ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 8 (2004) 205–275
Published: 12 February 2004

Nonpositively curved 2-complexes with isolated flats

G Christopher Hruska

Department of Mathematics, University of Chicago
5734 S University Ave, Chicago, IL 60637, USA

Email: `chruska@math.uchicago.edu`

Abstract

We introduce the class of nonpositively curved 2–complexes with the Isolated Flats Property. These 2–complexes are, in a sense, hyperbolic relative to their flats. More precisely, we show that several important properties of Gromov-hyperbolic spaces hold "relative to flats" in nonpositively curved 2–complexes with the Isolated Flats Property.

We introduce the Relatively Thin Triangle Property, which states roughly that the fat part of a geodesic triangle lies near a single flat. We also introduce the Relative Fellow Traveller Property, which states that pairs of quasigeodesics with common endpoints fellow travel relative to flats, in a suitable sense. The main result of this paper states that in the setting of CAT(0) 2–complexes, the Isolated Flats Property is equivalent to the Relatively Thin Triangle Property and is also equivalent to the Relative Fellow Traveller Property.

AMS Classification numbers Primary: 20F67

Secondary: 20F06, 57M20

Keywords: Word hyperbolic, nonpositive curvature, thin triangles, quasigeodesics, isolated flats

Proposed: Benson Farb — Received: 22 January 2003
Seconded: Walter Neumann, Martin Bridson — Revised: 12 February 2004

1 Introduction

The theory of δ–hyperbolic spaces has been enormously fruitful since it was first introduced by Gromov in his seminal article [21]. In that article, Gromov establishes two facts about δ–hyperbolic spaces which are used heavily in the proofs of many results in the theory. The first is that geodesic triangles are thin, and the second is that quasigeodesics with common endpoints (asynchronously) fellow travel.

In this article, we introduce nonpositively curved spaces with the *Isolated Flats Property*, generalizing δ–hyperbolic spaces. Spaces with the Isolated Flats Property can be studied using techniques analogous to those used in the study of δ–hyperbolic spaces. In fact, many results about δ–hyperbolic spaces have natural extensions to this new setting. The resulting theory shares much of the robust character of the δ–hyperbolic setting. In contrast, few methods are currently known for extending results from δ–hyperbolic spaces to arbitrary nonpositively curved spaces.

The main result of this article provides a starting point for the process of generalizing results from the δ–hyperbolic setting to the isolated flats setting. We introduce the Relatively Thin Triangle Property, which extends the notion of thin triangles, and the Relative Fellow Traveller Property, which generalizes the fellow travelling of quasigeodesics. The main theorem shows that these "relative" properties are each equivalent to the Isolated Flats Property in the 2–dimensional setting.

Theorem 1.1 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex. The following are equivalent.*

(1) *X has the Isolated Flats Property.*

(2) *X has the Relatively Thin Triangle Property.*

(3) *X has the Relative Fellow Traveller Property.*

The implications (2) $\Rightarrow$ (1) and (3) $\Rightarrow$ (1) are fairly straightforward. The converse implications (1) $\Rightarrow$ (2) and (1) $\Rightarrow$ (3) are more difficult and also have thus far provided more applications. Typically, the Isolated Flats Property is easier to detect than the Relatively Thin Triangle Property and the Relative Fellow Traveller Property. However, the latter two properties are quite useful in applications. For instance, one can start with an argument in the δ–hyperbolic setting which uses the thinness of triangles or the fellow travelling of quasigeodesics and try to convert it to an argument in the isolated flats setting which uses the corresponding relative properties. In Subsection 1.4, we discuss several applications of this nature.

In the course of the proof of Theorem 1.1, we prove a "quadratic divergence" theorem for geodesic rays in a 2–complex with the Isolated Flats Property (Proposition 10.3). This theorem can be interpreted as saying that given a pair of geodesic rays in such a 2–complex neither of which lingers very long near a single flat, the given rays diverge from each other at a rate which is at least quadratic. This divergence theorem is of independent interest because of its similarity to the exponential divergence theorem for δ–hyperbolic spaces proved by Cooper and Mihalik in [2, Theorem 2.19].

The exponential divergence theorem is a key ingredient of Lustig and Mihalik's proof that quasigeodesics track close to geodesics in hyperbolic spaces [2, Proposition 3.3]. Similarly, Proposition 10.3 implies that given any geodesic segment which does not linger very long near a single flat, that geodesic tracks close to any quasigeodesic connecting its endpoints (see Section 10).

It seems likely that the hypothesis that X is 2–dimensional can be dropped from Theorem 1.1. Some evidence for this conjecture is presented below in Subsection 1.3. However, the present methods are specific to the 2–dimensional setting. Our techniques depend heavily on the observation that a van Kampen diagram over a CAT(0) 2-complex is itself a CAT(0) space. In higher dimensions, completely new techniques would seem to be necessary.

1.1 The Isolated Flats Property

The Isolated Flats Property is defined precisely in Section 5. Roughly speaking, a CAT(0) space has the Isolated Flats Property if its isometrically embedded flat Euclidean subspaces diverge from each other in all directions, in the sense that their corresponding boundary spheres at infinity are disjoint. Note that the Isolated Flats Property is vacuously satisfied in any δ–hyperbolic CAT(0) space.

The prototypical example is the universal cover of a truncated finite volume hyperbolic manifold M. Such a space, called a *neutered space* or *core*, is obtained from hyperbolic space $\mathbb{H}^n$ by removing a family of disjoint open horoballs corresponding to the cusps of M. The neutered space (with the induced path metric) is a CAT(0) space whose only flat subspaces are the boundaries of the deleted horoballs, which are isolated.

The idea of studying spaces with isolated flats is implicit in work of Kapovich–Leeb [26] and of Wise [46, 47], and has also been studied by Kleiner (personal communication). In unpublished work, Wise has proved a Flat Triplane Theorem, which states that in the 2–dimensional setting the Isolated Flats Property is equivalent to an absence of isometrically embedded triplanes. A *triplane* is a space obtained by gluing three Euclidean half-planes together along their boundary lines. We provide Wise's proof of the Flat Triplane Theorem in Section 5.

Wise observed in [46, §4.0] that if X is any compact nonpositively curved 2–complex each of whose 2–cells is isometric to a regular Euclidean hexagon then the universal cover of X has the Isolated Flats Property. The reason is that a triplane cannot be built out of regular hexagons. In [4], Ballmann and Brin give explicit techniques for constructing CAT(0) hexagonal 2–complexes with arbitrary local data. The following theorem due to Moussong indicates some of the richness of this class of 2–complexes. The theorem is a special case of the main result of Moussong's thesis ([34], see also Haglund [22] and Benakli [6]).

Theorem 1.2 (Moussong) *For any simplicial graph L, there is a* CAT(0) *hexagonal 2–complex X such that the link of each vertex is isomorphic to the graph L. If L is finite, then the Coxeter group W defined by the graph L (labeled with a 3 on each edge) acts properly, cocompactly, and cellularly on X with a compact quotient.*

In the preceding theorem, we use the convention that a simplicial graph L (with all edges labeled by the number 3) defines a Coxeter system with one generator s_i of order two for each vertex v_i, and a relation $s_i s_j s_i = s_j s_i s_j$ whenever two vertices v_i and v_j are connected by an edge. The 2–complex arising in Theorem 1.2 is the Davis–Moussong geometric realization of the Coxeter system given by L.

In addition to hexagonal complexes, there are also many squared complexes with isolated flats. For instance, if L is a hyperbolic, prime, alternating link then $\pi_1(S^3 - L)$ is the fundamental group of a nonpositively curved squared 2–complex X whose universal cover has the Isolated Flats Property. This squared 2–complex was constructed by Dehn [13] and was shown to be nonpositively curved by Weinbaum [44]. The Isolated Flats Property for these 2-complexes follows from [24].

More generally, Aitchison has shown that every finite volume cusped hyperbolic 3–manifold deformation retracts onto a compact 2–complex that admits a piecewise Euclidean metric with nonpositive curvature [1]. The universal cover of this 2–complex has the Isolated Flats Property by [24]. The 2–complexes arising in Aitchison's construction typically have irregularly shaped cells.

Wise encountered spaces with isolated flats while investigating the $\mathbb{Z}\times\mathbb{Z}$ conjecture for $\mathrm{CAT}(0)$ groups. This conjecture states that, if a group acts properly and cocompactly by isometries on a $\mathrm{CAT}(0)$ space, then either the group is word hyperbolic or it contains a $\mathbb{Z}\times\mathbb{Z}$ subgroup. The $\mathbb{Z}\times\mathbb{Z}$ conjecture has been proved by Bangert–Schroeder in the case that the $\mathrm{CAT}(0)$ space is a real analytic manifold [5], however the general conjecture seems quite difficult. In fact, a theorem of Kari–Papasoglu [27] strongly indicates that the $\mathbb{Z}\times\mathbb{Z}$ conjecture may be false even for $\mathrm{CAT}(0)$ squared complexes.

Wise noticed that, in the presence of the Isolated Flats Property, the situation is much simpler, since one can then prove that all flats are periodic. Consequently the $\mathbb{Z}\times\mathbb{Z}$ conjecture is true in this setting ([46, Proposition 4.0.4], for a complete proof see [24]).

Theorem 1.3 (Wise) *Let G act properly, cocompactly, and isometrically on a $\mathrm{CAT}(0)$ space with the Isolated Flats Property. Then either G is word hyperbolic, or G contains a subgroup isomorphic to $\mathbb{Z}\times\mathbb{Z}$.*

The $\mathbb{Z}\times\mathbb{Z}$ conjecture was previously established by Ballmann–Brin in the special case that the $\mathrm{CAT}(0)$ space is a hexagonal 2-complex [4].

Sela has conjectured that the Isolated Flats Property is closely related to limit groups (or fully residually free groups), which arise in the study of the elementary theory of free groups. In particular, he has conjectured that each limit group acts properly and cocompactly on a $\mathrm{CAT}(0)$ space with the Isolated Flats Property [40].

1.2 Hyperbolicity relative to flats

The Relatively Thin Triangle Property and the Relative Fellow Traveller Property arise from an intuitive notion that spaces with the Isolated Flats Property are δ–hyperbolic "relative to flats."

For instance, a triangle is *δ–thin* if each side lies in a δ–neighborhood of the union of the other two sides. A geodesic space is *δ–hyperbolic* if every geodesic triangle in the space is δ–thin. In Section 6, we introduce the notion of a geodesic triangle being *δ–thin relative to a flat.* The idea is that each side of the triangle lies in a δ–neighborhood of the union of the other two sides and some flat, as illustrated in Figure 1. A space has the Relatively Thin Triangle Property if there is a constant δ such that every geodesic triangle is either δ–thin in the standard sense or δ–thin relative to some flat.

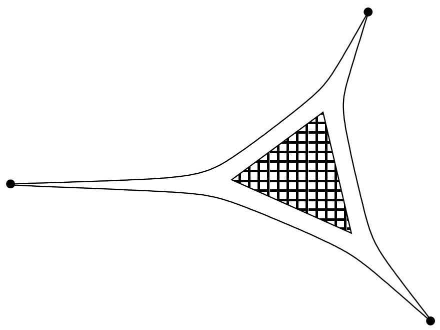

Figure 1: A triangle which is δ–thin relative to a flat

A lemma due to Morse states that in the hyperbolic plane any pair of quasigeodesics with common endpoints asynchronously fellow travel. Theorems of Gromov and Masur–Minsky ([21, Proposition 7.2.A] and [31, Lemma 7.2]) together show that this fellow traveller property is equivalent to δ–hyperbolicity. In Section 6, we introduce the more general notion of a pair of quasigeodesics that *fellow travel relative to flats.* The idea is that the curves alternate between two types of behavior: "tracking" close together and travelling near a common flat, as illustrated in Figure 2. A nonpositively curved space has the *Relative Fellow Traveller Property* if every pair of quasigeodesics with common endpoints fellow travels relative to flats.

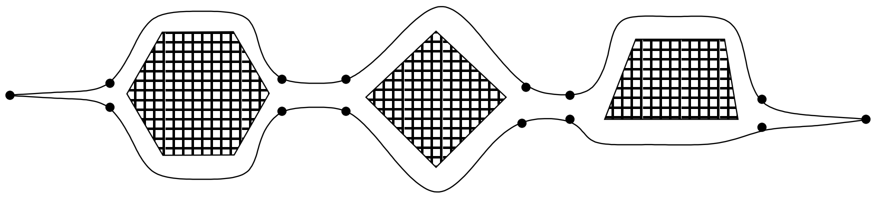

Figure 2: A pair of paths which fellow travel relative to flats

1.3 Historical background

Several results analogous to Theorem 1.1 have previously been shown for certain nonpositively curved manifolds. In [26], Kapovich and Leeb considered a class of nonpositively curved manifolds with isolated flats in which the flat subspaces are separated by regions of strict negative curvature. In that setting they proved that the Relatively Thin Triangle Property holds.

Epstein showed that an analogue of the Relative Fellow Traveller Property holds for real hyperbolic space $\mathbb{H}^n$ with a disjoint family of open horoballs removed [16, Theorem 11.3.1]. Lang generalized Epstein's result to manifolds with pinched sectional curvature $-a^2 \leq \kappa \leq -1$ for $1 \leq a < 2$ in [28]. Farb proved a result in the same spirit using a slightly different metric on the neutered space [17, Lemma 4.5]. Additionally, we remark that the Relative Fellow Traveller Property is similar to Farb's bounded coset penetration property. This similarity is exploited in [25] in the proof of the relative hyperbolicity theorem for groups acting on spaces with isolated flats.

Kapovich–Leeb and Epstein's results mentioned above provide additional examples of spaces with the Isolated Flats Property which also have either the Relatively Thin Triangle Property or the Relative Fellow Traveller Property. In light of these results, it seems likely that the hypothesis that X is 2–dimensional can be dropped from Theorem 1.1.

1.4 Applications of Theorem 1.1

Thus far, the main application of the Relatively Thin Triangle Property and the Relative Fellow Traveller Property has been to extend results from the δ–hyperbolic setting to the setting of CAT(0) spaces with the Isolated Flats Property. By way of example, we list below several immediate consequences of Theorem 1.1 whose proofs make use of either the Relatively Thin Triangle Property or the Relative Fellow Traveller Property.

Suppose a group G acts properly and cocompactly by isometries on a CAT(0) space X. In [24], the author shows that in the presence of the Isolated Flats Property and the Relative Fellow Traveller Property the boundary at infinity of X is an invariant of the group G. Recall that the *boundary* ∂X of X is the space of geodesic rays emanating from a fixed basepoint with the compact-open topology. Together with Theorem 1.1 we have the following consequence in the 2–dimensional setting.

Theorem 1.4 (Boundary is well-defined) *Suppose a group G acts properly and cocompactly by isometries on two* CAT(0) *spaces X and Y. Suppose further that X is a piecewise Euclidean 2–complex with the Isolated Flats Property. Then any equivariant quasi-isometry $X \to Y$ induces an equivariant homeomorphism $\partial X \to \partial Y$.*

The previous theorem was established in the word hyperbolic setting by Gromov in [21]. Croke and Kleiner showed in [11] that this theorem does not extend to the general CAT(0) setting by constructing two homeomorphic nonpositively curved 2–complexes whose universal covers have nonhomeomorphic boundaries. Julia Wilson has since shown that Croke and Kleiner's construction actually provides a continuous family of homeomorphic 2–complexes whose universal covers are pairwise nonhomeomorphic [45].[1]

A second consequence of Theorem 1.1 deals with quasiconvex subgroups. A subspace Z of X is *quasiconvex* if there is a constant ϵ so that every geodesic in X connecting two points of Z lies inside an ϵ–neighborhood of Z. If ρ is a proper, cocompact action of a group G by isometries on a CAT(0) space X, then a subgroup $H \leq G$ is *quasiconvex* with respect to ρ if for some $x \in X$, the orbit Hx is a quasiconvex subspace of X.

In [24] the author shows that in the presence of the Isolated Flats Property and the Relative Fellow Traveller Property the notion of a subgroup $H \leq G$ being quasiconvex does not depend on the choice of action ρ or on the choice of CAT(0) space X. Together with Theorem 1.1, we have the following consequence for 2–dimensional complexes.

Theorem 1.5 (Quasiconvexity is well-defined) *Let ρ and σ be two proper, cocompact actions of a group G by isometries on* CAT(0) *spaces X and Y. Suppose further that X is a piecewise Euclidean 2–complex with the Isolated Flats Property. For each subgroup $H \leq G$, the following are equivalent.*

(1) *H is quasiconvex with respect to ρ.*

(2) *H is quasiconvex with respect to σ.*

(3) *The inclusion $H \hookrightarrow G$ is a quasi-isometric embedding.*

The previous theorem was established by Short for word hyperbolic groups in [41]. The author shows in [24] that this result does not extend to the general CAT(0) setting.

In a subsequent article [25], the author will use Theorem 1.1 to prove the following result, which provides a precise group theoretic manifestation of the intuitive notion that a space with isolated flats is hyperbolic "relative to flats."

Theorem 1.6 (Isolated Flats $\Rightarrow$ Relatively Hyperbolic) *Suppose a group G acts properly and cocompactly by isometries on a* CAT(0) *2–complex with the Isolated Flats Property. Then G is hyperbolic relative to the collection of maximal virtually abelian subgroups of rank two.*

[1]The author has been informed that Kleiner has unpublished work related to the article [12] in which he proves Theorem 1.4 without the 2–dimensional hypothesis.

The previous theorem is stated using Gromov and Bowditch's terminology for relative hyperbolicity [21, 7]. Using the terminology of Farb [17], the conclusion is that the group in question is relatively hyperbolic with bounded coset penetration. The proof of Theorem 1.6 uses the fact that a 2–complex with isolated flats has the Relatively Thin Triangle Property.

Theorem 1.6 together with a result of Rebbechi [38] has the following immediate consequence.

Theorem 1.7 (Isolated Flats $\Rightarrow$ Biautomatic) *Suppose G acts properly and cocompactly by isometries on a* CAT(0) *2–complex with the Isolated Flats Property. Then G is biautomatic.*

It is unknown whether a group acting properly and cocompactly on an arbitrary CAT(0) space is necessarily biautomatic (or even automatic). Previously, biautomaticity has been proven only for CAT(0) complexes built from a small number of allowed shapes of polyhedral cells. For instance, Gersten–Short establish biautomaticity for 2–complexes built of squares and three shapes of triangles (specifically the 2–complexes of type $A_1 \times A_1$, A_2, B_2, and G_2) in [19, 20]. Niblo–Reeves establish biautomaticity for CAT(0) cube complexes in [35]. The main difference between Theorem 1.7 and these previous results is that Theorem 1.7 allows convex polygonal cells of arbitrary shapes.

Finally we mention that Theorem 1.6 and a result of Tukia [42] show that the Tits Alternative holds for 2–complexes with isolated flats.

Theorem 1.8 (Isolated Flats $\Rightarrow$ Tits Alternative) *Suppose G acts properly and cocompactly on a* CAT(0) *2–complex with the Isolated Flats Property. Then G satisfies the Tits Alternative. In other words, every subgroup of G is either virtually abelian or contains a free subgroup of rank two.*

Again it is unknown whether the previous theorem remains true if the Isolated Flats Property is dropped from the hypothesis.

If the hypothesis that X is 2–dimensional can be dropped from Theorem 1.1, as conjectured in the previous subsection, then the results listed above will apply to any CAT(0) space with the Isolated Flats Property.

1.5 Summary of the sections

We begin with a few sections of background material. In Section 2 we define CAT(0) spaces, and review several of their important properties. In Section 3, we discuss piecewise Euclidean complexes and their relation to nonpositive curvature by way of the Link Condition. In Section 4, we review definitions and basic results about diagrams, reduced maps, and the Combinatorial Gauss–Bonnet theorem, following the development of McCammond–Wise [32].

In Section 5, we give a definition of the Isolated Flats Property which is catered to the 2–dimensional setting. We also give Wise's proof of the Flat Triplane Theorem, which has not previously appeared in the literature. In Section 6, we state the Relatively Thin Triangle Property and the Relative Fellow Traveller Property. In the 2–dimensional setting, we use the Flat Triplane Theorem to show that each of these properties implies the Isolated Flats Property, thus establishing the implications (2) $\Rightarrow$ (1) and (3) $\Rightarrow$ (1) of Theorem 1.1.

The reverse implications (1) $\Rightarrow$ (2) and (1) $\Rightarrow$ (3) are more difficult and require a deeper analysis of the structure of disc diagrams corresponding to 2–complexes with isolated flats. In Section 7, we introduce the notion of a diagram which is *ruffled* along a boundary path. Such a diagram has the property that every point of the given boundary path is either close to a vertex with strictly negative curvature or close to some other part of the boundary. This notion generalizes the fact that in the negatively curved setting the fat part of any reduced disc diagram is filled with vertices at which the curvature is negative. We use the Combinatorial Gauss–Bonnet Theorem to relate ruffled diagrams to the usual notion of δ–thin triangles.

In Section 8, we study *preflats* in reduced disc diagrams. Preflats in a diagram $D \to X$ correspond to flats in X. We prove two results about preflats, which play a key role in the proof of Theorem 1.1. The first is Proposition 8.2, which states roughly that if a geodesic segment occurs in the boundary path of a nonpositively curved disc diagram, then either the diagram is ruffled along the geodesic or the geodesic is close to a large preflat. The second result is Proposition 8.7 which states that in the presence of the Isolated Flats Property, preflats in disc diagrams are surrounded by ruffles.

In Section 9, we use the ideas developed in the previous two sections to prove the implication (1) $\Rightarrow$ (2) of Theorem 1.1. The idea of the proof is that a triangular diagram either contains a large preflat surrounded by ruffles, or it is ruffled along one side. In the first case, the triangle is thin relative to a flat, and in the second case, the triangle is thin in the standard sense.

In Section 10 we turn our attention to the remaining implication (1) $\Rightarrow$ (3) of Theorem 1.1. In this section we prove the "quadratic divergence" theorem mentioned above. Cooper, Lustig, and Mihalik use an analogous exponential divergence theorem in their proof that quasigeodesics track close to geodesics in a δ–hyperbolic space [2, Proposition 3.3]. The main result of Section 10 is a similar fellow travelling result for 2–complexes with isolated flats in the presence of ruffles.

In order to prove (1) $\Rightarrow$ (3), it will be useful to understand the structure of the convex hull of a union of preflats in a disc diagram. In Section 11, we examine in detail the convex hull of a union of two preflats and also the convex hull of the union of one preflat and a point. We show that in the presence of the Isolated Flats Property, the convex hull of the two objects in question is essentially just the union of those objects together with a path of shortest length connecting them.

Section 12 contains the main part of the proof of (1) $\Rightarrow$ (3). In the isolated flats setting, we study disc diagrams with a geodesic segment along the boundary. We consider the convex hull of the union of that geodesic and all the preflats that come near the geodesic. We show that the boundary of this convex hull contains a path which fellow travels the geodesic relative to flats. We also show that this convex hull is surrounded by ruffles in the diagram.

In Section 13 we complete the proof of Theorem 1.1. We consider a disc diagram whose boundary consists of a geodesic and quasigeodesic. Once one removes the convex hull constructed in the previous section, one obtains a diagram which satisfies the hypothesis of the ruffled fellow travelling result from Section 10. We conclude that any geodesic and quasigeodesic with common endpoints fellow travel relative to flats.

Finally, we establish the Relative Fellow Traveller Property by considering the general case of a pair of quasigeodesics with common endpoints. We derive the general case from the special case by comparing each quasigeodesic with the geodesic connecting its endpoints and piecing together the two sequences of flats associated to each such pair.

1.6 Acknowledgements

The results in this article were originally published as part of my PhD dissertation at Cornell University. That dissertation was prepared under the guidance of Daniel Wise and Karen Vogtmann. I would like to thank Karen for supporting my desire to work with Dani on this project, although he was a postdoc at the time. I would also like to thank Dani for pointing me in the direction of such a fruitful research area, and for devoting such a large amount of time to encouraging this research even after he moved from Ithaca to Boston and eventually to Montreal. On many occasions Dani and his family graciously hosted me in their home so that I could talk mathematics with Dani for long hours at a stretch. I also benefited from conversations about this research with Martin Bridson, Marshall Cohen, Ilya Kapovich, Bruce Kleiner, Jon McCammond, John Meier and surely others I have forgotten to mention. Additionally, I would like to thank the referee for numerous helpful suggestions that I hope have improved the exposition of this article.

This research was partially supported by a grant from the National Science Foundation.

2 CAT(0) spaces

In this section, we review some definitions and several well-known results about CAT(0) spaces. We refer the reader to Bridson–Haefliger [10] for proofs of the results listed in this section, as well as for historical information about the origins of these ideas. We give precise theorem numbers for the corresponding statements in [10] when they do not have a common name that can be found in the index of that book.

A *geodesic* in a metric space X is an isometric embedding of an interval of $\mathbb{R}$ into X. A metric space X is *geodesic* if every pair of points in X is connected by at least one geodesic. If p and q are points in a geodesic space, we use the notation $[p,q]$ to denote a particular choice of geodesic segment connecting the points p and q.

Definition 2.1 CAT(0) Let X be a geodesic metric space. A *geodesic triangle* $\Delta(p,q,r)$ in X is the union of three geodesic segments $[p,q]$, $[q,r]$, and $[p,r]$ in X. A *comparison triangle* $\overline{\Delta} = \overline{\Delta}(\overline{p},\overline{q},\overline{r})$ for Δ is a triangle in the Euclidean plane $\mathbb{E}^2$ such that

$$d(p,q) = d(\overline{p},\overline{q}), \quad d(q,r) = d(\overline{q},\overline{r}), \quad \text{and} \quad d(p,r) = d(\overline{p},\overline{r}),$$

as illustrated in Figure 3. A point $\overline{x} \in [\overline{p},\overline{q}]$ is a *comparison point* for $x \in [p,q]$ provided that $d(\overline{p},\overline{x}) = d(p,x)$. Comparison points are defined similarly for points on the other sides $[q,r]$ and $[p,r]$.

Let Δ be a triangle in X, and let $\overline{\Delta}$ be a comparison triangle for Δ. We say that Δ satisfies the CAT(0) *inequality* if for all points $x,y \in \Delta$ and comparison points $\overline{x},\overline{y} \in \overline{\Delta}$ we have

$$d(x,y) \le d(\overline{x},\overline{y}).$$

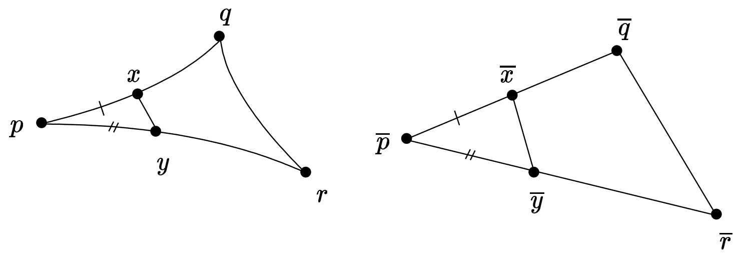

Figure 3: A comparison triangle and a pair of comparison points for the CAT(0) inequality

If every geodesic triangle in X satisfies the CAT(0) inequality, then X is called a CAT(0) *space*. The space X has *nonpositive curvature* if for each $x \in X$ there is an $\epsilon > 0$ so that $B(x,\epsilon)$ is CAT(0).

Definition 2.2 *Angles* Let X be a nonpositively curved space, and let $\gamma\colon [0,s] \to X$ and $\gamma'\colon [0,s'] \to X$ be geodesic segments with $p = \gamma(0) = \gamma'(0)$. The *angle* $\angle(\gamma,\gamma')$ between γ and γ' is defined by the formula

$$\angle(\gamma,\gamma') = \lim_{t\to 0} 2\arcsin \frac{1}{2t} d\bigl(\gamma(t),\gamma'(t)\bigr).$$

The above limit exists by [10, Proposition II.3.1]. When X is CAT(0), then geodesic segments are uniquely determined by their endpoints. In this case, if $x \neq p$ and $y \neq p$, then the angle between the segments $[p,x]$ and $[p,y]$ will frequently be denoted $\angle_p(x,y)$.

Note that the preceding definition agrees with the usual Euclidean angle between geodesic segments in the Euclidean plane $\mathbb{E}^2$.

Definition 2.3 *Comparison angles* Let $\Delta(p,q,r)$ be a geodesic triangle in a CAT(0) space and let $\overline{\Delta}(\overline{p},\overline{q},\overline{r})$ be a comparison triangle for Δ. The interior angle of $\overline{\Delta}$ at $\overline{p}$, denoted $\overline{\angle}_p(q,r)$, is called the *comparison angle* between q and r at p.

Theorem 2.4 ([10], II.1.7(4)) *Let X be a CAT(0) space. The angle between two sides of a geodesic triangle Δ in X with distinct vertices is no greater than the corresponding comparison angle in $\overline{\Delta}$.*

Theorem 2.5 (Convexity of the CAT(0) metric) *Let γ and γ' be two geodesic segments in a CAT(0) space X, each parametrized from 0 to 1 proportional to arclength. Then for each $t \in [0,1]$ we have*

$$d\bigl(\gamma(t),\gamma'(t)\bigr) \leq (1-t)\, d\bigl(\gamma(0),\gamma'(0)\bigr) + t\, d\bigl(\gamma(1),\gamma'(1)\bigr).$$

Theorem 2.6 (Cartan–Hadamard Theorem) *Let X be a complete metric space which is connected and simply connected. If X has nonpositive curvature, then X is a* $\mathrm{CAT}(0)$ *space.*

A map $f\colon Y \to X$ between metric spaces is a *local isometry* if for every $y \in Y$ there is an $\epsilon > 0$ such that the restriction of f to $B(y,\epsilon)$ is an isometry onto its image. A *local geodesic* in a metric space X is a local isometry from an interval of $\mathbb{R}$ into X.

The following two results are consequences of the Cartan–Hadamard Theorem.

Theorem 2.7 ([10], II.4.13) *Fix two points x_0 and x_1 in a complete, nonpositively curved metric space X. Then any homotopy class of paths between x_0 and x_1 in X is represented by a unique local geodesic.*

Theorem 2.8 ([10], II.4.14) *Let $f\colon Y \to X$ be a local isometry between two complete, connected nonpositively curved spaces. Then any lift of f to a map $\tilde{f}\colon \tilde{Y} \to \tilde{X}$ between their universal covers is an isometric embedding.*

3 Piecewise Euclidean 2–complexes

In this section, we collect definitions and basic results about piecewise Euclidean 2–complexes. For a more thorough development of these ideas see [10].

Definition 3.1 A *convex Euclidean polyhedron* P is the convex hull of a finite set of points in Euclidean space $\mathbb{E}^n$. The *dimension* of P is the minimal dimension of an affine subspace $E \subseteq \mathbb{E}^n$ containing P. If P is contained in one of the closed half-spaces determined by some hyperplane $H \subset \mathbb{E}^n$, then $H \cap P$ is called a *face* of P. We also consider P itself to be a face of P. Note that a face F of P is itself a convex polyhedron. The *vertices* of P are its 0–dimensional faces.

Definition 3.2 A *piecewise Euclidean complex* is a complex X formed from a disjoint union of convex Euclidean polyhedra by gluing isometric faces by isometries. A *metric graph* is a one dimensional piecewise Euclidean complex.

Note that a piecewise Euclidean complex X has a natural cell structure whose cells are the polyhedra of X. A piecewise Euclidean complex has a natural pseudometric where the distance between two points is the infimum of the lengths of paths connecting them. According to the following theorem due to Bridson, in many complexes of interest this infimum is actually realized by a geodesic path.

Theorem 3.3 [8] *If a connected, piecewise Euclidean complex X has only finitely many isometry types of cells, then X is a complete geodesic metric space.*

Definition 3.4 *Link* Let X be a piecewise Euclidean 2–complex, and let v be a vertex of X. The *link* $\mathrm{Lk}(v,X)$ of v in X is the space of all germs of geodesics in X based at v. The *radial projection* onto $\mathrm{Lk}(v,X)$ is the function

$$\pi\colon X-\{v\}\to \mathrm{Lk}(v,X)$$

sending each point x to the germ of the geodesic $[v,x]$. If X has only finitely many isometry types of cells, then by [10, Theorem I.7.39], the function d given by

$$d\big(\pi(x),\pi(y)\big)=\angle_v(x,y)$$

is a metric on $\mathrm{Lk}(v,X)$. Under this metric each link $\mathrm{Lk}(v,X)$ has a natural structure as a metric graph with one edge for each corner of a 2–cell incident at v. The length of each edge is equal to the angle of the corresponding corner.

Definition 3.5 A *locally geodesic loop of length L* in a metric space X is a local isometry $C\to X$ where C is a metric circle of length L.

Definition 3.6 *Link condition* A piecewise Euclidean 2–complex X satisfies the *link condition* if for every vertex $v\in X^{(0)}$, every locally geodesic loop $C\to \mathrm{Lk}(v,X)$ has length at least 2π.

The following theorem is due to Gromov [21, §4.2]. Ballmann provided a proof in the locally finite case [3], and Bridson in the general case [8].

Theorem 3.7 *Let X be a piecewise Euclidean 2–complex with finitely many isometry types of cells. Then X has nonpositive curvature if and only if it satisfies the link condition.*

Corollary 3.8 *Let X be a nonpositively curved, piecewise Euclidean 2–complex with finitely many isometry types of cells. Then any subcomplex Y of X is also nonpositively curved (using the induced path metric on Y).*

Corollary 3.9 *Let X be a piecewise Euclidean* $\mathrm{CAT}(0)$ *2–complex, and let Y be a connected subcomplex. Suppose that, for every vertex $v\in Y^{(0)}$, the link $\mathrm{Lk}(v,Y)$ has diameter at most π. Then the inclusion $Y\hookrightarrow X$ is an isometric embedding (using the induced path metric on Y).*

4 Diagrams and curvature

In this section we review some background on diagrams and curvature, including the Combinatorial Gauss–Bonnet Theorem. Diagrams are planar 2–complexes which often allow one to more clearly visualize and reason about combinatorial homotopies of paths. The early sections of [32] contain a concise and elegant development of many of the ideas in this section.

Definition 4.1 *Combinatorial maps and complexes* A map $Y \to X$ between CW complexes is *combinatorial* if its restriction to each open cell of Y is a homeomorphism onto an open cell of X. A CW complex is itself *combinatorial* if the attaching map of each cell is a combinatorial map (possibly after subdividing the given cell structures).

Notice that a convex Euclidean polyhedron is a combinatorial complex. Consequently a piecewise Euclidean complex is also combinatorial.

Convention 4.2 A piecewise Euclidean 2–complex is *cocompact* if it is cocompact as a combinatorial 2–complex. More precisely, the 2–complex X is cocompact if the group of combinatorial isometries acts on X with a compact quotient.

Note that cocompactness of a piecewise Euclidean 2–complex is a stronger notion than cocompactness of the underlying metric space. For instance, the Euclidean plane equipped with a Penrose tiling is not a cocompact 2–complex even though the underlying metric space $\mathbb{E}^2$ is cocompact.

Definition 4.3 A *diagram* D is a finite connected combinatorial 2–complex equipped with a fixed combinatorial embedding in the 2–sphere which misses at least one point. A simply connected diagram is a *disc diagram.*

A *subdiagram* E of D is a subcomplex such that the embedding $E \hookrightarrow S^2$ factors as $E \hookrightarrow D \hookrightarrow S^2$ where $D \hookrightarrow S^2$ is the given embedding of D into the sphere. Figure 4 shows a disc diagram and a subdiagram which is not simply connected.

Disc diagrams are often called *van Kampen diagrams*, particularly in the context of group presentations. Some authors reserve the term *disc diagram* for a diagram which is homeomorphic to a disc and use *singular disc diagram* for a contractible diagram. The author prefers the present terminology, as contractible diagrams occur with much higher frequency in applications and seem to be more natural objects than diagrams with the topology of a disc.

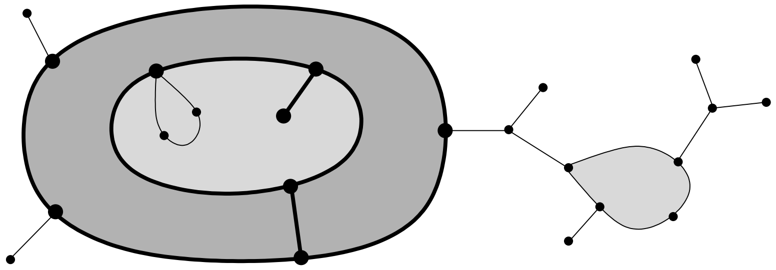

Figure 4: A disc diagram and a darkened subdiagram which is not simply connected

In practice, we often suppress mention of the embedding of a diagram into the sphere. However the boundary cycles of a diagram (defined below) will typically be different for different embeddings of the diagram into the sphere. The issue is that the diagram may have cut points.

Definition 4.4 *Boundary cycles* Let $D \hookrightarrow S^2$ be a diagram with its given combinatorial embedding in the sphere. Since D is connected, its complement in S^2 has a finite number of components $R_1, \ldots, R_k$, each homeomorphic to an open disc. Without loss of generality, we may assume that S^2 has a cell structure with a single 2–cell e_i for each region R_i. A choice of orientation for S^2 determines a collection of *boundary cycles* of D so that the boundary cycle corresponding to the region R_i is the attaching map of the 2–cell e_i. If D contains at least one 1–cell, then each boundary cycle is a combinatorial map $C_i \to D$ where C_i is a subdivided circle.

Definition 4.5 *Reduced maps* Let D be a disc diagram and $\phi\colon D \to X$ a combinatorial map. A pair of (not necessarily distinct) 2–cells C_1 and C_2 in D which meet along a 1–cell e is a *cancelable pair* with respect to ϕ if the boundary cycles of C_1 and C_2 beginning with e (in the same direction) are not identical in D, but are sent to identical paths in X by ϕ. Figure 5 shows a cancelable pair of 2–cells. The map ϕ is *reduced* if D does not contain a cancelable pair of 2–cells. We will often refer to a reduced map $D \to X$ with domain a disc diagram as a *reduced disc diagram.*

The following theorem concerning the existence of reduced disc diagrams was discovered by van Kampen [43] and independently by Lyndon [29]. For a complete proof, see for instance [32, Lemma 2.17].

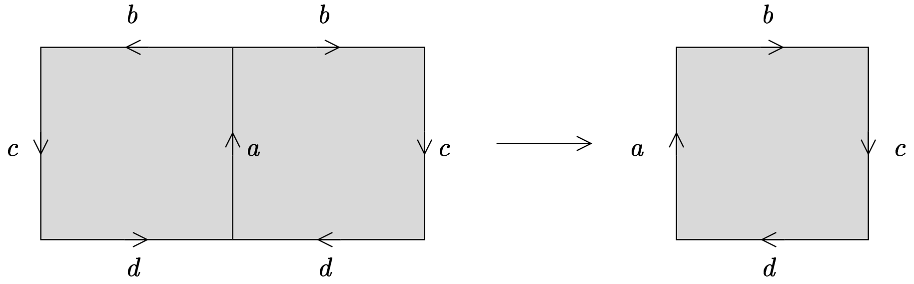

Figure 5: A combinatorial map which is not reduced because it contains a cancelable pair of 2–cells

Theorem 4.6 *If X is a combinatorial 2–complex and $P \to X$ is a combinatorial closed path which is nullhomotopic in X, then there exists a reduced disc diagram $D \to X$ so that $P \to X$ is the boundary path of D.*

We will often refer to the reduced disc diagram $D \to X$ obtained in the preceding theorem as a *reduced disc diagram for P*.

In the following theorem we show that nonpositive curvature pulls back under reduced disc diagrams. We emphasize that the metric constructed on the disc diagram is unrelated to any metric inherited from the given embedding of the diagram in the plane $\mathbb{R}^2$ (see Figure 6).

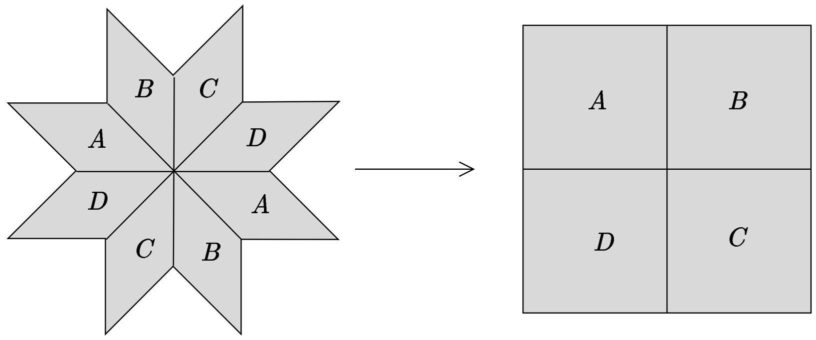

Figure 6: A reduced map from a disc diagram to a nonpositively curved squared 2–complex. The diagram on the left inherits a CAT(0) metric in which the eight 2–cells are Euclidean squares.

Proposition 4.7 *Let $\phi\colon D \to X$ be a reduced disc diagram, such that X is a nonpositively curved piecewise Euclidean 2–complex. Then D is a* CAT(0) *space.*

Proof Since D is connected and simply connected, we just need to show that D is nonpositively curved. Since ϕ is combinatorial, D has a natural piecewise Euclidean structure obtained by pulling back the Euclidean metrics on each cell of the image.

We will now show that for each vertex $v \in D^{(0)}$, the induced map

$$\phi_*\colon\ \mathrm{Lk}(v, D) \to \mathrm{Lk}\big(\phi(v), X\big)$$

is a local isometry. If not, then ϕ_* *folds* a pair of edges together. In other words, there is a pair of distinct oriented edges e_1 and e_2 in $\mathrm{Lk}(v, D)$ with the same initial vertex which map to the same oriented edge under ϕ_*. But such a pair of edges in $\mathrm{Lk}(v, D)$ corresponds to two distinct corners of of a cancelable pair of 2–cells of D, contradicting the fact that ϕ is reduced. Hence for every vertex $v \in D^{(0)}$ the map ϕ_* is a local isometry.

Thus if $C \to \mathrm{Lk}(v, D)$ is a locally geodesic loop, then the composition $C \to \mathrm{Lk}(v, D) \to \mathrm{Lk}\big(\phi(v), X\big)$ is a locally geodesic loop of the same length. Hence C has length at least 2π, and we see that D has nonpositive curvature. □

Definition 4.8 *Curvature* Let D be a piecewise Euclidean diagram. The *curvature at a vertex* v of D, denoted $\kappa(v)$, is defined by the formula

$$\kappa(v) = 2\pi - \pi\,\chi\big(\mathrm{Lk}(v, D)\big) \ - \sum_{e \in \mathrm{Edges}(\mathrm{Lk}(v,D))} \|e\| \,,$$

where χ denotes Euler characteristic and $\|e\|$ denotes the length of the edge e. We occasionally use the notation $\kappa_D(v)$ to emphasize the specific diagram D in which the curvature is measured. An alternate way to interpret curvature is to formally place an angle of size π at each corner of the complement of D in S^2 as illustrated in Figure 7. Then the curvature at v is equal to 2π minus the angle sum of all the corners at v including these "exterior" corners.

We now have two distinct notions of nonpositive curvature in a piecewise Euclidean diagram. The following lemma gives a simple correspondence between these two ideas.

Lemma 4.9 *Let D be a piecewise Euclidean diagram. Then D has nonpositive curvature if and only if the curvature at each interior vertex of D is nonpositive.*

Proof Notice that the link of an interior vertex is always a circle. The proof now follows immediately from the definition of curvature at a vertex and the Link Condition for D. □

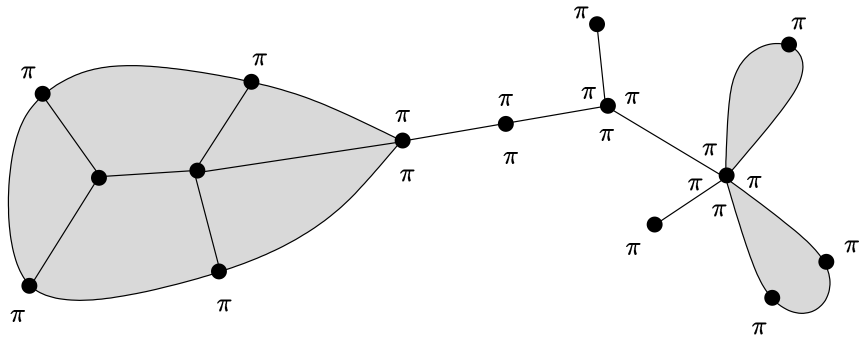

Figure 7: Formally place an angle of size π at each corner of $S^2 - D$. The curvature at a vertex is 2π minus the angle sum of all corners at v including these "exterior" corners.

The Combinatorial Gauss–Bonnet Theorem is based on Lyndon's Curvature Formula [30], and was proved for piecewise Euclidean disc diagrams independently by Gersten and Pride [18, 37]. It has since been generalized to arbitrary combinatorial 2–complexes by McCammond–Wise [32].

Theorem 4.10 (Combinatorial Gauss–Bonnet) *Let D be a piecewise Euclidean diagram. Then*

$$\sum_{v \in D^{(0)}} \kappa(v) = 2\pi\, \chi(D).$$

5 Isolated Flats and the Flat Triplane Theorem

Definition 5.1 A CAT(0) 2–complex X has the *Isolated Flats Property* if there is a function $\psi\colon \mathbb{R}_+ \to \mathbb{R}_+$ such that for every pair of distinct flat planes $F_1 \neq F_2$ in X and for every $k \geq 0$, the intersection $\mathcal{N}_k(F_1) \cap \mathcal{N}_k(F_2)$ of k–neighborhoods of F_1 and F_2 has diameter at most $\psi(k)$.

The definition of the Isolated Flats Property given above is catered to the two dimensional setting, and as such, is slightly simpler than the definition used in [24] and [25]. This simplification is due to the fact that several phenomema present in higher dimensions are absent in the 2–dimensional setting. For instance, in general a space with the Isolated Flats Property could contain subspaces of the form $\mathbb{E}^k \times K$ where K is a nontrivial compact set.

We note the following immediate consequence of the Isolated Flats Property, which will be useful in the sequel.

Proposition 5.2 *Let X be a CAT(0) space with the Isolated Flats Property. Then for every $k \geq 0$, each flat disc in X of radius at least $\psi(k)$ lies in a k–neighborhood of at most one flat.*

The Flat Plane Theorem states that a proper, cocompact CAT(0) space is δ–hyperbolic if and only if it does not contain an isometrically embedded flat plane.[2] Wise has proved an analogous result for two-complexes with the Isolated Flats Property. His Flat Triplane Theorem shows that in the two-dimensional case, the Isolated Flats Property is equivalent to an absence of triplanes. Since this useful result has not appeared in the literature, we provide Wise's proof in this section for the sake of completeness. For related results regarding the presence of triplanes in nonpositively curved 2–complexes, see [46, 47].

Definition 5.3 A *triplane* is the space formed from three closed half-planes by gluing their boundary lines (by isometries) to a common line. Notice that a triplane is a piecewise Euclidean CAT(0) 2–complex.

Fix a basepoint x on the singular line of a triplane T. Let T_r denote the closed ball in T of radius r around X. Notice that T_r looks like three half discs of radius r with their straight boundary sides glued together.

Theorem 5.4 (Flat Triplane Theorem) *Let X be a proper, cocompact piecewise Euclidean 2–complex. The following are equivalent:*

(1) *X has the Isolated Flats Property.*
(2) *There is a universal bound L on the diameter of the intersection of any two distinct flat planes in X.*
(3) *X does not contain an isometrically embedded triplane.*

The proof of Theorem 5.4 uses the following variant of the Arzelà–Ascoli theorem, which is proved in [10, Lemma II.9.34]. The statement given here is slightly stronger than the one given by Bridson and Haefliger. This additional strength follows immediately from their proof.

Lemma 5.5 *Let Y be a separable metric space with basepoint y_0, and let X be a proper, cocompact metric space. If for each $n \in \mathbb{N}$ there is an isometric embedding $\phi_n \colon B(y_0, n) \hookrightarrow X$ then there exists an isometric embedding $\phi \colon Y \hookrightarrow X$.*

[2] The Flat Plane Theorem was proved for Riemannian manifolds by Eberlein [14]. Gromov stated the theorem for general CAT(0) spaces in [21, §4.1]. Heber and Bridson independently provided proofs in this general setting [23, 9].

Furthermore, if we assume the existence of a compact set K such that for each n the point $\phi_n(y_0)$ lies in K, then we may take ϕ to be a pointwise limit of a subsequence $\{\phi_{n_i}\}$ of the original sequence of embeddings.

The proof of Theorem 5.4 also uses the Flat Strip Theorem [10], which states the following.

Theorem 5.6 (Flat Strip Theorem) *Let X be a CAT(0) space, and let $\gamma\colon \mathbb{R} \to X$ and $\gamma'\colon \mathbb{R} \to X$ be geodesic lines in X. If there is a constant K such that $d\big(\gamma(t), \gamma'(t)\big) \le K$ for all $t \in \mathbb{R}$, then the convex hull of $\gamma(\mathbb{R}) \cup \gamma'(\mathbb{R})$ is isometric to a flat strip $\mathbb{R} \times [0, L] \subset \mathbb{E}^2$.*

Proof of Theorem 5.4 The implications (1) $\Rightarrow$ (2) $\Rightarrow$ (3) are immediate. We now show that (3) implies (2). Suppose for each $k > 0$ there is a pair of distinct flat planes $E_k \neq F_k$ whose intersection I_k has diameter at least k. Since E_k and F_k are convex, their intersection I_k is isometric to a convex polygonal region (possibly unbounded) in the Euclidean plane. Since X is a cocompact 2–complex, it has only finitely many isometry types of cells (see Convention 4.2). So there are finitely many possible turning angles on the boundary of I_k. Thus there is a positive number θ (independent of k) so that every positive turning angle on ∂I_k is at least θ. It follows that there are at most $2\pi/\theta$ turns on the boundary of I_k. Thus for each $m > 0$, there is an $n > 0$ so that ∂I_n contains a straight line segment γ_m of length m. This segment lies in two distinct flat discs D_m and D'_m of diameter m which intersect in a half disc. Notice that $D_m \cup D'_m$ with its induced path metric is isometric to the space T_m (in the notation of Definition 5.3). But for every point $v \in T_m$, the link $\mathrm{Lk}(v, T_m)$ has diameter π, so $T_m \hookrightarrow X$ is an isometric embedding by Corollary 3.9. Since X contains an isometrically embedded copy of T_m for each $m > 0$, it follows from Lemma 5.5 that X contains an isometrically embedded triplane.

We have shown that (2) and (3) are equivalent. Now we show that (2) and (3) together imply (1). Suppose X does not have the Isolated Flats Property. Then there exists a constant $R > 0$ such that for every $k \in \mathbb{N}$ there are distinct flats E_k and F_k in X with geodesic segments $\gamma_k \subset E_k$ and $\gamma'_k \subset F_k$, each with length greater than k such that the Hausdorff distance between their images is at most R.

Consider the sequence of segments γ_k. By composing with a suitable isometry, we may assume that the midpoint x_k of each segment γ_k lies inside a given compact set K. Then Lemma 5.5 implies that there is a subsequence γ_{k_i} that converges pointwise to a geodesic γ. Applying Lemma 5.5 to the sequence of embedded flat planes E_{k_i} gives a further subsequence of integers $\{m_i\}$ such that the flat planes E_{m_i} converge to a flat plane E containing the geodesic γ. Letting x'_k denote the midpoint of γ'_k, we see that $d(x_k, x'_k) \leq 3R$, since the endpoints of γ_k and γ'_k are within $3R$ of each other. So x'_k lies in the closure K' of a $3R$–neighborhood of K, which is compact, since X is proper.

Continuing as above, we eventually obtain a sequence of integers $\{n_i\}$ such that E_{n_i} and F_{n_i} converge to isometrically embedded flat planes E and F containing geodesic lines γ and γ'. Furthermore, we may assume that the Hausdorff distance between the images of γ and γ' is at most R. If E and F are actually the same flat plane, then the intersections $E_{n_i} \cap F_{n_i}$ must have arbitrarily large diameter, contradicting (2). So we may assume that $E \neq F$.

By Theorem 5.6, we conclude that the convex hull S of $\operatorname{Im}(\gamma) \cup \operatorname{Im}(\gamma')$ is isometric to a flat strip $\mathbb{R} \times [0, L] \subset \mathbb{E}^2$. Since E, F, and S are convex, any intersection of them is also convex. It now follows that $E \cap S$ and $F \cap S$ are substrips $S_E = \mathbb{R} \times [0, s]$ and $S_F = \mathbb{R} \times [t, L]$ respectively. We now have two cases, depending on whether these substrips intersect.

Case 1 Suppose $s < t$. Then E and F are disjoint planes connected by a flat strip. Letting $U = E \cup F \cup S$, we see that for every point $v \in U$ the link $\operatorname{Lk}(v, U)$ has diameter π. So the inclusion $U \hookrightarrow X$ is an isometric embedding by Corollary 3.9. Since U contains an isometrically embedded triplane, we are done.

Case 2 Suppose $s \geq t$. Then $E \cap F$ is a closed convex set of the plane containing a line. It follows that $E \cap F$ is either a half-plane, or a (possibly degenerate) strip $\mathbb{R} \times [0, s - t]$. Let $U = E \cup F$. As in the previous case, U is isometrically embedded in X since every link $\operatorname{Lk}(v, U)$ has diameter π. Since U contains a triplane, we are done. □

6 Hyperbolicity relative to flats

In this section we give precise definitions of the Relative Fellow Traveller Property and the Relatively Thin Triangle Property. We show that in the 2–dimensional setting, each of these properties implies the Isolated Flats Property, establishing the implications (2) $\Rightarrow$ (1) and (3) $\Rightarrow$ (1) of Theorem 1.1.

Definition 6.1 *Relatively Thin Triangle Property* A geodesic triangle in a space is *δ–thin relative to the flat F* if each side of the triangle lies in a δ–neighborhood of the other two sides and the flat F, as illustrated in Figure 8.

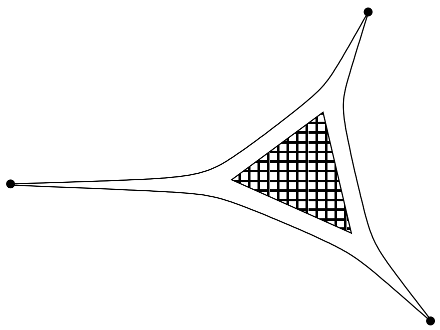

Figure 8: A triangle which is δ–thin relative to a flat

A space X has the *Relatively Thin Triangle Property* if there is a constant δ so that each triangle in X is either δ–thin in the usual sense or δ–thin relative to some flat.

It is not hard to see that any 2–complex with the Relatively Thin Triangle Property must also have the Isolated Flats Property.

Theorem 6.2 (Relatively Thin Triangle Property $\Rightarrow$ Isolated Flats Property) *Let X be a proper, cocompact, piecewise Euclidean* CAT(0) *2–complex satisfying the Relatively Thin Triangle Property. Then X also has the Isolated Flats Property.*

Proof Assume by way of contradiction that X does not have the Isolated Flats Property. Then by Theorem 5.4, X contains an isometrically embedded triplane T. Let us parametrize T as

$$T = \Big\{ (x,y,i) \;\Big|\; x,y \in \mathbb{R};\; y \geq 0; \text{ and } i \in \{1,2,3\} \Big\} \Big/ (x,0,j) \sim (x,0,k).$$

Now consider the triangle Δ_n with vertices

$$a = (0,n,1) \qquad b = (-2n,n,2) \qquad c = (2n,n,3)$$

illustrated in Figure 9. Notice that for any fixed δ, each side of Δ_n only intersects the δ–neighborhood of the union of the other two sides near the corners of the triangle. This intersection consists of two segments whose lengths are bounded by a constant which does not depend on the value of n. So for large values of n, the triangle Δ_n is not δ–thin.

But if n is sufficiently large, each side of Δ_n lies in a δ–neighborhood of at most one flat in X. Since the three sides of Δ_n lie in three distinct flats, it is

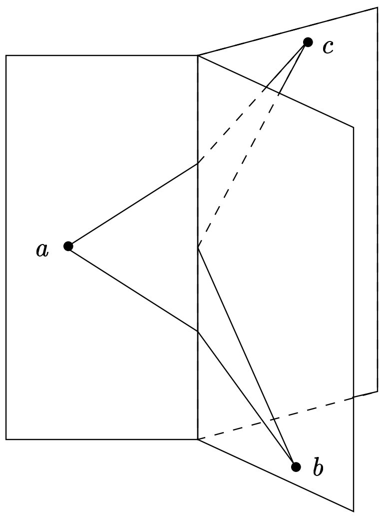

Figure 9: A fat triangle which does not lie close to a single flat. The two paths from b to c are quasigeodesics which are only close together near their endpoints.

clear that Δ_n is not δ–thin relative to any single flat in X. So X does not have the Relatively Thin Triangle Property. □

Definition 6.3 A (λ, ϵ)*–quasigeodesic* in a metric space X is a function $\alpha\colon [a,b] \to X$ for some real interval $[a,b]$ satisfying

$$\frac{1}{\lambda}\, d(s,t) - \epsilon \leq d\big(\alpha(s), \alpha(t)\big) \leq \lambda\, d(s,t) + \epsilon$$

for all $s,t \in [a,b]$. A map $\alpha\colon [a,b] \to X$ is a *quasigeodesic* if there exist constants λ and ϵ such that α is a (λ,ϵ)–quasigeodesic.

In hyperbolic geometry, quasigeodesics with common endpoints satisfy an asynchronous fellow traveller property. This fact was established for the hyperbolic plane by Morse [33] and for $\mathbb{H}^n$ by Efromovich–Tihomirova [15]. Gromov generalized the fellow traveller property to the following result about δ–hyperbolic spaces.

Theorem 6.4 ([21], Proposition 7.2.A) *Let α and β be a pair of (λ,ϵ)–quasigeodesics with common endpoints in a δ–hyperbolic space X. Then the Hausdorff distance between $\mathrm{Im}(\alpha)$ and $\mathrm{Im}(\beta)$ is at most L, where L depends only on the constants δ, λ, and ϵ.*

The concept of a pair of paths which *fellow travel relative to flats* generalizes the asynchronous fellow travelling described above. Roughly speaking, the idea is that the two paths alternate between tracking close together and travelling near a common flat as illustrated in Figure 10.

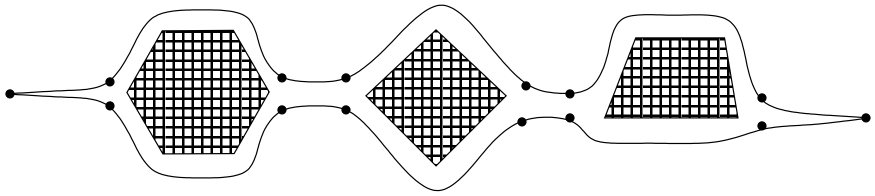

Figure 10: A pair of paths which fellow travel relative to flats

Definition 6.5 *Fellow travelling relative to flats* A pair of paths

$$\alpha\colon [0,a] \to X \quad \text{and} \quad \alpha'\colon [0,a'] \to X$$

in a CAT(0) space L*–fellow travel relative to a sequence of flats* $(F_1,\dots,F_n)$ if there are partitions

$$0 = t_0 \le s_0 \le t_1 \le s_1 \le \cdots \le t_n \le s_n = a$$

and

$$0 = t'_0 \le s'_0 \le t'_1 \le s'_1 \le \cdots \le t'_n \le s'_n = a'$$

so that for $0 \le i \le n$ the Hausdorff distance between the sets $\alpha\bigl([t_i,s_i]\bigr)$ and $\alpha'\bigl([t'_i,s'_i]\bigr)$ is at most L, while for $1 \le i \le n$ the sets $\alpha\bigl([s_{i-1},t_i]\bigr)$ and $\alpha'\bigl([s'_{i-1},t'_i]\bigr)$ lie in an L–neighborhood of the flat F_i.

We will frequently say that paths L*–fellow travel relative to flats* if they L–fellow travel relative to some sequence of flats.

Definition 6.6 A CAT(0) space X satisfies the *Relative Fellow Traveller Property* if for each choice of constants λ and ϵ there is a constant $L = L(\lambda,\epsilon,X)$ such that (λ,ϵ)–quasigeodesics in X with common endpoints L–fellow travel relative to flats.

In the 2–dimensional setting, it is easy to show that the Relative Fellow Traveller Property implies the Isolated Flats Property using the same counterexample as the proof of Theorem 6.2. The converse requires substantially more work and occupies most of the present article.

Theorem 6.7 (Relative Fellow Traveller Property $\Rightarrow$ Isolated Flats Property) *Let X be a proper, cocompact piecewise Euclidean* CAT(0) 2*–complex. If X has the Relative Fellow Traveller Property, then X also has the Isolated Flats Property.*

Proof Suppose X contains a triplane T. As in the proof of Theorem 6.2, consider the triangle Δ_n with vertices

$$a = (0, n, 1) \qquad b = (-2n, n, 2) \qquad c = (2n, n, 3)$$

illustrated in Figure 9. Notice that the union $[b, a] \cup [a, c]$ is a quasigeodesic when parametrized by arclength. Furthermore, the associated constants of this quasigeodesic are independent of n.

Also note that for each L, the geodesic $[b, c]$ intersects the L–neighborhood of the quasigeodesic only near their common endpoints. But $[b, a] \cup [a, c]$ does not lie in a C–neighborhood of any flat for any constant C which does not depend on n. So X does not have the Relative Fellow Traveller Property. $\square$

7 Ruffled boundaries and thin triangles

In the previous section, we established the implications (2) $\Rightarrow$ (1) and (3) $\Rightarrow$ (1) of Theorem 1.1. The remainder of this article is devoted to proving (1) $\Rightarrow$ (2) and (1) $\Rightarrow$ (3), which we will establish in Theorems 9.1 and 13.1 respectively.

Our objective for the next two sections is to establish some diagrammatic tools that will be useful in the proofs of these two theorems. In this section we introduce the notion of a diagram which is *ruffled* along a certain part of its boundary. This notion generalizes the fact that in the δ–hyperbolic setting, every point of a piecewise Euclidean diagram is either close to a point of negative curvature or lies near the boundary of the diagram. As a consequence, every δ–hyperbolic diagram is "ruffled" throughout its interior.

When studying the Isolated Flats Property, one frequently encounters diagrams with the property that negative curvature is distributed evenly along a certain portion of the boundary of the diagram except in places where the diagram is very thin. Such a diagram is *ruffled* along that boundary path.

Diagrams which are ruffled along part of their boundary have many features in common with δ–hyperbolic diagrams. For instance, we will see in Proposition 7.7 that a triangular diagram which is ruffled along one side must be δ–thin. In this section, we also prove results which describe various ways that ruffles can be inherited by a subdiagram.

Definition 7.1 *Ruffled* Let D be a nonpositively curved piecewise Euclidean diagram with boundary cycles $C_1, \dots, C_n$. Suppose C_1 is a concatenation of two paths α and β. Then the pair (D, α) is (R,θ)*–ruffled* for positive constants R and θ provided that for each point $\alpha(t)$, the open ball $B = B\big(\alpha(t), R\big)$ in D satisfies one of the following two properties:

(R-1) B contains a vertex v with curvature $\kappa_D(v) \leq -\theta$, or

(R-2) B intersects the image of at least one of the curves $C_2, \dots, C_n$, or β.

The pair (D, C_1) is (R,θ)*–ruffled* provided that for each point $C_1(t)$, the open ball $B = B\big(C_1(t), R\big)$ in D satisfied one of the following two properties:

(R$'$-1) B contains a vertex v with curvature $\kappa_D(v) \leq -\theta$, or

(R$'$-2) B intersects the image of at least one of the curves $C_2, \dots, C_n$.

Remark 7.2 Note that if (D, α) or (D, C_1) is (R_0, θ_0)–ruffled, then it is also (R, θ)–ruffled for any $R \geq R_0$ and any positive $\theta \leq \theta_0$.

The following lemma follows immediately from the definition of ruffled.

Lemma 7.3 (Subdiagrams) *Consider a subdiagram D' of a nonpositively curved piecewise Euclidean diagram D. Suppose there are boundary cycles C and C' of D and D' respectively so that C is a concatenation $\alpha\beta$ and C' is a concatenation $\alpha\beta'$, as illustrated in Figure 11*(a)*. If the pair (D, α) is (R, θ)–ruffled then so is the pair (D', α).*

Similarly suppose C is a boundary cycle of both D and D' as in Figure 11(b)*. If (D, C) is (R, θ)–ruffled, then (D', C) is also (R, θ)–ruffled.*

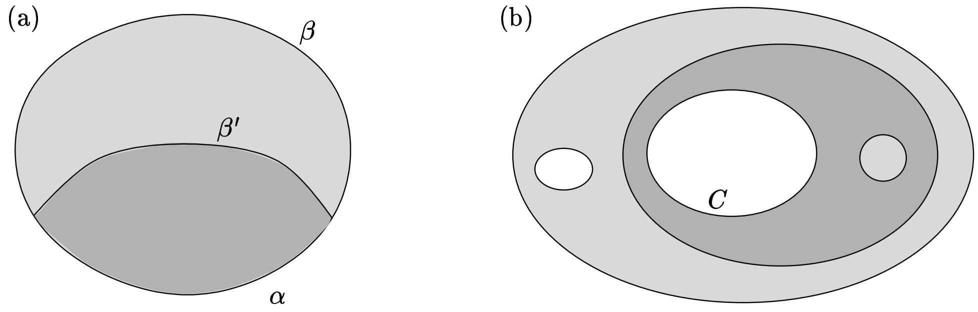

Figure 11: (a) A diagram and a darkened subdiagram which share a boundary arc α (b) A diagram and a darkened subdiagram which share an entire boundary cycle

The following lemma says that under many circumstances a concatenation of ruffled boundary arcs is also a ruffled boundary arc. In general the constants associated to the ruffles depend on the number of segments being concatenated. The issue is that one way for a boundary arc to be ruffled is for the arc to be very short. If each arc in the concatenation is known to be sufficiently long, then ruffling constants are obtained which do not depend on the number of concatenated segments.

Lemma 7.4 (Concatenations) *Let C be a boundary cycle of a nonpositively curved piecewise Euclidean diagram D. Suppose C is a concatenation $\alpha_1 \cdots \alpha_n \beta$ such that (D, α_i) is (R, θ)–ruffled for each i and $\alpha = \alpha_1 \cdots \alpha_n$ is a local geodesic in D.*

Suppose further that any (global) geodesic in D connecting two points of $\mathrm{Im}(\alpha)$ lies inside $\mathrm{Im}(C)$. Then

1. *(D, α) is (nR, θ)–ruffled.*
2. *If each α_i has length at least $2R$, then (D, α) is $(2R, \theta)$–ruffled.*
3. *If each α_i has length at least $2R$ and β has image a single point, then (D, C) is $(2R, \theta)$–ruffled.*

The requirement that any geodesic in D connecting two points of $\mathrm{Im}(\alpha)$ lies inside $\mathrm{Im}(C)$ is satisfied if D is a disc diagram and α is a geodesic boundary component, as shown in Figure 12(a), or alternately if D is an annular diagram and C is a locally geodesic boundary component as in Figure 12(b).

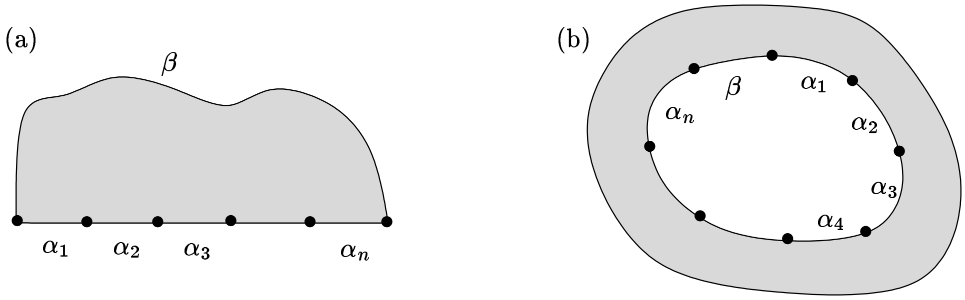

Figure 12: (a) A disc diagram such that the concatenation $\alpha_1 \cdots \alpha_n$ is a geodesic boundary arc (b) An annular diagram such that the concatenation $\alpha_1 \cdots \alpha_n \beta$ is a locally geodesic boundary component

Proof If every α_i has length less than $2R$, then every point of $\mathrm{Im}(\alpha)$ is within a distance nR of an endpoint of α, and hence (D, α) is (nR, θ)–ruffled. On the other hand, suppose some α_i has length at least $2R$, and x_i is a point on $\mathrm{Im}(\alpha_i)$ at least a distance R from both ends of α_i. Then the ball $B = B(x_i, R)$ intersects $\mathrm{Im}(\alpha)$ only in the interior of $\mathrm{Im}(\alpha_i)$. If B does not contain a vertex with curvature $\leq -\theta$, then B intersects $\mathrm{Im}(\beta)$ or some other boundary component of D. Every point of $\mathrm{Im}(\alpha)$ is within a distance $(n-1)R$ of such a point x_i. Therefore (1) holds.

If each α_i has length at least $2R$, then each contains a point x_i at least a distance R from both ends of $\mathrm{Im}(\alpha)$. Therefore, every point of $\mathrm{Im}(\alpha)$ is within a distance R of such a point x_i. Claim (2) now follows immediately. Furthermore, if β has image a single point, then that point is also within a distance R of some x_i, giving (3). □

Definition 7.5 *Triangular* A piecewise Euclidean CAT(0) disc diagram Δ is *triangular* if its boundary cycle is a concatenation of three geodesics α, β, and γ. These geodesics are the *sides* of the triangle, and their endpoints are its *corners*.

Lemma 7.6 *Let Δ be a geodesic triangle in a* CAT(0) *space X. Suppose one of the sides of Δ lies in a δ–neighborhood of the union of the other two sides. Then Δ is 2δ–thin.*

Proof Let a, b, and c be points lying on sides α, β, and γ respectively so that $d(a,b)$ and $d(a,c)$ are each less than δ. Then b lies in a 2δ–neighborhood of both α and γ. Similarly, c lies in a 2δ–neighborhood of both α and β. The result now follows from the convexity of the CAT(0) metric in X. □

The following proposition says that a triangular diagram which is ruffled along one side is δ–thin for some δ. The proof uses the Combinatorial Gauss–Bonnet Theorem to bound the total amount of negative curvature inside the triangle.

Proposition 7.7 (Thin triangles) *Let Δ be a piecewise Euclidean* CAT(0) *triangular diagram with sides α, β, and γ. Suppose R and θ are positive constants so that (Δ, α) is (R, θ)–ruffled. Then there is a constant $\delta = \delta(R, \theta)$ so that Δ is δ–thin.*

Proof By Theorem 4.10, the sum of the curvatures at all the vertices of Δ is exactly 2π. But by Lemma 4.9 the only positive curvature in Δ occurs at its three corners since Δ is nonpositively curved. Furthermore, the curvature at each corner is at most π. So the sum of all positive curvatures in Δ is at most 3π. It follows that the sum of all negative curvatures in Δ has magnitude at most π.

If α lies in an R–neighborhood of the union of β and γ, then by Lemma 7.6 setting $\delta = 2R$ completes the proof. Otherwise, let $[x,y]$ be the maximal subsegment of α which lies outside the open R neighborhood of $\beta \cup \gamma$. Setting $m = \lfloor d(x,y)/2R \rfloor$, we can choose points $x_1, \dots, x_m$ on $[x,y]$ so that the open balls $B_i = B(x_i, R)$ are disjoint. But each ball contains negative curvature with magnitude at least θ. So there can be at most π/θ such balls. Therefore $d(x,y)$ is bounded in terms of R and θ. The conclusion now follows from Lemma 7.6. □

Figure 13 illustrates a diagram D which is ruffled along a boundary arc α and a subdiagram D' so that α, β, and α' form a geodesic triangle. If the side β is sufficiently short, then the following lemma states that D' is ruffled along α'. As in the proof of Proposition 7.7, the reason is that the Combinatorial Gauss–Bonnet Theorem provides a bound on the total amount of negative curvature inside the triangle.

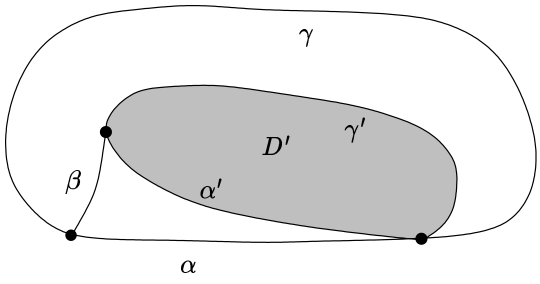

Figure 13: If (D,α) is ruffled and β is short, then (D',α') is also ruffled.

Lemma 7.8 *Given positive constants R, θ, and δ, there is a constant $R' = R'(R,\theta,\delta)$ such that the following property holds. Let D be a piecewise Euclidean* CAT(0) *disc diagram whose boundary cycle is a concatenation $\alpha\gamma$. Let D' and Δ be subdiagrams of D with disjoint interiors such that Δ is triangular with sides α, β, and α' and such that the boundary cycle of D' is a concatenation $\alpha'\gamma'$ as illustrated in Figure 13. If (D,α) is (R,θ)–ruffled and β has length at most δ, then (D',α') is $(R',\theta/2)$–ruffled.*

Proof If α has length at most δ, then we are done since α' then has length at most 2δ, and hence (D',α') is trivially (δ,θ)–ruffled. Thus we may assume that the length of α is more than δ.

Let x be the common endpoint of α and α'. For

$$1 \leq i \leq m = \lfloor (\ell(\alpha) - \delta)/2R \rfloor,$$

let x_i be the point on $\mathrm{Im}(\alpha)$ at a distance $(2i-1)R$ from x. Then the open balls $B_i = B(x_i, R)$ are pairwise disjoint. Since $\ell(\alpha) \geq 2mR + \delta$ and $\ell(\beta) \leq \delta$, it follows that the image of β does not intersect any B_i. Therefore, if B_i intersects $\mathrm{Im}(\gamma)$, it must also intersect $\mathrm{Im}(\gamma')$. Since (D, α) is (R,θ)–ruffled, if B_i does not intersect $\mathrm{Im}(\gamma')$, then B_i contains a vertex y_i with $\kappa_D(y_i) \leq -\theta$. If y_i is in the interior of D', then $\kappa_{D'}(y_i) = \kappa_D(y_i)$. Similarly, if $y_i \in \Delta - \mathrm{Im}(\alpha')$, then $\kappa_\Delta(y_i) = \kappa_D(y_i)$. If y_i lies on the image of α', it is a vertex in both Δ and D' satisfying

$$\kappa_D(y_i) = \kappa_\Delta(y_i) + \kappa_{D'}(y_i).$$

Call B_i *defiant* if it contains a vertex $y_i \in \Delta$ with $\kappa_\Delta(y_i) \leq -\theta/2$. As in the proof of Proposition 7.7, the triangular diagram Δ contains negative curvature with total magnitude at most π. So at most $2\pi/\theta$ of the B_i's are defiant. Each nondefiant ball B_i either intersects $\mathrm{Im}(\gamma')$ or contains a vertex $y_i \in D'$ with $\kappa_{D'}(y_i) \leq -\theta/2$.

Now fix any point $p' \in \mathrm{Im}(\alpha')$. Since the endpoints of α and α' are at most δ apart, p' is within a distance δ of some point $p \in \mathrm{Im}(\alpha)$. But p is within $\delta + 3R$ of the center x_i of some ball B_i. This x_i is within a distance $4R\pi/\theta$ of the center x_j of some nondefiant ball B_j. Thus p' is within a total distance $2\delta + 4R + (4R\pi/\theta)$ of either the image of γ' or some vertex $y_j \in D'$ with $\kappa_{D'}(y_j) \leq -\theta/2$. In other words, (D', α') is $(R', \theta/2)$–ruffled, where $R' = 2\delta + 4R + (4R\pi/\theta)$. □

8 Preflats in reduced disc diagrams

Let $\phi\colon D \to X$ be a reduced disc diagram where X is a piecewise Euclidean CAT(0) 2–complex. When X contains flat planes, we will frequently be interested in examining the preimages under ϕ of these various planes. These preimages of flats are especially useful in the presence of the Isolated Flats Property since distinct flats in X have only a small intersection.

Each component of such a preimage is composed of various topological discs and isolated edges glued together in a treelike fashion as illustrated in Figure 14. It is frequently more convenient to deal with these discs individually rather than considering the entire preimage as a unit. We call each such disc a *preflat.* The following definition makes this notion more precise.

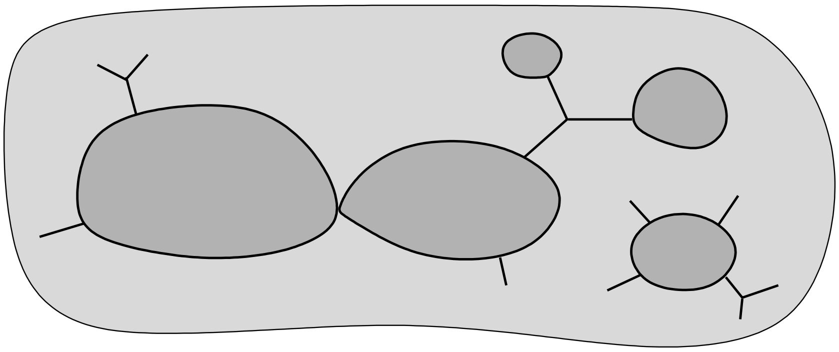

Figure 14: A reduced disc diagram $\phi\colon D \to X$ and the preimage under ϕ of some flat plane F in X. The preimage shown is disconnected and contains exactly five preflats.

Definition 8.1 *Preflat* Let $\phi\colon D \to X$ be a reduced disc diagram where X is a CAT(0) 2–complex with the Isolated Flats Property. A *preflat* P is the closure of a connected component of the interior of $\phi^{-1}(E)$ for some flat plane E in X.

In this section we see that both the presence and the absence of preflats in disc diagrams provides a source of ruffles. In Proposition 8.2 we show that if a geodesic segment occurs along the boundary of a reduced disc diagram $D \to X$, then either the diagram is ruffled along the geodesic or some preflat in X comes close to the geodesic. So in this case, the absence of preflats provides ruffles. In contrast, we see in Proposition 8.7 that in the presence of the Isolated Flats Property preflats themselves are surrounded by ruffles.

Proposition 8.2 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex. For each $M_0 > 0$, there are positive constants $R = R(M_0, X)$ and $\theta_0 = \theta_0(M_0, X)$ so that the following property holds for any positive $\theta \le \theta_0$.*

Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram where $\widehat{X}$ is a subdivision of X. Suppose the boundary cycle of D decomposes as a concatenation of two paths γ and α such that $\phi \circ \gamma$ is a geodesic in $\widehat{X}$. Then either

(1) *the pair (D, γ) is $(8R, \theta)$–ruffled, or*
(2) *some preflat P with inscribed radius at least M_0 intersects the R–neighborhood of $\mathrm{Im}(\gamma)$, while $\phi(P)$ is not contained in the $(R + M_0)$–neighborhood of $\mathrm{Im}(\phi \circ \gamma)$.*

The proof of the preceding proposition uses several lemmas which we will state and prove before giving the proof of the proposition. Before we get to these lemmas, let us reflect for a moment on why it is useful to allow subdivisions of X in the statement above. The complex X has a fixed cell structure such that the group of combinatorial isometries of X acts cocompactly. In practice, we need to consider disc diagrams arising from nullhomotopies of piecewise geodesic loops in X which do not respect the given cell structure of X. In general, a subdivision of X is required to ensure that a particular piecewise geodesic loop lies inside the 1–skeleton of X. We emphasize that the constants R and θ_0 obtained in the proposition depend only on the cell structure of X and not on the particular subdivision $\widehat{X}$ in question.

Our first lemma does not use the Isolated Flats Property or even nonpositive curvature.

Lemma 8.3 *Let X be a proper, cocompact, piecewise Euclidean 2–complex. For every $m > 0$ there is an $n = n(m, X)$ so that if B is a flat disc in X of radius n, then the central subdisc B' of radius m lies in some flat plane of X.*

Proof Our argument is a combinatorial version of the diagonal argument in the proof of Lemma 5.5. Suppose, by way of contradiction, that there is a number $m > 0$ so that for each integer $n \geq m$ there is a flat disc B_n of radius n whose central subdisc B'_n of radius m does not lie entirely within any flat plane of X. Consider the sequence of embeddings $B_n \hookrightarrow X$. By composing each embedding with a suitable isometry of X, we may assume that the center p_n of every disc B_n lies in a common finite subcomplex K.

We proceed by induction to construct a sequence

$$\varnothing = C_0 \subseteq C_1 \subseteq C_2 \subseteq \cdots$$

of subcomplexes of X and a sequence

$$\mathbb{N} = \mathcal{S}_0 \supseteq \mathcal{S}_1 \supseteq \mathcal{S}_2 \supseteq \cdots$$

of infinite sets such that C_i lies inside each disc of the family $\{\, B_n \mid n \in \mathcal{S}_i \,\}$ and such that the union $\bigcup_i C_i$ is a flat plane. Suppose $\mathcal{S}_{i-1}$ and C_{i-1} have already been constructed. Since X has finitely many isometry types of 2–cells, the cells of X have a universally bounded radius. It follows that for each sufficiently large $n \in \mathcal{S}_{i-1}$ there is a 2–cell inside $B_n - C_{i-1}$ which is a minimal distance R_{i-1} from the center p_n. The chosen 2–cells all lie in a bounded (hence compact) closed neighborhood of K. So there is an infinite set $\mathcal{S}_i \subseteq \mathcal{S}_{i-1}$ such that for $n \in \mathcal{S}_i$ the chosen 2–cells coincide in a single cell e_i. Let $C_i = C_{i-1} \cup \overline{e}_i$.

Notice that for $n \in \mathcal{S}_i$ the disc of radius R_i in B_n centered at p_n is contained in C_i. Furthermore, as i tends to infinity, the radii R_i become arbitrarily large.

It follows that the union $\bigcup_i C_i$ is a flat plane E. If we choose $n_i \in \mathcal{S}_i$, then B_{n_i} is a disc of radius n_i whose central subdisc of radius R_i lies in the flat plane E, which is a contradiction. □

The following lemma gives a lower bound on the size of certain turning angles occuring when a geodesic passes through a vertex.

Lemma 8.4 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex. For each $R \geq 0$, there is a positive constant $\theta = \theta(R)$ with the following property.*

Let x, y, and z be distinct vertices of X such that y lies on the geodesic $[x, z]$, and suppose that each of the distances $d(x, y)$ and $d(y, z)$ are less than R. Let $\pi\colon X - \{y\} \to \mathrm{Lk}(y, X)$ be the radial projection onto $\mathrm{Lk}(y, X)$. Then any immersed path in $\mathrm{Lk}(y, X)$ connecting $\pi(x)$ and $\pi(z)$ which has length strictly greater than π actually has length at least $\pi + \theta$.

Proof Since X is proper and cocompact we can bring the points x, y, and z into a fixed compact set $K = K(R)$ by applying a combinatorial isometry. But K contains only finitely many vertices. So for each R, only finitely many choices of x, y, and z concern us. Consequently it suffices to prove the result for a fixed triple of points x, y, and z.

Since $\mathrm{Lk}(y, X)$ is a finite metric graph, it contains only finitely many immersed paths which connect $\pi(x)$ and $\pi(z)$ and which have length less than 2π. Thus the paths with length greater than π have a minimum length as desired. □

Recall that the Monodromy Theorem from complex analysis states that an analytic continuation of an analytic function along a curve depends only on the homotopy class of the curve in question [39]. The next lemma is an easy consequence of the Monodromy Theorem.

Lemma 8.5 *Let S be a simply connected nonpositively curved surface which is locally isometric to the Euclidean plane $\mathbb{E}^2$. Then S admits a local isometry into $\mathbb{E}^2$.*

Proof By hypothesis, each point $x \in S$ is contained in an open disc D_x which admits an isometric embedding into $\mathbb{E}^2$. Fix a specific embedding of a single disc D_p. For each overlapping disc D_q, we can compose the given embedding $D_q \hookrightarrow \mathbb{E}^2$ with an isometry of $\mathbb{E}^2$ to make it agree with the chosen embedding

$D_p \hookrightarrow \mathbb{E}^2$ on the overlap $D_p \cap D_q$. This procedure provides an analytic function $D_p \cup D_q \hookrightarrow \mathbb{E}^2$ extending the original map $D_p \hookrightarrow \mathbb{E}^2$. Continuing in this fashion, a map $S \to \mathbb{E}^2$ can be defined by taking an analytic extension of the original embedding $D_p \hookrightarrow \mathbb{E}^2$ along various paths. By the Monodromy Theorem, these continuations fit together to give a globally defined local isometry $S \to \mathbb{E}^2$. □

The following lemma roughly states that if a geodesic occurs as part of the boundary of a disc diagram, then either the diagram is ruffled along the geodesic or some point of the geodesic has a neighborhood isometric to a Euclidean half-disc.

Lemma 8.6 *Let X be a piecewise Euclidean* CAT(0) *2–complex, and let $\phi\colon D \to X$ be a reduced disc diagram. Suppose the boundary path of D is a concatenation $\alpha\gamma$ where γ is a geodesic in D. Let $p = \gamma(t)$ be an arbitrary point in the image of γ. If we let $B = B(p, R)$, then either*

(1) *B contains a vertex with negative curvature, or*

(2) *B intersects the image of α, or*

(3) *B is isometric to a Euclidean half-disc.*

Proof Suppose conditions (1) and (2) fail for B. Then the curvature at every vertex of B is zero. So B is locally flat. In other words, each vertex of $\mathring{B}$ has a neighborhood isometric to a Euclidean disc, while each vertex of ∂B has a neighborhood isometric to a Euclidean half-disc.

Since B is a metric ball in a CAT(0) space, B is a convex subspace of D. So B is simply connected, and hence admits a local isometry into the Euclidean plane by Lemma 8.5. Since $\mathring{B}$ is convex, this local isometry is actually an isometry from B to a convex set in the Euclidean plane. Now B is easily seen to be isometric to a flat Euclidean half-disc of radius R. □

We are now ready to prove Proposition 8.2.

Proof of Proposition 8.2 By Lemma 8.3, we can choose $R = R(M_0, X)$ sufficiently large that for any flat disc of radius R in X the central subdisc of radius $2M_0$ lies in a flat plane. If some ball $B = B\big(\gamma(t), 2R\big)$ in D is isometric to a Euclidean half-disc, then B maps isometrically to X under ϕ. Furthermore, B contains a Euclidean disc of radius R whose central subdisc of radius $2M_0$ intersects the R–neighborhood of γ but is not contained in the $(R + M_0)$–neighborhood of γ, as illustrated in Figure 15. By our choice of R,

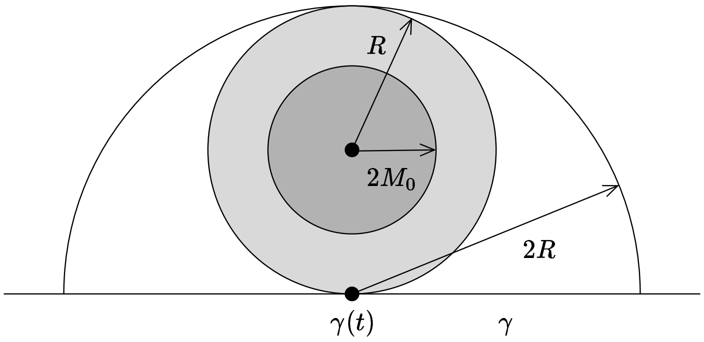

Figure 15: The inner disc intersects the R–neighborhood of γ but is not contained in the $(R+M_0)$–neighborhood of γ.

this subdisc lies in some preflat P. Since B maps isometrically to X under ϕ, it follows that $\phi(P)$ is not contained in the $(R+M_0)$–neighborhood of $\mathrm{Im}(\phi\circ\gamma)$.

We will now assume that no such ball is isometric to a flat half-disc. By Lemma 8.6, each ball $B\big(\gamma(t),2R\big)$ either intersects the image of α, or contains a vertex v with negative curvature. It remains to find a positive constant $\theta(M_0,X)$ so that (1) holds.

Let v be a vertex in the interior of D with $\kappa(v)<0$. Then $2\pi+|\kappa(v)|$ is the length of some locally geodesic loop ℓ in $\mathrm{Lk}\big(\phi(v),\widehat{X}\big)$ since ϕ is a reduced map. If $\phi(v)$ is not a vertex of X, then ℓ has length $n\pi$ for some integer $n>2$. On the other hand, if $\phi(v)$ is a vertex of X, then, since X has only finitely many isometry types of cells, there is a positive constant θ_1, depending only on the space X, such that ℓ has length at least $2\pi+\theta_1$. In either case, we see that $\kappa(v)\le-\theta_1$.

Now suppose v is a vertex on γ with $\kappa(v)<0$. We may assume that $\phi(v)$ is a vertex of X since otherwise $\kappa(v)$ would be a negative multiple of π. Let x and y be points on γ at a distance $4R$ on either side of v, or the endpoints of γ if v is too close to one of the ends. If either x or y is within $2R$ of α, then v is within $6R$ of α. Otherwise, there are vertices z and w with negative curvature so that $d(x,z)$ and $d(y,w)$ are each less than $2R$. If either z or w is in the interior of D, then its curvature is less than $-\theta_1$, and hence v is within $6R$ of a vertex with curvature less than $-\theta_1$. Now suppose that both z and w lie on γ. If either point does not map under ϕ to a vertex of X, then the curvature at that point is at least π as above, so v is within a distance $6R$

of a point with curvature at most $-\pi$. Finally assume that both z and w map under ϕ to vertices of X. Since v also maps to a vertex of X, and the distances from v to z and w are each at most $6R$, then the curvature at v has magnitude at least $\theta_2(R,X)$ by Lemma 8.4.

We have now shown that every vertex with negative curvature either in the interior of D or on γ is within $6R$ of either α or a vertex with negative curvature of magnitude at least

$$\theta_0 = \min\{\theta_1, \theta_2\} > 0.$$

Since every point on γ is within $2R$ of either α or a vertex with negative curvature (which does not lie on α), the pair (D,γ) must be $(8R,\theta_0)$–ruffled. The result now holds for any positive $\theta \leq \theta_0$. □

The second proposition of this section is the following result, stating roughly that in the presence of the Isolated Flats Property preflats are surrounded by ruffles. More specifically, if the interior of a preflat is removed from a disc diagram, then the resulting diagram will be ruffled along the new boundary cycle.

Proposition 8.7 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex with the Isolated Flats Property. There are positive constants $R(X)$ and $\theta(X)$ satisfying the following property. Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram where $\widehat{X}$ is a subdivision of X, and let $D' = D - \mathring{P}$ for some preflat P in D. Then the pair $(D', \partial P)$ is (R,θ)–ruffled.*

The proof of this proposition has many similarities with the proof of Proposition 8.2. Before giving this proof, we will state and prove some related lemmas describing important properties of preflats. First we need to define the notion of a *corner* of a subdiagram.

Definition 8.8 Let D be a disc diagram, and let S be a subdiagram. Let v be a vertex of D. A *corner of S at v* is a set of corners of 2–cells of S whose closure corresponds to an entire connected component of $\mathrm{Lk}(v,S)$ for some 0–cell v. An *exterior corner of S in D* is a set of corners of 2–cells of $D-S$ whose closure corresponds to a connected arc in $\mathrm{Lk}(v,D)$ which intersects $\mathrm{Lk}(v,S)$ only at the endpoints of the arc. The *angle* of a corner is the sum of the angles of its elements.

The first basic property of preflats that we need is their convexity, which is established in the following lemma.

Lemma 8.9 *Let X be a piecewise Euclidean* CAT(0) *2–complex, let $\phi\colon D \to X$ be a reduced disc diagram, and let E be any flat plane in X. Then each component of $\phi^{-1}(E)$ is a convex subspace of D. Furthermore, each preflat P is also convex in D.*

Proof It suffices to show that every exterior corner of $\phi^{-1}(E)$ in D has angle at least π, since then $\phi^{-1}(E)$ is locally convex. Let v be a vertex in $\partial\phi^{-1}(E)$, let I be a closed interval of $\mathbb{R}$, and let $I \to \mathrm{Lk}(v, D)$ be a local isometry such that the intersection $\mathrm{Im}(I) \cap \mathrm{Lk}\big(v, \phi^{-1}(E)\big)$ consists of only the endpoints of $\mathrm{Im}(I)$. Let $\mathrm{Lk}(v, D) \to \mathrm{Lk}\big(\phi(v), X\big)$ be the map induced by ϕ. Then the composition $I \to \mathrm{Lk}(v, D) \to \mathrm{Lk}\big(\phi(v), X\big)$ is a local isometry whose image intersects $\mathrm{Lk}\big(\phi(v), E\big)$ only at its endpoints. Notice that $\mathrm{Lk}\big(\phi(v), E\big)$ has diameter π. So I has length at least π, since otherwise there would be a locally geodesic loop $C \to \mathrm{Lk}\big(\phi(v), X\big)$ of length less than 2π contradicting the Link Condition for X (Theorem 3.7). It now follows easily that every exterior corner of $\phi^{-1}(E)$ in D has angle at least π. □

The next lemma is quite similar to Lemma 8.6. We will show that each preflat is either surrounded by ruffles or adjacent to a large Euclidean half-disc.

Lemma 8.10 *Let $\phi\colon D \to X$ be a reduced disc diagram where X is a piecewise Euclidean* CAT(0) *2–complex, and let P be a preflat. Let $D' = D - \mathring{P}$, and let p be any point on ∂P. If B' is the open ball of radius R in D' centered at p, then either*

(1) *B' contains a vertex with negative curvature in D', or*
(2) *B' intersects ∂D, or*
(3) *B' is isometric to a flat half-disc of radius R.*

Proof Suppose conditions (1) and (2) fail for B'. Then each vertex on the interior of B' has a neighborhood isometric to a Euclidean disc.

Now choose an arbitrary vertex v in $B' \cap \partial P$. By Lemma 8.9, every exterior corner of P in D has angle at least π. Since the curvature $\kappa_{D'}(v)$ is zero, v must have a neighborhood in B' isometric to a Euclidean half-disc. Furthermore, ∂P intersects B' in a collection of disjoint segments which are geodesic in B' as well as in D. In particular, note that P does not lie entirely inside the open ball B of radius R in D centered at p.

By Lemma 8.5 if B' is simply connected, it admits a local isometry into the Euclidean plane. Recall that any metric ball in a CAT(0) space is convex.

We will show that B' is contractible inside D', and hence lifts isometrically to the universal cover $\tilde{D}'$, where it is a metric ball in a CAT(0) space, and must therefore be convex and contractible.

Suppose B' is not contractible in D'. Notice that $\pi_1(D')$ is cyclic generated by the boundary cycle of D. So B' must contain a simple loop ℓ which encloses P. Since $B' \subset B$, it follows that ℓ also lies inside B, which is contractible. Therefore P lies entirely within B, which is a contradiction. It now follows that B' lifts isometrically to a convex subspace of $\tilde{D}'$, and therefore B' is itself CAT(0). Being simply connected and locally flat, B' admits a local isometry to the plane, which must actually be an isometric embedding since B' is CAT(0).

Since B' and P are both convex in D, their intersection is connected. Therefore B' is isometric to a Euclidean half disc of radius R. □

To complete the proof of Proposition 8.7 we need to show that in the presence of the Isolated Flats Property, the third case of the preceding lemma cannot occur.

Proof of Proposition 8.7 By Theorem 5.4, we can choose R sufficiently large that X does not contain an isometrically embedded copy of the space T_R (in the notation of Definition 5.3). Choose $\theta > 0$ so that for each vertex $v \in X^{(0)}$

(1) every combinatorial reduced path in $\mathrm{Lk}(v, X)$ with length greater than π actually has length at least $\pi + \theta$, and

(2) every locally geodesic loop in $\mathrm{Lk}(v, X)$ with length greater than 2π actually has length at least $2\pi + \theta$.

Such a θ exists since X has finitely many isometry types of 2–cells.

Now fix a reduced disc diagram $\phi\colon D \to \widehat{X}$ where $\widehat{X}$ is a subdivision of X, and let $D' = D - \mathring{P}$ for some preflat P in D. Notice that if v is a vertex in $D' - \partial D$ with $\kappa_{D'}(v) < 0$, then by our choice of θ in fact $\kappa_{D'}(v) \leq -\theta$. Thus by Lemma 8.10, either $(D', \partial P)$ is (R, θ)–ruffled or there is some point $p \in \partial P$ so that the ball $B = B(p, R)$ in D' is isometric to a Euclidean half disc.

Notice that ϕ maps the boundary edge of B isometrically to a geodesic segment in a flat plane E of X. Furthermore, $\phi^{-1}(E)$ does not have a local cut point anywhere along ∂B, since otherwise $\phi^{-1}(E)$ would fail to be convex, contradicting Lemma 8.9. It follows that $\phi(B)$ intersects E only along its boundary edge. But then X contains an isometrically embedded copy of the space T_R, which is a contradiction. □

9 2–complexes with isolated flats have the Relatively Thin Triangle Property

In this section we show that the Isolated Flats Property implies the Relatively Thin Triangle Property, establishing (1) $\Rightarrow$ (2) of Theorem 1.1. The proof combines the two main propositions from the previous section about preflats and ruffles.

Theorem 9.1 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex. If X satisfies the Isolated Flats Property, then X also satisfies the Relatively Thin Triangle Property.*

Lemma 9.2 *Let X be a proper, cocompact piecewise Euclidean* CAT(0) *2–complex with the Isolated Flats Property. There are positive constants μ and δ such that the following property holds. Let $\Delta(x,y,z)$ be a geodesic triangle in X. Form a subdivision $\widehat{X}$ of X such that the sides of Δ lie in the 1–skeleton of $\widehat{X}$. Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram for Δ. Suppose D contains a preflat P which maps into a flat F under ϕ. If P has inscribed radius at least μ, then the triangle Δ is δ–thin relative to the flat F.*

Proof To show that Δ is δ–thin relative to F for some δ, it suffices to show that each side of D lies in a δ–neighborhood of the union of the other two sides and P. If we let D' denote $D - \mathring{P}$, then Proposition 8.7 provides constants R and θ so that $(D', \partial P)$ is (R,θ)–ruffled.

Suppose the boundary of P has length L. For each i with

$$1 \leq i \leq m = \lfloor L/2R \rfloor,$$

let x_i be a point on ∂P so that the balls $B_i = B(x_i, R)$ in D' are pairwise disjoint. Call B_i *defiant* if it contains a vertex v with $\kappa_{D'}(v) \leq -\theta$. By the Combinatorial Gauss–Bonnet Theorem, the negative curvature inside D' has total magnitude at most π. So at most π/θ of the B_i are defiant, and each nondefiant B_i intersects ∂D.

We will now show that if P has a sufficiently large inscribed radius, then each of the three sides of the triangular diagram D intersects at least one of the B_i. For each side of D which intersects some ball B_i, the set of all such balls intersecting that side is a contiguous string of nondefiant balls. Thus the set of all balls B_i decomposes into alternate strings of defiant and nondefiant balls, each string of nondefiant balls containing at most π/θ balls.

If no side of D intersects any ball B_i then every B_i is defiant as in Figure 16(a). In this case, we see that ∂P has length less than $\epsilon = 2R(\pi/\theta+1)$, which bounds the inscribed radius of P.

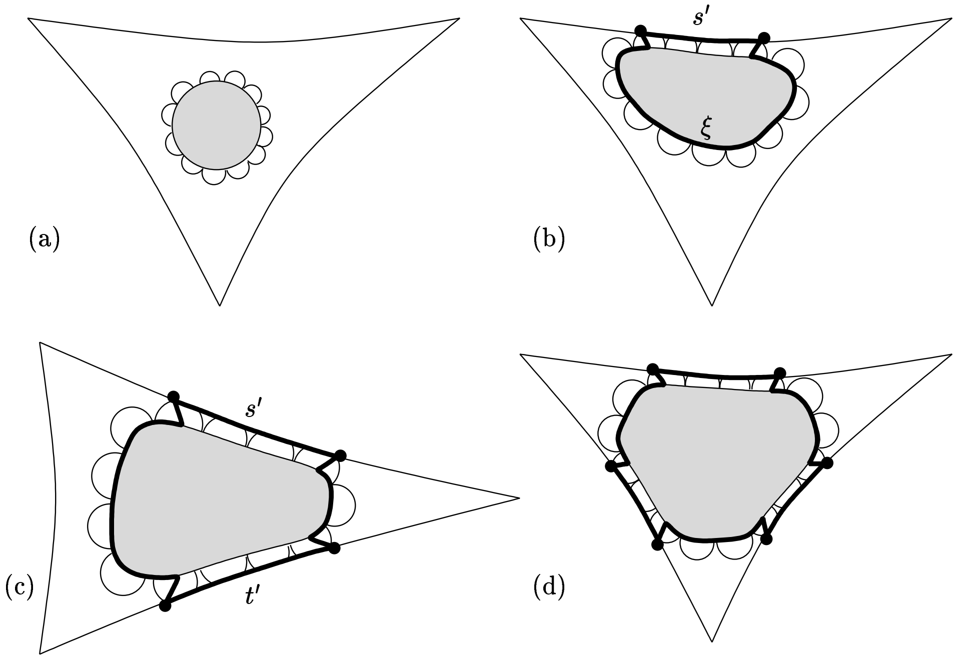

Figure 16: (a) No sides of D intersect the balls B_i (b) One side of D intersects the B_i (c) Two sides of D intersect the B_i (d) Three sides of D intersect the B_i. Only in this last case can the preflat have an arbitrarily large inscribed radius.

If exactly one side s of D intersects the balls B_i, let s' be the shortest subsegment of s containing $\bigcup_i(s \cap B_i)$. The two ends of s' are connected by a path ξ which travels around the "nondefiant" side of ∂P as illustrated in Figure 16(b). Since ξ has length at most $\epsilon+4R$, the geodesic s' also has length at most $\epsilon+4R$. So the inscribed radius of P is bounded by some constant which depends only on ϵ and R.

If exactly two sides s and t of D intersect the balls B_i, then as before, let s' (resp. t') denote the shortest subsegment of s (resp. t) containing

$$\bigcup_i s \cap B_i \quad \Bigl(\text{resp. } \bigcup_i t \cap B_i\Bigr).$$

As in the previous case, the endpoints of s' and t' are separated by a distance of at most $\epsilon+4R$, as shown in Figure 16(c). So again we have a bound on the inscribed radius of P.

Thus we see that for arbitrarily large inscribed radii, each side of D intersects some ball B_i. For each side s, let s' be the subsegment defined as above. It is clear that the endpoints of these three segments come in pairs each of which is a distance at most $\epsilon + 4R$ apart as in Figure 16(d). It now follows easily that the triangular diagram D is $(\epsilon + 4R)$–thin whenever its inscribed radius is larger than some constant μ depending only on the space X. $\square$

Proof of Theorem 9.1 Choose a geodesic triangle Δ in X, and let $D \to \widehat{X}$ be a reduced disc diagram with $\widehat{X}$ a subdivision of X. By the previous lemma, we have constants μ and δ so that if D contains a preflat P with inscribed radius at least μ, then Δ is δ–thin relative to some flat F.

Suppose D does not contain such a preflat. Then by Proposition 8.2, there are constants R and θ, depending on μ, so that for each side s of D the pair (D,s) is (R,θ)–ruffled. But then Proposition 7.7 gives that D is δ'–thin for some constant δ' depending only on R and θ. $\square$

10 The fellow travelling of quasigeodesics and ruffled geodesics

Our objective for the rest of this article is to prove Theorem 13.1, which states that CAT(0) 2–complexes with the Isolated Flats Property also have the Relative Fellow Traveller Property. In this section, we prove Proposition 10.4, which is a special case of Theorem 13.1. Let D be a CAT(0) disc diagram whose boundary is a concatenation of a geodesic and a quasigeodesic. Proposition 10.4 states that the geodesic and quasigeodesic track close together provided that D is ruffled along the geodesic.

Note that the conclusion of Proposition 10.4 is stronger than the conclusion of Theorem 13.1, where we only get that two paths fellow travel relative to flats. Recall that a δ–hyperbolic disc diagram is ruffled throughout. In particular if such a diagram contains a geodesic on its boundary, then the diagram will be ruffled along that geodesic. So from Proposition 10.4 we recover a 2–dimensional version of Theorem 6.4 which states that quasigeodesics and geodesics in a δ–hyperbolic space asynchronously fellow travel.

In fact the structure of our argument in this section is inspired by Cooper, Lustig, and Mihalik's proof in [2] of this fellow traveler theorem for δ–hyperbolic spaces. Their strategy is to first prove an "exponential divergence" theorem [2, Theorem 2.19] which roughly states that geodesics diverge at an exponential rate in a δ–hyperbolic space. They then use the fact that this divergence is *superlinear* to prove the fellow traveller property in [2, Proposition 3.3].

The main tool used in our proof of Proposition 10.4 is Proposition 10.3, which is essentially a "quadratic divergence" theorem analogous to the exponential divergence result alluded to above.

At this point, it may be useful to compare our divergence theorem to some related results in the literature. A *divergence function* for a geodesic space in the sense of [2] is, roughly speaking, a function that provides a lower bound on the rate of divergence of all pairs of geodesic rays in the space. The details of the definition are such that Euclidean space does not admit an unbounded divergence function. Papasoglu shows in [36] that a space admitting an unbounded divergence function admits an exponential divergence function, and is hence δ-hyperbolic by [2]. Since a space with isolated flats (that truly contains flats) fails to be δ-hyperbolic, it cannot admit a quadratic divergence function in this sense. The main distinction between Papasoglu's result and Proposition 10.3 is that we restrict the type of geodesic rays considered. Proposition 10.3 roughly says that a pair of "ruffled" geodesics must diverge at least quadratically. However, the proposition makes no conclusion about arbitrary pairs of rays.

The proof of Proposition 10.3 uses the notion of a *broom*, which is a type of disc diagram that occurs when two geodesics have a common initial segment and then separate from each other. After establishing Proposition 10.3, the proof of [2, Proposition 3.3] can be applied almost verbatim to prove Proposition 10.4.

Definition 10.1 A *broom* B is a piecewise Euclidean CAT(0) disc diagram whose boundary is a composition $\alpha\beta\gamma$ where α and β are geodesics, as illustrated in Figure 17. The *handle* of B is the intersection $\alpha \cap \beta$. The *tip* of B is the common endpoint of α and β. The *height* of B is the minimum of the lengths of α and β. The *branching angle* of B is the angle between α and β at the point where they first separate (in the degenerate case where either α or β is equal to the handle, then the branching angle is defined to be zero). The *outer path* of B is the boundary path γ.

Lemma 10.2 (Linear divergence of brooms) *Let B be a broom of height at least M, branching angle at least θ, a handle of length at most N, and an outer path γ. Then*

$$\ell(\gamma) \geq (M - N)\theta/2.$$

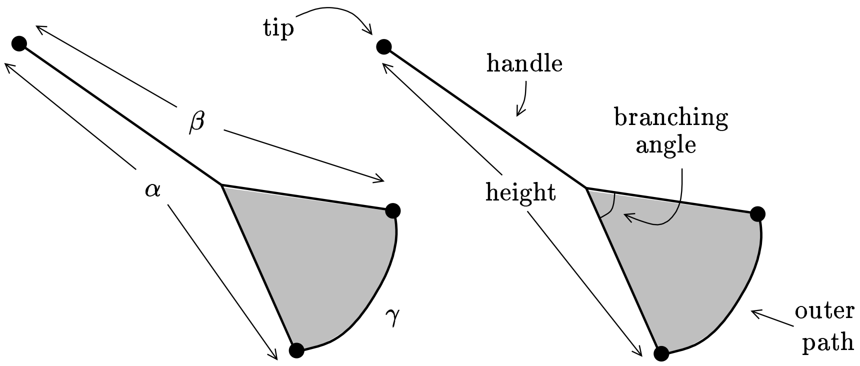

Figure 17: Two pictures of a broom with geodesic sides α and β. On the right, the different parts of the broom are labeled.

Proof Let α and β be the geodesic sides of B. Let $\alpha(0) = \beta(0) = p$ be the tip of B, let q be the other end of the handle, let x be the common endpoint of α and γ, and let y be the common endpoint of β and γ. Let Δ be the geodesic triangle with vertices q, x, and y, and let $\overline{\Delta}$ be a comparison triangle in the Euclidean plane with vertices $\overline{q}$, $\overline{x}$, and $\overline{y}$. Let $\overline{\theta}$ denote the angle at the vertex $\overline{q}$ of $\overline{\Delta}$.

Since α and β each have length at least M, we can set $x' = \alpha(M)$ and $y' = \beta(M)$, and let $\overline{x'}$ and $\overline{y'}$ be the corresponding comparison points on $\overline{\Delta}$. By the Law of Cosines,

$$d(\overline{x'}, \overline{y'}) \le d(\overline{x}, \overline{y}).$$

Using the fact that $\sin(\phi) \ge \phi/2$ whenever $0 \le \phi \le \pi/2$, we conclude that

$$\begin{aligned}
\ell(\gamma) &\ge d(x,y) \\
&= d(\overline{x}, \overline{y}) \\
&\ge d(\overline{x'}, \overline{y'}) \\
&\ge 2(M-N)\sin(\overline{\theta}/2) \\
&\ge (M-N)\overline{\theta}/2 \\
&\ge (M-N)\theta/2.
\end{aligned}$$

□

The following proposition is a quadratic divergence theorem analogous to the exponential divergence theorem for δ–hyperbolic spaces [2, Theorem 2.19]. Bridson and Haefliger reformulated this exponential divergence theorem as an exponential lower bound on the length of any path that a geodesic stays far away from [10, Proposition III.H.1.6]. The precise statement of our quadratic divergence result is closer to the form given by Bridson and Haefliger.

Proposition 10.3 (Quadratic divergence) *For each choice of positive constants R and θ, there is a quadratic function $Q\colon \mathbb{R}^+ \to \mathbb{R}$ so that the following property holds. Let D be any piecewise Euclidean* CAT(0) *disc diagram whose boundary is a concatenation $\alpha\gamma$, where γ is a geodesic, and (D,γ) is (R,θ)–ruffled. Then for any point $p \in \mathrm{Im}(\gamma)$, we have*

$$\ell(\alpha) \geq Q\big(d(p,\mathrm{Im}\,\alpha)\big).$$

Proof Fix a point $p \in \mathrm{Im}(\gamma)$, and let $r = d(p,\mathrm{Im}\,\alpha)$. For convenience, replace θ with $\min\{\theta,\pi\}$. We will show the existence of a quadratic function $Q(r)$ independent of our choice of diagram D. For any point $q \in D$, let $\mathrm{Sh}(q)$ denote the *shadow* of q on α, ie, the set of points $\alpha(t)$ such that the geodesic from p to $\alpha(t)$ passes through q. Notice that each shadow is connected.

Set $Q(r) = 0$ for $r < 2R$. Henceforth we assume that $r \geq 2R$. For $1 \leq i \leq k = \lfloor r/2R \rfloor$, let x_i be a point of $\mathrm{Im}(\gamma)$ with $d(p,x_i) = (2i-1)R$. Notice that the open balls $B_i = B(x_i,R)$ are pairwise disjoint and do not intersect the image of α, as illustrated in Figure 18. Since (D,γ) is (R,θ)–ruffled, B_i contains a point y_i with $\kappa(y_i) \leq -\theta$.

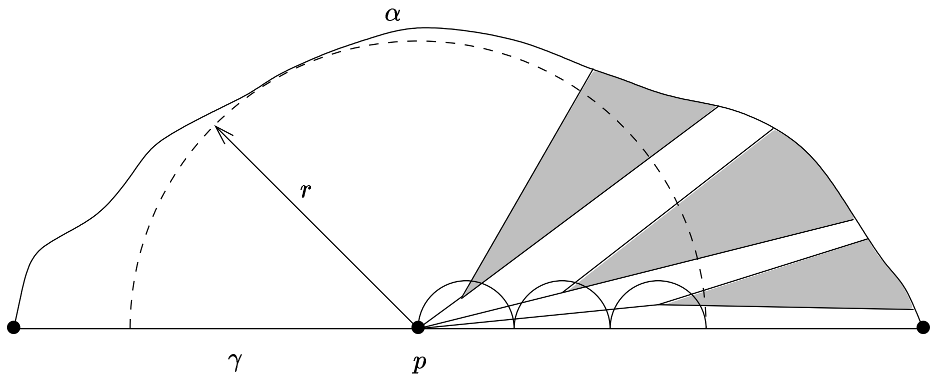

Figure 18: The length of α is bounded below by a quadratic function of r.

Since $\kappa(y_i) \leq -\theta$, we can find a pair of geodesics β_i and β_i' starting at p and passing through y_i, where they separate with an angle of θ. Since geodesics are extendible to the boundary in any CAT(0) disc diagram, we may assume that the given geodesics continue until they eventually hit points z_i and z_i' on the image of α.

Notice that the subdiagram B_i of D bounded by β_i, β_i', and $\mathrm{Sh}(y_i)$ is a broom with height at least $2Rk$, branching angle at least θ, a handle of length at most

$2Ri$, and outer path $\mathrm{Sh}(y_i)$. Applying Lemma 10.2 to B_i gives that

$$\ell\big(\mathrm{Sh}(y_i)\big) \geq R(k-i)\theta.$$

If we knew that the shadows of the y_i had pairwise disjoint interiors, then we would be done, since the sum of their lengths satisfies

$$\sum_{i=1}^{k} \ell\big(\mathrm{Sh}(y_i)\big) \geq R\theta \sum_{i=1}^{k} (k-i) = R\theta k(k-1)/2,$$

which is a quadratic function of $k = \lfloor r/2R \rfloor$, as desired. However, shadows are not in general disjoint.

If two distinct shadows intersect, then one of them is a subset of the other, say $\mathrm{Sh}(y_i) \subset \mathrm{Sh}(y_j)$. Subdividing B_j along the geodesics β_i and β_i' gives a decomposition of B_j into three brooms: B_i and two others C_j and C_j' which branch at y_j and have branching angles adding up to at least θ, as shown in Figure 19. Applying Lemma 10.2 to these three brooms separately, we see that

$$\ell\big(\mathrm{Sh}(y_j)\big) \geq R(k-i)\theta + \ell\big(\mathrm{Sh}(y_i)\big).$$

Repeatedly subdividing brooms in this manner until they all have disjoint interiors shows that, in fact,

$$\ell(\gamma) \geq R\theta k(k-1)/2,$$

as desired. □

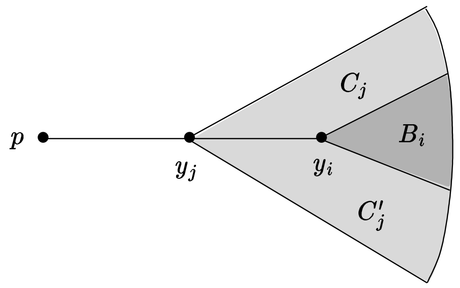

Figure 19: Two shadows that intersect

Proposition 10.4 (Ruffled Fellow Traveller Property) *Given positive constants R, θ, λ, and ϵ, there is a constant L such that the following property holds. Let D be any piecewise Euclidean* CAT(0) *disc diagram whose boundary is a concatenation $\alpha\gamma$, where γ is a geodesic and α is a (λ,ϵ)–quasigeodesic parametrized by arclength. Suppose also that (D,γ) is (R,θ)–ruffled. Then the Hausdorff distance between $\mathrm{Im}(\alpha)$ and $\mathrm{Im}(\gamma)$ is less than L.*

Proof We follow the proof of [2, Proposition 3.3] almost verbatim. Let Q be the quadratic function guaranteed by Proposition 10.3 corresponding to the ruffling constants R and θ. Fix a disc diagram D as in the statement of the proposition. Let

$$r = \sup_{x \in \operatorname{Im}\gamma} \bigl\{ d(x, \operatorname{Im}\alpha) \bigr\},$$

and choose a point $x \in \operatorname{Im}(\gamma)$ where this supremum is achieved. Then

$$\operatorname{Im}(\alpha) \cap B(x,r) = \varnothing.$$

Let y and z be points on $\operatorname{Im}(\gamma)$ at a distance r from x, and let y' and z' be points on $\operatorname{Im}(\gamma)$ at a distance $2r$ from x (or the endpoints of γ if these are closer to x than $2r$). Choose points u and v on $\operatorname{Im}(\alpha)$ such that $d(y',u) \le r$ and $d(z',v) \le r$, as in Figure 20. Notice that

$$\bigl([y',u] \cup [z',v]\bigr) \cap B(x,r) = \varnothing.$$

Following a path by way of y', x, and z', we see that $d(u,v) \le 6r$. However, since α is a (λ,ϵ)–quasigeodesic parametrized by arclength, we have that the length of α from u to v is at most $6\lambda r + \epsilon$. Hence there is a path β of length at most $4r + 6\lambda r + \epsilon$ from y to z which stays outside $B(x,r)$. Furthermore, β together with $[y,z]$ bounds a subdiagram D' of D such that $\bigl(D', [y,z]\bigr)$ is (R,θ)–ruffled by Lemma 7.3. But Proposition 10.3 says that β has length at least $Q(r)$. Therefore r is bounded above by some constant $L_0(R,\theta,\lambda,\epsilon)$, and we see that $\operatorname{Im}(\gamma) \subseteq \mathcal{N}_{L_0}\bigl(\operatorname{Im}(\alpha)\bigr)$.

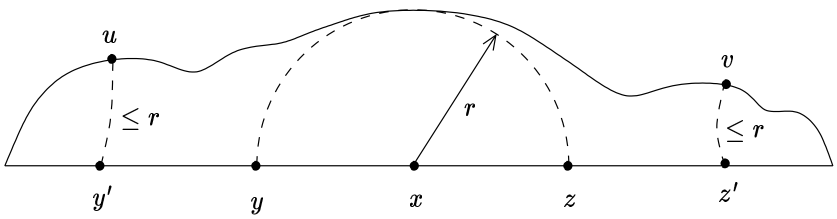

Figure 20: The Ruffled Fellow Traveller Property

Now suppose that $\operatorname{Im}(\alpha) \not\subseteq \mathcal{N}_{L_0}\bigl(\operatorname{Im}(\gamma)\bigr)$. Then each component of $\operatorname{Im}(\alpha) - \mathcal{N}_{L_0}\bigl(\operatorname{Im}(\gamma)\bigr)$ is a path α' with endpoints u and v at a distance L_0 from points y and z on $\operatorname{Im}(\gamma)$. Suppose α is a map $[0,a] \to D$, and $\alpha' = \alpha\big|[t_0,t_1]$. Then each point of $\operatorname{Im}(\gamma)$ is within L_0 of some point of $\alpha\bigl([0,t_0]\bigr) \cup \alpha\bigl([t_1,a]\bigr)$ by the first part of the proof. So there must be some point x on $\operatorname{Im}(\gamma)$ which is within L_0 of some point u_0 on $\alpha\bigl([0,t_0]\bigr)$ and also within L_0 of some point u_1 on $\alpha\bigl([t_1,a]\bigr)$. Thus $d(u_0,u_1) \le 2L_0$, and we see that the length of α from u_0 to u_1

is at most $2L_0\lambda + \epsilon$, which bounds the length of α'. It now follows that every point on α' is at most a distance L from $\mathrm{Im}(\gamma)$, where $L = L_0 + L_0\lambda + \epsilon/2$. □

11 Convex hulls involving preflats

Recall that we need to prove Theorem 13.1 in order to complete the proof of Theorem 1.1. Our strategy for proving Theorem 13.1 is to first consider the special case of a geodesic and quasigeodesic with common endpoints.

This special case consists of showing that a quasigeodesic stays close to the union of a geodesic and the flats that come near the geodesic. In order to prove this special case, it will be useful to understand in detail the structure of the convex hull of the union of two flats and also the convex hull of the union of a point and a flat. In the present diagrammatic setting, it suffices to examine the convex hulls of the corresponding objects inside a reduced disc diagram.

In Proposition 11.1 we give a detailed examination of the convex hull of the union of a preflat P and a point x outside the preflat. We determine that this convex hull lies in a δ–neighborhood of the union of P and the shortest geodesic connecting it to x. We also conclude that the portion of the convex hull close to the preflat is surrounded by ruffles.

We then examine the convex hull of the union of two preflats. In Proposition 11.2 we consider the case of two disjoint preflats, while in Proposition 11.3 we consider the case of two intersecting preflats. In each case, our conclusion is similar to the conclusion of Proposition 11.1. We determine that the convex hull lies in a δ–neighborhood of the union of the two preflats and a shortest geodesic connecting them. We also conclude that the portion of the convex hull close to either preflat is surrounded by ruffles.

Proposition 11.1 *Given a proper, cocompact piecewise Euclidean* $\mathrm{CAT}(0)$ *2–complex X with the Isolated Flats Property, there are positive constants $\delta_1(X)$, and $\theta_1(X)$ such that the following property holds for any positive $\delta \geq \delta_1$ and $\theta \leq \theta_1$.*

Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram, where $\widehat{D}$ is a subdivision of X. Let P be a preflat in D, let x be a point in $D - P$, and let H be the convex hull of $P \cup \{x\}$.

Then $x \in \partial H$, and the boundary path γ of H beginning and ending at the point x is a local geodesic in $D' = D - \mathring{P}$ (when considered as a map with domain an interval, rather than as a map $S^1 \to \partial H$). This local goedesic is a concatenation $\iota_0\omega\eta\upsilon\iota_1$ such that, as illustrated in Figure 21(a),

(1) *η is a subpath of the boundary cycle of P,*
(2) *$\mathrm{Im}(\upsilon)$ and $\mathrm{Im}(\omega)$ each lie in $\mathcal{N}_\delta(P)$,*
(3) *$(D - \mathring{H}, \omega\eta\upsilon)$ is (δ,θ)–ruffled, and*
(4) *the Hausdorff distance between $\mathrm{Im}(\iota_0)$ and $\mathrm{Im}(\iota_1)$ is at most δ.*

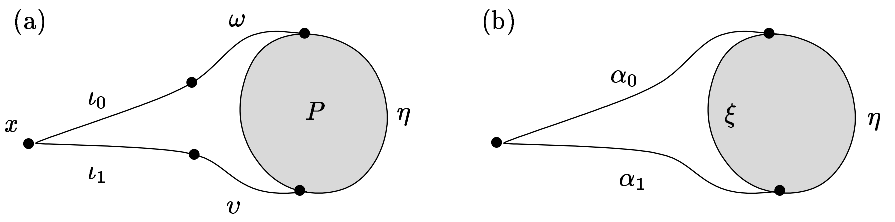

Figure 21: (a) The convex hull of a point and a preflat (b) The paths α_0, α_1 and ξ bound a triangular diagram which is δ–thin.

Proof We first check that $x \in \partial H$. Notice that, away from x, the boundary cycle C of H is a local geodesic in D'. Recall that preflats are convex by Lemma 8.9. If $x \notin \partial H$ then C and ∂P are homotopic local geodesics in D', which are therefore identical. But this is absurd since C encloses x and ∂P does not. So $x \in \partial H$ as desired.

Let γ denote the boundary cycle of H considered as a path which starts and ends at x (as opposed to C which is a map $S^1 \to H$). We have seen that γ is a local geodesic. We now check that $\mathrm{Im}(\gamma)$ intersects ∂P in a single connected arc. If $\mathrm{Im}(\gamma) \cap \partial P = \varnothing$, then γ is a local geodesic in D. But D is simply connected, so γ must be a geodesic, which contradicts the fact that γ is a closed nonconstant loop. So $\mathrm{Im}(\gamma) \cap \partial P$ consists of a finite (positive) number of arcs. Suppose there were more than one such arc. Then γ would contain a subpath which does not involve x and which intersects ∂P only at the endpoints of the subpath. Such a subpath would be a geodesic in D, contradicting the convexity of P.

We have now determined that γ is a concatenation $\alpha_0\eta\alpha_1$ such that η is the maximal subpath of γ lying inside ∂P, and such that α_i is a geodesic in D connecting x and P. Let ξ be the path so that the boundary cycle of P is a concatenation $\eta\xi$, as shown in Figure 21(b). By Proposition 8.7, there are global constants $R(X)$ and $\theta(X)$ so that (D',ξ) is (R,θ)–ruffled. The three D'–geodesics α_0, ξ, and α_1 bound a triangular diagram Δ, which is a subdiagram of a subdivision of D'. By Lemma 7.3, we have that (Δ,ξ) is also

(R,θ)–ruffled. Furthermore, we see by Proposition 7.7 that Δ is δ–thin for some constant $\delta = \delta(R,\theta)$. So there are points $y_i \in \mathrm{Im}(\alpha_i)$ with $d(y_0,y_1) < \delta$. Therefore, α_0 is a concatenation $\iota_0\omega$ and α_1 is a concatenation $\upsilon\iota_1$ such that (2) and (4) are satisfied.

It remains to verify (3). We already know that $(D - \mathring{H}, \eta)$ is (R,θ)–ruffled. We now show that $(D - \mathring{H})$ is (R',θ')–ruffled for some global constants R' and θ'. Let a be the common endpoint of υ and ξ, let b be the common endpoint of υ and ι_1, and let c be a point on $\mathrm{Im}(\xi)$ within a distance δ of b. The points a, b, and c are the vertices of a triangular subdiagram of Δ which has υ and a portion of ξ as two sides and with a third side $[b,c]$ of length less than δ. Lemma 7.8 now shows that $(D - \mathring{H}, \upsilon)$ is $(R',\theta/2)$–ruffled for some constant R' depending only on R, θ, and δ. By a nearly identical argument, we see that $(D - \mathring{H}, \omega)$ is also $(R',\theta/2)$–ruffled. Since $D - \mathring{H}$ is now known to be ruffled along each of ω, η, and υ, Lemma 7.4 shows that $(D - \mathring{H}, \omega\eta\upsilon)$ is $(3R',\theta/2)$–ruffled. Setting $\theta_1 = \theta/2$ and $\delta_1 = \max\{\delta, 3R'\}$ gives (4) for any positive $\delta \geq \delta_1$ and $\theta \leq \theta_1$, as desired. □

Proposition 11.2 *Given a proper, cocompact piecewise Euclidean* CAT(0) *2–complex X with the Isolated Flats Property, there are positive constants $\delta_2(X)$, and $\theta_2(X)$ such that the following property holds for any positive $\delta \geq \delta_2$ and $\theta \leq \theta_2$.*

Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram, where $\widehat{X}$ is a subdivision of X. Let P_0 and P_1 be disjoint preflats in D. Let H be the convex hull of $P_0 \cup P_1$.

Then the boundary cycle $C \to D$ of H is a local geodesic in $D' = D - \big(\mathring{P}_0 \cup \mathring{P}_1\big)$ which is a concatenation

$$\omega_0\eta_0\upsilon_0\iota_0\omega_1\eta_1\upsilon_1\iota_1$$

such that, as illustrated in Figures 22(a) and (b),

1. *η_i is a subpath of the boundary cycle of P_i,*
2. *$\mathrm{Im}(\upsilon_i)$ and $\mathrm{Im}(\omega_i)$ each lie in $\mathcal{N}_\delta(P_i)$,*
3. *$(D - \mathring{H}, \omega_i\eta_i\upsilon_i)$ is (δ,θ)–ruffled, and*
4. *either*
   - (a) *the Hausdorff distance between $\mathrm{Im}(\iota_0)$ and $\mathrm{Im}(\iota_1)$ is at most δ, or*
   - (b) *P_0 and P_1 correspond to distinct flats in X, and each ι_i has image a single point x_i.*

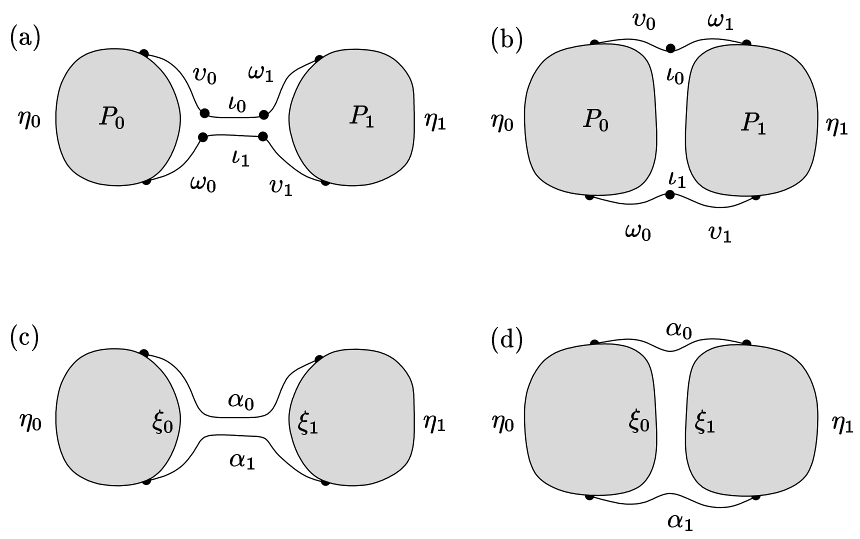

Figure 22: The convex hull of two disjoint preflats P_0 and P_1. In (a), the Hausdorff distance between ι_0 and ι_1 is at most δ. In (b), each ι_i has image a single point. In (c) and (d) the quadrilateral bounded by ξ_0, α_0, ξ_1, and α_1 is δ–thin. The two pictures indicate the two possible shapes of a thin quadrilateral.

Proof Since the present proof is extremely similar to the proof of Proposition 11.1, we omit the redundant details. The only new idea occurs during the verification of (4). By an argument similar to the one in the previous proof, we obtain a δ–thin quadrilateral Q with sides $\xi_0\alpha_0\xi_1\alpha_1$ as shown in Figures 22(c) and (d). Recall that δ–thin quadrilaterals have two possible shapes: either α_0 and α_1 come close together, or ξ_0 and ξ_1 come close together. This dichotomy accounts for the two distinct shapes in Figure 22 and described in the statement of condition (4).

We need to verify the assertion in (4) that if P_0 and P_1 correspond to the same flat in X then the Hausdorff distance between $\mathrm{Im}(\iota_0)$ and $\mathrm{Im}(\iota_1)$ is at most δ. It suffices to consider the following situation. Suppose P_0 and P_1 correspond to the same flat in X, and there exist points $x_i \in \mathrm{Im}(\alpha_i)$ so that for each i, j the distances $d\big(x_i, \mathrm{Im}(\xi_j)\big)$ are less than δ. In order to complete the proof of (4), we need to bound the distance $d(x_0, x_1)$. We can then set the constant δ_2 to be at least as large as this distance.

For the sake of notation, suppose x_0 is within a distance δ of points $p_j \in \mathrm{Im}(\xi_j)$ and x_1 is within a distance δ of points $q_j \in \mathrm{Im}(\xi_j)$, as shown in Figure 23. In

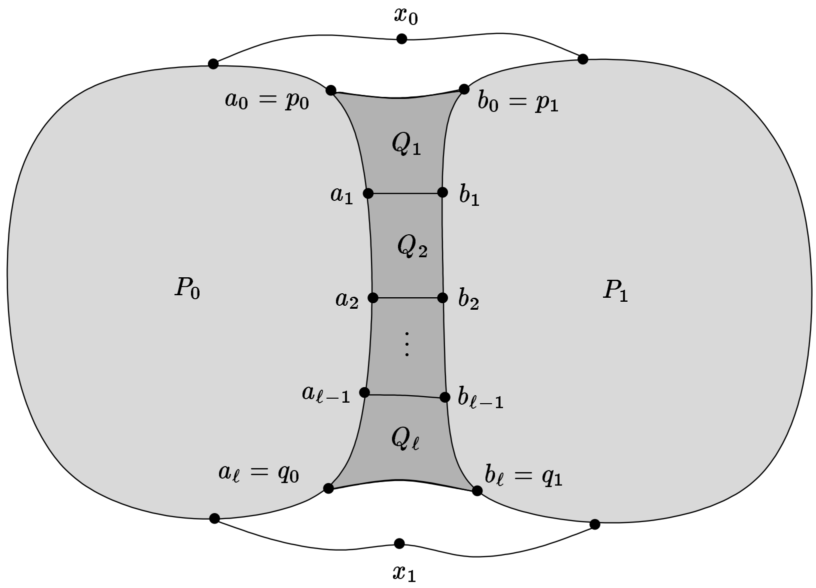

Figure 23: The preflats P_0 and P_1 map to the same flat in X. The darkened diagram is a geodesic quadrilateral Q with corners p_0, p_1, q_0, and q_1. This quadrilateral has been subdivided into ℓ smaller quadrilaterals Q_k.

order to bound $d(x_0, x_1)$, it suffices to bound the quantity

$$\ell = \left\lfloor \frac{\min\{d(p_0, q_0), d(p_1, q_1)\}}{4\delta} \right\rfloor .$$

If one of the distances $d(p_j, q_j)$ is less than 4δ, then $\ell = 0$ and we are done. Otherwise, let $a_0, \dots, a_\ell$ be equally spaced points on ξ_0 so that $a_0 = p_0$ and $a_\ell = q_0$. Similarly, let $b_0, \dots, b_\ell$ be equally spaced points on ξ_1 so that $b_0 = p_1$ and $b_\ell = q_1$. By our choice of ℓ, for $0 \le k \le \ell - 1$ we have

$$d(a_k, a_{k+1}) \ge 4\delta \quad \text{and} \quad d(b_k, b_{k+1}) \ge 4\delta. \tag{$*$}$$

Furthermore a_k, a_{k+1}, b_k, and b_{k+1} are the vertices of a quadrilateral diagram Q_k as shown in Figure 23.

We first rule out the degenerate case that Q_k is not a single line segment. Note that by Theorem 2.5, the distances $d(a_k, b_k)$ and $d(a_{k+1}, b_{k+1})$ are each less than 2δ. If the four corners of Q_k were colinear, then by $(*)$ the geodesics $[a_k, a_{k+1}]$ and $[b_k, b_{k+1}]$ would intersect, contradicting the fact that ξ_0 and ξ_1 are disjoint paths.

Consequently, each Q_k contains some vertex with negative curvature. For if not, then by Corollary 3.9 the diagram Q_k would map isometrically to X under ϕ. But that would mean that Q_k maps into the same flat that the preflats P_0 and P_1 map to, which is absurd.

By Theorem 4.10, the quadrilateral Q with vertices p_0, p_1, q_0, and q_1 has total curvature 2π. The only vertices in Q that could possibly have positive curvature are these four corners. At each of these, the curvature is at most π. So the positive curvature in Q totals at most 4π. As a result, the negative curvature in Q has total magnitude at most 2π. So the curvature at each negatively curved vertex of Q has magnitude at most 2π.

We now show that by our choice of ℓ, each vertex with negative curvature in Q intersects at most two of the quadrilaterals Q_k. Suppose Q_{k-1}, Q_k, and Q_{k+1} share a common point z. Then z is a cut point of Q_k at which the sides $[a_k, b_k]$ and $[a_{k+1}, b_{k+1}]$ meet. But by the triangle inequality, these sides cannot meet, since their lengths are each less than 2δ and we also have $(*)$.

Except along the paths $[a_0, b_0]$ and $[a_\ell, b_\ell]$, every point of negative curvature in Q has curvature with magnitude at least ψ, for some constant ψ depending only on X. So there are at most $2\pi/\psi$ points of negative curvature within the quadrilaterals $Q_2, \dots, Q_{\ell-1}$, which shows that $\ell \leq (4\pi/\psi) + 2$. □

Proposition 11.3 *Given a proper, cocompact piecewise Euclidean* CAT(0) *2–complex X with the Isolated Flats Property, there are positive constants $\delta_3(X)$ and $\theta_3(X)$ such that the following property holds for any positive $\delta \geq \delta_3$ and $\theta \leq \theta_3$.*

Let $\phi\colon D \to \widehat{X}$ be a reduced disc diagram, where $\widehat{X}$ is a subdivision of X. Let P_0 and P_1 be preflats in D with a nonempty intersection. Let H be the convex hull of $P_0 \cup P_1$.

Then the boundary cycle $C \to D$ of H is a local geodesic in $D' = D - \big(\mathring{P}_0 \cup \mathring{P}_1\big)$, which can be expressed as a concatenation

$$\omega_0 \eta_0 \upsilon_0 \omega_1 \eta_1 \upsilon_1 \cdots \omega_{2m-1} \eta_{2m-1} \upsilon_{2m-1}$$

such that, as illustrated in Figure 24(a),

1. *η_i is a subpath of the boundary cycle of $P_{(i \bmod 2)}$,*
2. *$\mathrm{Im}(\upsilon_i)$ and $\mathrm{Im}(\omega_i)$ each lie in $\mathcal{N}_\delta(P_{(i \bmod 2)})$, and*
3. *$(D - \mathring{H}, \omega_i \eta_i \upsilon_i)$ is (δ, θ)–ruffled.*

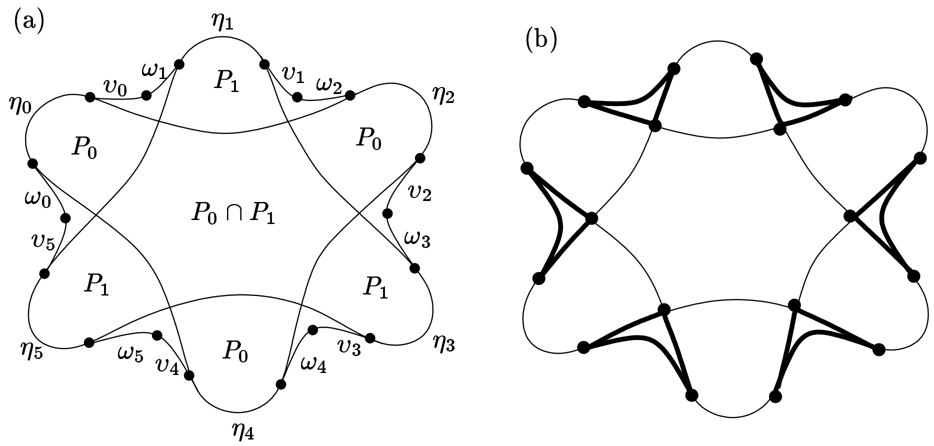

Figure 24: (a) The convex hull of the union of two intersecting preflats (b) The darkened triangles are δ–thin.

Proof The key idea is to observe that the triangles shown in Figure 24(b) are δ–thin. Using this observation, the proof is essentially the same as the proof of Proposition 11.1 without the additional complications that arose in the proof of Proposition 11.3. Since the details are nearly identical, we omit the verification. □

It is important to note that condition (3) does not in general imply that $(D - \mathring{H}, \partial H)$ is (R, ψ)–ruffled for constants R and ψ depending only on the 2–complex X. Although the boundary cycle of H is a concatenation of the paths $\omega_i \eta_i v_i$ for $0 \leq i \leq 2m-1$, Lemma 7.4 can be applied only if one has either a lower bound on the lengths of the concatenated paths or an upper bound on the number of these paths. In general we do not have control over either of these quantities.

12 The Flat Closure Lemma

Recall that we will prove Theorem 13.1 by first considering the special case of a geodesic and quasigeodesic with common endpoints, and then deriving the general case of two quasigeodesics from this special case.

This section is devoted to a proof of the following proposition, which will be used in the next section in the proof of Theorem 13.1. The reader should imagine that the diagram D in the statement of this proposition is a reduced disc diagram such that γ is a geodesic and α a quasigeodesic, although for the purposes of this section we do not need the path α to have any special properties.

The name Flat Closure refers to the fact that we construct a subdiagram of D containing all preflats that come within a certain distance of the geodesic α.

Proposition 12.1 (Flat Closure) *Let X be a proper, cocompact* CAT(0) *2–complex with the Isolated Flats Property. There are positive constants L, K and θ such that the following property holds.*

Let $\phi\colon D \to \widehat{X}$ be any reduced disc diagram where $\widehat{X}$ is a subdivision of X such that the boundary path of D is a concatenation $\alpha\gamma$ where $\phi\circ\gamma$ is a geodesic in X. Then there is a path β in D with the same endpoints as α and γ satisfying the following three conditions:

(I) *β is a geodesic in D_α, where D_α is the subdiagram of D bounded by β and α,*

(II) *$\phi\circ\gamma$ and $\phi\circ\beta$ L–fellow travel relative to flats, and*

(III) *(D_α, β) is (K,θ)–ruffled.*

Proof We begin by carefully choosing several constants which will be necessary for the construction of β. First let $\delta = \delta(X)$ be the largest of the three constants $\delta_1(X)$, $\delta_2(X)$, and $\delta_3(X)$ guaranteed by Propositions 11.1, 11.2, and 11.3 respectively. By the Isolated Flats Property, there is a constant $M(\delta, X) \geq 6\delta$ such that for any two flat planes E_1, E_2 in X, the intersection $\mathcal{N}_\delta(E_1)\cap\mathcal{N}_\delta(E_2)$ has diameter less than M. Using $M_0 = 2M$, choose constants $R(M_0, X)$ and $\theta_0(M_0, X)$ satisfying the conclusion of Proposition 8.2. Let $\theta = \min\{\theta_0, \theta_1, \theta_2, \theta_3\}$, where θ_1, θ_2, and θ_3 are the constants given by Propositions 11.1, 11.2 and 11.3.

The next step is the construction of β, which proceeds as follows. Let $\mathcal{P}$ be the set of all preflats in the diagram D which intersect $\mathcal{N}_R(\operatorname{Im}\gamma)$, where R is some positive constant to be determined in the course of the proof. Form a subdiagram D' from D by removing the interior of $\bigcup_{P\in\mathcal{P}} P$. In that subdiagram, let β be the local geodesic homotopic to α (rel. endpoints). Notice that β cuts D into two subdiagrams D_α and D_γ which intersect along β as shown in Figure 25. Notice that D_γ is the convex hull of $\operatorname{Im}(\gamma)\cup\big(\bigcup_{P\in\mathcal{P}} P\big)$. Condition (I) follows easily from the fact that D_γ is convex.

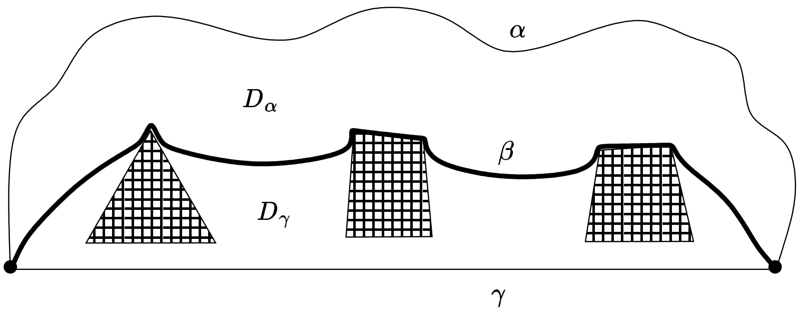

Figure 25: The path β cuts D into two subdiagrams D_α and D_γ that meet along β.

Claim 1 *β is a concatenation*

$$\tau_0\pi_1\rho_1\sigma_1\tau_1\pi_2\rho_2\sigma_2\tau_2\cdots\pi_n\rho_n\sigma_n\tau_n,$$

as shown in Figure 26, such that for some sequence of flats $(E_1,\dots,E_k)$ and some constant $\delta'(X)$, the following properties hold for all i:

A(i): $\operatorname{Im}(\phi\circ\rho_i)\subseteq\mathcal{N}_\delta(E_i)$*, and (D_α,ρ_i) is (δ,θ)–ruffled.*

B(i): $\operatorname{Im}(\phi\circ\sigma_i)\subseteq\mathcal{N}_\delta(E_i)$*, and (D_α,σ_i) is (δ,θ)–ruffled.*

C(i): $\operatorname{Im}(\phi\circ\tau_i)\subseteq\mathcal{N}_{R+M}\big(\operatorname{Im}(\phi\circ\gamma)\big)$*, and (D_α,τ_i) is $(\delta',\theta/4)$–ruffled.*

D(i): $\operatorname{Im}(\phi\circ\pi_{i+1})\subseteq\mathcal{N}_\delta(E_{i+1})$*, and (D_α,π_{i+1}) is (δ,θ)–ruffled.*

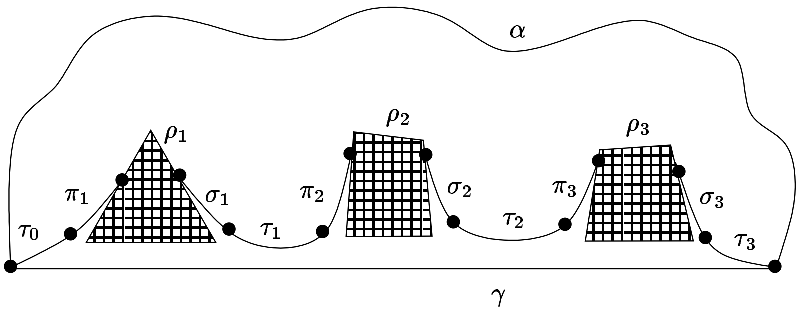

Figure 26: The path β stays close to the union of the preflats in $\mathcal{P}$ and $\operatorname{Im}(\gamma)$.

Proof of Claim 1 For each preflat $P\in\mathcal{P}$, the preimage $\beta^{-1}(P)$ consists of a finite number of closed intervals. Consider the collection of all such intervals for all choices of $P\in\mathcal{P}$. In general, these intervals intersect each other in a complicated fashion. To simplify matters, let $\{\,[a_i,b_i]\bigm| 1\le i\le n\,\}$ be a minimal subcollection of these intervals with the same union as the original collection. Then for each i, the path $\beta([a_i, b_i])$ lies in some P_i. By the minimality of the collection, no interval is contained within another, and there are no nonempty triple intersections among the intervals. We may assume that the intervals are ordered so that

$$a_1 < a_2 < \cdots < a_n$$

and

$$b_1 < b_2 < \cdots < b_n.$$

Our strategy is to verify properties A(i), B(i), C(i), and D(i) by examining in detail the path that β follows between two consecutive intervals $[a_i, b_i]$ and $[a_{i+1}, b_{i+1}]$. We will have two cases depending on whether these intervals intersect or are disjoint. The second case breaks into two cases depending on whether the associated preflats P_i and P_{i+1} intersect or are disjoint. The disjoint preflats case further subdivides into two cases depending on whether these preflats satisfy Condition (4a) or (4b) of Proposition 11.2.

We will also examine in detail what happens to β between the initial point $\beta(0)$ and the first interval $[a_1, b_1]$, as well as an identical case at the terminal end of β. We also consider the degenerate case when the set of intervals $[a_i, b_i]$ is empty. This degenerate case arises precisely when $\mathcal{P} = \varnothing$.

To summarize, we have the following situations to examine:

Case 0 $\{[a_i, b_i]\} = \varnothing$.

Case 1 Between $\beta(0)$ and $[a_1, b_1]$ (and similarly between $[a_n, b_n]$ and the terminal point of β).

Case 2 Between two consecutive intervals such that $P_i \cap P_{i+1} = \varnothing$ and these preflats satisfy Condition (4a) of Proposition 11.2.

Case 3 Between two consecutive intervals such that $P_i \cap P_{i+1} = \varnothing$ and these preflats satisfy Condition (4b) of Proposition 11.2.

Case 4 Between two disjoint consecutive intervals such that $P_i \cap P_{i+1} \neq \varnothing$.

Case 5 Between two intersecting intervals such that $P_i \cap P_{i+1} \neq \varnothing$.

Our examination of these cases will establish properties A(i), B(i), C(i), and D(i) for all values of i.

Case 0 We first consider the degenerate case in which the collection of intervals $\{[a_i, b_i]\}$ is empty. In other words, the image of β does not intersect any P_k. Then β is a local geodesic in the simply connected diagram D. It follows that β is a global geodesic which, therefore, coincides with γ. Since β and α are

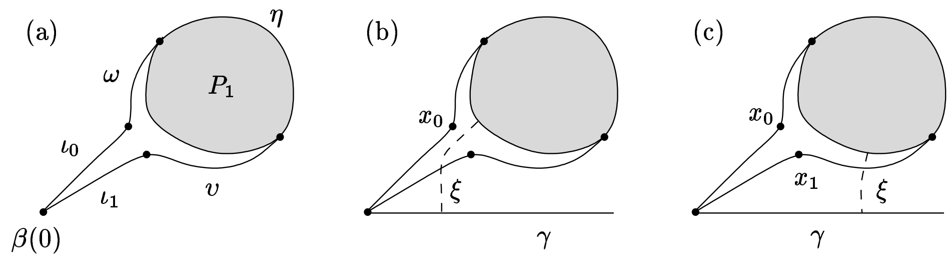

Figure 27: (a) The convex hull of the preflat P_1 and the point $\beta(0)$ (b) The path ξ intersects ι_1. (c) The path ξ intersects $\eta\upsilon$.

homotopic in D', it follows that D and D' are equal. In other words, there are no preflats intersecting $\mathcal{N}_R\big(\mathrm{Im}(\gamma)\big)$. Setting $\tau_0 = \beta = \gamma$, we see that C(0) follows from Proposition 8.2, using $\delta' = 8R$. Since $n = 0$, all other properties in the statement of Claim 1 are vacuous.

Case 1 Let H denote the convex hull of $\beta(0) \cup P_1$. By Proposition 11.1, we see that the boundary cycle of H is a local geodesic in $D - \mathring{P}_1$ except at the point $\beta(0)$. Furthermore, by our choice of δ and θ this boundary cycle is a concatenation $\iota_0\omega\eta\upsilon\iota_1$ as shown in Figure 27(a), satisfying

(1) η is a subpath of the boundary cycle of P_1,

(2) $\mathrm{Im}(\upsilon)$ and $\mathrm{Im}(\omega)$ each lie in $\mathcal{N}_\delta(P_1)$,

(3) $(D - \mathring{H}, \omega\eta\upsilon)$ is (δ,θ)–ruffled, and

(4) the Hausdorff distance between $\mathrm{Im}(\iota_0)$ and $\mathrm{Im}(\iota_1)$ is at most δ.

If we set $\tau_0 = \iota_0$ and $\pi_1 = \omega$, then D(0) follows easily from (2) and (3), since D_α is a subdiagram of $D - \mathring{H}$.

It remains to verify C(0). Recall that P_1 intersects the R–neighborhood of $\mathrm{Im}(\gamma)$. Let ξ be a geodesic from P_1 to $\mathrm{Im}(\gamma)$ with length at most R. Since P_1 and $\mathrm{Im}(\gamma)$ each lie inside the convex subdiagram D_γ, it follows that $\mathrm{Im}(\xi)$ also lies inside D_γ. Therefore $\mathrm{Im}(\xi)$ intersects the image of one of the paths ι_1, υ, and η. If $\mathrm{Im}(\xi)$ intersects $\mathrm{Im}(\iota_1)$ as in Figure 27(b), then $\mathrm{Im}(\xi)$ passes within a distance δ of the common endpoint x_0 of ι_0 and ω. So $d(x_0, \mathrm{Im}\,\gamma) \le R + \delta$. On the other hand, if $\mathrm{Im}(\xi)$ intersects the image of $\eta\upsilon$ as in Figure 27(c), then by the convexity of the metric, the common endpoint x_1 of ι_1 and υ lies within a distance R of $\mathrm{Im}(\gamma)$. So again, $d(x_0, \mathrm{Im}\,\gamma) \le R + \delta$. Thus, in either case, $\mathrm{Im}(\tau_0) \subseteq \mathcal{N}_{R+M}\big(\mathrm{Im}(\gamma)\big)$, since $M \ge \delta$.

We now only have to find a constant δ' so that (D_α, τ_0) is $(\delta', \theta/4)$–ruffled. Let x_2 be the point on $\mathrm{Im}(\gamma)$ which is closest to x_0, as illustrated in Figure 28, and let Δ be the triangular subdiagram of D with corners $\beta(0)$, x_0, and x_2. Let $\widehat{D} = D_\alpha \cup \Delta$, and let $\widehat{\gamma}$ be the subsegment of γ connecting $\beta(0)$ and x_2. Recall that by construction any preflat in D intersecting $\mathcal{N}_R\big(\mathrm{Im}(\gamma)\big)$ lies inside the subdiagram D_γ. The composition $\widehat{D} \hookrightarrow D \to X$ gives a natural notion of a preflat in $\widehat{D}$ with the property that every preflat of $\widehat{D}$ lies inside a preflat of D. So any preflat in $\widehat{D}$ intersecting $\mathcal{N}_R\big(\mathrm{Im}(\widehat{\gamma})\big)$ lies inside the subdiagram Δ. But Δ lies inside $\mathcal{N}_{R+2M}\big(\mathrm{Im}(\widehat{\gamma})\big)$, so we see that any preflat P in $\widehat{D}$ intersecting $\mathcal{N}_R\big(\mathrm{Im}(\widehat{\gamma})\big)$ lies inside $\mathcal{N}_{R+2M}\big(\mathrm{Im}(\widehat{\gamma})\big)$. Therefore, $\phi(P)$ lies inside $\mathcal{N}_{R+2M}\big(\mathrm{Im}(\phi \circ \widehat{\gamma})\big)$. By our choice of R, θ, and $M_0 = 2M$, Proposition 8.2 implies that $(\widehat{D}, \widehat{\gamma})$ is $(8R, \theta)$–ruffled. Since $[x_0, x_2]$ has length at most $R + \delta$, it follows from Lemma 7.8 that (D_α, τ_0) is $(\delta', \theta/2)$–ruffled for some constant $\delta'(R, \theta, \delta)$, which completes the proof of C(0).

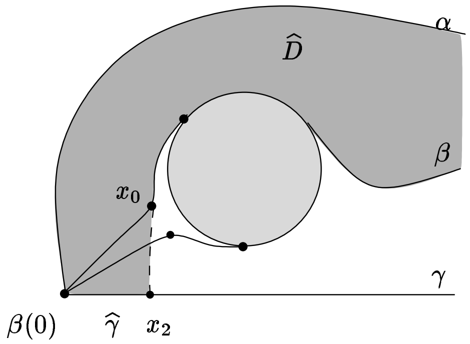

Figure 28: The darkened subdiagram $\widehat{D}$ consists of the union of D_α and the triangular diagram Δ with vertices $\beta(0)$, x_0, and x_2.

Case 2 Suppose $P_i \cap P_{i+1} = \varnothing$, and the preflats in question satisfy Condition (4a) of Proposition 11.2. Let H denote the convex hull of $P_i \cup P_{i+1}$. Then the boundary cycle of H is a local geodesic in $D - \big(\mathring{P}_i \cup \mathring{P}_{i+1}\big)$ freely homotopic to ∂D. By our choice of δ and θ, the cycle ∂H is a concatenation $\omega_0 \eta_0 \upsilon_0 \iota_0 \omega_1 \eta_1 \upsilon_1 \iota_1$ satisfying the conclusion of Proposition 11.2 including Condition (4a), as illustrated in Figure 22(a). Since D_α is a subdiagram of $D - \mathring{H}$, Conditions B(i) and D(i) follow easily if we set $\sigma_i = \upsilon_0$, $\tau_i = \iota_0$ and $\pi_{i+1} = \omega_1$.

It remains to verify C(i). Recall that P_i and P_{i+1} each intersect the R–neighborhood of $\mathrm{Im}(\gamma)$. For each $j \in \{0, 1\}$ let ξ_j be a D–geodesic from P_{i+j}

to $\mathrm{Im}(\gamma)$ with length at most R, and let C_j denote the path $\omega_j \eta_j \upsilon_j$. As in Case 1, the image of ξ_j is inside the convex subdiagram D_γ, so it must intersect the image of at least one of the paths C_0, C_1, or ι_1.

Suppose $\mathrm{Im}(\xi_j)$ intersects $\mathrm{Im}(C_{j+1})$ for some $j \in \mathbb{Z}_2$, as in Figure 29(a). Since H is convex, it follows that $\mathrm{Im}(\iota_0) \subseteq \mathcal{N}_\delta\big(\mathrm{Im}(\xi_j)\big)$. But ξ_j has length at most R, so ι_0 has length at most $R + 2\delta$ and $\mathrm{Im}(\iota_0) \subseteq B(z_j, R + \delta)$, where z_j is the point where the images of γ and ξ_j intersect. Since $\tau_i = \iota_0$, it now follows immediately that (D_γ, ι_0) is $(\delta + R/2, \theta/4)$–ruffled. Since $M \geq \delta$, we now have $\mathrm{C}(i)$, using $\delta' = \delta + R/2$.

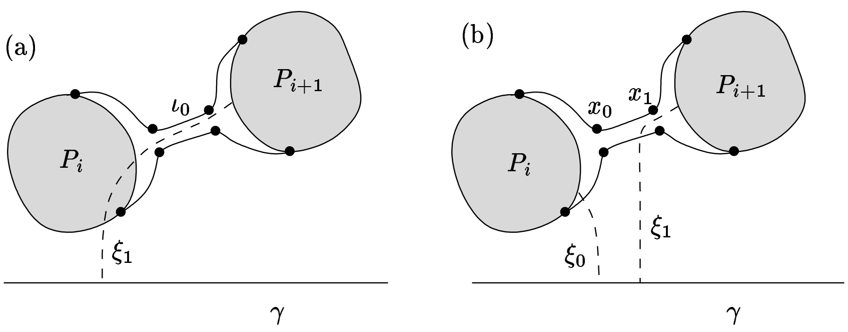

Figure 29: (a) The path ξ_1 intersects C_0. (b) For each $j \in \mathbb{Z}_2$, the path ξ_j intersects either C_j or ι_1.

On the other hand, suppose for each $j \in \mathbb{Z}_2$ that $\mathrm{Im}(\xi_j)$ intersects either ι_1 or C_j, as in Figure 29(b). Then as in Case 1, the convexity of the metric gives that the endpoints x_0 and x_1 of ι_0 lie within a distance $R + \delta$ of points y_0 and y_1 respectively on $\mathrm{Im}(\gamma)$ as shown in Figure 30. The proof of $\mathrm{C}(i)$ now follows from an argument which is essentially the same as that used in the proof of $\mathrm{C}(0)$ in Case 1 above.

Case 3 Suppose $P_i \cap P_{i+1} = \varnothing$, and the preflats in question satisfy Condition (4b) of Proposition 11.2. As in Case 2, let H be the convex hull of $P_i \cup P_{i+1}$. Then by our choice of δ and θ, the cycle ∂H is a concatenation $\omega_0 \eta_0 \upsilon_0 \iota_0 \omega_1 \eta_1 \upsilon_1 \iota_1$ satisfying the conclusion of Proposition 11.2 including Condition (4b), as illustrated in Figure 22(b). In particular, note that the map $\phi\colon D \to X$ sends the preflats P_i and P_{i+1} to distinct flats E_i and E_{i+1} in X, and each path ι_j has image a single point x_j. Conditions $\mathrm{B}(i)$ and $\mathrm{D}(i)$ follow as in Case 2 if we set $\sigma_i = \upsilon_0$, $\tau_i = \iota_0$, and $\pi_{i+1} = \omega_1$.

Since $\tau_i = \iota_0$ has length zero, the requirement that (D_α, τ_i) be ruffled is satisfied vacuously. Therefore, in order to prove $\mathrm{C}(i)$, it suffices to show that the distance

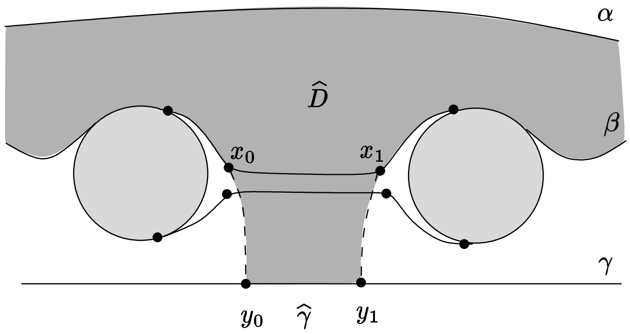

Figure 30: The shaded subdiagram $\widehat{D}$ consists of the union of D_α with the quadrilateral diagram $\square$ with corners x_0, x_1, y_1, and y_0.

in X between $\phi(x_0)$ and $\mathrm{Im}(\phi\circ\gamma)$ is at most $R+M$. As in Case 2, for each $j\in\mathbb{Z}_2$, let ξ_j denote a geodesic of length at most R connecting P_{i+j} with $\mathrm{Im}(\gamma)$, and let C_j denote the path $\omega_j\eta_j\upsilon_j$.

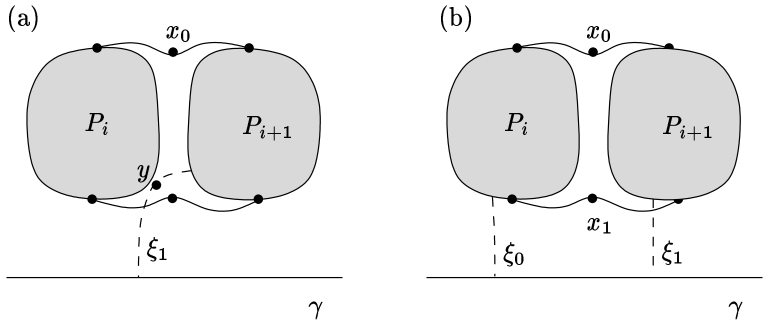

Figure 31: (a) The path ξ_1 intersects C_0. (b) For each $j\in\mathbb{Z}_2$, the path ξ_j intersects C_j.

Suppose for some $j\in\mathbb{Z}_2$ that $\mathrm{Im}(\xi_j)$ intersects $\mathrm{Im}(C_{j+1})$ as shown in Figure 31(a). Then there is a point $y\in\mathrm{Im}(\xi_j)$ with

$$y\in\mathcal{N}_\delta(P_i)\cap\mathcal{N}_\delta(P_{i+1}),$$

and $d(y,\mathrm{Im}\,\gamma)\le R$. But

$$x_0\in\mathcal{N}_\delta(P_i)\cap\mathcal{N}_\delta(P_{i+1}).$$

By our choice of M, the intersection $\mathcal{N}_\delta(E_i)\cap\mathcal{N}_\delta(E_{i+1})$ has diameter less than M. So $d\big(\phi(x_0), \phi(y)\big) \le M$, and hence

$$d\big(\phi(x_0), \mathrm{Im}(\phi \circ \gamma)\big) \le R + M.$$

On the other hand, if for each $j \in \mathbb{Z}_2$, the image of ξ_j intersects $\mathrm{Im}(C_j)$ in some point y_j as shown in Figure 31(b), then by convexity of the metric, it follows that x_1 is within a distance R of $\mathrm{Im}(\gamma)$. Since x_0 and x_1 each lie in $\mathcal{N}_\delta(P_i) \cap \mathcal{N}_\delta(P_{i+1})$, it follows that $d\big(\phi(x_0), \phi(x_1)\big) \le M$, and hence that $d\big(\phi(x_0), \mathrm{Im}(\phi \circ \gamma)\big) \le R + M$.

Case 4 Suppose the intervals $[a_i, b_i]$ and $[a_{i+1}, b_{i+1}]$ are disjoint and the corresponding preflats P_i and P_{i+1} intersect. As in the previous case, the path ι_i will have image a single point, and the map $\phi\colon D \to X$ sends the preflats to distinct flats in X. Conditions $\mathrm{B}(i)$, $\mathrm{C}(i)$, and $\mathrm{D}(i)$ follow almost exactly as in Case 3.

Case 5 The only difference between this case and the previous case is that the intervals $[a_i, b_i]$ and $[a_{i+1}, b_{i+1}]$ intersect. In this case, an argument similar to the one given in Case 3 shows that ϕ maps their intersection $[a_{i+1}, b_i]$ into $\mathcal{N}_{R+M}\big(\mathrm{Im}(\phi \circ \gamma)\big)$. The paths σ_i, τ_i, and π_i can be chosen to all have image the same single point, which can be any point in the intersection $[a_{i+1}, b_i]$. Conditions $\mathrm{B}(i)$, $\mathrm{C}(i)$, and $\mathrm{D}(i)$ now follow as before.

This completes the proof of Claim 1 □

Notice that Claim 1 establishes that β satisfies (II). It also gives a subdivision of β into subpaths such that D_α is ruffled along each subpath. In order to conclude that D_α is ruffled along β, we need to apply Lemma 7.4 which deals with concatenations of ruffled boundary paths. In that lemma, one can bound the constants associated to the ruffles if one has either an upper bound on the number of segments being concatenated or a lower bound on the lengths of the concatenated segments. In Claim 1, we have control over neither of these two quantities.

Our strategy for establishing (III) is to replace β with a smoother path $\check{\beta}$ which tracks close to β, using the following claim.

Claim 2 *There are universal constants $L(X)$ and $K(X)$ and a path $\check{\beta}$ in D with the same endpoints as β so that*

- ($\check{\mathrm{I}}$) *$\check{\beta}$ is a geodesic in $\check{D}_\alpha$, where $\check{D}_\alpha$ is the subdiagram of D bounded by $\check{\beta}$ and α,*
- ($\check{\mathrm{II}}$) *the paths $\phi \circ \check{\beta}$ and $\phi \circ \gamma$ L–fellow travel relative to some sequence of flats, and*

(I$\check{\text{I}}$I) $(D'_\alpha, \check{\beta})$ *is* $(K, \theta/4)$*–ruffled.*

The path $\check{\beta}$ will be a concatenation of *long* paths along each of which $\check{D}_\alpha$ is ruffled. The idea is to mimic the construction of β, this time using only those preflats that come close to $\mathrm{Im}(\gamma)$ and also extend far away from $\mathrm{Im}(\gamma)$. When some of the preflats involved intersect each other, this modification is not quite enough to prove Claim 2, but it is close in spirit to the modification we actually use in the proof of that lemma.

Proof of Claim 2 We use the subdivision of β provided by Claim 1 as a foundation for our construction of $\check{\beta}$. For each i, let us call the path $\pi_i\rho_i\sigma_i$ *nonwandering* if the image of $\phi \circ (\pi_i\rho_i\sigma_i)$ in X lies inside the $(R+2M)$–neighborhood of $\mathrm{Im}(\phi\circ\gamma)$ and call it *wandering* otherwise. The idea is that wandering paths correspond to "tall" preflats and nonwandering paths correspond to "short" preflats. An interval $[r,s] \subset \mathbb{Z}$ is a *maximal nonwandering interval* if it is a maximal interval such that for each i in $[r,s]$, the path $\pi_i\rho_i\sigma_i$ is nonwandering. The path $\check{\beta}$ is formed from β by cutting out each part of β corresponding to a maximal nonwandering interval and replacing it with a path ξ which is described below. The path ξ "smooths" out the nonwandering part of β. More precisely, for each maximal nonwandering interval $[r,s]$ replace the subpath $\sigma_{r-1}\tau_{r-1}\pi_r\cdots\sigma_s\tau_s\pi_{s+1}$ of β with a path ξ constructed as follows.

Case 1 Suppose the preflats P_{r-1} and P_{s+1} are disjoint. Let H be the convex hull of $P_{r-1}\cup P_{s+1}$. Then ξ is defined to be the subpath of ∂H connecting ρ_{r-1} with ρ_{s+1}, as shown in Figure 32. As in the proof of Cases 2 and 3 from Claim 1, the path ξ is a concatenation $\sigma\tau\pi$ so that $\mathrm{Im}(\sigma)\subseteq\mathcal{N}_\delta(P_{r-1})$, $\mathrm{Im}(\phi\circ\tau)\subseteq\mathcal{N}_{R+M}\big(\mathrm{Im}(\phi\circ\gamma)\big)$, and $\mathrm{Im}(\pi)\subseteq\mathcal{N}_\delta(P_{s+1})$. Furthermore, $(D-\mathring{H},\sigma)$ and $(D-\mathring{H},\pi)$ are (δ,θ)–ruffled and $(D-\mathring{H},\tau)$ is $(\delta',\theta/4)$–ruffled.

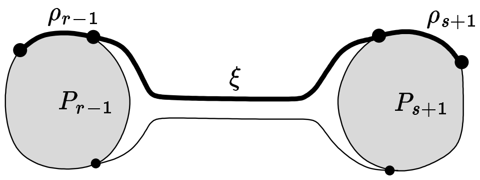

Figure 32: The new path $\check{\beta}$ follows ξ from ρ_{r-1} to ρ_{s+1}.

Case 2 Suppose the preflats P_{r-1} and P_{s+1} intersect. Let H be the convex hull of $P_{r-1}\cup P_{s+1}$. Then the boundary path of H is a concatenation

$$\omega_0\eta_0\upsilon_0\cdots\omega_{2m-1}\eta_{2m-1}\upsilon_{2m-1}$$

satisfying the conclusion of Proposition 11.3. We may assume that ρ_{r-1} is a subpath of η_0 and that ρ_{s+1} is a subpath of η_{2i-1} for some i. As shown in Figure 33, let σ denote the path which starts at the end of ρ_{r-1} and follows the rest of η_0 and then follows υ_0. Let π denote the path which follows ω_1 and then follows ∂P_{s+1} until it reaches the first endpoint of η_{2i-1}. Let τ be the constant path with image the common endpoint of σ and π. Then $\mathrm{Im}(\sigma) \subseteq \mathcal{N}_\delta(P_{r-1})$, $\mathrm{Im}(\phi\circ\tau) \subseteq \mathcal{N}_{R+M}\big(\mathrm{Im}(\phi\circ\gamma)\big)$, and $\mathrm{Im}(\pi) \subseteq \mathcal{N}_\delta(P_{s+1})$. Furthermore, (D'_α,σ) and (D'_α,π) are each (δ,θ)–ruffled, and (D'_α,τ) is $(\delta',\theta/4)$–ruffled.

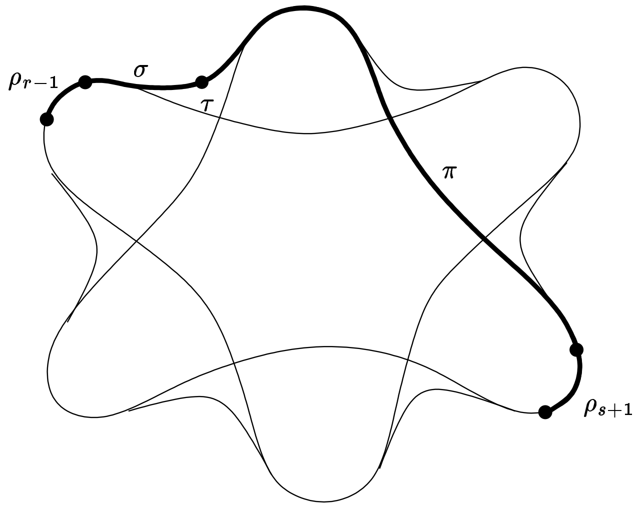

Figure 33: The new path $\check\beta$ follows $\xi = \sigma\tau\pi$ from ρ_{r-1} to ρ_{s+1}.

It now follows that in either case, we may relabel the subpaths of $\check\beta$ and the preflats $\{P_i\}$ so that $\check\beta$ is a concatenation

$$\tau_0\pi_1\rho_1\sigma_1\tau_1\cdots\pi_k\rho_k\sigma_k\tau_k$$

satisfying the following properties

(1) $(\widehat{D}_\alpha,\pi_i)$ is (δ,θ)–ruffled and $\mathrm{Im}(\phi\circ\pi_1) \subset \mathcal{N}_\delta(E_i)$,

(2) $(\widehat{D}_\alpha,\rho_i)$ is (δ,θ)–ruffled and $\mathrm{Im}(\phi\circ\rho_i) \subset \mathcal{N}_\delta(E_i)$,

(3) $(\widehat{D}_\alpha,\sigma_i)$ is (δ,θ)–ruffled and $\mathrm{Im}(\phi\circ\sigma_i) \subset \mathcal{N}_\delta(E_i)$,

(4) $(\widehat{D}_\alpha,\tau_i)$ is $(\delta',\theta/4)$–ruffled and $\mathrm{Im}(\phi\circ\tau_i) \subseteq \mathcal{N}_{R+M}\big(\mathrm{Im}(\phi\circ\gamma)\big)$, and

(5) $(\pi_i\rho_i\sigma_i)$ has endpoints in $\mathcal{N}_{R+M}\big(\mathrm{Im}(\phi\circ\gamma)\big)$ but its image does not lie entirely inside $\mathcal{N}_{R+2M}\big(\mathrm{Im}(\phi\circ\gamma)\big)$.

Clearly the paths $\phi \circ \check{\beta}$ and $\phi \circ \gamma$ L–fellow travel relative to the sequence $(E_1, \dots, E_k)$ of flats, where $L = R + M$. Since the pairs (D'_α, π_i), (D'_α, ρ_i), and (D'_α, σ_i) are each (δ, θ)–ruffled, it follows from Lemma 7.4 that $(D'_\alpha, \pi_i\rho_i\sigma_i)$ is $(6\delta, \theta)$–ruffled. Furthermore, by (5) it is clear that $(\pi_i\rho_i\sigma_i)$ has length at least $2M$, which is at least 12δ by our choice of M. A computation similar to the proof of Lemma 7.4 now easily shows that $(D'_\alpha, \check{\beta})$ is $(K, \theta/4)$–ruffled, where $K = 2M + 2\delta$, completing the proof of Claim 2. □

Replacing β with $\check{\beta}$ and θ with $\theta/4$ completes the proof of Proposition 12.1. □

13 2–complexes with isolated flats have the Relative Fellow Traveller Property

In this section, we are finally ready to prove the following theorem.

Theorem 13.1 *A proper, cocompact piecewise Euclidean* CAT(0) 2*–complex with the Isolated Flats Property satisfies the Relative Fellow Traveller Property.*

Recall that the Relative Fellow Traveller Property deals with the fellow travelling of a pair of quasigeodesics with common endpoints. The proof uses the following special case in which we consider a geodesic and quasigeodesic with common endpoints. This special case is an easy consequence of Propositions 10.4 and 12.1.

Proposition 13.2 *Let X be a proper, cocompact* CAT(0) *space with the Isolated Flats Property. Then for each fixed λ and ϵ there is a constant $R(\lambda, \epsilon, X)$ so that any geodesic and (λ, ϵ)–quasigeodesic in X with common endpoints R–fellow travel relative to flats.*

The proof of Proposition 13.2 uses the following standard technical result, which allows one to ignore the local pathologies of a quasigeodesic by approximating it with a piecewise geodesic path. For a proof of this lemma, see [10, Lemma III.H.1.11].

Lemma 13.3 (Taming quasigeodesics) *Let X be a geodesic space. Given any (λ, ϵ)–quasigeodesic α in X, one can find a continuous path α' satisfying the following properties:*

(1) *α and α' have the same endpoints,*

(2) *α' is piecewise geodesic,*

(3) *α' is a (λ', ϵ')–quasigeodesic when parametrized by arclength, where λ' and ϵ' depend only on λ and ϵ, and*

(4) *the Hausdorff distance between the images of α and α' is less than $\lambda + \epsilon$.*

Proof of Proposition 13.2 Let α be a (λ, ϵ)–quasigeodesic in X, and let γ be the geodesic connecting its endpoints. Let α' be a tame (λ', ϵ')–quasigeodesic as in Lemma 13.3. Let γ be the geodesic connecting the endpoints of α'. Notice that α' and γ' both lie in the 1–skeleton of some subdivision $\widehat{X}$ of X. Since the concatenation $\alpha'\overline{\gamma}$ is a nullhomotopic loop in X, Theorem 4.6 gives a reduced disc diagram $\phi\colon D \to \widehat{X}$ for this loop. Let $\tilde{\alpha}'$ and $\tilde{\gamma}$ be paths in D so that $\phi \circ \tilde{\alpha}' = \alpha'$ and $\phi \circ \tilde{\gamma} = \gamma$. Since ϕ is distance nonincreasing, it follows that $\tilde{\gamma}$ is a geodesic and $\tilde{\alpha}'$ is a (λ', ϵ')–quasigeodesic parametrized by arclength.

Let β be a path in D with the same endpoints as $\tilde{\alpha}'$ and $\tilde{\gamma}$ satisfying the conclusion of Proposition 12.1. Then applying Proposition 10.4 to the subdiagram D_α bounded by β and $\tilde{\alpha}'$, we see that the Hausdorff distance between the images of β and $\tilde{\alpha}'$ is at most $N = N(\lambda, \epsilon, X)$. Since $\phi\circ\beta$ and γ L–fellow travel relative to flats for some L depending only on X, it follows that γ and α R–fellow travel relative to flats, where R depends only on λ, ϵ, and X as desired. □

We are now ready to prove Theorem 13.1 using the special case proved in Proposition 13.2 and the Isolated Flats Property.

Proof of Theorem 13.1 We need to show that given constants λ and ϵ we can find a constant L so that any pair of (λ, ϵ)–quasigeodesics α and α' with common endpoints L–fellow travel relative to flats. So fix a pair α and α' of such quasigeodesics, and let γ denote the geodesic connecting the endpoints of α.

By Proposition 13.2, we know that the paths α and γ and the paths α' and γ each R–fellow travel relative to some sequence of flats, where $R = R(\lambda, \epsilon, X)$. The main difficulty with pasting together these two facts is that the sequences of flats involved may not be the same. In fact, the sequences of flats for each pair are not even well-defined. For instance, if a pair of paths travels for a sufficiently short distance in some flat, then that flat can be inserted or deleted from a sequence of flats without affecting whether the paths fellow travel relative to the sequence.

To circumvent the difficulties alluded to above, we construct a canonical sequence of flats using only properties of γ so that each of the given pairs fellow travels relative to the constructed sequence. Furthermore, the partition of γ in Definition 6.5 will be canonical, so that the two pairs can be pasted together coherently. It will then follow that α and α' fellow travel relative to the canonical sequence of flats.

By the Isolated Flats Property, there is a constant $K = K(R, X)$ such that given any pair of flat planes E_1 and E_2 in X the intersection $\mathcal{N}_R(E_1) \cap \mathcal{N}_R(E_2)$ has diameter less than K. Let us call a flat plane E in X an *essential γ–flat* if $\mathrm{Im}(\gamma)$ intersects the R–neighborhood of E in a segment ξ of length at least K. In this case, the segment ξ will be called a *maximal flat segment* of γ. Removing the subsegment of length $K/2$ from each end of ξ gives a shorter segment η called a *shortened flat segment* of γ. Notice that any two distinct shortened flat segments η_1 and η_2 of γ are disjoint, since otherwise there would be a pair of distinct flats E_1 and E_2 whose R–neighborhoods intersect in a set of diameter at least K.

Our canonical sequence of flats will be the essential γ–flats, and the canonical partition of γ will consist of the shortened flat segments of γ alternating with the segments of γ that connect two consecutive shortened flat segments. We need to verify that γ and α actually fellow travel relative to this canonical data using the given (non-canonical) sequence of flats and our given (non-canonical) partition of γ

Our argument consists of two directions. First we show that each flat of our given sequence appears in the canonical sequence, unless it is very small. Then we show that every flat of the canonical sequence which does not correspond to a flat of our given sequence may be added to that sequence without creating problems.

For the first direction, suppose the quasigeodesic α contains a subpath π whose endpoints x and y lie within a distance R of points w and z in the image of γ, and suppose further that the image of π lies in an R–neighborhood of some flat E. Finally suppose the segment $[w, z]$ lies in an R–neighborhood of E.

If the distance from w to z is less than K, then the distance from x to y is less than $K + 2R$. An easy computation using the definition of quasigeodesic then shows that $\mathrm{Im}(\pi)$ has diameter less than $\lambda^2(K + 2R) + 2\epsilon\lambda$. So in this case, $\mathrm{Im}(\pi)$ lies inside a R_1–neighborhood of $[w, z]$, where $R_1 = R + \lambda^2(K + 2R) + 2\epsilon\lambda$. Similarly, since the diameter of $[w, z]$ is less than K, it follows that $[w, z]$ lies in a R_2–neighborhood of $\mathrm{Im}(\pi)$, where $R_2 = R + K$. So the Hausdorff distance between $\mathrm{Im}(\pi)$ and $[w,z]$ is less than $R' = \max\{R_1, R_2\}$. On the other hand, if $d(w,z) \geq K$, then $[w,z]$ lies inside a unique maximal flat segment of γ.

Now for the second direction, suppose α contains a subpath β with endpoints x' and y' such that β is within a Hausdorff distance R from a subpath $[w',z']$ of γ. Suppose further that $[w',z']$ lies in an R–neighborhood of some flat plane E. Then $\mathrm{Im}(\beta)$ lies in a $2R$–neighborhood of E.

It is now clear that we can choose partitions

$$0 = t_0 \leq s_0 \leq t_1 \leq s_1 \leq \cdots \leq t_n \leq s_n = a$$

and

$$0 = t'_0 \leq s'_0 \leq t'_1 \leq s'_1 \leq \cdots \leq t'_n \leq s'_n = a'$$

of the domains of γ and α respectively, so that the subpaths $\gamma\big|[s_{i-1},t_i]$ are precisely the collection of all shortened flat segments of γ. Furthermore, the Hausdorff distance between $\gamma\big([t_i,s_i]\big)$ and $\alpha\big([t'_i,s'_i]\big)$ is at most $R''(\lambda,\epsilon,X)$, while the sets $\gamma\big([t_{i-1},s_i]\big)$ and $\alpha\big([t'_{i-1},s'_i]\big)$ lie in a $2R$–neighborhood of a unique essential γ–flat E_i.

Since the partition of γ and the sequence of essential γ–flats described above is independent of the choice of α, it follows that α and α' L–fellow travel relative to the sequence of flats $(E_1,\dots,E_n)$ for some constant L depending only on λ, ϵ, and X. □

References

[1] **I Aitchison**, *Canonical flat structures on 3-manifolds*, preprint

[2] **J M Alonso**, **T Brady**, **D Cooper**, **V Ferlini**, **M Lustig**, **M Mihalik**, **H Short**, *Notes on word hyperbolic groups*, (H Short, editor), from: "Group theory from a geometrical viewpoint (Trieste, 1990)", (É Ghys, A Haefliger, A Verjovsky, editors), World Sci. Publishing, River Edge, NJ (1991) 3–63

[3] **W Ballmann**, *Singular spaces of nonpositive curvature*, from: "Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988)", (É Ghys, P de la Harpe, editors), Birkhäuser Boston, Boston, MA (1990) 189–201

[4] **W Ballmann**, **M Brin**, *Polygonal complexes and combinatorial group theory*, Geom. Dedicata 50 (1994) 165–191

[5] **V Bangert**, **V Schroeder**, *Existence of flat tori in analytic manifolds of nonpositive curvature*, Ann. Sci. École Norm. Sup. (4) 24 (1991) 605–634

[6] **N Benakli**, *Polygonal complexes. I. Combinatorial and geometric properties*, J. Pure Appl. Algebra 97 (1994) 247–263

[7] **B H Bowditch**, *Relatively hyperbolic groups* (1999), preprint, University of Southampton

[8] **M R Bridson**, *Geodesics and curvature in metric simplicial complexes*, from: "Group theory from a geometrical viewpoint (Trieste, 1990)", (É Ghys, A Haefliger, A Verjovsky, editors), World Sci. Publishing, River Edge, NJ (1991) 373–463

[9] **M R Bridson**, *On the existence of flat planes in spaces of nonpositive curvature*, Proc. Amer. Math. Soc. 123 (1995) 223–235

[10] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Springer-Verlag, Berlin (1999)

[11] **C B Croke**, **B Kleiner**, *Spaces with nonpositive curvature and their ideal boundaries*, Topology 39 (2000) 549–556

[12] **C B Croke**, **B Kleiner**, *The geodesic flow of a nonpositively curved graph manifold*, Geom. Funct. Anal. 12 (2002) 479–545

[13] **M Dehn**, *Papers on group theory and topology*, Springer-Verlag, New York (1987), translated from the German by J. Stillwell

[14] **P Eberlein**, *Geodesic flows on negatively curved manifolds. II*, Trans. Amer. Math. Soc. 178 (1973) 57–82

[15] **V A Efromovich**, **E S Tihomirova**, *Continuation of an equimorphism to infinity*, Soviet Math. Dokl. 4 (1963) 1494–1496

[16] **D B A Epstein**, **J W Cannon**, **D F Holt**, **S V F Levy**, **M S Paterson**, **W P Thurston**, *Word processing in groups*, Jones and Bartlett Publishers, Boston, MA (1992)

[17] **B Farb**, *Relatively hyperbolic groups*, Geom. Funct. Anal. 8 (1998) 810–840

[18] **S M Gersten**, *Reducible diagrams and equations over groups*, from: "Essays in group theory", (S M Gersten, editor), Springer, New York (1987) 15–73

[19] **S M Gersten**, **H B Short**, *Small cancellation theory and automatic groups*, Invent. Math. 102 (1990) 305–334

[20] **S M Gersten**, **H B Short**, *Small cancellation theory and automatic groups. II*, Invent. Math. 105 (1991) 641–662

[21] **M Gromov**, *Hyperbolic groups*, from: "Essays in group theory", (S M Gersten, editor), Springer, New York (1987) 75–263

[22] **F Haglund**, *Les polyèdres de Gromov*, C. R. Acad. Sci. Paris Sér. I Math. 313 (1991) 603–606

[23] **J Heber**, *Hyperbolische geodatische Raume*, Diplomarbeit, Univ. Bonn (1987)

[24] **G C Hruska**, *Geometric invariants of spaces with isolated flats*, preprint available at `http://www.math.uchicago.edu/~chruska/papers`

[25] **G C Hruska**, *On the relative hyperbolicity of nonpositively curved groups with isolated flats*, in preparation

[26] **M Kapovich**, **B Leeb**, *On asymptotic cones and quasi-isometry classes of fundamental groups of* 3*-manifolds*, Geom. Funct. Anal. 5 (1995) 582–603

[27] **J Kari**, **P Papasoglu**, *Deterministic aperiodic tile sets*, Geom. Funct. Anal. 9 (1999) 353–369

[28] **U Lang**, *Quasigeodesics outside horoballs*, Geom. Dedicata 63 (1996) 205–215

[29] **R C Lyndon**, *On Dehn's algorithm*, Math. Ann. 166 (1966) 208–228

[30] **R C Lyndon**, **P E Schupp**, *Combinatorial group theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89, Springer-Verlag, Berlin (1977)

[31] **H A Masur**, **Y N Minsky**, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. 138 (1999) 103–149

[32] **J P McCammond**, **D T Wise**, *Fans and ladders in small cancellation theory*, Proc. London Math. Soc. (3) 84 (2002) 599–644

[33] **H M Morse**, *A fundamental class of geodesics on any closed surface of genus greater than one*, Trans. Amer. Math. Soc. 26 (1924) 25–60

[34] **G Moussong**, *Hyperbolic Coxeter groups*, PhD thesis, Ohio State Univ. (1988)

[35] **G A Niblo**, **L D Reeves**, *The geometry of cube complexes and the complexity of their fundamental groups*, Topology 37 (1998) 621–633

[36] **P Papasoglu**, *Strongly geodesically automatic groups are hyperbolic*, Invent. Math. 121 (1995) 323–334

[37] **S J Pride**, *Star-complexes, and the dependence problems for hyperbolic complexes*, Glasgow Math. J. 30 (1988) 155–170

[38] **D Y Rebbechi**, *Algorithmic properties of relatively hyperbolic groups*, PhD thesis, Rutgers Univ. (2001)

[39] **W Rudin**, *Real and complex analysis*, third edition, McGraw-Hill Book Co., New York (1987)

[40] **Z Sela**, *Diophantine geometry over groups: A list of research problems*, available at `http://www.ma.huji.ac.il/~zlil/`

[41] **H Short**, *Quasiconvexity and a theorem of Howson's*, from: "Group theory from a geometrical viewpoint (Trieste, 1990)", (É Ghys, A Haefliger, A Verjovsky, editors), World Sci. Publishing, River Edge, NJ (1991) 168–176

[42] **P Tukia**, *Convergence groups and Gromov's metric hyperbolic spaces*, New Zealand J. Math. 23 (1994) 157–187

[43] **E R van Kampen**, *On some lemmas in the theory of groups*, Amer. J. Math. 55 (1933) 268–273

[44] **C M Weinbaum**, *The word and conjugacy problems for the knot group of any tame, prime, alternating knot*, Proc. Amer. Math. Soc. 30 (1971) 22–26

[45] **J Wilson**, *A* CAT(0) *group with uncountably many distinct boundaries*, preprint

[46] **D T Wise**, *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*, PhD thesis, Princeton Univ. (1996)

[47] **D T Wise**, *Subgroup separability of the figure 8 knot group* (1998), preprint available at `http://www.math.mcgill.ca/wise/papers.html`