Commentarii Mathematici Helvetici

Quasi-isometries of groups, graphs and surfaces

Sylvain Maillot

Abstract. We give a characterization of virtual surface groups as groups quasi-isometric to complete simply-connected Riemannian surfaces. Results on the equivalence up to quasi-isometry of various bounded geometry conditions for Riemannian surfaces are also obtained.

Mathematics Subject Classification (2000). Primary 20F65, 20F69.

Keywords. Quasi-isometry, surface group, free group, isoperimetric inequality.

1. Introduction

A major question in geometric group theory is to determine which groups are quasi-isometric to a given metric space X. The answer is known in many cases, e.g. when X is a symmetric space (see for instance the survey [8]) or for some spaces arising in 3-manifold geometry [21, 22, 24] or combinatorial group theory [11, 9, 10, 27]. Here we address a related question: given a smooth manifold R, determine the class of groups that are quasi-isometric to *some* complete Riemannian metric on R. When R is the line $\mathbf{R}$ or the open annulus $S^1 \times \mathbf{R}$, the answer is obviously the class of 2-ended (i.e. virtually cyclic) groups, because the number of ends is a quasi-isometry invariant for proper geodesic metric spaces. Hence the first interesting case is that of the plane $\mathbf{R}^2$.

This case has been studied by G. Mess in his work on the Seifert fiber space conjecture [23], which unfortunately is still unpublished. To state his result, we need some terminology. A *Riemannian plane* is a Riemannian manifold diffeomorphic to $\mathbf{R}^2$. Such a manifold is *quasi-homogeneous* if all balls of given radius are isometric to some ball in a compact submanifold (the precise definition is given in section 3.2). A *virtual (closed) surface group* is a group Γ having a subgroup Γ' of finite index, such that Γ' is isomorphic to the fundamental group of some compact surface (without boundary). This is equivalent to saying that Γ is an extension of a finite group by an infinite 2-dimensional (closed) orbifold group.

Some results of [23], together with the convergence group theorem due to [26, 14, 5], give the following characterization of groups quasi-isometric to a complete,

quasi-homogeneous Riemannian plane.

Theorem 1.1. *Let* Γ *be a finitely generated group. If* Γ *is quasi-isometric to a complete, quasi-homogeneous Riemannian plane, then* Γ *is a virtual closed surface group.*

The main result of this paper generalizes theorem 1.1 in two directions. First the hypothesis of quasihomogeneity can be removed, as was stated in Mess's article. Second, we deal with simply-connected surfaces with geodesic boundary. This leads to the following statement:

Theorem 1.2. *Let* Γ *be a finitely generated group. Then* Γ *is a virtual surface group iff* Γ *is quasi-isometric to some complete simply-connected Riemannian surface with (possibly empty) geodesic boundary.*

Theorem 1.2 is a consequence of theorem 1.1 and the following two propositions:

Proposition 1.3. *Let* Γ *be a finitely generated group. If* Γ *is quasi-isometric to a complete Riemannian plane, then* Γ *is quasi-isometric to a complete, quasi-homogeneous Riemannian plane.*

Proposition 1.4. *A finitely generated group* Γ *is virtually free iff it is quasi-isometric to some complete simply-connected Riemannian surface* R *with nonempty geodesic boundary.*

We will also give a new proof of theorem 1.1 assuming the convergence group theorem. This proof is simpler than Mess's original proof and uses ideas of C. Pittet and T. Delzant.

Proposition 1.3 is obtained as a corollary of theorem 3.5, a more general result on large scale properties of complete Riemannian planar surfaces. The same method gives theorem 3.6, which characterizes quasi-homogeneous surfaces up to quasi-isometry in terms of bounds on the Gauss curvature. For more detailed statements, see section 3.2.

The paper is organized as follows. Sections 2 and 3 contain some general definitions and technical lemmas, as well as the statements of theorems 3.5 and 3.6. Sections 4–6 are devoted to the proofs of these theorems. In section 4, we introduce the notion of pseudo-triangulation, which is our most important tool. Definitions are given and technical results, including the crucial pseudo-triangulation theorem, are stated. In section 5, it is shown that the technical results imply theorems 3.5 and 3.6. The proof of the pseudo-triangulation theorem occupies all of section 6. It is strongly inspired by some work of Mess in [23].

Section 7 contains the proof of theorem 1.1 assuming the convergence group theorem. In section 8, we deal with surfaces with boundary and prove proposition 1.4 as a consequence of the more general, but rather technical theorem 8.1.

2. Nets and quasi-isometries

Let (X, d) be a metric space. We let $B(x, r)$ denote the metric ball of radius r centered at a point $x \in X$ and $N(Y, r)$ the metric r-neighborhood of a subset $Y \subset X$, i.e. $\bigcup_{x \in Y} B(x, r)$. We usually write d_Y for the metric on Y associated to the length structure induced by inclusion (which is in general different from the restriction of d to $Y \times Y$) and call it the *intrinsic metric* on Y. The diameter of Y with respect to the intrinsic metric is called the *intrinsic diameter* of Y. For instance, if $\xi \subset X$ is an embedded arc, its intrinsic diameter is equal to its length, which is in general different from the diameter of ξ as a subset of (X, d).

Let C_1, C_2 be positive real numbers. A subset $Y \subset X$ is said to be C_1-separated if balls of radius C_1 centered at distinct points of Y are disjoint. It is said to be C_2-*quasidense* if balls of radius C_2 centered at points of Y cover X, i.e. $N(Y, C_2) = X$. We say that Y is *uniformly discrete* if for all $r \geq 0$ there is a constant $n(r)$ such that for all $x \in X$, the cardinal of $Y \cap B(x, r)$ is not greater than $n(r)$.

Definition 2.1. *A* (C_1, C_2)-***net*** *in a metric space* X *is a subset* N *which is both* C_1-*separated and* C_2-*quasidense. A net is called* **uniform** *if it is uniformly discrete.*

Definition 2.2. *Let* (X_1, d_1) *and* (X_2, d_2) *be two metric spaces. We say that a map* $f : X_1 \to X_2$ *is a* **quasi-isometry** *if there exist positive real numbers* λ, C *such that:*

i *The inequality*

$$\lambda^{-1} d_1(x, x') - C \leq d_2(f(x), f(x')) \leq \lambda d_1(x, x') + C$$

holds for any $x, x' \in X_1$.

ii *The image of* f *is* C-*quasidense in* X_2.

It can be shown that if there is a quasi-isometry $f : X_1 \to X_2$, then there is a quasi-isometry $\bar{f} : X_2 \to X_1$ such that $f \circ \bar{f}$ (resp. $\bar{f} \circ f$) is a bounded distance from the identity of X_2 (resp. X_1). In that case, we say that X and Y are *quasi-isometric*. Being quasi-isometric is an equivalence relation between metric spaces. We assume the reader to be familiar with this concept and some basic properties that come readily from the definition. For example, if Y is a quasidense subset of a metric space X, the inclusion map $i : Y \to X$ is a quasi-isometry. In particular, if N_1 is a net in X_1 and N_2 is a net in X_2, then to prove that X_1 and X_2 are quasi-isometric it is enough to prove that N_1 and N_2 are quasi-isometric. For future reference, we state the following well-known lemma.

Lemma 2.3. [Rubber band principle, cf. [18, 11]] *Let* N *be a set, let* (X_1, d_1) *and* (X_2, d_2) *be geodesic metric spaces with one-to-one maps* $f_i : N \to X_i$ *such that*

$f_i(N)$ *is a net. Call* d_i *the distance functions on* N *induced by* f_i *for* $i = 1, 2$.
If d_1 *and* d_2 *are uniformly equivalent, i.e. there exist functions* $r \mapsto \alpha_1(r)$ *and*
$r \mapsto \alpha_2(r)$ *such that* $d_2(x, y) \leq \alpha_1(d_1(x, y))$ *and* $d_1(x, y) \leq \alpha_2(d_2(x, y))$ *for all*
$x, y \in N$, *then* X_1 *and* X_2 *are quasi-isometric.*

Proof. We need only show that the identity map between (N, d_1) and (N, d_2)
is a quasi-isometry. For $x, y \in N$, we choose a geodesic arc $\xi \subset X_1$ between
$f_1(x)$ and $f_1(y)$. We subdivide ξ into a concatenation of n subarcs $\xi_1 \cup_{x_1} \xi_2 \cup$
$\cdots \cup_{x_{n-1}} \xi_n$ with $n \leq d_1(x, y) + 1$ and each ξ_j has length at most 1. Let C
be a constant such that $f_1(N)$ is C-quasidense in X_1. Then every intermediate
point x_j is within C of some point $f_1(y_j)$. By the triangle inequality we get
$d_2(x, y) \leq n\alpha_1(2C + 1) \leq \alpha_1(2C + 1)(d_1(x, y) + 1)$. By a similar argument, one
gets $d_1(x, y) \leq \alpha_1(2C' + 1)(d_1(x, y) + 1)$ for some constant C′ coming from the
quasidensity of $f_2(N)$ in X_2.

3. Some notions of bounded geometry

3.1. Coarse bounded geometry

Let $\mathcal{G} = (N, \mathrm{E})$ be a (locally finite, unoriented) graph. We call *loops* the edges
whose vertices are equal. The underlying space of $\mathcal{G}$, still denoted $\mathcal{G}$, can be
endowed with the path metric where every edge has length 1. When no metric
is explicitly defined on $\mathcal{G}$, this one will be understood. The set of vertices N is
always a $(1/3, 1/2)$-net in $\mathcal{G}$.

Definition 3.1. *A graph* $\mathcal{G} = (N, \mathrm{E})$ *is said to have* **bounded geometry** *if there
is a constant* $n > 0$ *such that any vertex* $x \in N$ *belongs to at most* n *edges.*

If $\mathcal{G}$ has bounded geometry, then the net N is uniform. It turns out that if
X is a geodesic metric space, then X admits a uniform net iff there is a graph of
bounded geometry quasi-isometric to (X, d) (cf. [20]). This justifies the following
definition:

Definition 3.2. [cf. [3]] *We say that a metric space has* **coarse bounded geo-
metry** *if it admits a uniform net.*

Let X_1, X_2 be two metric spaces and $f : X_1 \to X_2$ be a quasi-isometry.
Assume that X_1 has coarse bounded geometry. The image of a uniform net
$N \subset X_1$ is then uniformly discrete and quasidense in X_2. Thus any maximal
1-separated subset of $f(N)$ is a uniform net in X_2. We have proved:

Lemma 3.3. *Being of coarse bounded geometry is invariant under quasi-isometry.*

Example 1. *Any complete Riemannian manifold whose Ricci curvature is bounded from below has coarse bounded geometry. This is a consequence of the Bishop-Gromov inequality.*

Example 2. *Every geodesic metric space quasi-isometric to a finitely generated group has coarse bounded geometry. Indeed, the Cayley graph of any finitely generated group has bounded geometry, hence coarse bounded geometry and we have seen that coarse bounded geometry is a quasi-isometry invariant.*

3.2. Quasi-homogeneity

Graphs of bounded geometry have an interesting property: given a real number $r > 0$, there are only finitely many isometry types of balls $B(x, r)$ such that $x \in N$. This leads us to the following definition, which was introduced by G. Mess [23] in a slightly different form.

Definition 3.4. *A metric space X is **quasi-homogeneous** if for all $r \geq 0$ there is a compact subset $Y_r \subset X$ such that for all $x \in X$ there is a point $y \in Y_r$ such that $B(x, r)$ is isometric to $B(y, r)$.*

Remarks.

- This property is **not** a quasi-isometry invariant, because one can take a quasi-homogeneous space X (e.g. the Euclidean plane $\mathbf{E}^2$), choose an unbounded sequence (x_n) and deform X in the neighborhood of each x_n in a different way without changing the quasi-isometry type.
- This notion is much weaker than what is called quasi-homogeneous in [21]. For instance, a space which is quasi-homogeneous in our sense need not have any other self-isometry than the identity.
- Note that nothing is assumed about the topology of the balls $B(x, r)$ and $B(y, r)$. There are contractible, quasi-homogeneous manifolds which are not uniformly contractible (cf. section 8).

Convention. For us, a *surface* will be a connected 2-manifold, in general non-compact.

Let R be an open surface. We say that R is *planar* if every simple closed curve in R is separating. We say that R is *planar at infinity* if there is a compact submanifold $K \subset R$ such that every component of $R - K$ is planar. (Intuitively, this means that R has at most a finite number of handles.) We are now in position to state our first results.

Consider the following properties, where R is an open surface and h a complete Riemannian metric on R.

(H) R admits a complete quasi-homogeneous Riemannian metric quasi-isometric to h.

(B) (R, h) has coarse bounded geometry.

(P) R is planar at infinity.

Theorem 3.5. *If* (R, h) *has property* (B) *and property* (P)*, then it also has property* (H)*.*

Since any geodesic metric space quasi-isometric to a finitely generated group satisfies property (B) (example 2), proposition 1.3 is a straightforward consequence of theorem 3.5.

Classical definitions of bounded geometry for Riemannian surfaces usually involve bounds on Riemannian invariants such as curvature or injectivity radius. The technique used to prove theorem 3.5 enables us to study the connections between those properties and quasihomogeneity when Riemannian surfaces are considered "up to quasi-isometry". Consider the following condition, which is apparently weakest possible:

(C) R admits a complete Riemannian metric with a lower curvature bound which is quasi-isometric to h.

Obviously (H) implies (C), because quasihomogeneity gives upper and lower bounds for any local invariant. It turns out that the converse is also true:

Theorem 3.6. *If* (R, h) *has property* (C) *then it also has property* (H)*.*

The main technical tool used in the proof of theorems 3.5 and 3.6 is the notion of pseudo-triangulation, which is introduced in the next section.

4. Pseudo-triangulations

4.1. Definitions

We begin by recalling a well-known definition:

Definition 4.1. *A **cell decomposition** of a smooth surface* R *is a collection* $\mathcal{D}$ *of pairs* (i, f) *where* $i \in \{0, 1, 2\}$ *and* $f : I^i \to R$ *is a smooth map, called an* i *-cell and abusively identified to its image, having the following properties:*

 i *Any* i *-cell for* $i < 2$ *is contained in some* $(i + 1)$ *-cell.*

 ii *Any two* 0 *-cells are disjoint.*

 iii *If* f *is an* i *-cell for* $i > 0$*, then* f *embeds* $\mathrm{Int}(I^i)$ *in the complement of the union* $\mathcal{D}^{(i-1)}$ *of the* $(i - 1)$ *-cells (henceforth called the* $(i - 1)$ *-**skeleton**) and maps* ∂I^i *to the* $(i - 1)$ *-skeleton.*

 iv *For* $i > 0$*, no two* i *-cells intersect each other outside the* $(i - 1)$ *-skeleton.*

 v *The* 2 *-skeleton is all of* R *.*

A 0-cell is simply a point. A 1-cell is either a simple arc or a simple closed

curve. The *combinatorial length* of a sequence of 1-cells is its cardinal. If f is a 2-cell, we can view its boundary as such a sequence, so we have a notion of *combinatorial boundary length*. Note that boundary 1-cells have to be counted "with multiplicity". We sometimes call f a n-*gon* when its combinatorial boundary length is n.

We say that $\mathcal{D}$ is *reduced* if no closed 1-cell bounds a 2-cell and no two 1-cells cobound a 2-cell. Thus $\mathcal{D}$ admits no 1-gon and admits a 2-gon only in the very special case where $\mathcal{D}$ consists of a single 2-cell with boundary a single one-sided curve. This can only happen in the projective plane, so we will not encounter it since we will work only with noncompact surfaces.

Thus the simplest 2-cells we will meet are 3-gons. A 3-gon is either a genuine triangle or some kind of "degenerate triangle" in which the number of vertices and/or edges is less than 3. There are four types of degenerate 3-gons, represented on figure 1: (1) two vertices, one loop and two edges; (2) two vertices, one loop and one double edge; (3) one vertex, three loops; (4) one vertex, one simple loop, one (one-sided) double loop.
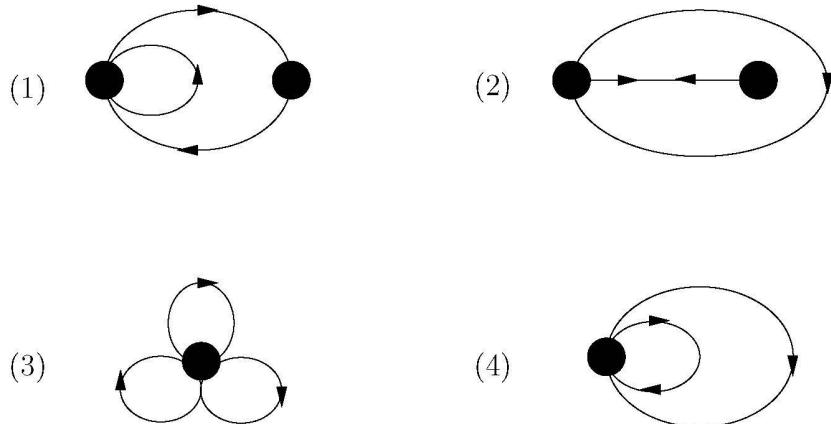
Figure 1. The four types of degenerate 3-gons.

Definition 4.2. *A cell decomposition is called a **pseudo-triangulation** if all 2-cells are 3-gons. A pseudo-triangulation is **nondegenerate** if no 2-cell is a degenerate 3-gon.*

Remark. Note that a nondegenerate pseudo-triangulation is not necessarily a triangulation in the usual sense, because the definition allows two 2-cells to have more than one edge in common.

Suppose now that the surface R is endowed with a geodesic metric d (for us, d will always be the distance function associated to a Riemannian metric,

possibly with conical singularities). The *modulus* $|\mathcal{D}|$ of a cell decomposition $\mathcal{D}$ (with respect to d) is then defined to be the (possibly infinite) supremum of the lengths of its 1-cells. We say that $\mathcal{D}$ is *uniformly locally finite* if for all $r \geq 0$ there is $n(r)$ such that any ball of radius r meets at most $n(r)$ cells.

Definition 4.3. *Let $\mathcal{D}$ be a pseudo-triangulation of a surface* R *with a (possibly singular) Riemannian metric. We say that $\mathcal{D}$ is* **uniform** *(with respect to the metric) if it satisfies the following conditions:*

i $\mathcal{D}^{(0)}$ *is a net.*

ii $\mathcal{D}$ *has finite modulus.*

iii $\mathcal{D}$ *is uniformly locally finite.*

Remarks.

- By (i) and (iii), $\mathcal{D}^{(0)}$ is a uniform net. However, (iii) is in general stronger than just requiring the uniformity of $\mathcal{D}^{(0)}$. This question will be dealt with by proposition 4.5.
- Let C_1, C_2 be constants such that $\mathcal{D}^{(0)}$ is a (C_1, C_2)-net. Then the intrinsic diameters of the non-closed 1-cells are bounded below by $2C_1$, because they contain at least two distinct points of N. Besides, (ii) gives a uniform upper bound for the intrinsic diameters of all 2-cells in the following way: the sum of the lengths of the 1-cells in the boundary of a given 2-cell is always bounded above by $3|\mathcal{D}|$, so the intrinsic diameter of a given 2-cell is always bounded above by $3|\mathcal{D}| + 2C_2$. In particular, any 1-cell or 2-cell meeting a compact K is contained in the $3|\mathcal{D}| + 2C_2$-neighborhood of K.

4.2. Main results on pseudo-triangulations

The following existence theorem for pseudo-triangulations is our main technical result:

Theorem 4.4. [Pseudo-triangulation theorem] *Let (R, h) be a complete open Riemannian surface. For each net N in (R, h), there exists a pseudo-triangulation $\mathcal{D}$ of* R *that has finite modulus with respect to h and whose 0-skeleton is exactly N.*

The proof of this theorem is deferred to section 6.

We intend to apply theorem 4.4 in the case where N is a uniform net. The next proposition shows that under some hypotheses on the topology or geometry of (R, h), the resulting pseudo-triangulation is actually uniform.

Proposition 4.5. *Let (R, h) be a complete open Riemannian surface. Assume either that* R *is planar at infinity or h has curvature bounded below. Let $\mathcal{D}$ be a*

pseudo-triangulation of R *that has finite modulus and such that* $\mathcal{D}^{(0)}$ *is a uniform net. Then* $\mathcal{D}$ *is uniformly locally finite, hence uniform.*

The end of this section is devoted to the proof of proposition 4.5. We start with the following lemma, which we state with somewhat general notation because it will also be needed in section 6.

Lemma 4.6. *Let* (F, h) *be a complete Riemannian surface (possibly with boundary), whose associated distance function is denoted by* d. *Let* C *be a positive real number and* Y *a* C *-quasidense closed subset of* F. *Let* $\xi \subset \mathrm{F}$ *be a simple closed curve which avoids* Y *and separates* F *into two closed subsurfaces* $\mathrm{F}_1, \mathrm{F}_2$. *Then the following assertions hold:*

  i *Given* $i \in \{1, 2\}$, *either* $\mathrm{F}_i \cap Y$ *is nonempty or* ξ *is* C *-quasidense in* F_i.
  ii *If both* $\mathrm{F}_1 \cap Y$ *and* $\mathrm{F}_2 \cap Y$ *are nonempty, then there exist* $x_1 \in \mathrm{F}_1 \cap Y$ *and* $x_2 \in \mathrm{F}_2 \cap Y$ *such that* $d(x_1, x_2) \leq 2\mathrm{C}$.

Proof. Assertion (i) is obvious from the definition of a quasidense subset. To prove assertion (ii), remark that by compacity of ξ, the distance between $\mathrm{F}_1 \cap Y$ and $\mathrm{F}_2 \cap Y$ is realized by a pair $(x_1, x_2) \subset \mathrm{F}_1 \cap Y \times \mathrm{F}_2 \cap Y$. If $d(x_1, x_2)$ were greater than $2\mathrm{C}$, then the midpoint of a minimizing geodesic between x_1 and x_2 would be at distance $> \mathrm{C}$ from Y, giving a contradiction.

Proof of proposition 4.5. Let N be the 0-skeleton of $\mathcal{D}$ and λ its modulus. We begin with the case where R is planar at infinity. Obviously, one may suppose that R is planar. We have to estimate the number of cells meeting a ball in function of the radius of the ball (and independently of its center). For the 0-cells, it is just the definition of uniform discreteness. If we do it for 1-cells, the estimate for 2-cells will follow, for a given 1-cell lies in the boundary of at most two 2-cells. Taking into account the fact that a 1-cell meeting a ball $\mathrm{B}(x, r)$ lies in the ball $\mathrm{B}(x, r + \lambda)$, we have reduced the problem to proving:

Claim. For all $r \geq 0$, there exists $\mathrm{C}(r)$ such that the number of 1-cells lying in a ball of radius r is at most $\mathrm{C}(r)$.

To prove the claim, we consider a ball $\mathrm{B}(x, r)$. Let $\mathcal{L}$ be the collection of 0-cells of $\mathcal{D}$ lying in $\mathrm{B}(x, r)$. By hypothesis, the cardinal of $\mathcal{L}$ is bounded independently of x. Generically, $\mathrm{B}(x, r)$ is a compact planar surface. Adding each component of $\mathrm{R} - \mathrm{B}(x, r)$ that is a disc containing no point of N, we get a subsurface Z, which by lemma 4.6 is contained in $\mathrm{B}(x, r + \mathrm{C}_2)$ where C_2 is the quasi-density constant of $N \subset \mathrm{R}$.

Consider the set $\mathcal{C}$ of arcs $\alpha : \mathrm{I} \to \mathrm{Z}$ such that

  i α is an embedding except possibly at its endpoints.
  ii $\mathrm{lg}(\alpha) \leq \lambda$.
  iii α meets N exactly at its endpoints (which therefore belong to $\mathcal{L}$).

We consider elements of $\mathcal{C}$ up to "proper" homotopy, saying that α and α' are equivalent when there is a homotopy between α and α' through elements of $\mathcal{C}$ (note that such a homotopy always fixes the endpoints since N is discrete.) Now in each equivalence class, there is at most one 1-cell of $\mathcal{D}$ (otherwise $\mathcal{D}$ would have a 2-gon), so we have reduced the claim to finding a uniform upper bound for the number of equivalence classes of $\mathcal{C}$.

This reduces to estimating the number of boundary components of the planar surface Z. Now by lemma 4.6, every component of R − Z contains a point $y \in N$ such that $d(x, y) \le r + 2C_2$. By uniformity of N, there is an upper bound for the number of such points y. This proves the claim, and therefore proposition 4.5 in the case where R is planar at infinity.

For the case where h has a curvature lower bound, we observe that the above argument would still work if, instead of being a planar surface, Z had uniformly bounded genus. But this will indeed be the case, for we have an upper bound for diam Z. This together with the lower bound on the curvature gives an upper bound on the area of Z, and then the Gauss-Bonnet formula yields a lower bound for the Euler-Poincaré characteristics of Z. Hence the above argument still applies. This completes the proof of proposition 4.5.

5. Proof of theorems 3.5 and 3.6

In this section, we prove theorems 3.5 and 3.6 assuming the existence theorem 4.4.

5.1. Regular piecewise Euclidean metrics

As we have remarked, graphs of bounded geometry are quasi-homogeneous. More generally, if X is a PL-manifold with a given triangulation $\mathcal{T}$ satisfying suitable (combinatorial) boundedness conditions, it can be endowed with a piecewise flat metric, which will be quasi-homogeneous. Below we make this construction precise in dimension 2. The reader can find a definition for n-dimensional PL-manifolds in a paper by O. Attie [1].

Definition 5.1. *Let X be a surface and $\mathcal{D}$ a nondegenerate, locally finite pseudo-triangulation of X. The **regular piecewise Euclidean metric** on $(X, \mathcal{D})$ is the path metric such that every 1-cell is a geodesic arc of length 1 and every 2-cell is isometric to a Euclidean equilateral triangle. It is uniquely defined up to choice of barycentric coordinates on each 2-cell.*

All regular piecewise Euclidean metrics are "Euclidean metrics with conical singularities" in the sense of M. Troyanov [25]. We call such a metric "regular" because the fact that all 2-cells are equilateral induces some restrictions, e.g. the cone angles are always multiples of $\pi/3$. We are mainly interested in the case

where the 1-skeleton is a graph of bounded geometry, or equivalently the cone angles are uniformly bounded above. In that case, the metric is quasi-homogeneous. Thus it is natural to consider the following property.

(E) R admits a complete, quasi-homogeneous regular piecewise Euclidean metric quasi-isometric to h.

Lemma 5.2. *Let* R *be an open surface with a geodesic metric d and $\mathcal{D}$ be a uniform pseudo-triangulation on (R, d). Let (X, d_X) be the 1-skeleton of $\mathcal{D}$ with the intrinsic metric. Then (X, d_X) is quasi-isometric to (R, d).*

Proof. Let N be the 0-skeleton. By hypothesis, N is a net in both (R, d), and (X, d_X), so we can use the rubber band principle. It is clear than $d(x, y) \leq d_X(x, y)$ for all $x, y \in N$. Conversely, let ξ be a geodesic arc in R between two points x, y. Call $D_1, \ldots, D_k$ the 2-cells which meet ξ. Since $\mathcal{D}$ is uniformly locally finite, there is an upper bound on k which depends only on $d(x, y)$. Decompose ξ into arcs $\xi_1 \cup \cdots \cup \xi_n$ where each ξ_i lies in some 2-cell $D_{j(i)}$. We can replace each ξ_i by a path in the boundary of $D_{j(i)}$ with the same endpoints. Piecing those paths together and getting rid of double points, we obtain an arc connecting x to y in X of length bounded above by $3k|\mathcal{D}|$. Then the rubber band principle gives the required quasi-isometry.

To prove that properties (B) and (P) (resp. property (C)) implies (H), we will go through (E). The following lemma proves that (E) implies (H).

Lemma 5.3. *Let* R *be a surface with a nondegenerate pseudo-triangulation $\mathcal{D}$ whose 1-skeleton has bounded geometry. Let d denote the corresponding regular piecewise Euclidean metric. Then (R, d) is quasi-homogeneous. Furthermore,* R *admits a complete (smooth) quasi-homogeneous Riemannian metric h' that is quasi-isometric to d.*

Proof. The bounded geometry condition on $\mathcal{D}^{(1)}$ implies that for a given $r \geq 0$, there are finitely many balls of radius r centered on $\mathcal{D}^{(0)}$, up to isometry. Since $\mathcal{D}^{(0)}$ is quasidense in (R, d), this proves that (R, d) is quasi-homogeneous. Now d corresponds to a smooth Riemannian metric in the complement of $\mathcal{D}^{(0)}$. Since there are finitely many balls of radius $1/10$ centered on $\mathcal{D}^{(0)}$ up to isometry, we can define a complete, quasi-homogeneous Riemannian metric h' on all of R by coherently smoothing the conical singularities in $N(\mathcal{D}^{(0)}, 1/10)$.

It remains to prove that (R, h') is quasi-isometric to (R, d). First remark that $\mathcal{D}$ is uniform with respect to each of these two metrics. By applying lemma 5.2 to both the singular and the smooth metric, we are left to consider $\mathcal{D}^{(1)}$ with the intrinsic metrics induced by respectively h' and d, which are easily seen to be quasi-isometric to each other.

Proof of theorems 3.5 and 3.6. We have already seen that property (E) implies property (H). Thus we need only prove the following two facts: property (C) implies property (E), and properties (B) and (P) together imply property (E).

We are going to prove these two facts at once: assume that either property (C) holds or properties (B) and (P) hold. Example 1 shows that in both cases property (B) holds, i.e. (R, h) has coarse bounded geometry. Let N be a uniform (C_1, C_2)-net in (R, h). By theorem 4.4, (R, h) admits a pseudo-triangulation $\mathcal{D}$ of finite modulus such that $\mathcal{D}^{(0)} = N$. Since either property (C) or property (P) holds, we can apply proposition 4.5, so $\mathcal{D}$ is uniform. Thus our task is to prove that the existence of this uniform pseudo-triangulation implies property (E).

Applying lemma 5.2, we obtain a quasi-isometry between (R, h) and the 1-skeleton of $\mathcal{D}$ with the intrinsic metric, which we denote (X, d_X). X is a graph and d_X a geodesic metric such that all loops have length bounded above by $|\mathcal{D}|$ and all edges that are not loops have length bounded above by $|\mathcal{D}|$ and below by $2C_1$. Hence we can renormalize d_X to get a metric d'_X quasi-isometric to d_X and such that all loops have length 2 and all edges that are not loops have length 1. (Note that d'_X and d_X cannot be made bi-Lipschitz in general since there could exist arbitrarily short loops.)

We will now subdivide X to get a new graph X' quasi-isometric to X and which is the 1-skeleton of a pseudo-triangulation $\mathcal{D}'$ of R without degenerate 3-gons. This is done as follows. First add a vertex in the middle of each loop. Then consider a degenerate 3-gon A. If A is of type (1) or (2), connect by an edge the new vertex to the vertex opposite to the loop. If A is of type (3), connect the three new vertices to each other pairwise. If A is of type (4), do the same thing as for type (3), except that one of the new vertices is double, so one of the three new edges is actually a loop l which belongs to two new 3-gons A_1 and A_2 of type (1). Subdivide l by adding a new vertex x and connect x to the opposite vertices in A_1 and A_2. Finally in this case we have subdivided A into six nondegenerate cells by adding four vertices and six edges.

Let d' be the regular piecewise Euclidean metric associated to $\mathcal{D}'$. Since X' has bounded geometry, (R, d') is quasi-homogeneous by lemma 5.3. By lemma 5.2, (R, d') and X' are quasi-isometric. Finally (R, d') is quasi-isometric to (R, h) by construction and we are done.

6. Proof of the pseudo-triangulation theorem

6.1. Blowing holes

Throughout this section, (R, h) is a complete Riemannian surface and $N \subset R$ is a (C_1, C_2)-net in R. In order to prove the pseudo-triangulation theorem, our goal is to produce a graph embedded in R, with vertex set N, which will be the 1-skeleton of the pseudo-triangulation. One might think of looking at pairs

$(x, y) \in N$ such that x, y are distinct and close to each other, and connect such pairs of points by geodesic arcs. But these arcs are to be chosen carefully: on the one hand, they must not intersect each other outside N, and on the other hand there must be sufficiently many of them to cut R into 3-gons.

For technical reasons, we will not actually connect points of N by geodesic arcs; instead we replace each $x \in N$ by a small topological disc around x and work in the complement of the union of the interiors of these discs, slightly perturbing the metric so that the resulting surface has strictly convex boundary, which ensures the existence of geodesic arcs.

More precisely, we choose for each $x \in N$ a triple $(r_1(x), r_2(x), r_3(x)) \in \mathbf{R}^3$ such that $0 < r_1(x) < r_2(x) < r_3(x) < \min(C_1/2, \mathrm{inj}(x))$ where $\mathrm{inj}(x)$ is the injectivity radius of h at x. Define $\mathrm{D}_x := \mathrm{B}(x, r_2(x))$, $\mathrm{C}_x := \partial \mathrm{D}_x$, $\mathrm{A}_x := \mathrm{B}(x, r_3(x)) - \mathrm{Int}\,\mathrm{B}(x, r_1(x))$. We modify h by a bump function with support in A_x so that the circle C_x is strictly convex in $\mathrm{R} - \mathrm{Int}\,\mathrm{D}_x$ and call h_N the modified metric.

Define $\mathrm{R}(N) := \mathrm{R} - \bigcup_{x \in N} \mathrm{Int}(\mathrm{D}_x)$. The boundary of $\mathrm{R}(N)$ is the disjoint union of the circles C_x, so it has strictly convex boundary with respect to h_N. By choosing r_2 sufficiently small, we may assume that there is a uniform upper bound C_3 for the h_N-lengths of the C_x's.

Let d denote the distance function on $\mathrm{R}(N)$ induced by h_N. There are constants $\mathrm{C}'_1, \mathrm{C}'_2 \geq 0$ such that $d(\mathrm{C}_x, \mathrm{C}_y) \geq \mathrm{C}'_1$ for all $x \neq y$ and $\partial \mathrm{R}(N)$ is C'_2-quasidense in $\mathrm{R}(N)$. From now on, we will work in the metric space $(\mathrm{R}(N), d)$.

We will use *shortest arcs* in $\mathrm{R}(N)$, that is by definition proper immersions $\mathsf{a} : \mathrm{I} \to \mathrm{R}(N)$ which are *essential* (i.e., not properly homotopic into $\partial \mathrm{R}(N)$) and are shortest in their proper homotopy classes. We do some usual abuses of language, such as omitting the word "proper" when talking of proper arcs and proper homotopies, and identifying a with its image in $\mathrm{R}(N)$.

In each (proper) homotopy class of essential arcs, there is a (generally not unique) shortest element. We say that two arcs $\mathsf{a}_1, \mathsf{a}_2$ are *homotopically disjoint* if there are arcs $\mathsf{a}'_1, \mathsf{a}'_2$ respectively homotopic to a_1 and a_2 and satisfying $\mathsf{a}'_1 \cap \mathsf{a}'_2 = \emptyset$. The following proposition contains the facts we will use about shortest arcs. The first and the second statements follow from the work of Freedman, Hass and Scott [13], while the third is a straightforward exercise using the exchange/roundoff trick.

Proposition 6.1.

   i *Any shortest arc that is properly homotopic to an embeddeding is an embedding.*

  ii *Two shortest arcs that are homotopically disjoint are disjoint or equal.*

 iii *Let $\mathsf{a}_1, \mathsf{a}_2$ be embedded shortest arcs such that there exists a (not properly) embedded disc D containing the images of $\mathsf{a}_1, \mathsf{a}_2$. Then a_1 and a_2 intersect in at most one point.*

We will write $\eta_1 \cup_w \eta_2$, or simply $\eta_1 \cup \eta_2$ for the concatenation of two arcs $\eta_1, \eta_2 \subset R(N)$ with a common endpoint $w \in R(N)$.

6.2. Geodesic divides

We now give definitions concerning what we call a geodesic divide. The reader should keep in mind that such an object is meant to correspond to the 1-skeleton of a cell decomposition, which explains some of the terminology.

A *geodesic divide* is a set E of shortest arcs called *edges*. We say that E is *embedded* if each edge is embedded and no two edges intersect each other. It is *reduced* if no two edges are homotopic. Its *modulus* $|E|$ is the supremum of the lengths of its edges. The *pieces* of $R(N)$ split along E (or simply E-pieces) are the components of $R(N) - \bigcup_{a \in E} a$. As in the case of cell decompositions, there is a notion of combinatorial boundary length for E-pieces (but this may be infinite). A simply-connected E-piece having combinatorial boundary length n will be called a n-gon.

Given two geodesic divides E, E', we say that E' is an *extension* of E if $E \subset E'$. If both E and E' are embedded, this amounts to insisting that the union of the edges of E is contained in the union of the edges of E'.

Proposition 6.2. *Define* $C_4 := 2C_2'$. *There exists an embedded, reduced geodesic divide* E_1 *in* $R(N)$ *such that* $|E_1| \leq C_4$ *and every* E_1-*piece is simply-connected.*

Proof. The idea of the proof is actually quite simple. Let $\mathcal{L}$ be a list of all nontrivial free homotopy classes of loops in $R(N)$ which admit an embedded element. We say that an arc a *kills* a class $[\xi] \in \mathcal{L}$ if a is not homotopically disjoint from ξ. If one takes the elements of $\mathcal{L}$ in an arbitrary order and kills them one by one by adding properly embedded arcs, one could expect that the use of shortest arcs guarantees that the divide obtained is embedded and has finite modulus. However, this does not work nicely. In particular, it does not seem easy to prove that the resulting divide is embedded, so our approach is somewhat different.

We say that an essential arc a *strongly kills* $[\xi]$ if its intersection number modulo 2 with ξ is 1 (this is independent of the choice of ξ in its homotopy class). Clearly this is (as the name tells) stronger than to say that a kills $[\xi]$ in the above sense. Let $\mathcal{L}' = [\xi_1], [\xi_2] \dots$ be a list of all elements of $\mathcal{L}$ that admit a strongly killing arc. We will first construct an embedded, reduced geodesic divide E_1 such that all elements of $\mathcal{L}'$ are strongly killed by an edge of E_1 and then show that in fact every element of $\mathcal{L}$ is killed.

The divide E_1 will be constructed inductively. Consider the family of all arcs that strongly kill $[\xi_1]$. By a standard compacity argument using the Ascoli-Arzela theorem and the fact that any such arc meets ξ_1, there exists an arc a_1 that has minimal length in this family. In particular, a_1 is a shortest arc (i.e. minimizes

length in its homotopy class).

Delete from the list $[\xi_1], [\xi_2], \dots$ all classes that are strongly killed by a_1. Then take the first class $[\xi_{n_2}]$ on the list which is *not* strongly killed by a_1. Choose an arc a_2 shortest among all arcs that strongly kill $[\xi_{n_2}]$ and so on. This construction gives a sequence of essential arcs a_n which defines a divide E_1. Each a_n is clearly shortest in its homotopy class, thus a geodesic (in the Riemannian sense). Moreover if a_n and a_m are homotopic, then the set of classes strongly killed by a_n is the same as a_m, so $n = m$. Hence E_1 is reduced.

Assume there is a $[\xi] \in \mathcal{L}$ which is not killed by any edge of E_1. Such a ξ must separate $R(N)$ for otherwise there would be a closed loop intersecting it transversely in one point, from which it is easy to construct a strongly killing arc. Let U, V be the components of $R(N)$ split along ξ. If both U and V contained a boundary circle of $R(N)$, an arc connecting them would strongly kill $[\xi]$. So we may assume that U does not meet $\partial R(N)$. In particular, U is compact by lemma 4.6.

Now U cannot be a disc, so U has at least one handle or one cross-cap. Let $\xi' \subset U$ be a nonseparating curve. Then $[\xi']$ appears on the list $\mathcal{L}'$, so some edge a of E_1 strongly kills $[\xi']$. But if a did not kill ξ, some arc a' homotopic to a would lie in V, contradicting the fact that a kills ξ'.

To complete the proof of proposition 6.2, we still have to show that the a_n's are embedded, do not intersect one another and have length $\leq C_4$.

If some a_n is not embedded, write $a_n = \eta_1 \cup_x \eta_2 \cup_x \eta_3$ for some multiple point x. By construction, there is a homotopy class $[\xi_p]$ such that a_n is shortest among arcs that strongly kill $[\xi_p]$. By general position, we may assume that ξ_p is self-transverse, intersects a_n transversely and that ξ_p avoids multiple points of a_n and vice-versa (in particular $x \notin \xi_p$). By hypothesis, $\#(\xi_p \cap a_n)$ is odd. If $\#(\xi_p \cap \eta_2)$ is even, then $\#(\xi_p \cap (\eta_1 \cup \eta_3))$ is odd, so smoothing out $\eta_1 \cup \eta_3$ yields a shorter arc killing $[\xi_p]$, a contradiction, so $\#(\xi_p \cap \eta_2)$ is odd. By symmetry we may assume that $\lg(\eta_1) \leq \lg(\eta_3)$. Rounding the corner of $\eta_1 \cup \eta_2 \cup \eta_1^{-1}$ yields an arc shorter than a_n and strongly killing $[\xi_p]$, again a contradiction. Thus every edge of E_1 is embedded.

If two edges a_n, a_m intersect in some point w, the intersection must be transverse. Write $a_n = \eta_1 \cup_w \eta_2$ and $a_m = \eta_3 \cup_w \eta_4$. Up to changing the indices, we may assume that η_3 is shortest among the η_i's, so in particular $\lg(\eta_3) \leq \lg(\eta_1)$ and $\lg(\eta_3) \leq \lg(\eta_2)$. Let $[\xi_p]$ be a homotopy class such that a_n is shortest among arcs that strongly kill $[\xi_p]$. Assume ξ_p is self-transverse and avoids w. Then by counting the intersection points of ξ_p with η_1 and η_2 one sees that exactly one of $\eta_3 \cup \eta_1$ and $\eta_3 \cup \eta_2$ strongly kills $[\xi_p]$, and we get a contradiction by rounding the corner of a suitable arc. We have shown that E_1 is embedded. Note that the same argument proves the following lemma.

Lemma 6.3. *Let a_n be an edge of E_1. Then the inequality*

$$\min(\lg(\eta_1), \lg(\eta_2)) \leq \min_{z \in N} d(C_z, w)$$

holds for every decomposition $a_n = \eta_1 \cup_w \eta_2$.

Remark. In particular, for all $w \in a_n$, one of the two subarcs η_1, η_2 minimizes the distance between its endpoints.

There remains to prove that $|E_1| \leq C_4$. Let a_n be an edge of E_1 . Write $a_n = \eta_1 \cup_w \eta_2$, where w is the midpoint of a_n . Since $\partial R(N)$ is C'_2 -quasidense, w must be C'_2 -close to some boundary component C_z , so the above lemma proves that w is C'_2 -close to one (hence both) of the endpoints of a_n . Therefore a_n has length at most $2C'_2 = C_4$. This completes the proof of proposition 6.2.

6.3. Further splitting

The goal of this section is to add sufficiently many edges to the divide E_1 so that its complement consists of topological open discs of uniformly bounded diameter. Since adding an edge in a simply-connected piece cuts it into two simply-connected pieces, we will restrict attention to divides E such that all E-pieces are simply-connected. This is not essential, but simplifies some statements.

Let X be a simply-connected piece of $R(N)$ split along some embedded, reduced geodesic divide. Its closure in $R(N)$ need not be simply-connected, for it could contain 1-sided loops or 2-sided edges both of whose sides lie in X . To avoid this problem, we consider the inverse image of X in the universal cover of $R(N)$: since X is simply-connected, it is a disjoint union of infinitely many copies of X . Let $\tilde{X}$ be one such component and $\hat{X}$ be the closure of $\tilde{X}$. Then $\partial \hat{X}$ is what we want to consider as the "proper" boundary of X .

To describe $\hat{X}$, we need to distinguish two cases: if $\hat{X}$ is compact, then $\partial \hat{X}$ is topologically a circle, consisting of arcs $a_0 \cup c_0 \cup a_1 \cup c_1 \cup \ldots \cup a_n \cup c_n$ glued together at their endpoints where the a_i 's project onto elements of E and the c_i 's onto subarcs of $\partial R(N)$. We will call the a_i 's *frontier arcs* and the c_i 's *boundary arcs*. Topologically, $\hat{X}$ is a disc. If $\hat{X}$ is noncompact, then it is homeomorphic to a disc minus some arcs in the boundary. Each component of $\partial \hat{X}$ is then a line with a similar decomposition $\ldots \cup a_{-1} \cup c_{-1} \cup a_0 \cup c_0 \cup a_1 \cup \ldots$ into frontier arcs and boundary arcs.

Given a constant $C \geq 0$, we say that X is C-*thin* if every point of $\hat{X}$ can be connected to some point of $\partial \hat{X}$ by a path in $\hat{X}$ of length $\leq C$ (measured with respect to the lift of h_N to the universal cover of $R(N)$). In other words, $\partial \hat{X}$ is C-quasidense in $\hat{X}$ with the intrinsic metric. This is an important property since it allows E-pieces to be cut into smaller pieces of uniformly bounded diameter. Fortunately, it is automatic thanks to the following lemma.

Lemma 6.4. *Let E be an embedded, reduced geodesic divide and X a simply-connected E-piece. Then X is C'_2-thin.*

Proof. Any point $x \in \operatorname{Int} \hat{X}$ projects to a point $y \in X$. There is a path $\xi \subset \mathrm{R}(N)$ of length $\leq \mathrm{C}'_2$ connecting y to $\partial \mathrm{R}(N)$. Let $\tilde{\xi}$ be the lift of ξ starting at x. Then if ξ cuts some edge $\mathsf{a} \in \mathrm{E}$ which lifts to a subarc of $\partial \hat{X}$, then $\tilde{\xi}$ cuts this subarc of $\partial \hat{X}$, so $d(x, \partial \hat{X}) \leq \mathrm{C}'_2$. Otherwise, ξ lies in X and the same conclusion holds.

Lemma 6.4 applies in particular to our divide E_1 and any embedded, reduced geodesic divide which is an extension of E_1.

Proposition 6.5. *There are constants* $\mathrm{C}_5, \mathrm{C}_6 \geq 0$ *and an embedded, reduced geodesic divide* E_2 *such that* $|\mathrm{E}_2| \leq \mathrm{C}_5$ *and for every* E_2 *-piece* X, $\hat{X}$ *is homeomorphic to a disc and its intrinsic diameter is* $\leq \mathrm{C}_6$.

First we need a technical lemma which ensures that, under some natural conditions, adding an edge to an embedded, reduced geodesic divide yields a divide with the same properties.

Lemma 6.6. *Let* E *be an embedded, reduced geodesic divide such that all* E *-pieces are simply-connected. Let* a *be an arc properly embedded in* $\mathrm{R}(N)$ *and disjoint from the edges of* E. *Let* X *be the* E *-piece that contains* a. *Assume that* a *is not parallel rel* $\partial \mathrm{R}(N)$ *in* $\bar{X}$ *to an arc in* $\partial \mathrm{R}(N)$.

  i a *is essential.*
 ii *Let* $\tilde{\mathsf{a}}$ *be a shortest arc homotopic to* a. *If* a *is not parallel rel* $\partial \mathrm{R}(N)$ *in* $\bar{X}$ *to an edge of* E, *then* $\tilde{\mathsf{a}}$ *lies in* X *and adding it to* E *yields an embedded, reduced divide.*

Proof. To prove assertion (i), observe that since a is not parallel rel $\partial \mathrm{R}(N)$ in $\bar{X}$ to an arc in $\partial \mathrm{R}(N)$, a is not parallel into $\partial \mathrm{R}(N)$, because if it were, the product region could not contain an edge of E and thus would lie entirely in X.

We now prove assertion (ii): by proposition 6.1, $\tilde{\mathsf{a}}$ is embedded and does not intersect any edge of E. If $\tilde{\mathsf{a}}$ did not lie in $\bar{X}$, it would be disjoint from a, so a and $\tilde{\mathsf{a}}$ would be parallel, and the same argument as above involving a product region shows that a would be homotopic to an edge, a contradiction. Hence $\tilde{\mathsf{a}}$ does lie in $\bar{X}$, and again by the same argument it is not homotopic to an edge of E. It follows that the geodesic divide obtained by adding $\tilde{\mathsf{a}}$ to E is embedded and reduced.

Proof of 6.5. Here is a sketch of the proof: we start with the divide E_1 given by proposition 6.2. We construct by transfinite induction an increasing family of divides E'_λ such that $\mathrm{E}'_0 = \mathrm{E}_1$. At each step, we choose a piece X and split $\partial \hat{X}$ into frontier arcs a_n and boundary arcs c_n as explained before. The thinness of X allows us to find two boundary arcs $\mathsf{c}_{n_1}, \mathsf{c}_{n_2}$ such that $n_1 - n_2 > 1$ and there is an arc γ of controlled length in $\tilde{X}$ connecting them. However, the bound on

$\lg(\gamma)$ may depend on the lengths of the neighboring frontier arcs, so since we want to iterate the construction, we need to distinguish carefully between the edges of E_1 and the new edges. The basic strategy is to choose a frontier arc, say a_0, and find a collection of *cutting arcs* $\gamma_1, \ldots, \gamma_p$ which are sufficiently far from a_0 and split $\hat{X}$ into $p+1$ components, one of which has controlled diameter, and such that the others have only one cutting arc in their boundary. Adding the projections to X of the cutting arcs to the divide yields a larger divide E'. The cutting arcs become frontier arcs of E'-pieces, so we can iterate the construction to cut those new pieces, taking a_0 to be the corresponding new frontier arc.

Define $C_5 := 2C_2' + 2C_4$ and $C_6 := 3C_5 + 2C_3 + C_2' + 1$. All the divides E_λ' will have the following properties: E_λ' is an extension of every divide E_μ' for $\mu < \lambda$; the modulus of E_λ' is at most C_5 and for every piece X of $R(N)$ split along E_λ', either the intrinsic diameter of X is at most C_6, or X admits at most one frontier arc which is not already an edge of E_1.

Clearly $E_0' = E_1$ has all these properties. Next is the description of the inductive step. Assume that some piece X of E_λ' is noncompact or has intrinsic diameter greater than C_6. If necessary, we number $\partial \hat{X}$ so that a_0 is the lift of the new arc. Let U denote the arc $\mathsf{c}_{-1} \cup \mathsf{a}_0 \cup \mathsf{c}_0$ and x, y denote its endpoints. Consider the metric neighborhood $Y = N(U, C_5) \subset \hat{X}$ with respect to the intrinsic metric on $\hat{X}$. By hypothesis, $Y \neq \hat{X}$. Let F be a small neighborhood of Y in $\hat{X}$ which lies in $N(U, C_5 + 1)$ and is a compact, connected planar surface. Let γ be the boundary component of F which meets $\partial \hat{X}$. Clearly, γ contains U. Moreover, γ does not fill $\partial \hat{X}$, for $\gamma = \partial \hat{X}$ would imply $\operatorname{diam}(X) \leq C_5 + C_2' + 1$.

We can split $(\gamma - U) \cup \{x, y\}$ into $\gamma_1 \cup \xi_1 \cup \gamma_2 \cup \cdots \cup \xi_{p-1} \cup \gamma_p$ where $\gamma_i \subset \partial \hat{X}$ and $\xi_i \subset \hat{X} - N(U, C_5)$. Now X is C_2'-thin, so every point of any ξ_i is a distance at most C_2' of some point of $\partial \hat{X}$. Let $Z_1, \ldots, Z_q$ be the components of the closure of $\partial \hat{X} \cap N(\gamma, C_2') - U$. Each γ_i lies in some $Z_{j(i)}$. Consider the (abstract) graph $\mathcal{G}$ whose vertices are the Z_j's and where there is an edge between Z_j and Z_k iff $N(Z_j, C_2') \cap N(Z_k, C_2')$ is not contained in $N(U, C_5)$.

Lemma 6.7. *There is an edge path in $\mathcal{G}$ connecting $Z_{j(i)}$ and $Z_{j(i+1)}$ for all i.*

Proof. If $j(i) = j(i+1)$, it is obvious. Otherwise we can write

$$\xi_i = (\xi_i \cap N(Z_{j(i)}, C_2')) \cup \bigcup_{k \neq j(i)} (\xi_i \cap N(Z_k, C_2')).$$

This gives ξ_i as a union of two nonempty closed subsets. Since ξ_i is connected, this union cannot be disjoint. This means that there is a $k \neq j(i)$ such that $\xi_i \cap N(Z_{j(i)}, C_2') \cap N(Z_k, C_2') \neq \emptyset$. In particular, $Z_{j(i)}$ and Z_k are connected by an edge in $\mathcal{G}$. If $j(i+1) = k$, we are done. If not, we can repeat the same argument, putting together $j(i)$ and k on one side and everything else on the other side. We get an edge between $j(i)$ or k and some k'. After finitely many steps, we obtain an edge path in $\mathcal{G}$ connecting $Z_{j(i)}$ and $Z_{j(i+1)}$.

Call W the union of all the a_n's and c_n's that do not lie entirely within U and meet $N(Z_j, C'_2)$ for at least one j. The components of W will be denoted $W_1, \ldots, W_t$. Each Z_j is contained in some $W_{k(j)}$.

Lemma 6.8. *Given indices* j, j' *such that there is an edge of* $\mathcal{G}$ *between* Z_j *and* $Z_{j'}$, *either* $k(j) = k(j')$ *holds or there are boundary arcs* $c_n \subset W_{k(j)}$, $c_m \subset W_{k(j')}$ *with* $|n - m| \geq 2$ *and a shortest arc of length* $\leq C_5$ *connecting* c_n *to* c_m *in* $\hat{X}$.

Proof. By hypothesis, there is a point $z \in \hat{X}$ such that $d(z, Z_j) \leq C'_2$, $d(z, Z_{j'}) \leq C'_2$ and $d(z, U) > C_5$. Assume that $k(j) = k(j')$ does not hold. Then there is an arc b of length $\leq 2C'_2$ connecting Z_j to $Z_{j'}$. The endpoints of b lie either on c_n's or on a_n's. If they both lie on c_n's, we are done. Otherwise we can extend b along the boundary, getting an arc b' of length at most $2C'_2 + 2C_4 = C_5$ connecting some $c_n \subset W_{k(j)}$ to some $c_m \subset W_{k(j')}$ with $n \neq m$. If $|n - m| < 2$, the arc we are looking for exists already: it is one of the a_i's. Otherwise, lemma 6.6 applied to b' gives the required arc.

Combining the two previous lemmas, we get the following conclusion: up to reindexing the W_k's, there exists a finite sequence $b'_1, \ldots, b'_r$ having the following properties:

i Every b'_i is a shortest arc and its length is at most $\leq C_5$.

ii W_1 contains x, W_r contains y and for all i, b'_i connects boundary arcs $c_{n(i)} \subset W_i$ and $c_{m(i)} \subset W_{i+1}$.

The problem is that some of the b'_i's might intersect each other. This is the purpose of the next lemma.

Lemma 6.9. *Possibly after reindexing again the* W_k's, *there is a sequence* $b''_1, \ldots$ b''_s *satisfying the same properties, and in addition the* b''_i's *are pairwise disjoint.*

Proof. Let $(b''_1, \ldots, b''_s)$ be a finite sequence of arcs of length $\leq C_5$ such that each b''_i is either a frontier arc or a cutting arc and connects boundary arcs $c_{n(i)} \subset W_{p(i)}$ and $c_{m(i)} \subset W_{q(i)}$ with the properties that $W_{p(1)}$ contains x, $W_{q(s)}$ contains y and $p(i+1) = q(i)$. The *complexity* of the sequence is defined as the pair (s, η), lexicographically ordered, where s is the cardinal of the sequence and η is the sum of the lengths of the b''_i's.

The work already done shows that such a sequence exists. Since the b''_i's are shortest arcs, the set of all complexities of sequences consisting only of shortest arcs and of given cardinal is finite, so we can choose one of least complexity. To show that the elements of such a sequence do not intersect each other, we use again exchange/roundoff arguments.

Assume there are indices $i < j$ such that $b''_i \cap b''_j$ is nonempty. Then b''_i and b''_j are cutting arcs and by lemma 6.1(iii) they intersect transversely in one single

point. Let w be this intersection point and $\mathsf{d}_1, \mathsf{d}_2, \mathsf{d}_3, \mathsf{d}_4$ be the arcs connecting w to respectively $\mathsf{c}_{n(i)}$, $\mathsf{c}_{m(i)}$, $\mathsf{c}_{n(j)}$ and $\mathsf{c}_{m(j)}$. If $j - i > 1$, we find a sequence of cardinal $< s$ by replacing b_j'' by a shortest arc homotopic to $\mathsf{d}_1 \cup \mathsf{d}_4$ (connecting $\mathsf{c}_{(n_i)}$ to $\mathsf{c}_{(m_j)}$ and suppressing every b_k'' for $i < k < j$. Thus j must be equal to $i + 1$ and $q(i) = p(j)$ holds.

If $\lg(\mathsf{d}_2) \geq \lg(\mathsf{d}_3)$, replace b_i'' by a shortest arc homotopic to $\mathsf{d}_1 \cup \mathsf{d}_3$. Otherwise replace b_j'' by a shortest arc homotopic to $\mathsf{d}_2 \cup \mathsf{d}_4$. In each case the resulting sequence has same cardinal and smaller η, hence smaller complexity.

Since $\operatorname{diam} X > \mathrm{C}_6$, there is at least one cutting arc among the b_i'''s given by lemma 6.9. Projecting the cutting arcs to X gives shortest arcs, and by lemma 6.6, adding those arcs to E'_λ yields an embedded, reduced geodesic divide satisfying the required properties.

The result follows by transfinite induction: the construction has to stop for some countable ordinal λ. The corresponding E'_λ has the additional property that all E'_λ-pieces have intrinsic diameter at most C_6.

6.4. Getting the pseudo-triangulation

Proposition 6.10. *There is an embedded, reduced geodesic divide* E_3 *such that* $|\mathrm{E}_3| \leq \mathrm{C}_6$ *and all* E_3 *-pieces are* 3 *-gons.*

Proof. Start with the divide E_2 given by proposition 6.5. Recall that the fact that E_2 is reduced implies that there are no 1-gons, and that 2-gons can only occur in a very special case in the projective plane, which is excluded here by the noncompacity of R. Thus all E_2-pieces have combinatorial boundary length at least 3.

For each E_2-piece X having combinatorial boundary length $n > 3$, consider its boundary $\mathsf{a}_0 \cup \mathsf{c}_0 \cup \mathsf{a}_1 \cup \mathsf{c}_1 \cup \ldots \cup \mathsf{a}_{n-1} \cup \mathsf{c}_{n-1}$ and choose a finite sequence of pairwise disjoint arcs $b_2, \ldots, \mathsf{b}_{n-2}$ such that each b_i connects c_0 to c_i. Project those arcs to X and choose a shortest arc $\tilde{\mathsf{b}}_i$ in the homotopy class of the projection for each i. By proposition 6.1, the $\tilde{\mathsf{b}}_i$'s are embedded, pairwise disjoint and their lengths are certainly bounded by the intrinsic diameter of X, hence by C_6. The $\tilde{\mathsf{b}}_i$'s cut X into a finite number of (possibly degenerate) 3-gons, so the proof is complete.

Connecting each endpoint of an edge of E_3 to the center of the corresponding D_x gives a cell decomposition $\mathcal{D}$ of R in which all 2-cells are 3-gons and whose 0-skeleton is exactly N. All 1-cells have length (with respect to h) bounded by $\lambda = \mathrm{C}_6 + \mathrm{C}_1$. This completes the proof of the pseudo-triangulation theorem.

7. Proof of theorem 1.1

The goal of this section is to prove theorem 1.1 modulo the convergence group theorem. For this, it is enough to prove:

Proposition 7.1. *Let Γ be a finitely generated group which is quasi-isometric to a complete quasi-homogeneous Riemannian plane* (R, h). *Then either Γ is virtually* $\mathbf{Z}^2$ *or* R *is quasi-isometric to the hyperbolic plane* $\mathbf{H}^2$.

Indeed, let Γ be a finitely generated group quasi-isometric to the hyperbolic plane. Then Γ is word-hyperbolic with boundary S^1, hence a uniform convergence group on S^1. By the convergence groups theorem [26, 14, 5], it is virtually the fundamental group of a hyperbolic surface.

Proposition 7.1 is one of the results of G. Mess's paper [23]. Our proof, outlined below, is completely different and considerably shorter.

Let R be a Riemannian plane. By Riemann's uniformization theorem, R has the conformal structure of either the Euclidean plane $\mathbf{E}^2$ or the hyperbolic plane $\mathbf{H}^2$. So the proof of proposition 7.1 falls into two cases, which we call respectively the "Euclidean" and the "hyperbolic" case. The two cases will be dealt with independently and by completely different methods. Next is a short discussion of these methods.

In the Euclidean case (i.e. under the assumption that the plane R is conformally equivalent to $\mathbf{E}^2$) the goal is to prove that Γ has subcubic growth, which allows us to conclude using M. Gromov's celebrated theorem on groups of polynomial growth. The bound on the growth of Γ is obtained by comparing the isoperimetric dimensions of R and $\mathbf{E}^2$.

In the hyperbolic case (i.e. when R is conformally equivalent to $\mathbf{H}^2$) the proof is based on the conformal version of a renormalization principle, which gives bounds on the dilatation of conformal diffeomorphisms between certain Riemannian surfaces. In our case, it gives a lower bound for g in terms of the hyperbolic metric. To get an upper bound, we use a trick due to Gromov (cf. [6]).

7.1. The Euclidean case

In this section, we assume that R is conformally equivalent to $\mathbf{E}^2$ and prove that Γ is virtually $\mathbf{Z}^2$. We fix a finite set of generators for Γ and let $V_\Gamma(n)$ denote the growth function of Γ with respect to these generators.

We need some more notation. If $\Omega \subset R$ is a bounded domain with sufficiently smooth boundary, we denote by $|\Omega|$ its area and $|\partial\Omega|$ the length of its boundary. It is convenient to use the same notation for the corresponding discrete notions in Γ, i.e. when Ω is a finite subset of Γ, we define $\partial\Omega$ to be the set $\{g \in \Omega \mid \exists g' \in \Gamma - \Omega,\ d_S(g, g') = 1\}$ and we let $|\Omega|$ denote the cardinal of Ω.

Proposition 7.2. *If there is a constant* $C > 0$ *such that* $V_\Gamma(n) \geq Cn^3$ *, then there is a constant* $C' > 0$ *such that the following isoperimetric inequality holds for any bounded domain* $\Omega \subset \mathrm{R}$ *:*

$$|\Omega|^{2/3} \leq C'|\partial\Omega|.$$

Proof. By theorem 1 of [7], there is a constant $C'' > 0$ such that the discrete isoperimetric inequality

$$|\Omega|^{2/3} \leq C''|\partial\Omega|$$

holds for any finite subset $\Omega \subset \Gamma$. By lemma 4.2 of [20], the same discrete inequality (up to a multiplicative constant) holds in some net $N \subset \mathrm{R}$, so by lemma 4.5 of [20] the required inequality holds in R.

By the Ahlfors lemma ([19], 6.9), the fact that R is conformally equivalent to $\mathbf{E}^2$ implies that there is a sequence of domains $\mathrm{D}_i \subset \mathrm{R}$ such that

$$\lim_{i \to +\infty} \frac{|\partial\mathrm{D}_i|^{3/2}}{|\mathrm{D}_i|} = 0.$$

Together with proposition 7.2, this implies that there is no constant C such that $V_\Gamma(n) \geq Cn^3$ for all n. Thus by Gromov's theorem [16], Γ has at most quadratic growth and is virtually nilpotent. Applying the result of [2], we deduce that Γ has a subgroup of finite index Γ' such that Γ' is nilpotent, there is at most one nontrivial quotient in the lower central series of Γ', and this quotient has rank at most 2. This means that Γ' has a subgroup of finite index isomorphic to $\mathbf{Z}$, $\mathbf{Z}^2$ or the trivial group. But we know that Γ is infinite, and Γ cannot be virtually $\mathbf{Z}$ for the following reason: in that case, the Cayley graph of Γ would have two ends, which contradicts the fact that the number of ends is a quasi-isometry invariant for proper geodesic metric spaces. This completes the proof of proposition 7.1 in the Euclidean case.

7.2. The hyperbolic case

From now on we assume that R is conformally equivalent to $\mathbf{H}^2$. We will make the further hypothesis that Γ is not virtually $\mathbf{Z}^2$ and prove that R is quasi-isometric to $\mathbf{H}^2$. In fact, Mess shows that no complete Riemannian plane conformally equivalent to $\mathbf{H}^2$ can be quasi-isometric to $\mathbf{Z}^2$, but there does not seem to be a short, direct proof of this fact.

Our proof of the hyperbolic case relies on the following lemma.

Lemma 7.3. [Half-minimum lemma, cf. [17], p. 256]
Let (X, d) *be a complete metric space and* $h : X \to \mathbf{R}_+^*$ *a function which is locally bounded away from zero. Let* x *be a point of* X *such that* $0 < h(x) < \frac{1}{2}$ *. Then there exists a point* $x' \in X$ *such that* $d(x, x') \leq 2$ *,* $h(x') \leq h(x)$ *and*

x' is a so-called half-minimum, i.e. for all $y \in X$, if $d(y, x') \leq \frac{1}{2}\sqrt{h(x')}$ then $h(y) \geq \frac{1}{2}h(x')$.

Proof. Seeking a contradiction, assume there is no such x'. In particular x itself is not a half-minimum. Then by definition, there is a $x_1 \in X$ such that $d(x_1, x) \leq \frac{1}{2}\sqrt{h(x)}$ and yet $h(x_1) < \frac{1}{2}h(x)$. But x_1 cannot be a half-minimum either, so there is a x_2 such that $d(x_2, x_1) \leq \frac{1}{2}\sqrt{h(x_1)}$ and $h(x_2) < \frac{1}{2}(x_1)$. This implies that $d(x, x_2) \leq \frac{1}{2\sqrt{2}} + \frac{1}{4} \leq 2$ and $h(x_2) \leq \frac{1}{4}h(x)$. This process never stops and yields a Cauchy sequence (x_n) satisfying $d(x, x_n) \leq 2$ and $h(x_n) \leq \frac{1}{2^n}h(x)$. Set $y = \lim x_n$. Then h cannot be locally bounded away from zero at y, a contradiction.

Using the half-minimum lemma, we deduce:

Lemma 7.4. [cf. [17], p. 255] *Let* (X, g_X) *and* (V, g_V) *be complete Riemannian surfaces. Suppose that* X *is quasi-homogeneous and* $\mathrm{Isom}(V)$ *is cocompact. Let* $c : X \to V$ *be a conformal map. Define a function* $\mu : X \to \mathbf{R}^*_+$ *by* $g_X = \mu^2 c^* g_V$. *If there is a sequence* $x_n \in X$ *such that* $\mu(x_n) \to 0$, *then there exists a conformal map* $f : \mathbf{E}^2 \to V$.

Proof. Applying the half-minimum lemma to the function μ at each x_n, we get a sequence y_n satisfying

$$\lim \mu(y_n) = 0$$

$$\forall y \in X \qquad d_X(y, y_n) \leq \frac{1}{2}\sqrt{\mu(y_n)} \implies \mu(y) \geq \frac{1}{2}\mu(y_n).$$

Since X is quasi-homogeneous, there exist constants r, λ and for each n a conformal chart $\phi_n : \mathrm{B}(0, r) \to X$ satisfying $\phi_n(0) = y_n$, $\|\mathrm{D}\phi_n(0)\| = 1$ and $\sup_{a \in \mathrm{B}(0,r)} \|\mathrm{D}\phi_n(a)\| = 1/2\lambda$.

Set $\mathrm{B}_n = \mathrm{B}(0, \lambda/\sqrt{\mu(y_n)}) \subset \mathbf{E}^2$. Then for n sufficiently large so that $2\sqrt{\mu(y_n)} \leq r$, we can define a map $z_n : \mathrm{B}_n \to \mathrm{B}(0, r)$ by $z_n(a) = \mu(y_n)a$. Next we use the cocompacity of $\mathrm{Isom}(V)$: by postcomposing c with suitable isometries of V, we get conformal maps $c_n : X \to V$ such that for each n, $c_n(y_n)$ belongs to some fixed compact $\mathrm{K} \subset V$. Finally define $f_n = c_n \circ \phi_n \circ z_n$.

We need to estimate $\sup \|\mathrm{D}f_n\|$, which we do by giving upper bounds for $\|\mathrm{D}c_n\|$, $\|\mathrm{D}\phi_n\|$ and $\|\mathrm{D}z_n\|$.

First we see that $\|\mathrm{D}z_n(a)\| = \mu(y_n)$ and $\|\mathrm{D}\phi_n(z_n(a))\| \leq 1/2\lambda$ for all $a \in \mathrm{B}_n$. Thus if $a \in \mathrm{B}_n$, then $\phi_n(z_n(a)) \in \mathrm{B}(y_n, \frac{1}{2}\sqrt{\mu(y_n)})$ and the half-minimum property asserts that $\mu(\phi_n(z_n(a))) \geq \frac{1}{2}\mu(y_n)$. Now recall that c_n is the composition of an isometry and of c, which is conformal of dilatation $1/\mu$. Hence $\|\mathrm{D}c_n(\phi_n(z_n(a)))\| \leq 2/\mu(y_n)$. We deduce:

$$\|\mathrm{D}f_n(a)\| \le \|\mathrm{D}c_n(\phi_n(z_n(a)))\| \cdot \|\mathrm{D}\phi_n(z_n(a))\| \cdot \|\mathrm{D}z_n(a)\|$$

$$\le \frac{2}{\mu(y_n)} \cdot \frac{1}{2\lambda} \cdot \mu(y_n)$$

$$\le \frac{1}{\lambda}.$$

Thus for every n, the sequence $\{\tilde{f}_p\}_{p \ge n}$ obtained by restricting each f_p to B_n is equicontinuous on B_n and for all p we know that $\tilde{f}_p(0) \in \mathrm{K}$ and $\|\mathrm{D}\tilde{f}_n(0)\| = 1$. By Ascoli's theorem, $\{\tilde{f}_p\}_{p \ge n}$ admits a convergent subsequence. By diagonal extraction, we obtain a sequence of conformal maps $g_n : \mathrm{B}_n \to \mathrm{V}$ which converges uniformly on compacts to some map $g : \mathbf{E}^2 \to \mathrm{V}$. By standard properties of conformal mappings in dimension 2, the map g is conformal or constant, but the latter possibility is ruled out by the fact that $\|\mathrm{D}g(0)\| = \lim \|\mathrm{D}g_n(0)\| = 1$.

Next we turn to the proof of proposition 7.1 in the hyperbolic case. To simplify notations, let ds^2 denote the hyperbolic metric on $\mathbf{H}^2$ and consider a quasi-homogeneous conformal metric $g = \mu^2 ds^2$ such that $(\mathbf{H}^2, g)$ is quasi-isometric to a finitely generated group Γ. Let d denote the hyperbolic distance and d_g the distance with respect to g. Our goal is to prove that $(\mathbf{H}^2, d)$ and $(\mathbf{H}^2, d_g)$ are quasi-isometric.

Taking $(X, g_X) = (\mathbf{H}^2, g)$, $(\mathrm{V}, g_\mathrm{V}) = (\mathbf{H}^2, ds^2)$ and c the identity map, lemma 7.4 implies that the function μ is bounded away from zero, since Liouville's theorem tells us that there is no conformal map $f : \mathbf{E}^2 \to \mathbf{H}^2$. Thus we have a linear upper bound for d in terms of d_g. The remaining task is to obtain an inequality in the reverse direction. For this, we will again use the half-minimum lemma.

First we need some notation. For given $x \in \mathbf{H}^2$, $\rho \ge 0$, let $\mathrm{D}(x, \rho)$ denote the hyperbolic disc around x of radius ρ, $\mathrm{A}(x, \rho)$ the area of this disc with respect to g, $\mathrm{C}(x, \rho)$ its boundary circle, $l(\rho)$ (resp. $\mathrm{L}(x, \rho)$) the length of $\mathrm{C}(x, \rho)$ with respect to ds^2 (resp. to g). (Note that by homogeneity, $l(\rho)$ is independent of x, while $\mathrm{L}(x, \rho)$ may not be.)

Besides homogeneity, the only property of the hyperbolic metric we will need is the following fact.

Fact. There is a small positive constant R_0 such that $\rho \le \mathrm{R}_0$ implies $l(\rho) \le 20\rho$.

Given $x \in \mathbf{H}^2$ and $\mathrm{A} \ge 0$, let $r(x, \mathrm{A})$ be the infimum of the numbers $\rho > 0$ such that $\mathrm{A}(x, \rho) \ge \mathrm{A}$. For fixed $\mathrm{A} > 0$, the function $x \mapsto r(x, \mathrm{A})$ takes positive values and is locally bounded away from zero.

Lemma 7.5. *Fix* $\mathrm{A} \ge 0$. *If* $0 < r \le \mathrm{R}_0$ *and* $\mathrm{A}(x, r) \le \mathrm{A}$ *then there exists*

$\rho \in [r/2, r]$ *such that* $\mathrm{L}(x, \rho) \leq 10\sqrt{\mathrm{A}}$.

Proof. Assume the conclusion does not hold. Then by the Cauchy-Schwarz inequality and the above fact, we have for all $\rho \in [r/2, r]$:

$$100\mathrm{A} \leq \left(\int_{\mathrm{C}(x,\rho)} \mu \, ds \right)^2 \leq l(\rho) \cdot \int_{\mathrm{C}(x,\rho)} \mu^2 \, ds$$

$$\leq 20\rho \cdot \int_{\mathrm{C}(x,\rho)} \mu^2 \, ds.$$

Dividing by ρ and integrating between $r/2$ and r , we get:

$$100\mathrm{A} \int_{r/2}^{r} \frac{d\rho}{\rho} \leq 20\mathrm{A}(x, r) \leq 20\mathrm{A}$$

$$5\ln 2 \leq 1.$$

Lemma 7.6. *If* $\inf\{r(x, \mathrm{A}) \mid x \in \mathbf{H}^2\} = 0$ *for some* $\mathrm{A} > 0$, *then there is a closed curve of length* $\leq 800\sqrt{\mathrm{A}}$ *enclosing a domain of area* $\geq \mathrm{A}$.

Proof. We apply the half-minimum lemma to the function $x \mapsto r(x, \mathrm{A})$. We can find x_A such that $r_\mathrm{A} = r(x_\mathrm{A}, \mathrm{A}) < \mathrm{R}_0$ and for all x , $d(x, x_\mathrm{A}) \leq \frac{1}{2}\sqrt{r_\mathrm{A}}$ implies $r(x, \mathrm{A}) \geq \frac{1}{2}r_\mathrm{A}$. Let x_A be chosen so that $r_\mathrm{A} < \frac{1}{4}\sqrt{r_\mathrm{A}}$. Then for all $x \in \mathrm{D}(x_\mathrm{A}, 2r_\mathrm{A})$ and all $r \leq r_\mathrm{A}$, we have $\mathrm{A}(x, r) \leq \mathrm{A}$.

By lemma 7.5, we can associate to any point x such that $d(x, x_\mathrm{A}) = r_\mathrm{A}$ a $\rho(x) \in [r_\mathrm{A}/4, r_\mathrm{A}/2]$ such that $\mathrm{L}(x, \rho(x)) \leq 10\sqrt{\mathrm{A}}$. Let x_1 be a point on $\mathrm{C}(x_\mathrm{A}, r_\mathrm{A})$. Then $\mathrm{C}(x_1, \rho(x_1))$ and $\mathrm{C}(x_\mathrm{A}, r_\mathrm{A})$ intersect in two points (see figure 2). Let x_2 be one of them.

Then $\mathrm{C}(x_1, \rho(x_1))$ cannot lie in the exterior of $\mathrm{C}(x_2, \rho(x_2))$, and if it was contained in its interior, then we would get $\rho(x_1) < \rho(x_2)/2$, which cannot happen. So these two circles intersect. We can further construct x_3 , x_4 and so on, stopping when $\mathrm{C}(x_n, \rho(x_n))$ intersects $\mathrm{C}(x_1, \rho(x_1))$. We obtain

$$\frac{nr_\mathrm{A}}{4} \leq \sum_{i=1}^{n} \rho(x_i) \leq l(r_\mathrm{A}) \leq 20r_\mathrm{A},$$

from which we deduce $n < 80$. Piecing together the outward arcs of the $\mathrm{C}(x_i, \rho(x_i))$'s, we get a curve of length less than $800\sqrt{\mathrm{A}}$ bounding a domain of area greater than A .

Lemma 7.7. *There is a constant* $\mathrm{A}_0 > 0$ *such that* $\inf\{r(x, \mathrm{A}_0) \mid x \in \mathbf{H}^2\} > 0$.

Proof. In order to obtain a contradiction, take a sequence $\mathrm{A}_i \to +\infty$ such that $\inf\{r(x, \mathrm{A}_i) \mid x \in \mathbf{H}^2\} = 0$. Lemma 7.6 then gives a sequence of domains Ω_i with

$\mathrm{C}(x_1, \rho(x_1))$

$\mathrm{C}(x_2, \rho(x_2))$

x_1 x_2

$\mathrm{C}(x_3, \rho(x_3))$

x_3

x_A

$\mathrm{C}(x_\mathrm{A}, r_\mathrm{A})$
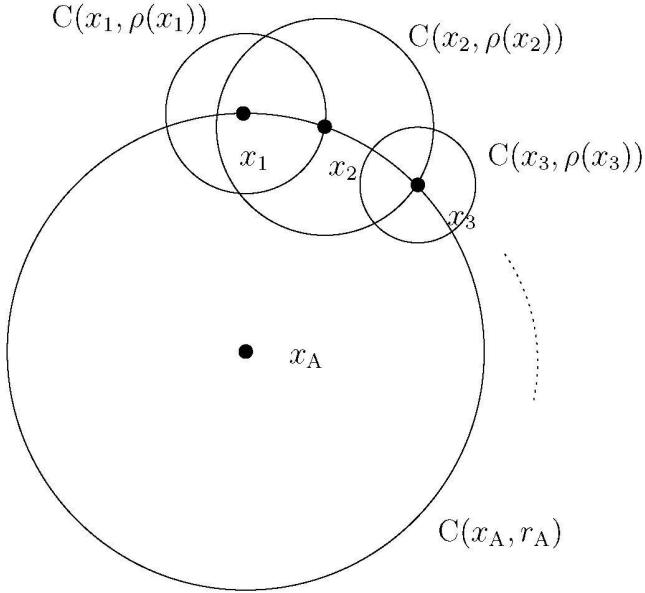
Figure 2.

$\frac{|\partial\Omega_i|^2}{|\Omega_i|}$ bounded and $|\Omega_i| \to +\infty$. (All lengths and areas are taken with respect to the metric g).

Now assume (again by contradiction) that for some constant $\mathrm{C} > 0$, the inequality

$$|\Omega|^{2/3} \leq \mathrm{C}|\partial\Omega|$$

holds for all bounded domains. Then there is a constant $\mathrm{C}' > 0$ such that

$$|\Omega_i|^{2/3} \leq \mathrm{C}|\partial\Omega_i| \leq \mathrm{C}'|\Omega_i|^{1/2},$$

which is impossible since $|\Omega_i| \to +\infty$.

So the above isoperimetric inequality does not hold. By proposition 7.2, the growth function of Γ cannot be at least cubic. As in the Euclidean case, it follows that Γ is virtually $\mathbf{Z}^2$, which contradicts the hypothesis made at the beginning of this subsection.

Lemma 7.8. *Let* V *be a complete, quasi-homogeneous Riemannian manifold. Then for any* $\mathrm{A} > 0$ *there is a* $\tau(\mathrm{A}) > 0$ *such that for all* $x \in \mathrm{V}$, *the volume of* $\mathrm{B}(x, \tau(\mathrm{A}))$ *is* $> \mathrm{A}$.

Proof. Let $a > 0$ be a lower bound for the volumes of balls of radius 1 and $n(\mathrm{A}) \in \mathbf{N}$ such that $\mathrm{A} \leq n(\mathrm{A})a$. Take $\tau(\mathrm{A})$ large enough so that any ball of radius $\tau(\mathrm{A})$ contains at least $n(\mathrm{A}) + 1$ disjoint balls of radius 1. (e. g., $\tau(\mathrm{A}) = 3n(\mathrm{A}) + 1$.)

We are now in position to complete the proof of proposition 7.1. Define $r_0 = \min(\mathrm{R}_0, \inf\{r(x, \mathrm{A}_0) \mid x \in \mathbf{H}^2\})$. Using lemma 7.5, we define a function $\rho : \mathrm{R} \to [r_0/2, r_0]$ such that the inequalities $\mathrm{A}(x, \rho(x)) \leq \mathrm{A}_0$ and $\mathrm{L}(x, \rho(x)) \leq 10\sqrt{\mathrm{A}_0}$ hold for all $x \in \mathbf{H}^2$. From now on, we write $\mathrm{C}(x) = \mathrm{C}(x, \rho(x))$.
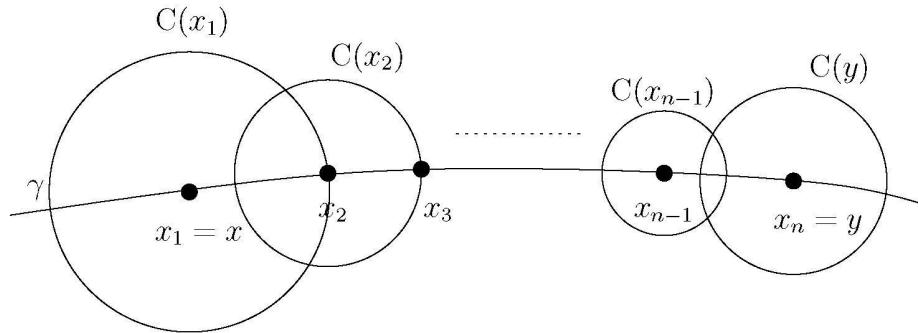

Figure 3.

Let x, y be points of $\mathbf{H}^2$. Let γ be a hyperbolic geodesic passing through x and y. We want to construct a finite sequence $(x_1, \ldots, x_n)$ of points of γ such that $x_1 = x$, $x_n = y$ and consecutive $\mathrm{C}(x_i)$'s intersect. We will use the same technique as in the proof of lemma 7.6: we take $x_1 = x$ and for $i > 1$, x_i will be one of the two intersection points between γ and $\mathrm{C}(x_{i-1})$ (see fig. 3). If we look at γ in such a way that x lies on the left of y, then at each step we can describe the choice made by saying that we pick either the "leftmost" or "rightmost" intersection point.

If $\mathrm{C}(x)$ and $\mathrm{C}(y)$ intersect, we are done. If they are exterior to each other, move toward the right. For some m, $\mathrm{C}(x_{m-1})$ lies in the exterior of $\mathrm{C}(y)$ and $\mathrm{C}(x_m)$ does not. If $\mathrm{C}(y)$ lies in the interior of $\mathrm{C}(x_m)$, then we have $2\rho(y) < \rho(x_m)$, so $\mathrm{C}(y)$ and $\mathrm{C}(x_m)$ intersect. If $\mathrm{C}(x)$ and $\mathrm{C}(y)$ are disjoint, but not exterior to each other, move first to the left until $\mathrm{C}(x_m)$ is exterior to $\mathrm{C}(y)$, then to the right as in the previous case. In each case, we find a $\mathrm{C}(x_n)$ intersecting $\mathrm{C}(y)$ with n bounded by an affine function of $d(x, y)$.

We can now connect x and y by a path $\alpha \cup \beta \cup \delta$ where α is the shortest g-geodesic connecting x to a point of $\mathrm{C}(x)$, δ the shortest g-geodesic connecting y to a point of $\mathrm{C}(y)$ and $\beta \subset \bigcup \mathrm{C}(x_i)$. Each of the $\mathrm{C}(x_i)$'s has length less than $10\sqrt{\mathrm{A}_0}$, so the length of β with respect to g is bounded by an affine function of $d(x, y)$.

We are left with the following problem: given a point $x \in \mathbf{H}^2$, find an upper bound for the g-distance between x and $\mathrm{C}(x, \rho(x))$ which does not depend on x. Lemma 7.8 provides a solution. Indeed, let $\tau_0 = \tau(\mathrm{A}_0)$ be the constant given by that lemma. Assume all points of $\mathrm{C}(x, \rho(x))$ are distant of x by more than τ_0.

Let z be a point of the g-ball centered at x of radius τ_0. If $z \notin D(x, \rho(x))$, a connectivity argument shows that some minimizing g-geodesic segment connecting x and z must cut $C(x, \rho(x))$. If u is an intersection point, then it satisfies $d_g(x, u) \leq \tau_0$, a contradiction.

We have just proved that the g-ball centered at x of radius τ_0 is contained in $D(x, \rho(x))$, whose area is less than A_0. This contradiction completes the proof of theorem 7.1.

8. Surfaces with nonempty boundary

The section is devoted to the proof of proposition 1.4, which says that a finitely generated group Γ is virtually free iff it is quasi-isometric to some complete simply-connected surface R with nonempty geodesic boundary. The "only if" part is easy (just take the universal covering of a compact surface with geodesic boundary of the adequate genus), so we concentrate on the "if" part. It is a consequence of the following technical result:

Theorem 8.1. *Let R be a simply-connected, noncompact surface, with nonempty boundary, $\mathcal{D}$ a nondegenerate pseudo-triangulation of R and d the regular piecewise Euclidean metric induced by $\mathcal{D}$. Let Γ be a finitely generated group quasi-isometric to (R, d). Then Γ is virtually free.*

Proof of proposition 1.4 assuming theorem 8.1. Let (R, h) be a complete simply-connected Riemannian surface with nonempty geodesic boundary. The proof of the pseudo-triangulation theorem works in this case as well as in the boundary-free case, starting with a net N that has a subnet N' that is contained in ∂R and is quasidense in it. So the proof that (B) + (P) implies (E) goes through as in the boundary-free case and shows that (R, h) is quasi-isometric to some regular piecewise Euclidean simply-connected surface. Then theorem 8.1 implies that Γ is virtually free.

The proof of theorem 8.1 is based on the following lemma.

Lemma 8.2. ∂R *is quasidense in* R.

The proof of lemma 8.2 is rather technical, so we first show that it implies theorem 8.1.

Proof of theorem 8.1 assuming lemma 8.2. Let N be the set of vertices of the pseudo-triangulation that lie in ∂R. Lemma 8.2 implies that N is quasidense. Perturb the metric so that it is smooth in the interior of R and each point $x \in N$ is replaced by a convex boundary arc. Now we are in the setting of section 6. Using the same terminology, R consists of just one thin, simply-connected piece. Hence

we may apply proposition 6.5. This yields an embedded geodesic divide which splits R into pieces of uniform size. Let T be the dual graph of this splitting. Then T is a tree and is quasi-isometric to R. It follows that Γ is quasi-isometric to a tree, hence a word-hyperbolic group whose boundary has dimension 0, so it is virtually free [15].

The end of the paper is devoted to the proof of lemma 8.2. Let (R, d) and Γ satisfy the hypotheses of theorem 8.1. The distance function on G will also be denoted by d. There are constants $\lambda', C' \geq 0$ and (λ', C')-quasi-isometries $f : \Gamma \to R$, $\bar{f} : R \to \Gamma$ such that $\bar{f}$ is a coarse inverse of f, i.e. for all $g \in \Gamma$ we have $d(f(\bar{f}(g)), g) \leq C'$ and for all $x \in R$ we have $d(\bar{f}(f(x)), x) \leq C'$.

A (bi-infinite) (λ, C)-*quasi-geodesic* in R is a map $\alpha : \mathbf{R} \to R$ that is a (λ, C)-quasi-isometry.

Lemma 8.3. *There are constants* $\lambda, C \geq 0$ *such that*

i *For every* $x, y \in R$ *there is a* (λ, C)-*quasi-isometry* $f_{x,y} : R \to R$ *such that* $f_{x,y}(x)$ *lies in the* C-*neighborhood of* y.

ii *For every* $x \in R$ *there is a* (λ, C)-*quasi-geodesic* α_x *such that* x *lies in the* C-*neighborhood of* α_x.

Proof. The key is that (i) and (ii) are coarse versions of properties that Γ satisfies because it is a group, namely: Γ acts transitively on itself by isometries and for all $g \in \Gamma$ there is a geodesic in the Cayley graph of Γ that passes through g. Thus the corresponding coarse properties hold in any metric space quasi-isometric to Γ. For instance, let x, y be points R and $h : \Gamma \to \Gamma$ be an isometry which sends $\bar{f}(x)$ to $\bar{f}(y)$. Then $f \circ h \circ \bar{f}$ is a quasi-isometry and sends x close to y. A straightforward computation shows that the various constants depend only on λ' and C'. The proof of (ii) is similar.

The idea of the next lemma is due to B. Bowditch. (cf. [4]).

Lemma 8.4. (R, d) *is* ***uniformly simply-connected,*** *which means that there is a function* $\rho : \mathbf{R}_+ \to \mathbf{R}_+$ *such that every loop* $\gamma \subset R$ *bounds a singular disc of diameter less than* $\rho(\operatorname{diam} \gamma)$.

Proof. Define
$$\rho(r) = \lambda^2(r+1) + C(\lambda + 1).$$

Assume this function does not work. Then by planar topology it does not even work for embedded loops, which means that there is a $r \geq 0$ and an embedded disc Ω such that $\operatorname{diam} \partial \Omega = r$ and $\operatorname{diam} \Omega > \rho(r)$.

Let x be a point of Ω such that $d(x, \partial \Omega) > \rho(r)/2$. By lemma 8.3(ii), there is a (λ, C)-quasi-geodesic α_x which meets $B(x, C)$. We may assume that α_x is continuous. Thus it must meet $\partial \Omega$ in at least two points $x_1 = \alpha_x(t_1)$ and

$x_2 = \alpha_x(t_2)$, where t_1 and t_2 have the property that there is a $t_0 \in [t_1, t_2]$ such that $d(\alpha(t_0), x) \leq C$. But then the quasi-geodesicity of α implies

$$t_2 - t_1 = (t_2 - t_0) + (t_0 - t_1) \geq \frac{2(\rho(r) - C)}{\lambda}.$$

It follows that

$$d(x_2, x_1) \geq \frac{2(\rho(r) - C) - C\lambda}{\lambda^2} > r,$$

giving a contradiction.

In the next two lemmas, we exploit the uniform simple connectivity of R to obtain "connect-the-dots" type results (cf. [12]).

Lemma 8.5. *For all* $\mu, D \geq 0$ *there exists* $D' \geq 0$ *such that for every* (μ, D) - *quasi-isometry* $h : R \to R$, *there is a continuous* (μ, D')- *quasi-isometry* $h' :$ $R \to R$ *such that* $d(h(x), h'(x)) \leq D'$ *for all* $x \in R$. *The constant* D' *depends only on* μ, D *and the function* ρ.

Proof. The continuous quasi-isometry h' is constructed by induction over the skeleta of the pseudo-triangulation. First set $h' = h$ on the 0-skeleton. Then map each edge to a geodesic. Finally use lemma 8.4 to define h' on the 2-skeleton. A straightforward computation shows that h' satisfies the above properties with $D' = 2(\mu + 1)\rho(2\mu + 2D) + 2\mu + D + 2$.

Lemma 8.6. *For all* $D \geq 0$ *there exists* $\epsilon(D) \geq 0$ *such that any continuous map* $h : S^1 \times \{0, 1\} \to R$ *satisfying* $d(h(t, 0), h(t, 1)) \leq D$ *for all* $t \in S^1$ *can be extended to a map* $h : S^1 \times I \to R$ *satisfying* $\operatorname{diam} h(t \times I) \leq \epsilon(D)$.

Proof. Similar to that of lemma 8.5, using a triangulation of $S^1 \times I$.

Proof of lemma 8.2. Seeking a contradiction, assume there is a point $x \in R$ such that $d(x, \partial R) > \delta$, where δ is to be determined. Let y be any point of ∂R. By lemma 8.3(i), there is a (λ, C)-quasi-isometry $f_{x,y}$ which sends x into $B(y, C)$. Let $\bar{f}_{x,y}$ be a coarse inverse for $f_{x,y}$. By increasing the constants, we may suppose that $f_{x,y}$ and $\bar{f}_{x,y}$ are continuous (lemma 8.5).

Let U be a topological disc containing $B(f_{x,y}(x), \lambda(C + 1) + C)$ and such that $\operatorname{diam} U \leq \rho(2\lambda(C + 1) + 2C)$. Since $d(y, f_{x,y}(x)) \leq C$, the disc U meets ∂R, which implies that every component of $R - U$ is simply-connected.

Let Ω be a large disc such that $x \in \Omega$ and $d(x, \partial\Omega) > \delta$. If δ is large enough so that $(\delta - C)/\lambda > \operatorname{diam} U$, then the loop $f_{x,y}(\partial\Omega)$ will avoid U. Since every component of $R - U$ is simply-connected, $f_{x,y}(\partial\Omega)$ bounds a singular disc V such that $U \cap V = \emptyset$. For all $z \in V$, we have $d(f_{x,y}(x), z) \geq \lambda(C + 1) + C$, whence $d(\bar{f}_{x,y}(z), x) \geq 1$. This proves that $\bar{f}_{x,y}(f_{x,y}(\partial\Omega))$ is null-homotopic in $R - \{x\}$.

Let $h : S^1 \times \{0, 1\} \to R$ be defined by $h_0 = \partial\Omega$ and $h_1 = \bar{f}_{x,y}(f_{x,y}(\partial\Omega))$. Lemma 8.6 provides a homotopy between $\partial\Omega$ and $\bar{f}_{x,y}(f_{x,y}(\partial\Omega))$ such that $\operatorname{diam} h(t \times I) \leq \epsilon(C)$ for all $t \in S^1$. If δ is chosen so that $\epsilon(C) < \delta$ then this homotopy avoids x, which gives the expected contradiction.

References

[1] O. Attie, Quasi-isometry classification of some manifolds of bounded geometry. *Math. Z.* **216**(4) (1994), 501–527.

[2] H. Bass, The degree of polynomial growth of finitely generated nilpotent groups. *Proc. London Math. Soc.* **25** (1982), 603–614.

[3] J. Block, S. Weinberger, Aperiodic tilings, positive scalar curvature and amenability of spaces. *J. Amer. Math. Soc.* **5**(4) (1992), 907–918.

[4] B. H. Bowditch, Planar groups and the Seifert conjecture. Preprint 1999.

[5] A. Casson and D. Jungreis, Convergence groups and Seifert fibered 3-manifolds. *Invent. Math.* **118** (1994), 441–456.

[6] M. Coornaert, T. Delzant, A. Papadopoulos, *Géométrie et théorie des groupes*. Les groupes hyperboliques de Gromov. Lecture Notes in Mathematics, Vol. 1441. Springer-Verlag, Berlin, 1990.

[7] T. Coulhon, L. Saloff-Coste, Isopérimétrie pour les groupes et les variétés. *Revista Math. Iberoamericana* **9** (1993), 293–314.

[8] B. Farb, The quasi-isometry classification of lattices in semisimple Lie groups. *Math. Res. Lett.* **4**(5) (1997), 705–717.

[9] B. Farb, L. Mosher, Quasi-isometric rigidity for the solvable Baumslag-Solitar groups. II. *Invent. Math.* **137**(3) (1999), 613–649.

[10] B. Farb, L. Mosher, On the asymptotic geometry of abelian-by-cyclic groups. *Acta Math.* **184**(2) (2000), 145–202.

[11] B. Farb, L. Mosher, A rigidity theorem for the solvable Baumslag-Solitar groups. *Invent. Math.* **131**(2) (1998), 419–451. With an appendix by Daryl Cooper.

[12] B. Farb, R. Schwartz, The large-scale geometry of Hilbert modular groups. *J. Differential Geom.* **44**(3) (1996), 435–478.

[13] M. Freedman, J. Hass, P. Scott, Closed geodesics on surfaces. *Bull. London Math. Soc.*, **14** (1982), 385–391.

[14] D. Gabai, Convergence groups are Fuchsian groups. *Annals of Math.*, **136** (1992), 447–510.

[15] É. Ghys, P. de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov*. Birkhäuser Boston Inc., Boston, MA, 1990. Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988.

[16] M. Gromov, Groups of polynomial growth and expanding maps. *Publ. Math.* IHES **53** (1981), 53–73.

[17] M. Gromov, Foliated Plateau problem, part II: harmonic maps of foliations. *GAFA* **1** (1991), 253–320.

[18] M. Gromov, Asymptotic invariants of infinite groups. In *Geometric group theory, 2 (Sussex, 1991)*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[19] M. Gromov, J. Lafontaine, P. Pansu, *Structures métriques pour les variétés riemanniennes*. CEDIC, Paris, 1981.

[20] M. Kanai, Rough isometries and combinatorial approximations of geometries of non-compact riemannian manifolds. *J. Math. Soc. Japan* **37**(3) (1985), 391–413.

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

[22] M. Kapovich, B. Leeb, Quasi-isometries preserve the geometric decomposition of Haken man-ifolds. *Invent. Math.* **128**(2) (1997), 393–416.

[23] G. Mess, The Seifert conjecture and groups which are coarse quasiisometric to planes. Preprint.

[24] E. G. Rieffel, *Groups coarse quasi-isometric to* $\mathbf{H}^2 \times \mathbf{R}$. PhD thesis, UCLA, 1991.

[25] M. Troyanov, Les surfaces euclidiennes à singularités coniques. *Enseign. Math. (2)* **32**(1-2) (1986), 79–94.

[26] P. Tukia, Homeomorphic conjugates of Fuchsian groups. *J. Reine Angew. Math.* **391** (1988), 1–54.

[27] K. Whyte, The quasi-isometry types of the higher Baumslag-Solitar groups. To appear in *GAFA*.

Sylvain Maillot
Laboratoire Mathématiques Émile Picard
Université Paul Sabatier
118, route de Narbonne
F-31062 Toulouse Cédex 4
France
e-mail: maillot@borel.ups-tlse.fr