TRIANGULATING SURFACES QUASI-ISOMETRICALLY

AGELOS GEORGAKOPOULOS AND FEDERICO VIGOLO

Abstract. We prove that every complete Riemannian surface (Σ, d_Σ) admits a triangulation D whose 1-skeleton, when endowed with the inherited length metric, is quasi-isometric to (Σ, d_Σ). Moreover, the faces of D have intrinsic diameters uniformly bounded by an arbitrarily small constant.

Keywords: complete Riemannian surface, triangulation, quasi-isometry, planar graph.

MSC 2020 Classification: 51F30, 57M15, 57Q15, 57R05, 05C10

1. Introduction

A classical result of Rado states that every closed surface admits a triangulation, and it is well-known that this extends to the non-compact case [1]. It is also known that every complete Riemannian surface (Σ, d_Σ) contains a subspace G homeomorphic to a graph, such that the metric that G inherits from Σ via the lengths of its edges is quasi-isometric to (Σ, d_Σ) [2, Lemma 8.2]. Our main result combines the two aforementioned statements: we prove that G can be chosen so as to triangulate Σ. Moreover, we can ensure that the faces of the triangulation have arbitrarily small diameters. To state this, let us first recall that any (arc-connected) subspace G of a Riemannian manifold (Σ, d_Σ) inherits a *length metric* d_ℓ via

$$d_\ell(x, y) := \inf\{\ell(\gamma) \mid \gamma \text{ is a } x\text{–}y \text{ arc in } G\},$$

whereby $\ell(\gamma)$ denotes the length of γ.

Theorem 1.1. *Let (Σ, d_Σ) be a complete Riemannian surface, and $\Theta \in \mathbb{R}_{\geq 0}$. Then there is a triangulation D of (Σ, d_Σ) each 2-cell of which has diameter at most Θ (with respect to the inherited length metric). Moreover, the identity map from the 1-skeleton G of D, endowed with the length metric d_ℓ, to (Σ, d_Σ) is (1-Lipschitz and) a quasi-isometry.*

The lengths of the edges of this G obey a uniform upper-bound $C\Theta$, but no uniform lower bound. This is unavoidable: James Davies (private communication) constructed complete Riemannian surfaces which are not quasi-isometric to any *simplicial graph*, i.e. a graph with all edges having length 1, embeddable into Σ. However, if (Σ, d_Σ) is quasi-isometric to some graph of bounded degree, equivalently, if (Σ, d_Σ) has a uniform net [?], then we can strengthen Theorem 1.1 to ensure that G has its edge-lengths uniformly bounded above and below, and therefore its simplicial graph metric is quasi-isometric to (Σ, d_Σ):

Theorem 1.2. *If a complete Riemannian surface (Σ, d_Σ) has a uniform net, then Σ admits a triangulation whose (simplicial) 1-skeleton is quasi-isometric to Σ.*

AG is supported by EPSRC grant EP/V009044/1.

By allowing degenerate triangles in such a 'triangulation' we can ensure that the 0-skeleton (i.e. the vertices) of the triangulation coincides with any uniform net of Σ. Moreover, the 1-skeleton of our triangulation will have bounded vertex degrees up to ignoring parallel edges. Similar statements had been proved by Maillot [10], and we obtain simpler proofs for some such results (see Theorem 4.1).

The first author & Papasoglu [6] asked whether for every graph G embedded into a Riemannian plane (Σ, d_Σ), the metric (G, d_ℓ) is quasi-isometric to that of a simplicial planar graph. An equivalent formulation is whether every planar metric graph is quasi-isometric to a simplicial planar graph. This question was answered in the affirmative by Davies [4]. Questions of this type currently attract interest in the context of coarse graph theory. Nguyen, Scott & Seymour [11] observe that there are planar (simplicial) graphs G, H such that there is a (M, A)-quasi-isometry from G to H for some (multiplicative, resp. additive) constants M, A but no $(1, A')$-quasi-isometry. Similarly, one can ask whether the additive constant A can be set to 0 by modifying the quasi-isometry between fixed planar graphs G, H, i.e. whether G, H are bi-Lipschitz equivalent. In Section 5 we combine our results with a theorem of Burago & Kleiner [3] to provide a strong negative answer:

Corollary 1.3. *There are plane graphs G, H (with bounded degrees and face-boundary sizes) which are quasi-isometric but not bi-Lipschitz equivalent.*

An interesting open question is whether every (non-planar) graph G which is quasi-isometric to a planar graph H must be $(1, A)$-quasi-isometric to some planar graph H'; see [11] for more.

2. Notation and preliminaries

2.1. **Quasi-isometries.** Let (X, d_X) and (Y, d_Y) be metric spaces. For $M \in \mathbb{R}_{\geq 1}$ and $A \in \mathbb{R}_{\geq 0}$, an (M, A)-*quasi-isometry* from (X, d_X) to (Y, d_Y) is a map $\varphi \colon X \to Y$ such that

(1) for every $x, y \in X$ we have
$$\frac{1}{M} \cdot d_X(x, y) - A \leq d_Y(\varphi(g), \varphi(h)) \leq M \cdot d_X(x, y) + A,$$
(2) for every $y \in Y$ there is $x \in X$ such that $d_Y(y, \varphi(x)) \leq A$.

If there is such a map, (X, d_X) and (Y, d_Y) are *quasi-isometric* (it is a well-know exercise that this notion is a symmetric and transitive).

A subset $Y \subset X$ is θ-separated if $d_X(x, x') > \theta$ for every $x \neq x' \in X$. It is Θ-dense if for every $y \in X$ there is $x \in X$ with $d_X(x, y) < \theta$. It is a *net* if it is θ-separated and Θ-dense for some $\theta, \Theta > 0$. Obviously, the inclusion $(Y, d_X|_Y) \hookrightarrow (X, d_X)$ is a quasi-isometry if and only if Y is Θ-dense for some $\Theta > 0$.

2.2. **Curves.** A *curve* in a metric space (X, d_X) is a continuos map $\gamma \colon I \to X$ where $I \subseteq \mathbb{R}$ is some interval. If $I = [a, b]$, $\gamma(a)$ and $\gamma(b)$ are the *endpoints* of γ. A curve is *simple* if it is injective, it is a *simple closed curve* if it endpoints coincide, but it is injective otherwise.

A curve γ is *rectifiable* if it has finite length, which we denote by $|\gamma|$. By *geodesic* we mean a curve $\gamma \colon [a, b] \to X$ that realizes the distance between its endpoints. We will often omit the parametrization and identify a curve with its image $\gamma(I) \subseteq X$. If we do refer to a parametrization for a rectifiable curve, we will generally mean the arc-length parametrization.

2.3. **Riemannian surfaces.** A *surface (with boundary and corners)* is a topological surface X (with boundary ∂X) together with a smooth atlas so that every point has a neighbourhood that is diffeomorphic to an open set in $[0,\infty)\times[0,\infty)\subset\mathbb{R}^2$. A curve $\gamma\colon I\to X$ is *piecewise smooth* if it is smooth except perhaps at a discrete set of points in I. In particular, every connected component of ∂X can be seen as a piecewise smooth curve in X.

Convention 2.1. Throughout, we denote by Σ a complete connected Riemannian surface, possibly with boundary and corners, and by d_Σ the metric on Σ.

At places, we will need to work with contours of subsets of Σ, and we would need them to consists of unions of simple curves. This can be problematic in general, but one can always find arbitrarily small neighbourhoods that do have this property. This fact is very well-known, but we are having difficulties to find an appropriate reference.

Lemma 2.2. *For every subset $S\subset\Sigma$ and every $\epsilon>0$ there is a subsurface $\Sigma'\subseteq\Sigma$ that contains S and is contained in its ϵ-neighbourhood.*

Sketch of proof. For every $s\in S$ there is a small neighbourhood $B_s\subset B(s;\epsilon)$ that is diffeomorphic to the radius-one ball $B(0;1)$ in $\mathbb{R}^2$; $[0,\infty)\times\mathbb{R}$ or $[0,\infty)\times[0,\infty)$. For $t\leq 1$, let $B_s^t\subseteq B_s$ be the subset correspondig to $B(0;t)$. Since Σ is locally compact, there is a locally finite set $\{s_n\mid n\in\mathbb{N}\}\subseteq S$ such that $S\subseteq\bigcup_{\mathbb{N}}B_{s_n}^{1/2}$. We may then exploit the local finiteness of $\{s_n\mid n\in\mathbb{N}\}$ to choose radii $1/2<t_n<\epsilon$ so that closures of the neighbourhoods $B_{s_n}^{t_n}$ may only intersect in a generic way. That is, the boundary curves of the $B_{s_n}^{t_n}$'s are never tangent to one another. The union $\bigcup_{n\in\mathbb{N}}B_{s_n}^{t_n}$ is then a surface as required. ☐

2.4. **Graphs and triangulations.** Throughout, graphs will be connected and may have loops and multiple edges. A *metric graph* is a graph G together with an assignment of positive lengths to each edge. Such a graph is a metric space when equipped with the natural path metric. Given $I\subseteq(0,\infty)$, an *I-metric* path is a metric path where the lengths of the edges belong to I. A metric graph G is *locally finite* if it is proper as a metric space. Equivalently, if every ball of finite radius meets finitely many edges.

An embedding of a graph G in a surface Σ is a topological embedding of the corresponding 1-complex. In the presence of an embedding, we will sometime abuse notation and treat G as a subset of Σ, and say that $G\subset\Sigma$ is an an *embedded graph*. A *face* of an embedded graph G is a connected component of $\Sigma\smallsetminus G$. The *intrinsic metric* in a face F is the length metric inherited from Σ, and the *intrinsic diamater* is the diameter with respect to the intrinsic metric. Note that the intrinsic metric may be arbitrarily larger than the restriction of the metric of Σ to F.

Remark 2.3. If $G\subset\Sigma$ is an embedded graph such that every edge $e\in E(G)$ is a rectifiable curve, then G can be made into a metric graph by assigning to every $e\in E(G)$ its length $|e|$ in (Σ,d_Σ). The arc-length parametrizations of the edges then define a 1-Lipschitz embedding $G\hookrightarrow\Sigma$. In particular, the difficult part in finding embeddings that are quasi-isometric is the uniform lower bound on the metric distortion (1).

Remark 2.4. Let G be a metric graph and $\sigma\colon G\to\Sigma$ an embedding that is also Lipschitz. If the family of edges $\{\sigma(e)\mid e\in E(G)\}$ is locally finite (*i.e.* for every ball of finite radius $B \subseteq \Sigma$ there are only finitely many $e \in E(G)$ with $\sigma(e) \cap B \neq \emptyset$), then G is locally finite as a metric graph. If σ is also a quasi-isometric embedding, then the converse holds as well.

A *cell decomposition* of Σ can be described as a locally finite embedded graph $G \subset \Sigma$ such that $\partial\Sigma \subseteq G$ and every face F is homeomorphic to the unit open ball $B_1 \subset \mathbb{R}^2$ via a homeomorphism $B_1 \to \Sigma$ that extends to a continuous map of the disc $\mathbb{D}^2 = \overline{B_1}$ mapping $\partial\mathbb{D}^2$ into G (the image of $\partial\mathbb{D}^2$ is ∂F). This mapping defines a *boundary path* in G, and the *boundary size* of F is the number of edges in its boundary path, counted with multiplicity. We say that F is a *n-gon* if it has boundary size n. A n-gon F is *non-degenerate* if the boundary path is a simple curve (i.e. it consists of n distinct vertices and edges), otherwise it is *degenerate*. A cell decomposition is a *3-gonal decomposition* if every face is a 3-gon. By *triangulation* we mean a 3-gonal decomposition where every face is non-degenerate.

Remark 2.5. 3-gonal decomposition are called *pseudo-triangulations* in [10]. We preferred to use the former because pseudo-triangulations of surfaces have a different meaning in computational geometry. In Hatcher's notation [8], a choosing a 3-gonal decomposition would be equivalent to realizing Σ as a (unordered) Δ-complex. Other authors would call this a "generalized triangulation".

Remark 2.6. If $G \subset \Sigma$ is a 3-gonal decomposition, we can realize the barycentric subdivision G' as an embedded graph $G \subset G' \subset \Sigma$, so that G' is a triangulation of Σ. Moreover, if G is a metric graph such that the embedding $G \hookrightarrow \Sigma$ is a quasi-isometry, then we can arrange that the same holds for $G' \hookrightarrow \Sigma$.

For later reference, we prove below two simple lemmas.

Lemma 2.7. *Given a cell decomposition $G \subset \Sigma$ where every face has boundary of size at most three and $|V(G)| \geq 3$, we may choose a subgraph $H \subset \Sigma$ that defines a 3-gonal decomposition. Moreover, the following metric statements hold:*

(1) if G is a metric graph and $G \hookrightarrow \Sigma$ is a quasi-isometric embedding, then $H \hookrightarrow \Sigma$ is a quasi-isometric embedding as well;

(2) if the faces of G have intrinsic diameter bounded by R, then the faces of H have intrinsic diameter bounded by $3R$.

Proof. We simply proceed by removing loops and parallel edges until no more are left. To do so, we enumerate $V(G)$ as $v_1, v_2, \ldots$, and do the following procedure for $i = 1, 2, \ldots$, inductively assuming that after step $i-1$ each face incident with $v_1, \ldots, v_{i-1}$ has boundary of size at least three.

The 'bad' faces have boundary of size 0, 1 or 2, and can come in six types:

(1) 0-gons (∂F is one vertex);
(2) 1-gons;
(3) 2-gons walking twice along a single loop;
(4) 2-gons walking along two distinct loops on one vertex;
(5) 2-gons walking back and forth along a single edge;
(6) 2-gons on two distinct edges.

We start by observing that we will never meet faces of type (1), (3) or (5). In fact, either of those cases implies that $\overline{F} = \Sigma$, and hence $|V(G)| \leq 2$ (cases (1) and (5) only happen if $\Sigma = \mathbb{S}^2$, while (3) implies that $\Sigma = \mathbb{PR}^2$).

We may now describe the inductive step. Suppose there is a bad face F incident with v_i. It must be of type (2), (4), or (6). If F is of type (2), we simply remove the loop e contained in ∂F. Since e is only walked along once in ∂F, we see that by removing it we are merging F with another face F' having e on its boundary. This results in a face F'' that is still homeomorphic to an open ball, and has smaller boundary size than F'. This F'' is not a 3-gon yet, but since it is still incident with v_i it will be dealt with in a subsequent step of this process.

If F is of type (4) or (6), let e, e' be the two parallel edges in ∂F. We remove the longest between e and e', say e'. As before, this merges F with another face F' having e' on its boundary. Note that the resulting face is an open ball that has the same boundary size as F'.

We repeat this procedure until there are no more bad faces incident with v_i. Note that eventually we must stop, because G is locally finite and hence v_i is incident with only finitely many edges, and hence faces. At that point all faces incident with v_i (including those incident with any v_j, with $j < i$) are 3-gons. As the process goes through all i, it converges to a subgraph $H \subseteq G$ which is a 3-gonal decomposition of Σ.

The two metric statement are quite clear. For (1) observe that $d_G(v_i, v_j) = d_H(v_i, v_j)$ for any choice of vertices. In fact, it is clear that $\leq$ holds, and for $\geq$ it suffices to observe that removing loops does not have any effect on the path metric, and every time we removed a non-loop edge e' we left there a shorter parallel edge e.

For (2), observe that a face F of H with vertices v_1, v_2, v_3 is obtained by glueing (finitely many) faces F_i of G whose vertex-sets are contained in $\{v_1, v_2, v_3\}$. Since all the F_i have intrinsic diameter bounded by R, we certainly have that $d(v_i, v_j) \leq R$ for every $i, j = 1, 2, 3$ and that $\{v_1, v_2, v_3\}$ is R-dense in F. □

2.5. **Resolving intersections of curves.** We will later need a few technical results to control how curves in a surface intersect. For the most part, the following arguments could be substituted by a sentence along the lines of "taking a small perturbation, it is easy to see that. . . ". However, we thought it would do no harm to include a few more details.

Lemma 2.8. *Let Σ be a complete Riemannian surface, possibly with boundary and corners, and let $T \subset \Sigma$ be an embedded finite rooted tree consisting of piecewise smooth curves of finite length. Let o be the root and $\{p_i \mid i = 1, \dots n\}$ be the leaves. Then for every $\epsilon > 0$ we may find piecewise smooth curves γ_i joining p_i to o such that:*

(1) different γ_i only meet at o;
(2) $\bigcup_i \gamma_i$ is within Hausdorff distance ϵ from T;
(3) $||\gamma_i| - d_T(x_i, o)| \leq \epsilon$;
(4) the curves γ_i can only meet $\partial\Sigma$ at o.

Moreover, we may also replace (4) by

(4') $(\bigcup_i \gamma_i) \cap \partial\Sigma = T \cap \partial\Sigma$.

Proof. Orient the edges of T so that they point away from the origin. Subdividing edges in T if necessary, we may assume that every edge $e \in T$ is smooth. Moreover, we may assume that e only meets $\partial\Sigma$ on one side, so that we can construct a nearby path completely contained in the interior of Σ by "pushing it" in the other direction.

Formally, this can be done by appropriately choosing a smooth variation of e. Say that e is given a constant speed parametrization $e\colon [0, 1] \to \Sigma$, with $e(0)$ closer to the root $o \in T$. We consider a unit vector field $V(t)$ along e that is perpendicular to

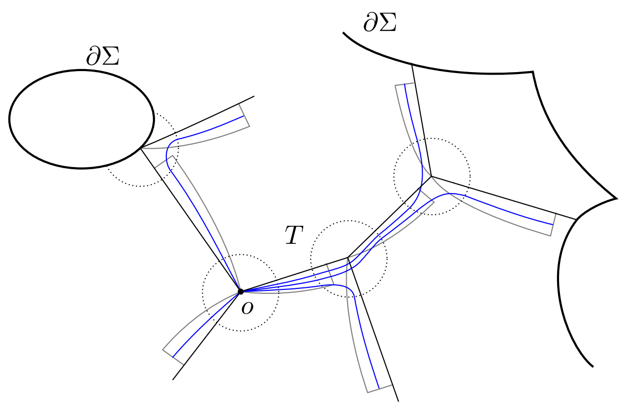

FIGURE 1. Transforming a rooted tree T to a collection of disjoint paths meeting at the root (blue lines in the picture).

$\dot{e}(t)$ and points towards the interior of Σ. Fix a very small $\epsilon' > 0$ (to be determined later). For every $t \in [0, 1-\epsilon']$, if s is small enough, $\exp(stV(t))$ is a point in Σ.[1] It follows that if we fix $0 < \delta \leq \epsilon'$ small enough, the mapping $(s,t) \mapsto \exp(stV(t))$ defines a smooth variation $E\colon [0, 1-\epsilon'] \times [0,\delta) \to \Sigma$ ([5, Chapter 5-4]) such that if we let $e_s(t) \coloneqq E(t,s)$ then:

- E is a smooth embedding on $(0, 1-\epsilon'] \times [0,\delta)$;
- $e(t) = e_0(t)$ for every $t \in [0, 1-\epsilon']$;
- $E(0,s) = e(0)$ for every $s \in [0,\delta)$;
- e_s is at Hausdorff distance at most $2\epsilon'$ from e;
- $||e_s| - |e|| < 2\epsilon'$.

Note that, as $s > 0$ varies, the smooth curves e_s only meet at $e(0)$ and are disjoint from $\partial\Sigma$ except perhaps at $e(0)$. Moreover, taking δ small enough we may also arrange that the curves e_s and $e'_{s'}$ are disjoint as the edges e, e' vary.

Pick $0 < s_1 < \cdots < s_{n_e}$, where n_e is the number of paths in T going from the leaves to the root o that cross e. As $e \in E(T)$ and $i = 0, \ldots, n_e$ vary, we may now arbitrarily join up all the curves e_{s_i} with appropriately chosen piecewise smooth paths contained in $2\epsilon'$-balls around the vertices of T (see Figure 1). The difference in length between these new paths and the paths in T is bounded by $6|V(T)|\epsilon'$.

The "moreover" part of the statement is proven by letting $s_1 = 0$ for every edge e and judiciously joining the curves at their extremities. □

2.6. **Arcs in minimal positions and ϵ-geodesics.** Suppose that Σ is homeomorphic to the disc $\mathbb{D}^2$. A curve $\gamma\colon [0,1] \to \Sigma$ is an *arc* if $\{0,1\} = \gamma^{-1}(\partial\Sigma)$. An *simple arc* is an arc that is a simple curve (it may be a simple closed curve meeting $\partial\Sigma$ at its endpoint). We say that two simple arcs α, β are in *minimal position* if $\alpha \cap \beta \smallsetminus \partial\Sigma$ consists of one point if the endpoints of α and β separate each other in $\partial\Sigma \cong \mathbb{S}^1$, and empty otherwise (this includes the case where α and β share an extremity).

A curve γ connecting two points $x, y \in \Sigma$ is an ϵ-*geodesic* if it is piecewise smooth and $|\gamma| < d(x,y) + \epsilon$.

[1] We only define the variation on $(0, 1-\epsilon']$ to avoid difficulties if $e(1)$ is a sharp corner in $\partial\Sigma$ and to ensure that different edges give rise to disjoint variations.

Lemma 2.9. *Let Σ be homeomorphic to $\mathbb{D}^2$ and α, β be piecewise smooth simple arcs that are ϵ_α resp. ϵ_β-geodesics. Further let $\gamma_1, \dots, \gamma_n$ be simple arcs that are in minimal position with respect to both α and β. Then we may find for every $\delta > 0$ a piecewise smooth simple arc β' with the same endpoints of β that is an $(\epsilon_\alpha + \epsilon_\beta + \delta)$-geodesic and is in minimal position with every α and all of the γ_i's.*

Proof. If $\alpha \cap \beta \smallsetminus \partial P$ is empty there is nothing to do. Otherwise, let β'' be the arc obtained by following β until it first meets α, then following α until its last intersection with β, and then resuming to follow β until its other end. This is an $(\epsilon_\alpha + \epsilon_\beta)$-geodesic that intersects α in one subsegment $\xi \subseteq \alpha$.

We claim that β'' is in minimal position with each γ_i. Let $\beta'' = \beta_- \cup \xi \cup \beta_+$, and $\beta = \beta_- \cup \beta_0 \cup \beta_+$. If $\gamma_i \cap \xi = \emptyset$, then $\beta'' \cap \gamma_i \subseteq \beta \cap \gamma_i$ and there is nothing to show. Otherwise, $\gamma_i \cap \xi = \gamma_i \cap \alpha$ consists of exactly one point and the endpoints of alpha separate those of γ_i. But then γ_i must intersect β_0, because the path following α until it reaches β_0, then β_0, and then α again separates the endpoints of γ_i. Since β is in minimal position with γ_i, the same is true for β''.

At this point we can apply the same technique of Lemma 2.8 to "push β'' away from α": this way find an arbitrarily small variation β' of β'' that is in minimal position with α. This can be done so that $|\beta' \cap \gamma_i| = |\beta'' \cap \gamma_i|$ for every $i = 1, \dots, n$. □

Corollary 2.10. *Let $\alpha_1, \dots, \alpha_n$ be arcs in a surface Σ homeomorphic to $\mathbb{D}^2$. Then there exist piecewise smooth ϵ-geodesics β_i which are simple arcs with the same endpoints as α_i and are pairwise in minimal position.*

Proof. Let $\overline{\alpha}_i$ be geodesics in P with the same endpoints of α_i. To begin with, we apply Lemma 2.8 to each curve separately find arbitrarily small variations thereof that only meet ∂P at their extremities. These are simple arcs that are δ-geodesic for arbitrarily small $\delta > 0$.

We may then iteratively fix $i = 1, \dots, n-1$ and apply Lemma 2.9 to the pair $\overline{\alpha}_i$, $\overline{\alpha}_j$ to put all the $\overline{\alpha}_j$ with $j > i$ in minimal position with $\overline{\alpha}_i$, while maintaining it in minimal position with all the $\overline{\alpha}_k$ with $k < i$. Doing this for every i concludes the proof. □

3. The main result

The following is our main technical result.

Theorem 3.1. *Let Σ be a complete Riemannian surface, possibly with boundary and corners. For every $\Xi > 0$ there exists a locally finite $(0, \Xi]$-metric graph G with a 1-Lipschitz embedding $G \hookrightarrow \Sigma$ such that:*

(1) $G \hookrightarrow \Sigma$ is a quasi-isometry;
(2) every face has diameter at most Ξ;
(3) G defines a 3-gonal decomposition of Σ.

By Remark 2.6, applying a barycentric subdivision to the 3-gonal decomposition of Theorem 3.1, we obtain:

Corollary 3.2. *Under the same hypotheses, we may further impose that G is a $(0, \Xi]$-metric graph triangulating Σ.*

The proof of Theorem 3.1 occupies the entirety of this section.

3.1. **Constructing a quasi-isometrically embedded metric graph.** The construction begins with a few arbitrary choices.

(3.1) Fix small parameters θ and Θ with $0 < \theta < \Theta/2$.

To all intents and purposes, one may once and for all fix *e.g.* $\theta = \Theta/3$ (and take $\Theta = \Xi/138$, where Ξ is the constant of Theorem 3.1), but we prefer to keep the parameters independent as they represent two different quantities: the denseness of vertices vs. the maximal length of edges.

The second choice is a pre-selection of vertices:

(3.2) let $X \subset \Sigma$ be a θ-dense net such that $X \cap \partial\Sigma$ is θ-dense on each connected component of $\partial\Sigma$ with respect to the intrinsic metric (this condition is vacuous if $\partial\Sigma = \emptyset$).

We consider a graph $\overline{G}$ obtained by taking $\partial\Sigma$ together with a union curves of length at most Θ joining points in X. The key point is to choose sufficiently many curves to represent every (local) homotopy type of short curves.

Formally, enumerate the pairs of points $\{x, y\} \subset X$ with $d_\Sigma(x, y) < \Theta$, and for every such pair choose some compact subsurface (with boundary) $\Sigma_{\{x,y\}} \subseteq \Sigma$ that contains $B(x; \Theta) \cap B(y; \Theta)$ and is itself contained in a $\theta/2$-neighbourhood of it (Lemma 2.2). Consider homotopy classes of curves $\gamma\colon [0, 1] \to \Sigma_{\{x,y\}}$ joining x to y. Observe that there are only finitely many such classes, because $\Sigma_{\{x,y\}}$ is a compact surface. Enumerate the homotopy classes that contain a curve γ of length less than Θ.

We will now iteratively choose a family $\Gamma_{\{x,y\}}$ of representatives for these homotopy classes. Specifically, for every such homotopy class, we choose a representative $\overline{\gamma}$ that is a piecewise smooth curve of length less than Θ and intersects all the previosly chosen representatives (including those belonging to $\Gamma_{\{x',y'\}}$ for previously considered pairs $\{x', y'\} \subset X$) and $\partial\Sigma$ in finitely many points. Let

$$\Gamma_\Sigma \coloneqq \bigcup \bigl\{ \Gamma_{\{x,y\}} \bigm| \{x, y\} \subset X,\ d_\Sigma(x, y) < \Theta \bigr\}.$$

Since X is a net and each $\Gamma_{\{x,y\}}$ is a finite collection of curves, the collection Γ_Σ is locally finite in Σ.

We let $\overline{G}$ be the graph traced by $\Gamma \cup \partial\Sigma$. Specifically, we let

$$V(\overline{G}) \coloneqq X \cup \bigcup \bigl\{ \overline{\gamma} \cap \overline{\gamma}' \bigm| \overline{\gamma}, \overline{\gamma}' \in \Gamma \bigr\} \subset \Sigma$$

and let $E(\overline{G})$ be the set of all the subsegments of curves in Γ or $\partial\Sigma$ that connect points in $V(\overline{G})$ and do not contain any other vertex in their interior. Note that the edges coming from $\partial\Sigma$ have length bounded by Θ because $X \cap \partial\Sigma$ is θ-dense in the intrinsic metric.

This construction defines a graph, which will generally have loops and multiple edges (especially if Θ is large) and is—by construction—embedded in Σ. By construction, when equipped with its arc-length metric, $\overline{G}$ is a locally finite $(0, \Theta]$-metric graph with a 1-Lipschitz embedding $\overline{G} \hookrightarrow \Sigma$ (Remark 2.3).

The following is the key property of $\overline{G}$.

Lemma 3.3. *Let $\gamma\colon [0, 1] \to \Sigma$ be a rectifiable curve with endpoints in $X \subseteq V(\overline{G})$. Then there exists a path $\bar{\gamma}$ in $\overline{G} \subset \Sigma$ such that*

(1) γ and $\bar{\gamma}$ are homotopic within the 2Θ-neighbourhood of γ;

(2) $|\bar{\gamma}|$ *satisfies*

$$|\bar{\gamma}| \leq \frac{\Theta}{\Theta - 2\theta}|\gamma| + \Theta.$$

Proof. Parameterise γ by arc length. Let $n = \lfloor |\gamma|/(\Theta - 2\theta) \rfloor$ and for every $0 \leq k \leq n$ let $p_k := \gamma(k(\Theta - 2\theta))$. Further let $p_{n+1} := y$. For every $0 \leq k \leq n+1$ let x_k be a point in X nearest to p_k (in particular, $x_0 = p_0$ and $x_{n+1} = p_{n+1}$). Further let α_k be a geodesic path from p_k to x_k.

Let γ_k denote the segment of γ between p_k and p_{k+1}. The curve $\alpha_k^{-1}\gamma_k\alpha_{k+1}$ has length less than Θ. By construction, $\alpha_k^{-1}\gamma_k\alpha_{k+1}$ is contained in the subsurface $\Sigma_{\{x_k, x_{k+1}\}}$, and is hence homotopic within $\Sigma_{\{x_k, x_{k+1}\}}$ to one of the fixed representatives $\overline{\gamma}_k \in \Gamma_{\{x_k, x_{k+1}\}}$. In particular, the homotopy between $\alpha_k^{-1}\gamma_k\alpha_{k+1}$ and $\overline{\gamma}_k$ is entirely contained in the $(\Theta + \theta)$-ball centered at p_k. By construction, $\bar{\gamma}_k$ is entirely contained in $\overline{G}$.

We may then join the curves $\bar{\gamma}_k$ to obtained a curve $\bar{\gamma}$ that is homotopic to γ. Moreover, we see that

$$|\bar{\gamma}| \leq (n+1)\Theta \leq \frac{\Theta}{\Theta - 2\theta}|\gamma| + \Theta. \qquad \square$$

Corollary 3.4. *The embedding* $\overline{G} \hookrightarrow \Sigma$ *is a quasi-isometry.*

Proof. Let x, y be points in X. Applying Lemma 3.3 to a geodesic connecting them in Σ, we deduce that

$$(1 - 2\theta/\Theta) d_{\overline{G}}(x, y) - \Theta + 2\theta \leq d_\Sigma(x, y) \leq d_{\overline{G}}(x, y). \tag{3.3}$$

The claim follows, because X is coarsely dense in both $\overline{G}$ and Σ. $\square$

Remark 3.5. We observe that when the parameter θ goes to zero the quasi-isometry constants improve and the embedding becomes ever closer to be a rough isometry. On the flip side, the resulting graph $\overline{G}$ will have many more vertices and edges.

Remark 3.6. If Σ is a surface where the surfaces $\Sigma_{\{x,y\}}$ can be chosen to be simply connected (*e.g.* because they are convex), and X is a uniform net, then the graph $\overline{G}$ has bounded degree.

3.2. **Bounding the diameter of the faces.** We shall now show that $\Sigma \smallsetminus \overline{G}$ has connected components of uniformly bounded diameter, which hence have compact closure.

Proposition 3.7. *Each face of* $\overline{G} \subset \Sigma$ *has* d_Σ*-diameter bounded by* 13Θ*.*

Proof. Assume by contradiction that there is a face F with diameter greater than 13Θ. We can hence pick two points $p, p' \in F$ with $d_\Sigma(p, p') = 13\Theta$ and a smooth, simple curve $c \subset F$ joining them. We want to reach a contradiction by showing that there is an edge in $\overline{G}$ crossing c.

The constant 13Θ is large enough that we can split c into sub-segments $c = c_- \cup c_m \cup c_+$ such that

1. $d_\Sigma(c_-, c_+) \geq 2\theta$;
2. both the endpoints of c_- and those of c_+ are at d_Σ-distance greater than $(4\Theta + 3\theta)$ from one another.

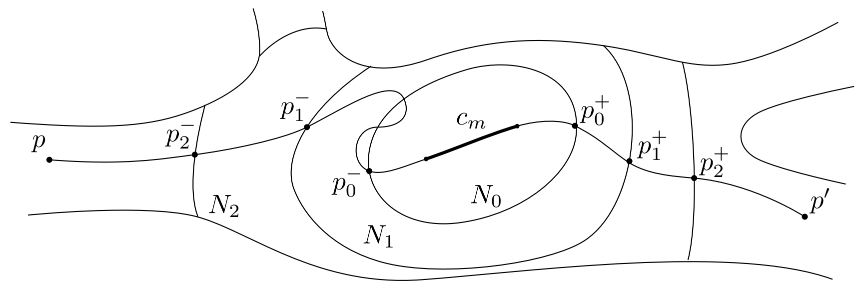

FIGURE 2. Constructing nested surfaces around the middle portion of a very long curve c.

Let N_0 be the closed 2θ-neighbourhood of c_m. Enlarging it a little if necessary, we assume that it is a (compact) subsurface of Σ (Lemma 2.2). Let c_0 be the closure of the connected component of $c \smallsetminus \partial N_0$ containing c_m.

Let N_1 be the closed $(2\Theta + \theta)$-neighbourhood of N_0. We again assume that it is a submanifold of Σ and let c_1 be the closure of the connected component of $c \smallsetminus \partial N_1$ containing c_0. Similarly, N_2 is the 2Θ-neighbourhood of N_1 and c_2 is the closure of the component of $c \smallsetminus \partial N_2$ containing c_1. Note that condition (2) implies that the curves c_i are properly nested, and that for each $i = 0, 1, 2$ the curve c_i meets the boundary ∂N_i in two points $p_i^- \in c_-$ and $p_i^+ \in c_+$. The construction is pictured in Figure 2.

We now argue by cases depending on whether $N_1 \smallsetminus c_1$ is connected or not.

Case I. Assume $N_1 \smallsetminus c_1$ is connected. Since $X = V(\overline{G})$ is θ-dense, we can choose a point $x \in X \cap N_1$. We can also choose a simple closed curve $\gamma \subset N_1$ that starts at x and intersects c_1 at exactly one point (head towards c_1, cross it, and return to x using that $N_1 \smallsetminus c_1$ is path connected). Let $\bar{\gamma}$ be the path in $\overline{G}$ given by Lemma 3.3. By construction, the homotopy between γ and $\bar{\gamma}$ takes place within N_2. Since c_2 is a simple arc joining two points in ∂N_2 and $\gamma \cap c_2 = \gamma \cap c_1$ is one point, it follows that $\bar{\gamma}$ must also intersects c_2 in an odd number of points (this can be seen in several ways, for instance with a homological argument, or by taking the double of N_2 in order to prolong c_2 to a closed curve and apply standard intersection theory see e.g. [7, Section 2.4]). Hence c crosses an edge of $\overline{G}$, contradiction.

Case II (Figure 3). Assume $N_1 \smallsetminus c_1$ is disconnected. Then $N_0 \smallsetminus c_0$ is a fortiori disconnected (every point in N_1 is connected to N_0 via a path avoiding c). It follows that p_0^+ and p_0^- must belong to the same component C of ∂N_0 (if a simple arc c' in a connected surface Σ' meets a boundary component C' at a single point, then $\Sigma' \smallsetminus c'$ is path connected because so is $C' \smallsetminus c'$).

Note that C is a circle and p_0^+ and p_0^- cut it into two segments, each of which connects c_- to c_+. By continuity, (1) implies that each of these segments contain points that are at distance greater than θ from both c_- and c_+. Let q and q' be two such points, one for each segment. Since every point on C is at distance greater than θ from c_m, it follows that q and q' are at distance greater than θ from $c = c_- \cup c_m \cup c_+$. Let x_- and x_+ be points in X closest to q_- and q_+ respectively. Since the balls $B(q;\theta)$ and $B(q';\theta)$ are entirely contained in $N_1 \smallsetminus c$, the points x and x' belong to different components of $N_1 \smallsetminus c_1$.

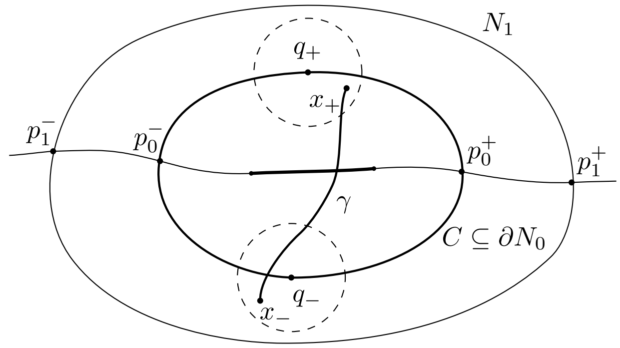

FIGURE 3. Constructing a curve γ with endpoints in X that cut c in the case that $N_1 \setminus c_1$ is disconnected.

We arbitrarily choose a curve γ joining x_- with x_+ without leaving the (closed) θ-neighbourhood of N_0, and apply Lemma 3.3 to obtain a curve $\bar{\gamma}$ in $\overline{G}$ homotopic to it. The homotopy between γ and $\bar{\gamma}$ takes place in N_1 by construction, and in particular $\bar{\gamma}$ is a curve connecting x and x' in N_1. Since x and x' are in two distinct components of $N_1 \setminus c_1$, it follows that $\bar{\gamma}$ intersects c. $\square$

3.3. **Proving that faces are discs.** The construction of $\overline{G} \subset \Sigma$ of Section 3.1 is designed to give rise to a cell decomposition of Σ (Section 2.4). To prove it, we start with a few preliminary observations, which will be of use both here and in the next section.

Suppose $G \subset \Sigma$ is some embedded graph whose edges are piecewise smooth and have finite length. Given a face F of G with compact closure, we can consider it with its own intrinsic path-metric d_F and denote by $\widehat{F}$ its completion. Observe that $\widehat{F}$ will generally differ from the closure $\overline{F} \subseteq \Sigma$, and it is a compact Riemannian surface with boundary and corners (taking the completion of the interior with respect to the intrinsic metric has the effect of "opening up" non-trivial glueings of the boundary). In particular, $\partial\widehat{F} \coloneqq \widehat{F} \setminus F$ is homeomorphic to a disjoint union of loops. Note that the inclusion $F \hookrightarrow \Sigma$ extends to a 1-Lipschitz surjective map $p\colon (\widehat{F}, d_F) \to (\overline{F}, d_\Sigma)$ such that $p^{-1}(\partial F) = \partial\widehat{F}$. In turn, p descends to a homeomorphism when quotienting out the boundary

$$\widehat{F}/\partial\widehat{F} \cong \overline{F}/\partial F \cong \Sigma/(\Sigma \setminus F).$$

Let now F be a face of the embedded graph $\overline{G}$ constructed in Section 3.1. To prove that $\overline{G}$ is a cell decomposition it is enough to show that $\widehat{F}$ is homeomorphic to the disc $\mathbb{D}^2$. Let $\pi\colon \widehat{F} \to \widehat{F}/\partial\widehat{F}$ be the quotient map. We observe the following.

Lemma 3.8. *If $\gamma\colon I \to \widehat{F}$ is a path with endpoints in $\partial\widehat{F}$, then the closed loop $\pi \circ \gamma\colon I \to \widehat{F}/\partial\widehat{F}$ is null-homotopic.*

Proof. Consider the path $p \circ \gamma$ in Σ and note that if we identify $\widehat{F}/\partial\widehat{F} \cong \Sigma/(\Sigma \setminus F)$ then $\pi \circ \gamma = \pi_\Sigma \circ p \circ \gamma$, where π_Σ is the quotient map $\Sigma \to \Sigma/(\Sigma \setminus F)$. We may prolong $p \circ \gamma$ along $\overline{G}$ to obtain a path γ' with endpoints in X, and we observe that $\pi \circ p \circ \gamma$ and $\pi \circ \gamma'$ are homotopic as closed loops. By Lemma 3.3, γ' is homotopic to a path $\overline{\gamma}' \subseteq \overline{G} \subseteq \Sigma \smallsetminus F$. But then we are done, because $\pi \circ \overline{\gamma}'$ is constant in $\Sigma/(\Sigma \smallsetminus F)$. □

Given the the classification of compact surfaces, the following fact is an exercise in algebraic topology (which can be solved using either fundamental groups or homology computations).

Fact 3.9. *A connected compact surface Δ with $\partial\Delta \neq \emptyset$ is homeomorphic to a disk if and only if the image of every path $\gamma\colon I \to \Delta$ with endpoints in $\partial\Delta$ under the quotient map $\pi\colon \Delta \to \Delta/\partial\Delta$ is null-homotopic (*i.e. $\pi \circ \gamma = 0 \in \pi_1(\Delta/\partial\Delta)$*).*

Combining Fact 3.9 with Lemma 3.8, we obtain the claimed result:

Corollary 3.10. *$\overline{G} \subset \Sigma$ defines a cell decomposition of Σ.*

3.4. **Cutting to 3-gonal decompositions.** Now that we know that $\overline{G} \subset \Sigma$ gives a cell decomposition, it is a trivial matter to improve the state of affairs to obtain a 3-gonal decomposition $G \subset \Sigma$ by adding extra edges if necessary. The only issue is that, since Proposition 3.7 does not bound the *intrinsic* diameter, the newly added edges may be too long, thus spoiling the metric properties of the embedding. To address this we need to work a little more.

Let F be a face, $\widehat{F}$ its completion in the intrinsic metric, and $p\colon \widehat{F} \to \overline{F}$ as in Section 3.3. Observe that the graph structure of $\partial F \subseteq \overline{G}$ lifts to a graph structure on $\partial\widehat{F}$, where vertices are preimages of vertices and edges are lifts of an edges. Moreover, every edge $e \in \partial\widehat{F}$ has the same length of its image $p(e) \in E(\overline{G})$. Note that $\widehat{F}$ is a *piecewise smooth polygon* with edge-length bounded by Θ. That is, it is a disc whose boundary (considered as a graph) is a cycle consisting of finitely many edges, each of which is piecewise smooth and has length bounded by Θ. Moreover, since X is θ-dense, the polygon $\widehat{F}$ is also *θ-thin*. That is, every point in $\widehat{F}$ is within distance θ from $\partial\widehat{F}$. To conclude the proof of Theorem 3.1, it is then enough to explain how to triangulate such a polygon with edges of controlled length.

Recall that a curve is an ϵ-geodesic if it length realizes the distance of its endpoints up to ϵ Section 2.6.

Lemma 3.11. *If $e \subset \partial P$ is an edge in a θ-thin polygon, and e is an ϵ-geodesic, then every point in e is within distance $2\epsilon + 3\theta$ from $\partial P \smallsetminus e$.*

Proof. Given $x \in e$, let N be a subsurface with $B(x, \epsilon + 2\theta) \subseteq N \subseteq B(x, 2\epsilon + 2\theta)$ (Lemma 2.2). If N intersects $\partial P \smallsetminus e$, there is nothing to do. If that is not the case, note that ∂N is contained in $e \cup (P \smallsetminus B(x, \epsilon + 2\theta))$ (here ∂N denotes the surface boundary, not the topological boundary of the subset $N \subset P$). Let e_0 be the component of $e \cap B(x, \epsilon/2 + \theta)$ containing x, and decompose e into (non-empty) segments $e_- \cup e_0 \cup e_+$.

Observe that no point in $\partial N \smallsetminus e$ can be within distance θ from both e_- and e_+, because otherwise e would not be an ϵ-geodesic (since $|e_0| \geq 2\theta + \epsilon$). On the other hand the component of ∂N containing x is a simple closed curve γ that contains e_0. It follows that $\gamma \smallsetminus e$ must have a component γ' that joins a point in e_- to a point in e_+. Since γ' is contained in $P \smallsetminus B(x, \epsilon + 2\theta)$, it contains points at distance greater than θ from e_- and e_+ respectively, and since it is connected it must also contain a point that is at distance greater than θ from both e_- and e_+ simultanously. That point is also at distance greater than θ from $e_0 \subset B(x, \epsilon/2 + \theta)$, so it is not θ-close to e at all, and must hence be in the θ-neighbourhood of $\partial P \smallsetminus e$. □

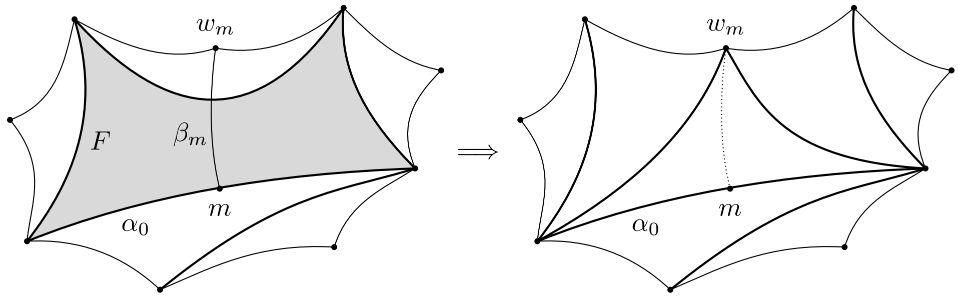

FIGURE 4. The thick lines are disjoint ϵ-geodesics. In the case where the curves crossing the short path β_m are not connected, removing them and adding ϵ-geodesics from their extremities to v_m increases the cardinality of the family

Proposition 3.12. *Let $\theta < \Theta/2$ and let P be a θ-thin piecewise smooth n-gon with edge-length bounded by Θ and $n \geq 3$. We may triangulate P by adding finitely many arcs of length at most $12(3\theta + \Theta)$ and no extra vertices.*

Proof. For convenience, let $\kappa := 3\theta + \Theta$. We arbitrarily fix a small $0 < \epsilon < \Theta/4$. Let $A = \{\alpha_i \mid i = 1, \dots, n\}$ be a family of curves of maximal cardinality such that

- each α_i is a piecewise smooth ϵ-geodesic of length less than 12κ that join non-adjacent vertices in ∂P and is otherwise contained in the interior of P;
- no pair of vertices in ∂P is joined by more than one curve in A;
- the curves α_i do not intersect in the interior of P.

We claim that A yields the required triangulation of P.

We prove the claim by contradiction, assuming that $P \smallsetminus A$ has a face F that is not a triangle. Observe that

(3.4) ∂F must contain an edge α_0 of length at least 6κ,

because otherwise we may further cut it using ϵ-geodesics of length less than 12κ by joining non-consecutive vertices in ∂P at distance less than 12κ (Corollary 2.10). Since the edges in ∂P have length at most Θ, α_0 must be one of the curves in A.

The ϵ-geodesics α_0 cuts P into two piecewise smooth polygons, P' and P'', both of which are θ-thin. We may suppose that F is contained in P'. Let $v_-, v_+ \in \partial P'$ be the endpoints of α_0. By Lemma 3.11, every point $m \in \alpha_0$ is within distance $3\theta + \epsilon$ from $\partial P' \smallsetminus \alpha_0$. In particular, there must be a vertex in $\partial P'$ that is at distance at most

$$3\theta + \Theta/2 + \epsilon < \kappa \tag{3.5}$$

from m. Let $w_m \in \partial P'$ be a vertex in P' that is closest to m.

Observe that if m is chosen so that both $d(v_-, m)$ and $d(m, v_+)$ are at least κ, then w_m is neither v_- nor v_+, and is hence some other vertex in $\partial P' \smallsetminus \alpha_0$. Given such a point m, let β_m be a piecewise smooth δ-geodesic connecting m and w_m in P'. If $\delta > 0$ is small enough, we may apply Lemma 2.9 to modify curves in $A \smallsetminus \{\alpha_0\}$ so that they are still ϵ-geodesic and are in minimal position with respect to β_m. Moreover, by (3.5), δ can be chosen small enough so that

$$|\beta_m| + \epsilon \leq (3\theta + \Theta/2 + \epsilon) + \delta + \epsilon < \kappa = 3\theta + \Theta. \tag{3.6}$$

Let $B_m \subseteq A$ be the set of curves that intersect β_m outside ∂P. We observe that—as a graph—B_m contains no closed loop (*i.e.* it is a forest), and it is non-empty because α_0 belongs to it. By Euler characteristic,

$$|V(B_m)| = |B_m| + |\{\text{components of } B_m\}| \geq\geq |B_m| + 1.$$

Also note that if $\alpha \in B_m$ intersects β_m in some point x which splits α as $\alpha_- \cup \alpha_+$, then

(3.7) both $|\alpha_-|$ and $|\alpha_+|$ are at least as large as $d(x, w_m)$,

because otherwise w_m would not be a closest vertex to m. One can hence perform the following construction:

(3.8) Remove every curve in B_m from A. In their place, add for every vertex in $V(B_m)$ which is not adjacent to w_m in $\partial P'$ one arc connecting it with w_m. Then add back α_0 and apply Corollary 2.10, to make these curves into piecewise smooth ϵ-geodesic arcs in minimal position (Figure 4). Denote by A_m be the resulting family of curves.

By (3.7), the curves in A_m can be taken to have length less than 12κ: if v is an endpoint of $\alpha \in B_m$, it can be joined to w_m by following α until it reaches β_m and then following the latter until w_m. We would like to reach a contraction by showing that $|A_m| > |A|$.

We argue by cases. If $w_m \in \partial F$ then B_m consisted uniquely of the curve α_0. Since F is not a triangle, at least one between v_- and v_+ is was not adjacent to w_m in ∂F, thus A_m now has strictly more curves than A, so we are done. We may hence assume that $w_m \notin \partial F$, and hence $|B_m| > 1$. Since w_m is adjacent to at most two vertices, we see that

$$|A_m| \geq |A \smallsetminus B_m| + (|V(B_m)| + 1) - 2 \geq |A|. \tag{3.9}$$

If B_m is not connected, the above inequality is strict.

It remains to deal with the case where $|B_m| > 1$ and B_m is connected. This case is the most delicate, and to deal with it we may have to make different choices of m (see below).

Let $\alpha_m \in A_m$ be the first curve encountered by β after α_0. Since α_m disconnects P' and β intersects it exactly once (because they are in minimal position), we see that α_m must meet α_0 at one of its extremities. For otherwise B_m would not be connected. Let $v_m \in \partial F$ denote the other end point of α_m.

If $\alpha_m \cap \alpha_0 = v_+$, then v_m is not adjacent to v_- in ∂F, so if $d_F(v_-, v_m)$ was to be less than 12κ we could add another curve to A_m and conclude the proof. The symmetric argument applies if $\alpha_m \cap \alpha_0 = v_-$. This reduces to the following case: for every $m \in \alpha_0$ with $d(v_-, m), d(m, v_+) \geq \kappa$, β_m meets a curve $\alpha_m \subset \partial F$ connecting one of v_- or v_+ to a point v_m at d_F-distance at least 12κ from the other one.

We claim that if $d_F(v_-, m) = \kappa$, then α_m must intersect α_0 at v_-. Suppose that was not the case, and let $x_m = \alpha_m \cap \beta_m$. Then

$$\begin{aligned} d_F(x_m, v_+) &\geq d_F(v_+, m) - d_F(x_m, m) \\ &\geq (|\alpha_0| - \epsilon - \kappa) - |\beta_m| \\ &\geq 6\kappa - \kappa - (|\beta| + \epsilon) \geq 4\kappa, \end{aligned}$$

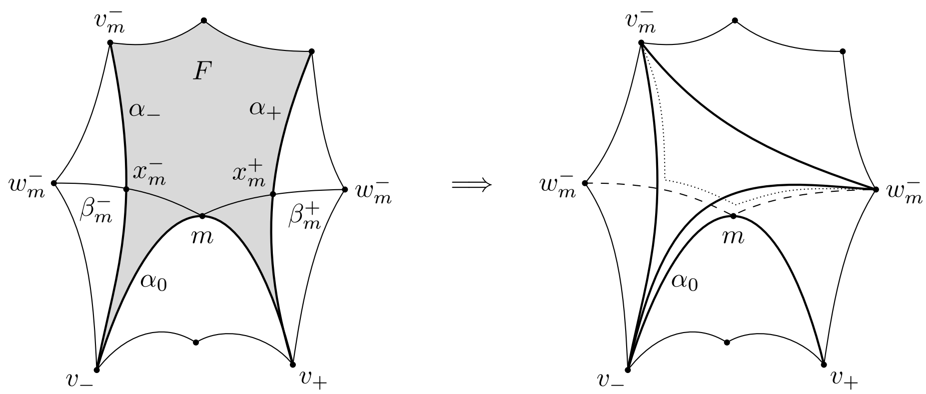

FIGURE 5. The thick lines are disjoint ϵ-geodesics. If there is a point m connected to ∂P via two short paths $\beta_m^\pm$ that intersect ϵ-geodesics connected to the starting curve α_0, we increase the number of disjoint ϵ-geodesic by removing one of them and adding ϵ-geodesics from the extremities of the other one.

where we used ϵ-geodesicity of α_0 at the second step. But then

$$\begin{aligned} d_F(v_-, v_m) &\leq d_F(v_-, x_m) + d_F(x_m, v_m) \\ &\leq (d_F(v_-, m) + d_F(m, x_m)) + (|\alpha_m| - d_F(x_m, v_+)) \\ &\leq (\kappa + |\beta|) + (12\kappa - 4\kappa) < 12\kappa. \end{aligned}$$

The symmetric holds if $d_F(v_+, m) = \kappa$.

By continuity, this implies that there must be some point $m \in \alpha_0$ for which there are two possible choices of w_m, one of which yields as α_m a curve $\alpha_- \in A$ containing v_-, and the other yields as α_m a curve $\alpha_+ \in A$ containing v_+. We denote those vertices w_m^- and w_m^+ respectively, and similarly denote β_m^-, β_m^+ (Figure 5).

We may assume that $d(m, v_+) \leq d(m, v_-)$, which implies that

$$d(m, v_-) \geq (6\kappa - \epsilon)/2.$$

We then do the construction (3.8) with respect to w_m^+ to obtain A_m. Let x_m^- be the intersection point of α_- and β_m^-, and let v_m^- be the endpoint of α_- away from v_-. Observe that

$$\begin{aligned} d_F(v_m^-, x_m^-) &\leq |\alpha_m^-| - d_F(v_-, x_m^-) \\ &\leq |\alpha_m^-| - (d_F(v_-, m) - d_F(m, x_m^-)) \\ &\leq |\alpha_m^-| - (6\kappa - \epsilon)/2 + |\beta_m^-| \\ &\leq |\alpha_m^-| - 2\kappa \end{aligned}$$

the curve following α_- from v_m^- to x_m^- and then going to w_m following β_m^- and β_m^+ has length bounded by

$$d_F(v_m^-, x_m^-) + |\beta_m^-| + |\beta_m^+| \leq |\alpha_m^-|.$$

This means we can add one extra curve to A_m, which completes the proof. □

At this point it is easy to complete the proof of Theorem 3.1.

Proof of Theorem 3.1. The construction in Section 3.1 yields a locally finite embedded graph $\overline{G} \subset \Sigma$ such that $\overline{G} \subset \Sigma$ is a cell-decomposition (Corollary 3.10), and $\overline{G}$ is a $(0, \Theta]$-metric graph where the embedding $\overline{G} \hookrightarrow \Sigma$ is a 1-Lipschitz (Remark 2.3) quasi-isometry (Corollary 3.4).

For every face F of $\overline{G} \subset \Sigma$, the completion $\widehat{F}$ with respect to the intrinsic metric is a θ-thin piecewise smooth polygon $P := \widehat{F}$ with edge-length bounded by Θ. Apply Proposition 3.12 to every such polygon to obtain a new piecewise smoothly embedded graph $G' \subset \Sigma$ whose edges are ϵ-geodesic (for any arbitrarily fixed $0 < \epsilon < \Theta/4$) of length at most $12(3\theta + \Theta) < 30\Theta$, and such that $G' \subset \Sigma$ is a cell decomposition where all the faces are θ-thin polygons with boundary of size at most 3. In particular, the intrinsic diameter of every face is bounded by $3/2 \cdot 30\Theta + \theta < 46\Theta$.

Observe that G' remains locally finite, because each face contributes finitely many new edges. Since G' is obtained from $\overline{G}$ by adding extra edges, the 1-Lipschitz embedding $G' \hookrightarrow \Sigma$ must be a quasi-isometry as well (adding extra edges in this fashion can only improve the lower bound in (3.3)). Finally, we may apply Lemma 2.7 to obtain a 3-gonal decomposition $G \subseteq G' \subset \Sigma$ so that the embedding $G \hookrightarrow \Sigma$ is a quasi-isometry and the faces have intrinsic diameter bounded by 138Θ. □

Remark 3.13. Note that at no point in the construction of G Theorem 3.1 we needed to add extra vertices to $\overline{G}$. In particular, every vertex in G is at distance at most $\Theta/2$ from a point in the originating net X.

4. Applications

In [10], the *modulus* of a cell decomposition $G \subset \Sigma$ is defined to be the supremum of the length of the 1-cells. In particular, if $G \hookrightarrow \Sigma$ is a 1-Lipschitz embedded $(0, \Theta]$-metric graph defining a cell decomposition, then the decomposition has modulus bounded by Θ.

Theorem 4.1. *Let (Σ, d_Σ) be a complete Riemannian surface possibly with boundary and corners. For each net X in (Σ, d_Σ) such that $X \cap \partial\Sigma$ is a net in each component of $\partial\Sigma$ with its intrinsic metric, there exists a 3-gonal decomposition $\mathcal{D}$ of Σ that has finite modulus and whose* 0*-skeleton is exactly X.*

Moreover, if X is uniform, then $\mathcal{D}$ can be chosen so that its 1*-skeleton equipped with the simplicial metric is quasi-isometric to (Σ, d_Σ).*

If $\partial\Sigma = \emptyset$, the first statement is [10, Theorem 4.4] (called the 'main technical result' of that paper). The analogous fact with $\partial\Sigma \neq \emptyset$ is used in [10, Theorem 8.1], where it is remarked that their proof of [10, Theorem 4.4] works in this case as well as in the boundary-free case.

The second statement bypasses Proposition 4.5 and Lemma 5.2 of [10] without assuming planarity at infinity or curvature bounded below.

Proof. Suppose that X is a θ-densee net such that $X \cap \partial\Sigma$ is θ-dense in each component with its intrinsic metric. Use it in the construction of Section 3.1 with respect to some fixed $\Theta > 2\theta$. Then Theorem 3.1 yields a 3-gonal decomposition $G \subseteq \Sigma$ of finite modulus with $X \subseteq V(G)$ and such that $(G, d_G) \hookrightarrow (\Sigma, d_\Sigma)$ is a 1-Lipschitz quasi-isometry.

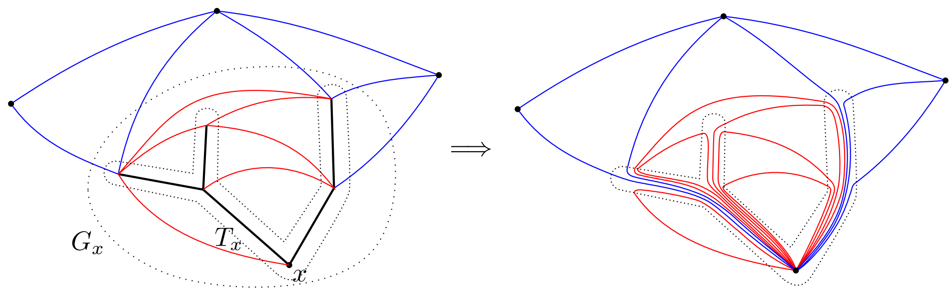

FIGURE 6. Contracting the spanning tree T_x to the point x. The subgraph $G_x \subset G$ is the union of T_x with the red edges in the left hand side.

Let $\mathcal{V} = \{V_x \mid x \in X\}$ be the Voronoi decomposition of $V(G)$ with respect to X in the metric d_G, *i.e.* $V_x \coloneqq \{v \in V(G) \mid d_G(v, X) = d_G(v, x)\}$. By perturbing d_G slightly we can break all ties, *i.e.* ensure that the V_x are pairwise disjoint.

Let G_x be the subgraph of G induced by the vertices in V_x. Note that each G_x is connected. As a matter of fact, we even have

(4.1) $\operatorname{diam}(G_x) \leq \Theta$ for every $x \in X$,

because each vertex of G has distance at most $\Theta/2$ from X (Remark 3.13).

For each $x \in X$ we pick a *geodetic spanning tree* T_x of V_x rooted at x.[2] We can further choose an ϵ-neighbourhood N_x of T_x so that $G \cap N_x$ is a tree rooted at x (this has the effect of adding to T_x the terminal part of each edge $e \in G$ which ends in T_x without being contained in it). We can then apply Lemma 2.8 to $G \cap N_x$ in such a way that $\partial N_x \cap G$ is preserved. Pictorially, this has the effect of 'dragging' T_x onto x inside a small enough neighbourhood N_x of T_x in Σ, pulling all edges incident with T_x as we drag (Figure 6).

By applying this procedure to each $x \in X$, we modify G into a graph $\widetilde{G}$ embedded into Σ with $V(\widetilde{G}) = X$. Note that there is an one-to-one correspondence between the faces of G and the faces of $\widetilde{G}$. The latter generally fails to be a 3-gonal decomposition, because contracting the edges of the T_x will decrease the boundary size of some of the faces. However, we may immediately apply Lemma 2.7 to obtain a 3-gonal decomposition $\mathcal{D}$ of Σ. Observe that

(4.2) The lengths of the edges in $\widetilde{G}$ are uniformly bounded above.

In fact, by construction every edge in $\widetilde{G}$ has length within ϵ from the length of a path consisting of the concatenation an edge in G and two branches in some spanning trees T_x (Lemma 2.8), and the latter is bounded because of (4.1). This completes the proof of the first statement. (Note that at this point it is not yet clear whether the embedding $\widetilde{G} \hookrightarrow \Sigma$ is a quasi-isometry.)

For the second one, let $\widetilde{G}^{(1)}$ denote the 1-skeleton with its simplicial metric, *i.e.* the graph $\widetilde{G}$ where every edge is given length one. Observe that the lengths of the non-loop edges in $\widetilde{G}$ are uniformly bounded below because they join distinct points

[2] This means that for each $y \in V_x$ we have $d_{V_x}(x, y) = d_{T_x}(x, y)$. Such a T_x can be obtained by ordering the vertices of V_x according to their distance from x, and recursively joining the next closest vertex to the tree constructed so far.

in X, which is a net. Since the length of all the edges is uniformly bounded above by (4.2), we see that

(4.3) $\widetilde{G}^{(1)}$ and $\widetilde{G} \subset \Sigma$ are a quasi-isometric via the identity map.

By construction, the embedding $\widetilde{G} \hookrightarrow \Sigma$ is a 1-Lipschitz map which is coarsely surjective. Together with (4.3), it is thus enough to show that if the net X is uniform, then there are constants L, A such that for every $x, y \in X$ we have a bound $d_{\widetilde{G}^{(1)}}(x,y) \leq L d_\Sigma(x,y) + A$.

Observe that contracting each G_x into a vertex we obtain a minor H of G, which we equip with its simplicial metric. Note that, as a abstract graphs, H is a subgraph of $\widetilde{G}^{(1)}$ obtained by removing loops and some parallel edges. In particular, the identity map on X defines a quasi-isometry between H and $\widetilde{G}^{(1)}$.

Given $x, y \in X$, let γ be a x–y geodesic in Σ, and $\bar\gamma$ a x–y geodesic in G. Lemma 3.3, gives a uniform affine upper bound on $|\bar\gamma|$ in terms of $d_\Sigma(x,y)$. Moreover, $\bar\gamma$ induces a x–y path p in H, and we have $|p| \leq C|\bar\gamma|$ for some constant C, because each unit of length of $\bar\gamma$ meets a uniformly bounded number of cells V_i by the uniformity of X. Thus $d_H(x,y) \leq |p| \leq C|\bar\gamma|$ has a uniform affine upper bound in terms of $d_\Sigma(x,y)$. The claim follows. □

The *essential degree* of a vertex x of a graph G is the number of $y \in V(G)$ such that G contains an x–y edge (this is smaller than the degree of x if G has several x–y edges). The following refines Theorem 1.2.

Corollary 4.2. *The following statements are equivalent for a complete Riemannian surface* (Σ, d_Σ)*:*

(1) Σ *is quasi-isometric to a graph of bounded degree;*
(2) Σ *has a uniform net;*
(3) Σ *has a 3-gonal decomposition of bounded essential degree whose* 1*-skeleton is with the simplicial metric is quasi-isometric to* Σ*.*

Proof. The equivalence of (1) and (2) is well-known, and holds in the greater generality where Σ is a geodesic metric space [9]. It is easy to deduce (1) from (3) by removing loops and all but one x–y edges for each pair of adjacent vertices x, y. The implication (2) $\implies$ (3) is established by Theorem 4.1: the fact that $\mathcal{D}$ is of bounded essential degree follows from the fact that all neighbours of $x \in X = V(\mathcal{D})$ in the 1-skeleton of $\mathcal{D}$ are within bounded distance from x, and so their cardinality is bounded by the fact that X is uniform. □

Using Remark 2.6, we can modify the 3-gonal decomposition of Corollary 4.2(3) into a triangulation by subdividing edges and adding vertices as needed. This suffices to complete the proof of Theorem 1.2. (Note however that the essential degree may become unbounded in this process.)

5. Quasi-isometric planar graphs that are not bi-Lipschitz equivalent

In this section we will combine Theorem 4.1 with a construction of Burago & Kleiner [3] in order to prove Corollary 1.3, which we restate for convenience:

Corollary 5.1. *There are plane graphs* H_1, H_2*, with bounded degrees and face-boundary sizes, which are quasi-isometric to each other but not bi-Lipschitz equivalent.*

Proof. Burago & Kleiner [3] constructed a net X_1 in $\mathbb{R}^2$ which is not bi-Lipschitz equivalent to the 'integer' net $X_2 := \mathbb{Z}^2$. It is not hard to see that every net in $\mathbb{R}^2$ is uniform by a volume argument. Thus, we can apply Theorem 4.1 to obtain 3-gonal decompositions $\mathcal{D}_1, \mathcal{D}_2$ of $\mathbb{R}^2$ with vertex sets X_1, X_2 and 1-skeletons H_1, H_2 quasi-isometric to $\mathbb{R}^2$ with their simplicial graph metric. (For H_2 we can just use the standard Cayley graph of $\mathbb{Z}^2$.) Thus, H_1, H_2 are quasi-isometric to each other. Easily, they have bounded face-boundary sizes, and bounded degrees since X_i is uniform (Remark 3.6).

By (4.3), the quasi-isometry to $\mathbb{R}^2$ is defined by the identity $(X_i, d_{H_i}) \to (X_i, d_{\mathbb{R}^2})$. Since this is a bijection, we deduce that (X_i, d_{H_i}) is bi-Lipschitz equivalent to $(X_i, d_{\mathbb{R}^2})$ for $i = 1, 2$. Thus if H_1, H_2 are bi-Lipschitz equivalent, then so are $(X_1, d_{\mathbb{R}^2}), (X_2, d_{\mathbb{R}^2})$, a contradiction. □

References

[1] Lars V. Ahlfors and Leo Sario, *Riemann Surfaces*, Princeton University Press, 1960.

[2] M. Bonamy, N. Bousquet, L. Esperet, C. Groenland, C.-H. Liu, F. Pirot, and A. Scott, *Asymptotic Dimension of Minor-Closed Families and Assouad-Nagata Dimension of Surfaces*, J. Eur. Math. Soc. **26** (2023), no. 10, 3739–3791.

[3] D. Burago and B. Kleiner, *Separated nets in Euclidean space and Jacobians of bi-Lipschitz maps*, GAFA **8** (1998), 273–282.

[4] J. Davies, *String graphs are quasi-isometric to planar graphs*. arXiv:2510.19602.

[5] Manfredo Perdigao do Carmo, *Differential geometry of curves and surfaces*, Prentice-Hall, Inc, 1976.

[6] A. Georgakopoulos and P. Papasoglu, *Graph minors and metric spaces*, Combinatorica **45** (2025), 33.

[7] Victor Guillemin and Alan Pollack, *Differential topology*, Vol. 370, American Mathematical Society, 2025.

[8] Allen Hatcher, *Algebraic Topology*, Cambrigde Univ. Press, 2002.

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

[10] S. Maillot, *Quasi-isometries of groups, graphs and surfaces*, Commentarii Mathematici Helvetici **76** (2001), no. 1, 29–60.

[11] T. Nguyen, A. Scott, and P. Seymour, *Asymptotic structure. II. Path-width and additive quasi-isometry.* ArXiv:2509.09031.

(Agelos Georgakopoulos) Mathematics Institute, University of Warwick, CV4 7AL, UK

(Federico Vigolo) Mathematisches Institut, Georg-August-Universität Göttingen, Bunsenstr. 3-5, 37073 Göttingen, Germany.

Email address: `federico.vigolo@uni-goettingen.de`