Γ_A edge-incident with B is contained in some component T of $\Gamma_A \setminus B$ that satisfies Statements 1-3 of the theorem. Hence as by Lemma 3.2.4 all consolidated edges between any red blob and any green face of Γ_A have length at most one, B contains edges on both boundaries of Γ_A, and by the previous paragraph either Γ_A is island-free and $\Gamma_A \setminus B$ is a disjoint union of components $T_1, \ldots, T_k$ that satisfy Statements 1-3 of the theorem; or $B \subseteq E$ for some island E of Γ_A, and $E \setminus B$ is a disjoint union of components $T_1, \ldots, T_k$ that satisfy Statements 1-3 of the the theorem. In particular, if B is contained in the island E, then no internal green face of E is incident with any endpoint of E, and as each internal green face of Γ_A contains a boundary edge and all consolidated edges between any red blob and any green face of Γ_A have length at most one, we deduce that in E, no $\rho \in \{\omega, \tau\}$ contains two consecutive red edges.

Suppose for a contradiction that Γ_A is island-free. Let $t = \mathbf{Area}(B)$, $d = |\partial(B) \cap \partial(\Gamma_A)|$, and $k = |\partial(B)|$. By the previous paragraph for each boundary edge e with $e \subseteq \partial(B)$ we can associate two distinct interior edges f_1 and f_2 with $\overline{f}_1 \cap \overline{e} \neq \emptyset \neq \overline{f}_2 \cap \overline{e}$, and such that $f_1 \subseteq \partial(B) \cap \partial(T_1)$ and $f_2 \subseteq \partial(B) \cap \partial(T_2)$, where T_1 and T_2 are components of $\Gamma_A \setminus B$. Observe that for distinct boundary edges e_1 and e_2 with $e_1, e_2 \subseteq \partial(B)$ the sets of edges associated to e_1 and e_2 respectively as above are pairwise disjoint. Thus $d \leq k/3$. Since B is simply-connected, by Lemma 3.2.8 we have

$$\chi(B, \Gamma_A) = -3/4 = \frac{-t}{2(t-d)+4},$$

rearranging gives $d = t/3 + 2$. By Lemma 3.2.6 we have $k = t + 2$, hence

$$k/3 \geq d = t/3 + 2 > t/3 + 2/3 = k/3,$$

a contradiction. The theorem follows. ∎

Definition 4.6.8. Let B be a simply-connected red blob of an annular diagram Γ_A. Then B is called *highly hyperbolic* if B contains precisely two boundary edges e_1 and e_2, with $e_1 \subseteq \omega$ and $e_2 \subseteq \tau$.

Lemma 4.6.9. *Suppose that $\Gamma_A \in \mathcal{T}$ contains a simply-connected red blob B that is not complicated. Then $\chi(B, \Gamma_A) \geq -1/2$, and $\chi(B, \Gamma_A) = -1/2$ if and only if B is highly hyperbolic.*

Proof. Since B is not complicated, B has at most one edge on a single boundary of Γ_A. Hence by Lemma 3.2.8 we have $\chi(B, \Gamma_A) \geq -1/2$, and $\chi(B, \Gamma_A) = -1/2$ if and only if B is highly hyperbolic ∎

In the next chapter we shall use extensively the following proposition to simplify cases where a simply-connected red blob contains at most one boundary edge.

Proposition 4.6.10. *Let B be a simply-connected red blob of $\Gamma_A \in \mathcal{T}$, and let $\rho \in \{\omega, \tau\}$. Suppose that B has an edge on ρ, and that B is tricky for Γ_A with respect to ρ. Then B contains at least two boundary edges.*

Hence if B' is a simply-connected red blob with exactly one boundary edge, on ρ say, then $\partial(B') \cap \rho$ is a sub-path of $\partial(B')$.

Proof. By Lemma 4.6.5 we have $\chi(B, \Gamma_A) \leq -1/2$, hence by Lemma 3.2.8 B contains at least two boundary edges. For the second statement, note that by the first statement of the lemma B' is not tricky with respect to ρ, hence $\partial(B') \cap \rho$ is a single consolidated edge. So $\partial(B') \cap \rho$ is a sub-path of $\partial(B')$. ∎

Lemma 4.6.11. *Let B be a highly hyperbolic red blob of $\Gamma_A \in \mathcal{T}$, and assume that all red blobs of Γ_A are simply-connected. Then $|\partial(B)| \leq 6$.*

Proof. Let e_1 and e_2 be the edges of B with $e_1 \subseteq \omega$ and $e_2 \subseteq \tau$. We claim that there are two distinct simple sub-paths l and r of $\partial(B)$ that contain no edge of $\partial(\Gamma_A)$, and the endpoints of l and r lie on $\overline{e_1}$ and $\overline{e_2}$. To prove the claim, suppose first that $\partial(B)$ does not pass more than once through any vertex. Then $\partial(B)$ is a simple closed path, so $\partial(B) = e_1 l e_2 r$, where l and r have the stated properties.

Suppose instead that $\partial(B)$ passes through some vertex v more than once. Then by Lemma 4.2.3 v is the only vertex with this property, $\partial(B)$ passes through v twice, and by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Hence again $\partial(B) = e_1 l e_2 r$, where l and r have the required properties.

Note that it suffices to show $|l|, |r| \leq 2$, so suppose for a contradiction that $|l| > 2$. Let v_1 and v_2 be the endpoints of l with $v_1 \in \overline{e_1}$. Then there is an internal green face F that shares an edge e with B such that $e \subseteq l$ and $\overline{e} \cap (\{v_1\} \cup \{v_2\}) = \emptyset$. By Axiom T_6, F contains a boundary edge. Without loss of generality assume that it lies on ω. Let p be the sub-path of l with endpoints v_1 and $v \in \overline{e}$, where v is such that $e \subseteq p$. Let F_1 be the internal green face such that $\partial(F_1) \cap \partial(B)$ contains the edge incident with v_1, and let F_2 be an internal green face with an edge on p distinct from F_1.

Suppose that B is edge-incident more than once with F_2. By Proposition 4.2.8 B is edge-incident twice with F_2, hence by Lemma 4.6.9 we have $\chi(B, F_2, \Gamma_A) = -1$. By Axiom T_6 we have $\kappa_{\Gamma_A}(F_2) = 0$, so $\mathcal{S}_{F_2} = \{B\}$. By Proposition 4.2.8 $\Gamma_A \setminus (\overline{F_2} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Hence as $\overline{e} \cap (\{v_1\} \cup \{v_2\}) = \emptyset$ and $\mathcal{S}_{F_2} = \{B\}$, it follows that F_2 has no boundary edge, contradicting Axiom T_6. So B is edge-incident once with F_2, hence $\overline{e} \cap (\{v_1\} \cup \{v_2\}) = \emptyset$ implies $B \in \mathcal{S}_{F_2} \setminus \mathcal{B}_{F_2}$.

Since B contains edges on the opposite boundaries of Γ_A, by Part 3 of Lemma 4.2.3 no vertex of $\mathcal{S}_{F_2}$ is incident more than once with F_2; and as all red blobs of Γ_A are simply-connected, by Proposition 4.2.8 no red blob of $\mathcal{S}_{F_2}$ is edge-incident more than once with F_2. Hence by Part 3 of Lemma 4.3.4 we have $|\mathcal{B}_{F_2}| \geq 2$. By Part 2 of Lemma 4.3.4 each element of $\mathcal{B}_{F_2}$ gives F_2 curvature of at most $-1/4$. Hence $|\mathcal{S}_{F_2}| = 3$, and each $x \in \mathcal{B}_{F_2}$ is either a red triangle that contains one boundary edge, or a vertex with $\delta_G(x) = 3$ and $x \notin \omega \cap \tau$. In particular, all interior vertices $v_2 \in \partial(F_2)$ satisfy $\delta_G(v_2) = 2$.

Since $\partial(F_1) \cap \rho \neq \emptyset$, by Theorem 4.5.13 $\partial(F_1) \cap \rho$ is a single consolidated edge. Hence by the previous paragraph we have $\mathcal{S}_{F_1} = \{B, x\} = \mathcal{B}_{F_1}$, where B is edge-incident once with F_1, and x is either a red triangle edge-incident once with F_1 that contains one boundary edge, or a vertex of green degree 3 incident once with F_1 and not on $\omega \cap \tau$. Hence

$$\kappa_{\Gamma_A}(F_1) = 1 + \chi(B, F_1, \Gamma_A) + \chi(x, F_1, \Gamma_A) = 1 - 1/2 - 1/4 = 1/4 > 0,$$

contradicting Axiom T_6. So $|l| \leq 2$, and similarly $|r| \leq 2$. The lemma follows. ∎

4.6.2 Bad red blobs

In this subsection we characterize structure of diagrams in $\mathcal{T}$ that contain bad red blobs (see Theorems 4.6.13-4.6.14). This will enable us to solve the conjugacy problem in quadratic time. Recall that a red blob B is annular if $\mathbb{R}^2 \setminus B^\circ$ is comprised of two components.

Lemma 4.6.12. *Let B be an annular red blob of $\Gamma_A \in \mathcal{T}$ with boundary length l and* $\mathbf{Area}(B) =$

t. Then $l = t$. Hence letting $d = |\partial(B) \cap \partial(\Gamma_A)|$, we have

$$\chi(B, \Gamma_A) = \frac{-t}{2(t-d)}.$$

In particular, $\chi(B, \Gamma_A) \leq -1/2$.

Proof. We use ideas from the proof of [34, Lemma 4.12]. Let B^* be the induced sub-graph of the dual graph Γ_A^* of Γ_A on those vertices that corresponds to triangles in B. Note that a vertex of degree one in B^* corresponds to a triangle of B having two edges on $\partial(B)$. If we delete this triangle from B, then we decrease both the number of triangles in B and the number of edges in $\partial(B)$ by one while B^* remains connected. Furthermore, the vertices of degree two correspond to triangles of B with exactly one edge on $\partial(B)$.

Let B_1^* be the graph obtained from B^* by repeatedly removing vertices of degree 1. By graph duality the vertices of B° correspond to cycles of B^*. Hence as B is annular and by Axiom T_3 all vertices of B lie on $\partial(B)$, it follows that B_1^* is a cycle graph. Let B_1 be the sub-diagram of Γ_A consisting of the closures of the red triangles of B that correspond to the vertices of B_1^*. By the previous paragraph $|\partial(B_1)| = \mathbf{Area}(B_1)$, hence $l = t$.

Now by the description of $\mathbf{ComputeRSym}$ (see Algorithm 2.6.4) we have

$$\chi(B, \Gamma_A) = \frac{-t}{2|\partial(B) \setminus \partial(\Gamma_A)|},$$

so the second statement follows from $l = t$. ■

Theorem 4.6.13. *Let B be a bad red blob of $\Gamma_A \in \mathcal{T}$ with an edge on $\partial(\Gamma_A)$. Then $\Gamma_A \setminus B$ is a disjoint union of components T_i (see Figure 4.12) that satisfy Statements 1-3 of Theorem 4.6.2. Moreover, B contains an edge on both boundaries of Γ_A, and each $\rho \in \{\omega, \tau\}$ satisfies $|\rho| \geq 2$.*

Proof. Let B have an edge e on $\rho \in \{\omega, \tau\}$. By Lemma 4.2.4 Γ_A is island-free, B is the only bad red blob of Γ_A, and by definition of a bad red blob there exist annular diagrams S and S' such that $\Gamma_A \setminus B^\circ = S \cup S'$ and $S \cap S' = \emptyset$, where $\rho \subseteq S$ (say). If S has no internal green faces then S has area 0. Hence as by Lemma 3.2.4 all consolidated edges between B and the external faces of Γ_A have length at most one, we deduce that $\rho = \overline{e}$; and the endpoint v of e has $\delta_G(v) = 1$, contradicting Axiom T_3. So $\mathbf{CArea}(S) \geq (1, 0)$. By Axiom T_6 each internal green face of Γ_A contains a boundary edge, hence $|\rho| \geq 2$.

We next show that $\chi(B, \Gamma_A) = -3/4$. By Axiom T_1 and Lemma 4.2.4 the boundaries of S are simple closed paths, hence as B contains an edge on ρ, S is a union of islands and bridges. Therefore, there is a sub-diagram K of Γ_A well-bounded by B and ρ that is well-connected (see Definition 4.4.2), with corners v and w such that v lies on an edge of $\partial(B) \cap \rho$, and $\partial(B)$ does not pass more than once through any vertex of $\partial(B) \cap \partial(K)$. Since $|\partial(K) \cap \partial(B)| \geq 1$, we have $\mathbf{CArea}(K) \geq (1, 0)$, hence K is well-contained (see Definition 4.4.3). By Axiom T_6 and Theorem 4.5.13 $\partial(F) \cap \rho$ is a single consolidated edge for all green faces $F \subseteq K$; and as $\partial(B)$ does not pass more than once through any vertex of $\partial(B) \cap \partial(K)$, by Lemma

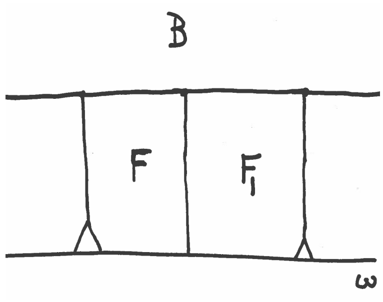

Figure 4.15: A bad red blob B with no boundary edge, see Theorem 4.6.14.

4.6.3 $\delta_G(v) = 2$. So Lemma 4.4.8 implies $\delta_G(w) = 2$, and therefore Lemma 4.4.9 gives $\chi(B, \Gamma_A) = -3/4$, as claimed.

Hence Lemma 4.6.7 holds for all internal green faces edge-incident with B, so B has an edge on both boundaries of Γ_A. Therefore, similarly as for S we can show that $\mathbf{CArea}(S') \geq (1, 0)$, and that the $\rho' \in \{\omega, \tau\}$ with $\rho' \neq \rho$ has $|\rho'| \geq 2$. Hence as all consolidated edges between B and the external faces of Γ_A have length at most one, and Lemma 4.6.7 holds for all internal green faces edge-incident with B, $\Gamma_A \setminus B$ is a disjoint union of components T_i that satisfy Statements 1-3 of Theorem 4.6.2, as required. ∎

Theorem 4.6.14. *Let B be a bad red blob of $\Gamma_A \in \mathcal{T}$. Then both of the following two statements hold.*

1. *Suppose that B contains no edge of $\partial(\Gamma_A)$. Then all internal green faces F of Γ_A satisfy the following 3 statements (see Figure 4.15).*

   (i) $\mathcal{S}_F = \{B, x, y\} = \{B\} \cup \mathcal{B}_F$.

   (ii) B is edge-incident once with F.

   (iii) Each element $z \in \mathcal{B}_F$ is either a red triangle that contains one boundary edge, or a vertex with $\delta_G(z) = 3$ and $z \notin \omega \cap \tau$.

   Furthermore, each of ω and τ contains a green edge.

2. *Each internal green face of Γ_A is edge-incident with B, no $\rho \in \{\omega, \tau\}$ contains two consecutive red edges, and either $|\rho| \geq 2$, or ρ contains a green edge.*

Proof. Part 1. By Lemma 4.2.4 Γ_A is island-free, and B is the only bad red blob of Γ_A, hence by Lemma 4.2.5 all red blobs of Γ_A other than B are simply-connected. By definition of a bad red blob there exist annular diagrams S and S' such that $\Gamma_A \setminus B^\circ = S \cup S'$ and $S \cap S' = \emptyset$, where $\omega \subseteq S$ (say). By symmetry, it suffices to show that $\mathbf{CArea}(S) \geq (1, 0)$ (as then by Axiom T_6, ω contains a green edge), and that the theorem holds for all internal green faces of S.

Since B contains no edge of $\partial(\Gamma_A)$, it follows that S contains an internal green face F edge-incident with B, and note that $B \in \mathcal{S}_F \setminus \mathcal{B}_F$. We first show that the theorem holds for

F. By Lemma 4.6.12 we have $\chi(B, \Gamma_A) = -1/2$. Suppose that B is edge-incident more than once with F. By Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$, hence $\mathcal{S}_F = \{B\}$. But by Part 3 of Lemma 4.3.4 we have $|\mathcal{B}_F| \geq 1$, a contradiction. Hence B is edge-incident once with F.

By Axiom T_6, F has a consolidated edge $e \subseteq \omega$ with $|e| \geq 1$. Let v and w be the endpoints of e. By Part 1 of Lemma 4.3.4 there are $x, y \in \mathcal{B}_F$ that are curvature incident with F at v and w respectively. Suppose first that $x = y$ is a red blob. As Γ_A is island-free, by Theorem 4.6.2 x is not complicated. Hence as $S \cap S' = \emptyset$, x contains one boundary edge, so by Proposition 4.6.10 $\partial(x) \cap \omega$ is a single consolidated edge f with $|f| = 1$. By Theorem 4.5.13 we have $\partial(F) \cap \omega = e$. Hence $\delta_G(v) = 2 = \delta_G(w)$ and $\omega = ef$. Therefore, we have $\mathcal{B}_F = \{x\}$, so by Part 3 of Lemma 4.3.4 x is edge-incident exactly twice with F, and thus by Lemma 3.2.8 the theorem holds for F. Next assume that $v = x = w$ is a vertex. Then x is incident more than once with F, hence by Lemma 2.6.10 F satisfies the theorem. Finally, assume that $x \neq y$. Then by Part 2 of Lemma 4.3.4 we have $\chi(x, F, \Gamma_A), \chi(y, F, \Gamma_A) \leq -1/4$, so apply Lemmas 2.6.10 and 3.2.8 to deduce that the theorem holds for F.

Therefore, all interior vertices $v \in \partial(F)$ satisfy $\delta_G(v) = 2$. So if $|\partial(B) \cap \partial(S)| = 1$, then the theorem holds for all internal green faces of S. Otherwise, there are two (not necessarily distinct) internal green faces F_1 and F_2 incident with F, and F_1 and F_2 are edge-incident with B. Since F_1 and F_2 are edge-incident with B, similarly as for F we deduce that the theorem holds for F_1 and F_2. Hence by induction we can show that all internal green faces of S satisfy the theorem. By symmetry the same holds for S'. So Part 1 holds.

Part 2. By Theorem 4.6.13 and Part 1 each internal green face of Γ_A is edge-incident with B. Let $\rho \in \{\omega, \tau\}$. Since each internal green face of Γ_A contains a boundary edge and by Lemma 3.2.4 all consolidated edges between any red blob and any green face of Γ_A have length at most one, by Theorem 4.6.13 and Part 1, ρ does not contain two consecutive red edges. By Theorem 4.6.13 $|\rho| \geq 2$ if B contains a boundary edge, else by Part 1, ρ contains a green edge, so we are done. ■

4.6.3 Thickness of complicated and bad red blobs

Throughout this whole subsection we label all edges of the boundary of a red blob B with respect to the orientation from B, and we let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R} \rangle$ be a valid pregroup presentation (see Definition 2.6.14) for a group G. We shall show that if a diagram in $\mathcal{T}$ contains a complicated or a bad red blob (see Definitions 4.1.1 and 4.6.1), then its boundary words are conjugate in G by an element of X^* of length at most $2r + 1$, where $r = \max\{|R| : R \in \mathcal{R}\}$ (see Proposition 4.6.18).

Recall Definition 3.1.3 of a minimal coloured area of an annular and simply-connected diagram. The next lemma holds only under valid pregroup presentations.

Lemma 4.6.15. *Let Γ be a coloured van Kampen diagram over $\mathcal{P}$ with a simply-connected red blob B. Suppose that both of the following statements hold.*

(a) Γ has minimal coloured area.

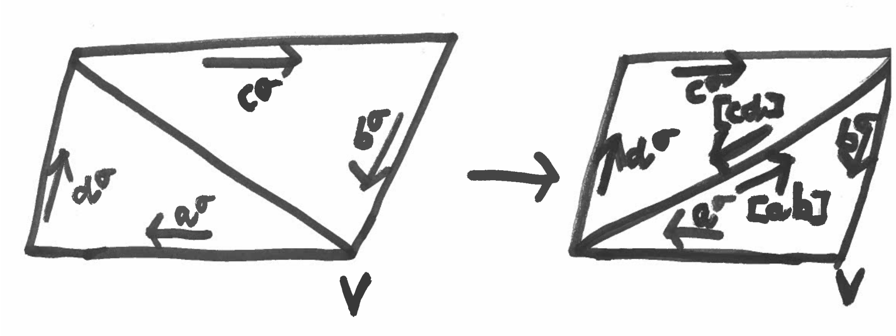

Figure 4.16: Retriangulation of Φ at v, see the proof of Lemma 4.6.17.

(b) All vertices of B lie on $\partial(B)$.

Then all of the following statements hold.

1. *No proper sub-word of any (cyclic) boundary word w of B is equal to 1 in $U(P)$, w is cyclically σ-reduced, and if $a, b \in X^\sigma$ and ab is a sub-word of w, then $(a, b) \in D(P)$.*

2. *For each vertex $v \in \partial(B)$, there exists a retriangulation B_1 of B with precisely one triangle incident with v.*

Proof. Since Γ has minimal coloured area and each vertex of B lies on $\partial(B)$, the proof of Lemma 3.2.7 shows that no proper sub-word of w is equal to 1 in $U(P)$; by Lemma 3.2.6 we have $|\partial(B)| \geq 3$: so combining this with the previous statement we deduce that w is cyclically σ-reduced; and as $\mathcal{P}$ satisfies trivial-interleaving, by Lemma 2.5.13 we have $(a, b) \in D(P)$. Hence Part 1 follows.

For Part 2, if there is a single triangle of B incident with v, then take $B_1 = B$. So assume that there are at least two triangles of B incident with v. Let e and f be consecutive edges on $\partial(B)$, meeting at v, and labelled by a and b respectively. Since $\mathbf{RSym}$ succeeds on $\mathcal{P}$, by Proposition 2.6.12 Γ contains no loops, so the vertices of each red triangle of Γ are pairwise distinct. Also, as Γ has minimal coloured area, by [34, Proposition 3.8] Γ is semi-σ-reduced. Therefore, as the edges of B incident with v other than e and f are all interior to B (since ef is a sub-path of $\partial(B)$), by the proof of [34, Lemma 3.15] we can retriangulate B at v, without changing its boundary and area, to reduce the number of triangles of B incident with v to two. So let Φ be the sub-diagram of Γ equal to the union of the two triangles of B incident with v. Then Φ has boundary word $w := abcd$. By Part 1 we have $a \neq b^\sigma$ and $(a, b) \in D(P)$. Now $w =_{U(P)} 1$, and by Theorem 2.3.11 $[ab] \neq_{U(P)} 1$, so $c \neq d^\sigma$, and $cd =_P [ab]^\sigma$. By Axiom P2 we have $(b^\sigma, a^\sigma), (d^\sigma, c^\sigma) \in D(P)$, and hence we can retriangulate Φ at v (see Figure 4.16), without changing its boundary and area, to reduce the number of triangles of B incident with v to one, as required. ■

Recall Definition 3.1.1 of cutting an annular diagram Γ_A open along a path p, resulting in a simply-connected diagram Γ. From the definition it follows that an image of a bad red blob of Γ_A is a simply-connected red blob of Γ; that if p passes through the interior of a simply-connected red blob $B \subseteq \Gamma_A$, then the image of B are two simply-connected red blobs of Γ; and that the labels of the images of all edges not lying on p do not change. If $x \in \Gamma_A$, then we shall denote its image in Γ by x' (we shall never work with more than one object in the image of x). Recall Definition 2.5.14 of the 1-skeleton of a coloured diagram.

Lemma 4.6.16. *Let Γ_A be an annular diagram of minimal coloured area defined over a finite pregroup presentation, with boundary words w_1 and w_2, and with a path $p \in \Gamma_A^1$ intersecting both ω and τ. Let Γ be the simply-connected diagram resulted by cutting Γ_A open along p. Then Γ has minimal coloured area.*

Proof. Since p intersects both ω and τ, it follows that p has label some $\alpha \in X^*$ such that Γ has boundary word $W_\alpha := \alpha w_2 \alpha^\sigma w_1^\sigma$. Suppose that Γ does not have minimal coloured area. Then there exists a coloured simply-connected diagram Δ with boundary word W_α such that $\mathbf{CArea}(\Delta) < \mathbf{CArea}(\Gamma)$. Let π_α and π_{α^σ} be sub-paths of $\partial(\Delta)$ that make up α and α^σ respectively. Then it is possible to identify π_α with π_{α^σ} to obtain a coloured annular diagram Δ_A with boundary words w_1 and w_2. But

$$\mathbf{CArea}(\Delta_A) = \mathbf{CArea}(\Delta) < \mathbf{CArea}(\Gamma) = \mathbf{CArea}(\Gamma_A),$$

a contradiction. ■

Lemma 4.6.17. *Let Γ_A be a green-rich coloured annular diagram over $\mathcal{P}$, with boundary words w_1 and w_2, and of minimal coloured area. Assume that the following statements hold.*

(a) The boundaries of Γ_A are simple closed paths.

(b) Γ_A contains a red blob B satisfying one of the following statements.

(i) B is bad, so $\Gamma_A \setminus B^\circ$ decomposes as a disjoint union of annular diagrams S and S'.

(ii) B is simply-connected, B is contained in an island E of Γ_A, B contains both endpoints of E, and $E \setminus B^\circ$ decomposes as an edge-disjoint union of two annular or simply-connected diagrams S and S' such that $\partial(B) \cap \partial(S)$ and $\partial(B) \cap \partial(S')$ are simple paths.

Then for each vertex $v \in \partial(B) \cap \partial(S)$, there exists a retriangulation B_1 of B with an internal edge e such that $v \in \partial(e)$ and $\overline{e} \cap \partial(S') \neq \emptyset$.

Proof. First note that since Γ_A is green-rich, all vertices of B lie on $\partial(B)$. We claim that there is an internal edge of B meeting both $\partial(S)$ and $\partial(S')$. Suppose first that B is bad. Then since Γ_A is connected and $S \cap S' = \emptyset$, there is such an edge.

Now assume that B is simply-connected. Let v_1 be an endpoint of E. Since B contains both endpoints of E and $E \setminus B^\circ$ decomposes as S and S', we have $v_1 \in \partial(B)$ and $v_1 \in \partial(S) \cap \partial(S')$. Since Γ_A is green-rich and B is simply-connected, by Lemma 3.2.6 we have $|\partial(B)| \geq 3$, hence either $|\partial(S) \cap \partial(B)| \geq 2$ or $|\partial(S') \cap \partial(B)| \geq 2$. It follows that if v_1 is incident with exactly one triangle T of B, then T contains an internal edge in B, as required. So suppose that v_1 is incident with at least two triangles of B. Then there is an internal edge e in B with endpoints v_1 and some vertex of $\partial(S) \cup \partial(S')$. So as $v_1 \in \partial(S) \cap \partial(S')$, e meets both $\partial(S)$ and $\partial(S')$.

Let $v \in \partial(B) \cap \partial(S)$ be a vertex. By our claim there is a vertex of $\partial(B) \cap \partial(S)$ for which $B_1 = B$ satisfies the lemma. The proof is by induction on the length n of a path in $\partial(B) \cap \partial(S)$ from v to such a v_1. The base case is $n = 0$, so $v = v_1$. Assume that the lemma holds for $n-1$, and that there is a vertex v_1 of $\partial(B) \cap \partial(S)$ such that there is a path $p \subseteq \partial(B) \cap \partial(S)$ of length n between v and v_1, and such that there is a retriangulation B_1 of B with an internal edge e such that $v_1 \in \partial(e)$ and $\overline{e} \cap \partial(S') \neq \emptyset$. Let Δ_A be the diagram with B replaced by B_1. Since S and S' are connected, there is a path $r \in \Delta_A^1$ meeting both boundaries of Δ_A, such that $r \cap B_1^\circ = e$, and cutting Δ_A open along r gives us a coloured simply-connected diagram Γ. Let $B_1' \subseteq \Gamma$ be the red blob in the image of B_1 that contains p'. Let v_1' be the image of v_1 with $v_1' \in B_1'$, let e' be the image of e with $e' \subseteq B_1'$, and let $f' \subseteq p'$ be an edge incident with v_1'. If $\partial(B) \cap \partial(S)$ is a loop, then $v = v_1$, so we can assume that $|\partial(B) \cap \partial(S)| \geq 2$. Then note that f is an edge incident with e, and since by Lemma 4.2.4 and Condition (ii) of the lemma $\partial(B) \cap \partial(S)$ is a simple path, from $r \cap B_1^\circ = e$ we deduce that $e'f'$ or $f'e'$ is a sub-path of $\partial(B_1')$.

Now Δ_A has minimal coloured area, so by Lemma 4.6.16 Γ has also minimal coloured area; and as all vertices of B lie on $\partial(B)$, the same holds for B_1 and B_1'. Thus, by Lemma 4.6.15 we can without loss of generality assume that there is precisely one triangle of B_1' incident with v_1'. Hence as $e'f'$ or $f'e'$ is a sub-path of $\partial(B_1')$, the pre-image v_2 of the endpoint of f' distinct from v_1' satisfies the lemma, and as $f' \subseteq p'$, it follows that v_2 is at distance at most $n-1$ from v in $(\partial(B) \cap \partial(S))^1$, so by induction the lemma holds. ■

We can now present the main result of this subsection. Recall that we label all edges of the boundary of a red blob B with respect to the orientation from B.

Proposition 4.6.18. *Let $\Gamma_A \in \mathcal{T}$ be defined over $\mathcal{P}$, and of minimal coloured area. Assume that Γ_A contains a bad or complicated red blob B, and B is edge-incident with an internal green face F with an edge on ω.*

Then there exists a retriangulation B_1 of B in which there is an internal edge f incident with a common edge e of F and B, where the labels y of f and x^σ of e satisfy both of the following statements.

1. *ef is a path, and $x \neq y^\sigma$ and $(x, y) \in D(P)$.*

2. *If the other endpoint of f is not on τ, then there is an internal green face F_1 with an edge $g \subseteq \partial(B)$, such that fg is a path, the label z^σ of g satisfies $(y, z) \in D(P)$, and no sub-word of xyz is trivial in $U(P)$. Moreover, F_1 has an edge on τ.*

In particular, there is a path between ω and τ of length at most $2r+1$, where $r = \max\{|R| : R \in \mathcal{R}\}$ is the length of the longest green relator.

Proof. By Theorems 4.6.2, 4.6.13 and 4.6.14 $\partial(B) \cap \partial(F)$ is a single edge e (labelled x^σ, say), and B satisfies Condition (b) of Lemma 4.6.17. Hence as by Axiom T_1 the boundaries of Γ_A are simple closed paths, by Lemma 4.6.17 there exists a retriangulation B_1 of B in which there is an internal edge f incident with e such that ef is a path, and f meets the boundary of the diagram of the decomposition of $Y \setminus B^\circ$ (for some $Y \in \{\Gamma_A, E\}$) that does not contain F. Let Δ_A be the diagram with B replaced by B_1. Applying Theorems 4.6.2, 4.6.13 and 4.6.14 again shows that there is a path $pef \in \Delta_A^1$, where $p \subseteq \partial(F)$ and $p \cap \omega \neq \emptyset$; and if $f \cap \tau = \emptyset$, then there is a path $p_1 \in \Delta_A^1$ such that fp_1 is a path, and $p_1 \cap \partial(B)$ and $p_1 \cap \tau$ are both single vertices.

Suppose first that $f \cap \tau \neq \emptyset$. Cut Δ_A open along pef to obtain a simply-connected diagram Γ. Then there is a simply-connected red blob $B_1' \subseteq \Gamma$ in the image of B_1 that contains e'. Hence let y be the label of f such that xy is a sub-word of a boundary word of B_1'. Since Δ_A is green-rich, all vertices of B_1 lie on $\partial(B_1)$, and the same holds for B_1'. Also, as Δ_A has minimal coloured area, by Lemma 4.6.16 Γ has also minimal coloured area, so by Lemma 4.6.15 we have $x \neq y^\sigma$ and $(x, y) \in D(P)$, as required.

Now assume that $f \cap \tau = \emptyset$. Since $p_1 \cap \partial(B)$ and $p_1 \cap \tau$ are both single vertices, there is an internal green face F_1 with edges $g \subseteq \partial(B)$ and $h \subseteq \tau$, such that fg is a path; and cutting Δ_A open along $pefp_1$ results in a simply-connected diagram Γ, where there is a red blob $B_1' \subseteq \Gamma$ in the image of B_1 that contains e' and g'. Then similarly as in the previous paragraph we can show that Γ and B_1' satisfy assumptions of Lemma 4.6.15, hence by Lemma 4.6.15 no proper sub-word of any (cyclic) boundary word of B_1' is equal to 1 in $U(P)$. Since e and g lie on internal green faces, we have $|\partial(B_1')| \geq 4$. So letting y and z^σ be the labels of f and g such that xyz is a sub-word of a boundary word of B_1', we have that no sub-word of xyz is trivial in $U(P)$. The proposition follows. ■

Chapter 5

Thickness of diagrams in $\mathcal{T}$

5.1 Results and key definitions

Recall Definitions 4.6.1 and 4.6.8 of what it means for a red blob to be complicated, or highly hyperbolic. Recall also Definitions 4.3.2 and 4.3.4 of the curvature neighbourhoods $\mathcal{S}_F$ and $\mathcal{B}_F$ of an internal green face. Our main goal in this chapter is to prove the following theorem. Throughout this chapter we will prove results about boundaries that are symmetric in ω and τ, so we let $\{\rho, \rho'\} = \{\omega, \tau\}$.

Theorem 5.1.1. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$ with an edge on ρ, and suppose that all red blobs of Γ_A are simply-connected, and none of them are complicated. Then at least one of the following statements holds.*

1. $\partial(F) \cap \rho' \neq \emptyset$.
2. *F is edge-incident with a red blob B with an edge on ρ'. Moreover, either B is highly hyperbolic and $|\partial(B)| \leq 6$, or B contains precisely one boundary edge and $|\partial(B)| \leq 5$.*
3. *There is an internal green face F' with $\partial(F') \cap \partial(F) \neq \emptyset$, with an edge on ρ', and either F' and F are edge-incident, or $\mathcal{S}_{F'} \cap \mathcal{S}_F$ contains a red blob B with no boundary edge and $|\partial(B)| \leq 4$.*

Throughout this chapter we shall study faces that satisfy the following two definitions.

Definition 5.1.2. Let $\Gamma_A \in \mathcal{T}$. We say that a face F of Γ_A is *thin with respect to* ρ if F is an internal green face, F has an edge on ρ, and both of the following conditions hold:

(a) $\partial(F) \cap \rho' = \emptyset$.

(b) F is not edge-incident with any red blob that has an edge on ρ'.

We say that F is *thin* if F is thin with respect to some $\rho \in \{\omega, \tau\}$.

If F is a thin face with respect to ρ, then a *a boundary link* of F is an internal green face F' with $\partial(F') \cap \partial(F) \neq \emptyset$ and $\partial(F') \cap \rho' \neq \emptyset$.

Definition 5.1.3. Let $\Gamma_A \in \mathcal{T}$. We say that an internal green face F is *pre-neighbourly with respect to* ρ if F is thin with respect to ρ, no element of $\mathcal{S}_F$ is curvature incident more than once with F, and each vertex of $\mathcal{B}_F$ has green degree 3.

We say that F is *neighbourly with respect to* ρ if F is thin with respect to ρ, and F has no boundary links.

We say that F is *(pre-)neighbourly* if F is (pre-)neighbourly with respect to some $\rho \in \{\omega, \tau\}$.

The chapter is structured as follows. In Section 5.2 we prove some auxiliary results about faces F curvature incident more than once with some element of $\mathcal{S}_F$. Suppose that all red blobs of $\Gamma_A \in \mathcal{T}$ are simply-connected, and that none of them are complicated. In Section 5.3 we describe curvature neighbourhoods of neighbourly and pre-neighbourly faces of Γ_A, and prove that all neighbourly faces of Γ_A are pre-neighbourly. In Section 5.4 we prove that Γ_A has no neighbourly faces. In Section 5.5 we prove Theorem 5.1.1. Finally, we apply Theorem 5.1.1 to prove the Three Face Theorem in Section 5.6.

5.2 Faces which are curvature incident more than once with a blob or vertex

In this section we collect several preliminary results that we shall use throughout this chapter.

Lemma 5.2.1. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Then the following statements hold.*

1. *Suppose that F is incident more than once with a vertex v. By Lemma 4.2.3 $\Gamma_A \setminus (\overline{F})^{\circ}$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2. Assume further that for some $i \in \{1, 2\}$, the only face in Γ_i incident with v is a red blob B. Then B is edge-incident twice with F.*

2. *Suppose that F satisfies at least one of the following two conditions.*

   (i) F is edge-incident with itself.

   (ii) $\partial(F) \cap \partial(B)$ contains two edges that are consecutive on $\partial(B)$ for some blob B, or F is incident more than once with a vertex of green degree 3.

   Then F is curvature incident more than once with at least two elements of $\mathcal{S}_F$.

Proof. For Part 1, by Proposition 4.2.8 B is edge-incident at most twice with F. So suppose for a contradiction that B is edge-incident once with F. Then as all faces of Γ_i incident with v are contained in B, the common edge of F and B is an interior loop labelled by a V^{σ}-letter, contradicting Axiom T_5. Hence Part 1 follows.

For Part 2, assume first that F is a edge-incident with itself. Then by Lemma 4.2.3 there is a consolidated edge $f \subseteq \partial(F)$ with distinct endpoints v_1 and v_2 that are incident twice with F. Now for $1 \leq i \leq 2$, either $v_i \in \mathcal{S}_F$, or $\delta_G(v_i) = 2$, and by Part 1 there is a red blob incident with v_i and edge-incident twice with F, so we are done.

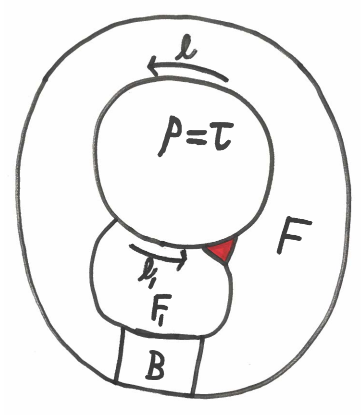

Figure 5.1: The face F_1 with $\mathcal{S}_{F_1} = \mathcal{S}_F$, see the proof of Lemma 5.2.3.

Now assume that F is not edge-incident with itself. By Part 1 if F is incident more than once with a vertex of green degree 3, then Part 2 holds. So assume that $\partial(F) \cap \partial(B)$ contains two edges that are consecutive on $\partial(B)$ for some blob B. Then by Lemma 4.2.6 F is incident more than once with some vertex $v \in \partial(B)$. Hence if $\delta_G(v) \geq 3$, then the lemma holds, so assume $\delta_G(v) = 2$. Then by Part 1 there is a red blob $B_1 \neq B$ edge-incident twice F. ■

Lemma 5.2.2. *Suppose that F is an internal green face of $\Gamma_A \in \mathcal{T}$ that is curvature incident more than once with some $x \in \mathcal{S}_F \setminus \mathcal{B}_F$, and no element of $\mathcal{B}_F$ is curvature incident more than once with F. Then $|\mathcal{B}_F| = 2$.*

Proof. Since no element of $\mathcal{B}_F$ is curvature incident more than once with F, Part 3 of Lemma 4.3.4 shows that $|\mathcal{B}_F| \geq 2$. By Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$. Hence as by Part 2 of Lemma 4.3.4 each $y \in \mathcal{B}_F$ has $\chi(y, F, \Gamma_A) \leq -1/4$, and by Lemmas 2.6.10 and 3.2.8 we have $\chi(x, F, \Gamma) \leq -1/3$, it follows that $|\mathcal{B}_F| \leq 2$. ■

Lemma 5.2.3. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, and that Γ_A contains an internal green face F edge-incident more than once with some red blob. Then there is at least one element of $\mathcal{B}_F$ curvature incident exactly twice with F.*

Proof. Let B be the blob from the statement. By Axiom T_6, F has a boundary (consolidated) edge l on $\rho \in \{\omega, \tau\}$ with $|l| \geq 1$. By Proposition 4.2.8 Γ_A is island-free: so by Theorem 4.5.13 we have $\partial(F) \cap \rho = l$; and every element of $\mathcal{S}_F$ is curvature incident at most twice with F. Hence suppose for a contradiction that all elements of $\mathcal{S}_F$ that are curvature incident twice with F are in $\mathcal{S}_F \setminus \mathcal{B}_F$, so in particular $B \in \mathcal{S}_F \setminus \mathcal{B}_F$. Then Part 4 of Lemma 4.3.4 implies $\chi(B, F, \Gamma_A) \geq -1/2$, and Lemma 3.2.8 gives $\mathbf{Area}(B) \leq 2$.

By Proposition 4.2.8 $\Gamma_A \setminus (\overline{F} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Let Γ_1 be the diagram that contains ρ, and let $\Gamma = \Gamma_1 \cup \overline{F} \cup B$. By Lemma 5.2.2 we can let $\mathcal{B}_F = \{x, y\}$, and let e and f be the common edges of B and F. Assume first that $\mathbf{Area}(B) = 2$. Then by Lemmas 3.2.6 and 3.2.8 we have $|\partial(B)| = 4$ and $\chi(B, F, \Gamma_A) \leq$

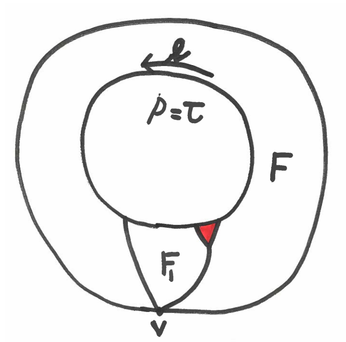

Figure 5.2: The face F_1 with $\mathcal{S}_{F_1} \subseteq \{v, x, y\}$, see the proof of Lemma 5.2.4.

$-1/2$. Since $\kappa_{\Gamma_A}(F) = 0$, by Lemmas 2.6.10 and 3.2.8 we have

$$\mathcal{S}_F = \{x, y, B\},$$

B contains no boundary edge, and each $z \in \{x, y\}$ is either a red triangle with one boundary edge, or a vertex of green degree 3. In particular, by Part 1 of Lemma 4.3.4 z lies on or is incident with l. Suppose first that e and f are not consecutive edges on $\partial(B)$, and let g be an edge of $\partial(B)$ with $g \subseteq \Gamma^\circ$ distinct from e and f. Note that g is unique since $|\partial(B)| = 4$. Let F_1 be the green face that shares g with B. We have $F_1 \subseteq \Gamma_1$, so by Axiom T_6, F_1 has a consolidated edge l_1 on ρ with $|l_1| \geq 1$. By Theorem 4.5.13 we have $\partial(F_1) \cap \rho = l_1$. Hence as each $z \in \{x, y\}$ is either a red triangle, or a vertex with $\delta_G(z) = 3$, we have $\mathcal{S}_{F_1} = \mathcal{S}_F$ (see Figure 5.1). But B is edge-incident only once with F_1, so

$$\kappa_{\Gamma_A}(F_1) = 1 - 1/4 - 1/4 - 1/4 = 1/4 < 0,$$

a contradiction.

Hence e and f are consecutive on $\partial(B)$. Then by Lemma 5.2.1 F is curvature incident more than once with at least two elements of $\mathcal{S}_F$, contradicting $\mathcal{S}_F = \{x, y, B\}$. This concludes the argument when $\mathbf{Area}(B) = 2$.

Suppose that B is a red triangle. Then e and f are consecutive on $\partial(B)$, so by Lemma 5.2.1 F is curvature incident more than once with at least two elements of $\mathcal{S}_F \setminus \mathcal{B}_F$. But these give F curvature of at most $-1/3 - 1/3 = -2/3$, contradicting Part 4 of Lemma 4.3.4. ■

Lemma 5.2.4. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, and that Γ_A contains an internal green face F curvature incident more than once with some element of $\mathcal{S}_F$. Then at least one element of $\mathcal{B}_F$ is curvature incident exactly twice with F.*

Proof. By Axiom T_6, F has a boundary (consolidated) edge l on $\rho \in \{\omega, \tau\}$ with $|l| \geq 1$. Now $\Gamma_A \in \mathcal{T} \setminus \mathcal{U}$, so if Γ_A contains an island E, then by Proposition 4.5.1 $E = \Gamma_A$ and some green face contained in E is incident with both of its endpoints, hence by Lemma 4.5.2 the lemma holds. So we may assume that Γ_A is island-free. Hence by Theorem 4.5.13 we

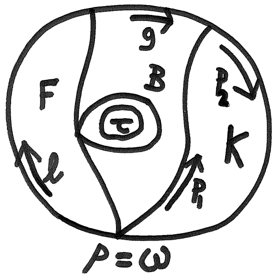

Figure 5.3: K bounded by the path p_1p_2, see the proof of Lemma 5.2.6.

have $\partial(F) \cap \rho = l$. By Proposition 4.2.8 every element of $\mathcal{S}_F$ is curvature incident at most twice with F, so suppose for a contradiction that no element of $\mathcal{B}_F$ is curvature incident twice with F. Then by Lemma 5.2.3 no red blob is edge-incident twice with F. So our assumption and Part 4 of Lemma 4.3.4 imply that there is an interior vertex v incident twice with F and with $\chi(v, F, \Gamma_A) \geq -1/2$, so by Lemma 2.6.10 $\delta_G(v) \leq 4$. By Lemma 5.2.2 we can let $\mathcal{B}_F = \{x, y\}$.

By Lemma 4.2.3 $\Gamma_A \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of two annular diagrams, so let Γ_1 be the one with $\rho \subseteq \Gamma_1$, and let $\Gamma = \Gamma_1 \cup \overline{F}$. Suppose that $\delta_G(v) = 2$. Then by Lemma 5.2.1 there is a red blob edge-incident twice with F, contradicting the last paragraph. So $\delta_G(v) \in \{3, 4\}$.

Assume next that $\delta_G(v) = 4$. Then $\chi(v, F, \Gamma_A) = -1/2$. Since $\kappa_{\Gamma_A}(F) = 0$, by Lemmas 2.6.10 and 3.2.8 we have $\mathcal{S}_F = \{v, x, y\}$, and each $z \in \{x, y\}$ is either a red triangle or a vertex of green degree 3. Suppose that v is the only vertex incident twice with F. Since $\delta_G(v) = 4$, there is an internal green face $F_1 \subseteq \Gamma_1$ incident with v. As each $z \in \{x, y\}$ is either a red triangle or a vertex of green degree 3, we deduce that $\mathcal{S}_{F_1} \supseteq \{v, x, y\}$ (see Figure 5.2). By Theorem 4.5.13 $\partial(F_1) \cap \rho$ is a single consolidated edge, hence $\mathcal{S}_{F_1} = \{v, x, y\}$. But v is incident only once with F_1, so

$$\kappa_{\Gamma_A}(F_1) = 1 - 1/4 - 1/4 - 1/4 = 1/4 > 0,$$

a contradiction.

Assume instead that there exists a vertex $v_1 \neq v$ incident twice with F. By Lemma 4.2.3 F is edge-incident with itself. Hence by Lemma 5.2.1 F is curvature incident more than once with at least two elements of $\mathcal{S}_F$, contradicting $\mathcal{S}_F = \{v, x, y\}$.

Finally, suppose that $\delta_G(v) = 3$. Then by Lemma 5.2.1 F is curvature incident more than once with at least two elements of $\mathcal{S}_F \setminus \mathcal{B}_F$, which give F curvature of at most $-1/3 - 1/3 = -2/3$, contradicting Part 4 of Lemma 4.3.4. ■

Recall that if we say that a closed path p is of the form $p_1p_2 \ldots p_n$, then p is a sequence of simple sub-paths p_i. Recall also Definition 2.5.14 of the 1-skeleton of a coloured diagram.

In the statement of the next lemma we allow tracing boundaries of faces in both directions.

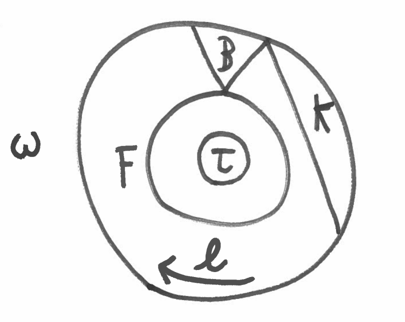

Figure 5.4: Depiction of a case where $\partial(F) \cap \rho$ is not a single consolidated edge, see the proof of Lemma 5.2.6.

Lemma 5.2.5. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains an internal green face F and a red blob $B \in \mathcal{S}_F$ satisfying all of the following conditions.*

1. *F contains a consolidated edge $l \subseteq \rho \in \{\omega, \tau\}$, and B contains at most two boundary edges.*

2. *$\Gamma_A \setminus (\overline{F} \cup B)^\circ$ contains a simply-connected sub-diagram K with $\mathbf{CArea}(K) \geq (1,0)$, that is bounded by a closed path $p \in \Gamma_A^1$ of the form $p_1p_2p_3$, where p_1 is a sub-path of $\partial(F)$ with $|p_1| \geq 1$, p_2 is a sub-path of $\partial(B)$ with $|p_2| \geq 1$, and p_3 is a sub-path of ρ such that $p_1 \cap p_3 \cap l$ contains an endpoint of l (see Figure 5.5 with p_1, p_2 and p_3 replaced by r_1, r_2 and p_2 respectively).*

3. *$\partial(B) \cap \rho$ contains an edge g such that $p_2 g$ is a sub-path of $\partial(B)$, and $\partial(F) \cap \partial(B)$ contains an edge e such that $p_1 e$ is a sub-path of $\partial(F)$ and ep_2 is a sub-path of $\partial(B)$.*

Then B contains two boundary edges.

Proof. Let v be an endpoint of l with $v \in p_1 \cap p_3 \cap l$. Since K is simply-connected and $\mathbf{CArea}(K) \geq (1,0)$, by Axiom T_6 we have $|p_3| \geq 1$. Hence as $g \subseteq \rho$, the common endpoint of g and p_2 does not lie on $\omega \cap \tau$; and if there is no red blob in $\mathcal{B}_F \cap K$ edge-incident with F at v, then v is incident with some internal green face of K, and $\delta_G(v) \geq 3$. In particular, there is an $x \in \mathcal{B}_F \cap K$ curvature incident with F at v. We first show that F, B, x are a neighbourhood of ρ (see Definition 4.5.10). Suppose first that $x = v$. Then K satisfies all assumptions of Lemma 4.5.12, hence by Lemma 4.5.12 F, B, v are a neighbourhood of ρ, and K is a sub-diagram bounded by F, B, v and ρ.

Now suppose that x is a red blob B_1. Since $B_1 \in K$, we deduce that B_1 is not highly hyperbolic. Hence as B_1 is not complicated, B_1 contains at most one boundary edge, so by Proposition 4.6.10 $\partial(B_1) \cap \rho$ is a single edge g_1. Furthermore, since K is simply-connected and $p_1 = \partial(K) \cap \partial(F)$ is a sub-path of $\partial(F)$, by Lemma 4.5.8 $\partial(B_1) \cap \partial(F)$ is a single edge e_1. Let

$$K_1 := K \setminus B_1.$$

Note that K_1 is simply-connected, and $\mathbf{CArea}(K) \geq (1,0)$ implies that $\mathbf{CArea}(K_1) \geq (1,0)$. So as by Axiom T_6 each internal green face contains a boundary edge and $\partial(B_1) \cap \rho$ and $\partial(B_1) \cap \partial(F)$ are single edges, it follows that K_1 is bounded by a closed path $p' \in \Gamma_A^1$ of the form $r_1 r_2 p_2 r_3$, where r_1 is a sub-path of $\partial(B_1)$ with $|r_1| \geq 1$, r_2 is a sub-path of p_1 such that if $|r_2| < 1$, then F is incident precisely once with r_2, and r_3 is a sub-path of p_3 with $|r_3| \geq 1$. Furthermore, $g_1 r_1$ and $r_1 e_1$ are sub-paths of $\partial(B_1)$, $e_1 r_2$ and $r_2 e$ are sub-paths of $\partial(F)$, and the common endpoint of g_1 and r_1 does not lie on $\omega \cap \tau$, since $g_1 \subseteq \rho$ and $|r_3| \geq 1$. So we can apply Lemma 4.5.12 to deduce that F, B, B_1 are a neighbourhood of ρ, and that K_1 is a sub-diagram bounded by F, B, B_1 and ρ.

Since the sub-diagram L bounded by F, B, x and ρ is simply-connected, by Axiom T_6 and Theorem 4.5.13 $\partial(D) \cap \rho$ is a single consolidated edge for all green faces $D \subseteq L$. Thus, by Lemma 4.5.11 B contains two boundary edges. ■

Lemma 5.2.6. *Let F be an internal green face of Γ_A. Assume that the following statements hold.*

1. *All red blobs of Γ_A are simply-connected, and none of them are complicated.*

2. *There is a red blob $B \in \mathcal{B}_F$ edge-incident exactly twice with F, and with one boundary edge.*

Let $g \subseteq \rho$ be the boundary edge of B. Then $\rho = \overline{g} \cup l$, where l is a single consolidated edge of F.

Proof. Since B contains one boundary edge and $B \in \mathcal{B}_F$, we have $\partial(F) \cap \rho \neq \emptyset$. Hence by Theorem 4.5.13 $\partial(F) \cap \rho$ is a single consolidated edge l. By Proposition 4.2.8 Γ_A is island-free, and $\Gamma_A \setminus (\overline{F} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams whose boundaries are simple closed paths. Let Γ be the one with $\rho \subseteq \Gamma$. Suppose for a contradiction that $\rho \neq \overline{g} \cup l$. Since $g \subseteq \rho$ and the boundaries of Γ are simple, Γ is a union of islands and bridges (see Definition 2.5.10). Hence by assumption on ρ it follows that $\mathbf{Area}(\Gamma) > 0$, so Γ contains a simply-connected sub-diagram K bounded by a closed path $p \in \Gamma_A^1$ of the form $p_1 p_2$, where $p_1 \subseteq \partial(F) \cup \partial(B)$ and $|p_1| \geq 1$, and p_2 is a sub-path of ρ such that $p_1 \cap p_2 \cap l$ contains an endpoint of l (see Figure 5.3).

As B contains at most one boundary edge, by Proposition 4.6.10 $\partial(B) \cap \rho = \overline{g}$, and $\partial(B) \cap \rho$ is a sub-path of $\partial(B)$. Hence $|p_1 \cap \partial(F)| \geq 1$. Also, if $|p_1 \cap \partial(B)| < 1$ (see Figure 5.4), then $\partial(F) \cap \rho \neq l$, a contradiction. So write $p_1 = r_1 r_2$, where r_1 is a sub-path of $\partial(F)$ with $|r_1| \geq 1$, and r_2 is a sub-path of $\partial(B)$ with $|r_2| \geq 1$.

By Proposition 4.2.8 $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at most one bounded component, hence $\partial(F) \cap \partial(B)$ contains an edge e such that $r_1 e$ is a sub-path of $\partial(F)$ and $e r_2$ is a sub-path of $\partial(B)$ (see Figure 5.5). Since $|r_2| \geq 1$ and K is simply-connected, we have $\mathbf{CArea}(K) \geq (1,0)$. Also, as $\partial(B) \cap \rho = \overline{g}$ and $\partial(B) \cap \rho$ is a sub-path of $\partial(B)$, we deduce that $r_2 g$ is a sub-path of $\partial(B)$. So by Lemma 5.2.5 B contains two boundary edges, a contradiction. ■

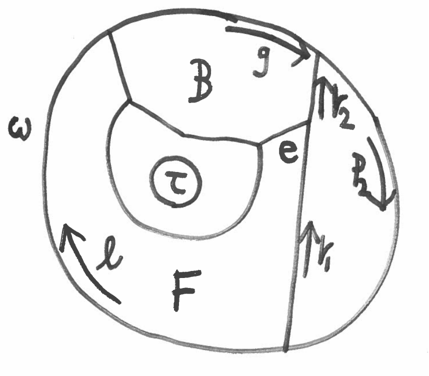

Figure 5.5: K bounded by the path $r_1r_2p_2$, see the proof of Lemma 5.2.6.

5.3 Neighbourly and pre-neighbourly faces

We will prove Theorem 5.1.1 by describing curvature neighbourhoods of a pre-neighbourly face (see Definition 5.1.3). Let us therefore dedicate this section to their analysis: our main result is Theorem 5.3.8. Recall Definition 2.6.8 that $\chi(x, \Gamma)$ is the curvature that x gives to a single internal green face across each curvature incidence.

Lemma 5.3.1. *Let $\Gamma_A \in \mathcal{T}$ contain a thin face F with respect to ρ. Then the following statements hold.*

1. *$\partial(F) \cap \partial(\Gamma_A)$ is a single consolidated edge l with $|l| \geq 1$, and F is not incident with any vertex of $\rho \cap \rho'$.*

2. *If F is curvature incident more than once with some element of $\mathcal{S}_F$, then Γ_A is island-free.*

3. *If all red blobs of Γ_A are simply-connected and none of them are complicated, then each red blob B of $\mathcal{S}_F$ contains at most one boundary edge; if B contains an edge g on $\rho_1 \in \{\omega, \tau\}$, then $\partial(B) \cap \rho_1 = \overline{g}$ and $\partial(B) \cap \rho_1$ is a sub-path of $\partial(B)$; and all $x \in \mathcal{S}_F$ satisfy $\chi(x, \Gamma_A) > -1/2$.*

Proof. Part 1. Follows immediately from Theorem 4.5.13 and definition of a thin face (see 5.1.2).

Part 2. Suppose for a contradiction that Γ_A contains an island E. Then by Proposition 4.5.1 $\Gamma_A = E$, and E contains a green face incident with both of its endpoints. But then by Lemma 4.5.2 F is incident with an endpoint of E, so F is not thin, a contradiction.

Part 3. Assume that $B \in \mathcal{S}_F$ is a red blob. Since F is thin, B is not highly hyperbolic. Hence as B is not complicated, B contains at most one boundary edge; and if B contains an edge on $\rho_1 \in \{\omega, \tau\}$, then by Proposition 4.6.10 $\partial(B) \cap \rho_1$ is a single edge. The last statement follows by Lemmas 3.1.9 and 4.6.9. ∎

Lemma 5.3.2. *Let $\Gamma_A \in \mathcal{T}$ contain a thin face F with respect to ρ. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that $\mathcal{B}_F$ contains an element x curvature incident more than once with F.*

Then for some consolidated edge l of F we have $\rho = l$ if x is a vertex, else $\rho = l \cup \overline{g}$, where g is a boundary edge of x. In particular, F is the only internal green face with an edge on ρ.

Proof. By Lemma 5.3.1 $\partial(F) \cap \partial(\Gamma_A) = \partial(F) \cap \rho$ is a single consolidated edge l. Assume first that x is a vertex v. Then v is an endpoint of l, so by Lemma 4.5.16 $\rho = l$. Now assume that x is a red blob B. By Proposition 4.2.8 B is edge-incident exactly twice with F, and by Lemma 5.3.1 B contains at most one boundary edge. So the lemma follows from Lemma 5.2.6. ∎

Lemma 5.3.3. *Let F be a neighbourly face of $\Gamma_A \in \mathcal{T}$, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then F is pre-neighbourly.*

Proof. We first show that no element of $\mathcal{B}_F$ is curvature incident more than once with F. Suppose a contradiction that there is such an $x \in \mathcal{B}_F$. Since F is thin, by Lemma 5.3.1 Γ_A is island-free. Now F has an edge on $\rho \in \{\omega, \tau\}$, and by Lemma 5.3.2 F is the only internal green face with an edge on ρ, so as F is neighbourly, F is not incident with any other internal green face.

Assume that x is a vertex. Then by Lemma 4.2.3 $\Gamma_A \setminus (\overline{F})^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Hence since Γ_A is island-free and F is not incident with any other internal green face, it follows that F is edge-incident with an annular red blob, a contradiction.

Now suppose that x is a red blob B. By Proposition 4.2.8 $\Gamma_A \setminus (\overline{F} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams. Hence as $\partial(F) \cap \rho' = \emptyset$, F is incident with an internal green face $F' \neq F$, a contradiction.

Thus, no element of $\mathcal{B}_F$ is curvature incident more than once with F, so by Lemma 5.2.4 no element of $\mathcal{S}_F$ is curvature incident more than once with F. Now if v is a vertex in $\mathcal{B}_F$, then as F is neighbourly, Corollary 4.5.17 shows that $\delta_G(v) = 3$. ∎

Lemma 5.3.4. *Let F be a pre-neighbourly face of $\Gamma_A \in \mathcal{T}$, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then $|\mathcal{B}_F| = 2$; $|\mathcal{S}_F| \geq 3$; and no red blob of $\mathcal{S}_F \setminus \mathcal{B}_F$ contains a boundary edge.*

Proof. Since F is thin, by Lemma 5.3.1 $\partial(F) \cap \partial(\Gamma_A)$ is a single consolidated edge $l \subseteq \rho \in \{\omega, \tau\}$ with endpoints $v, w \notin \rho'$. As by assumption no element of $\mathcal{S}_F$ is curvature incident more than once with F, we have that $\partial(F)$ does not pass more than once through any vertex. Hence as $\partial(F) \cap \partial(\Gamma_A) \subseteq \rho \setminus \rho'$, we have $|\mathcal{B}_F| \leq 2$. Furthermore, by Part 3 of Lemma 4.3.4 we have $|\mathcal{B}_F| \geq 2$, so $|\mathcal{B}_F| = 2$. By Lemma 5.3.1 $\chi(x, F, \Gamma_A) > -1/2$ for all $x \in \mathcal{S}_F$, and by Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$, hence $|\mathcal{S}_F| \geq 3$.

To prove the last statement, suppose for a contradiction that there is such a blob B. By Lemma 5.3.1 B contains precisely one boundary edge, and $\partial(B) \cap \rho$ is a single edge g. By

Figure 5.6: A case where $\partial(B)$ passes more than once through some vertex, see the proof of Lemma 5.3.4.

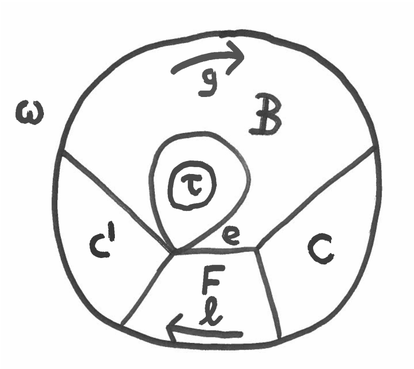

Lemma 4.2.3 $\partial(B)$ does not pass more than twice through any vertex, and it passes more than once through at most one vertex. Suppose first that such a vertex w exists. Then by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams Γ_1 and Γ_2 with $F \subseteq \Gamma_1$ (say), and the boundaries of Γ_1 are simple closed paths. Since $\partial(B) \cap \rho = \overline{g}$, Γ_1 is a union of islands and bridges (see Definition 2.5.10), so as $\partial(F)$ does not pass more than once through any vertex; $\partial(B) \cap \rho$ is a single edge; and $\partial(F) \cap \partial(\Gamma_A) = l$, we deduce that $\Gamma_A \setminus (\overline{F} \cup B)$ contains two components C such that $\overline{C}$ is a simply-connected sub-diagram of Γ_A bounded by a closed path $p \in \Gamma_A^1$ (see Definition 2.5.14) of the form $p_1p_2p_3$, where p_1 is a sub-path of $\partial(F)$ with $|p_1| \geq 1$, $p_2 \subseteq \partial(B)$ and $|p_2| \geq 1$, and p_3 is a sub-path of ρ such that $p_1 \cap p_3 \cap l$ contains an endpoint of l (see Figure 5.6). Now the closure of one of these components C does not contain w, hence for this C the path p_2 is in fact a sub-path of $\partial(B)$. Moreover, since by Proposition 4.2.8 $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at most one bounded component, it follows that $\partial(F) \cap \partial(B)$ contains an edge e such that p_1e is a sub-path of $\partial(F)$ and ep_2 is a sub-path of $\partial(B)$; and as $\partial(B) \cap \rho$ is just g, p_2g is a sub-path of $\partial(B)$. The simply-connectedness of $\overline{C}$ and $|p_2| \geq 1$ imply that $\mathbf{CArea}(\overline{C}) \geq (1,0)$, so by Lemma 5.2.5 B contains two boundary edges, a contradiction.

Now suppose that B does not pass more than once through any vertex. Then $\partial(B)$ is a simple closed path, hence as $\mathbb{R}^2 \setminus (\overline{F} \cup B)$ contains at most one bounded component (and if it exists, then by Proposition 4.2.8 it contains the external face with boundary τ), we have that $\Gamma_A \setminus (\overline{F} \cup B)$ contains a component C such that $\overline{C}$ is a simply-connected sub-diagram of Γ_A bounded by a closed path $p \in \Gamma_A^1$ of the form $p_1p_2p_3$, where p_1 and p_3 have the same properties as before, p_2 is a sub-path of $\partial(B)$ with $|p_2| \geq 1$, p_2g is a sub-path of $\partial(B)$, and $\partial(F) \cap \partial(B)$ contains an edge e such that p_1e is a sub-path of $\partial(F)$ and ep_2 is a sub-path of $\partial(B)$ (see Figure 5.7). So Lemma 5.2.5 gives us a contradiction again. ■

Lemma 5.3.5. *Let F be a pre-neighbourly face of $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that $\mathcal{B}_F$ contains a red blob B*

Figure 5.7: A case where $\partial(B)$ is a simple closed path, see the proof of Lemma 5.3.4.

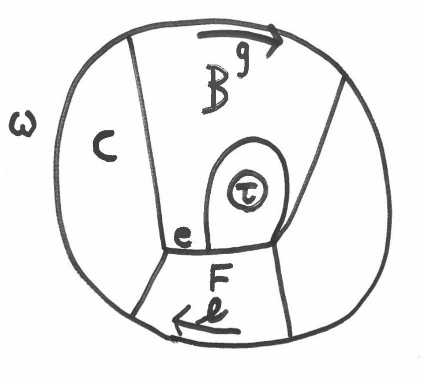

with $\mathbf{Area}(B) \geq 2$*. Then* $\partial(B)$ *does not pass more than once through any vertex,* $\mathbf{Area}(B) = 2$*, and there is an internal green face* F_1 *with the following properties.*

1. *There are edges* $e \subseteq \partial(F)$ *and* $f \subseteq \partial(F_1)$ *such that* ef *or* fe *is a sub-path of* $\partial(B)$*.*

2. $B \notin \mathcal{B}_{F_1}$*.*

Proof. Since F is thin, by Lemma 5.3.1 $\partial(F) \cap \partial(\Gamma_A)$ is a single consolidated edge $l \subseteq \rho \in \{\omega, \tau\}$ with $|l| \geq 1$. Suppose for a contradiction that $\partial(B)$ passes more than once through some vertex v. By Lemma 4.2.3 v is unique and $\partial(B)$ passes through v twice, and by Lemma 4.2.5 $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2, with $F \subseteq \Gamma_1$ (say). By Lemma 5.3.1 B contains one boundary edge and $\partial(B) \cap \rho$ is a single edge h. Since $B \in \mathcal{B}_F$, $\partial(B) \cap \partial(F)$ contains an edge e such that eh or he is a sub-path of $\partial(B)$ (without loss of generality assume that it is eh), and as F is pre-neighbourly, e is the only edge of $\partial(B) \cap \partial(F)$. By Lemma 3.2.6 we have $|\partial(B)| \geq 4$.

Suppose that $|\partial(B)| = 4$. By Lemma 5.3.4 we have $\mathcal{S}_F \setminus \mathcal{B}_F \neq \emptyset$. Hence there is an edge $g \subseteq \partial(B) \cap \partial(\Gamma_1)$ with $g \notin \{e, h\}$. Since $|\partial(B)| = 4$ and $|\partial(B) \cap \partial(\Gamma_2)| \neq 0$, we deduce that $\partial(B) \cap \partial(\Gamma_2)$ is a single edge g_1; and as F is thin, $g_1 \not\subseteq \rho'$. So g_1 is an interior loop labelled by a V^σ-letter, contradicting Axiom T_5.

Hence $|\partial(B)| \geq 5$, so $\mathbf{Area}(B) \geq 3$. Suppose that $v \in \partial(F)$. There is an internal green face $F' \subseteq \Gamma_2$ such that $\partial(F') \cap \partial(B)$ contains an edge with endpoint v (see Figure 5.8). Since $F' \subseteq \Gamma_2$, by Axiom T_6, F' has an edge on ρ', so by Theorem 4.5.13 we have $\partial(F') \cap \rho' = l'$, where l' is a consolidated edge of F' with $|l'| \geq 1$. As $|\partial(B) \cap \partial(F)| = 1$ and $\mathcal{S}_F \setminus \mathcal{B}_F \neq \emptyset$, we have $\delta_G(v) \geq 3$. Now by Corollary 4.5.15 there are elements in $\mathcal{B}_{F'}$ that lie on or are incident with l' that collectively give F' curvature of at most $-1/2$; and as $v \notin \rho'$ and B has no edge on ρ', none of these elements are v or B. Hence by Lemmas 2.6.10 and 3.2.8 we have

$$\begin{aligned}\kappa_{\Gamma_A}(F') &\leq 1 + \chi(B, F', \Gamma_A) + \chi(v, F', \Gamma_A) - 1/2 \\ &\leq 1 - 3/8 - 1/6 - 1/2 = -1/24,\end{aligned}$$

contradicting Axiom T_6.

Hence $v \notin \partial(F)$. Since $|\partial(B) \cap \partial(F)| = 1$ and $\mathcal{S}_F \setminus \mathcal{B}_F \neq \emptyset$, by our choice that eh is a sub-path of $\partial(B)$, we deduce that there is an internal green face $F' \subseteq \Gamma_1$ and an edge $f \subseteq \partial(F') \cap \partial(B)$ such that fe is a sub-path of $\partial(B)$. Suppose that $B \notin \mathcal{B}_{F'}$. Since fe is a sub-path of $\partial(B)$ and $\mathcal{S}_F \setminus \mathcal{B}_F \neq \emptyset$, there is an $x \in \mathcal{S}_{F'} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$. By Lemma 5.3.4 we have $x \notin \mathcal{B}_{F'}$. Hence as $B \notin \mathcal{B}_{F'}$ and

$$\chi(B, F', \Gamma_A) + \chi(x, F', \Gamma_A) \leq -3/8 - 1/6 = -13/24 < -1/2,$$

B and x contradict Part 4 of Lemma 4.3.4.

So $B \in \mathcal{B}_{F'}$. As $F' \subseteq \Gamma_1$, it follows that F' cannot have any edge on ρ', hence by Axiom T_6, F' has an edge on ρ, and by Theorem 4.5.13 $\partial(F') \cap \rho$ is a single consolidated edge l'. Assume first that F' is edge-incident once with B. Then as fe is a sub-path of $\partial(B)$ and $B \in \mathcal{B}_{F'}$, it follows that hf is a sub-path of $\partial(B)$. So since eh is a sub-path of $\partial(B)$, we have $\partial(B) \cap \partial(\Gamma_1) = ehf$, and ehf is a sub-path of $\partial(B)$ (see Figure 5.9). But then $\partial(B)$ cannot pass through v more than once, a contradiction. So by Proposition 4.2.8 F' is edge-incident exactly twice with B. But then since B contains one boundary edge, Lemma 5.2.6 implies $\rho = h \cup l'$, contradicting $\partial(F) \cap \partial(\Gamma_A) = l$.

Hence the first statement of the lemma holds, so $\partial(B)$ is a simple closed path. Applying $\mathcal{S}_F \setminus \mathcal{B}_F \neq \emptyset$ again shows that there is an internal green face F_1 and an edge $f \subseteq \partial(F_1) \cap \partial(B)$ such that ef or fe is a sub-path of $\partial(B)$. Assume for a contradiction that $B \in \mathcal{B}_{F_1}$. Since $|\partial(B)| \geq 4$ and B contains one boundary edge, we have $|\partial(B) \setminus \partial(\Gamma_A)| \geq 3$. Therefore, as ef or fe is a sub-path of $\partial(B)$, by Proposition 4.2.8 B is edge-incident exactly twice with F_1. Hence as B contains one boundary edge, Lemma 5.2.6 gives us a contradiction. So $B \notin \mathcal{B}_{F_1}$. As ef or fe is a sub-path of $\partial(B)$, there is an $x \in \mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$, and by Lemma 5.3.4 we have $x \notin \mathcal{B}_{F_1}$. Hence applying Part 4 of Lemma 4.3.4 to F_1, B, x shows that $\mathbf{Area}(B) = 2$. The lemma follows. ■

Lemma 5.3.6. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains green faces F and F_1 satisfying both of the following conditions.*

1. *F is pre-neighbourly with respect to ρ.*

2. *$(\mathcal{S}_F \setminus \mathcal{B}_F) \cap \mathcal{S}_{F_1} = \{x, y\}$ has size two, and x and y are curvature incident precisely once with F_1 and satisfy $\chi(x, F_1, \Gamma_A) + \chi(y, F_1, \Gamma_A) = -1/2$.*

Then $|\mathcal{B}_{F_1}| = 2$, and each element of $\mathcal{B}_{F_1}$ is curvature incident once with F_1, and is either a red triangle containing one boundary edge, or a boundary vertex of green degree 3 that is not on $\rho \cap \rho'$.

Proof. By Lemma 5.3.4 we have $x, y \notin \mathcal{B}_{F_1}$. Hence as by Axiom T_6 we have $\kappa_{\Gamma_A}(F_1) = 0$, by Part 2 of Lemma 4.3.4 we have $|\mathcal{B}_{F_1}| \leq 2$. So assume for a contradiction that $|\mathcal{B}_{F_1}| = 1$,

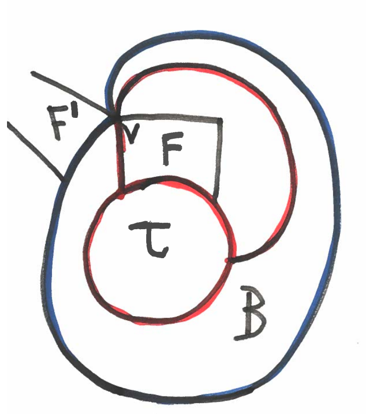

Figure 5.8: A case where $F' \subseteq \Gamma_2$: the red curves are the boundaries of Γ_1, and the blue curve is the inner boundary of Γ_2, see the proof of Lemma 5.3.5.

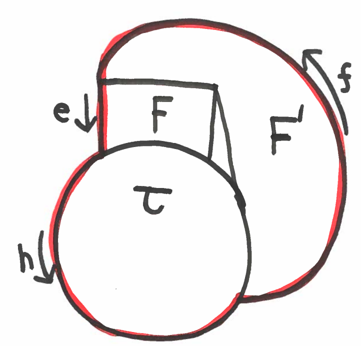

Figure 5.9: A case where $F' \subseteq \Gamma_1$ and B is edge-incident once with F': the red curve depicts the path ehf, see the proof of Lemma 5.3.5.

and let $\mathcal{B}_{F_1} = \{z\}$. Then by Part 3 of Lemma 4.3.4 z is curvature incident exactly twice with F_1. Hence $\mathcal{S}_{F_1} = \{x, y, z\}$ and $\chi(z, F_1, \Gamma_A) = -1/2$, and z is either a red triangle or a vertex of green degree 3. Therefore, F_1 satisfies Condition (ii) of Lemma 5.2.1, so F_1 is curvature incident more than once with at least two elements of $\mathcal{S}_{F_1}$, contradicting $\mathcal{S}_{F_1} = \{x, y, z\}$.

Hence $|\mathcal{B}_{F_1}| = 2$, so by Part 2 of Lemma 4.3.4 each $z \in \mathcal{B}_{F_1}$ has $\chi(z, F_1, \Gamma_A) = -1/4$, and thus the lemma follows from Lemmas 2.6.10 and 3.2.8. ∎

The next lemma characterizes red blobs in the boundary curvature neighbourhood of a pre-neighbourly face.

Lemma 5.3.7. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains a pre-neighbourly face F such that $\mathcal{B}_F$ contains a red blob B. Then B is a red triangle.*

Proof. Assume for a contradiction that $\mathbf{Area}(B) \geq 2$. By Lemma 5.3.1 B contains one boundary edge, on ρ say. By Lemma 5.3.5 $\partial(B)$ does not pass more than once through any vertex; $\mathbf{Area}(B) = 2$; and there exists an internal green face F_1 such that $B \notin \mathcal{B}_{F_1}$, and there are edges $e \subseteq \partial(F)$ and $f \subseteq \partial(F_1)$ such that ef or fe is a sub-path of $\partial(B)$: hence there is an $x \in \mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$. By Lemmas 3.2.8 and 5.3.4 we have $\chi(B, \Gamma_A) = -1/3$ and $x \notin \mathcal{B}_{F_1}$, hence applying Part 4 of Lemma 4.3.4 shows that B and x are curvature incident once with F_1, and that x is a red triangle or a vertex of green degree 3. So $\chi(B, F_1, \Gamma_A) + \chi(x, F_1, \Gamma_A) = -1/2$, and hence by Lemma 5.3.6 $|\mathcal{B}_{F_1}| = 2$, and each element of $\mathcal{B}_{F_1}$ is either a red triangle containing one boundary edge, or a boundary vertex of green degree 3 that is not on $\rho \cap \rho'$. In particular, $\mathcal{S}_{F_1} = \mathcal{B}_{F_1} \cup \{B, x\}$. By Lemma 3.2.6 $|\partial(B)| = 4$, hence $|\partial(B) \setminus \partial(\Gamma_A)| = 3$. So as B is curvature incident once with F_1 and ef or fe is a sub-path of $\partial(B)$, there is an internal green face $F_2 \notin \{F, F_1\}$ edge-incident with B, and $B \in \mathcal{B}_{F_2}$. In particular, $\partial(F_2) \cap \rho \neq \emptyset$, so by Theorem 4.5.13 $\partial(F_2) \cap \rho$ is a single consolidated edge l. Now as each element of $\mathcal{B}_{F_1}$ is either a red triangle containing one boundary edge, or a vertex of green degree 3 that is not on $\rho \cap \rho'$, there is a $y \in \mathcal{B}_{F_1} \cap \mathcal{B}_{F_2}$. Suppose that F_1 has an edge on ρ. Then from $\partial(F_2) \cap \rho = l$ we have $\mathcal{S}_{F_2} = \{B, y\}$. Hence $\kappa_{\Gamma_A}(F_2) = 1 - 1/3 - 1/4 = 5/12 > 0$, contradicting Axiom T_6.

So F_1 has an edge on ρ'. Since $y \in \mathcal{B}_{F_1} \cap \mathcal{B}_{F_2}$, we have $\partial(F_2) \cap \rho' \neq \emptyset$, hence by Theorem 4.5.13 $\partial(F_2) \cap \rho$ is a single consolidated edge l'. From Part 1 of Lemma 4.3.4 we therefore deduce that at least one of the following cases holds.

(i) B_{F_2} contains a vertex on $\rho \cap \rho'$.

(ii) $|\mathcal{B}_{F_2} \setminus \{B, y\})| \geq 2$.

(iii) B_{F_2} contains a highly hyperbolic red blob.

But since $\mathcal{S}_{F_2} \supseteq \{B, y\}$, by Lemmas 2.6.10 and 3.2.8 in all Cases (i)-(iii) we have

$$\kappa_{\Gamma_A}(F_2) \leq 1 - 1/3 - 1/4 - 1/2 = -1/12 < 0,$$

a contradiction. ∎

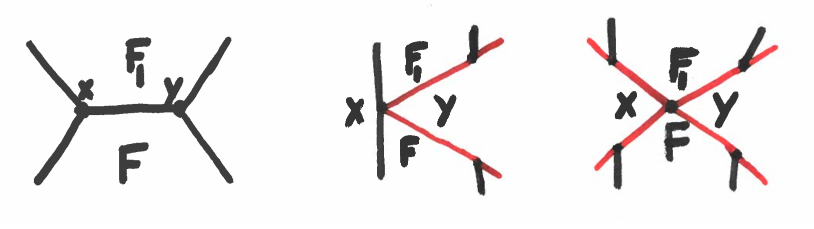

Figure 5.10: Depiction of ceiling neighbours, see Definition 5.4.1.

We are now ready prove the main result of this section.

Theorem 5.3.8. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains a pre-neighbourly face F. Then $|\mathcal{B}_F| = 2$, no $x \in \mathcal{B}_F$ is a red blob with $\mathbf{Area}(x) \geq 2$, and no red blob of $\mathcal{S}_F \setminus \mathcal{B}_F$ contains a boundary edge. Furthermore, the multiset of curvature values of $\mathcal{S}_F \setminus \mathcal{B}_F$ is $\{-1/4, -1/4\}$, or $\{-1/6, -1/6, -1/6\}$, or $\{-1/3, -1/6\}$.*

Proof. By Lemma 5.3.4 we have $|\mathcal{B}_F| = 2$ and no red blob of $\mathcal{S}_F \setminus \mathcal{B}_F$ contains a boundary edge. By Lemma 5.3.7 each red blob of $\mathcal{B}_F$ is a red triangle.

So it remains to prove the final statement. By definition of a pre-neighbourly face, each vertex of $\mathcal{B}_F$ has green degree 3. Hence the elements of $\mathcal{B}_F$ collectively give F curvature precisely $-1/2$. By Lemma 5.3.1 each $x \in \mathcal{S}_F$ has $\chi(x, \Gamma_A) > -1/2$, so the lemma follows from Lemmas 2.6.10 and 3.2.8. ■

5.4 Two face thickness

The main result of this section is Proposition 5.4.11, which plays a central part in the proof of Theorem 5.1.1. Recall that a consolidated edge between faces F and F' is a non-empty path of maximal length that is a sub-path of both $\partial(F)$ and $\partial(F')$, and that a path may consist of a single vertex.

In the next definition we allow tracing boundaries of faces in both directions.

Definition 5.4.1. Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. A *corner of F* is a green face $F' \neq F$ such that $\mathcal{B}_{F'} \cap \mathcal{B}_F \neq \emptyset$.

Suppose that $|\mathcal{S}_F \setminus \mathcal{B}_F| \geq 2$. A *ceiling of F* is a green face F_1 such that $\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$ contains distinct elements x and y, and F_1 and F are incident by a consolidated edge e that has the following properties (see Figure 5.10).

1. e contains or is incident with each of x and y.
2. If x and y are vertices, then $|e| \geq 1$.
3. If some $z \in \{x, y\}$ is a red blob, then there are edges $f \subseteq \partial(F)$ and $g \subseteq \partial(F_1)$ such that ef is a sub-path of $\partial(F)$, eg is a sub-path of $\partial(F_1)$, and fg is a sub-path of $\partial(z)$.

We shall call x and y *ceiling neighbours of* F *and* F_1.

Lemma 5.4.2. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Suppose that $\partial(F) \cap \omega$ and $\partial(F) \cap \tau$ are single consolidated edges. Then the elements of $\mathcal{B}_F$ collectively give F curvature of at most $-3/4$.*

Proof. By Axiom T_6, F contains a boundary edge, so without loss of generality assume that $|\partial(F) \cap \omega| \geq 1$. By Part 1 of Lemma 4.3.4 there are $x, y \in \mathcal{B}_F$ that lie on or are incident with $\partial(F) \cap \omega$. Suppose first that $x = y$. If F is incident with some vertex v on $\omega \cap \tau$ distinct from x, then by Lemmas 2.6.10 and 3.2.8 we have $\chi(x, F, \Gamma_A) + \chi(v, F, \Gamma_A) \leq -1/4 - 1/2 = -3/4$. So we may assume that no such v exists. Then since $\partial(F) \cap \tau \neq \emptyset$, x is curvature incident more than once with F. Hence by Part 2 of Lemma 4.3.4 $\chi(x, F, \Gamma_A) \leq -1/2$. If x is a red blob with at least two boundary edges, or a vertex on $\omega \cap \tau$, then by Lemmas 2.6.10 and 3.2.8 we have $\chi(x, F, \Gamma_A) \leq -1$. Hence we may assume that x is either a red blob with exactly one boundary edge, or a vertex on $\omega \setminus \tau$. If x is a red blob with an edge on τ, then by Part 1 of Lemma 4.3.4 some endpoint of $\partial(F) \cap \omega$ lies on $\omega \cap \tau$, a contradiction. Hence we can assume that x is not a red blob with an edge on τ. Hence there is $z \in \mathcal{B}_F \setminus \{x\}$ such that z is either a red blob that meets F with an edge on τ, or a vertex on τ. By Part 2 of Lemma 4.3.4 $\chi(z, F, \Gamma_A) \leq -1/4$, so the lemma follows.

Now assume that $x \neq y$. If x or y is a red blob with at least two boundary edges, or a vertex on $\omega \cap \tau$, then $\chi(x, F, \Gamma_A) + \chi(y, F, \Gamma_A) \leq -1/2 - 1/4 = -3/4$; and if x or y is a red blob with an edge on τ, then some endpoint of $\partial(F) \cap \omega$ lies on $\omega \cap \tau$. So we may assume that x and y are either red blobs with precisely one boundary edge that lies on $\omega \setminus \tau$, or vertices on $\omega \setminus \tau$. Then there is $z \in \mathcal{B}_F \setminus \{x, y\}$ with the same properties as in the previous paragraph, so we are done. ∎

Definition 5.4.3. Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. We denote by $\mathbf{T}_F$ a cyclic tuple of elements of $\mathcal{S}_F$.

Lemma 5.4.4. *Let F be a pre-neighbourly face of $\Gamma_A \in \mathcal{T}$ with respect to ρ, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then the following statements hold.*

1. *F has a ceiling.*
2. *Let F_1 be a ceiling of F. Then the following statements hold.*
   - *(i) $\partial(F_1)$ intersects exactly one of the boundaries of Γ_A.*
   - *(ii) If $\partial(F_1) \cap \rho \neq \emptyset$, then all elements of $\mathcal{S}_F \setminus \mathcal{B}_F$ are curvature incident once with F_1.*

Proof. By Theorem 5.3.8 $|\mathcal{S}_F \setminus \mathcal{B}_F| \geq 2$, and no red blob of $\mathcal{S}_F \setminus \mathcal{B}_F$ contains a boundary edge. Let x and y be distinct elements of $\mathcal{S}_F \setminus \mathcal{B}_F$ that are adjacent in $\mathbf{T}_F$. Since x and y are either interior vertices or red blobs with no boundary edge, there is an internal green face F_1 such that F_1 and F are incident by a consolidated edge e that contains or is incident with each of x and y, and since x and y are adjacent in $\mathbf{T}_F$ and $x \neq y$, we can choose F_1 such that if x and y are vertices, then $|e| \geq 1$; and such that if some $z \in \{x, y\}$ is a red blob, then Condition 3 of Definition 5.4.1 holds. So Part 1 follows.

For Part 2 (i), let F_1 be a ceiling of F. By Axiom T_6, F_1 has an edge on $\partial(\Gamma_A)$, so suppose for a contradiction that $\partial(F_1)$ intersects both boundaries of Γ_A. Applying Theorem 4.5.13 shows that both $\partial(F_1) \cap \omega$ and $\partial(F_1) \cap \tau$ are single consolidated edges, hence by Lemma 5.4.2 the elements of $\mathcal{B}_{F_1}$ give F_1 curvature of at most $-3/4$. Let x and y be ceiling neighbours of F and F_1, so that $x, y \in \mathcal{S}_F \setminus \mathcal{B}_F$. As neither x nor y is a red blob with a boundary edge, we have $x, y \notin \mathcal{B}_{F_1}$. By Lemmas 2.6.10 and 3.2.8 $\chi(x, F_1, \Gamma_A) + \chi(y, F_1, \Gamma_A) \leq 2 \cdot (-1/6) = -1/3$, hence $\kappa_{\Gamma_A}(F_1) \leq 1 - 3/4 - 1/3 = -1/12$, contradicting Axiom T_6. Hence Part 2 (i) follows.

For Part 2 (ii), assume for a contradiction that $\partial(F_1) \cap \rho \neq \emptyset$, and that some $z \in \mathcal{S}_F \setminus \mathcal{B}_F$ is curvature incident more than once with F_1. Since F_1 is a ceiling, we have $(\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)) \setminus \{z\} \neq \emptyset$, so by Lemmas 2.6.10 and 3.2.8 the elements of $\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$ collectively give F_1 curvature of at most $3 \cdot (-1/6) = -1/2$, and as no red blob of $\mathcal{S}_F \setminus \mathcal{B}_F$ contains a boundary edge, we have $(\mathcal{S}_F \setminus \mathcal{B}_F) \cap \mathcal{B}_{F_1} = \emptyset$. Hence by Part 4 of Lemma 4.3.4 we have $\mathcal{S}_{F_1} = \{x, y\} \cup \mathcal{B}_{F_1}$, where x and y are ceiling neighbours of F and F_1, x and y are red triangles or vertices of green degree 3, and the element of $\{x, y\} \setminus \{z\}$ is curvature incident once with F_1. In particular, F_1 is curvature incident more than once with a vertex of green degree 3 or a red triangle, so F_1 satisfies Condition (ii) of Lemma 5.2.1, hence by Lemma 5.2.1 there is $t \in \mathcal{S}_{F_1} \setminus \{z\}$ curvature incident more than once with F_1, and so $t \in \mathcal{B}_{F_1}$.

Applying Part 2 of Lemma 4.3.4 shows that $\chi(t, F_1, \Gamma_A) \leq -1/2$. By the previous paragraph the elements of $\mathcal{S}_{F_1} \setminus \mathcal{B}_{F_1}$ collectively give F_1 curvature of at most $-1/2$, hence $\chi(t, F_1, \Gamma_A) = -1/2$, and so t is a red triangle or a vertex with green degree 3. It follows that $\mathcal{S}_{F_1} = \{x, y, t\}$. Hence as z and t are red triangles or vertices with green degree 3, F is incident with itself by a consolidated edge that contains or is incident with z and t. Hence by Part 3 of Lemma 4.2.3 $\Gamma_A \setminus (\overline{F_1})^\circ$ decomposes as an edge-disjoint union of two annular diagrams, so either t is a red triangle with an edge on ρ', or a vertex on ρ'. But then since $t \in \mathcal{B}_{F_1}$, we have $\partial(F_1) \cap \rho' \neq \emptyset$, contradicting Part 2 (i) of the lemma. The result follows. ■

Lemma 5.4.5. *Suppose that all red blobs of $\Gamma_A \in \mathcal{T}$ are simply-connected, that none of them are complicated, and that Γ_A contains a pre-neighbourly face F with respect to ρ, so that by Lemma 5.4.4 F has a ceiling F_1. If F_1 has an edge on ρ, then F has distinct corners F_2 and F_3, and $F_1 \notin \{F_2, F_3\}$.*

In particular, if F' is a neighbourly face with respect to ρ, then F' has two distinct corners, and a ceiling, and any ceiling of F' is not a corner of F'.

Proof. We first prove that F has (not necessarily) distinct corners F_2 and F_3 and that $F_1 \notin \{F_2, F_3\}$, and then apply $F_1 \notin \{F_2, F_3\}$ to show that F_2 and F_3 are in fact distinct. As F is pre-neighbourly, F is thin: so no red blob of $\mathcal{B}_F$ has an edge ρ', and no vertex of $\mathcal{B}_F$ lies on ρ'; and all elements of $\mathcal{S}_F$ are curvature incident once with F. Hence as by Theorem 5.3.8 we have $|\mathcal{B}_F| = 2$ and the elements of $\mathcal{B}_F$ are red triangles or vertices of green degree 3, there are

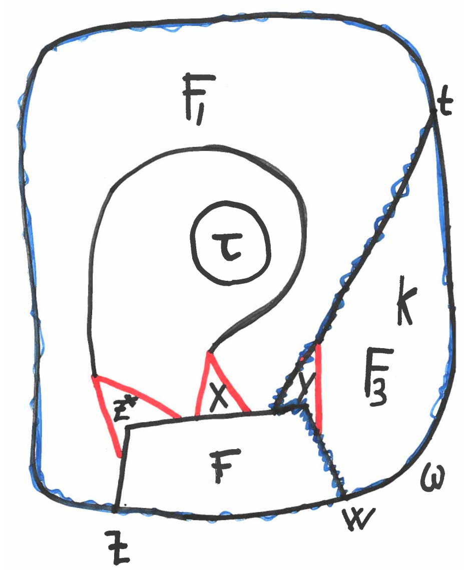

Figure 5.11: Depiction of F_3 contained in the sub-diagram K. The blue curves are the boundaries of Γ, see the proof of Lemma 5.4.5.

green faces $F_2 \neq F$ and $F_3 \neq F$ with $\mathcal{B}_{F_i} \cap \mathcal{B}_F \neq \emptyset$ for $2 \leq i \leq 3$; so $\partial(F_i) \cap \rho \neq \emptyset$; each element of $\mathcal{B}_F$ is curvature incident at most once with F_i; and for each $z \in \mathcal{B}_{F_i} \cap \mathcal{B}_F$, there is an element $z^* \in \mathcal{S}_{F_i} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$ adjacent to z in $\mathbf{T}_F$.

Let x and y be ceiling neighbours of F and F_1. Since F_1 has an edge on ρ and $\partial(F) \cap \rho' = \emptyset$, by Theorem 4.5.13 $\partial(F_1) \cap \rho$ is a single consolidated edge l with $|l| \geq 1$. Suppose that $F_1 = F_2$. If some $t \in \{x, y\}$ is adjacent in $\mathbf{T}_F$ to an element of $\mathcal{B}_{F_1} \cap \mathcal{B}_F$, then by Definition 5.4.1 F_1 is curvature incident more than once with t, contradicting Lemma 5.4.4 since $\partial(F_1) \cap \rho \neq \emptyset$. So if $F_2 = F_3$, then $|\mathcal{B}_{F_1} \cap \mathcal{B}_F| = 2$, and by the previous paragraph we have $|\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)| \geq 4$. But by Theorem 5.3.8 $|\mathcal{S}_F \setminus \mathcal{B}_F| \leq 3$, a contradiction.

Hence $F_1 \neq F_3$. Then $|\mathcal{B}_{F_1} \cap \mathcal{B}_F| = 1$, say $\mathcal{B}_{F_1} \cap \mathcal{B}_F = \{z\}$, and l contains or is incident with z. By the first paragraph there exists $z^* \in \mathcal{S}_{F_1}$ that is adjacent in $\mathbf{T}_F$ to z. As no element of $\mathcal{S}_F \setminus \mathcal{B}_F$ is curvature incident more than once with F_1, we have $z^* \notin \{x, y\}$. Hence $|\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)| \geq 3$, and therefore by Theorem 5.3.8 we have $\mathcal{S}_F \setminus \mathcal{B}_F = \{z^*, x, y\}$, so

$$\mathbf{T}_F = (z, z^*, x, y, w),$$

say, and all elements of $\mathcal{S}_F$ are red triangles or vertices of green degree 3. Since $|l| \geq 1$ and z is curvature incident once with F_1, there is $t \in \mathcal{B}_{F_1} \setminus \{z\}$ that lies on or is incident with l. As by Theorem 5.3.8 no red blob of $\{z^*, x, y\}$ contains a boundary edge, we have $\{z, t\} \cap \{z^*, x, y\} = \emptyset$. By Part 2 of Lemma 4.3.4 we have $\chi(t, F_1, \Gamma_A) \leq -1/4$, so as by Axiom T_6 we have $\kappa_{\Gamma_A}(F_1) = 0$, we deduce that $\mathcal{S}_{F_1} = \{z, z^*, x, y, t\}$, and $\chi(t, F_1, \Gamma_A) = -1/4$. In particular, all elements of $\mathcal{S}_{F_1}$ are curvature incident once with F_1, and are red triangles or vertices of green degree 3.

Since $\{z, z^*, x, y\} \subseteq \mathcal{S}_{F_1}$ and all elements of $\mathcal{S}_{F_1}$ are curvature incident once with F_1, it follows that $\mathbb{R}^2 \setminus (\overline{F} \cup \overline{F_1})$ contains a bounded component C. Suppose that $\rho = \tau$. As $F_1 \neq F_3$, we have $F_3 \subseteq C$. Now assume that $\rho = \omega$. Then as all elements of $\mathcal{S}_F \setminus \mathcal{B}_F$ are interior vertices or red blobs with no boundary edge, we deduce that $\mathbf{CArea}(\overline{C}) \geq (1, 0)$. Hence we showed that C contains an internal green face F'. By Axiom T_6, F' contains a boundary edge, hence C contains the external face with boundary τ. Therefore, $\Gamma_A \setminus (\overline{F} \cup \overline{F_1})^\circ$ contains an annular sub-diagram Γ with $\partial(\Gamma) = \{\rho, \rho_1\}$, where $\rho_1 \subseteq \partial(F) \cup \partial(F_1)$ and $|\rho_1| \geq 1$. In particular, $F_3 \subseteq \Gamma$. By Axiom T_1, ρ is a simple closed path. Since F and F_1 are not incident more than once with any vertex, $\partial(F)$ and $\partial(F_1)$ are simple closed paths, so ρ_1 is also a simple closed path. Hence as F has an edge on ρ, Γ is a union of islands and bridges (see Definition 2.5.10), and therefore F_3 is contained in a simply-connected sub-diagram $K \subseteq \Gamma_A \setminus (\overline{F} \cup \overline{F_1})^\circ$ bounded by a closed path $p \in \Gamma_A^1$ (see Definition 2.5.14) of the form $p_1 p_2$, where $p_1 \subseteq \partial(F) \cup \partial(F_1)$ and $|p_1| \geq 1$, and p_2 is a sub-path of ρ (see Figure 5.11 with $\rho = \omega$).

Since $F_3 \neq F_1$, we have $|\mathcal{B}_{F_3} \cap \mathcal{B}_F| = 1$, so $\mathcal{B}_{F_3} \cap \mathcal{B}_F = \{w\}$; and as $\partial(F_3) \cap \rho \neq \emptyset$ and $\partial(F) \cap \rho' = \emptyset$, by Theorem 4.5.13 $\partial(F_3) \cap \rho$ is a single consolidated edge l_1 that contains or is incident with w. Since $\mathbf{T}_F = (z, z^*, x, y, w)$, by the first paragraph $y \in \mathcal{S}_{F_3}$, and as y is a red triangle or a vertex of green degree 3, F_1 and F_3 are incident by a consolidated edge l_2 that contains or is incident with y. Since $\mathcal{S}_{F_1} = \{z, z^*, x, y, t\}$ and all elements of $\mathcal{S}_{F_1}$ are red triangles or vertices of green degree 3 curvature incident once with F_1, we deduce from $F_3 \subseteq K$ and $t \in l$ that l_2 contains or is incident with t. Therefore, as $\partial(F_3) \cap \rho = l_1$, we have $\mathcal{S}_{F_3} = \{w, y, t\}$. Hence all elements of $\mathcal{S}_{F_3}$ are red triangles or vertices of green degree 3, and therefore they are all curvature incident once with F_3. So

$$\kappa_{\Gamma_A}(F_3) = 1 - 1/4 - 1/4 - 1/6 = 1/3 > 0,$$

a contradiction. Hence we showed that $F_1 \notin \{F_2, F_3\}$.

Now if $F_2 = F_3$, then by Theorem 4.5.13 F_1 cannot have any edge on ρ, a contradiction.

Finally, suppose that F' is a neighbourly face with respect to ρ. By Lemma 5.3.3 F' is pre-neighbourly with respect to ρ, and as F' has no boundary links (see Definition 5.1.2), every ceiling of F' has an edge on ρ, so the result follows. ■

Corollary 5.4.6. *Let F be a pre-neighbourly face of $\Gamma_A \in \mathcal{T}$ with respect to ρ, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then for any ceiling F_1 of F that has an edge on ρ, we have $\mathcal{S}_{F_1} \cap \mathcal{S}_F = \{x, y\}$, where x and y are the ceiling neighbours of F and F_1.*

Proof. By Theorem 5.3.8 all elements of $\mathcal{B}_F$ are red triangles or vertices of green degree 3, hence as by Lemma 5.4.5 F_1 is not a corner of F, we have $\mathcal{S}_{F_1} \cap \mathcal{B}_F = \emptyset$. So suppose for a contradiction that $|\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)| \geq 3$. Then by Theorem 5.3.8 $|\mathcal{S}_F \setminus \mathcal{B}_F| = 3$, and each element of $\mathcal{S}_F$ is either a red triangle or a vertex of green degree 3. But then the middle element of $\mathcal{S}_F \setminus \mathcal{B}_F$ in $\mathbf{T}_F$ is curvature incident twice with F_1, contradicting Part 2 (ii) of Lemma 5.4.4. ■

Recall that two internal faces F and F are incident if and only if $\overline{F} \cap \overline{F'} \neq \emptyset$.

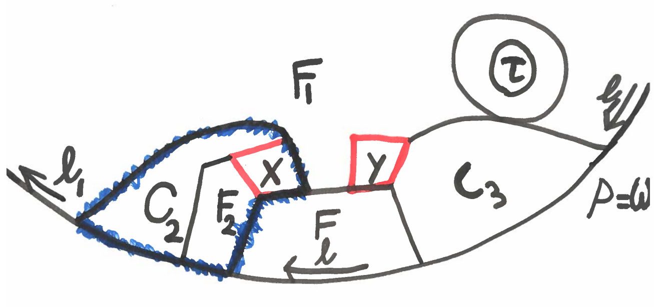

Figure 5.12: Depiction of the ceiling sub-diagram $\overline{C_2}$. The blue curve is $\partial(C_2)$, see the proof of Lemma 5.4.9.

Definition 5.4.7. Let F be a neighbourly face of $\Gamma_A \in \mathcal{T}$ with respect to ρ, and let $\mathbf{K}_F$ be the sub-diagram of Γ_A which is the union of all internal green faces incident with F. The *neighbourly sub-diagram* $\mathbf{N}_F$ *of* F is the union of $\mathbf{K}_F$ and all components C of $\Gamma_A \setminus \mathbf{K}_F$ such that $\overline{C} \cap \rho' = \emptyset$.

In the next definition we allow tracing boundaries of faces in both directions.

Definition 5.4.8. Suppose that all red blobs of $\Gamma_A \in \mathcal{T}$ are simply-connected, that none of them are complicated, and that Γ_A contains a pre-neighbourly face F with respect to ρ, and a ceiling F_1 of F that has an edge on ρ. Since $\partial(F) \cap \rho' = \emptyset$, by Theorem 4.5.13 $\partial(F) \cap \rho$ and $\partial(F_1) \cap \rho$ are single consolidated edges l and l_1 respectively. Then a *ceiling sub-diagram of* F *by* F_1 is the closure $\mathbf{C}_F$ of some component of $\Gamma_A \setminus (\overline{F} \cup \overline{F_1})$ that satisfies all of the following statements (see Figure 5.12 with $\mathbf{C}_F = \overline{C_2}$, and Figure 5.13 with $\mathbf{C}_F = \overline{C}$).

1. $\mathbf{C}_F$ is simply-connected and $\mathbf{C}_F$ contains precisely one corner of F.

2. The concatenation of l and $\partial(\mathbf{C}_F) \cap \partial(F)$ is a sub-path of $\partial(F)$, and the concatenation of $\partial(\mathbf{C}_F) \cap \partial(F_1)$ and l_1 is a sub-path of $\partial(F_1)$.

3. Either F and F_1 are edge-incident by a consolidated edge e such that the concatenation of $\partial(\mathbf{C}_F) \cap \partial(F)$ and e is a sub-path of $\partial(F)$ and the concatenation of e and $\partial(\mathbf{C}_F) \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$, or some ceiling neighbour z of F and F_1 is a red blob, and there are edges $e \subseteq \partial(F)$ and $f \subseteq \partial(F_1)$ such that the concatenation of $\partial(\mathbf{C}_F) \cap \partial(F)$ and e is a sub-path of $\partial(F)$, ef is a sub-path of $\partial(z)$, and the concatenation of f and $\partial(\mathbf{C}_F) \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$.

Lemma 5.4.9. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then the following statements hold.*

1. *Assume that F is pre-neighbourly with respect to ρ, and that F has a ceiling F_1 with an edge on ρ. Then F_1 defines a ceiling sub-diagram $\mathbf{C}_F$.*

Let F_2 be the corner of F with $F_2 \subseteq \mathbf{C}_F$. Then $\mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$ contains all internal green faces incident with F_2. In particular, F_2 is neighbourly with respect to ρ, and if the neighbourly sub-diagram of F_2 is simply-connected, then it is contained in $\mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$.

2. *Assume that F is neighbourly with respect to ρ. Then every ceiling of F defines a ceiling sub-diagram, and every ceiling sub-diagram of F is contained in the neighbourly sub-diagram $\mathbf{N}_F$ of F.*

   Let $\mathbf{C}_F$ be a ceiling sub-diagram of F given by F_1, and containing a corner F_2 of F. Then every ceiling sub-diagram of F_2 is contained in $\mathbf{C}_F$, and if $\mathbf{N}_F$ is simply-connected, then the neighbourly sub-diagram of F_2 is also simply-connected and is contained in $\mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$.

Proof. Part 1. We have $\partial(F) \cap \rho' = \emptyset$, and by Lemma 5.4.4 $\partial(F_1) \cap \rho' = \emptyset$. Hence by Theorem 4.5.13 $\partial(F) \cap \rho$ and $\partial(F_1) \cap \rho$ are single consolidated edges l and l_1 with $|l|, |l_1| \geq 1$. By Lemma 5.4.5 F has distinct corners F_2 and F_3, and $F_1 \notin \{F_2, F_3\}$. Let $\mathbf{K}_{F_2}$ be the sub-diagram of Γ_A which is the union of all internal green faces incident with F_2, and let x and y be ceiling neighbours of F and F_1. Since F_1 has an edge on ρ, by Lemma 5.4.4 x and y are curvature incident once with F_1, and by Corollary 5.4.6 we have $\mathcal{S}_{F_1} \cap \mathcal{S}_F = \{x, y\}$.

Suppose first that $\partial(F_1)$ passes more than once through some vertex. By Part 3 of Lemma 4.2.3 $\Gamma_A \setminus (F_1)^\circ$ decomposes as an edge-disjoint union of annular diagrams Γ_1 and Γ_2 with $F \subseteq \Gamma_1$ (say), the boundaries of Γ_1 are simple closed paths, and Γ_1 contains precisely one vertex incident more than once with F_1. Since F_1 has an edge on ρ, Γ_1 is a union of islands and bridges. Hence for each $i \in \{2, 3\}$: there is a component C_i of $\Gamma_A \setminus (\overline{F} \cup \overline{F_1})$ such that $\overline{C_i}$ is a simply-connected sub-diagram of Γ_A with $F_i \subseteq \overline{C_i}$, and $F_j \not\subseteq \overline{C_i}$ for $j \in \{2, 3\} \setminus \{i\}$. As Γ_1 contains precisely one vertex incident more than once with F_1, we deduce that one of the $\overline{C_i}$'s contains no vertex incident more than once with F_1, say it is C_2, so $\partial(C_2) \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$, and as by definition of a pre-neighbourly face no vertex is incident more than once with F, $\partial(C_2) \cap \partial(F)$ is a sub-path of $\partial(F)$.

We show that $\overline{C_2}$ is a ceiling sub-diagram of F by F_1. By the previous paragraph $\overline{C_2}$ is simply-connected and containing precisely one corner of F. As $\partial(F) \cap \rho = l$ and $\partial(F_1) \cap \rho = l_1$, we deduce that the concatenation of l and $\partial(C_2) \cap \partial(F)$ is a sub-path of $\partial(F)$, and the concatenation of $\partial(C_2) \cap \partial(F_1)$ and l_1 is a sub-path of $\partial(F_1)$. Moreover, since $\mathcal{S}_{F_1} \cap \mathcal{S}_F = \{x, y\}$, and x and y are curvature incident once with F_1, from Definition 5.4.1 of ceiling neighbours we deduce that either F and F_1 are edge-incident by a consolidated edge e such that the concatenation of $\partial(C_2) \cap \partial(F)$ and e is a sub-path of $\partial(F)$ and the concatenation of e and $\partial(C_2) \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$ (see Figure 5.12), or some $z \in \{x, y\}$ is a red blob, and there are edges $e \subseteq \partial(F)$ and $f \subseteq \partial(F_1)$ such that the concatenation of $\partial(C_2) \cap \partial(F)$ and e is a sub-path of $\partial(F)$, ef is a sub-path of $\partial(z)$, and the concatenation of f and $\partial(C_2) \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$. Hence $\overline{C_2}$ is the desired ceiling sub-diagram.

To prove the rest of Part 1, first note that if D is an internal face incident with F_2 such that $D \not\subseteq \overline{C_2}$, then either $D \in \{F, F_1\}$, or D is contained in some red blob of $\{x, y\}$. Hence all internal green faces incident with F_2 are contained in $\overline{C_2} \cup \overline{F} \cup \overline{F_1}$, and as $\overline{C_2}$ is simply-

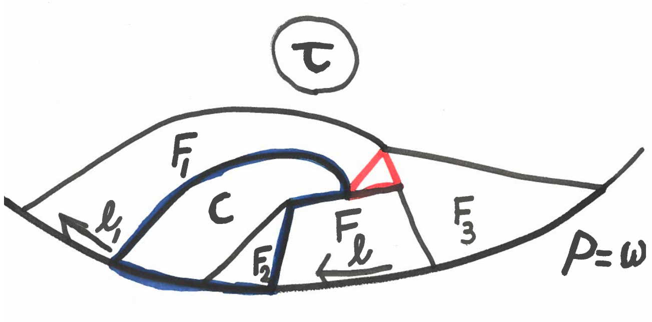

Figure 5.13: Depiction of the ceiling sub-diagram $\overline{C}$. The blue curve is $\partial(C)$, see the proof of Lemma 5.4.9.

connected and $\partial(F) \cap \rho' = \emptyset = \partial(F_1) \cap \rho'$, it follows that F_2 is neighbourly with respect to ρ. Suppose that the neighbourly sub-diagram $\mathbf{N}_{F_2}$ of F_2 is simply-connected. Then by Definition 5.4.7 of a neighbourly sub-diagram, $\mathbf{K}_{F_2} \cap \Gamma^1_A$ (see Definition 2.5.14) contains no closed path p such that $\overline{C} \cap \Gamma_A$ is an annular sub-diagram of Γ_A for each of the two components C of $\mathbb{R}^2 \setminus p$. In particular, F_2 is not incident with F_1. So all components C of $\Gamma_A \setminus \mathbf{K}_{F_2}$ with $\overline{C} \cap \rho' = \emptyset$ satisfy $C \subseteq \overline{C_2} \cup \overline{F} \cup \overline{F_1}$, and therefore $\mathbf{N}_{F_2} \subseteq \overline{C_2} \cup \overline{F} \cup \overline{F_1}$.

Now suppose that $\partial(F_1)$ does not pass more than once through any vertex. Then $\partial(F_1)$ is a simple closed path. Hence as F_1 has an edge on ρ, there is a component C of $\Gamma_A \setminus (\overline{F} \cup \overline{F_1})$ such that $\overline{C}$ is simply-connected, $\overline{C}$ contains precisely one corner of F, say $F_2 \subseteq \overline{C}$, and $\overline{C} \cap \partial(F_1)$ is a sub-path of $\partial(F_1)$. Since F is not incident more than once with any vertex, $\overline{C} \cap \partial(F)$ is a sub-path of $\partial(F)$. So as $\mathcal{S}_{F_1} \cap \mathcal{S}_F = \{x, y\}$; x and y are curvature incident once with F_1; and $\partial(F) \cap \rho = l$, $\partial(F_1) \cap \rho = l_1$, we deduce that C is a ceiling sub-diagram of F by F_1 (see Figure 5.13).

So similarly as in the previous case we can show that F_2 is neighbourly with respect to ρ, and $\overline{C} \cup \overline{F} \cup \overline{F_1}$ contains all internal green faces incident with F_2. As $\partial(F)$ and $\partial(F_1)$ are simple closed paths, $\overline{F}$ and $\overline{F_1}$ are simply-connected, hence all components C' of $\Gamma_A \setminus \mathbf{K}_{F_2}$ with $\overline{C'} \cap \rho' = \emptyset$ satisfy $C' \subseteq \overline{C} \cup \overline{F} \cup \overline{F_1}$, and so $\mathbf{N}_{F_2} \subseteq \overline{C} \cup \overline{F} \cup \overline{F_1}$. Hence Part 1 follows.

Part 2. By Lemma 5.3.3 F is pre-neighbourly with respect to ρ, and since F has no boundary links, by Axiom T_6 every ceiling F_1 of F has an edge on ρ. So by Part 1, F_1 defines a ceiling sub-diagram $\mathbf{C}_F$. Let $\mathbf{K}_F$ be the sub-diagram of Γ_A that is the union of all internal green faces incident with F. Since $\partial(F) \cap \rho' = \emptyset$ and by Lemma 5.4.4 $\partial(F_1) \cap \rho' = \emptyset$, from Definition 5.4.8 of a ceiling sub-diagram it follows that $\mathbf{C}_F \cap \rho' = \emptyset$. Hence as $\mathbf{K}_F \subseteq \mathbf{N}_F$ and $\mathbf{N}_F$ contains all components C of $\Gamma_A \setminus \mathbf{K}_F$ with $\overline{C} \cap \rho' = \emptyset$, we have $\mathbf{C}_F \subseteq \mathbf{N}_F$. Let F_2 be the corner of F with $F_2 \subseteq \mathbf{C}_F$. By Part 1, F_2 is neighbourly with respect to ρ, hence similarly as for $\mathbf{C}_F$ we deduce that all ceiling sub-diagrams $\mathbf{C}_{F_2}$ of F_2 satisfy $\mathbf{C}_{F_2} \cap \rho' = \emptyset$. Hence as $F_2 \subseteq \mathbf{C}_F$, we have $\mathbf{C}_{F_2} \subseteq \mathbf{C}_F$. Finally, suppose that $\mathbf{N}_F$ is simply-connected. By Part 1 all internal green faces incident with F_2 are contained in $\mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$, hence the neighbourly sub-diagram of F_2 is also simply-connected, and is contained in $\mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$. The lemma follows. ■

The following definition will be used repeatedly in the next proof.

Definition 5.4.10. Let F be a neighbourly face of $\Gamma_A \in \mathcal{T}$, and suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Let $\mathbf{N}_F$ be the neighbourly sub-diagram of F. A *neighbourly region of* F is a sub-diagram $\mathbf{R}_F \subseteq \mathbf{N}_F$ such that if $\mathbf{N}_F$ is simply-connected, then $\mathbf{R}_F = \mathbf{N}_F$, else $\mathbf{CArea}(\mathbf{R}_F)$ is minimal subject to the following two conditions.

1. $\mathbf{R}_F$ contains all internal green faces incident with F.

2. $\mathbf{R}_F$ contains a ceiling sub-diagram of F given by some ceiling of F.

Proposition 5.4.11. *Suppose that all red blobs of* $\Gamma_A \in \mathcal{T}$ *are simply-connected, and that none of them are complicated. Then* Γ_A *has no neighbourly faces.*

Proof. Assume for a contradiction that Γ_A has a neighbourly face F with respect to ρ, and let $\mathbf{N}_F$ and $\mathbf{R}_F$ be the neighbourly sub-diagram and a neighbourly region of F. Furthermore, choose F so that $\mathbf{CArea}(\mathbf{R}_F)$ is minimal among all neighbourly regions corresponding to neighbourly faces with respect to ρ. By Lemmas 5.4.5 and 5.4.9 F has distinct corners F_2 and F_3, a ceiling F_1 with $F_1 \notin \{F_2, F_3\}$, and a ceiling sub-diagram $\mathbf{C}_F$ defined by F_1. By Lemma 5.3.3 F is pre-neighbourly, hence by Theorem 5.3.8 the elements of $\mathcal{B}_F$ are red triangles or vertices of green degree 3, so F_2 and F_3 are incident with F. Without loss of generality let $F_2 \subseteq \mathbf{C}_F$. Then $F_3 \nsubseteq \mathbf{C}_F$, and by Lemma 5.4.9 F_2 is neighbourly with respect to ρ, so let $\mathbf{N}_{F_2}$ and $\mathbf{R}_{F_2}$ be the neighbourly sub-diagram and a neighbourly region of F_2.

Suppose that $\mathbf{N}_F$ is simply-connected. Then by Lemma 5.4.9 $\mathbf{C}_F \subseteq \mathbf{N}_F$, $\mathbf{N}_{F_2}$ is simply-connected and $\mathbf{R}_{F_2} = \mathbf{N}_{F_2} \subseteq \mathbf{C}_F \cup \overline{F} \cup \overline{F_1}$. But since $\mathbf{N}_F$ contains all internal green faces incident with F, we have $\mathbf{R}_F = \mathbf{N}_F \supseteq \mathbf{C}_F \cup \overline{F} \cup \overline{F_1} \cup \overline{F_3}$, so as $F_3 \nsubseteq \mathbf{C}_F$, we have $\mathbf{CArea}(\mathbf{R}_{F_2}) < \mathbf{CArea}(\mathbf{R}_F)$, a contradiction.

Hence $\mathbf{N}_F$ is not simply-connected. By Definition 5.4.10 $\mathbf{R}_F$ contains a ceiling sub-diagram $\mathbf{C}'_F$ given by some ceiling F'_1 of F. Without loss of generality assume that $F_2 \subseteq \mathbf{C}'_F$, so $F_3 \nsubseteq \mathbf{C}'_F$. By Lemmas 5.4.5 and 5.4.9 we have $F'_1 \notin \{F_2, F_3\}$, F_2 is neighbourly with respect to ρ, and $\mathbf{C}'_F \cup \overline{F} \cup \overline{F'_1}$ contains all internal green faces incident with F_2. By Definition 5.4.10 $\mathbf{R}_F$ contains all internal green faces incident with F, so as $\mathbf{R}_F \supseteq \mathbf{C}'_F$, we have $\mathbf{R}_F \supseteq \mathbf{C}'_F \cup \overline{F} \cup \overline{F'_1} \cup \overline{F_3}$.

Suppose that $\mathbf{N}_{F_2}$ is simply-connected. Then by Lemma 5.4.9 we have $\mathbf{R}_{F_2} = \mathbf{N}_{F_2} \subseteq \mathbf{C}'_F \cup \overline{F} \cup \overline{F'_1}$, hence $\mathbf{CArea}(\mathbf{R}_{F_2}) < \mathbf{CArea}(\mathbf{R}_F)$, a contradiction. So $\mathbf{N}_{F_2}$ is not simply-connected. By Lemma 5.4.9 every ceiling sub-diagram $\mathbf{C}_{F_2}$ of F_2 satisfies $\mathbf{C}_{F_2} \subseteq \mathbf{C}'_F$. Hence by Definition 5.4.10 $\mathbf{R}_{F_2} \subseteq \mathbf{C}'_F \cup \overline{F} \cup \overline{F'_1}$, so $\mathbf{CArea}(\mathbf{R}_{F_2}) < \mathbf{CArea}(\mathbf{R}_F)$. ∎

Corollary 5.4.12. *Let* $\Gamma_A \in \mathcal{T}$. *Suppose that all red blobs of* Γ_A *are simply-connected, that none of them are complicated, and that* Γ_A *contains a pre-neighbourly face* F *with respect to* ρ. *If* F_1 *is a ceiling of* F, *then* F_1 *has an edge on* ρ'.

Proof. By Axiom T_6, F_1 has a boundary edge, so assume for a contradiction that it lies on ρ. Then by Lemma 5.4.9 some corner of F is neighbourly, contradicting Proposition 5.4.11. ∎

By using Proposition 5.4.11 we can strengthen Theorem 5.3.8.

Theorem 5.4.13. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains a pre-neighbourly face F with respect to ρ. Then both of the following two statements hold.*

1. *If $|\mathcal{S}_F \setminus \mathcal{B}_F| = 2$, then each element of $\mathcal{S}_F \setminus \mathcal{B}_F$ is either a red blob with area 2 and boundary length 4, or a vertex with green degree 4.*

2. *There is an internal green face F_1 with an edge on ρ' such that either F_1 and F are edge-incident, or $\mathcal{S}_{F_1} \cap \mathcal{S}_F$ contains a red blob B with no boundary edge and $|\partial(B)| \leq 4$.*

Proof. First note that by Lemma 5.4.5 there is a ceiling F_1 of F, and by Corollary 5.4.12 F_1 has an edge on ρ', so by Lemma 5.4.4 $\partial(F_1) \cap \rho = \emptyset$.

Part 1. It suffices to show that the multiset of curvature values of $\mathcal{S}_F \setminus \mathcal{B}_F$ is not $\{-1/3, -1/6\}$. So assume for a contradiction that $\mathcal{S}_F \setminus \mathcal{B}_F = \{x, y\}$, where x is either a red blob with area 4 and boundary length 6, or a vertex with green degree 6, and y is either a red triangle, or a vertex with green degree 3. Note that x and y are ceiling neighbours of F and F_1. For $1 \leq i \leq 6$, reading clockwise around x, let F_i be an internal green face with $x \in \mathcal{S}_{F_i}$, and without loss of generality assume that $F_6 = F$. By Axiom T_6 we can let l_i be a boundary consolidated edge of F_i with $|l_i| \geq 1$, where $l_1 \subseteq \rho'$ and $l_6 \subseteq \rho$.

By Corollary 4.5.15 there are $x_i, y_i \in \mathcal{B}_{F_i}$ that lie on or are incident with l_i that collectively give F_i curvature of at most $-1/2$, and since by Theorem 5.3.8 neither x nor y is a red blob with a boundary edge, we have $x, y \notin \mathcal{B}_{F_i}$. We will work through increasing i, showing that F_i is pre-neighbourly with respect to ρ' for $1 \leq i \leq 2$, that F_i is pre-neighbourly with respect to ρ for $3 \leq i \leq 4$, and that for $1 \leq i \leq 4$ the multiset of curvature values of $\mathcal{S}_{F_i} \setminus \mathcal{B}_{F_i}$ and $\mathcal{B}_{F_i}$ is $\{-1/3, -1/6\}$ and $\{-1/4, -1/4\}$ respectively. Using this we will then deduce that F_5 is a ceiling of F_4 with $\partial(F_5) \cap \rho \neq \emptyset \neq \partial(F_5) \cap \rho'$, contradicting Lemma 5.4.4.

Since $\chi(x, \Gamma_A) = -1/3$ and $\chi(y, \Gamma_A) = -1/6$ (see Definition 2.6.8), by Part 4 of Lemma 4.3.4 each $z \in \{x, y\}$ is curvature incident once with F_1, so by Lemma 5.3.6 $|\mathcal{B}_{F_1}| = 2$, and each element of $\mathcal{B}_{F_1}$ is curvature incident once with F_1, and is either a red triangle containing one boundary edge, or a boundary vertex with green degree 3 that is not on $\rho \cap \rho'$. In particular, no red blob of $\mathcal{S}_{F_1}$ has an edge on ρ. So as $\partial(F_1) \cap \rho = \emptyset$ and all elements of $\mathcal{S}_{F_1}$ are curvature incident once with F_1, we have that F_1 is pre-neighbourly with respect to ρ', and F_1 has the properties stated in the previous paragraph.

Suppose that for some $2 \leq i \leq 5$, $\partial(F_i)$ passes more than once through some vertex v. Then by Lemma 4.2.3 $\partial(F_i)$ passes through v twice, and $\Gamma_A \setminus (\overline{F_i})^\circ$ decomposes as an edge-disjoint union of two annular diagrams. But this is impossible since $l_1 \subseteq \rho'$, $l_6 \subseteq \rho$, and x and y are ceiling neighbours of F and F_1. Now suppose that for some $2 \leq i \leq 5$, F_i is edge-incident more than once with some blob B. By Proposition 4.2.8 Γ_A is island-free, B is edge-incident twice with F_i, and $\Gamma_A \setminus (\overline{F_i} \cup B)^\circ$ decomposes as an edge-disjoint union of two annular diagrams. This is only possible if $B \in \{x, y\}$, F and F_1 are incident at some vertex $v \in \{x, y\}$ and are not edge-incident, and $\Gamma_A \setminus B^\circ$ decomposes as an edge-disjoint union of two annular diagrams. But then the two common edges of F_i and B are consecutive on $\partial(B)$. So as F_i is not incident with any vertex more than once, we contradict Lemma 4.2.6. Hence we showed that for all $1 \leq i \leq 6$, all elements of $\mathcal{S}_{F_i}$ are curvature incident once with F_i.

By the choice of F_1 and F_2 there is $z \in \mathcal{B}_{F_1} \cap \mathcal{S}_{F_2}$. Since z is a red triangle or a vertex with $\delta_G(z) = 3$, z has 3 curvature incidences, so $z \in \mathcal{B}_{F_2}$, and F_2 has a consolidated edge $l_2' \subseteq \rho'$ that contains or is incident with z. By Theorem 4.5.13 $\partial(F_2) \cap \rho' = l_2'$. Since all elements of $\mathcal{S}_{F_2}$ are curvature incident once with F_2 and z is a red triangle or a vertex with $\delta_G(z) = 3$, we deduce that there is $t \in \mathcal{B}_{F_2} \setminus \{z\}$ that lies on or is incident with l_2'. By Part 2 of Lemma 4.3.4 we further have $\chi(t, F_2, \Gamma_A) \leq -1/4$. Now $\kappa_{\Gamma_A}(F_2) = 0$, hence as $\chi(z, F_2, \Gamma_A) = -1/4$ and $\chi(x, F_2, \Gamma_A) = -1/3$, by Lemmas 2.6.10 and 3.2.8 there are two possible cases.

(a) We have $|\mathcal{S}_{F_2}| = 3$ and $\chi(t, F_2, \Gamma_A) = -5/12$.

(b) We have $|\mathcal{S}_{F_2}| = 4$, $\chi(t, F_2, \Gamma_A) = -1/4$, and there is $w \in \mathcal{S}_{F_2}$ with $\chi(w, F_2, \Gamma_A) = -1/6$.

Assume that Case (a) holds. By Lemmas 2.6.10 and 3.2.8 t is either a vertex with $\delta_G(t) = 7$ and $t \notin \rho \cap \rho'$, or a red blob with $|\partial(t) \cap \partial(\Gamma_A)| = 1$ and $\mathbf{Area}(t) = 5$: and by Lemma 3.2.6 $|\partial(t)| = 7$. Since t is curvature incident once with F_2, by the choice of F_3 we have $t \in \mathcal{S}_{F_3}$, so t is curvature incident once with F_3, and therefore if t is a red blob, then $t \notin \mathcal{B}_{F_3}$, and if t is a vertex, then $t \notin \{x_3, y_3\}$. By Part 4 of Lemma 4.3.4 the elements of $\mathcal{B}_{F_3}$ collectively give F_3 curvature of at most $-1/2$, so as $x \notin \mathcal{B}_{F_3}$, and either $t \notin \mathcal{B}_{F_3}$ or $t \notin \{x_3, y_3\}$, we have

$$\begin{aligned}\kappa_{\Gamma_A}(F_3) &\leq 1 + \chi(x, F_3, \Gamma_A) + \chi(t, F_3, \Gamma_A) - 1/2 \\ &= 1 - 1/3 - 5/12 - 1/2 = -1/4 < 0,\end{aligned}$$

a contradiction.

So Case (b) holds. Hence as $\chi(t, F_2, \Gamma_A) = -1/4$, by Lemmas 2.6.10 and 3.2.8 t is a red triangle containing one boundary edge or a boundary vertex with green degree 3 that is not on $\rho \cap \rho'$. In particular, if t is a red triangle, then it has an edge on ρ'. From $\chi(w, F_2, \Gamma_A) = -1/6$ we deduce that w is a red triangle with no boundary edge, or a vertex with $\delta_G(w) = 3$ and $w \notin \partial(\Gamma_A)$. Hence no red blob of $\mathcal{S}_{F_2}$ has an edge on ρ, and no vertex of $\mathcal{S}_{F_2}$ lies on ρ. Suppose that $\partial(F_2) \cap \rho \neq \emptyset$. Then as no vertex of $\mathcal{S}_{F_2}$ lies on ρ, there is a vertex of F_2 of green degree 2 on ρ, so $\mathcal{S}_{F_2}$ contains a red blob with an edge on ρ, a contradiction. Hence F_2 is thin with respect to ρ', and since all elements of $\mathcal{S}_{F_2}$ are curvature incident once with F_2 and each vertex of $\mathcal{B}_{F_2}$ has green degree 3, F_2 is pre-neighbourly with respect to ρ'.

By the choice of F_3 we have $\mathcal{S}_{F_3} \subseteq \{x, w\}$, and note that F_3, x, w satisfy Conditions 1-3 of Definition 5.4.1. Hence F_3 is a ceiling of F_2, so by Corollary 5.4.12 we have $l_3 \subseteq \rho$, and similarly as for F_1 we deduce that the multiset of curvature values of $\mathcal{S}_{F_3} \setminus \mathcal{B}_{F_3}$ and $\mathcal{B}_{F_3}$ is $\{-1/3, -1/6\}$ and $\{-1/4, -1/4\}$ respectively, and that F_3 is pre-neighbourly with respect to ρ. Then similarly as for F_2 we deduce that the multiset of curvature values of $\mathcal{S}_{F_4} \setminus \mathcal{B}_{F_4}$ and $\mathcal{B}_{F_4}$ is $\{-1/3, -1/6\}$ and $\{-1/4, -1/4\}$ respectively, and that F_4 is pre-neighbourly with respect to ρ. So similarly as for F_2 and F_3 we deduce that F_5 is a ceiling of F_4, hence by Corollary 5.4.12 $l_5 \subseteq \rho'$, and by Lemma 5.4.4 $\partial(F_5) \cap \rho = \emptyset$. But by the choice of F_5 we have $\mathcal{S}_{F_5} \cap \mathcal{B}_F \neq \emptyset$, and as the elements of $\mathcal{B}_F$ are red triangles or vertices with green degree 3, we have $\partial(F_5) \cap \rho \neq \emptyset$, a contradiction. Hence Part 1 follows.

Part 2. By Definition 5.4.1 if F_1 and F are not edge-incident, then $\mathcal{S}_{F_1} \cap (\mathcal{S}_F \setminus \mathcal{B}_F)$ contains a red blob B. By Theorem 5.3.8 and Part 1, B contains no boundary edge and $|\partial(B)| \leq 4$. ■

5.5 Proof of Theorem 5.1.1

In this section we prove Theorems 5.1.1.

Proposition 5.5.1. *Let F be a thin face of $\Gamma_A \in \mathcal{T}$ with respect to ρ. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that no element of $\mathcal{S}_F$ is curvature incident more than once with F. Then there is an internal green face E with $\partial(E) \cap \partial(F) \neq \emptyset$, with an edge on ρ', and either E and F are edge-incident, or $\mathcal{S}_E \cap \mathcal{S}_F$ contains a red blob B with no boundary edge and $|\partial(B)| \leq 4$.*

Proof. Assume first that each vertex of $\mathcal{B}_F$ has green degree 3. Then F is pre-neighbourly (see Definition 5.1.3), hence by Part 2 of Theorem 5.4.13 the proposition follows.

Now suppose that there is a vertex $v \in \mathcal{B}_F$ with $\delta_G(v) \geq 4$. By Lemma 2.6.10 we have $\chi(v, F, \Gamma_A) \leq -1/3$, and since F is thin, we have $v \notin \rho \cap \rho'$. Hence by Corollary 4.5.17 there is an internal green face E incident with v, with an edge e on ρ', and v is not incident with any boundary edge of E. Moreover, note that we can choose E so that either E and F are edge-incident, or $\mathcal{S}_E \cap \mathcal{S}_F$ contains a red blob B. If E and F are edge-incident, then we are done, so assume they are not.

Suppose that E contains a vertex $v' \in \rho \cap \rho'$. Then as v is not incident with any boundary edge of E, we have that $\partial(E) \cap \rho$ is not a single consolidated edge, contradicting Theorem 4.5.13. By Corollary 4.5.15 there are $x, y \in \mathcal{B}_E$ that lie on or are incident with the consolidated edge of $\partial(E) \cap \rho'$ that collectively give E curvature of at most $-1/2$. In particular, $v \notin \{x, y\}$. Since E is not incident with any vertex of $\rho \cap \rho'$ and F is thin, by the second statement of Corollary 4.5.15 it follows that $B \notin \{x, y\}$. Hence as $\kappa_{\Gamma_A}(E) = 0$, and v, x, y collectively give E curvature of at most $-1/3 - 1/2 = -5/6$, we have $\chi(B, E, \Gamma_A) \geq -1/6$. So by Lemma 3.2.8 B is a red triangle with no boundary edge. ■

The next theorem provides the main argument for the proof of Theorem 5.1.1.

Theorem 5.5.2. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, that none of them are complicated, and that Γ_A contains a thin face F with respect to ρ. Then there is an internal green face E with $\partial(E) \cap \partial(F) \neq \emptyset$, with an edge on ρ', and either E and F are edge-incident, or $\mathcal{S}_E \cap \mathcal{S}_F$ contains a red blob B with no boundary edge and $|\partial(B)| \leq 4$.*

Proof. By Proposition 5.5.1 we may assume that there is an element of $\mathcal{S}_F$ curvature incident more than once with F. Then by Lemma 5.2.4 there is $x \in \mathcal{B}_F$ curvature incident exactly twice with F. By Lemma 5.3.1 $\partial(F) \cap \partial(\Gamma_A) = \partial(F) \cap \rho$ is a single consolidated edge l with $|l| \geq 1$, Γ_A is island-free, and if $B \in \mathcal{S}_F$ is a red blob, then B contains at most one boundary edge.

We show that x is a red triangle or a vertex of green degree 3. Suppose not. By Lemmas 2.6.10; 3.2.8; and 5.3.1 we have

$$-1 < \chi(x, F, \Gamma_A) \leq -2 \cdot (1/3) = -2/3.$$

By Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$, hence there is $y \in \mathcal{S}_F \setminus \{x\}$. Now by Lemma 5.3.2 no element of $\mathcal{S}_F \setminus \{x\}$ is a red blob with an edge on ρ, or a vertex on ρ. Hence as F is thin, each such z is either a red blob with no boundary edge, or an interior vertex.

Suppose first that $\chi(x, F, \Gamma_A) < -2/3$. Since $\kappa_{\Gamma_A}(F) = 0$, by Lemmas 2.6.10 and 3.2.8 we have $\mathcal{S}_F = \{x, y\}$. Now there is an internal green face $F_1 \neq F$ with $\{x, y\} \in \mathcal{S}_{F_1}$. Since y is a red blob with no boundary edge, or an interior vertex, we have $y \in \mathcal{S}_{F_1} \setminus \mathcal{B}_{F_1}$. By Axiom T_6, F_1 has a boundary consolidated edge l_1 with $|l_1| \geq 1$, and by Corollary 4.5.15 there are $z, t \in \mathcal{B}_{F_1}$ that lie on or are incident with l_1 such that z and t collectively give F_1 curvature of at most $-1/2$. By Lemma 5.3.2 F is the only internal green face with an edge on ρ, hence $l_1 \subseteq \rho'$. So as Γ_A is island-free, by the second statement of Corollary 4.5.15 we have $x \notin \{z, t\}$ since x contains no edge on ρ'. Hence

$$\begin{aligned}\kappa_{\Gamma_A}(F_1) &\leq 1 + \chi(x, F_1, \Gamma_A) + \chi(y, F_1, \Gamma_A) - 1/2 \\ &< 1 - 1/3 - 1/6 - 1/2 = 0,\end{aligned}$$

a contradiction.

Now suppose that $\chi(x, F, \Gamma_A) = -2/3$. Then by Lemmas 2.6.10 and 3.2.8 x is either a vertex with $\delta_G(x) = 4$, or a red blob with $\mathbf{Area}(x) = 2$: and by Lemma 3.2.6 we have $|\partial(x) \setminus \partial(\Gamma_A)| = 3$. Since each element of $\mathcal{S}_F \setminus \{x\}$ is either a red blob with no boundary edge, or an interior vertex, by Lemmas 2.6.10 and 3.2.8 either there is $z \in \mathcal{S}_F \setminus \{x, y\}$ and $\chi(y, F, \Gamma_A) = -1/6 = \chi(z, F, \Gamma_A)$, or no such z exists and $\chi(y, F, \Gamma_A) = -1/3$. Since either $\delta_G(x) = 4$ or $|\partial(x) \setminus \partial(\Gamma_A)| = 3$, in both cases there is an internal green face F_1 that receives curvature $-1/3$ from elements of $\mathcal{S}_F \setminus \{x\}$ and curvature $-1/3$ from x. Hence similarly as in the previous case we deduce that

$$\kappa_{\Gamma_A}(F_1) \leq 1 - 1/3 - 1/3 - 1/2 = -1/6 < 0,$$

a contradiction. Hence we showed that x is a red triangle or a vertex of green degree 3.

Now F satisfies Condition (ii) of Lemma 5.2.1, hence F is curvature incident more than once with some element $y \in \mathcal{S}_F \setminus \{x\}$. So as $\kappa_{\Gamma_A}(F) = 0$ and $\chi(x, F, \Gamma_A) = -1/2$, by Lemmas 2.6.10 and 3.2.8 one of the following cases holds.

1. $\mathcal{S}_F = \{x, y\}$ has size 2 and $\chi(y, F, \Gamma_A) = -1/2$.

2. $\mathcal{S}_F = \{x, y, z\}$ has size 3 and $\chi(y, F, \Gamma_A) = -1/3$, $\chi(z, F, \Gamma_A) = -1/6$.

Assume that Case (a) holds. By the second paragraph y is either an interior vertex, or a red blob with no boundary edge, hence by Lemmas 2.6.10 and 3.2.8 y is either a vertex of green degree 4, or a red blob with area 2: so by Lemma 3.2.6 with boundary length 4. Hence there is an internal green face E with $y \in \mathcal{S}_E$; which by Lemma 5.3.2 has an edge on ρ'; and if E and F are not edge-incident, then y is a red blob, so we are done.

Suppose that Case (b) holds instead. By Lemmas 2.6.10 and 3.2.8 y and z are red triangles or vertices with green degree 3. So there is an internal green face E with $y, z \in \mathcal{S}_E$; with an edge on ρ'; and if E and F are not edge-incident, then at least one of y or z is a red blob. Hence the theorem follows. ∎

To prove Theorem 5.1.1 we need two additional lemmas.

Lemma 5.5.3. *Let $\Gamma_A \in \mathcal{T}$. Suppose that all red blobs of Γ_A are simply-connected, and that none of them are complicated. Then there is no green face F of Γ_A that satisfies all of the following statements.*

1. *F has an edge on at most one of the boundaries of Γ_A.*

2. *$\mathcal{S}_F = \{B, x, y\}$ has size 3, and each element of $\mathcal{S}_F$ is curvature incident once with F.*

3. *B is a red blob with an edge on the opposite boundary from F.*

4. *We have $\chi(B, F, \Gamma_A) = -5/12$, $\chi(x, F, \Gamma_A) = -1/3$ and $\chi(y, F, \Gamma_A) = -1/4$.*

Proof. Assume for a contradiction that there is such an F. Since Γ_A contains no complicated red blobs, all red blobs of Γ_A contain at most two boundary edges; and if a blob $B' \in \mathcal{S}_F$ attains this bound, then by Lemma 3.2.8 $\chi(B', \Gamma_A) = -1/2$ (see Definition 2.6.8) and Statement 4 fails to hold. Hence all blobs $B' \in \mathcal{S}_F$ contain at most one boundary edge.

By Axiom T_6, F has a consolidated edge $l \subseteq \rho \in \{\omega, \tau\}$ with $|l| \geq 1$, and by Corollary 4.5.15 there are $t, w \in \mathcal{B}_F$ that lie on or are incident with l. As by Lemma 3.1.9 each vertex $v \in \rho \cap \rho'$ satisfies $\chi(v, \Gamma_A) = -1/2$, there is no such v in $\partial(F)$. So as B has an edge on ρ', by the second statement of Corollary 4.5.15 we have $B \notin \{t, w\}$, hence $\{t, w\} = \{x, y\}$. By Lemmas 2.6.10 and 3.2.8 the following statements hold.

(i) We have $\mathbf{Area}(B) = 5$, so by Lemma 3.2.6 $|\partial(B)| = 7$,

(ii) x is either a vertex with $\delta_G(x) = 4$, or a red blob with $\mathbf{Area}(x) = 2$: so $|\partial(x)| = 4$,

(iii) y is either a vertex with $\delta_G(y) = 3$, or a red triangle.

Suppose that $B \in \mathcal{B}_F$. Then $\partial(F) \cap \rho' \neq \emptyset$. Hence by Theorem 4.5.13 $\partial(F) \cap \rho'$ is a single consolidated edge l' consisting of a single vertex v. Since $v \notin \rho \cap \rho'$, we have $v \notin \{x, y\}$, and therefore $\delta_G(v) = 2$. So as all elements of $\mathcal{S}_F$ are curvature incident once with F, it follows that there is a red blob $B' \in \mathcal{B}_F \setminus \{B\}$ incident with l'. Further, since $v \notin \rho \cap \rho'$, we have $B' \notin \{x, y\}$, a contradiction.

Hence there are internal green faces F_1 and F_2 that share edges e_1 and e_2 with B, where e_1 and e_2 are incident with the common edge of F and B, and $B, x \in \mathcal{S}_{F_1}$ and $B, y \in \mathcal{S}_{F_2}$. Since each element of $\mathcal{S}_F$ is curvature incident once with F, and F and B are edge-incident and have edges on the opposite boundaries of Γ_A, by Part 3 of Lemma 4.2.3 no internal green face of Γ_A is incident more than once with a vertex; and if some internal green face of Γ_A is edge-incident more than once with a blob B', then by Proposition 4.2.8 $B' = B$. Hence $\chi(x, F_1, \Gamma_A) = -1/3$, and if x is a blob, then $x \notin \mathcal{B}_{F_1}$. Since $\kappa_{\Gamma_A}(F_1) = 0$, we deduce that B is edge-incident once with F_1, so all elements of $\mathcal{S}_{F_1}$ are curvature incident once with F_1. Hence as by Lemmas 2.6.10 and 3.2.8 each $t \in \mathcal{S}_{F_1}$ has $\chi(t, \Gamma_A) \leq -1/6$, we have that there is $t \in \mathcal{S}_{F_1} \setminus \{B, x\}$ such that $\chi(t, F_1, \Gamma_A) = -1/4$ and $\mathcal{S}_{F_1} = \{B, x, t\}$.

Assume that $B \notin \mathcal{B}_{F_1}$. By Axiom T_6, F_1 has a consolidated edge $l_1 \subseteq \partial(\Gamma_A)$ with $|l_1| \geq 1$, and by Corollary 4.5.15 there are $x_1, y_1 \in \mathcal{B}_{F_1}$ that lie on or are incident with l_1 that collectively give F_1 curvature of at most $-1/2$. If $x_1 = y_1$, then $\chi(x_1, F_1, \Gamma_A) \leq -1/2$, a contradiction. So $x_1 \neq y_1$. Since $x_1, y_1 \in \mathcal{B}_{F_1}$, we have $B \notin \{x_1, y_1\}$, and $x \notin \{x_1, y_1\}$ if x is a red blob since we deduced $x \notin \mathcal{B}_{F_1}$ in that case. Hence x is a vertex. But then $\delta_G(x) = 4$ and $x \notin \rho \cap \rho'$, and by above x is incident once with F_1, so again $x \notin \{x_1, y_1\}$. Thus, $|\mathcal{S}_{F_1}| \geq 4$, a contradiction.

Hence $B \in \mathcal{B}_{F_1}$. By Statement (iii) above y has 3 curvature incidences, so $y \in \mathcal{B}_{F_2}$ and $\partial(F_2) \cap \rho \neq \emptyset$, and hence F_2 has a consolidated edge $l_2 \subseteq \rho$ that contains or is incident with y. By Theorem 4.5.13 we have $\partial(F_2) \cap \rho = l_2$. As by Axiom T_6 we have $\kappa_{\Gamma_A}(F_2) = 0$, and as each vertex $v \in \rho \cap \rho'$ has $\chi(v, \Gamma_A) = -1/2$, no such v lies on $\partial(F_2)$. So as y is curvature incident once with F_2, there is $w \in \mathcal{B}_{F_2} \setminus \{B, y\}$ that lies on or is incident with l_2. By Part 2 of Lemma 4.3.4 we have $\chi(w, F_2, \Gamma_A) \leq -1/4$. Since each $t' \in \mathcal{S}_{F_2}$ has $\chi(t', \Gamma_A) \leq -1/6$, we deduce that B is edge-incident once with F_2, $\mathcal{S}_{F_2} = \{B, y, w\}$, and $\chi(w, F_2, \Gamma_A) = -1/3$. Now applying $B \in \mathcal{B}_{F_1}$ and $|\partial(x) \setminus \partial(\Gamma_A)| = 6$ shows that there is an internal green face F_3 that shares an edge with B incident with e_2, and $B, w \in \mathcal{S}_{F_3}$. Similarly as for F_1 we deduce that all elements of $\mathcal{S}_{F_3}$ are curvature incident once with F_3, and that $\mathcal{S}_{F_3} = \{B, w, r\}$ has size 3, where $\chi(r, F_3, \Gamma_A) = -1/4$. In particular, $B \notin \mathcal{B}_{F_3}$, and $w \notin \mathcal{B}_{F_3}$ if w is a red blob (since then $|\partial(w)| = 4$). So by repeating the analysis from the previous paragraph we deduce that $|\mathcal{S}_{F_3}| \geq 4$, a contradiction. ■

The next lemma will be used in the proofs of Theorems 1 and 5.1.1.

Lemma 5.5.4. *Let F be an internal green face of $\Gamma_A \in \mathcal{T}$. Assume that the following statements hold.*

- *All red blobs of Γ_A are simply-connected, and none of them are complicated.*
- *F has an edge on at most one of the boundaries of Γ_A, and F is edge-incident with a red blob B with an edge on the opposite boundary from F.*

Then either B is highly hyperbolic and $|\partial(B)| \leq 6$, or B contains precisely one boundary edge and $|\partial(B)| \leq 5$.

Proof. Since no blobs are complicated, either B is highly hyperbolic, or B contains precisely one boundary edge. If B is highly hyperbolic, then by Lemma 4.6.11 we have $|\partial(B)| \leq 6$. So assume that B contains precisely one boundary edge. By Axiom T_6, F has a consolidated edge $l \subseteq \rho \in \{\omega, \tau\}$ with $|l| \geq 1$.

Let v and v_1 be the endpoints of l. If $v, v_1 \in \rho \cap \rho'$, then by Lemma 4.5.2 $\mathcal{S}_F$ contains no red blobs, a contradiction. Assume that $v \in \rho \cap \rho'$, so $v_1 \notin \rho \cap \rho'$. By Lemma 3.1.9 we have $\chi(v, F, \Gamma_A) = -1/2$. Since $v_1 \notin \rho \cap \rho'$, if $\delta_G(v_1) = 2$, then there is a red blob $B_1 \in \mathcal{B}_F$ incident with v_1 and with an edge on ρ. As B contains precisely one boundary edge, which is on ρ', we have $B_1 \neq B$. Otherwise, $v_1 \in \mathcal{B}_F \setminus \{v, B\}$, hence there is $x \in \mathcal{B}_F \setminus \{v, B\}$. By Part 2 of Lemma 4.3.4 we have $\chi(x, F, \Gamma_A) \leq -1/4$, and by Axiom T_6 we have $\kappa_{\Gamma_A}(F) = 0$. Hence $\chi(B, F, \Gamma_A) \geq -1/4$, so by Lemma 3.2.8 B is a red triangle, and the lemma follows.

Hence without loss of generality assume that $v, v_1 \notin \rho \cap \rho'$. Then by Corollary 4.5.15 there are $x, y \in \mathcal{B}_F \setminus \{B\}$ that lie on or are incident with l such that x and y collectively give F curvature of at most $-1/2$. Assume first that $\mathcal{S}_F = \{B, x, y\}$. Since $\kappa_{\Gamma_A}(F) = 0$, we have $\chi(B, F, \Gamma_A) \geq -1/2$, so if B is edge-incident more than once with F, then by Lemma 3.2.8 B is a red triangle, and the lemma follows. Hence we may assume that B is edge-incident once with F. Then by Lemma 3.2.8 we have $\chi(B, F, \Gamma_A) > -1/2$. So as $\kappa_{\Gamma_A}(F) = 0$, by Lemmas 2.6.10 and 3.2.8 x and y collectively give F curvature of at most $-1/4 - 1/3 = -7/12$. Suppose that this upper bound is attained. Then $\chi(B, F, \Gamma_A) = -5/12$, and since $\chi(x, F, \Gamma_A) = -1/3$, $\chi(y, F, \Gamma_A) = -1/4$ or $\chi(x, F, \Gamma_A) = -1/4$, $\chi(y, F, \Gamma_A) = -1/3$, applying Lemmas 2.6.10 and 3.2.8 shows that x and y are curvature incident once with F. But now Lemma 5.5.3 give us a contradiction. Hence by Lemmas 2.6.10 and 3.2.8 x and y collectively give F curvature of at most $-1/4 - 3/8 = -5/8$. So $\mathbf{Area}(B) \leq 3$, and by Lemma 3.2.6 $|\partial(B)| \leq 5$, as claimed.

Now assume that $\mathcal{S}_F \supsetneq \{B, x, y\}$. Then the elements of $\mathcal{S}_F \setminus \{B\}$ collectively give F curvature of at most $-1/6 - 1/2 = -2/3$. So $\mathbf{Area}(B) \leq 2$ and $|\partial(B)| \leq 4$. ∎

Proof of Theorem 5.1.1. We may assume that $\partial(F) \cap \rho' = \emptyset$. Suppose first that F is not thin. Then there is a red blob B in $\mathcal{S}_F$ with an edge on ρ'. Hence by Lemma 5.5.4 either B is highly hyperbolic and $|\partial(B)| \leq 6$, or B contains precisely one boundary edge and $|\partial(B)| \leq 5$, so the theorem holds. Now assume that F is thin. Then by Theorem 5.5.2 there is an internal green face F' with $\partial(F') \cap \partial(F) \neq \emptyset$, with an edge on ρ', and either F' and F are edge-incident, or $\mathcal{S}_{F'} \cap \mathcal{S}_F$ contains a red blob B with no boundary edge and $|\partial(B)| \leq 4$, hence the theorem follows. ∎

5.6 Proof of the Three Face Theorem

Recall Definition 3.1.3 of the minimal coloured area of an annular diagram; and Definitions 4.1.1, 4.6.1 and 4.6.8 of what it means for a red blob to be good, complicated, or highly hyperbolic. We shall now prove the Three Face Theorem (see Theorem 1). The Three Face Theorem will follow from Theorem 5.6.1, where we describe the structure of diagrams in $\mathcal{T}$ in more detail: we shall use Theorem 5.6.1 to make $\mathbf{IsConjugate}$ efficient. Throughout this section let $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$ be a valid pregroup presentation (see Definition 2.6.14), and let G be the group defined by $\mathcal{P}$.

Let Γ_A be an annular diagram with the external face O with boundary ω. Recall from Definition 2.5.4 that O is oriented counter-clockwise, and that all other faces of Γ_A are oriented clockwise.

Theorem 5.6.1. *Let $w_1 = a_1a_2\dots a_{n-1}a_n \in X^*$ and $w_2 \in X^*$ be such that no cyclic permutation of w_2 is equal in G to w_1, or to $w_1' = a_2\dots a_na_1$. Assume that there exists a diagram $\Gamma_A \in \mathcal{T}$ over $\mathcal{P}$ with $\mathbf{CArea}(\Gamma_A) \geq (1,0)$, and with boundaries ω and τ, labelled by w_1 and w_2 respectively. Then one of the following two statements holds.*

1. *There is an internal green face F with a consolidated edge on ω with a label containing a_1, and at least one of the following statements holds.*

   (i) $\partial(F) \cap \tau \neq \emptyset$.

   (ii) F is edge-incident with a simply-connected red blob B with an edge on τ. Moreover, either B is highly hyperbolic and $|\partial(B)| \leq 6$, or B contains precisely one boundary edge and $|\partial(B)| \leq 5$.

   (iii) There is an internal green face F_1 with an edge on τ, and either F_1 and F are edge-incident, or $\mathcal{S}_{F_1} \cap \mathcal{S}_F$ contains a simply-connected red blob B, where B does not contain a boundary edge and $|\partial(B)| \leq 4$.

   (iv) There is a red blob B edge-incident with F, and either B is complicated, or B is bad. Furthermore, if Γ_A has minimal coloured area, then there exists a retriangulation B_1 of B as described in Proposition 4.6.18.

2. *There is a red blob B with an edge on ω labelled by a_1, and one of the following statements holds.*

   (i) There is an internal green face F satisfying Statement 1, but with a_1 replaced by a_2.

   (ii) There is an internal green face F edge-incident with B, and with an edge on τ. Furthermore, B is simply-connected and either B is highly hyperbolic and $|\partial(B)| \leq 6$, or B contains precisely one boundary edge and $|\partial(B)| \leq 5$.

Proof. Since no cyclic permutation of w_2 is equal in G to w_1 or to w_1', the two endpoints v and v_1 of the edge e labelled by a_1 satisfy $v, v_1 \notin \omega \cap \tau$.

Suppose first that there is an internal green face F with a consolidated edge on ω that contains e. We show that Statement 1 of the theorem holds. Suppose first that no blobs of Γ_A are complicated, and assume further that Γ_A contains a bad red blob B. Then by Theorem 4.6.14 F is edge-incident with B, hence if Γ_A has minimal coloured area, then there is a triangulation

B_1 of B satisfying Proposition 4.6.18, so F satisfies Statement 1 (iv) of the theorem. Assume instead that all blobs of Γ_A are good. Then by Lemma 4.2.5 they are all simply-connected, so by Theorem 5.1.1 at least one of the Statements 1 (i)-(iii) of the theorem holds for F. Next suppose that F is contained in an island with a complicated red blob B. Then by Theorem 4.6.2 F is edge-incident with B, so if Γ_A has minimal coloured area, then there is a triangulation B_1 of B satisfying Proposition 4.6.18, hence F satisfies Statement 1 (iv). Finally, assume that F is not contained in any island with a complicated red blob, and that Γ_A contains complicated red blobs. Then by Theorem 4.6.2 Γ_A contains at least two islands, hence applying Lemma 3.1.10 repeatedly we can delete all islands containing complicated red blobs from Γ_A to obtain a diagram $\Delta_A \in \mathcal{T}$ containing the island E with $F \subseteq E$. By Lemma 4.2.4 all blobs of Δ_A are good, so by Lemma 4.2.5 they are all simply-connected, and therefore by Theorem 5.1.1 at least one of the Statements 1 (i)-(iii) of the theorem holds for F in Δ_A, so also in Γ_A.

Now assume that no internal green face of Γ_A contains a consolidated edge on ω that contains e. We show that Statement 2 of the theorem holds. Since $v, v_1 \notin \omega \cap \tau$, there is a red blob B with $e \subseteq B$. Suppose first that $|w_1| = 1$. If Γ_A contains an island, then $v, v_1 \in \omega \cap \tau$, a contradiction. Hence by Theorem 4.6.2 no blobs of Γ_A are complicated. If Γ_A contains a bad red blob, then by Theorem 4.6.14 either $|\omega| \geq 2$, or ω contains a green edge, a contradiction. Hence all red blobs of Γ_A are good, so by Lemma 4.2.5 they are all simply-connected. Suppose that there is no internal green face incident with v and edge-incident with B. By Axiom T_3 we have $\delta_G(v) \geq 2$, hence as by Axiom T_1, ω and τ are simple closed paths, we have $v \in \omega \cap \tau$, a contradiction. So there is an internal green face F incident with v and edge-incident with B. As the unique face with an edge on ω is red, by Axiom T_6, F has an edge on τ. Thus, by Lemma 5.5.4 Statement 2 (ii) of the theorem holds.

Suppose instead that $|w_1| > 1$, and assume further that there is an internal green face F with a consolidated edge on ω with a label containing a_2. Then by repeating the arguments from the second paragraph we deduce that Statement 2 (i) of the theorem holds. Otherwise, as $v, v_1 \notin \omega \cap \tau$, there is a red blob with an edge $e_1 \subseteq \omega$ labelled by a_2. Since e and e_1 are both red and consecutive on ω, by Theorem 4.6.2 B is not contained in any island with a complicated red blob; and by Theorem 4.6.14 all red blobs of Γ_A are good, so by Lemma 4.2.5 they are all simply-connected. Suppose first that no blobs of Γ_A are complicated. Let $u \in \overline{e} \cap \overline{e_1}$. Since $u \in \{v, v_1\}$ and $\delta_G(u) \geq 2$, similarly as in the previous paragraph we deduce that there is an internal green face F incident with u and edge-incident with B. By Theorem 4.5.13 we have $\partial(F) \cap \omega = \{u\}$. Hence by Axiom T_6, F has an edge on τ. Now applying Lemma 5.5.4 shows that Statement 2 (ii) holds. Assume instead that Γ_A contains complicated red blobs. Then Γ_A contains at least two islands, so applying Lemma 3.1.10 repeatedly we obtain a diagram $\Delta_A \in \mathcal{T}$ with no complicated red blobs, and containing the island E with $B \subseteq E$. Then similarly as before we deduce that there is an internal green face F of Δ_A edge-incident with B and with an edge on τ, and by Lemma 5.5.4 Statement 2 (ii) holds for B in Δ_A. Hence B and F satisfy Statement 2 (ii) in Γ_A. ∎

Proof of Theorem 1. Let $e \in \omega$ be an edge, let E be the external face with boundary τ, and let $w_1 = a_1 \ldots a_n$ and w_2 be labels of ω and τ respectively. Suppose that Γ_A has no green faces. Then by Lemma 3.2.10 Γ_A is island-free. Hence as by Axiom T_1, ω and τ are simple closed paths, by Axiom T_3, Γ_A has no faces, so $e \subseteq \tau$. Now if some cyclic permutation of w_2 is equal in G to w_1, or to $a_2 \ldots a_n a_1$, then some endpoint of e lies on τ, hence $\overline{e} \cap \tau \neq \emptyset$. Therefore, without loss of generality assume that $\mathbf{CArea}(\Gamma_A) \geq (1, 0)$, and that w_1 and w_2 satisfy assumptions of Theorem 5.6.1. Then Theorem 5.6.1 holds for Γ_A.

Assume first that Statement 1 of Theorem 5.6.1 holds. If at least one of the Statements (i)-(iii) holds, then there is an internal green face F that contains e such that either $\partial(F) \cap \tau \neq \emptyset$, or F is at dual distance at most three from E. So without loss of generality assume that Statement (iv) holds. Let F and B be as in Statement (iv). Note that $e \subseteq \overline{F}$. If B is complicated, then by Theorem 4.6.2 B contains an edge on τ, so F is at dual distance at most two from E. Hence without loss of generality assume that B is bad. If B contains a boundary edge, then by Theorem 4.6.13 B contains an edge on τ, so we are done. Otherwise, by Theorem 4.6.14 F is at dual distance at most three from E.

Now assume that Γ_A satisfies Statement 2 of Theorem 5.6.1. If Statement (i) holds, then by the previous paragraph Statement 2 of the theorem holds for e. So suppose that Statement (ii) holds. Then there is an internal face that contains e at dual distance at most two from E. ■

Chapter 6

Conjugacy problem solver

In this chapter we describe the conjugacy problem solver $\mathbf{IsConjugate}$, and prove Theorem 1.0.3. The conjugacy tests of $\mathbf{IsConjugate}$ run on valid pregroup presentations (see Definition 2.6.14), and all input pregroup presentations to $\mathbf{IsConjugate}$ are already assumed to be by sound (see Definition 2.6.14). To check that a sound pregroup presentation $\mathcal{P}$ is valid, we only need to check that $\mathcal{P}$ is proper (see Definition 2.6.14):

Definition 6.0.1. The procedure $\mathbf{IsProper}(\mathcal{P})$ returns `true` if $\mathcal{P}$ is proper. Otherwise, $\mathbf{IsProper}(\mathcal{P})$ returns `false`.

Throughout the rest of this chapter let $\mathcal{P} = \langle X^\sigma \mid V_P \cup \mathcal{R} \rangle$ be a valid pregroup presentation, and let G be defined by $\mathcal{P}$. Recall that we assume the RAM model of computation, in which the basic arithmetical operations on integers can be carried out in constant time. Furthermore, we assume throughout this chapter that the products and inverses in the pregroup can also be computed in constant time.

6.1 Cyclic $\mathcal{P}$-reduction

In this section we describe a modification of $\mathbf{RSymSolve}$ (see [34, Section 8]), $\mathbf{RSymSolveSimpler}$, that cyclically $\mathcal{P}'$-reduces (see Definition 3.1.15) a given word when $\mathcal{P}'$ satisfies trivial-interleaving (see Definition 2.3.23).

Since in the description of $\mathbf{RSymSolve}$ the subscripts are interpreted cyclically, we note that $\mathbf{RSymSolve}$ is already performing cyclic $\mathcal{P}$-reduction. We thus give its simplification convenient for our purposes. Similarly as in [34, Section 8], we compute a list $\mathcal{L}_1$, which has entries pairs of words $(u, v) = (u_1 \ldots u_k, v_1 \ldots v_l) \in X^* \times X^*$, where $u_1 u_2 \ldots u_{k-1} u_k (v_1 v_2 \ldots v_{l-1} v_l)^{-1}$ is a cyclic conjugate of some $R \in (V_P \cup \mathcal{R}^{\pm 1})$ and $k = \lceil (|R| + 1)/2 \rceil$. We allow $R \in V_P$ because the input word is not assumed to be cyclically P-reduced. We interpret all subscripts cyclically, so that $x_{n+1} = x_1$. Also, we let $r := \max\{|R| : R \in \mathcal{R}\}$ be the length of the longest green relator.

Algorithm 6.1.1. $\mathbf{RSymSolveSimpler}(w = x_1 \ldots x_n)$:

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 Search for $\alpha \leq i \leq n$ such that $x_i x_{i+1} =_{F(X^\sigma)} 1$.

(i) Let $i, u := x_i x_{i+1}$ be the first such found, if any. Let $v := \eta$ be the empty word, let $m = 2$, and go to Step 5.

(ii) If none such exists then w is already cyclically σ-reduced. In that case go to Step 4.

Step 4 For each $\alpha \leq i \leq n$ do: search for $m \in \{1, \ldots, \lceil (r+1)/2 \rceil\}$ and $(u, v) \in \mathcal{L}_1$ such that

$$x_i x_{i+1} \ldots x_{i+m-2} x_{i+m-1} =_{F(X^\sigma)} u.$$

(i) Let $i, m, v := v_1 \ldots v_l$ be the first such found, if any.

(ii) If none such exist then w is already cyclically $\mathcal{P}$-reduced. Terminate and return w.

Step 5 Put a pointer `CutStart` to x_{i-1} and a pointer `CutEnd` to x_{i+m}.

Step 6 Replace u by v in w and update the links in the list describing w so that v is inserted into the correct place in $x_1 \ldots x_n$, yielding a word w_1.

Step 7 Let j be the position in w_1 to which `CutStart` points, and let $\alpha := \max\{1, j - \lceil (r+1)/2 \rceil + 1\}$. Replace n by $|w_1|$, and go to Step 3 with w_1 in place of w.

Step 8 Repeat the process above until no further reductions are found, resulting with a word w'. Return w'.

In the statement of the next proposition we treat $|X|$, $|\mathcal{R}|$ and r as constants.

Proposition 6.1.2. *For all $n \in \mathbb{N}$, and for all $x_1 \ldots x_n \in X^*$,* $\mathbf{RSymSolveSimpler}(w = x_1 \ldots x_n)$ *finds a cyclically $\mathcal{P}$-reduced word that is G-conjugate to w in time $O(n)$.*

Proof. By Remark 2.6.15 we can solve the word problem in G by the standard Dehn algorithm using the length reducing rewrite rules derived from $V_P \cup \mathcal{R}$ (i.e. $\mathcal{P}$ is a Dehn presentation). Hence if w is not cyclically $\mathcal{P}$-reduced, then at least one of the following statements holds.

1. $x_i x_{i+1} =_{F(X^\sigma)} 1$ for some $1 \leq i \leq n$;

2. some cyclic permutation of w contains a contiguous sub-word $u = u_1 u_2 \ldots u_{k-1} u_k$, such that there exists a cyclic conjugate $uv = u_1 u_2 \ldots u_k (v_1 v_2 \ldots v_l)^{-1}$ of some $R \in (V_P \cup \mathcal{R}^{\pm 1})$ and $k = \lceil (|R| + 1)/2 \rceil$.

Hence either i, or the word u will be found by $\mathbf{RSymSolveSimpler}$. To find such i or u, $\mathbf{RSymSolveSimpler}$ runs

(i) $O(n)$ tests $w_1 =_{F(X^\sigma)} 1$, where w_1 is a length two sub-word of w; and

(ii) $O(n)$ tests of equality in $F(X^\sigma)$ of words $w_1 = t_1 \ldots t_m$ with $m \leq (r+1)/2$ that are sub-words of w, with the first entry of each pair in $\mathcal{L}_1$.

By Corollary 2.3.12 we can solve the word problem in $U(P)$ (and hence also in $F(X^\sigma)$) in linear time, hence each such test takes time $O(1)$. Moreover, after each replacement of a sub-string of w we backtrack at most $O(1)$ letters, and each such replacement shortens $|w|$, so at most $O(n)$ replacements are carried out. Hence the overall complexity of $\mathbf{RSymSolveSimpler}(w)$ is as stated. ■

Remark 6.1.3. We note that if the input word w to $\mathbf{RSymSolveSimpler}$ has $w =_G 1$, then $\mathbf{RSymSolveSimpler}(w) = \eta$. We implemented $\mathbf{RSymSolveSimpler}$ using the two stack model, the principle is described for example in [22].

6.2 Algorithms for conjugacy diagrams in $\mathcal{T}$

In this section we present algorithms for analysing diagrams in $\mathcal{T}$. Let $w_1 = a_1 \dots a_n \in X^*$ and $w_2 = b_1 \dots b_m \in X^*$.

Definition 6.2.1. [34, Definition 7.3] Let $R \in \mathcal{R}^{\pm}$, and fix a word $w = x_1 x_2 \dots x_{|w|}$ such that $R = w^k$ with k maximal amongst such expressions for R. We call w the *root* of R. A *location* on R is an ordered triple (i, a, b), denoted by $R(i, a, b)$, where $i \in \{1, \dots, |w|\}$, $a = x_{i-1}$ (or $x_{|w|}$ if $i = 1$), and $b = x_i$.

Definition 6.2.2. [34, Definition 7.11] We call a letter $x \in X^{\sigma}$ an $\mathcal{R}$-letter if x occurs in $R \in \mathcal{R}^{\pm}$. Observe that since we ignore the internal structure of red blobs, if $x \in X^{\sigma}$ is a non-$\mathcal{R}$-letter, then x can appear only on $\partial(\Gamma_A)$.

Throughout the rest of this section let `lr` be the length of the root of w_2. $\mathbf{IsConjugate}$ uses various auxiliary sub-routines. In this section, we present algorithms that check whether there exists a diagram in $\mathcal{T}$ with boundary words w_1 and w_2 (see Theorem 6.2.15). We describe them via pseudocodes (see Algorithms 6.2.3-6.2.13).

Our algorithms seek to find $x \in X^*$ such that $x w_2' x^{-1} w_1^{-1} =_G 1$ for some cyclic permutation w_2' of w_2 (note that it suffices to consider only cyclic permutations of w_2 that start in position $1 \leq x \leq \texttt{lr}$). If they succeed, then by concatenating the inverse of the appropriate prefix of w_2 with x they return $x_1 \in X^*$ such that $x_1 w_2 x_1^{-1} =_G w_1$. We prove their correctness at the end of this section, see Lemma 6.2.14 and Theorem 6.2.15. To check that a given word is equal to the identity in G, we use $\mathbf{RSymSolveSimpler}$ as it gives the correct answer by Remark 6.1.3.

The sub-routine $\mathbf{CyclicConj}$ checks (and finds a conjugating word) whether there is a cyclic permutation of w_2 equal in G to w_1, or to $w_1' = a_2 \dots a_n a_1$. Note that $\mathbf{CyclicConj}$ checks whether w_1 and w_2 satisfy assumptions of Theorem 5.6.1.

Algorithm 6.2.3. $\mathbf{CyclicConj}(w_1, w_2 = b_1 b_2 \dots b_m)$:

Step 1 For each $x \in [1 \dots \texttt{lr}]$ do:

(i) Let $w_2' = b_x \dots b_m b_1 \dots b_{x-1}$. If $\mathbf{RSymSolveSimpler}(w_2' w_1^{\sigma}) = \eta$ then return `true`, $(b_1 b_2 \dots b_{x-1})^{\sigma}$.

(ii) If $\mathbf{RSymSolveSimpler}(w_2' w_1'^{\sigma}) = \eta$ then return `true`, $a_1 (b_1 b_2 \dots b_{x-1})^{\sigma}$.

Step 2 Return `false`.

We next compute a list $\mathcal{B}$ of cyclic words $w \in X^*$ that satisfy all of the following conditions:

1. w is equal to 1 in $U(P)$.

2. $3 \leq |w| \leq 6$.

3. No proper non-empty sub-word of w is equal to 1 in $U(P)$.

4. Each consecutive pair of letters in w multiply.

5. w contains at most two non-$\mathcal{R}$-letters.

Lemma 6.2.4. *Assume that that there exists a diagram $\Gamma_A \in \mathcal{T}$ with boundary words w_1 and w_2 that satisfies Theorem 5.6.1.*

If one of the Statements 1 (ii)-(iii), or Statement 2 (ii) of Theorem 5.6.1 holds for Γ_A, then the list $\mathcal{B}$ contains all potential boundary words of the red blob B.

Proof. Let w be the boundary word of B. Since B is simply-connected, w is equal to 1 in $U(P)$, so w satisfies Condition 1. As B contains at least one red triangle, we have $|w| \geq 3$. Hence Condition 2 holds for w.

By Axiom T_4 (see Definition 3.1.8) Condition 3 also holds for w, and by Lemma 2.5.13 w satisfies Condition 4. Since a highly hyperbolic red blob (see Definition 4.6.8) contains two boundary edges, B contains at most two boundary edges. Hence w contains at most two non-$\mathcal{R}$-letters. Thus, $\mathcal{B}$ satisfies the lemma. ■

Definition 6.2.5. [34, 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.

(i) There exists a σ-reduced annular or simply-connected diagram Γ (see Definition 2.5.15) 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.

We shall work only with instantiable places. If $C = \mathbf{G}$, then as 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 shall work only with the case where f_2 is simply-connected, hence by Axiom T_4 the boundary words of f_2 are cyclically σ-reduced, and therefore by Lemma 2.5.13 (b^σ, c) is an intermult pair (see Definition 2.3.13).

Before running the algorithms described below, we compute all instantiable places, as follows. We first compute an array of all intermult pairs, and then find all locations $R(i, a, b)$ with $R \in \mathcal{R}^{\pm}$. For each such location $R(i, a, b)$, each $c \in X$, and each $C \in \{\mathbf{R}, \mathbf{G}\}$, if $C = \mathbf{R}$ then we check that b^σ intermults with c, else we check if there exists a location $R'(j, b^\sigma, c)$ with the property that a simply-connected diagram equal to the union of faces labelled by R and R' that share the edge with label b is σ-reduced.

We now present several sub-routines, where each of them tries to find a diagram in $\mathcal{T}$ with boundary words w_1 and w_2 that satisfies Theorem 5.6.1. If a sub-routine returns returns `false`, then there does not exist such diagram. Else it returns `true`; and if it returns `true`, then it also returns an $x \in X^*$ such that $w_1 =_G w_2^{x^{-1}}$. In our pseudocodes `Rels` is a record containing all words in $\mathcal{R}^{\pm 1}$, and the root for each $R \in \mathcal{R}^{\pm 1}$.

We first consider Case 1 of Theorem 5.6.1. The sub-routine **EdgeWithFirstLetter** stores all $R \in \mathcal{R}^{\pm}$ (and their cyclic permutations) starting with a_1. We are only checking the first letter of w_1 since there does not have to be longer common sub-words. Note that **EdgeWith FirstLetter** stores all labels of internal green faces F satisfying Statement 1 of Theorem 5.6.1.

Algorithm 6.2.6. **EdgeWithFirstLetter**$(w_1, \texttt{Rels})$:

Step 1 Initialize $L := [\,]$.
Step 2 For each $R \in \mathcal{R}^{\pm 1}$ do:

- (i) Let $w = w_1 \dots w_\iota$ be the root of R. Write $R = w^k$.
- (ii) For each $1 \le i \le \iota$ do:
  - (a) If $w_i = a_1$ then append $w_i \dots w_\iota w^{k-1} w_1 \dots w_{i-1}$ to L.

Step 3 Return L.

In descriptions of the remaining algorithms we let $L = \textbf{EdgeWithFirstLetter}(w_1, \texttt{Rels})$. The next sub-routine, **BothBoundaries**, checks whether there exists a diagram in $\mathcal{T}$ that satisfies Statement 1 (i) of Theorem 5.6.1. It does so by checking whether any relator found by **EdgeWithFirstLetter** has a prefix x such that $w_1 =_G w_2^{x^{-1}}$.

Algorithm 6.2.7. **BothBoundaries**$(w_1, w_2 = b_1 b_2 \dots b_m, L, \texttt{Rels})$:

Step 1 For each $x \in [1 \dots \texttt{lr}]$ for each $R \in L$ and for each prefix c of R do:

- (i) If **RSymSolveSimpler**$(c b_x \dots b_{x-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{x-1})^\sigma$.

Step 2 Return `false`.

The function **ConjByTwoLabels** checks if there exists a diagram in $\mathcal{T}$ that satisfies Statement 1 (ii) of Theorem 5.6.1, or a diagram in $\mathcal{T}$ with edge-incident green faces F and F_1 that satisfy Statement 1 (iii) of Theorem 5.6.1. For all relators R (and their cyclic permutations) that potentially label an internal green face F that satisfies Statement 1 of Theorem 5.6.1, it checks whether any internal green face or a red blob edge-incident with F has an edge on τ. To find the boundary word of such a blob, by Lemma 6.2.4 we can use $\mathcal{B}$.

Algorithm 6.2.8. **ConjByTwoLabels**$(w_1, w_2, L, \mathcal{B}, \texttt{Rels})$:

Step 1 For each $x \in [1 \dots \texttt{lr}]$ for each $R = r_1 r_2 \dots r_k \in L$ and for each $2 \le i \le k$ do:

- (i) For each location $R_1(j, r_i^\sigma, d)$ instantiating a green place $(R(\iota, r_{i-1}, r_i), d, \mathbf{G})$ do:

(a) For each letter u of R_1 with $b_x = u^\sigma$ do: write $R_1 = s_1 r_i^\sigma d s_2 u s_3$, for some $s_1, s_2, s_3 \in X^*$, and let $c := (r_1 r_2 \dots r_{i-1}) d s_2 u$.

If $\mathbf{RSymSolveSimpler}(c b_x \dots b_{x-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{x-1})^\sigma$.

(ii) For each red place $(R(\iota, r_{i-1}, r_i), d, \mathbf{R})$ and for each $3 \leq l \leq 6$ do:

(a) For each $B = d m_2 \dots m_{l-1} r_i^\sigma \in \mathcal{B}$ and for each letter m of B with $b_x = m^\sigma$ do:

Let $c := (r_1 r_2 \dots r_{i-1}) d m_2 \dots m$. If $\mathbf{RSymSolveSimpler}(c\, b_x \dots b_{x-1}\, c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{x-1})^\sigma$.

Step 2 Return `false`.

Using the same ideas as in our previous sub-routines we now present a sub-routine $\mathbf{ConjBy\ ThreeLabels}$, which checks if there exists a diagram in $\mathcal{T}$ with green faces F, F_1 and a red blob B that satisfy Statement 1 (iii) of Theorem 5.6.1.

Algorithm 6.2.9. $\mathbf{ConjByThreeLabels}(w_1, w_2, L, \mathcal{B}, \mathtt{Rels})$:

Step 1 For each $x \in [1 \dots \mathtt{lr}]$ for each $R = r_1 r_2 \dots r_k \in L$ and for each $2 \leq i \leq k$ do:

(i) For each red place $(R(\iota, r_{i-1}, r_i), d, \mathbf{R})$ and for each $3 \leq l \leq 4$ do:

(a) For each $B = m_1 \dots m_l = d m_2 \dots m_{l-1} r_i^\sigma \in \mathcal{B}$ and for each $1 \leq j \leq l$ do:

(A) For each location $R_1(\iota_1, m_j^\sigma, e)$ and for each letter u of R_1 with $b_x = u^\sigma$ do:

Write $R_1 = s_1 m_j^\sigma e s_2 u s_3$, for some $s_1, s_2, s_3 \in X^*$, and let $c := (r_1 r_2 \dots r_{i-1})(d m_2 \dots m_{j-1})\, e s_2 u$.

If $\mathbf{RSymSolveSimpler}(c b_x \dots b_{x-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{x-1})^\sigma$.

Step 2 Return `false`.

The sub-routine $\mathbf{ComplicatedRedBlobs}$ checks if there exists a diagram in $\mathcal{T}$ that satisfies Statement 1 (iv) of Theorem 5.6.1. As we shall see later, if such a diagram exists, then we can assume that it has minimal coloured area (see Definition 3.1.3). Hence $\mathbf{ComplicatedRedBlobs}$ checks if there is a word $xy \in X^*$ or $xyz \in X^*$ that satisfies Proposition 4.6.18. $\mathbf{ComplicatedRedBlobs}$ uses a list $\mathcal{L}_{23}$ that contains all $xy \in X^*$ such that $x \neq y^\sigma$ and $(x, y) \in D(P)$; and all $xyz \in X^*$ such that $(x, y) \in D(P)$, $(y, z) \in D(P)$, and no sub-word of xyz is trivial in $U(P)$.

Algorithm 6.2.10. $\mathbf{ComplicatedRedBlobs}(w_1, w_2, L, \mathcal{L}_{23}, \mathtt{Rels})$:

Step 1 For each $t \in [1 \dots \mathtt{lr}]$ for each $R = r_1 r_2 \dots r_k \in L$ and for each $2 \leq i \leq k$ do:

(i) For each location $R(\iota, r_{i-1}, r_i)$ and for each $w = m_1 \dots m_n \in \mathcal{L}_{23}$ with $m_1 = r_i$ do:

(a) Let $c := (r_1 r_2 \dots r_{i-1}) w$. If $\mathbf{RSymSolveSimpler}(c b_t \dots b_{t-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{t-1})^\sigma$.

(b) If $|w| = 3$ then for each location $R_1(\iota_1, m_3, e)$ and for each letter u of R_1 with $b_t = u^\sigma$ do:

(A) Write $R_1 = s_1 m_3 e s_2 u s_3$, for some $s_1, s_2, s_3 \in X^*$, and let $c := (r_1 r_2 \dots r_{i-1})(m_1 m_2)(m_3 e s_2)u$.

If **RSymSolve Simpler**$(cb_t \dots b_{t-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{t-1})^\sigma$.

Step 2 Return `false`.

Finally, we consider Case 2 of Theorem 5.6.1. The algorithm **StartAtIthLetter** takes an input $1 \leq i \leq |w_1| = n$, and runs Algorithms 6.2.6-6.2.10 on input w_1' and w_2, where $w_1' = a_i \dots a_n a_{i-1}$. In the case $i = 2$, **StartAtIthLetter** checks if there exists a diagram in $\mathcal{T}$ that satisfies Statement 2 (i) of Theorem 5.6.1.

Algorithm 6.2.11. **StartAtIthLetter**$(i, w_2, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$:

Step 1 $w_1' := a_i \dots a_n a_{i-1}$; $L :=$ **EdgeWithFirstLetter**$(w_1', \texttt{Rels})$ (see Algorithm 6.2.6).

Step 2 $\texttt{conj}, c :=$ **BothBoundaries**$(w_1', w_2, L, \texttt{Rels})$ (see Algorithm 6.2.7). If `conj` then return `true`, $a_1 \dots a_{i-1} c$.

Step 3 $\texttt{conj}, c :=$ **ConjByTwoLabels**$(w_1', w_2, L, \mathcal{B}, \texttt{Rels})$ (see Algorithm 6.2.8). If `conj` then return `true`, $a_1 \dots a_{i-1} c$.

Step 4 $\texttt{conj}, c :=$ **ConjByThreeLabels**$(w_1', w_2, L, \mathcal{B}, \texttt{Rels})$ (see Algorithm 6.2.9). If `conj` then return `true`, $a_1 \dots a_{i-1} c$.

Step 5 $\texttt{conj}, c :=$ **ComplicatedRedBlobs**$(w_1', w_2, L, \mathcal{L}_{23}, \texttt{Rels})$ (see Algorithm 6.2.10). If `conj` then return `true`, $a_1 \dots a_{i-1} c$.

Step 6 Return `false`.

It remains to consider Case 2 (ii) of Theorem 5.6.1. Using the same ideas as before the sub-routine **StartWithRedBlob** checks if there exists a diagram in $\mathcal{T}$ with a red blob B and a green face F that satisfy Statement 2 (ii) of Theorem 5.6.1.

Algorithm 6.2.12. **StartWithRedBlob**$(w_1, w_2, \mathcal{B}, \texttt{Rels})$:

Step 1 For each $x \in [1 \dots \texttt{lr}]$ for each $B = m_1 m_2 \dots m_t \in \mathcal{B}$ with $m_1 = a_1$ and for each $2 \leq j \leq t$ do:

(i) For each location $R(\iota, m_j^\sigma, d)$ and for each letter u of R with $b_x = u^\sigma$ do:

(a) Write $R = m_j^\sigma d s_1 u s_2$, for some $s_1, s_2 \in X^*$, and let $c := m_1 m_2 \dots m_{j-1} d s_1 u$.

If **RSymSolveSimpler**$(cb_x \dots b_{x-1} c^\sigma w_1^\sigma) = \eta$ then return `true`, $c(b_1 b_2 \dots b_{x-1})^\sigma$.

Step 2 Return `false`.

Finally, we present the algorithm **ConjInT** that uses the sub-routines presented above to check if there exists a diagram in $\mathcal{T}$ with boundary words w_1 and w_2.

Algorithm 6.2.13. **ConjInT**$(w_1, w_2, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$:

Step 1 Find the length of the root of w_2. $\texttt{conj}, c :=$ **CyclicConj**(w_1, w_2) (see Algorithm 6.2.3). If `conj` then return `true`, c.

Step 2 $\texttt{conj}, c :=$ **StartAtIthLetter**$(1, w_2, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$ (see Algorithm 6.2.11). If $\texttt{conj}$ then return $\texttt{true}$, c.

Step 3 $\texttt{conj}, c :=$ **StartWithRedBlob**$(w_1, w_2, \mathcal{B}, \texttt{Rels})$ (see Algorithm 6.2.12). If $\texttt{conj}$ then return $\texttt{true}$, c.

Step 4 If $|w_1| > 1$ then $\texttt{conj}, c :=$ **StartAtIthLetter**$(2, w_2, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$. If $\texttt{conj}$ then return $\texttt{true}$, c.

Step 5 Return $\texttt{false}$.

We shall show that if **ConjInT** returns $\texttt{false}$, then there does not exist a diagram in $\mathcal{T}$ with boundary words w_1 and w_2. Before doing so, we prove the following auxiliary lemma.

Lemma 6.2.14. *Let* $r := \max\{|R| : R \in \mathcal{R}\}$. *All of the following statements hold.*

1. *The lists* $\mathcal{L}_{23}$ *and* $\mathcal{B}$ *can be constructed in time* $O(|X|^3)$ *and* $O(|X|^6)$ *respectively, and* $|\mathcal{L}_{23}| = O(|X|^3)$, $|\mathcal{B}| = O(|X|^6)$.

2. *The length of the root of* w_2 *can be found in time* $O(|w_2|^2)$.

3. *If* **CyclicConj**(w_1, w_2) *returns* $\texttt{true}$*, then it returns an* $x \in X^*$ *such that* $w_1 =_G w_2^{x^{-1}}$*, and* w_1 *and* w_2 *do not satisfy assumptions of Theorem 5.6.1 if and only if* **CyclicConj**(w_1, w_2) *returns* $\texttt{true}$*.*

   The running time of **CyclicConj**(w_1, w_2) *is* $O((|w_1| + |w_2|) \cdot |w_2|)$.

4. **EdgeWithFirstLetter**$(w_1, \texttt{Rels})$ *stores a list of size* $O(r|\mathcal{R}|)$ *containing all potential labels of internal green faces* F *that satisfy Statement 1 of Theorem 5.6.1.*

   EdgeWithFirstLetter$(w_1, \texttt{Rels})$ *has time complexity* $O(r^2|\mathcal{R}|)$.

Proof. Proof of 1. The proof of complexity is very similar to the proof of complexity of Step 6 of **RSymVerify** (see [34, Procedure 7.19]), see the proof of [34, Theorem 7.22]. Since all $w \in \mathcal{L}_{23}$ have $|w| \leq 3$ and all $w \in \mathcal{B}$ have $|w| \leq 6$, it follows that $|\mathcal{L}_{23}| = O(|X|^3)$ and $|\mathcal{B}| = O(|X|^6)$.

Proof of 2. We first find w that maximises the value of k for which $w_2 = w^k$. For $2 \leq l \leq |w_2|/2$, we let w be the prefix of w_2 of length l, and test whether $w^{|w_2|/l} = w_2$, in time $O(|w_2|^2)$.

Proof of 3. If **CyclicConj** returns $\texttt{true}$, then the $x \in X^*$ returned by it satisfies $w_1 =_G w_2^{x^{-1}}$, and the cyclic permutation w_2' of w_2 is equal in G to w_1 or to $w_1' = a_2 \ldots a_n a_1$. The reverse implication follows similarly. The complexity statement follows from Proposition 6.1.2 since **RSymSolveSimpler** runs in linear time.

Proof of 4. **EdgeWithFirstLetter** stores all $R \in \mathcal{R}^{\pm}$ (and their cyclic permutations) starting with a_1, hence the first statement follows. As each of the $O(r|\mathcal{R}|)$ checks of equality $w_i = a_1$ performed by **EdgeWithFirstLetter** takes constant time, and adding a word to the list L takes time $O(r)$, we see that **EdgeWithFirstLetter** runs in the stated time. ■

The following theorem constitutes the main result of this section.

Theorem 6.2.15. *Assume that there exists $\Gamma_A \in \mathcal{T}$ over $\mathcal{P}$ with boundary words w_1 and w_2, with* $\mathbf{CArea}(\Gamma_A) \geq (1,0)$*, and of minimal coloured area. Then* $\mathbf{ConjInT}(w_1, w_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ *returns* `true`*; and if it returns* `true`*, then it returns an $x \in X^*$ such that $w_1 =_G w_2^{x^{-1}}$.*

Let $r := \max\{|R| : R \in \mathcal{R}\}$. The running time of $\mathbf{ConjInT}(w_1, w_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ *is $O((r^4|\mathcal{R}|^2|X|^3 + r^2|\mathcal{R}||X|^5) \cdot (|w_1| + |w_2|) \cdot |w_2|)$.*

Proof. Assume that there exists a diagram Γ_A with properties as in the statement of the theorem. By Part 3 of Lemma 6.2.14 we can without loss of generality assume that $\mathbf{CyclicConj}(w_1, w_2)$ returns `false`. Then w_1 and w_2 satisfy assumptions of Theorem 5.6.1, hence Theorem 5.6.1 holds for Γ_A. Furthermore, by Part 4 of Lemma 6.2.14 $\mathbf{EdgeWithFirstLetter}(w_1, \mathtt{Rels})$ stores a list L containing all potential labels of internal green faces F of Γ_A that satisfy Statement 1 of Theorem 5.6.1.

Suppose first that Γ_A satisfies Statement 1 of Theorem 5.6.1. Assume that Γ_A satisfies Statement (i). Then by construction $\mathbf{BothBoundaries}(w_1, w_2, L, \mathtt{Rels})$ returns `true`; and if it returns `true`, then it returns an $x \in X^*$ such that $w_1 =_G w_2^{x^{-1}}$. Assume next that Γ_A satisfies Statement (ii), or that there are edge-incident green faces F and F_1 that satisfy Statement (iii). Then by construction $\mathbf{ConjByTwoLabels}(w_1, w_2, L, \mathcal{B}, \mathtt{Rels})$ returns `true`; and if it returns `true`, then it returns an $x \in X^*$ satisfying the theorem. If Γ_A does not satisfy any of the previous two assumptions, then either: (a) There are green faces $F, F_1 \subseteq \Gamma_A$ and a red blob $B \subseteq \Gamma_A$ that satisfy Statement (iii), or (b) Γ_A satisfies Statement (iv). Suppose first that Case (a) holds for Γ_A. Then $\mathbf{ConjByThreeLabels}(w_1, w_2, L, \mathcal{B}, \mathtt{Rels})$ returns `true`; and if it returns `true`, then it returns an $x \in X^*$ with the desired properties. Now suppose that Case (b) holds for Γ_A. Then as Γ_A has minimal coloured area, there exists a retriangulation B_1 of B satisfying Proposition 4.6.18. Hence by construction, $\mathbf{ComplicatedRedBlobs}(w_1, w_2, L, \mathcal{L}_{23}, \mathtt{Rels})$ returns `true`; and if it returns `true`, then it returns an $x \in X^*$ such that $w_1 =_G w_2^{x^{-1}}$. Thus, if Γ_A satisfies Statement 1 of Theorem 5.6.1, then $\mathbf{StartAtIthLetter}(1, w_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ returns `true`. Also, if $\mathbf{StartAtIthLetter}(i, w_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ returns `true`, then it returns an $x \in X^*$ with $w_1 =_G w_2^{x^{-1}}$.

Assume next that Γ_A satisfies Statement 2 (i) of Theorem 5.6.1. Then by above $\mathbf{StartAtIthLetter}(2, w_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ returns `true`. Finally, assume that Γ_A satisfies Statement 2 (ii). Then by construction $\mathbf{StartWithRedBlob}(w_1, w_2, \mathcal{B}, \mathtt{Rels})$ returns `true`; and if it returns `true`, then it returns an $x \in X^*$ with $w_1 =_G w_2^{x^{-1}}$. Thus, the first statement of the theorem holds.

To prove the second statement, first note that by Lemma 6.2.14 it suffices to analyse time complexity of Algorithms 6.2.7-6.2.10 & 6.2.12.

Time complexity of $\mathbf{BothBoundaries}(w_1, w_2, L, \mathtt{Rels})$. By Part 4 of Lemma 6.2.14 $|L| = O(r|\mathcal{R}|)$. Now for each $1 \leq x \leq \mathtt{lr}$ and for each $R \in L$, finding c takes time $O(r)$. Hence as by Proposition 6.1.2 $\mathbf{RSymSolveSimpler}$ runs in linear time, it follows that $\mathbf{BothBoundaries}$ runs in time $O((r^2|\mathcal{R}|) \cdot ((|w_1| + |w_2|) \cdot |w_2|))$.

Time complexity of $\mathbf{ConjByTwoLabels}(w_1, w_2, L, \mathcal{B}, \mathtt{Rels})$ *and of* $\mathbf{ConjByThree Labels}(w_1, w_2, L, \mathcal{B}, \mathtt{Rels})$. Each red and green place on each $R \in L$ can be found in time $O(r|X|)$, and each green place has $O(r|\mathcal{R}|)$ locations instantiating it. For relator R_1 the letter u can be found in time $O(r)$. Hence as **RSymSolveSimpler** runs in linear time, Part (i) of **ConjByTwoLabels** takes time $O((r^4\,|\mathcal{R}^2||X|) \cdot (|w_1| + |w_2|) \cdot |w_2|))$.

In Part (ii) of **ConjByTwoLabels**, by Part 1 of Lemma 6.2.14 the boundary word B can be found in time $O(|X|^6)$, and there are at most $O(|X|^4)$ such words B (since the first and the last letter of B is fixed). Furthermore, the letter m can be found in constant time (since there are at most 6 possibilities). Hence Part (ii) of **ConjByTwoLabels** takes time $O((r^2|\mathcal{R}||X|^5) \cdot (|w_1| + |w_2|) \cdot |w_2|))$. The overall complexity of **ConjByTwoLabels** is therefore $O((r^4|\mathcal{R}|^2|X| + r^2|\mathcal{R}||X|^5) \cdot ((|w_1| + |w_2|) \cdot |w_2|))$.

Using the same ideas as in the previous two paragraphs we deduce that **ConjByThree Labels** has time complexity $O((r^4|\mathcal{R}|^2|X|^3) \cdot (|w_1| + |w_2|) \cdot |w_2|)$.

Time complexity of $\mathbf{ComplicatedRedBlobs}(w_1, w_2, L, \mathcal{L}_{23}, \mathtt{Rels})$. Observe that we obtain the time complexity of **ComplicatedRedBlobs** by analysing Part (b) of Step 1. The word $w = m_1m_2m_3 \in \mathcal{L}_{23}$ with $m_1 = r_i$ can be found in time $O(|X|^3)$, and there are at most $O(|X|^2)$ such words. So similarly as for Algorithms 6.2.8-6.2.9 we deduce that **ComplicatedRedBlobs** runs in time $O((r^4|\mathcal{R}|^2|X|^2) \cdot (|w_1| + |w_2|) \cdot |w_2|)$.

Time complexity of $\mathbf{StartWithRedBlob}(w_1, w_2, \mathcal{B}, \mathtt{Rels})$. By Part 1 of Lemma 6.2.14 the word B can be found in time $O(|X|^6)$, and there are at most $O(|X|^5)$ such words. A location $R(\iota, m_j^\sigma, d)$ can be found in time $O(r|R|)$, and the letter u of R can be found in time $O(r)$. Hence **RSymSolveSimpler** performs at most $O(r^2|\mathcal{R}||X|^5)$ tests, so **StartWithRedBlob** runs in time $O((|r^2|\mathcal{R}||X|^5) \cdot ((|w_1| + |w_2|) \cdot |w_2|))$.

Hence as each of the Algorithms 6.2.7-6.2.10 & 6.2.12 is run a finite number of times by **ConjInT**, the theorem follows. ■

6.3 Algorithms for conjugacy diagrams in $\mathcal{S}$

This section presents procedures for analysing minimal conjugacy diagrams (see Definition 3.1.4) that contain loops labelled by single letters. We describe a procedure, **ConjLetters** (see Procedure 6.3.7), which finds conjugacy classes of single letters in G.

We begin with the following observation that holds under weaker assumptions on $\mathcal{P}$ than being valid.

Lemma 6.3.1. *Let G be given by a finite pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \mid \mathcal{R}\rangle$, where $\mathcal{P}$ satisfies trivial-interleaving, and* **RSym** *succeeds on $\mathcal{P}$. Let $t_1, t_2 \in X^\sigma$ be distinct. Then $t_1 \neq_G t_2$.*

Proof. Suppose for a contradiction that $t_1 =_G t_2$. If $(t_1, t_2^\sigma) \in D(P)$, then by uniqueness of inverses in P and Theorem 2.3.11 (iv) we have $[t_1, t_2^\sigma] \neq 1$, hence $[t_1, t_2^\sigma] = t \in X^\sigma$. Since **RSym** succeeds on $\mathcal{P}$, by Proposition 2.6.12 we have $[t_1, t_2^\sigma] = t \neq_G 1$, a contradiction. Hence $(t_1, t_2^\sigma) \notin D(P)$. Similarly, $(t_2^\sigma, t_1) \notin D(P)$. Since $\mathcal{P}$ satisfies trivial-interleaving (see

Definition 2.3.23), by [34, Proposition 6.10] there exists a diagram $\Gamma \in \mathcal{D}$ with boundary word $t_1 t_2^\sigma$. But by [34, Lemma 6.11] $\mathbf{RSym}$ does not succeed on Γ, so $\mathbf{RSym}$ does not succeed on $\mathcal{P}$, a contradiction. We conclude that $t_1 \neq_G t_2$. ∎

From Lemma 6.3.1 it follows that all conjugating elements in G between distinct $t_1, t_2 \in X^\sigma$ are not equal to 1 in P. Also, if t_1, t_2 are conjugate in $U(P)$, then by Theorem 3.2.11 there exists $c \in X^\sigma$ such that $ct_2c^\sigma t_1^\sigma =_{U(P)} 1$.

Definition 6.3.2. The array `LettersArray` is indexed by $t_1, t_2 \in X^\sigma$ with $t_1 \neq t_2$. We set `LettersArray`$(t_1, t_2) := c$ if there exists $c \in X^\sigma$ such that $ct_2c^\sigma t_1^\sigma =_G 1$. If there is no such c, then we set `LettersArray`$(t_1, t_2) := 0$.

Recall Definition 2.5.14 of the 1-skeleton of a coloured diagram and Definition 3.1.6 of a layer.

Definition 6.3.3. The procedure $\mathbf{ConjLetters}$ constructs a directed labelled simple graph D, with the following properties. The vertex-set of D is X^σ. Let $t_1, t_2 \in X^\sigma$ be distinct. There is an arc $e = (t_1, t_2)$ if there exists $c \in X^*$ such that $ct_2c^\sigma t_1^\sigma =_G 1$, and at least one of the following two conditions holds.

1. $c \in X^\sigma$.

2. There exists a minimal coloured conjugacy diagram Γ_A for t_1 and t_2 such that c labels a path $p \in (\Gamma_A)^1$ with endpoints lying on the opposite boundaries of Γ_A, and Γ_A is a single layer.

The label of e is any $d \in X^*$ such that $dt_2d^\sigma t_1^\sigma =_G 1$ (note that we may have $d \neq c$).

Proposition 6.3.4. *The components of D are conjugacy classes in G of single letters.*

Proof. Let t_1 and t_2 be distinct elements of X^σ that are conjugate in G. If there exists $c \in X^\sigma$ such that $ct_2c^\sigma t_1^\sigma =_G 1$, then by Definition 6.3.3 there is an arc (t_1, t_2) in D. So we can assume that no such c exists. Then by Theorem 3.2.11 t_1 and t_2 are not $U(P)$-conjugate. Moreover, as $\mathbf{RSym}$ succeeds on $\mathcal{P}$, by Proposition 2.6.12 t_1 and t_2 are non-trivial in G. Therefore, by Proposition 3.1.5 there exists a minimal conjugacy diagram (see Definition 3.1.4) Γ_A for t_1 and t_2. By Lemma 3.1.7 Γ_A is a face-disjoint union of finitely many layers, and as Γ_A is a minimal conjugacy diagram, the same holds for each of its layers. Hence there is path in D with endpoints t_1 and t_2.

One the other hand, if distinct $t_1, t_2 \in X^\sigma$ are in the same connected component of D, then by Definition 6.3.3 t_1 and t_2 are conjugate in G. ∎

Recall Definition 3.3.1 of a decomposable annular diagram.

Lemma 6.3.5. *Let $\Gamma_A \in \mathcal{S} \setminus \mathcal{T}$. Assume that both of the following statements hold:*

- *Γ_A is a single layer.*

- *The boundary words of Γ_A are not single letters G-conjugate by any single letter.*

Then Γ_A is decomposable and Γ_A satisfies Statements (i)-(ii) of Definition 3.3.2. Let Γ be the core of Γ_A and let B be a boundary red blob of Γ_A. Then for the label t of the loop of B, there exists a pair $(c, d) \in (X^\sigma)^2$ such that

1. *cd labels $\partial(\Gamma) \cap \partial(B)$;*
2. *$(c, d) \in D(P)$; $c \neq d^\sigma$;*
3. *there exists $\alpha \in X^\sigma$ such that $\alpha[cd]\alpha^\sigma t^\sigma =_{U(P)} 1$.*

Proof. By the assumptions and Definition 3.3.2 either $\Gamma_A \in \mathcal{T}$, or Γ_A is decomposable and Γ_A satisfies Statements (i)-(ii) of Definition 3.3.2. Hence as $\Gamma_A \in \mathcal{S} \setminus \mathcal{T}$, the latter holds for Γ_A, so we can let Γ and B be the core and a boundary red blob of Γ_A respectively. Let cd be a label of $\partial(\Gamma) \cap \partial(B)$. Then by Definition 3.3.2 we have $(c, d) \in D(P)$; $c \neq d^\sigma$; and $t \neq_P [cd]$: so by Theorem 3.2.11 and Lemma 6.3.1, t is $U(P)$-conjugate to $[cd]$ by some $\alpha \in X^\sigma$. The lemma follows. ∎

Notation 6.3.6. Let D be the graph from Definition 6.3.3. We denote by $e^c_{(t_1,t_2)}$ the arc in D with initial vertex t_1 and terminal vertex t_2, labelled by c.

Recall the lists $\mathcal{B}$, $\mathcal{L}_{23}$ and the record `Rels` from Section 6.2.

Procedure 6.3.7. ConjLetters($\mathcal{B}, \mathcal{L}_{23}$, `Rels`):

Step 1 Let D be the null graph on X^σ vertices, and construct `LettersArray`.

Step 2 For $t_1, t_2 \in X^\sigma$ with $t_1 \neq t_2$ do: if there is not an arc in D between t_1 and t_2 then

(i) If `LettersArray`$(t_1, t_2) = \gamma$, then add arcs $e^\gamma_{(t_1,t_2)}$ and $e^{\gamma^\sigma}_{(t_2,t_1)}$ to D and go to the beginning of Step 2.

(ii) `conj`$, \gamma :=$ **ConjInT**$(t_1, t_2, \mathcal{B}, \mathcal{L}_{23}$, `Rels`$)$ (see Algorithm 6.2.13). If `conj` then add arcs $e^\gamma_{(t_1,t_2)}$ and $e^{\gamma^\sigma}_{(t_2,t_1)}$ to D and go to the beginning of Step 2.

(iii) For $c_1, d_1 \in (X^\sigma)^2$ with $(c_1, d_1) \in D(P)$ and $c_1 \neq d_1^\sigma$ do:

If there exists $\alpha \in X^\sigma$ such that $\alpha[c_1d_1]\alpha^\sigma t_1^\sigma =_{U(P)} 1$ then

`conj`$, \gamma :=$ **ConjInT**$(c_1d_1, t_2, \mathcal{B}, \mathcal{L}_{23}$, `Rels`$)$. If `conj` then add arcs $e^{\alpha\gamma}_{(t_1,t_2)}$ and $e^{(\alpha\gamma)^\sigma}_{(t_2,t_1)}$ to D and go to the beginning of Step 2.

(a) For $(c_2, d_2) \in (X^\sigma)^2$ with $(c_2, d_2) \in D(P)$ and $c_2 \neq d_2^\sigma$ do:

If there exists $\beta \in X^\sigma$ such that $\beta[c_2d_2]\beta^\sigma t_2^\sigma =_{U(P)} 1$ then

`conj`$, \gamma :=$ **ConjInT**$(t_1, c_2d_2, \mathcal{B}, \mathcal{L}_{23}$, `Rels`$)$. If `conj` then add arcs $e^{\gamma\beta^\sigma}_{(t_1,t_2)}$ and $e^{(\gamma\beta^\sigma)^\sigma}_{(t_2,t_1)}$ to D and go to the beginning of Step 2.

`conj`$, \gamma :=$ **ConjInT**$(c_1d_1, c_2d_2, \mathcal{B}, \mathcal{L}_{23}$, `Rels`$)$. If `conj` then add arcs $e^{\alpha\gamma\beta^\sigma}_{(t_1,t_2)}$ and $e^{(\alpha\gamma\beta^\sigma)^\sigma}_{(t_2,t_1)}$ to D and go to the beginning of Step 2.

Step 3 Return `LettersArray`, D.

Proposition 6.3.8. *The graph D from Definition 6.3.3 can be constructed by* $\mathbf{ConjLetters}$ *in time $O(r^4|\mathcal{R}|^2|X|^9 + r^2|\mathcal{R}||X|^{11})$, where $r := \max\{|R| : R \in \mathcal{R}\}$ is the length of the longest green relator.*

Proof. To prove that $\mathbf{ConjLetters}$ constructs D, we show that given $t_1, t_2 \in X^\sigma$ with $t_1 \neq t_2$, if there is a c that satisfies Definition 6.3.3, then $\mathbf{ConjLetters}$ adds an arc $e = (t_1, t_2)$ with label d according to Definition 6.3.3. So assume that such a c exists. If $c \in X^\sigma$, then $\mathtt{LettersArray}(t_1, t_2) = d$ for some $d \in X^\sigma$, hence in Step 2 (i) $\mathbf{ConjLetters}$ adds a correct arc to D.

Now suppose that c satisfies Condition 2 of Definition 6.3.3. By the previous paragraph we can assume that t_1 and t_2 are not G-conjugate by any single letter. Then by Theorem 3.2.11 t_1 and t_2 are not $U(P)$-conjugate. By Assumption 2.3.15 (which states that no $R \in \mathcal{R}$ satisfies $|R| \in \{1, 2\}$) t_1 and t_2 are cyclically $\mathcal{P}$-reduced, and as $\mathbf{RSym}$ succeeds on $\mathcal{P}$, by Proposition 2.6.12 t_1 and t_2 are non-trivial in G. Hence as Γ_A is a single layer, by Theorem 3.3.3 there is a minimal conjugacy diagram $\Delta_A \in \mathcal{S}$ (see Definition 3.3.2) with boundary words t_1 and t_2, and consisting of a single layer. Since t_1 and t_2 are not $U(P)$-conjugate, by Theorem 3.1.2 $\mathbf{CArea}(\Delta_A) \geq (1, 0)$; and by Definition 3.1.4 Δ_A has minimal coloured area. Assume first that $\Delta_A \in \mathcal{T}$. Then by Theorem 6.2.15 in Step 2 (ii) $\mathbf{ConjLetters}$ adds a correct arc to D.

Assume instead that $\Delta_A \notin \mathcal{T}$. Then by Lemma 6.3.5 Δ_A is decomposable, Δ_A satisfies Statements (i)-(ii) of Definition 3.3.2, and for the label t of the loop of each boundary red blob of Δ_A there exists a pair $(c, d) \in (X^\sigma)^2$ that satisfies Statements 1-3 of the lemma. Since Δ_A has minimal coloured area, the same is true for its core Γ. Also, $\mathbf{CArea}(\Delta_A) \geq (1, 0)$ implies $\mathbf{CArea}(\Gamma) \geq (1, 0)$, hence by Theorem 6.2.15 in Step 2 (iii) $\mathbf{ConjLetters}$ adds again a correct arc to D. Thus, we showed that $\mathbf{ConjLetters}$ constructs D.

Note that to prove the final statement, it suffices to analyse the time complexity of Step 2 (iii). By Theorem 6.2.15 $\mathbf{ConjInT}$ runs in time $O(r^4|\mathcal{R}|^2|X|^3 + r^2|\mathcal{R}||X|^5)$ when called by $\mathbf{ConjLetters}$. Hence as $\mathbf{ConjLetters}$ may need to run through all pairs $(c_1, d_1), (c_2, d_2) \in (X^\sigma)^2$, Step 2 (iii) of $\mathbf{ConjLetters}$ takes time $O(r^4|\mathcal{R}|^2|X|^7 + r^2|\mathcal{R}||X|^9)$. Now $\mathbf{Conj}$ $\mathbf{Letters}$ runs Step 2 at most $O(|X|^2)$ times, so the overall complexity is as stated. ∎

Let $w_1, w_2 \in X^*$ be cyclically $\mathcal{P}$-reduced and non-trivial in G. The next algorithm checks if there exists a minimal conjugacy diagram for w_1 and w_2 that contains a loop.

Algorithm 6.3.9. $\mathbf{LoopyDiagram}(w_1, w_2, \mathtt{LettersArray}, D, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$:
// input: D – the graph from Definition 6.3.3.

Step 1 For $1 \leq i \leq 2$ and $t_i \in X^\sigma$ do:

(i) If $\mathbf{RSymSolveSimpler}(w_i t_i^\sigma) = \eta$ (see Algorithm 6.1.1) then let $\gamma_i := \eta$ be the empty word and go to the beginning of Step 1.

(ii) $\mathtt{conj}, \gamma := \mathbf{ConjInT}(w_i, t_i, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$. If $\mathtt{conj}$ then $\gamma_i := \gamma$ and go to the beginning of Step 1.

(iii) For $c, d \in (X^\sigma)^2$ with $(c, d) \in D(P)$ and $c \neq d^\sigma$ do:

If there exists $\alpha \in X^\sigma$ such that $\alpha[cd]\alpha^\sigma t_i^\sigma =_{U(P)} 1$ then

$\texttt{conj}, \gamma := \mathbf{ConjInT}(w_i, cd, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$. If $\texttt{conj}$ then $\gamma_i := \gamma\alpha^\sigma$ and go to the beginning of Step 1.

Step 2 If t_1 and t_2 are in the same component of D, then let γ_3 be the label of p, a shortest path in D from t_1 to t_2. If γ_1 and γ_2 are defined, then return $\texttt{true}$, $\gamma_1\gamma_3\gamma_2^\sigma$.

Step 3 Return $\texttt{false}$.

Theorem 6.3.10. *Suppose that* $\texttt{LettersArray}$ *and the graph* D *from Definition 6.3.3 have been constructed, and that either* w_1 *and* w_2 *are* G*-conjugate and are both equal to single letters in* G*, or there exists a minimal conjugacy diagram* Γ_A *for* w_1 *and* w_2 *that contains a loop. Then* $\mathbf{LoopyDiagram}(w_1, w_2, \texttt{LettersArray}, D, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$ *returns* $\texttt{true}$*; and if it returns* $\texttt{true}$*, then it returns an* $x \in X^*$ *such that* $w_1 =_G w_2^{x^{-1}}$.

Let $r := \max\{|R| : R \in \mathcal{R}\}$. *The running time of* $\mathbf{LoopyDiagram}(w_1, w_2, \texttt{Letters Array}, D, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$ *is* $O((r^4|\mathcal{R}|^2|X|^6 + r^2|\mathcal{R}||X|^8) \cdot (\max\{|w_1|, |w_2|\}))$.

Proof. Let $i \in \{1, 2\}$. Assume first that $w_i =_G t$ for some $t \in X^\sigma$. Then by Proposition 6.1.2 in Step 1 (i) of $\mathbf{LoopyDiagram}$ we have $\mathbf{RSymSolveSimpler}(w_i t^\sigma) = \eta$.

Assume instead that $w_i \neq_G t$ for all $t \in X^\sigma$, so that Γ_A exists. Since Γ_A contains a loop, by Lemma 3.1.7 Γ_A is a face-disjoint union of finitely many layers. Hence for each $i \in \{1, 2\}$, there exists $t_i \in X^\sigma$ and a layer Γ_i of Γ_A with boundary words w_i and t_i. Since Γ_A is a minimal conjugacy diagram, the same holds for Γ_i. Suppose that w_i and t_i are $U(P)$-conjugate. Then by Theorem 3.2.11 $|w_i| = 1$, contradicting our assumption, so they are not $U(P)$-conjugate. Similarly as in the proof of Proposition 6.3.8 we can use Assumption 2.3.15 and the fact that $\mathbf{RSym}$ succeeds on $\mathcal{P}$ to deduce that t_i is cyclically $\mathcal{P}$-reduced and non-trivial in G. Hence as by the assumption the same holds for w_i and Γ_i is a single layer, by Theorem 3.3.3 there exists a minimal conjugacy diagram $\Delta_A \in \mathcal{S}$ with boundary words w_i and t_i, and consisting of a single layer. Furthermore, by Theorem 3.1.2 we have $\mathbf{CArea}(\Delta_A) \geq (1, 0)$ since w_i and t_i are not $U(P)$-conjugate; and by Definition 3.1.4 Δ_A has minimal coloured area.

Assume first that $\Delta_A \in \mathcal{T}$. Then by Theorem 6.2.15 in Step 1 (ii) of $\mathbf{LoopyDiagram}$, for some $t_i' \in X^\sigma$ the algorithm $\mathbf{ConjInT}(w_i, t_i', \mathcal{B}, \mathcal{L}_{23}, \texttt{Rels})$ returns $\texttt{true}$; and if it returns $\texttt{true}$, then it returns $\gamma_i \in X^*$ such that $w_i =_G t_i'^{\gamma_i^{-1}}$.

Assume instead that $\Delta_A \notin \mathcal{T}$. Then Δ_A satisfies assumptions of Lemma 6.3.5. Hence as $|w_i| \geq 2$, by Lemma 6.3.5 Δ_A is decomposable with precisely one boundary red blob B, Δ_A satisfies Statements (i)-(ii) of Definition 3.3.2, and for the label t of the loop of B there exists a pair $(c, d) \in (X^\sigma)^2$ that satisfies Statements 1-3 of the lemma. Since Δ_A has minimal coloured area, the same holds for its core Γ. Now $\mathbf{CArea}(\Delta_A) \geq (1, 0)$ implies $\mathbf{CArea}(\Gamma) \geq (1, 0)$, hence by Theorem 6.2.15 in Step 1 (iii) for some $t_i' \in X^\sigma$, $\mathbf{LoopyDiagram}$ returns $\texttt{true}$; and if it returns $\texttt{true}$, then it returns $\gamma_i \in X^*$ such that $w_i =_G t_i'^{\gamma_i^{-1}}$.

Thus, for each $i \in \{1,2\}$: in Step 1 **LoopyDiagram** finds $t_i \in X^\sigma$ that is G-conjugate to w_i; and if any of the sub-routines returns `true`, then (for some $t_i \in X^\sigma$) it returns $\gamma_i \in X^*$ such that $w_i =_G t_i^{\gamma_i^{-1}}$. By Proposition 6.3.4 t_1 and t_2 are G-conjugate if and only if t_1 and t_2 are connected in D. Hence **LoopyDiagram** returns `true`. Furthermore, if **LoopyDiagram** returns `true`, then it also returns an $x \in X^*$ with the desired properties. Hence the first statement follows.

Now in Step 3 of **LoopyDiagram** we use Dijkstra's algorithm (see [20]) to find the path p. This algorithm runs in time $O(|V|\log|V| + |E|)$, on a graph with $|V|$ vertices and $|E|$ edges, so $O(|X|^2)$ in our case. Then note that p does not contain repeating vertices, so we can define γ_3 in time $O(|X|)$. Therefore, to prove the final statement, we need to analyse the time complexity of Step 1 (iii). By Theorem 6.2.15 for each $i \in \{1,2\}$, **ConjInT** runs in time $O((r^4|\mathcal{R}|^2|X|^3 + r^2|\mathcal{R}||X|^5) \cdot (|w_i|))$ when called by **LoopyDiagram**. So as **LoopyDiagram** may need to run through all $(c,d) \in (X^\sigma)^2$, Step 1 (iii) takes time $O((r^4|\mathcal{R}|^2|X|^5 + r^2|\mathcal{R}||X|^7) \cdot (|w_i|))$. Hence the theorem follows as Step 1 (iii) is computed at most $O(|X|)$ times. ∎

6.4 The IsConjugate algorithm and the proof of Theorem 1.0.3

In this section we present the conjugacy problem solver, **IsConjugate**, and prove Theorem 1.0.3. Before running **IsConjugate**, we precompute data via a procedure **ConjPreprocess**.

ConjPreprocess takes as input a sound pregroup presentation $\mathcal{P}$ (see Definition 2.6.14), and a list `Rec` that contains data computed by Steps 1-4 of **RSymVerify** (see [34, Section 7.6]) on input $\mathcal{P}$: intermult table, roots of green relators, locations, places and the vertex graph.

Procedure 6.4.1. ConjPreprocess$(\mathcal{P}, \texttt{Rec})$:

Step 1 If not **IsProper**$(\mathcal{P})$ (see Definition 6.0.1) then return `error`.
Step 2 Construct $\mathcal{B}$ and $\mathcal{L}_{23}$ (see Section 6.2).
Step 3 `LettersArray`, $D :=$ **ConjLetters** (see Definition 6.3.3).
Step 4 Return $\mathcal{B}$, $\mathcal{L}_{23}$, `LettersArray`, D.

Let $w_1, w_2 \in X^*$ be cyclically $\mathcal{P}$-reduced and such that $|w_1| = |w_2|$. In the description below of **IsConjugate**, we write **KMP**(w_1, w_2^2) when running the Knuth-Morris-Pratt algorithm on input w_1 and w_2^2. From [2, Section 9.1], **KMP**(w_1, w_2^2) returns `true` if and only if w_1 is a sub-word of w_2^2, hence **KMP**(w_1, w_2^2) returns `true` if and only if w_1 is a cyclic conjugate of w_2. If so, then **KMP**(w_1, w_2) also returns the corresponding inverse x of a prefix of w_2 such that $w_1 =_G w_2^{x^{-1}}$. Moreover, from [2, Section 9.1], **KMP**(w_1, w_2^2) runs in time $O(|w_1| + |w_2|)$.

Algorithm 6.4.2. IsConjugate$(w_1, w_2, \mathcal{B}, \mathcal{L}_{23}, \texttt{Rec})$:
// Input: w_1, w_2 – input words tested for conjugacy.

Step 1 Cyclically $\mathcal{P}$-reduce w_1 and w_2 via **RSymSolveSimpler** (see Algorithm 6.1.1), and let the resulting words be r_1 and r_2 respectively. If ($r_1 = \eta$ and $r_2 \neq \eta$) or if ($r_2 = \eta$ and $r_1 \neq \eta$) then return `false`. If $r_1 = \eta = r_2$ then return `true`, η, η, $x := \eta$.
Step 2 If $|r_1| = |r_2|$ then let `conj`, $x :=$ **KMP**(r_1, r_2^2). If `conj` then return `true`, r_1, r_2, x.

Step 3 $\mathtt{conj}, x$:= **LoopyDiagram**$(r_1, r_2, \mathtt{LettersArray}, D, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ (see Algorithm 6.3.9). If `conj` then return `true`, r_1, r_2, x.

Step 4 $\mathtt{conj}, x$:= **ConjInT**$(r_1, r_2, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ (see Algorithm 6.2.13). If `conj` then return `true`, r_1, r_2, x.

Step 5 Return `false`.

We are ready to prove Theorem 1.0.3. For time complexity of **IsConjugate** we treat $|X|$, $|\mathcal{R}|$ and r as constants.

Proof of Theorem 1.0.3. By Proposition 3.2.13, Lemma 6.2.14 and Proposition 6.3.8 **ConjPreprocess** runs in time $O(r^4|\mathcal{R}|^2|X|^9 + r^2|\mathcal{R}||X|^{11})$. This constitutes the first claim about construction of **IsConjugate**.

By Proposition 6.1.2 for each $i \in \{1, 2\}$, r_i := **RSymSolveSimpler**(w_i) is a cyclically $\mathcal{P}$-reduced G-conjugate of w_i. So to prove the correctness of **IsConjugate**, we need to show that if w_1 and w_2 are G-conjugate, then **IsConjugate** returns `true`, else it returns `false`; and if **IsConjugate** returns `true`, then it returns an $x \in X^*$ such that $r_2 =_G r_1^x$.

Assume first that w_1 and w_2 are not G-conjugate. Then by Proposition 6.1.2 we cannot have $r_1 = \eta = r_2$. Furthermore, **KMP**(r_1, r_2^2) returns `false` if $|r_1| = |r_2|$; by Theorem 6.3.10 **LoopyDiagram**$(r_1, r_2, \mathtt{LettersArray}, D, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ returns `false`; and by Theorem 6.2.15 **ConjInT**$(r_1, r_2, L, \mathcal{B}, \mathcal{L}_{23}, \mathtt{Rels})$ returns `false`. We conclude that **IsConjugate** returns `false`.

Now assume that w_1 and w_2 are G-conjugate. Then we cannot we have $r_1 = \eta$ and $r_2 \neq \eta$ or $r_2 = \eta$ and $r_1 \neq \eta$. If $r_1 = \eta = r_2$, then **IsConjugate** returns `true` in Step 1. So assume that $r_1 \neq \eta \neq r_2$. Then by Proposition 6.1.2 we have $r_1 \neq_G 1 \neq_G r_2$. If r_1 and r_2 are cyclic conjugate or are both equal to single letters in G, then by [2, Section 9.1] and Theorem 6.3.10 **IsConjugate** returns `true` in Steps 2 or 3. So assume that they are neither cyclic conjugate nor both equal to the single letters in G. Then by Corollary 3.2.12 they are not $U(P)$-conjugate. Hence by Proposition 3.1.5 there exists a minimal conjugacy diagram Γ_A for r_1 and r_2, and by Definition 3.1.4 Γ_A has minimal coloured area. If Γ_A contains a loop, then by Theorem 6.3.10 in Step 3 **IsConjugate** returns `true`. Hence assume that Γ_A is loop-free. Then by Theorem 3.2.2 $\Gamma_A \in \mathcal{T}$. Also, as r_1 and r_2 are not $U(P)$-conjugate, by Theorem 3.1.2 we have **CArea**$(\Gamma_A) \geq (1, 0)$. Therefore, by Theorem 6.2.15 in Step 4 **IsConjugate** returns `true`.

Finally, assume that **IsConjugate** returns `true`. If $r_1 = \eta = r_2$, or if r_1 and r_2 are cyclic conjugates, then in Steps 1 or 2 **IsConjugate** returns an $x \in X^*$ with $r_2 =_G r_1^x$. Otherwise, by Theorems 6.2.15 and 6.3.10 in Steps 3 or 4 **IsConjugate** returns an $x \in X^*$ with the desired properties.

The final complexity claim follows from Proposition 6.1.2 and Theorems 6.2.15 and 6.3.10, since for each $i \in \{1, 2\}$ we have $|r_i| \leq |w_i|$. ∎

Chapter 7

Experiments

We implemented **IsConjugate**, as **IsConjugate**, in the computer algebra system `MAGMA` (see [6]). We used the code of the implementation **IsHyperbolic** of **RSymVerify** (see [34, Procedure 7.19]) to produce the pregroup multiplication table; to produce the lists `Rec`, $\mathcal{B}$ and $\mathcal{L}_{23}$ from Section 6.4; to find sound pregroup presentations (see Definition 2.6.14); and then modify it to construct **RSymSolveSimpler** (see Algorithm 6.1.1). The remaining sub-routines of **IsConjugate** are computed by new code.

7.1 Accuracy

In this section we describe accuracy of our implementation. We want to demonstrate that on input a valid (see Definition 2.6.14) pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \cup \mathcal{R} \rangle$, **IsConjugate** returns `true` on input $w_1 \in X^*$ and $w_2 \in X^*$ if and only if w_1 and w_2 are conjugate in the group defined by $\mathcal{P}$. If our implementation of **IsConjugate** returns `true`, then it should also return a conjugating word, hence proving that the input words are indeed conjugate. Hence let us test our implementation on input words that are conjugate, and check if it returns `true` and a conjugating word.

Our aim was to come up with cases that tests the correctness of all sub-routines of **IsConjugate** that perform conjugacy tests on input words. We constructed the examples by hand: in all of them the input words are already cyclically $\mathcal{P}$-reduced, so they are returned by **IsConjugate**. Also, in all examples the group $F(X^\sigma)$ is a free group, so we use x^{-1} to denote the inverse of $x \in X$ in $F(X^\sigma)$ instead of x^σ to be consistent with inputs and outputs of the implementation. Finally, to verify that the input words w_1 and w_2 are conjugate, we used Derek Holt's implementation (in `GAP`) of Marshall's (see [43]) conjugacy problem solver for hyperbolic groups, and to verify that the word c returned by **IsConjugate** is a conjugating word, we used the `MAGMA`'s KBMAG package (by constructing an automatic structure of the input group and verifying that the product $cw_2c^{-1}w_1^{-1}$ is the identity word).

Test 1: Our first example is the small cancellation group $C'(1/7) - T(4)$ defined as

$$\mathcal{P} = \langle a, b, c, d \mid \{\emptyset\} \mid \{R := [a, b][c, d]\} \rangle.$$

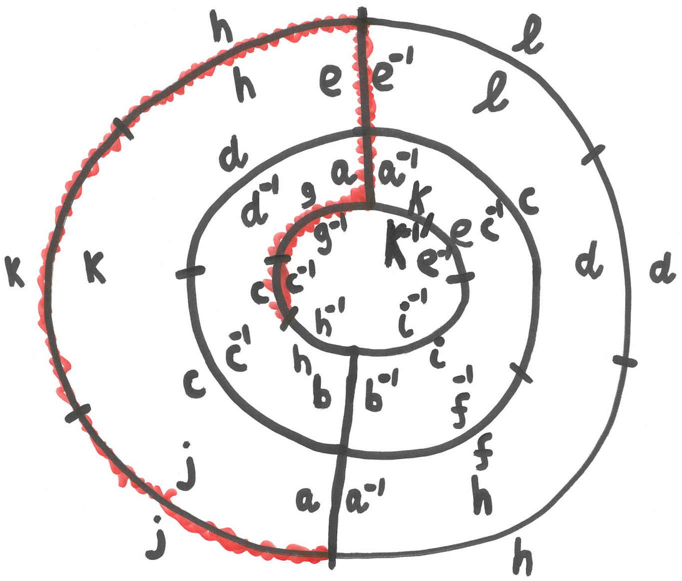

Figure 7.1: The annular diagram Γ with the outer boundary labelled by w_1 and the inner boundary labelled by w_2, see Test 2. The red curve depicts the path labelled by the conjugating word $jkheagc$

.

We verified that $\mathcal{P}$ is valid. Let $w_1 = d^{-1}cdbc^{-1}d^{-1}$ and $w_2 = b^{-1}a^{-1}bbad^{-1}$. Checking by hand we found that there is an annular diagram Γ (see Definition 3.1.4) with boundary words w_1 and w_2 containing two internal green faces labelled by R and R^{-1} that share an edge with label a and have edges on both boundaries of Γ. As expected, **IsConjugate** returned `true`, where **BothBoundaries** (see Algorithm 6.2.7) found conjugating word $d^{-1}cda^{-1}b^{-1}ab$.

Test 2 : The next example is a presentation of the form

$$\begin{aligned}\mathcal{P} = \langle &a, b, c, d, e, f, g, h, i, j, k, l \mid \{\emptyset\} \mid \\ &\mathcal{R} := \{ajkhedc, e^{-1}ldha^{-1}fc, c^{-1}f^{-1}b^{-1}ieka^{-1}, agchbc^{-1}d^{-1}\}\rangle.\end{aligned}$$

The reason why $\mathcal{P}$ is defined on so many letters is to ensure that the steps have short length, making $\mathcal{P}$ sound. Let $w_1 = jkhldh$ and $w_2 = c^{-1}g^{-1}k^{-1}e^{-1}i^{-1}h^{-1}$. Then there is an annular diagram Γ for w_1 and w_2 with $\mathbf{CArea}(\Gamma) = (4, 0)$, where the set of labels of the four internal faces of Γ is $\mathcal{R}$, all internal faces share the same two interior vertices, and there are two pairs (D, D') of them such that D and D' share an edge and have edges on the opposite boundaries of Γ (see Figure 7.1). **IsConjugate** returned `true` on input w_1 and w_2, where **ConjByTwoLabels** (see Algorithm 6.2.8) found $jkheagc$.

In the remaining examples we use the set of red relators from Example 2.4.1, and subsequently add additional letters and green relators to define valid pregroup presentations.

Test 3: The first presentation formed in this way satisfies

1. $X = \{a, b, c, d, e, f, g, h, i, j, k, l, m\}$;

2. $V_P = \{abd^{-1}, bce^{-1}, db^{-1}a^{-1}, ec^{-1}b^{-1}, aef^{-1}, dcf^{-1}\}$;

3. $\mathcal{R} = \{amhdia^{-1}lbg, c^{-1}jbdm^{-1}a^{-1}ghk\}$.

We aimed to test the correctness of **StartWithRedBlob** (see Algorithm 6.2.12), so we chose $w_1 = bjbdhdi$ and $w_2 = fk^{-1}h^{-1}g^{-2}b^{-1}l^{-1}$. Then there is an annular diagram for w_1 and w_2 that satisfies Statement 2 (ii) of Theorem 5.6.1. **IsConjugate** returned `true` on input w_1 and w_2, and **StartWithRedBlob** returned $bd^{-1}lbg^2hkf^{-1}$.

Test 4: We took a presentation with

1. $X = \{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, r, s, t, u, v, x\}$;

2. $V_P = \{abd^{-1}, bce^{-1}, db^{-1}a^{-1}, ec^{-1}b^{-1}, aef^{-1}, dcf^{-1}\}$;

3. $\mathcal{R} = \{ghijkb^{-1}, c^{-1}lmnh^{-1}g^{-1}, ft^{-1}s^{-1}uvx, a^{-1}oprst\}$.

By taking $w_1 = ijklmn$ and $w_2 = r^{-1}p^{-1}o^{-1}x^{-1}v^{-1}u^{-1}$, both conjugate to the letter e, we were able to test the correctness of **LoopyDiagram** (see Algorithm 6.3.9). As desired, **IsConjugate** returned `true`, and **LoopyDiagram** returned $ijkb^{-1}a^{-1}\ opr$. We also took w_1^5 and w_2^5 for testing **ConjByThreeLabels** (see Algorithm 6.2.9). In this case **IsConjugate** returned `true` and **ConjByThreeLabels** returned $ijkct^{-1}s^{-1}uvxopr$.

Test 5: Our final presentation is for testing **ComplicatedRedBlobs** (see Algorithm 6.2.10). We were unable to find diagrams containing bad or complicated red blobs (see Definitions 4.1.1 and 4.6.1). However, we were able to test **ComplicatedRedBlobs** by finding an annular diagram Γ containing two internal green faces that share an edge with the same red triangle, and have edges on the opposite boundaries of Γ. (We also let **ConjInT** (see Algorithm 6.2.13) skip running **ConjByTwoLabels** and **ConjByThreeLabels** on input w_1 and w_2, as otherwise they would find solution before running **ComplicatedRedBlobs**.) We took

1. $X = \{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, r, s, t, u, v, w, x, y, z, q\}$;

2. $V_P = \{abd^{-1}, bce^{-1}, db^{-1}a^{-1}, ec^{-1}b^{-1}, aef^{-1}, dcf^{-1}\}$;

3. $\mathcal{R} = \{fs^{-1}y^{-1}tqwz^{-1}r^{-1}, ysoprhgijk, d^{-1}rzlmng^{-1}h^{-1}uvx\}$;

4. $w_1 = wlmnijktq$ and $w_2 = x^{-1}v^{-1}u^{-1}r^{-1}p^{-1}o^{-1}c^{-1}$.

Then **IsConjugate** returned `true` on input w_1 and w_2, and **ComplicatedRedBlobs** returned $wz^{-1}r^{-1}fc^{-1}$. We also took w_1 and a cyclic permutation $mnijktqwl$ of w_1 to test the correctness of our implementation of the Knuth-Morris-Pratt algorithm; and two cyclic permutations $lmnijktqw$ and $ijktqwlmn$ of w_1 to test the correctness of **CyclicConj** (see Algorithm 6.2.3). In both cases the algorithms returned correctly `true` and correct conjugating words $q^{-1}t^{-1}k^{-1}j^{-1}i^{-1}n^{-1}m^{-1}$ and $w^{-1}q^{-1}t^{-1}k^{-1}j^{-1}i^{-1}$ respectively (and with run time only 0.01 seconds).

Thus, we found examples for testing the correctness of all sub-routines of **IsConjugate** that perform conjugacy tests on input words; and in addition, in all Tests 1-5, **IsConjugate**

Table 7.1: Run times of Examples (a)-(e) from Section 7.1.

Example	**ConjPreprocess**	**IsConjugate**
Test 1	0.070s	0.010s
Test 2	0.72s	0.010s
Test 3	7.26s	2.03s
Test 4 (i)	3.83s	0.08s
Test 4 (ii)	3.83s	0.96s
Test 5	4.55s	0.29s

returned `true` on all pairs of cyclic conjugates of the input words, and conjugating words were found by different sub-routines. Finally, note that in all the examples above, by taking powers of the input words one can find conjugate words with arbitrary length.

7.2 Efficiency

In this section we report the run times of our experiments (in seconds). Tests 1-5 from Section 7.1 are described in Table 12.1. Test 4 is split into Parts (i)-(ii), where in Part (i) **IsConjugate** run on input w_1 and w_2, while in Part (ii) it run on w_1^5 and w_2^5. To analyse the worst-case complexity of **IsConjugate**, we used the pregroup presentations from Tests 2-5 of Section 7.1, and chose the input words as follows:

(*) For $1 \leq i \leq 100$, we run **IsConjugate** on input w_1^* and w_2^*, where w_1^* and w_2^* are random freely cyclically reduced words of lengths $10i$ and $10i - 5$ respectively.

We were interested in the growth of run times of **IsConjugate** against the increase of $|w_1^* + w_2^*|$. This relationship is depicted in Figure 7.2 (where Cases 1-4 correspond to Tests 2-5 respectively). There is a large difference between the run times in Case 2 and the other cases, however, even in Case 2 the plot suggests a low-degree polynomial time growth, quite possibly quadratic.

To analyse this further, we assumed that the growth is quadratic, and of the form $Ax^2 + Bx$. To estimate the constants A and B, we picked two points $(x_j, y_j)_{j=1}^2$: one for a low value ($w_1^* = 50$; $w_2^* = 45$) of $|w_1^* + w_2^*|$ and one for a middle value ($w_1^* = 500$; $w_2^* = 495$) of $|w_1^* + w_2^*|$, and solved the system of equations $y_j = Ax_j^2 + Bx_j$ for $1 \leq j \leq 2$. As we can see in Figure 7.3, the estimated quadratic function seems to bound the run times well.

Similarly, we tried to estimate the growth with a cubic function $Ax^3 + Bx^2 + Cx$ by choosing three points $(x_j, y_j)_{j=1}^3$ and solving the system of equations $y_j = Ax_j^3 + Bx_j^2 + Cx_j$ for $1 \leq j \leq 3$. However, all estimated values A_0 of A had $A_0 < 0$, which suggests that the growth is rather quadratic than cubic. This agrees with Theorem 1.0.3.

Figure 7.2: Growth of run times of $\mathbf{IsConjugate}$ against $|w_1^* + w_2^*|$.

(a) Case 1

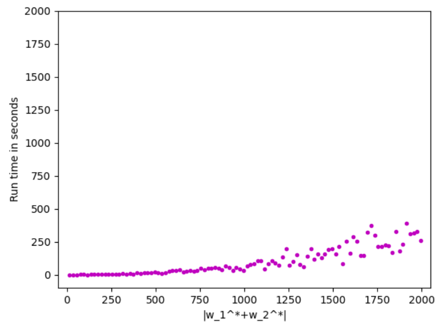

(b) Case 2

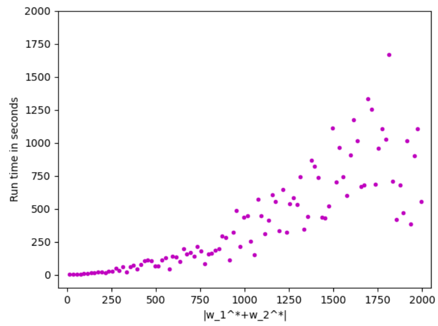

(c) Case 3

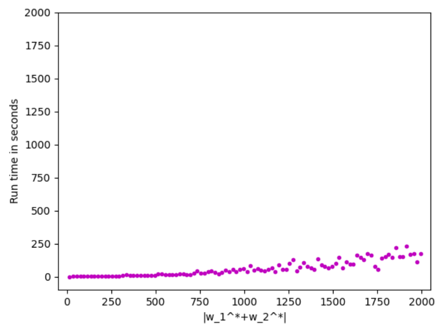

(d) Case 4

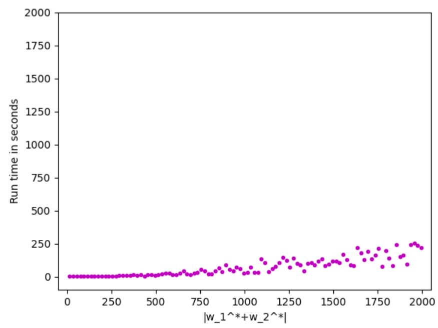

Figure 7.3: Growth of run times of $\mathbf{IsConjugate}$ (purple dots) compared to the estimated quadratic growth (blue line).

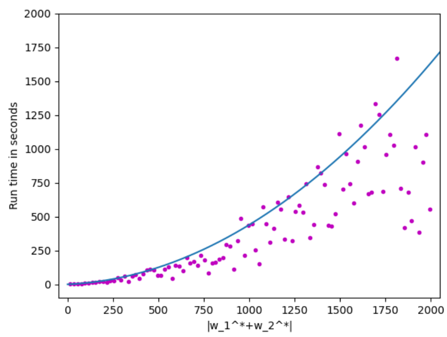

7.3 Examples where IsConjugate works and existing methods do not

Let $G = \langle a, b \mid a^\ell, b^m, (ab)^n \rangle$ with $2 \leq \ell \leq m \leq n$ and $1/\ell + 1/m + 1/n < 1$. By [34, Proposition 9.5], G can be defined by a sound pregroup presentation (see Definition 2.6.14). In particular, consider a group G defined by a pregroup presentation

$$\mathcal{P} = \langle a, b \,|\, V_P = \{b^3\} \,|\, \mathcal{R} := \{(ab)^7\} \rangle,$$

where the relation $a = a^\sigma$ is required in the pregroup. Then $\mathcal{P}$ is valid, so **IsConjugate** solves the conjugacy problem in G.

Let $w_1 = (ab)^5$ and $w_1 = (ab)^{-2}$. Then as $(ab)^7 \in \mathcal{R}$, we have $w_1 =_G w_2$ and $w_1^7 =_G 1 =_G w_2^7$. **IsConjugate** returned correctly `true` on input w_1 and w_2. However, even on this simple example, Marshall's (see [43]) algorithm returned an error as it does not solve the conjugacy problem for elements with finite order.

Finally, consider a group G defined by a pregroup presentation

$$\mathcal{P} = \langle a, b \,|\, \{\emptyset\} \,|\, \mathcal{R} := \{a^{\ell=100}, b^{m=153}, (ab)^{n=157}\} \rangle.$$

Then $\mathcal{P}$ is valid, and **IsConjugate** solves the conjugacy problem in G. However, in both `GAP` and `MAGMA`, the KBMAG algorithm failed to precompute an automatic structure for G (since too many computational steps were required to be carried out), hence all previous implementations of conjugacy problem solvers for hyperbolic groups cannot be constructed on input $\mathcal{P}$ – Wakefield's implementation (see [55, Chapters 5 & 6]) of the algorithm of Gersten and Short (see [27]) and Marshall's algorithm.

We expect the same outcome for all ℓ, m, n with $\ell \leq m \leq n$ and $\ell \geq 100, m \geq 153$ and $n \geq 157$, though the construction of **IsConjugate** performed by **ConjPreprocess** (see Procedure 6.4.1) might take a very long time for large values of ℓ, m, n.

Chapter 8

Improvements and generalizations of IsConjugate

The final chapter of Part 1 includes suggestions for making the construction of **IsConjugate** more efficient, and for its generalizations. Let G be given by a valid pregroup presentation $\mathcal{P} = \langle X^\sigma \mid V_P \cup \mathcal{R} \rangle$ (see Definition 2.6.14). Then it might be possible to simplify Definition 3.3.2 of the set $\mathcal{S}$ as follows.

Definition 8.0.1. We define $\mathcal{S}'$ to be the set of all coloured annular diagrams Γ_A over $\mathcal{P}$ each of which is a face-disjoint union of layers Γ with area at least 1 and such that: for non-boundary layers Γ, both boundaries of Γ have length 1, and one of the following 2 statements holds for Γ.

1. The boundary words of Γ are single letters G-conjugate by some single letter.
2. $\Gamma \in \mathcal{T}$.

For boundary layers Γ the same holds, but at least one of the boundaries of Γ is a boundary of Γ_A, so may have length greater than 1.

With this definition, the construction of **IsConjugate** takes time $O(r^4|\mathcal{R}|^2|X|^5 + r^2|\mathcal{R}||X|^7)$ instead of $O(r^4|\mathcal{R}|^2|X|^9 + r^2|\mathcal{R}||X|^{11})$.

The proof of showing an existence of a diagram in $\mathcal{S}'$ would be similar as for $\mathcal{S}$. Assume that that $w_1 \in X^*$ and $w_2 \in X^*$ are cyclically $\mathcal{P}$-reduced, non-trivial in G, and G-conjugate but not $U(P)$-conjugate; and that Γ_A^m is a minimal conjugacy diagram for w_1 and w_2 that contains a loop l. Let Γ be a layer of Γ_A^m with $l \subseteq \Gamma$, and let v be the endpoint of l. Then as in the proof of Theorem 3.3.5, we can show that either $\delta_G(v, \Gamma) = 2$ or $\delta_G(v, \Gamma) = 1$ and $\delta_R(v, \Gamma) = 4$, and that in the second case, the length 2 boundary of the boundary red blob of Γ containing l is non-trivial in P, and is not P-reduced – using this fact, we can retriangulate one more time at v, to reduce $\delta_R(v, \Gamma)$ to 3, and subsequently apply the minimality (see Definition 3.1.4) of Γ_A^m to deduce that Γ is a union of two layers satisfying one of the Conditions 1-2 of Definition 8.0.1. The problem with this approach is that reducing $\delta_R(v, \Gamma)$ from 4 to 3 creates an additional loop in the diagram, however, we believe one might still show that if Γ_A^m contains

n layers, then the final diagram $\Gamma_A \in \mathcal{S}'$ with boundary words w_1 and w_2 contains at most $3n$ layers, and is a smaller sibling of Γ_A^m (see Definition 3.1.13).

Another interesting question is the following open problem:

Question 1: Does there exist a quadratic time and practically implementable algorithm that solves the conjugacy problem for all hyperbolic groups defined by finite pregroup presentations that do not necessarily satisfy trivial-interleaving (see Definition 2.3.23)?

We presume that the input pregroup presentation still needs to have all the other properties of valid pregroup presentations.

Part II

A new method for showing hyperbolicity

Brief outline

Part II of this thesis is structured as follows. In Chapter 9 we describe a curvature distribution scheme $\mathbf{RSymVert}$. Descriptions of sub-routines of $\mathbf{VerifyHypVertex}$ are given in Chapter 10. In Chapter 11 we prove Theorem 1.0.4 stated in Chapter 1, which is the main result of Part II, as follows. In Section 11.1 we show that the success of $\mathbf{RSymVert}$ on a pregroup presentation $\mathcal{P}$ establishes an explicit linear bound on the Dehn function of $\mathcal{P}$, thus proving that the group defined by $\mathcal{P}$ is hyperbolic. In Section 11.2 we present $\mathbf{VerifyHypVertex}$, and show that if $\mathbf{VerifyHypVertex}$ returns `true`, then $\mathbf{RSymVert}$ succeeds. This proves the first statement of Theorem 1.0.4. In Section 11.3 we complete the proof of Theorem 1.0.4. In Chapter 12 we describe experiments with our implementation. Finally, Chapter 13 includes examples of future curvature distribution schemes that might be useful for showing hyperbolicity.

Throughout the whole Part II, let G be a group defined by a finite pregroup presentation $\mathcal{P} = \langle X \mid V_P \cup \mathcal{R} \rangle$ such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.20), let $\varepsilon \in \mathbb{R}_{>0}$, and let $h \in \mathbb{Z}_{\geq 1}$.

Chapter 9

The RSymVert scheme

In this chapter we shall present the curvature distribution scheme $\mathbf{RSymVert}$. We start by defining the enhanced vertex graph $\mathcal{E}$ of $\mathcal{P}$, which is an enhanced version of the vertex graph $\mathcal{G}$ given by [34, Definition 7.6], where each green $\mathcal{E}$-vertex $(R(i,a,b),\mathbf{G})$ corresponds to precisely one location $R(i,a,b)$ (see Definition 6.2.1). (Note the difference from [34, Definition 7.6] where there can be more than one location $R'(j,c,d)$ instantiating the same green $\mathcal{G}$-vertex $(c,d,\mathbf{G})$.) Before doing so, we summarize the following graph-theoretic terminology that will be used in our work.

Definition 9.0.1. Let G be a weighted directed simple graph. A walk $W = v_1v_2\ldots v_n$ is a sequence of vertices such that for all $1 \leq i \leq n-1$, we have $(v_i, v_{i+1}) \in E(G)$. We say that W is *open* if $v_1 \neq v_n$, otherwise we say that W is *closed* or a *circuit*. The sum of weights of all edges of W is denoted by $\omega(W)$.

Recall Definition 2.3.13 of an intermult pair.

Definition 9.0.2. The *enhanced vertex graph* $\mathcal{E}$ of $\mathcal{P}$ has two types of vertices. There is a green $\mathcal{E}$-vertex $(\mathcal{L},\mathbf{G})$ for each location $\mathcal{L}$. There is a red $\mathcal{E}$-vertex $(a,b,\mathbf{R})$ for each intermult pair (a,b).

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

Recall Definition 2.5.17 of the set $\mathcal{D}$ of coloured van Kampen diagrams.

Remark 9.0.3. There is a surjection from the vertices of the enhanced vertex graph to the vertices of the vertex graph. If $(R(i,a,b),\mathbf{G})$ is a green $\mathcal{E}$-vertex, the corresponding $\mathcal{G}$-vertex is $(a,b,\mathbf{G})$. On the other hand, a red $\mathcal{E}$-vertex $\nu = (a,b,\mathbf{R})$ is also a $\mathcal{G}$-vertex.

Also, applying Lemma 2.5.13 shows that each interior vertex v in a diagram $\Gamma \in \mathcal{D}$ is represented by a circuit in $\mathcal{E}$.

Definition 9.0.4. Let $\mathcal{W}$ be the set of all walks ν_1, ν, ν_2 in $\mathcal{E}$ of length 2, and let $\mathcal{W}_R$ be the set of all walks in $\mathcal{W}$ with middle vertex a green $\mathcal{E}$-vertex with relator R. Let F be an internal green face in a diagram $\Gamma \in \mathcal{D}$, incident with a vertex v such that $\delta(v, \Gamma) \geq 3$. Denote by $w(F, v, \Gamma)$ the multiset of all walks in $\mathcal{W}$ around v through a location of F – corresponding to three consecutive faces $F_1, F, F_2 \subseteq \Gamma$ incident with F at v.

Let $\mathcal{C}$ be the set of all circuits C in $\mathcal{E}$ with $\omega(C) \leq 5$ and $|C| \geq 4$. Further, for $w \in \mathcal{W}$, let $\mathcal{C}^w$ be the set of all circuits in $\mathcal{C}$ that contain w as a sub-walk.

Note that if $F \subseteq \Gamma \in \mathcal{D}$ is an internal green face incident with an (interior) vertex v of degree at least 3, or if $F \subseteq \Gamma \in \mathcal{D}$ is an (interior) green face incident with a vertex v with $\delta(v, \Gamma) \geq 3$, then every incidence of F with v is described by some walk $w \in w(F, v, \Gamma)$.

In the description below of **ComputeRSymVert**, the functions $\mathcal{M} : \mathcal{R}^{\pm 1} \times \mathbb{R} \times \{1, \ldots, h\} \to \mathbb{R} \times \{\mathtt{true}, \mathtt{false}\}$ and $\chi_{\text{out}} : \mathcal{W} \times \mathbb{R} \times \{1, \ldots, h\} \to \mathbb{R}$ are not yet defined. They are calculated by **VerifyHypVertex** (see Algorithm 11.2.1), and we shall give their definitions in Sections 10.3-10.4. For now we inform the reader that if there exists $\Gamma \in \mathcal{D}$ and an interior green face $F \subseteq \Gamma$ labelled by $R \in \mathcal{R}^{\pm 1}$, then $\mathcal{M}(R, \varepsilon, i)[2] = \mathtt{true}$, and if $\mathcal{M}(R, \varepsilon, i)[2] = \mathtt{true}$, then $\mathcal{M}(R, \varepsilon, i)[1]$ is an upper bound over all $\Gamma \in \mathcal{D}$ on $\alpha_\Gamma^{\varepsilon,i}(F) + \varepsilon$, where $\alpha_\Gamma^{\varepsilon,i}(F)$ is the curvature of an interior face $F \subseteq \Gamma$ labelled by R after the i^{th} iteration of $\mathbf{RSymVert}(\varepsilon, h)$ (see Algorithm 9.0.5 below): we shall prove this in Section 11.2, see Proposition 11.2.3. Furthermore, if $\mathcal{M}(R, \varepsilon, i)[2] = \mathtt{false}$ then $\mathcal{M}(R, \varepsilon, i) = (0, \mathtt{false})$.

Algorithm 9.0.5. $\mathbf{ComputeRSymVert}(\Gamma \in \mathcal{D}, \varepsilon, h)$:

For each iteration $i \in [1, \ldots, h]$ do:

- Step 1 If $i > 1$ then go to Step 2. Otherwise, let $\alpha_\Gamma^{\varepsilon,1} := \kappa_\Gamma = \mathbf{ComputeRSym}(\Gamma)$ (see Algorithm 2.6.4), and proceed to the next iteration.
- Step 2 For each internal green face F of Γ, denote the label of F by R. If $\mathcal{M}(R, \varepsilon, i-1)[1] \leq 0$, then for each interior vertex v with $\delta(v, \Gamma) \geq 3$ incident with F, if there is an interior green face incident with v with label R_1 such that $\mathcal{M}(R_1, \varepsilon, i-1)[1] > 0$, then let F give v curvature $\chi_{\text{out}}(w, \varepsilon, i-1)$ across each incidence described by some walk $w \in w(F, v, \Gamma)$.
- Step 3 For each interior vertex v of Γ, let $\mathcal{F}_v^i$ be the multiset of interior green faces incident with v and labelled by relators R that satisfy $\mathcal{M}(R, \varepsilon, i-1)[1] > 0$. Distribute v's curvature equally amongst the faces of $\mathcal{F}_v^i$.
- Step 4 Let $\alpha_\Gamma^{\varepsilon,i}$ be the function from $V(\Gamma) \cup E(\Gamma) \cup F(\Gamma)$ to $\mathbb{R}$, such that $\alpha_\Gamma^{\varepsilon,i}(x)$ is the current curvature of x. If $i = h$, then return $\alpha_\Gamma^{\varepsilon,h}$.

Definition 9.0.6. We define $\mathbf{RSymVert}(\varepsilon, h)$ to be the map from $\mathcal{D}$ to $\{\alpha_\Gamma^{\varepsilon,h}(x) : \Gamma \in \mathcal{D}\}$.

Let $2 \leq i \leq h$. We denote by $\Omega(F, w, i)$ the curvature that a green face F gives to a vertex v in Step 2 of the i^{th} iteration of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$ across the incidence described by some $w \in w(F, v, \Gamma)$ (we emphasize that if F gives no such curvature to v, then $\Omega(F, w, i) = 0$), and similarly, $\Pi(v, w, i)$ the curvature that v gives to F in Step 3 of the i^{th} iteration of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$ across the incidence described by some

$w \in w(F, v, \Gamma)$. Finally, for a vertex v of degree at least 3 incident with an interior green face F labelled by R, we define for each incidence of F with v described by some $w \in w(F, v, \Gamma)$:

$$\eta(w, i) = \begin{cases} -\Omega(F, w, i) & \text{if } v \text{ is interior, } \mathcal{M}(R, \varepsilon, i-1)[1] \leq 0, \text{and } \mathcal{F}_v^i \neq \emptyset, \\ \Pi(v, w, i) & \text{if } v \text{ is interior and } \mathcal{M}(R, \varepsilon, i-1)[1] > 0, \\ 0 & \text{otherwise.} \end{cases}$$

We shall think of $\eta(w, i)$ as the curvature moved from v to F across the incidence described by $w \in w(F, v, \Gamma)$ in iteration i (noting that this value can be positive), and of $\chi(v, \Gamma) + \sum_{n=2}^{i} \eta(w, n)$ (recall Definition 2.6.8 that $\chi(v, \Gamma)$ is the curvature that v gives to F across each incidence in the first iteration of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$) as the total curvature moved from v to F across the incidence described by $w \in w(F, v, \Gamma)$ over the first i iterations.

Proposition 9.0.7. $\mathbf{RSymVert}(\varepsilon, h)$ *is a curvature distribution scheme on* $\mathcal{D}$.

Proof. Let $\Gamma \in \mathcal{D}$. By Proposition 2.6.6 $\mathbf{RSym}$ is a curvature distribution scheme on $\mathcal{D}$, hence by Definition 2.6.1 we have $\kappa(\Gamma) = 1$. As in Steps 2 and 3 of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$ the curvature is neither created nor destroyed, we have $\alpha_{\Gamma}^{\varepsilon,h}(\Gamma) = 1$. The result follows. ∎

Lemma 9.0.8. *Let* $F \subseteq \Gamma \in \mathcal{D}$ *be a green face labelled by* R *and incident with an interior vertex* v *with* $\delta(v, \Gamma) \geq 3$*, let* $w \in w(F, v, \Gamma)$ *be a walk describing an incidence of* F *with* v*, and let* $i \geq 2$*. Suppose that* $\mathcal{M}(R, \varepsilon, i-1)[1] \leq 0$*. Then* $\Omega(F, w, i) = \chi_{out}(w, \varepsilon, i-1)$ *if and only if* $\mathcal{F}_v^i \neq \emptyset$.

Proof. Suppose that $\Omega(F, w, i) = \chi_{\text{out}}(w, \varepsilon, i-1)$. Then by Step 2 of Algorithm 9.0.5, v is incident with an interior green face with label R_1 such that $\mathcal{M}(R_1, \varepsilon, i-1)[1] > 0$, hence $\mathcal{F}_v^i \neq \emptyset$. Suppose instead that $\mathcal{F}_v^i \neq \emptyset$. Then by Step 2 of Algorithm 9.0.5 we have $\Omega(F, w, i) = \chi_{\text{out}}(w, \varepsilon, i-1)$. ∎

Finally, we define the success of $\mathbf{RSymVert}(\varepsilon, h)$.

Definition 9.0.9. We say that $\mathbf{RSymVert}(\varepsilon, h)$ *succeeds* on a diagram $\Gamma \in \mathcal{D}$ if $\alpha_{\Gamma}^{\varepsilon,h}(F) \leq -\varepsilon$ for all interior green faces F of Γ.

We say that $\mathbf{RSymVert}(\varepsilon, h)$ *succeeds on* $\mathcal{P}$ if this is true for every $\Gamma \in \mathcal{D}$.

There are infinitely many diagrams in $\mathcal{D}$, hence to test each of them to see whether $\mathbf{RSym}$ $\mathbf{Vert}(\varepsilon, h)$ succeeds is not possible. By using the fact that there are only finitely many elements in $\mathcal{R}$, in Chapters 10-11 we show how to test all such diagrams at once.

Chapter 10

The sub-routines of **VerifyHypVertex**

Recall Algorithm 9.0.5 for computing $\mathbf{RSymVert}(\varepsilon, h)$. Our main procedure, **VerifyHyp Vertex** (see Procedure 11.2.1), will be introduced in Section 11.2. The procedure takes as input $\mathcal{P}$, ε and h, where h is the maximum number of iterations, and returns either `true` or `fail`. If $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ returns `true`, then $\mathbf{RSymVert}(\varepsilon, h)$ is guaranteed to succeed on $\mathcal{P}$. If `fail` is returned, then this does not mean that $\mathbf{RSymVert}(\varepsilon, h)$ does not succeed on $\mathcal{P}$. The user can test again with a smaller value of ε or a larger value of h. In this chapter we shall describe the sub-routines of **VerifyHypVertex**. Throughout the whole chapter let $1 \leq \iota \leq h$.

Lemma 10.0.1. *Let $\Gamma \in \mathcal{D}$ contain an interior green face F. Then*

$$\alpha_\Gamma^{\varepsilon,\iota}(F) = 1 + \sum_{v \in \partial(F)} \sum_{w \in w(F,v,\Gamma)} \left(\chi(v,\Gamma) + \sum_{n=2}^{\iota} \eta(w,n) \right) + \sum_B n(B,F) \cdot \chi(B,\Gamma),$$

where the last sum is over all red blobs B that are edge-incident $n(B,F)$ times with F for some $n(B,F) \geq 1$.

Proof. Let $R \in \mathcal{R}^{\pm 1}$ be the label of F. By Algorithm 9.0.5 we have

$$\alpha_\Gamma^{\varepsilon,1}(F) = 1 + \sum_{v \in \partial(F)} \sum_{w \in w(F,v,\Gamma)} \chi(v,\Gamma) + \sum_B n(B,F) \cdot \chi(B,\Gamma),$$

and if $2 \leq n \leq \iota$, the by Algorithm 9.0.5 and Definition 9.0.6, if $x \in \Gamma$ is not an interior vertex of degree at least 3, then x neither gives curvature to nor receives curvature from F in iteration n; else the curvature moved from x to F across the incidence described by $w \in w(F, x, \Gamma)$ is equal to $\eta(w, n)$. Hence

$$\alpha_\Gamma^{\varepsilon,n}(F) - \alpha_\Gamma^{\varepsilon,n-1}(F) = \sum_{v \in \partial(F)} \sum_{w \in w(F,v,\Gamma)} \eta(w,n),$$

and so

$$\begin{aligned}
\alpha_\Gamma^{\varepsilon,\iota}(F) &= \alpha_\Gamma^{\varepsilon,1}(F) + \sum_{n=2}^{\iota}\left(\alpha_\Gamma^{\varepsilon,n}(F) - \alpha_\Gamma^{\varepsilon,n-1}(F)\right) \\
&= 1 + \sum_{v\in\partial(F)}\sum_{w\in w(F,v,\Gamma)} \chi(v,\Gamma) + \sum_B n(B,F)\cdot\chi(B,\Gamma) + \sum_{n=2}^{\iota}\left(\sum_{v\in\partial(F)}\sum_{w\in w(F,v,\Gamma)}\eta(w,n)\right) \\
&= 1 + \sum_{v\in\partial(F)}\sum_{w\in w(F,v,\Gamma)}\left(\chi(v,\Gamma) + \sum_{n=2}^{\iota}\eta(w,n)\right) + \sum_B n(B,F)\cdot\chi(B,\Gamma).
\end{aligned}$$

■

For each $R \in \mathcal{R}^{\pm 1}$, **VerifyHypVertex** seeks to find the maximum value of the curvature $\alpha_\Gamma^{\varepsilon,\iota}(F)$ over all $\Gamma \in \mathcal{D}$ of an interior face $F \subseteq \Gamma$ labelled by R. By Lemma 10.0.1 it suffices to bound $\alpha_\Gamma^{\varepsilon,\iota}(F)$ by bounding $\chi(v,\Gamma) + \sum_{n=2}^{\iota}\eta(w,n)$ and $n(B,F)\cdot\chi(B,\Gamma)$, over all diagrams in $\mathcal{D}$ that contain an interior green face F labelled by R.

Splitting up $\partial(F)$ into consolidated edges that F shares with its edge-incident faces induces a decomposition $R' = w_1 \dots w_k$, where R' is a cyclic conjugate of R, and each w_m labels the corresponding consolidated edge of F. We attach a colour $C_m \in \{\mathbf{G}, \mathbf{R}\}$ to each w_m, where C_m is the colour of the adjacent region (green face or a red blob). Since diagrams in $\mathcal{D}$ are green-rich, if $C_m = \mathbf{R}$, then $|w_m| = 1$.

Definition 10.0.2. If $C_m = \mathbf{G}$ then let ϵ_m^ι be the maximum value of

$$\chi(v_m,\Gamma) + \sum_{n=2}^{\iota}\eta(w,n),$$

considered over all possible diagrams $\Gamma \in \mathcal{D}$ in which w_m labels a maximal green consolidated edge on F, v_m is the vertex at the end of w_m, and $w \in w(F, v_m, \Gamma)$ is a walk determined by faces edge-incident with F in Γ at v_m. If $C_m = \mathbf{R}$ then let ϵ_m^ι be the maximum value of

$$\chi(v_m,\Gamma) + \sum_{n=2}^{\iota}\eta(w,n) + \chi(B,\Gamma),$$

considered over all possible diagrams $\Gamma \in \mathcal{D}$ on which w_m labels a red consolidated edge of F, v_m and $w \in w(F, v_m, \Gamma)$ are as in Case **G**, and all possible edge-incident red blobs B at w_m.

Definition 10.0.3. [34, Definition 7.2] Let $R' = w_1 w_2 \dots w_k$ be a coloured decomposition. A *step* consists either of a single sub-word w_m, or of two consecutive sub-words $w_m w_{m+1}$ that are determined as follows (interpret subscripts cyclically):

(i) If $C_m = \mathbf{G}$ and $C_{m+1} = \mathbf{R}$, then $w_m w_{m+1}$ is a step.

(ii) If neither $w_{m-1}w_m$ nor $w_m w_{m+1}$ is a step by Condition (i), then w_m is a step.

Let ϵ_m^ι be the curvature from Definition 10.0.2. A *stepwise curvature* χ_j^ι *with respect to* ι of a step s_j is $\chi_j^\iota = \epsilon_m^\iota$ when $s_j = w_m$ and $\chi_j^\iota = \epsilon_m^\iota + \epsilon_{m+1}^\iota$ when $s_j = w_m w_{m+1}$.

The *length* of a step is the number of letters of R' that it comprises.

We do not allow the combination $C_1 = \mathbf{R}$ and $C_k = \mathbf{G}$: if R' has $C_1 = \mathbf{R}$ and $C_k = \mathbf{G}$, then we consider the decomposition $w_2 \ldots w_k w_1$ instead. Hence it is clear from Definition 10.0.3 that for each such decomposition $R' = w_1 \ldots w_k$, there is a decomposition $R' = s_1 \ldots s_l$, where each s_j labels a step and is equal either to a single w_m, or to some $w_m w_{m+1}$ with w_m green and w_{m+1} red.

Lemma 10.0.4. *We have* $\alpha_\Gamma^{\varepsilon,\iota}(F) \leq 1 + \sum_{j=1}^{l} \chi_j^\iota$.

Proof. By Definitions 10.0.2-10.0.3 we have

$$\chi_j^\iota \geq \sum_v \left(\chi(v,\Gamma) + \sum_{n=2}^{\iota} \eta(w,n) \right) + \sum_B \chi(B,\Gamma),$$

where the first sum is over all vertices v that are at the end of a maximal consolidated edge of F labelled by a sub-word of s_j (and where $w \in w(F,v,\Gamma)$ is determined by faces edge-incident with F in Γ at v), and the last sum is over all red blobs edge-incident with F at a red consolidated edge labelled by a letter of s_j (there is always at most one such blob). Hence the lemma follows from Lemma 10.0.1. ∎

10.1 Vertex curvature

Let $R \in \mathcal{R}^{\pm 1}$, let F be an interior green face labelled by R in a diagram $\Gamma \in \mathcal{D}$, incident with a vertex v with $\delta(v,\Gamma) \geq 3$, and let F_1, F, F_2 be three consecutive faces incident with F at v. Recall [34, Algorithm 7.7, Section 7.2] for creating the *vertex function* $\mathbf{Vertex}$, which takes as input three $\mathcal{G}$-vertices corresponding to locations of F_1, F, F_2 at v, and returns an upper bound on $\chi(v,\Gamma)$. We extend this idea and create functions $\mathcal{Y}$ and $\mathcal{B}$, which take as input a walk $w \in w(F,v,\Gamma)$ (see Definition 9.0.4) determined by F_1, F, F_2; and ι, and return an upper bound on $\chi(v,F) + \sum_{j=2}^{\iota} \eta(w,j)$ (see Definition 9.0.6), where $\mathcal{Y}$ assumes that v is interior, and $\mathcal{B}$ assumes that $v \in \partial(\Gamma)$. We shall prove this in Section 11.2, see Lemma 11.2.2.

Definition 10.1.1. Let $w = \nu_1, \nu, \nu_2 \in \mathcal{W}_R$. We define $\mathcal{Y}(w,1) = \mathbf{Vertex}(\nu_1', \nu', \nu_2')$, where ν_1', ν', ν_2' are the $\mathcal{G}$-vertices that correspond to ν_1, ν, ν_2 (see Remark 9.0.3).

The value of $\mathcal{B}(w,1)$ is: $-1/3$ if ν_1 and ν_2 are both green; $-1/4$ if exactly one is green; and 0 otherwise.

In Section 10.4 we shall define $\mathcal{Y}(w,\iota)$ and $\mathcal{B}(w,\iota)$ for $\iota > 1$.

Lemma 10.1.2. *If* $v \in \partial(\Gamma)$*, then* $\chi(v,\Gamma) + \sum_{j=2}^{\iota} \eta(w,j) \leq \mathcal{B}(w,1)$.

Proof. Assume that $v \in \partial(\Gamma)$, and let $w = \nu_1, \nu, \nu_2$. By Definition 9.0.6 $\chi(v,\Gamma) + \sum_{j=2}^{\iota} \eta(w,j) = \chi(v,\Gamma)$. Now note that if ν_1 and ν_2 are both green, then $\delta_G(v) \geq 4$; $\delta_G(v) \geq 3$ if exactly one is green; and $\delta_G(v) \geq 2$ otherwise. Hence by Lemma 2.6.10 $\chi(v,\Gamma) \leq \mathcal{B}(w,1)$. ∎

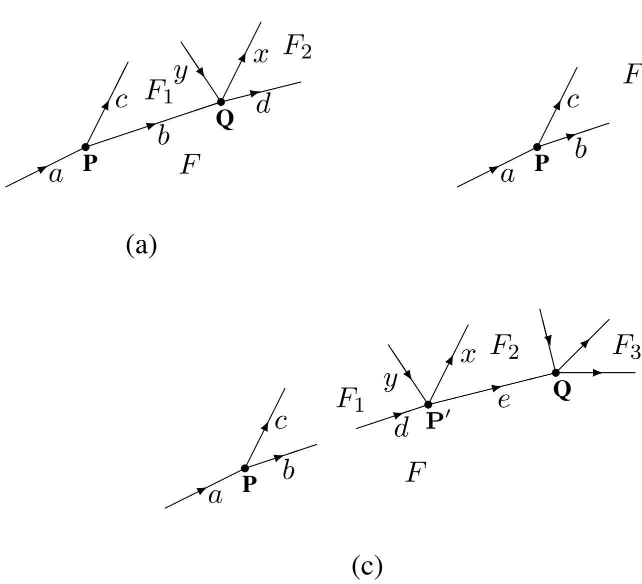

Figure 10.1: Illustration of the three types of Steps. ([34, Figure 5, page 30]).

10.2 One-step reachable places and the `OneStepVert` lists

In this section we describe how to find all places (see Definition 6.2.5) that can be reached from a fixed place by a single step.

Definition 10.2.1. [34, 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 a red edge between $\mathbf{P}'$ and $\mathbf{Q}$.

The algorithm $\mathbf{ComputeOneStepVert}(\mathbf{P}, \mathcal{Y}, \mathcal{B}, \mathbf{Blob}, \iota)$ takes as input a place $\mathbf{P}$; the functions $\mathcal{Y}$ and $\mathcal{B}$ (see Section 10.1); the function $\mathbf{Blob}$ (see [34, Algorithm 7.12, Section 7.4]), which takes as input three letters a, b, c and returns an upper bound on the curvature $\chi(B, \Gamma)$ (see Definition 2.6.8) that a red blob B with sub-word abc of its boundary word can give (across each edge-incidence) to an internal green face F edge-incident with B at b, considered over all diagrams in $\mathcal{D}$; and ι.

$\mathbf{ComputeOneStepVert}(\mathbf{P}, \mathcal{Y}, \mathcal{B}, \mathbf{Blob}, \iota)$ computes a list $\mathtt{OneStepVert}(\mathbf{P}, \iota)$ consisting of triples $(\mathbf{Q}, l, \chi)$, where $\mathbf{Q}$ is a one-step reachable place at distance l from $\mathbf{P}$, and

χ is an upper-bound on the stepwise curvature with respect to ι (see Definition 10.0.3) of the step between $\mathbf{P}$ and $\mathbf{Q}$.

In the description below of **ComputeOneStepVert**, by *including* an item $(\mathbf{Q}, l, \chi)$ in $\texttt{OneStepVert}(\mathbf{P}, \iota)$, we mean append it to the list if there is no entry of the form $(\mathbf{Q}, l, \chi')$ already or, if there is such an entry with $\chi > \chi'$, then replace that entry with $(\mathbf{Q}, l, \chi)$. (If there is such an entry with $\chi \leq \chi'$, we do nothing).

Algorithm 10.2.2. $\mathbf{ComputeOneStepVert}(\mathbf{P} = (R(i, a, b), c, C), \mathcal{Y}, \mathcal{B}, \mathbf{Blob}, \iota)$:

Step 1 Initialise $\texttt{OneStepVert}(\mathbf{P}, \iota)$ 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 10.1(a).)
Let $\nu_1 = (y, b^\sigma, \mathbf{R})$ and $\nu = (R(i+1, b, d), \mathbf{G})$. Let $\chi_1 := \mathbf{Blob}(y, b^\sigma, c)$, and for each out-neighbour ν_2 of ν, let $\chi_2 := \max\{\mathcal{Y}(\nu_1, \nu, \nu_2, \iota), \mathcal{B}(\nu_1, \nu, \nu_2, \iota)\}$. Include $(\mathbf{Q}, 1, \chi_1 + \chi_2)$ in $\texttt{OneStepVert}(\mathbf{P}, \iota)$.

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 of length l between R and R', let ν_1 be the $\mathcal{E}$-vertex corresponding to the location on R' at the end of α. For each out-neighbour ν_2 of $\nu = (R(j, d, e), \mathbf{G})$, compute $\chi_1 := \max\{\mathcal{Y}(\nu_1, \nu, \nu_2, \iota), \mathcal{B}(\nu_1, \nu, \nu_2, \iota)\}$.

1. If $\mathbf{P}'$ is green, then include $(\mathbf{P}', l, \chi_1)$ in $\texttt{OneStepVert}(\mathbf{P}, \iota)$. (See Figure 10.1(b).)
2. If $\mathbf{P}'$ is red, then find all places $\mathbf{Q}$ that are one letter further along R than $\mathbf{P}'$. (See Figure 10.1(c).) Then proceed in a similar way as in Case **R**, and compute the combined maximum curvature χ_2 returned by the red blob between $\mathbf{P}'$ and $\mathbf{Q}$ and the vertex at $\mathbf{Q}$. Include $(\mathbf{Q}, l+1, \chi_1 + \chi_2)$ in $\texttt{OneStepVert}(\mathbf{P}, \iota)$.

Lemma 10.2.3. *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.*

1. *The place $\mathbf{Q}$ is one-step reachable from $\mathbf{P}$ at distance l.*
2. *There exists a coloured decomposition of a cyclic conjugate R' of R such that a sub-word w_k or $w_k w_{k+1}$ of R' between the location of $\mathbf{P}$ and the location of $\mathbf{Q}$ is a step of length l, the face edge-incident with a face F labelled R at w_k has colour C, and edge after w_k^{-1} labelled c, and the face edge-incident with F at the letter after the end of the step has colour C' and next letter c'.*
3. *There exists χ such that $(\mathbf{Q}, l, \chi) \in \texttt{OneStepVert}(\mathbf{P}, \iota)$.*

Assume that for each interior green face $F' \subseteq \Gamma \in \mathcal{D}$, for each vertex $v \in \partial(F')$ of degree at least 3 and for each incidence of F' with v described by some walk $w \in w(F', v, \Gamma)$, we have

$$\chi(v, \Gamma) + \sum_{j=2}^{\iota} \eta(w, j) \leq \max\{\mathcal{Y}(w, \iota), \mathcal{B}(w, \iota)\}.$$

Then if $(\mathbf{Q}, l, \chi) \in$ `OneStepVert`$(\mathbf{P}, \iota)$*, then* χ *is an upper bound on the step curvature (with respect to* ι*).*

Proof. The equivalence of Statements 1-2 is [34, Lemma 7.17]. From the description of Algorithm 10.2.2 it follows that **ComputeOneStepVert** finds all one-step reachable places, and does not find any places that are not one-step reachable. Hence Statements 1 and 3 are also equivalent.

To prove the last statement, note that the step curvature is the sum of the total curvature moved (across a single incidence) to F from at most two (not necessarily distinct) vertices (of degree at least 3), and the curvature given (across a single edge-incidence) to F by at most one red blob. By the assumption the $\mathcal{Y}$ and $\mathcal{B}$ functions return an upper bound on $\chi(v, \Gamma) + \sum_{j=2}^{\iota} \eta(w, j)$, and by [34, Lemma 7.13] the **Blob** function returns an upper bound on $\chi(B, \Gamma)$. Hence χ is an upper bound on the step curvature. ∎

10.3 Finding an upper bound on curvature of green relators

For each $R \in \mathcal{R}^{\pm 1}$, we shall check whether there exists $\Gamma \in \mathcal{D}$ and an interior green face $F \subseteq \Gamma$ labelled by R. For those R for which the answer is positive, we shall find an upper bound on $\alpha_{\Gamma}^{\varepsilon,\iota}(F)$ (see Algorithm 9.0.5) over all such $\Gamma \in \mathcal{D}$.

The procedure **VertexVerifyAtPlace**$(\mathbf{P}, \varepsilon, \iota)$ takes as input a place $\mathbf{P}$ on a relator $R \in \mathcal{R}^{\pm 1}$; ε; and ι. If there exists a coloured decomposition $R' = s_1 \ldots s_l$ of some cyclic conjugate R' of R beginning at $\mathbf{P}$ such that each s_j labels a step, then **VertexVerifyAtPlace**$(\mathbf{P}, \varepsilon, \iota)$ returns a real number $c_{max} = \chi_{max} + (1+\varepsilon)$, where χ_{max} is the largest possible total curvature over all coloured decompositions p of a cyclic conjugate of R beginning at $\mathbf{P}$ arising from the steps of p; and `true`. Otherwise, the procedure returns 0 and `false`.

VertexVerifyAtPlace$(\mathbf{P}, \varepsilon, \iota)$ creates a list L whose elements are quadruples $(\mathbf{Q}, l, k, \psi)$. The first three components describe a place $\mathbf{Q}$ at distance l from $\mathbf{P}$ along R that can be reached from $\mathbf{P}$ in k steps. The fourth component ψ is $(1+\varepsilon)l/|R| + \chi$, where χ is an upper bound on the total curvature arising from these k steps.

Similarly to Algorithm 10.2.2, by *including* an entry $(\mathbf{Q}, l + l', i, \phi_1)$ in a list L, we mean appending it to L if there is no entry $(\mathbf{Q}, l + l', j, \phi_2)$ in L or, if there is such an entry with $\phi_1 > \phi_2$, then replacing it by $(\mathbf{Q}, l + l', i, \phi_1)$.

Procedure 10.3.1. VertexVerifyAtPlace$(\mathbf{P} = (R(i, a, b), c, C), \varepsilon, \iota)$:

Step 1 Initialise $L := [(\mathbf{P}, 0, 0, 0)]$ and `Curvs`$(\iota) := [\ ]$.

Step 2 For each $i := 1$ to $|R|$ do:

- (*) For each $(\mathbf{P}, l, i-1, \Psi)$ in L with $l < |R|$, and for each $(\mathbf{Q}, l', \chi) \in$ `OneStepVert`$(\mathbf{P}, \iota)$ with $l + l' \leq |R|$ do:
  - (i) Let $\Psi' := \Psi + \chi + (1+\varepsilon)l'/|R|$.
  - (ii) If $l + l' = |R|$ and $\mathbf{Q} \neq \mathbf{P}$, then do nothing.
  - (iii) Else if $l + l' = |R|$ and $\mathbf{Q} = \mathbf{P}$, then append Ψ' to `Curvs`(ι).
  - (iv) Else include $(\mathbf{Q}, l + l', i, \Psi')$ in L.

Step 4 If $\texttt{Curvs}(\iota) = [\ ]$, then return $0, \texttt{false}$. Else return $\max\{x : x \in \texttt{Curvs}(\iota)\}, \texttt{true}$.

Even though $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ (see Procedure 11.2.1) might not compute all h iterations, we purposefully let ι be the input to its sub-routines, so that we can work over all $i \in \{1, \dots, h\}$.

Lemma 10.3.2. *Let $R \in \mathcal{R}^{\pm 1}$. Then the following holds.*

1. *Suppose that there exists an interior green face $F \subseteq \Gamma \in \mathcal{D}$ labelled by R. Then for $1 \leq i \leq h$, there exists a place $\mathbf{P}$ on R such that $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)[2] = \texttt{true}$.*

2. *Suppose that there is a place $\mathbf{P}$ on R such that $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)[2] = \texttt{true}$ for some $1 \leq i \leq h$. Then for all $1 \leq k \leq h$ we have $\mathbf{VertexVerifyAtPlace}$ $(\mathbf{P}, \varepsilon, k)[2] = \texttt{true}$.*

Proof. Part 1. Let $R' = w_1 w_2 \dots w_k$ be a coloured decomposition of some cyclic conjugate R' of R, such that each w_j labels a consolidated edge e_j of F with $|e_j| \geq 1$, and has an associated colour $C_j \in \{\mathbf{G}, \mathbf{R}\}$, which is the colour of the face edge-incident with F at e_j. Since we do not allow the combination $C_1 = \mathbf{R}$, $C_k = \mathbf{G}$, there is a decomposition $R' = s_1 s_2 \dots s_l$, where each s_m 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 $\mathbf{P}$ be the place on R at the beginning of R', and note that $\mathbf{P}$ is instantiable since F is interior and $\Gamma \in \mathcal{D}$. By Lemma 10.2.3 each step s_m corresponds to a pair $\mathbf{P}, \mathbf{Q}$ of places on R, and there exist l, χ such that $(\mathbf{Q}, l, \chi) \in \texttt{OneStepVert}(\mathbf{P}, i)$. Hence by the description of Procedure 10.3.1 we have $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)[2] = \texttt{true}$.

Part 2. Let $\mathbf{P}_0 = \mathbf{P}$. Since $\mathbf{VertexVerifyAtPlace}(\mathbf{P}_0, \varepsilon, i)[2] = \texttt{true}$, by the description of Procedure 10.3.1, there exists a sequence $(\mathbf{P}_j, l_j, \chi_j)_{j=0}^{c}$ satisfying

(*) $(\mathbf{P}_j, l_j, \chi_j) \in \texttt{OneStepVert}(\mathbf{P}_{j-1}, i), \sum_{m=1}^{c} l_m = |R|,$ and $\mathbf{P}_c = \mathbf{P}_0$.

Furthermore, if a sequence $(\mathbf{P}_j, l_j, \chi_j)_{j=0}^{c}$ satisfies Statement (*), then by Lemma 10.2.3 for all $1 \leq j \leq c$, the place $\mathbf{P}_j$ is one-step reachable from $\mathbf{P}_{j-1}$ at distance l_j. Hence Part 2 holds. ∎

We now define the function $\mathcal{M}(R, \varepsilon, \iota)$ from the description of $\mathbf{ComputeRSymVert}$. Since ι we chosen to be arbitrary, we emphasize that we define $\mathcal{M}(R, \varepsilon, i)$ for all $i \in \{1, \dots, h\}$.

Definition 10.3.3. Let $R \in \mathcal{R}^{\pm 1}$. If there exists a place $\mathbf{P}$ on R such that $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, \iota)[2] = \texttt{true}$, then define

$$\mathcal{M}(R, \varepsilon, \iota) := (\max\{\mathbf{VertexVerifyAtPlace}(\mathbf{Q}, \varepsilon, \iota)[1] : \mathbf{Q} \in R \text{ and } \mathbf{VertexVerifyAtPlace}(\mathbf{Q}, \varepsilon, \iota)[2] = \texttt{true}\}, \texttt{true}).$$

Else define $\mathcal{M}(R, \varepsilon, \iota) := (0, \texttt{false})$.

Corollary 10.3.4. *Let $R \in \mathcal{R}^{\pm 1}$. The both of the following holds.*

1. *Suppose that* $\mathcal{M}(R,\varepsilon,i)[2] = \texttt{true}$ *for some* $1 \le i \le h$*. Then for all* $1 \le k \le h$ *we have* $\mathcal{M}(R,\varepsilon,k)[2] = \texttt{true}$*.*
2. *Suppose that there exists an interior green face* $F \subseteq \Gamma \in \mathcal{D}$ *labelled by* R*. Then for all* $1 \le k \le h$ *we have* $\mathcal{M}(R,\varepsilon,k)[2] = \texttt{true}$*.*

Proof. Part 1. By Definition 10.3.3 there exists a place $\mathbf{P}$ on R such that **VertexVerifyAtPlace** $(\mathbf{P},\varepsilon,i)[2] = \texttt{true}$. Hence by Lemma 10.3.2 for all $1 \le k \le h$ we have **VertexVerifyAtPlace** $(\mathbf{P},\varepsilon,k)[2] = \texttt{true}$. So $\mathcal{M}(R,\varepsilon,k)[2] = \texttt{true}$.

Part 2. By Lemma 10.3.2 for all $1 \le k \le h$, there exists a place $\mathbf{P}$ on R such that **VertexVerifyAtPlace**$(\mathbf{P},\varepsilon,k)[2] = \texttt{true}$, hence $\mathcal{M}(R,\varepsilon,k)[2] = \texttt{true}$. ∎

The procedure **VertexVerify**(ε,ι) runs **VertexVerifyAtPlace**$(\mathbf{P},\varepsilon,\iota)$ at each place $\mathbf{P}$. If no call returns m with $m[1] > 0$, then **VertexVerify**(ε,ι) returns `true`, 0, 0. Otherwise, **VertexVerify**(ε,ι) calculates values $\mathcal{M}(R,\varepsilon,\iota)$ for the given ι, and returns `fail`, the function $\mathcal{M}$, and set $S_{\text{Rich}} = \{\nu = (R(j,c,d),\mathbf{G}) \in V(\mathcal{E}) : \mathcal{M}(R,\varepsilon,\iota)[1] \le 0\}$.

In the description below, by *including* an item (R,χ) in L_{MaxCurvs}, we mean appending it to L_{MaxCurvs} if there is no entry $(R,\chi') \in L_{\text{MaxCurvs}}$ or, if there is such an entry with $\chi > \chi'$, then replacing it by (R,χ). (If there is such an entry with $\chi \le \chi'$, we do nothing).

Procedure 10.3.5. VertexVerify(ε,ι):

Step 1 Initialise $S_{\text{Rich}} = \emptyset$ and $L_{\text{MaxCurvs}} := [\,]$.
Step 2 For each $\mathbf{P} = (R(i,a,b),c,C)$, let $m = $ **VertexVerifyAtPlace**$(\mathbf{P},\varepsilon,\iota)$.

(*) If $m[2] = \texttt{true}$, then include $(R, m[1])$ in L_{MaxCurvs}.

Step 3 If ($L_{\text{MaxCurvs}} = [\,]$) or ($L_{\text{MaxCurvs}} \ne [\,]$ and $\{(R,\chi) \in L_{\text{MaxCurvs}} : \chi > 0\} = \emptyset$), then return `true`, 0, 0.
Step 4 Initialize $S = \emptyset$. For each $(R,\chi) \in L_{\text{MaxCurvs}}$, $S := S \cup \{R\}$, and set

$$\mathcal{M}(R,\varepsilon,\iota) := (\chi, \texttt{true}).$$

Step 5 For each $R \in \mathcal{R}^{\pm 1} \setminus S$, set $\mathcal{M}(R,\varepsilon,\iota) := (0, \texttt{false})$.
Step 6 For each location $R(i,a,b)$, if $\mathcal{M}(R,\varepsilon,\iota)[1] \le 0$, then

$$S_{\text{Rich}} := S_{\text{Rich}} \cup (R(i,a,b),\mathbf{G}).$$

Step 7 Return `fail`, $\mathcal{M}$, S_{Rich}.

10.4 Updating the vertex curvature

Throughout the whole section let $R \in \mathcal{R}^{\pm 1}$, and assume that $h \ge 2$ and $\iota \le h-1$. In Definition 10.1.1 we defined values $\mathcal{Y}(w,1)$ and $\mathcal{B}(w,1)$ for each $w \in \mathcal{W}_R$. In this section we shall define $\mathcal{Y}(w,i)$ and $\mathcal{B}(w,i)$ for all $i \in \{2,\dots,h\}$.

We first define the function $\chi_{\text{out}}(w,\varepsilon,\iota)$ (where $w \in \mathcal{W}_R$) from the description of **Compute RSymVert** (see Algorithm 9.0.5). We emphasize that $\chi_{\text{out}}(w,\varepsilon,\iota)$ is defined only in cases $\mathcal{M}(R,\varepsilon,\iota)[1] \le 0$: we shall use it to define $\mathcal{Y}(w,\iota+1)$ and $\mathcal{B}(w,\iota+1)$ for such cases.

Definition 10.4.1. Let $w \in \mathcal{W}_R$. Assume that $\mathcal{M}(R,\varepsilon,\iota)[1] \leq 0$. If $\mathcal{M}(R,\varepsilon,\iota)[2] = \texttt{false}$ and $\iota = 1$, then define

$$\chi_{\text{out}}(w,\varepsilon,\iota) := \mathcal{Y}(w,1),$$

and if $\iota > 1$, then define $\chi_{\text{out}}(w,\varepsilon,\iota) := 0$.

If $\mathcal{M}(R,\varepsilon,\iota)[2] = \texttt{true}$, then let $m \geq 1$ be minimal subject to $\mathcal{M}(R,\varepsilon,m)[1] \leq 0$, so $1 \leq m \leq \iota$. If $m = \iota$, then define

$$\chi_{\text{out}}(w,\varepsilon,\iota) := \max\{\frac{\mathcal{M}(R,\varepsilon,\iota)[1]}{|R|}, \mathcal{Y}(w,1)\},$$

else if $\iota > m$ and $\mathcal{Y}(w,\iota) \leq \mathcal{Y}(w,m) - \mathcal{Y}(w,1)$, then define

$$\chi_{\text{out}}(w,\varepsilon,\iota) := \max\{\frac{\mathcal{M}(R,\varepsilon,\iota)[1]}{|R|}, \mathcal{Y}(w,\iota) - (\mathcal{Y}(w,m) - \mathcal{Y}(w,1))\},$$

else define

$$\chi_{\text{out}}(w,\varepsilon,\iota) := 0.$$

Finally, in all cases define

$$\mathcal{Y}(w,\iota+1) := \mathcal{Y}(w,\iota) - \chi_{\text{out}}(w,\varepsilon,\iota)$$

and

$$\mathcal{B}(w,\iota+1) := \mathcal{B}(w,\iota) - \chi_{\text{out}}(w,\varepsilon,\iota).$$

By Lemma 9.0.8, if $F \subseteq \Gamma \in \mathcal{D}$ is a green face labelled by R with $\mathcal{M}(R,\varepsilon,\iota)[\iota] \leq 0$ that is incident with an interior vertex v such that $\mathcal{F}_v^{\iota+1} \neq \emptyset$, then $-\Omega(F,w,\iota+1) = -\chi_{\text{out}}(w,\varepsilon,\iota)$ (see Definition 9.0.6) for each incidence of F with v described by some walk $w \in w(F,v,\Gamma)$. Hence that is why we define $\mathcal{Y}(w,\iota+1)$ and $\mathcal{B}(w,\iota+1)$ as in Definition 10.4.1.

Recall Definition 9.0.4 of $\mathcal{C}$ and of $\mathcal{C}^w$. We now define $\mathcal{Y}(w,\iota+1)$ and $\mathcal{B}(w,\iota+1)$ for the case $\mathcal{M}(R,\varepsilon,\iota)[1] > 0$. To do so, we define auxiliary functions $\xi : \mathcal{W} \times \mathcal{C} \times \mathbb{R} \times \{1,\ldots,h\} \to \mathbb{R}$ and $\chi_{\text{in}} : \mathcal{W} \times \mathbb{R} \times \{1,\ldots,h\} \to \mathbb{R}$, such that $\xi(w,C,\varepsilon,\iota)$ and $\chi_{\text{in}}(w,\varepsilon,\iota)$ return an upper bound on $\eta(w,\iota+1)$ over all $\Gamma \in \mathcal{D}$ that contain an interior green face F labelled by R, incident with an interior vertex v such that $w \in w(F,v,\Gamma)$, where $\xi(w,C,\varepsilon,\iota)$ requires correspondence between locations of the faces around v and the $\mathcal{E}$-vertices of C, and $\chi_{\text{in}}(w,\varepsilon,\iota)$ requires $\delta_G(v,\Gamma) \leq 5$.

Definition 10.4.2. Let $w \in \mathcal{W}_R$, and assume that $\mathcal{M}(R,\varepsilon,\iota)[1] > 0$, so $\mathcal{M}(R,\varepsilon,\iota)[2] = \texttt{true}$. Define

$$\mathcal{B}(w,\iota+1) := \mathcal{B}(w,\iota).$$

Let $C \in \mathcal{C}^w$, with m green $\mathcal{E}$-vertices with relators R_1 such that $\mathcal{M}(R_1,\varepsilon,\iota)[1] > 0$. If there is no green $\mathcal{E}$-vertex on C with relator R_1 such that $\mathcal{M}(R_1,\varepsilon,\iota)[1] \leq 0$, then define

$\xi(w, C, \varepsilon, \iota) = 0$. Otherwise, define

$$\xi(w, C, \varepsilon, \iota) := \sum_{w_1 \in C} \frac{\chi_{\text{out}}(w_1, \varepsilon, \iota)}{m}.$$

Next define

$$\chi_{\text{in}}(w, \varepsilon, \iota) := \max_{C \in \mathcal{C}^w} \xi(w, C, \varepsilon, \iota).$$

Finally, define

$$\mathcal{Y}(w, \iota + 1) := \mathcal{Y}(w, \iota) + \chi_{\text{in}}(w, \varepsilon, \iota).$$

We now present a procedure that calculates $\mathcal{Y}(w, i)$ and $\mathcal{B}(w, i)$ for $i \in \{2, \ldots, h\}$.

Procedure 10.4.3. $\mathbf{VertCurvsModify}(\mathcal{W}_G, \mathcal{M}, \mathcal{Y}, \mathcal{B}, S_{\text{Rich}}, \mathcal{C}, \iota)$:

// Input: $\mathcal{W}_G$ – the set of all walks in $\mathcal{W}$ with green middle vertex.
// $\mathcal{M}$: the function $\mathcal{M}(R, \varepsilon, i)$ with already defined values for each
// $R \in \mathcal{R}^{\pm 1}$ and each $i \le \iota$.
// $\mathcal{Y}, \mathcal{B}$: the functions $\mathcal{Y}(w, i)$ and $\mathcal{B}(w, i)$ with already defined values
// for each $w \in \mathcal{W}_G$ and each $i \le \iota$.
// $S_{\text{Rich}} = \mathbf{VertexVerify}(\varepsilon, \iota)[3]$ (see Procedure 10.3.5).

Step 1 For each $w \in \mathcal{W}_G$ with the middle vertex in S_{Rich}, use the $\mathcal{M}$, $\mathcal{Y}$ and $\mathcal{B}$ functions to set

$$\mathcal{Y}(w, \iota + 1) := \mathcal{Y}(w, \iota) - \chi_{\text{out}}(w, \varepsilon, \iota)$$

and

$$\mathcal{B}(w, \iota + 1) := \mathcal{B}(w, \iota) - \chi_{\text{out}}(w, \varepsilon, \iota).$$

Step 2 For each $w \in \mathcal{W}_G$ with the middle vertex not in S_{Rich} do:

(i) Set
$$\mathcal{B}(w, \iota + 1) := \mathcal{B}(w, \iota)$$
and initialise
$$\mathcal{Y}(w, \iota + 1) := \mathcal{Y}(w, \iota).$$

(ii) Let L_1 be a list of all $C \in \mathcal{C}$ containing w as a sub-walk.

(iii) If $L_1 \neq [\,]$, then initialize InCurvs $= \emptyset$, and for each $C \in L_1$ do:

  (A) Use the $\mathcal{M}$, $\mathcal{Y}$ and $\mathcal{B}$ functions to calculate $\xi(w, C, \varepsilon, \iota)$ as in Definition 10.4.2.

  (B) InCurvs := InCurvs $\cup \{\xi(w, C, \varepsilon, \iota)\}$.

(iv) If InCurvs $\neq \emptyset$, then update

$$\mathcal{Y}(w, \iota + 1) := \mathcal{Y}(w, \iota) + \max\{x : x \in \text{InCurvs}\}.$$

Step 3 Return $\mathcal{Y}, \mathcal{B}$.

Proposition 10.4.4. *Let $R \in \mathcal{R}^{\pm 1}$. Assume that $\mathcal{M}(R, \varepsilon, i)[2] =$ `true` and $\mathcal{M}(R, \varepsilon, i)[1] \le 0$ for some $1 \le i \le h - 1$. If $i < j \le h$, then $\mathcal{M}(R, \varepsilon, j)[1] \le 0$.*

Proof. Let $m \geq 1$ be minimal subject to $\mathcal{M}(R, \varepsilon, m)[1] \leq 0$, and let $w \in \mathcal{W}_R$. By Definition 10.4.1 one of the following statements hold.

1. $\chi_{\text{out}}(w, \varepsilon, i) = \max\{\frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}, \mathcal{Y}(w, 1)\} \geq \frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}$.

2. $\chi_{\text{out}}(w, \varepsilon, i) = \max\{\frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}, \mathcal{Y}(w, i) - (\mathcal{Y}(w, m) - \mathcal{Y}(w, 1))\} \geq \frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}$.

3. $\chi_{\text{out}}(w, \varepsilon, i) = 0 \geq \frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}$.

Hence $\chi_{\text{out}}(w, \varepsilon, i) \geq \frac{\mathcal{M}(R,\varepsilon,i)[1]}{|R|}$, so by Definition 10.4.1 we have

$$\max\{\mathcal{Y}(w, i+1), \mathcal{B}(w, i+1)\} \leq \max\{\mathcal{Y}(w, i), \mathcal{B}(w, i)\} - \frac{\mathcal{M}(R, \varepsilon, i)[1]}{|R|}. \tag{10.1}$$

By Corollary 10.3.4, $\mathcal{M}(R, \varepsilon, i+1)[2] = \texttt{true}$, so let $\mathbf{P}$ be a place on R such that **VertexVerify AtPlace**$(\mathbf{P}, \varepsilon, i+1)[2] = \texttt{true}$. Let $\Psi \in \texttt{Curvs}(i+1)$ (see Procedure 10.3.1). Then Ψ is calculated by finding a sequence $(\mathbf{P}_j, l_j, \chi_j)_{j=0}^{c}$ satisfying

(*) $(\mathbf{P}_j, l_j, \chi_j) \in \texttt{OneStepVert}(\mathbf{P}_{j-1}, i+1), \sum_{n=1}^{c} l_n = |R|,$ and $\mathbf{P}_c = \mathbf{P}_0 = \mathbf{P}$,

and setting

$$\Psi = \sum_{j=1}^{c} \chi_j + (1+\varepsilon) \cdot \frac{l_j}{|R|} = (1+\varepsilon) \cdot (\frac{\sum_{j=1}^{c} l_j}{|R|}) + \sum_{j=1}^{c} \chi_j = 1 + \varepsilon + \sum_{j=1}^{c} \chi_j.$$

By Lemma 10.2.3 for all $1 \leq j \leq c$, the place $\mathbf{P}_j$ is one-step reachable from $\mathbf{P}_{j-1}$ at distance l_j, hence there is a sequence $(\mathbf{P}_j, l_j, \chi'_j)_{j=0}^{c}$ satisfying Statement (*) when replacing $i+1$ by i; and $\Psi' \in \texttt{Curvs}(i)$ with $\Psi' = 1 + \varepsilon + \sum_{j=1}^{c} \chi'_j$. By Algorithm 10.2.2 for all $1 \leq j \leq c$,

$$\chi_j = \sum_{w \in S_j \subseteq \mathcal{W}_R} \max\{\mathcal{Y}(w, i+1), \mathcal{B}(w, i+1)\} + \sum_{(a,b,d) \in S'_j \subseteq X^3} \mathbf{Blob}(a, b, d),$$

where $|S_j| \leq l_j$; and $\chi'_j \geq \chi''_j$, where

$$\chi''_j = \sum_{w \in S_j} \max\{\mathcal{Y}(w, i), \mathcal{B}(w, i)\} + \sum_{(a,b,d) \in S'_j} \mathbf{Blob}(a, b, d).$$

Hence by (10.1) we have

$$\chi_j \leq \chi''_j - \frac{l_j}{|R|}\mathcal{M}(R, \varepsilon, i)[1] \leq \chi'_j - \frac{l_j}{|R|}\mathcal{M}(R, \varepsilon, i)[1].$$

So

$$\begin{aligned}\Psi = 1+\varepsilon+\sum_{j=1}^{c}\chi_j \le 1+\varepsilon+\sum_{j=1}^{c}(\chi'_j - \frac{l_j}{|R|}\mathcal{M}(R,\varepsilon,i)[1])\\ = 1+\varepsilon+\sum_{j=1}^{c}\chi'_j - \frac{\sum_{j=1}^{c} l_j}{|R|}\mathcal{M}(R,\varepsilon,i)[1]\\ = 1+\varepsilon+\sum_{j=1}^{c}\chi'_j - \frac{|R|}{|R|}\mathcal{M}(R,\varepsilon,i)[1]\\ = \Psi' - \mathcal{M}(R,\varepsilon,i)[1].\end{aligned}$$

By Procedure 10.3.1 and Definition 10.3.3 we have

$$\Psi' \le \mathbf{VertexVerifyAtPlace}(\mathbf{P},\varepsilon,i)[1] \le \mathcal{M}(R,\varepsilon,i)[1],$$

and so $\Psi \le \Psi' - \mathcal{M}(R,\varepsilon,i)[1] \le 0$. Since Ψ was arbitrary, by Procedure 10.3.1 we have $\mathbf{VertexVerifyAtPlace}(\mathbf{P},\varepsilon,i+1)[1] \le 0$. So as $\mathbf{P}$ was arbitrary, by Definition 10.3.3 we have $\mathcal{M}(R,\varepsilon,i+1)[1] \le 0$. By induction the lemma holds for j. ■

The next lemma shows that the curvatures $\chi_{\text{out}}(w,\varepsilon,\iota)$ are always non-positive. This will enable us to prove that for each $\Gamma \in \mathcal{D}$ and for each interior green face $F \subseteq \Gamma$ with $\alpha_\Gamma^{\varepsilon,\iota}(F) > -\varepsilon$, we have $\alpha_\Gamma^{\varepsilon,\iota+1}(F) \le \alpha_\Gamma^{\varepsilon,\iota}(F)$.

Lemma 10.4.5. *Let $R \in \mathcal{R}^{\pm 1}$, and assume that $\mathcal{M}(R,\varepsilon,\iota)[1] \le 0$. Then for all $w \in \mathcal{W}_R$, both of the following two statements hold.*

1. *Assume that $\mathcal{M}(R,\varepsilon,\iota)[2] = \texttt{true}$, and let $m \in \mathbb{Z}$ be minimal subject to $\mathcal{M}(R,\varepsilon,m)[1] \le 0$, so $1 \le m \le \iota$. Then for all $m \le j \le \iota+1 \le h$ we have $\mathcal{Y}(w,j) \le \mathcal{Y}(w,m) - \mathcal{Y}(w,1)$.*

2. *We have $\chi_{out}(w,\varepsilon,\iota) \le 0$.*

Proof. By Definition 10.1.1 and [34, Algorithm 7.7, Section 7.2] we have $\mathcal{Y}(w,1) \le 0$.

Part 1. Proof is by induction on j. Base case $j = m$. We have $\mathcal{Y}(w,j) = \mathcal{Y}(w,m) \le \mathcal{Y}(w,m) - \mathcal{Y}(w,1)$. Assume by induction that $\iota+1 \ge j > m$ and $\mathcal{Y}(w,j-1) \le \mathcal{Y}(w,m) - \mathcal{Y}(w,1)$. By Proposition 10.4.4 we have $\mathcal{M}(R,\varepsilon,s)[1] \le 0$ for all $s \in \{m+1,\ldots,\iota+1\}$. If $j-1 = m$, then by Definition 10.4.1

$$\chi_{\text{out}}(w,\varepsilon,j-1) = \max\{\frac{\mathcal{M}(R,\varepsilon,j-1)[1]}{|R|}, \mathcal{Y}(w,1)\} \ge \mathcal{Y}(w,1)$$

and

$$\mathcal{Y}(w,j) = \mathcal{Y}(w,j-1) - \chi_{\text{out}}(w,\varepsilon,j-1) \le \mathcal{Y}(w,m) - \mathcal{Y}(w,1),$$

or $j-1>m$, and again by Definition 10.4.2

$$\begin{aligned}\chi_{\text{out}}(w,\varepsilon,j-1) &= \max\{\frac{\mathcal{M}(R,\varepsilon,j-1)[1]}{|R|}, \mathcal{Y}(w,j-1)-(\mathcal{Y}(w,m)-\mathcal{Y}(w,1))\}\\ &\geq \mathcal{Y}(w,j-1)-(\mathcal{Y}(w,m)-\mathcal{Y}(w,1))\end{aligned}$$

and

$$\begin{aligned}\mathcal{Y}(w,j) &= \mathcal{Y}(w,j-1)-\chi_{\text{out}}(w,\varepsilon,j-1)\\ &\leq \mathcal{Y}(w,j-1)-(\mathcal{Y}(w,j-1)-(\mathcal{Y}(w,m)-\mathcal{Y}(w,1))\\ &= \mathcal{Y}(w,j-1)-\mathcal{Y}(w,j-1)+\mathcal{Y}(w,m)-\mathcal{Y}(w,1)=\mathcal{Y}(w,m)-\mathcal{Y}(w,1).\end{aligned}$$

So Part 1 holds.

Part 2. Suppose first that $\mathcal{M}(R,\varepsilon,\iota)[2]=\texttt{false}$. If $\iota=1$, then by Definition 10.4.1 we have $\chi_{\text{out}}(w,\varepsilon,\iota)=\mathcal{Y}(w,1)\leq 0$, and if $\iota>1$, then by Definition 10.4.1 $\chi_{\text{out}}(w,\varepsilon,\iota)=0$. So Part 2 holds.

Assume instead that $\mathcal{M}(R,\varepsilon,\iota)[2]=\texttt{true}$, and let $m\in\mathbb{Z}$ be minimal subject to $\mathcal{M}(R,\varepsilon,m)[1]\leq 0$. If $\iota=m$, then by Definition 10.4.1 we have

$$\chi_{\text{out}}(w,\varepsilon,\iota)=\max\{\frac{\mathcal{M}(R,\varepsilon,\iota)[1]}{|R|},\mathcal{Y}(w,1)\}.$$

So $\chi_{\text{out}}(w,\varepsilon,\iota)\leq 0$ since $\mathcal{M}(R,\varepsilon,\iota)[1]\leq 0$ by assumption, and $\mathcal{Y}(w,1)\leq 0$. By Part 1 we have $\mathcal{Y}(w,\iota)\leq\mathcal{Y}(w,m)-\mathcal{Y}(w,1)$, so if $\iota>m$, then by Definition 10.4.1 we have

$$\chi_{\text{out}}(w,\varepsilon,\iota)=\max\{\frac{\mathcal{M}(R,\varepsilon,\iota)[1]}{|R|},\mathcal{Y}(w,\iota)-(\mathcal{Y}(w,m)-\mathcal{Y}(w,1))\},$$

so as $\mathcal{Y}(w,\iota)\leq\mathcal{Y}(w,m)-\mathcal{Y}(w,1)$ implies $\mathcal{Y}(w,\iota)-(\mathcal{Y}(w,m)-\mathcal{Y}(w,1))\leq 0$, we have $\chi_{\text{out}}(w,\varepsilon,\iota)\leq 0$. ■

Chapter 11

Proof of Theorem 1.0.4

In this chapter we shall prove Theorem 1.0.4. Recall that G is a group given by a finite pregroup presentation $\mathcal{P}$ such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.20), that $\varepsilon \in \mathbb{R}_{>0}$, and that $h \in \mathbb{Z}_{\geq 1}$.

11.1 The success of $\mathbf{RSymVert}$ shows hyperbolicity

In this section we shall show that if $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$ (see Algorithm 9.0.5 and Definition 9.0.9), then G is hyperbolic. Recall Definition 2.6.8 that $\chi(x, \Gamma)$ is the curvature that x gives to a single internal green face across each curvature incidence in the first iteration of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$, and that $\chi(x, F, \Gamma)$ is the total curvature that x gives to an internal green face F in the first iteration of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$.

Lemma 11.1.1. *Let $\Gamma \in \mathcal{D}$. Then*

1. *if $T \subseteq \Gamma$ is a red triangle, then $\alpha_\Gamma^{\varepsilon,h}(T) \leq 0$;*
2. *if $\mathbf{Area}(\Gamma) > 1$, and F is a boundary green face of Γ, then $\alpha_\Gamma^{\varepsilon,h}(F) \leq 1/2$.*

Proof. Since the curvature of T is not altered in Steps 2 and 3 of $\mathbf{ComputeRSymVert}(\Gamma, \varepsilon, h)$, we have $\alpha_\Gamma^{\varepsilon,h}(T) = \kappa_\Gamma(T)$, so Part 1 follows from [34, Lemma 6.8].

For Part 2, let $R \in \mathcal{R}^{\pm 1}$ be the label of F. By Algorithm 9.0.5, F gives no curvature to its edge-incident red blobs and incident boundary vertices. Let χ_1 be the total curvature given to F by them. By Algorithm 9.0.5, we have

$$\chi_1 = \sum_{v \in \partial(F) \cap \partial(\Gamma)} \chi(v, F, \Gamma) + \sum_B \chi(B, F, \Gamma),$$

where the last sum is over all red blobs edge-incident with F. Therefore, we can apply the same arguments as in the first two paragraphs of the proof of [34, Lemma 6.9] to show that $\chi_1 \leq -1/2$. Hence since F gives no curvature to its edge-incident red blobs and incident boundary vertices, if there is no interior vertex v such that $v \in \partial(F)$ and $\delta(v, \Gamma) \geq 3$, then by Algorithm 9.0.5, $\alpha_\Gamma^{\varepsilon,h}(F) = \kappa_\Gamma(F) \leq 1 + \chi_1 \leq 1 - 1/2 = 1/2$. So assume that such a

v exists, and let $w \in w(F, v, \Gamma)$ (see Definition 9.0.4) be a walk describing an incidence of F with v.

By Definition 10.1.1 $\mathcal{Y}(w, 1) = \mathbf{Vertex}(\nu_1, \nu, \nu_2)$, where ν_1, ν, ν_2 are the $\mathcal{G}$-vertices that correspond to the $\mathcal{E}$-vertices of w (see Remark 9.0.3). Hence by [34, Lemma 7.8] we have $\chi(v, \Gamma) \leq \mathcal{Y}(w, 1)$. If $h = 1$ then $\alpha_\Gamma^{\varepsilon,h}(F) = \kappa_\Gamma(F)$; and as F is boundary, if $h \geq 2$ and $\mathcal{M}(R, \varepsilon, i)[1] > 0$ for all $1 \leq i \leq h - 1$, then by Algorithm 9.0.5 in each iteration $i \in \{2, \ldots, h\}$, no $x \in \Gamma$ gives F curvature, and no such x receives curvature from F, so again $\alpha_\Gamma^{\varepsilon,h}(F) = \kappa_\Gamma(F)$. Hence assume that $h \geq 2$, and that there is $1 \leq i \leq h-1$ such that $\mathcal{M}(R, \varepsilon, i)[1] \leq 0$. Let $m \in \mathbb{Z}$ be minimal subject to $\mathcal{M}(R, \varepsilon, m)[1] \leq 0$, so $1 \leq m \leq h - 1$. Let $\Delta(w, h) = \sum_{j=m+1}^{h} \Omega(F, w, j)$ (see Definition 9.0.6). We show that $\Delta(w, h) \geq \chi(v, \Gamma)$.

For each $m + 1 \leq j \leq h$, by Lemma 9.0.8 we have $\Omega(F, w, j) = \chi_{\text{out}}(w, \varepsilon, j - 1)$ if and only if $\mathcal{F}_v^j \neq \emptyset$, and by Lemma 10.4.5 we have $\chi_{\text{out}}(w, \varepsilon, j - 1) \leq 0$. Hence $\Delta(w, h) = \sum_{j=m+1}^{h} \Omega(F, w, j) \geq \sum_{j=m}^{h-1} \chi_{\text{out}}(w, \varepsilon, j)$.

Case $\mathcal{M}(R, \varepsilon, 1)[2] = \texttt{false}$. Suppose that $\mathcal{M}(R, \varepsilon, j)[2] = \texttt{true}$ for some $2 \leq j \leq h$. Then by Corollary 10.3.4 $\mathcal{M}(R, \varepsilon, 1)[2] = \texttt{true}$, a contradiction. Hence $\mathcal{M}(R, \varepsilon, j)[2] = \texttt{false}$ for all $1 \leq j \leq h$ (and note that $m = 1$). So by Definition 10.4.1 we have $\chi_{\text{out}}(w, \varepsilon, 1) = \mathcal{Y}(w, 1)$ and $\chi_{\text{out}}(w, \varepsilon, j) = 0$ if $2 \leq j \leq h - 1$. Therefore, by the previous two paragraphs

$$\Delta(w, h) \geq \sum_{j=m}^{h-1} \chi_{\text{out}}(w, \varepsilon, j) = \mathcal{Y}(w, 1) \geq \chi(v, \Gamma).$$

Case $\mathcal{M}(R, \varepsilon, 1)[2] = \texttt{true}$. By Corollary 10.3.4 we have $\mathcal{M}(R, \varepsilon, j)[2] = \texttt{true}$ for all $1 \leq j \leq h$, hence by Proposition 10.4.4 $\mathcal{M}(R, \varepsilon, j)[1] \leq 0$ for all $j \in \{m + 1, \ldots, h\}$. Therefore, by Definition 10.4.1

$$\begin{aligned}\mathcal{Y}(w, h) &= \mathcal{Y}(w, m) + \sum_{j=m}^{h-1} \mathcal{Y}(w, j + 1) - \mathcal{Y}(w, j)\\ &= \mathcal{Y}(w, m) + \sum_{j=m}^{h-1} (-\chi_{\text{out}}(w, \varepsilon, j)),\end{aligned}$$

so $\sum_{j=m}^{h-1} \chi_{\text{out}}(w, \varepsilon, j) = \mathcal{Y}(w, m) - \mathcal{Y}(w, h)$. Hence

$$\Delta(w, h) \geq \sum_{j=m}^{h-1} \chi_{\text{out}}(w, \varepsilon, j) = \mathcal{Y}(w, m) - \mathcal{Y}(w, h).$$

By Lemma 10.4.5

$$\mathcal{Y}(w, h) \leq \mathcal{Y}(w, m) - \mathcal{Y}(w, 1),$$

hence

$$\begin{aligned}\Delta(w, h) &\geq \mathcal{Y}(w, m) - (\mathcal{Y}(w, m) - \mathcal{Y}(w, 1))\\ &= \mathcal{Y}(w, 1) \geq \chi(v, \Gamma),\end{aligned}$$

where the last inequality follows from the third paragraph. Hence we showed that $\Delta(w, h) \geq \chi(v, \Gamma)$.

By Algorithm 9.0.5, for $i \in \{2, \ldots, h\}$, v gives F no curvature in iteration i, and in no iteration $i \in \{1, \ldots, m\}$, F gives v curvature. Hence as both v and $w \in w(F, v, \Gamma)$ were chosen to be arbitrary; by Lemma 2.6.9 $\chi(v, \Gamma) \leq 0$ since $\delta_G(v) \geq 2$ (because $\Gamma \in \mathcal{D}$); and F gives no curvature to its edge-incident red blobs and incident boundary vertices, we have $\alpha_\Gamma^{\varepsilon,h}(F) \leq 1 + \chi_1 \leq 1/2$, as required. ∎

Lemma 11.1.2. *Assume that no $R \in \mathcal{R}^{\pm 1}$ has length 1 or 2 and that 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. Let Γ be a diagram in $\mathcal{D}$ with boundary length 2. Then $\mathbf{RSymVert}(\varepsilon, h)$ does not succeed on Γ.*

Proof. Suppose that $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on Γ. Since each $R \in \mathcal{R}^{\pm 1}$ has $|R| \geq 3$, we have $\mathbf{Area}(\Gamma) > 1$. By Lemma 11.1.1 each boundary face F has $\alpha_\Gamma^{\varepsilon,h}(F) \leq 1/2$ if F is green, and $\alpha_\Gamma^{\varepsilon,h}(F) \leq 0$ if it is red. Since $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on Γ, all positive curvature of $\alpha_\Gamma^{\varepsilon,h}$ lies on the boundary green faces, and sums to at least 1. Hence Γ has exactly two boundary faces, F_1 and F_2 say, both green and such that $\alpha_\Gamma^{\varepsilon,h}(F_1) = 1/2 = \alpha_\Gamma^{\varepsilon,h}(F_2)$. As any other internal green face F satisfies $\alpha_\Gamma^{\varepsilon,h}(F) < 0$, no such face exists. Hence by Algorithm 9.0.5, if $h \geq i > 1$, then no curvature is redistributed through any interior vertex of Γ in iteration i, so $\alpha_\Gamma^{\varepsilon,h} = \kappa_\Gamma = \mathbf{ComputeRSym}(\Gamma)$. Therefore, as by [34, Lemma 6.11] $\mathbf{RSym}$ does not succeed on Γ, it follows that $\mathbf{RSymVert}(\varepsilon, h)$ does not succeed on Γ, a contradiction. ∎

Theorem 11.1.3. *Assume that no $R \in \mathcal{R}^{\pm 1}$ has length 1 or 2 and that 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. If $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$, then no V^σ-letter is trivial in G.*

Proof. Suppose for a contradiction that there exists such a V^σ-letter x, and let Γ be a coloured van Kampen diagram over $\mathcal{P}$ with boundary word x, and with smallest possible coloured area for simply-connected diagrams with boundary word a single V^σ-letter. We do *not* assume that $\Gamma \in \mathcal{D}$. We show that Γ does not exist.

In the proof of [34, Theorem 6.12] (which shows that under the same assumptions as in this theorem, if $\mathbf{RSym}$ succeeds on $\mathcal{P}$, then no V^σ-letter is trivial in G), the only facts about $\mathbf{RSym}$ that are used are that $\mathbf{RSym}$ is a curvature distribution scheme on $\mathcal{D}$; that each boundary green face F of $\Gamma' \in \mathcal{D}$ satisfies $\kappa_{\Gamma'}(F) \leq 1/2$ if $\mathbf{Area}(\Gamma') > 1$; and that $\mathbf{RSym}$ does not succeed on $\Gamma' \in \mathcal{D}$ if $|\partial(\Gamma')| = 2$. Now Proposition 9.0.7; Lemma 11.1.1 and Lemma 11.1.2 give us the analogous results for $\mathbf{RSymVert}(\varepsilon, h)$. Hence the proof of [34, Theorem 6.12] shows that Γ does not exist, a contradiction. ∎

Definition 11.1.4. [34, Definition 5.5] The *pregroup Dehn function* $\mathrm{PD}(n) : \mathbb{Z}_{\geq 0} \to \mathbb{Z}$ of $\mathcal{P}$ is defined as follows. For each σ-reduced word $w \in X^*$ with $w =_G 1$, let $A(w)$ be the smallest number of internal faces of a coloured van Kampen diagram over $\mathcal{P}$ with boundary word w. Then $\mathrm{PD}(n) := \max\{A(w) : w \in X^*, w =_G 1, |w| \leq n\}$.

We now present the main result of this section.

Theorem 11.1.5. *Assume that no $R \in \mathcal{R}$ has length 1 or 2 and that no two distinct cyclic conjugates of relators $R, S \in \mathcal{R}^{\pm}$ have a common prefix consisting of all but one letter of R or S. Let r be the maximum length of a relator in $\mathcal{R}$.*

If $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$, then the pregroup Dehn function of $\mathcal{P}$ is bounded above by

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

In particular, G is hyperbolic, and an explicit linear bound on the Dehn function of G can be calculated.

Proof. We prove the theorem by showing that $\mathbf{RSymVert}(\varepsilon, h)$ satisfies all conditions of [34, Theorem 5.9]. By Proposition 9.0.7 $\mathbf{RSymVert}(\varepsilon, h)$ is a curvature distribution on $\mathcal{D}$. Let $\Gamma \in \mathcal{D}$, and let $F_R(\Gamma)$ be the set of all red faces of Γ. We first show that $\alpha_\Gamma^{\varepsilon,h}(x) \leq 0$ for all $x \in V(\Gamma) \cup E(\Gamma) \cup F_R(\Gamma)$. By Lemma 11.1.1 $\alpha_\Gamma^{\varepsilon,h}(x) \leq 0$ for each $x \in F_R(\Gamma)$, and clearly $\alpha_\Gamma^{\varepsilon,h}(x) = 0$ for each $x \in E(\Gamma)$. So let $x \in V(\Gamma)$. If only the first iteration of $\mathbf{ComputeRSymVert}(\Gamma,\, \varepsilon, h)$ is computed, then $\alpha_\Gamma^{\varepsilon,h}(x) \leq 0$ by Lemma 2.6.9 since $\delta_G(x, \Gamma) \geq 2$ (because $\Gamma \in \mathcal{D}$). Otherwise, in Step 2, x receives curvature only if x is interior and $\mathcal{F}_x^i \neq \emptyset$, and in Step 3, x gives all its curvature to the faces of $\mathcal{F}_x^i$, so after Step 3, x has curvature 0. Hence $\alpha_\Gamma^{\varepsilon,h}(x) \leq 0$, as claimed. Thus, $\mathbf{RSymVert}(\varepsilon, h)$ satisfies Condition (a) of [34, Theorem 5.9].

As $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$, it satisfies Condition (b). By Lemma 11.1.1 $\mathbf{RSym}$ $\mathbf{Vert}(\varepsilon, h)$ satisfies Condition (c) with $m = 1/2$. By Theorem 11.1.3 no V^σ-letter is trivial in G, hence all coloured van Kampen diagrams over $\mathcal{P}$ are loop-minimal (see [34, Definition 3.13]), so by [34, Proposition 3.17] all diagrams in $\mathcal{D}$ satisfy Condition (d) with $\lambda = 3 + r$ and $\mu = 1$. We can thus apply [34, Theorem 5.9 & Corollary 5.10] to show that if $\Gamma \in \mathcal{D}$ has boundary length n and $\mathbf{Area}(\Gamma) > 1$, then

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

Since all $R \in \mathcal{R}$ satisfy $|R| \geq 3$, a diagram of area 1 has boundary length $n \geq 3$. Now for $n \geq 3$, the above bound evaluates to at least 1, hence this area bound holds for all diagrams in $\mathcal{D}$.

By [34, Proposition 6.10] if $w =_G 1$, then some $w' \in \mathcal{I}(w)$ (see Definition 2.3.20) is a boundary word of a diagram Γ in $\mathcal{D}$, and by Lemma 2.5.7 there exists a coloured van Kampen diagram over $\mathcal{P}$ with boundary word w of area at most $\mathbf{Area}(\Gamma) + 2n$. Therefore, the given formula gives the desired bound on the pregroup Dehn function of $\mathcal{P}$.

Since we have an explicit linear bound on the pregroup Dehn function of G, [34, Lemma 5.8] gives us such a bound on the standard Dehn function of G. Thus, G is hyperbolic. ■

11.2 The main VerifyHypVertex algorithm

In this section we shall assume that all preprocessing steps from [34, Section 7.1] have been performed on $\mathcal{P}$, so that $\mathcal{P}$ satisfies assumptions of Theorem 11.1.5. Our aim is to show that if $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ returns `true`, then $\mathbf{RSymVert}(\varepsilon, h)$ (see Algorithm 9.0.5) succeeds on $\mathcal{P}$. This together with Theorem 11.1.5 proves the first statement of Theorem 1.0.4.

We shall now present $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$. In each iteration i and for each $R \in \mathcal{R}^{\pm 1}$, $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ finds an upper bound on $\alpha_\Gamma^{\varepsilon,i}(F)$ over all $\Gamma \in \mathcal{D}$ of an interior face $F \subseteq \Gamma$ labelled by R. If all these bounds are smaller than $-\varepsilon$, then $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ returns `true`. Otherwise, $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ either returns `fail`, or it proceeds to the next iteration.

In the description below, $\mathbf{FindCircuits}$ is a sub-routine that constructs $\mathcal{C}$ (see Definition 9.0.4) on input $\mathcal{E}$ and 5.

Procedure 11.2.1. $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$:

Step 1 Compute Steps 1-3 & 5-6 from the description of $\mathbf{RSymVerify}$ in [34, Section 7.6] to compute the intermult table, roots of green relators, locations, places, and to create the $\mathbf{Vertex}$ and $\mathbf{Blob}$ functions.

Step 2 Construct the enhanced vertex graph $\mathcal{E}$ (see Definition 9.0.2).

Step 3 Construct the set $\mathcal{W}_G$ of all walks in $\mathcal{W}$ (see Definition 9.0.4) with a green middle vertex.

Step 4 For all $w \in \mathcal{W}_G$, define $\mathcal{Y}(w, 1)$ and $\mathcal{B}(w, 1)$ (see Definition 10.1.1).

Step 5 For i in $[1, \ldots, h]$ do:

(i) For each place $\mathbf{P}$, run $\mathbf{ComputeOneStepVert}(\mathbf{P}, \mathcal{Y}, \mathcal{B}, \mathbf{Blob}, i)$ (see Algorithm 10.2.2). Store the list `OneStepVert`$(\mathbf{P}, i)$, for each such place $\mathbf{P}$.

(ii) Let $b, \mathcal{M}, S_{\text{Rich}} = \mathbf{VertexVerify}(\varepsilon, i)$ (see Procedure 10.3.5). If $b =$ `true`, then return `true`. Otherwise, if $i = h$ or if $S_{\text{Rich}} = \emptyset$, then return `fail`.

(iii) If $i = 1$, then compute $\mathcal{C} := \mathbf{FindCircuits}(\mathcal{E}, 5)$.

(iv) Let $\mathcal{Y}, \mathcal{B} = \mathbf{VertCurvsModify}(\mathcal{W}_G, \mathcal{M}, \mathcal{Y}, \mathcal{B}, S_{\text{Rich}}, \mathcal{C}, i)$ (see Procedure 10.4.3).

The next lemma shows that the $\mathcal{Y}$ and $\mathcal{B}$ functions provide correct bounds on the vertex curvature. Recall Definition 10.3.3 of $\mathcal{M}(R, \varepsilon, i)$, and Definition 9.0.4 that given an internal green face $F \subseteq \Gamma \in \mathcal{D}$ and a vertex $v \in \partial(F)$ of degree at least 3, $w(F, v, \Gamma)$ is the multiset of all walks in $\mathcal{W}$ around v through a location of F that correspond to three consecutive faces $F_1, F, F_2 \subseteq \Gamma$ incident with F at v.

Lemma 11.2.2. *Let $F \subseteq \Gamma \in \mathcal{D}$ be an interior green face incident with a vertex v of degree at least 3, let $w \in w(F, v, \Gamma)$ be a walk describing an incidence of F with v, and let $1 \leq i \leq h$. Then*

$$\chi(v, \Gamma) + \sum_{j=2}^{i} \eta(w, j) \leq \max\{\mathcal{Y}(w, i), \mathcal{B}(w, i)\}. \tag{11.1}$$

Proof. If $v \in \partial(\Gamma)$, then setting $\iota = 1$ in the statement of Lemma 10.1.2 gives $\chi(v, \Gamma) \leq \mathcal{B}(w, 1)$. If v is interior, then by Definition 10.1.1 $\mathcal{Y}(w, 1) = \mathbf{Vertex}(\nu_1, \nu, \nu_2)$, where

ν_1, ν, ν_2 are the $\mathcal{G}$-vertices that correspond to the $\mathcal{E}$-vertices of w (see Remark 9.0.3). Hence by [34, Lemma 7.8] $\chi(v, \Gamma) \leq \mathcal{Y}(w, 1)$, and so (11.1) holds for $i = 1$.

Hence let $i \geq 2$. Let $R \in \mathcal{R}^{\pm 1}$ be the label of F. By Corollary 10.3.4 $\mathcal{M}(R, \varepsilon, j)[2] =$ `true` for all $1 \leq j \leq h$.

Case $\mathcal{M}(R, \varepsilon, i-1)[1] > 0$. If $j < i-1$ and $\mathcal{M}(R, \varepsilon, j)[1] \leq 0$, then by Proposition 10.4.4 $\mathcal{M}(R, \varepsilon, i-1)[1] \leq 0$, a contradiction. So $\mathcal{M}(R, \varepsilon, j)[1] > 0$ for all $j \leq i-1$. Hence by Definition 10.4.2 $\mathcal{B}(w, j) = \mathcal{B}(w, 1)$ for all $j \leq i$. So if v is boundary, then the lemma holds by Lemma 10.1.2. Hence assume that v is interior. By Definition 9.0.6 for each $2 \leq j \leq i$ we have $\eta(w, j) = \Pi(v, w, j)$. By Lemma 10.4.5 for each $R_1 \in \mathcal{R}^{\pm 1}$ and each length 3 walk w_1 we have $\chi_{\text{out}}(w_1, \varepsilon, j-1) \leq 0$. Hence as by the description of **ComputeRSymVert**(Γ, ε, h) (see Algorithm 9.0.5), $\Pi(v, w, j)$ is a sum of finitely many values $\frac{\chi_{\text{out}}(w_1, \varepsilon, j-1)}{m}$ with $m > 0$, we have $\Pi(v, w, j) \leq 0$. So if $\delta_G(v, \Gamma) \geq 6$, then by Lemma 2.6.10 and Definition 10.1.1 we have

$$\chi(v, \Gamma) + \sum_{j=2}^{i} \Pi(v, w, j) \leq \chi(v, \Gamma) \leq -1/3 \leq \mathcal{B}(w, 1) = \mathcal{B}(w, i),$$

so we can assume that $\delta_G(v, \Gamma) \leq 5$. Then by Remark 9.0.3 there exists $C \in \mathcal{C}^w$ with $\mathcal{E}$-vertices corresponding to locations of faces around v.

Let $2 \leq j \leq i$, and suppose that there is a face F_1 incident with v in location $\mathcal{L}$ such that $\mathcal{M}(R_1, \varepsilon, j-1)[1] \leq 0$. Let $w_1 \in w(F_1, v, \Gamma)$ be the sub-walk of C around v through $\mathcal{L}$. Further, let m be the number of green $\mathcal{E}$-vertices of C with relators R' such that $\mathcal{M}(R', \varepsilon, j-1)[1] > 0$. In Step 2 of the j^{th} iteration of **ComputeRSymVert**(Γ, ε, h), F_1 gives v curvature $\chi_{\text{out}}(w_1, \varepsilon, j-1)$ across the incidence described by w_1. On the other hand, in Step 3 of the j^{th} iteration of **ComputeRSymVert**(Γ, ε, h), we have $|\mathcal{F}_v^j| \leq m$ since $\mathcal{F}_v^j$ contains only interior green faces. Hence in Step 3 of the j^{th} iteration of **ComputeRSymVert**(Γ, ε, h), v gives at most $\frac{\chi_{\text{out}}(w_1, \varepsilon, j-1)}{m}$ of curvature to F across the incidence described by w. So by Definition 10.4.2

$$\begin{aligned} \Pi(v, w, j) &\leq \sum_{w_1 \in C} \frac{\chi_{\text{out}}(w_1, \varepsilon, j-1)}{m} \\ &= \xi(w, C, \varepsilon, j-1), \end{aligned}$$

and therefore by Definition 10.4.2

$$\begin{aligned} \mathcal{Y}(w, j) - \mathcal{Y}(w, j-1) &= \chi_{\text{in}}(w, \varepsilon, j-1) \\ &= \max_{C' \in \mathcal{C}^w} \xi(w, C', \varepsilon, j-1) \\ &\geq \xi(w, C, \varepsilon, j-1) \\ &\geq \Pi(v, w, j). \end{aligned}$$

If there is no face F_1 with $\mathcal{M}(R_1, \varepsilon, j-1)[1] \leq 0$, as above, then $\Pi(v, w, j) = 0$, and by

Definition 10.4.2 $\xi(w, C, \varepsilon, j-1) = 0$, so by 10.4.2

$$\begin{aligned}\mathcal{Y}(w,j) - \mathcal{Y}(w,j-1) &= \max_{C' \in \mathcal{C}^w} \xi(w, C', \varepsilon, j-1) \\ &\geq \xi(w, C, \varepsilon, j-1) = 0 = \Pi(v, w, j).\end{aligned}$$

Hence as by the first paragraph $\chi(v, \Gamma) \leq \mathcal{Y}(w, 1)$, we have

$$\begin{aligned}\chi(v,\Gamma) + \sum_{j=2}^{i} \eta(w,j) &= \chi(v,\Gamma) + \sum_{j=2}^{i} \Pi(v,w,j) \\ &\leq \mathcal{Y}(w,1) + \sum_{j=2}^{i} (\mathcal{Y}(w,j) - \mathcal{Y}(w,j-1)) \\ &= \mathcal{Y}(w,i),\end{aligned}$$

as required.

Case $\mathcal{M}(R, \varepsilon, i-1)[1] \leq 0$. Let $m \in \mathbb{Z}$ be minimal subject to $\mathcal{M}(R, \varepsilon, m)[1] \leq 0$, so $1 \leq m \leq i-1$. The fact that the result is proved for $\mathcal{M}(R, \varepsilon, k)[1] > 0$ shows that

$$\chi(v,\Gamma) + \sum_{j=2}^{m} \eta(w,j) \leq \max\{\mathcal{Y}(w,m), \mathcal{B}(w,m)\}.$$

By Proposition 10.4.4 we have $\mathcal{M}(R, \varepsilon, j-1)[1] \leq 0$ for all $j \in \{m+1, \ldots, i\}$, so by Definition 9.0.6 either $\eta(w,j) = 0$, or v is interior and $\mathcal{F}_v^j \neq \emptyset$, and by Lemma 9.0.8

$$\eta(w,j) = -\Omega(F, w, j) = -\chi_{\text{out}}(w, \varepsilon, j-1).$$

Hence as by Lemma 10.4.5 $\chi_{\text{out}}(w, \varepsilon, j-1) \leq 0$, by Definition 10.4.1

$$\begin{aligned}\mathcal{Y}(w,j) - \mathcal{Y}(w,j-1) &= \mathcal{B}(w,j) - \mathcal{B}(w,j-1) \\ &= -\chi_{\text{out}}(w, \varepsilon, j-1) \\ &\geq \eta(w,j).\end{aligned}$$

So if $\mathcal{Y}(w,m) \geq \mathcal{B}(w,m)$, then

$$\begin{aligned}\chi(v,\Gamma) + \sum_{j=2}^{i} \eta(w,j) &= \chi(v,\Gamma) + \sum_{j=2}^{m} \eta(w,j) + \sum_{j=m+1}^{i} \eta(w,j) \\ &\leq \mathcal{Y}(w,m) + \sum_{j=m+1}^{i} (\mathcal{Y}(w,j) - \mathcal{Y}(w,j-1)) \\ &= \mathcal{Y}(w,i) \leq \max\{\mathcal{Y}(w,i), \mathcal{B}(w,i)\}.\end{aligned}$$

Similarly, if $\mathcal{B}(w, m) \geq \mathcal{Y}(w, m)$, then

$$\chi(v, \Gamma) + \sum_{j=2}^{i} \eta(w, j) \leq \mathcal{B}(w, i) \leq \max\{\mathcal{Y}(w, i), \mathcal{B}(w, i)\}.$$

The lemma follows. ∎

Proposition 11.2.3. *Let $F \subseteq \Gamma \in \mathcal{D}$ be an interior green face with label $R \in \mathcal{R}^{\pm 1}$, and let $1 \leq i \leq h$. Then $\mathcal{M}(R, \varepsilon, i)[1] \geq \alpha_{\Gamma}^{\varepsilon,i}(F) + \varepsilon$.*

Proof. As in the proof of Lemma 10.3.2 we can let $R' = s_1 s_2 \dots s_l$, where R' is a cyclic conjugate of R, and each s_m labels a step. Let $\mathbf{P}$ be the instantiable place on R at the beginning of R', let χ_m^i be the stepwise curvature with respect to i given to F by the step corresponding to s_m, and let l_m be the length of s_m. By Lemma 10.2.3 each step s_m corresponds to a pair $\mathbf{P}, \mathbf{Q}$ of places on R; and as by Lemma 11.2.2 the $\mathcal{Y}$ and $\mathcal{B}$ functions satisfy assumptions of Lemma 10.2.3, by 10.2.3 there exists χ_m such that $(\mathbf{Q}, l_m, \chi_m) \in \mathtt{OneStepVert}(\mathbf{P}, i)$ and $\chi_m \geq \chi_m^i$. Hence from the description of $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)$ (see Procedure 10.3.1) we have $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)[2] = \mathtt{true}$, and by Lemma 10.0.4 and Definition 10.3.3 we have

$$\begin{aligned}
\mathcal{M}(R, \varepsilon, i)[1] \geq \mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)[1] &\geq \sum_{m=1}^{l} \chi_m + (1+\varepsilon) \cdot \frac{l_m}{|R|} \\
&\geq \sum_{m=1}^{l} \chi_m^i + (1+\varepsilon) \cdot \frac{l_m}{|R|} = \left(\frac{\sum_{m=1}^{l} l_m}{|R|} + \sum_{m=1}^{l} \chi_m^i\right) + \varepsilon \cdot \frac{\sum_{m=1}^{l} l_m}{|R|} \\
&= \left(\frac{|R|}{|R|} + \sum_{m=1}^{l} \chi_m^i\right) + \varepsilon \cdot \frac{|R|}{|R|} = \left(1 + \sum_{m=1}^{l} \chi_m^i\right) + \varepsilon \geq \alpha_{\Gamma}^{\varepsilon,i}(F) + \varepsilon,
\end{aligned}$$

as required. ∎

The following theorem proves the first statement of Theorem 1.0.4.

Theorem 11.2.4. *If $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ returns $\mathtt{true}$, then $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$. Hence G is hyperbolic, and an explicit linear bound on the Dehn function of G can be calculated.*

Proof. We first show that $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on $\mathcal{P}$. Let $\Gamma \in \mathcal{D}$. We show that $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on Γ. Suppose that Γ contains no interior green faces. Then if $h \geq i > 1$, then by Algorithm 9.0.5, no curvature is redistributed through any interior vertex of Γ in iteration i. Hence $\alpha_{\Gamma}^{\varepsilon,h} = \kappa_{\Gamma}$, and $\mathbf{RSymVert}(\varepsilon, h)$ succeeds on Γ. Therefore, we can assume that Γ contains an interior green face, so let F be such face. We show that $\alpha_{\Gamma}^{\varepsilon,h}(F) \leq -\varepsilon$.

Let $R \in \mathcal{R}^{\pm 1}$ be the label of F. Since $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ returns $\mathtt{true}$, there exists $1 \leq i \leq h$ such that $\mathbf{VertexVerify}(\varepsilon, i)$ return $\mathtt{true}$. Hence by Step 3 of Procedure 10.3.5 one of the following statements holds.

1. $L_{\text{MaxCurvs}} = [\,]$.
2. $L_{\text{MaxCurvs}} \neq [\,]$ and $\{(R', \chi) \in L_{\text{MaxCurvs}} : \chi > 0\} = \emptyset$.

Suppose that Statement 1 holds. Then there is no place $\mathbf{P}$ such that **VertexVerifyAtPlace** $(\mathbf{P}, \varepsilon, i)[2] = \texttt{true}$ (see Procedure 10.3.1), hence by Lemma 10.3.2, Γ contains no interior green faces, a contradiction.

So Statement 2 holds. Therefore, for each place $\mathbf{P}$ on R we have **VertexVerifyAtPlace** $(\mathbf{P}, \varepsilon, i)[1] \leq 0$, hence by Definition 10.3.3 we have $\mathcal{M}(R, \varepsilon, i)[1] \leq 0$. By Corollary 10.3.4 we have $\mathcal{M}(R, \varepsilon, i)[2] = \texttt{true}$, so if $i < h$, then by Proposition 10.4.4 we have $\mathcal{M}(R, \varepsilon, h)[1] \leq 0$. Thus, by Proposition 11.2.3 we have $\alpha_\Gamma^{\varepsilon,h}(F) + \varepsilon \leq 0$, so $\alpha_\Gamma^{\varepsilon,h}(F) \leq -\varepsilon$, as required. Hence **RSymVert**(ε, h) succeeds on every $\Gamma \in \mathcal{D}$, so **RSymVert**(ε, h) succeeds on $\mathcal{P}$.

The final statement follows directly from Theorem 11.1.5. ∎

11.3 Complexity of VerifyHypVertex

Recall that $\mathcal{P}$ is a finite pregroup presentation such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$ (see Definition 2.3.20), and that we assume the RAM model of computation, in which the basic arithmetical operations on integers can be computed in constant time. In this section we shall show that **VerifyHypVertex**$(\mathcal{P}, \varepsilon, h)$ runs in time $O(r^9|\mathcal{R}|^9|X|^9)$, where $r := \max\{|R| : R \in \mathcal{R}\}$ is the length of the longest green relator. We shall assume that all preprocessing steps from [34, Section 7.1] have been performed on $\mathcal{P}$. This process involves comparing sub-words of cyclic conjugates of the relators, and any simplification reduces the total length of the presentation, so takes polynomial time.

Before presenting our complexity results, we shall describe the algorithm **FindCircuits** (D, k), which constructs $\mathcal{C}$ (see Definition 9.0.4) on input the enhanced vertex graph $\mathcal{E}$ and 5. It uses a modified depth-first search, where we use a trivial fact that any walk $W = (v_1, v_2, \ldots, v_m)$ is composed of the walk $W' = (v_1, v_2, \ldots, v_{m-1})$ and the edge (v_{m-1}, v_m). Hence we can generate all walks of length m from walks of length $m - 1$.

Algorithm 11.3.1. FindCircuits(D, k):

// Input: D – a directed weighted simple graph, with all weights non-negative,
// and no circuits of weight zero.
// k: positive integer.
// Output: a list L_k of all circuits C in D with $\omega(C) \leq k$ and $|C| \geq 4$.

Step 1 Initialize $L_k := [\,]$.
Step 2 Find all connected components of D.
Step 3 For each connected component C do:
- (i) Let $Q := [v_1, \ldots, v_n]$ be the vertices of C, considered as walks of length 0.
- (ii) While $Q \neq [\,]$ do:
  - (a) Let $W := (v_{i_1}, \ldots, v_{i_m})$ be a walk in Q, set $Q := Q \setminus \{W\}$.
  - (b) If $\omega(W) \leq k$ then

(A) if $(v_{i_m}, v_{i_1}) \in E(D)$ and $\omega(W)+\omega(v_{i_m}, v_{i_1}) \leq k$, then add $(v_{i_1}, \ldots, v_{i_m}, v_{i_1})$ to Q, and if $m \geq 3$, then add $(v_{i_1}, \ldots, v_{i_m})$ to L_k.

(B) Let $S := \texttt{OutNeighbours}(v_{i_m}) \setminus \{v_j \mid j \leq i_1\}$. For all $v_{i_{m+1}} \in S$, if $\omega(W) + \omega(v_{i_m}, v_{i_{m+1}}) \leq k$, then add $(v_{i_1}, \ldots, v_{i_m}, v_{i_{m+1}})$ to Q.

Step 5 Return L_k.

Let $C \in L_k$ have v_{i_1} as its lowest numbered vertex. Then $\mathbf{FindCircuits}(D, k)$ only produces cyclic rotations of C starting at v_{i_1}. Furthermore, by Remark 9.0.3 we have that $\mathbf{FindCircuits}(\mathcal{E}, 5)$ finds all interior vertices $v \in \Gamma \in \mathcal{D}$ with $\delta_G(v, \Gamma) \leq 5$ and $\delta(v, \Gamma) \geq 3$.

Lemma 11.3.2. $\mathbf{FindCircuits}(\mathcal{E}, 5)$ *runs in time* $O(r^7|\mathcal{R}|^7|X|^6 + r^6|\mathcal{R}|^6|X|^8)$*, where* $r := \max\{|R| : R \in \mathcal{R}\}$ *is the length of the longest green relator, and* $|\mathcal{C}| \leq O(r^6|\mathcal{R}|^6|X|^5 + r^5|\mathcal{R}|^5|X|^7)$.

Proof. By Definition 9.0.2 $\mathcal{E}$ contains $O(r|\mathcal{R}|)$ green and $O(|X|^2)$ red $\mathcal{E}$-vertices. Then there are $O(r^2|\mathcal{R}|^2)$ $\mathcal{E}$-edges between green $\mathcal{E}$-vertices, $O(r|\mathcal{R}||X|)$ $\mathcal{E}$-edges from a green and to a red $\mathcal{E}$-vertex, and $O(r|\mathcal{R}||X|^2)$ $\mathcal{E}$-edges from a red to a green $\mathcal{E}$-vertex

We first find, for all $v \in V(\mathcal{E})$, the set of out-neighbours of v. We store $\mathcal{E}$ as an adjacency matrix M, so this takes time $O(|V(\mathcal{E})|^2) = O(r^2|\mathcal{R}|^2 + r|\mathcal{R}||X|^2 + |X|^4)$. The time complexity of finding connected components of an undirected graph D is $O(|V(D)| + |E(D)|)$ (see [38]). Since we store $\mathcal{E}$ as M, we can construct the underlying undirected graph by defining an adjacency matrix M' with $|V(\mathcal{E})|$ rows and $|V(\mathcal{E})|$ columns, and such that $M'[i][j] = 1$ if and only if $M[i][j] = 1$ or $M[j][i] = 1$. Hence Step 2 takes time $O(|V(\mathcal{E})|^2 + |E(\mathcal{E})|) = O(r^2|\mathcal{R}|^2 + r|\mathcal{R}||X|^2 + |X|^4)$. For Step 3, we may assume that $\mathcal{E}$ is connected. Then the complexity is bounded by $O(t)$, where t is the number of (closed or open) walks in $\mathcal{E}$ of weight at most 6. Since there are no edges between red $\mathcal{E}$-vertices, every walk W in $\mathcal{E}$ with $\omega(W) \leq m$ contains at most $m+1$ green (if it starts at a green $\mathcal{E}$-vertex, then it might have $m+1$ green $\mathcal{E}$-vertices; else it has at most m such vertices) and at most $m+1$ red $\mathcal{E}$-vertices (if it starts at a red $\mathcal{E}$-vertex, then it might have $m+1$ red $\mathcal{E}$-vertices; else it has at most m such vertices), and if it attains this bound, then the colours of the vertices of W alternate. Now each green $\mathcal{E}$-vertex has at most $|X|$ red out-neighbours (since one letter is fixed), and each $\mathcal{E}$-vertex has at most $O(r|\mathcal{R}|)$ green out-neighbours (since there are at most $O(r|\mathcal{R}|)$ locations). Hence as every walk starts at a green or a red $\mathcal{E}$-vertex, we have

$$\begin{aligned} t =& O(r|\mathcal{R}| \cdot (r^6|\mathcal{R}|^6|X|^6) + |X|^2 \cdot (r^6|\mathcal{R}|^6|X|^6)) \\ =& O(r^7|\mathcal{R}|^7|X|^6 + r^6|\mathcal{R}|^6|X|^8), \end{aligned}$$

So Step 3 takes time $O(r^7|\mathcal{R}|^7|X|^6 + r^6|\mathcal{R}|^6|X|^8)$. Hence the overall complexity of **Find Circuits**$(\mathcal{E}, 5)$ is as stated.

Finally, as $|\mathcal{C}|$ is bounded by the number of (closed or open) walks in $\mathcal{E}$ of weight at most 5, applying the arguments from the previous paragraph shows that $|\mathcal{C}| \leq O(r^6|\mathcal{R}|^6|X|^5 + r^5|\mathcal{R}|^5|X|^7)$. ∎

The next theorem proves the second statement of Theorem 1.0.4. In its proof we assume that the products and inverses in the pregroup can be computed in constant time.

Theorem 11.3.3. *Let $\mathcal{P} = \langle X \mid V_P \cup \mathcal{R}\rangle$ be a finite pregroup presentation such that $\mathcal{I}(\mathcal{R}) = \mathcal{R}$, let $\varepsilon \in \mathbb{R}_{>0}$, and let $h \in \mathbb{Z}_{\geq 1}$. Then $\mathbf{VerifyHypVertex}(\mathcal{P}, \varepsilon, h)$ runs in time $O(r^9|\mathcal{R}|^9|X|^9)$, where $r := \max\{|R| : R \in \mathcal{R}\}$ is the length of the longest green relator.*

Proof. From the proof of [34, Theorem 7.22] Step 1 can be computed in time $O(|X|^5 + r^2|\mathcal{R}|^2|X|)$. Now $|V(\mathcal{E})| = O(r|\mathcal{R}| + |X|^2)$ and $|E(\mathcal{E})| = O(r^2|\mathcal{R}|^2 + r|\mathcal{R}||X|^2)$, so Step 2 takes time $O(r^2|\mathcal{R}|^2 + r|\mathcal{R}||X|^2)$. (We store $\mathcal{E}$ as an adjacency matrix and $V(\mathcal{E})$ as an indexed set.)

In Step 3, since $|V(\mathcal{E})| = O(r|\mathcal{R}|+|X|^2)$ and each $\mathcal{E}$-vertex has $\mathcal{E}$-degree at most $O(r|\mathcal{R}|+|X|)$, we deduce that $|\mathcal{W}| = O((r|\mathcal{R}| + |X|^2) \cdot (r|\mathcal{R}| + |X|)^2)$, so $O(r^3|\mathcal{R}|^3 + r^2|\mathcal{R}|^2|X|^2 + r|\mathcal{R}||X|^3 + |X|^4)$ is the time complexity of Step 3. Furthermore, $|\mathcal{W}_G| = O(r|\mathcal{R}| \cdot (r|\mathcal{R}| + |X|)^2) = O(r^3|\mathcal{R}|^3 + r^2|\mathcal{R}|^2|X| + r|\mathcal{R}||X|^2)$. (We store each $w \in \mathcal{W}_G$ as a sequence of indices of elements of $V(\mathcal{E})$).

In Step 4, for each of the $O(r^3|\mathcal{R}|^3 + r^2|\mathcal{R}|^2|X| + r|\mathcal{R}||X|^2)$ walks $w \in \mathcal{W}_G$, we can look up the corresponding walk in $\mathcal{G}$ in time $O(1)$. We then define $\mathcal{Y}(w, 1)$ and $\mathcal{B}(w, 1)$ in constant time. Hence Step 4 takes time $O(r^3|\mathcal{R}|^3 + r^2|\mathcal{R}|^2|X| + r|\mathcal{R}||X|^2)$.

In Step 5 (i) the only difference compared to the proof of [34, Theorem 7.22] when deriving the time complexity of Step 7 of $\mathbf{RSymVerify}$ is that instead of performing $O(|X|)$ calls to the $\mathbf{Vertex}$ function, we need to perform $O(|X|)$, or $O(r|\mathcal{R}|)$ calls to $\mathcal{Y}$ and $\mathcal{B}$ functions when the newly found $\mathcal{E}$-vertex is red, or green. Hence the time complexity of Step 5 (i) is $O(r^3|\mathcal{R}|^2|X|^3 \cdot (r|\mathcal{R}| + |X|)) = O(r^4|\mathcal{R}|^3|X|^3 + r^3|\mathcal{R}|^2|X|^4)$.

For each place $\mathbf{P}$, the length of the list L constructed by $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)$ (see Procedure 10.3.1) is at most $O(r|X|)$, and each item on L is considered at most r times by $\mathbf{VertexVerifyAtPlace}(\mathbf{P}, \varepsilon, i)$. Hence as $\mathbf{VertexVerify}$ (see Procedure 10.3.5) runs $\mathbf{VertexVerifyAtPlace}$ at each of the $O(r|\mathcal{R}||X|)$ places, Step 2 of $\mathbf{VertexVerify}$ takes time $O(r^3|\mathcal{R}||X|^2)$. Now there are at most $O(|\mathcal{R}|)$ elements in the list L_{MaxCurvs} and at most $O(r|\mathcal{R}|)$ locations, hence Step 5 (ii) has time complexity $O(r^3|\mathcal{R}||X|^2)$.

By Lemma 11.3.2 Step 5 (iii) takes time $O(r^7|\mathcal{R}|^7|X|^6 + r^6|\mathcal{R}|^6|X|^8)$. (We store each $C \in \mathcal{C}$ as a sequence of indices of elements of $V(\mathcal{E})$).

Step 5 (iv). We use the $\mathcal{M}$ function to check whether a given $\mathcal{E}$-vertex lies in S_{Rich}. Each such look-up takes time $O(1)$, hence Step 1 of $\mathbf{VertCurvsModify}$ (see Procedure 10.4.3) takes time $O(|\mathcal{W}_G|)$. To check whether $w \in \mathcal{W}_G$ is a sub-walk of $C \in \mathcal{C}$, we use the Knuth-Morris-Pratt ($\mathbf{KMP}$) string-searching algorithm, by noting that if S and T are two sequences of integers, then S is a (cyclic) contiguous sub-sequence of T if and only if S is a contiguous sub-sequence of T^2, where T^2 is the concatenation of T with itself. Therefore, by [2, Section 9.1], $\mathbf{KMP}(S, T^2)$ returns `true` on input S and T^2 if and only if S is a (cyclic) contiguous sub-sequence of T. Hence $\mathbf{KMP}(w, C^2)$ returns `true` if and only if w is a sub-walk of C. Now by [2, Section 9.1], $\mathbf{KMP}(w, C^2)$ runs in time $O(|w| + |C|)$. So as every circuit in $\mathcal{C}$ contains at most 11 vertices (since there are no edges between red $\mathcal{E}$-vertices), for a given $w \in \mathcal{W}_G$:

Part (ii) of **VertCurvsModify** (see Procedure 10.4.3) takes time $O(|\mathcal{C}|)$. Let $C \in \mathcal{C}$, and assume that $w \in \mathcal{W}_G$ is a sub-walk of C. As we use the $\mathcal{M}$ function to check whether a given $\mathcal{E}$-vertex lies in S_{Rich}, and C contains at most 11 vertices, we can find the number of $\mathcal{E}$-vertices of C with relators R such that $\mathcal{M}(R, \varepsilon, i)[1] > 0$ in time $O(1)$. So we can use the $\mathcal{M}$, $\mathcal{Y}$ and $\mathcal{B}$ functions to calculate $\xi(w, C, \varepsilon, \iota)$ (see Definition 10.4.2) also in time $O(1)$. Hence for a given $w \in \mathcal{W}_G$: Part (iii) of **VertCurvsModify** runs in time $O(|\mathcal{C}|)$. Thus, as by Lemma 11.3.2 we have $|\mathcal{C}| \leq O(r^6|\mathcal{R}|^6|X|^5 + r^5|\mathcal{R}|^5|X|^7)$, Step 2 of **VertCurvsModify** has time complexity

$$\begin{aligned} O(|\mathcal{W}_G||\mathcal{C}|) &= O((r|\mathcal{R}| \cdot (r|\mathcal{R}| + |X|)^2) \cdot (r^6|\mathcal{R}|^6|X|^5 + r^5|\mathcal{R}|^5|X|^7)) \\ &= O(r^9|\mathcal{R}|^9|X|^9). \end{aligned}$$

So $O(r^9|\mathcal{R}|^9|X|^9)$ is the overall complexity. ■

Proof of Theorem 1.0.4. Follows directly from Theorems 11.2.4 and 11.3.3.

Chapter 12

Implementation

We implemented **VerifyHypVertex**, in the computer algebra system MAGMA (see [6]). We used the code of the implementation **IsHyperbolic** of **RSymVerify** (see [34, Procedure 7.19]) to produce the pregroup multiplication table and to compute Step 1 of **VerifyHypVertex** (see Procedure 11.2.1), and then modify it to compute Steps 2-4 and Step 5 (i)-(ii). Parts (iii)-(iv) of Step 5 are computed by new code.

In this chapter we describe experiments with our implementation, and report the run times. Since in the first iteration **VerifyHypVertex** uses the same vertex and blob curvature bounds as **RSymVerify**, on all examples on which **IsHyperbolic** returns `true`, **VerifyHyp Vertex** succeeds in the first iteration. Hence we are particularly interested in examples on which **VerifyHypVertex** succeeds in the i^{th} iteration for some $i > 1$. We used $\varepsilon = 1/100$ and $h = 4$ in our tests. To test the correctness of our implementation, on the examples on which **VerifyHypVertex** succeeds and **RSymVerify** fails, we run KBMAG to check whether the input presentation defines a hyperbolic group: we found that on examples on which **RSymVerify** fails, KBMAG often succeeds, and **VerifyHypVertex** seems to shrink this gap by succeeding in higher iterations. We have not found a better way to test the accuracy of our implementation, and we are aware of this limitation of it.

The first example on which **IsHyperbolic** returned `fail` and **VerifyHypVertex** succeeded is the presentation of the form $\mathcal{P} = \langle a, b, c, d, e, x \mid abcde, axbxcxdxex, x^2 \rangle$, constructed as a quotient of the free group of rank 6, where $\{abcde, axbxcxdxex\}$ is the set of green relators, and the relation $x = x^\sigma$ is required in the pregroup. Using KBMAG we verified that the group defined by $\mathcal{P}$ is hyperbolic. The reason for failure of **IsHyperbolic** is that an interior green face F labelled by $abcde$ might have five interior vertices of green degree 3, each giving F curvature $-1/6$, resulting in $\kappa(F) = 1/6$. **VerifyHypVertex**, however, succeeded on $\mathcal{P}$ in the second iteration, with run time 0.13 seconds.

Next we considered presentations with randomly chosen relators. For random quotients of free groups, we choose random, freely cyclically reduced words of a given length as green relators, and leave the set of red relators to be empty. For random quotients of free products of two groups we choose random non-trivial group elements alternating between the two factors. Finally, for random quotients of three finite groups, we choose a factor at random (other than the previous factor) and then a random non-trivial element from that factor.

Interesting examples were found when the random quotients were over the free groups of ranks 10 and 20. The results are described in Table 12.1, where $i = k$ means success in iteration k, the numbers represent the number of random presentations with given properties, the penultimate entry (success for $i = 1$) is the average run time for successes in the first iteration, and the last entry is the average run time over all other cases. The table demonstrates that there exist presentations on which $\mathbf{VerifyHypVertex}$ succeeds but $\mathbf{RSymVerify}$ fails.

We also tried to run KBMAG on them and found that, for example, for random quotients over the free group of rank 20, there are 29 (out of 140) presentations $\mathcal{P}$ on which $\mathbf{VerifyHypVertex}$ succeeds, but KBMAG (with default input values) fails to precompute an automatic structure of $\mathcal{P}$, hence fails to show that $\mathcal{P}$ is hyperbolic; and 33 (out of 140) presentations on which KBMAG succeeds but $\mathbf{VerifyHypVertex}$ fails. This suggests that the two methods complement each other well. Furthermore, there was only one random quotient presented in Table 12.1 on which $\mathbf{VerifyHypVertex}$ succeeded but both $\mathbf{RSymVerify}$ and KBMAG failed, hence all but one quotients from Table 12.1 on which $\mathbf{VerifyHypVertex}$ succeeds in higher iterations have been verified to be hyperbolic by KBMAG.

We are confident that there are many additional presentations on which $\mathbf{VerifyHypVertex}$ succeeds but $\mathbf{RSymVerify}$ fails. For example, we found at least 20 such presentations over the free product of three cyclic groups of order 3. However, the presentations were very large and on each of them the procedure returned `true` after several days. The implementation in `MAGMA` would stop earlier and return `fail`, hence we have decided not to present the results here.

Table 12.1: Run times and successes of $\mathbf{VerifyHypVertex}$ on random presentations.

A free group of rank 20 with m relators of length n						
$m = 10$ & $n = 5$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	17	1	0	2	0.03 seconds	0.38 seconds
$m = 20$ & $n = 6$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	7	4	1	8	0.29 seconds	79.25 seconds
$m = 30$ & $n = 7$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	8	9	0	3	1.13 seconds	2218.12 seconds $\approx$ 37 minutes
$m = 34$ & $n = 8$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	17	3	0	0	2.62 seconds	8152.15 seconds $\approx$ 2 hours & 15 minutes
A free group of rank 10 with m relators of length n						
$m = 10$ & $n = 8$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	16	3	1	0	0.14 seconds	268 seconds $\approx$ 4 minutes & 28 seconds
$m = 10$ & $n = 20$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	20	0	0	0	1 second	NA
$m = 10$ & $n = 30$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	20	0	0	0	1 second	NA
$m = 20$ & $n = 10$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	11	9	0	0	1.07 seconds	93089 seconds $\approx$ 25 hours & 52 minutes
$m = 20$ & $n = 20$						
Total	$i = 1$	$i = 2$	$i > 2$	Fails	Success for $i = 1$	Else
20	20	0	0	0	1 second	NA

Chapter 13

Future curvature distributions schemes for showing hyperbolicity

The final chapter of Part 2 includes examples of several curvature distributions schemes that might be useful for showing hyperbolicity, but that we were unable to develop completely. The reader might notice that $\mathbf{RSymVert}$ (see Algorithm 9.0.5) redistributes curvature only through interior vertices.

Question 2: Could we come up with an (iterative) curvature distribution scheme that allows redistribution of curvature through boundary vertices?

In that case we will not obtain analogous result to Lemma 11.1.1 since the final curvature of a boundary green face might be greater than $1/2$. For this reason, we were unable to show that the curvature distribution scheme fails on diagrams with boundary length two (the analogous result to Lemma 11.1.2), and hence unable to show that its success implies non-triviality of V^σ-letters (as in Theorem 11.1.3): essential result for proving hyperbolicity.

Let $\varepsilon > 0$. Similarly as for vertices, we can redistribute curvature through red blobs: each interior green face with curvature less than $-\varepsilon$ and each boundary green face gives some of its negative curvature to edge-incident red blob B, and then B gives its negative curvature to edge-incident interior green faces with curvature greater than $-\varepsilon$. However, we expect that such scheme would not be more general then $\mathbf{RSymVert}$, and finding green faces with the aforementioned properties edge-incident with B is computationally significantly more expensive than for green faces incident with a given vertex, so we did not proceed with this idea.

We also tried to redistribute curvature across consolidated edges (see Definition 2.5.6): given a consolidated edge e common to internal green faces F_1 and F_2, if F_1 is boundary or the curvature of F_1 is less than $-\varepsilon$, and if F_2 is interior and the curvature of F_2 is greater than $-\varepsilon$, let F_1 give some of its negative curvature to F_2. The same problem occurred as for boundary vertices: since F_1 might be boundary, we were unable to ensure that a boundary green face has final curvature of at most $1/2$, hence unable to guarantee non-triviality of V^σ-letters.

We chose to present $\mathbf{RSymVert}$ in this thesis because it can be tested in polynomial time, but if one is willing to accept a higher degree polynomial cost, or perhaps an exponential cost in the length of the longest green relator, then one might come up with schemes that show hyperbolicity of a much wider classes of finite presentations.

Bibliography

[1] S. I. Adian, Algorithmic unsolvability of problems of recognition of certain properties of groups, *Doklady Akademii Nauk SSSR*, **103** 1955, 533–535.

[2] A. V. Aho, J. E. Hopcroft and J. D. Ullman, *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.

[3] 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.

[4] G. Baumslag, *Topics in combinatorial group theory*. Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 1993.

[5] W. W. Boone, The word problem. PNAS, **44** 1958, 1061–1065.

[6] W. Bosma, J. Cannon and C. Playoust, The Magma algebra system. I. The user language. *J. Symbolic Comput.*, **24** 1997, 235–265.

[7] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Springer-Verlag Berlin Heidelberg, 1999.

[8] D.J. Buckley and D.F. Holt, The Conjugacy problem in hyperbolic groups for finite lists of group elements. *Int. J. Algebra Comput.*, **23** 2013, 1127-1150.

[9] I. Bumagin, The conjugacy problem for relatively hyperbolic groups. *Algebr. Geom. Topol*, **4** 2004, 1013-1040.

[10] J. Cannon, D. Epstein and D.F. Holt, S. Levy, M. Paterson and W. Thurston, *Word Processing in Groups*. A K Peters, Natick, 1992.

[11] J. Convay, *A course in point set topology*. Undergraduate Texts in Mathematics, Springer, Cham, 2014.

[12] M. Coornaert, A. Papadopoulos and T. Delzant, *Géométrie et théorie des groupes : les groupes hyperboliques de Gromov.* Lecture Notes in Mathematics, **1441**, Springer, 1990.

[13] F. Dahmani and V. Guirardel, Foliations for solving equations in groups: free, virtually free, and hyperbolic groups. *J. Topol.*, **3(2)** 2010, 343–404.

[14] F. Dahmani and V. Guirardel, The isomorphism problem for all hyperbolic groups. *Geom. Funct. Anal.*, **21(2)** 2011, 223–300.

[15] M. Dehn, Über unendliche diskontinuierliche Gruppen. *Math. Ann.*, **71** 1911, 116-144.

[16] M. Dehn, Transformation der Kurven auf zweiseitigen Flächen. *Math. Ann.*, **72** 1912, 413-421.

[17] V. Diekert, A.J. Duncan and A. Myasnikov, Cyclic rewriting and conjugacy problems. *Groups Complex. Cryptol.*, **4** 2012, 321-355.

[18] V. Diekert, J. Laun and A. Ushakov, Efficient algorithms for highly compressed data: the word problem in Higman's group is in P. *Int. J. Algebra Comput.*, **22(8)**, 2012.

[19] V. Diekert, A.Duncan and A. G. Myasnikov, Geodesic rewriting systems and pregroups. *Combinatorial and Geometric Group Theory*, Dortmund and Carleton Conferences. Series: Trends in Mathematics O. Bogopolsk, I. Bumagin, O. Kharlampovich, E. Ventura, 2010, 55-91.

[20] E. W. Dijkstra, A note on two problems in connexion with graphs. *Numerische Mathematik*, **1**(1) 1959, 269–271.

[21] W. Dison, E. Einstein and T. R. Riley, Ackermannian integer compression and the word problem for hydra groups. *In Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science*, MFCS 2016, volume **58** of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik 2016, 30:1–30:14.

[22] B. Domanski and M. Anshell, The complexity of Dehn's algorithm for word problems in groups. *J. Algebra*, **6** 1985, 543–549.

[23] A. Eisenberg and A. Piggott, Gilman's conjecture. *J. Algebra*, **517** 2019, 167-185.

[24] D. Epstein and D.F. Holt, The linearity of conjugacy problem in word-hyperbolic groups. *Int. J. Algebra Comput.*, **16(2)** 2006, 287-306.

[25] B. Farb, Relatively hyperbolic groups. *GAFA*, **8** 1998, 810–840.

[26] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.12.1*; 2022, `https://www.gap-system.org`.

[27] S. M. Gersten and H. B. Short, Rational subgroups of biautomatic groups. *Ann. Math.*, **134** 1991, 125–158.

[28] E. Ghys and P. de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov.* Progress in Mathematics, **83**, Springer, 1990.

[29] R.H. Gilman. Computations with rational subsets of confluent groups. In: Fitch, J. (eds) EUROSAM 84 (Cambridge, 1984). *Lecture Notes in Comput. Sci.*, **174**, Springer, pages 207-212, Berlin, Heidelberg, 1984.

[30] M. Greendlinger, Dehn's algorithm for the word problem. *Comm. Pure Appl. Math.*, **13** 1959, 67–83.

[31] M. Gromov, *Hyperbolic groups.* In: Gersten S.M. (eds) Essays in Group Theory. Mathematical Sciences Research Institute Publications, **8**, Springer, pages 75-263, New York, NY, 1987.

[32] A. H. M. Hoare, Pregroups and length functions. *Math. Proc. Cambridge Philos. Soc.* **104(1)** 1988, 21-30.

[33] D. F. Holt, "KBMAG - Knuth–Bendix in monoids and automatic groups", software package (1995), available from
`http://homepages.warwick.ac.uk/~mareg/download/kbmag2/`

[34] D.F. Holt, S. Linton, M. Neunhoeffer, R. Parker, M. Pfeiffer and C. M. Roney-Dougal, Polynomial-time proofs that groups are hyperbolic. *J. Symb. Comput.* **104** 2021, 419-475.

[35] D.F. Holt, M. Lohrey and S. Schleimer, Compressed decision problems in hyperbolic groups. *In: 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)*, Berlin, Germany, 13–16 Mar 2019. Published in: LIPIcs, 2019, 34:1-34:16.

[36] D. Holt and S. Rees, The compressed word problem in relatively hyperbolic groups, *J. Algebra* (2022), https://doi.org/10.1016/j.jalgebra.2022.01.001.

[37] D. F. Holt, S. Rees and C. E. Röver, *Groups, languages and automata*. London Mathematical Society Student Texts, **88**, Cambridge University Press, 2017.

[38] Hopcroft, J. and Tarjan, R., Algorithm 447: efficient algorithms for graph manipulation. *Commun. ACM*, **16** 1973, 372–378.

[39] D. B. Johnson, Efficient algorithms for shortest paths in sparse networks. *J. ACM*, **24** 1977, 1-13.

[40] M. Lohrey, *The compressed word problem for groups*. SpringerBriefs in Mathematics, Springer, 2014.

[41] R. C. Lyndon, On Dehn's algorithm. *Math. Ann.*, **166** 1966, 208–228.

[42] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*. Classics in Mathematics, **89**, Springer-Verlag Berlin Heidelberg, 1977.

[43] J. Marshall, Computation problems in hyperbolic groups. *Int. J. Algebra Comput.*, **15(1)** 2005, 1-13.

[44] A. G. Myasnikov, A. Ushakov and D. W. Won, The word problem in the Baumslag group with a non-elementary Dehn function is polynomial time decidable. *J. Algebra*, **345**(1) 2011, 324–342.

[45] A. G. Myasnikov, A. Ushakov and D. W. Won, Power circuits, exponential algebra, and time complexity. *Int. J. Algebra Comput.*, **22**(6), 2012.

[46] P.S. Novikov, On the algorithmic unsolvability of the word problem in group theory, *Proc. Steklov Inst. Math.*, **44** 1955, 3-143.

[47] D. Osin, Relatively hyperbolic groups: Intrinsic geometry, algebraic properties, and algorithmic problems, *Mem. Amer. Math. Soc.*, **179**(843), 2006.

[48] M. O. Rabin, Recursive unsolvability of group theoretic problems. *Ann. Math.*, **67** 1958, 172–194.

[49] F. Rimlinger, Pregroups and Bass-Serre theory, *Mem. Amer. Math. Soc.*, **65**(361), 1987.

[50] P. E. Schupp, On Dehn's algorithm and the conjugacy problem. *Math. Ann.*, **178** 1968, 119–130.

[51] J.-P. Serre. *Trees*. 2nd Edition, Springer, 2003.

[52] J.R. Stallings, On torsion-free groups with infinitely many ends. *Ann. Math.*, **88** 1968, 312–334.

[53] J.R. Stallings. *Group theory and three-dimensional manifolds.* Yale University Press, New Haven, Conn., 1971.

[54] J.R. Stallings, Groups of dimension 1 are locally free. *Bull. Am. Math. Soc*, **74** 1968, 361–364.

[55] P. Wakefield, *Procedures for automatic groups*. Thesis (Ph.D.) – University of Newcastle upon Tyne, 1998, 164 pages.