Isometric PL Embedding of Surfaces

Francis Lazarus

2020

Contents

The purpose of this lecture is to extend the Nash-Kuiper Theorem on C^1 isometric embeddings of Riemannian surfaces to polyhedral surfaces. These notes are based on the work of Burago and Zalgaller [BZ95]. As usual, $\mathbb{E}^d$ denotes the d-dimensional Euclidean space.

1 Polyhedral surfaces

Here, the objects of interest are **polyhedral surfaces** which are compact topological surfaces endowed with a polyhedral metric. Those can be obtained by considering a set of Euclidean triangles in the plane, gluing their sides according to a partial oriented pairing. This pairing should be such that each side appears at most once in the pairs and two sides in a pair should have the same length. The pair orientation specifies one of the two isometries between its sides. Note that two sides of a same triangle may well be glued together. The resulting surface is *closed*, i.e., without boundary, when each side appears in one pair, i.e., when the pairing is complete.

Exercise 1.1. Prove that the above construction always results in a topological surface.

Recall that a **simplicial triangulation** of a surface is a decomposition into triangles[1] such that any two (closed) triangles are either disjoint or intersect along a common vertex or a common edge.

[1] In general, one call a simplicial triangulation the homeomorphic image of the carrier of an abstract simplicial complex. A well-known result of Radó [Rad25, DM68] states that every topological surface has a simplicial triangulation. In these notes we only consider *geometric* triangulations composed of Euclidean triangles.

Exercise 1.2. The gluing of triangles may not define a simplicial triangulation of the resulting surface, for instance when two edges of a same triangle are paired. Assuming a pairing (each edge belongs to at most one pair) that excludes this case, do you always get a simplicial triangulation? Show that any gluing of triangles admits a simplicial subdivision, *i.e.*, that the triangles can be subdivided into a finite union of triangles in order to get a simplicial triangulation.

The gluing of Euclidean triangles induces an **intrinsic metric** on the resulting polyhedral surface: the distance between any two points is the infimum of the lengths of the paths connecting the two points, where paths are finite concatenations of paths contained in a single triangle and the length of a path is the sum of the Euclidean length of these triangle paths.

Exercise 1.3. Prove that the intrinsic metric is indeed a metric.

There is an intrinsic definition of polyhedral surfaces that does not assume any specific triangulation. Formally, a **polyhedral metric** on a surface is a metric such that every point has a neighborhood isometric to a neighborhood of the apex of a Euclidean cone, where we ask that the isometry sends the considered point to the apex of the cone. In turn, a (2-dimensional) **Euclidean cone** is defined by coning a rectifiable simple (non self-intersecting) curve on the unit sphere in $\mathbb{E}^3$ from the origin. The length of this curve is the **total angle** of the cone; it determines the geometry of the cone up to a length preserving map. A point whose conic neighborhood has total angle different from 2π is called a **singular vertex**. Note that in any triangulation of a polyhedral surface by Euclidean triangles the singular vertices must be vertices of the triangles.

Exercise 1.4. Show that the above definitions based on triangles or on conic neighborhoods are indeed equivalent. See [LP15] for a generalisation of this equivalence to higher dimensional polyhedral spaces.

Let S be a polyhedral surface. A map $f : S \to \mathbb{E}^3$ is said **piecewise linear** (PL) if S admits a triangulation such that the restriction of f to any triangle is *linear*, i.e., it preserves barycentric coordinates. f is **piecewise distance preserving** if S admits a triangulation such that the restriction of f to any triangle is distance preserving, i.e., $|f(x) - f(y)| = d_S(x, y)$ for any x, y in a same triangle. Here, $|\cdot|$ is the Euclidean norm and d_S is the metric on S. In particular, f must be PL.

2 The PL isometric embedding theorem of Burago and Zalgaller

A map $f : S \to \mathbb{E}^3$ is C-Lipschitz if $|f(x) - f(y)| \leq C d_S(x, y)$ for all $x, y \in S$. A C-Lipschitz map is said **contracting**, or **short** when $C < 1$, and **nonexpanding** when $C = 1$.

As a topological surface, a polyhedral surface admits a unique smooth structure compatible with the conic charts at the non-singular points (the local isometries are used as coordinate maps). We can thus speak of a C^2-immersion of S. Burago and Zalgaller [BZ95] proved a PL version of the Nash-Kuiper theorem on C^1 isometric immersions. We recall that $f : X \to Y$ is a (topological) **embedding** if $f : X \to f(X)$ is a homeomorphism, where $f(X) \subset Y$ is given the topology induced by Y. f is an **immersion** if it is a local embedding, i.e., every $x \in X$ has a neighborhood the restriction to which f is an embedding. Note that an immersion may have "self-intersections" as opposed to an embedding. A piecewise distance preserving embedding is also called a **PL isometric embedding**.

Theorem 2.1 (Burago and Zalgaller, 1996)**.** *Let S be a polyhedral surface. Every short C^2-immersion of S in $\mathbb{E}^3$ can be approximated by a piecewise distance preserving immersion in $\mathbb{E}^3$. The same is true, replacing immersion by embedding.*

Here, the approximation by a piecewise distance preserving map means that for any $\varepsilon > 0$ there is such a map whose C^0 distance is less than ε. We recall that the C^0 distance of two maps $f, g : S \to \mathbb{E}^3$ is $\sup_{s \in S} |f(s) - g(s)|$.

Remark 2.2. This theorem implies that every polyhedral surface has a piecewise distance preserving immersion in 3-space. In fact, every orientable surface and every surface with non-empty boundary is isometric to a PL surface embedded in $\mathbb{E}^3$! Indeed, it is well-known that every (compact) closed non-orientable surface can be smoothly immersed in 3-space while all other surfaces embeds smoothly in 3-space. One can compose such an immersion or embedding with a homothety whose ratio is small enough to get a short map. Applying the above theorem to this map allows to conclude.

Remark 2.3. The approximation result in the theorem tells that we can approximately prescribe the shape of the immersion as long as it is short. For instance, we can find a PL isometric embedding of a unit cube as close as desired to a cube of half size. An even more surprising consequence is that the unit cube – and in fact any polyhedron in $\mathbb{E}^3$ – has another PL isometric embedding enclosing a larger volume! See [Pak06] for the general case. The case of a cube has actually a simple solution [Pak08] independent of the theorem of Burago and Zalgaller.

3 The basic case

Before dealing with general polyhedral surfaces we consider the simplest case of a surface with boundary reduced to a single triangle T and embedded into $\mathbb{E}^3$ by a linear short map $T \to t$. In other words, we ask that

- (1) the sides of the image triangle t are shorter than the corresponding ones in T.

We also assume that

- (2) T and t are **acute triangles**, meaning that the angle at each vertex is less than the right angle. Equivalently, the circumcenter of each triangle is interior to the triangle.

- (3) The distance of the circumcenter Ω to each side of T is larger than the corresponding distance in t, *i.e.*, than the distance of its circumcenter ω to the corresponding side.

Consider one of the two right prisms with base t in $\mathbb{E}^3$. Let us call it the prism *above* t. Let PQ be a side of T and let pq be the corresponding side in t. Embed PQ isometrically as an equilateral broken line pmq inside the lateral face of the prism above pq and embed the two other sides of T in a similar manner in the corresponding lateral faces. See the figure below.

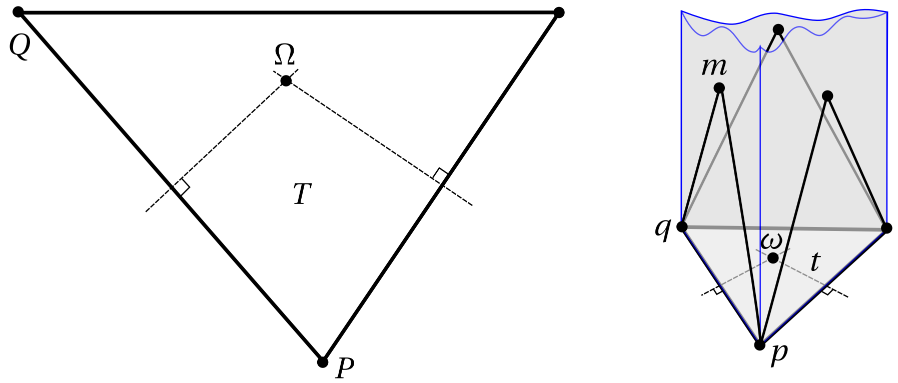

Lemma 3.1. *The above embedding of the sides of T extends to a PL isometric embedding of T lying inside the prism above t. Moreover, refining this isometric embedding we can enforce that its C^0 distance to the linear embedding $T \to t$ is arbitrarily small.*

Proof. Let ω' the point vertically above ω such that $|p\omega'| = |P\Omega|$. Refer to Figure 2 for an illustration. Note that ω' is well-defined since by the assumptions (1) and (2) the circumradius $|P\Omega|$ of T is larger than the circumradius $|p\omega|$ of t. Subdivide T into three subtriangles by cutting along the circumradii joining Ω to the vertices of T.

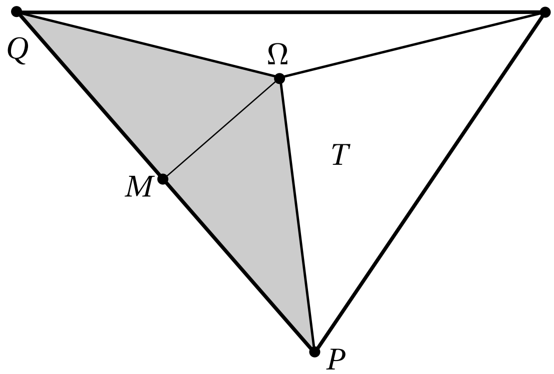

We show below how to fold the subtriangle $[PQ\Omega]$ in the prism above t so that its boundary fits the broken line $pmq\omega'p$. Similar constructions apply to the other two subtriangles so that putting together the three constructions we obtain the desired embedding of T. The second part of the lemma will follow after subdividing both T and t uniformly into sufficiently small triangles as on Figure 1. We can indeed apply the same three constructions to each pair of corresponding small triangles. The deviation of the whole construction from the base triangle t can be made arbitrarily small by using finer and finer subdivisions.

Let M be the midpoint of PQ. Fold $[PQ\Omega]$ along its height $M\Omega$ so as to apply its side PQ along pmq. The folding segment $M\Omega$ now coincides with a horizontal segment $m\omega''$. See Figure 2.

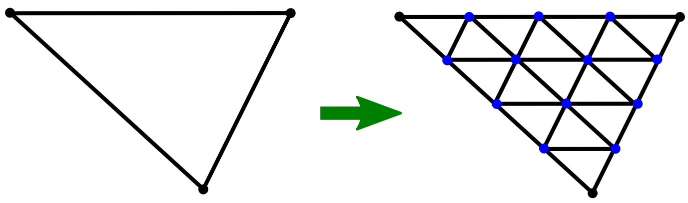

Figure 1: Uniform subdivision of a triangle. The vertices of the subdivision have barycentric coordinates $(i/n, j/n, k/n)$ for $i, j, k \in \mathbb{N}$ and $i + j + k = n$ for some fixed n.

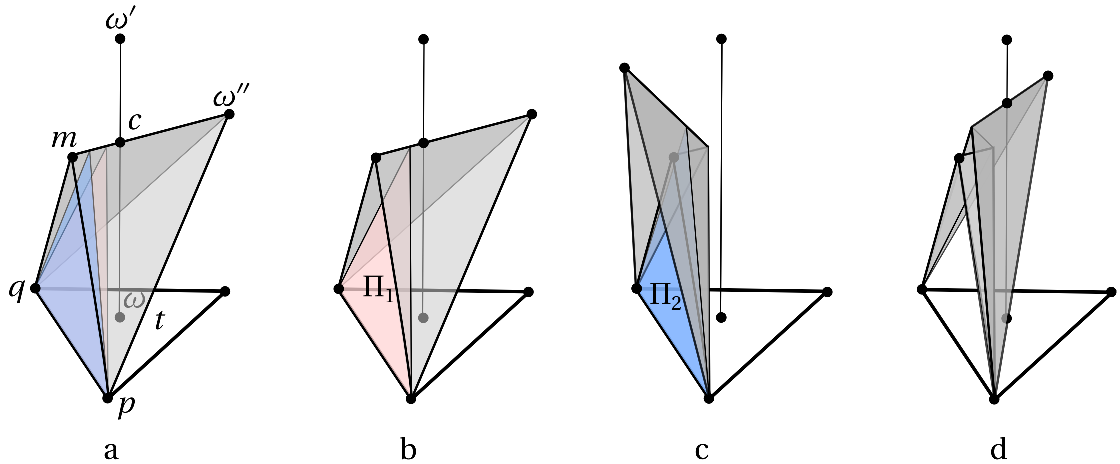

Figure 2: a: fold of $[PQ\Omega]$ along its height $M\Omega$ and the two planes Π_1 and Π_2. b: the plane Π_1. c: the plane Π_2 and reflection in Π_1. d: reflection in Π_2.

By the above assumption (3) the segment $m\omega''$ cuts $\omega\omega'$ at some point c. Let Π_1 and Π_2 be two planes in the sheaf generated by the line pq so that they cut mc in the order $m, \Pi_2 \cap mc, \Pi_1 \cap mc, c$. We further fold $[PQ\Omega]$ by reflecting across Π_1 its part lying below Π_1 (i.e., in the halfspace bounded by Π_1 and containing ω''). We then reflect along Π_2 the part above Π_2 and continue this way, alternating between reflections across Π_1 and Π_2. By a suitable choice of Π_1 and Π_2 we can ensure that after a finite and even number of reflections ω'' is sent to ω'. The resulting folding of $[PQ\Omega]$ defines an isometric embedding bounded by $pmq\omega'p$ as desired. □

Exercise 3.2. Propose an explicit folding of $[PQ\Omega]$ as in the above proof, possibly varying the reflection planes Π_1 and Π_2. Can you estimate the required number of reflections in your algorithm as a function of some appropriate geometric quantities?

This construction allows for some flexibility. The prism wall above pq may be slightly tilted around pq as long as $m\omega''$ crosses $\omega\omega'$. The same is true for the other two walls. The maximum angle of rotation depends on the maximum length ratio between the edges of T and the corresponding edges of t, and on the minimum and maximum angles at the vertices of T. It also depends on the degree of similarity between T and t, as the above condition (3) is trivially satisfied when T and t are similar. Playing with those parameters one can get uniform conditions for applying this construction to a collection of pairs (T_i, t_i).

4 Proof of the Burago and Zalgaller theorem

We shall assume once for all that the surface S in Theorem 2.1 is orientable and without boundary. The non-orientable or boundary cases need non-trivial special treatments and we refer to the original paper [BZ95] for the details. Denote by $f : S \to \mathbb{E}^3$ the short C^2 map in Theorem 2.1. The strategy for the proof is the following. In view of the construction in the previous section, suppose that S is triangulated so that each triangle is acute. By applying a uniform subdivision as on Figure 1 we can assume that the largest edge length of this triangulation, call it $\mathscr{T}$, is as small as desired. Consider the **PL approximation F of f with respect to $\mathscr{T}$** mapping a triangle $T = [PQR]$ of $\mathscr{T}$ to the triangle $F(T) := [f(P)f(Q)f(R)]$ in $\mathbb{E}^3$.

- As f is short, if T is small enough then the pair $(T, F(T))$ satisfies Condition (1) in Section 3.
- Since f is C^2 and S is compact, adjacent triangles are mapped to triangles having a dihedral angle uniformly close to π.

Suppose in addition that

- every triangle of $\mathscr{T}$ is acute.
- S has no singular vertex, so that its polyhedral metric, say μ, is flat and C^∞.
- f is almost conformal, meaning that μ and the pullback metric $f^*\langle\cdot,\cdot\rangle_{\mathbb{E}^3}$ are almost proportional at every point. In other words, for any point $s \in S$ and every tangent vectors u, v at s we have $\mu(u, v) \approx \lambda_s^2 \langle \mathrm{d}f_s.u, \mathrm{d}f_s.v\rangle_{\mathbb{E}^3}$ for some **conformal factor** $\lambda_s > 0$ independent of u, v.

Then every small enough triangle T of $\mathscr{T}$ is approximately similar to its linear image $F(T)$. We are thus in the uniform conditions evoked at the end of Section 3 and we can apply the tilted isometric embedding construction to each triangle T above $F(T)$ as described there. Since S is orientable we can orient all its triangles consistently so that the tilted embedding of an edge coincides for its two incident triangles. The individual triangle embeddings thus fit together to form a PL isometric immersion. This would conclude the proof of Theorem 2.1 noting that when f is an embedding, a sufficiently fine subdivision of an acute triangulation of S ensures that the embeddings of the individual triangles do not intersect, leading to a PL embedding as desired. The difficulty of the proof of Theorem 2.1 thus resides in removing the above assumptions:

- proving that any polyhedral surface has an acute triangulation,
- dealing with singular vertices on a polyhedral surface, and
- replacing f by an almost conformal map.

Exercise 4.1. Prove by simple counting arguments, without the help of the Gauss–Bonnet theorem, that a closed orientable polyhedral surface without singular vertices is a flat torus.

The fact that any polyhedral surface has an acute triangulation is of independent interest and is the subject of the next section. Concerning the conformality of f, we can invoke the Nash-Kuiper Theorem, or more precisely a simpler construction of Kuiper [Kui55]. We refer the reader to Kuiper's original paper (eq. (5.3)) or to the course on the h-principle in the Master program for a proof of the next result.

Theorem 4.2 (Kuiper'55)**.** *Any short C^1 immersion (embedding) $f:(S,\mu)\to\mathbb{E}^3$ of a surface S, possibly with boundary, endowed with a C^1 metric μ can be approximated by a C^∞ almost isometric immersion (embedding) $g:S\to\mathbb{E}^3$, i.e., satisfying $(1-\varepsilon)\mu < g^*\langle\cdot,\cdot\rangle_{\mathbb{E}^3} < \mu$ with ε arbitrarily small. Moreover, if f is isometric on the boundary of S (and short inside S), we can enforce $g=f$ on the boundary.*

Thanks to this lemma we can approximate f with an almost isometric immersion[2] g which is *a fortiori* almost conformal. Moreover, replacing μ by $\alpha\mu$, with $\alpha<1$, so that f is still short for $\alpha\mu$ we ensure that g is short for μ. It remains to deal with singular vertices. The singular vertices with total angle smaller or larger than 2π are dealt with separately. We first introduce certain maps between cones.

The standard conformal map. Let C_φ denote the Euclidean cone with total angle φ. Fixing a generating line ℓ on C_φ we get polar coordinates (r,θ) for a point at distance $r>0$ from the apex, such that the generating line through the point makes an angle $\theta\in[0,\varphi)$ with ℓ. The **standard conformal map** $f_{\varphi,\psi,\lambda}:C_\varphi\to C_\psi$ sends apex to apex and the point with polar coordinates (r,θ) to the point with polar coordinates $(\lambda r^{\frac{\psi}{\varphi}},\frac{\psi}{\varphi}\theta)$, where $\lambda>0$ is a fixed parameter. This map is conformal (apart from the apex) with conformal factor $\lambda\frac{\psi}{\varphi}r^{\frac{\psi}{\varphi}-1}$.

Exercise 4.3. Prove that $f_{\varphi,\psi,\lambda}$ is indeed conformal with the claimed conformal factor.

Solution:
In local charts, the map is just $f:z\mapsto\lambda z^{\frac{\psi}{\varphi}}$ which is holomorphic, hence conformal with factor $|f'|=\lambda\frac{\psi}{\varphi}r^{\frac{\psi}{\varphi}-1}$.
Other longer proof: The 1-forms $\mathrm{d}x,\mathrm{d}y$ are related to the 1-forms $\mathrm{d}r,\mathrm{d}\theta$ by differentiating $x=r\cos\theta$ and $y=r\sin\theta$. Those forms are dual to the bases $(\partial x,\partial y)$ and $(\partial r,\partial\theta)$ respectively. We have the metric at the source: $g:=\mathrm{d}x^2+\mathrm{d}y^2=(c\,\mathrm{d}r-rs\,\mathrm{d}\theta)^2+(s\,\mathrm{d}r-rc\,\mathrm{d}\theta)^2=\mathrm{d}r^2+r^2\mathrm{d}\theta^2$ where $c=\cos\theta$ and $s=\sin\theta$. At the target we have the metric: $G:=\mathrm{d}X^2+\mathrm{d}Y^2=\mathrm{d}R^2+R^2\mathrm{d}\Theta^2$. For a general metric m, we have

$$m=m(\partial r,\partial r)\mathrm{d}r^2+m(\partial\theta,\partial\theta)\mathrm{d}\theta^2+2m(\partial r,\partial\theta)\mathrm{d}r\,\mathrm{d}\theta$$

Hence, writing $f=(R,\Theta)$ (with $R=\lambda r^{\frac{\psi}{\varphi}}$ and $\Theta=\frac{\psi}{\varphi}\theta$)) for $f_{\varphi,\psi,\lambda}$,

$$f^*M=G(\mathrm{d}f.\partial r,\mathrm{d}f.\partial r)\mathrm{d}r^2+G(\mathrm{d}f.\partial\theta,\mathrm{d}f.\partial\theta)\mathrm{d}\theta^2+G(\mathrm{d}f.\partial r,\mathrm{d}f.\partial\theta)\mathrm{d}r\,\mathrm{d}\theta$$

On the other hand we compute $\mathrm{d}f.\partial r=\frac{\partial R}{\partial r}\partial R+\frac{\partial\Theta}{\partial r}\partial\Theta=\lambda\frac{\psi}{\varphi}r^{\frac{\psi}{\varphi}-1}\partial R$ and $\mathrm{d}f.\partial\theta=\frac{\partial R}{\partial\theta}\partial R+\frac{\partial\Theta}{\partial\theta}\partial\Theta=\frac{\psi}{\varphi}\partial\Theta$. It ensues that

$$f^*M=(\lambda\frac{\psi}{\varphi}r^{\frac{\psi}{\varphi}-1})^2\mathrm{d}r^2+(\lambda r^{\frac{\psi}{\varphi}})^2(\frac{\psi}{\varphi})^2\mathrm{d}\theta^2=((\lambda\frac{\psi}{\varphi}r^{\frac{\psi}{\varphi}-1})^2g$$

[2] The C^1 (exact) isometric immersion in the Nash-Kuiper Theorem is obtained as the limit of a converging sequence of such approximations.

as desired.

Dealing with singular vertices of total angle smaller than 2π. Let $v \in S$ be a singular vertex with total angle $\varphi < 2\pi$. Let $B_{v,\rho}$ be the (conic) disk with center v and radius ρ in S. We modify f in $B_{v,\rho}$ for some small ρ so that its restriction to $B_{v,\rho'}$, for some $\rho' < \rho$, coincides with the standard conformal map $f_{\varphi,2\pi,\lambda}$ where the image cone $C_{2\pi}$ is identified with the plane tangent to $f(v)$ with "apex" $f(v)$ and λ is chosen small enough so that $f_{\varphi,2\pi,\lambda}$ is contracting in $B_{v,\rho'}$. We further extend $f_{\varphi,2\pi,\lambda}$ inside $B_{v,\rho}$ so that the overall modification of f remains short and C^2.

Dealing with singular vertices of total angle larger than 2π. Let $v \in S$ be a singular vertex with total angle $\varphi > 2\pi$. We modify f in $B_{v,\rho}$ for some small ρ so that for some $\rho' < \rho$:

1. its restriction to $B_{v,\rho'/2}$ expressed in polar coordinates is the map $(r,\theta) \to (r, \frac{2\pi}{\varphi}\theta)$ where we again identify the flat cone $C_{2\pi}$ with the plane tangent to $f(v)$. This map is isometric in the radial direction and uniformly contracting in the orthogonal direction.

2. its restriction to the annulus $B_{v,\rho'} \setminus B_{v,\rho'/2}$ is the standard conformal map $f_{\varphi,2\pi,\lambda}$ with $\lambda = (\rho'/2)^{1-\frac{2\pi}{\varphi}}$. This choice of λ implies that the conformal factor of $f_{\varphi,2\pi,\lambda}$ is bounded by $\frac{2\pi}{\varphi} < 1$ outside $B_{v,\rho'/2}$.

3. its restriction to $B_{v,\rho} \setminus B_{v,\rho'}$ is smooth, short, and connects to f at the boundary of $B_{v,\rho}$ in a C^2 manner.

Note that the modified f is not short on the disk $B_{v,\rho'/2}$ and is only C^1 at its boundary. We surround v in S with a regular k-gone $N_v(k)$ inscribed in a disk of radius $\rho'/2$, where k is large and may be fixed later. We triangulate $N_v(k)$ by coning its boundary from its center v. We next slightly enlarge $N_v(k)$ to a neighborhood $N'_v = N'_v(k)$ to form a cogged disk obtained by attaching equilateral triangles to the k sides of $N_v(k)$. The reason for this enlargement is to allow for the uniform subdivision of the complement of N'_v. Indeed, this complement needs to be triangulated and possibly subdivided uniformly, say ℓ times. This subdivision can easily be extended to N'_v by changing $N_v(k)$ for $N_v(\ell k)$, as shown on the figure below.

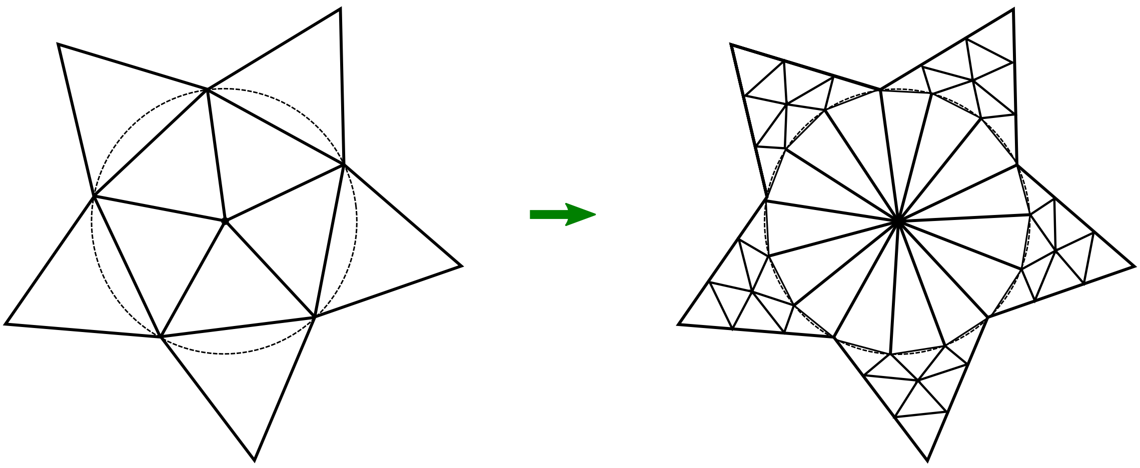

Replacing $N_v(5)$ by $N_v(15)$ allows to extend the uniform subdivision of the boundary of the cogged disk $N'_v(5)$.

Putting the pieces together. In the above local modifications of f, the radii ρ are chosen small enough so that the disks $B_{v,\rho}$ are pairwise disjoint and the modified map, say f_1, remains close to f. Let V_+, V_- be the set of singular vertices of S with total angle respectively smaller and larger than 2π. Set $V = V_+ \cup V_-$ for the set of singular vertices of S. We shall now invoke Theorem 4.2 to replace f_1 on $S \setminus \cup_{v\in V} B_{v,\rho'/2}$ by a close immersion f_2 which is both almost conformal and short with respect to the polyhedral metric μ. To this end, we first consider outside the disks $B_{v,\rho'/2}$ a contracting scaling $\alpha\mu$ of μ, $\alpha < 1$ so that f_1 is still short for $\alpha\mu$. We next consider the metric μ' on $S \setminus \cup_{v\in V} B_{v,\rho'/2}$ defined by

$$\mu' = \begin{cases} \alpha\mu & \text{outside the disks } B_{v,\rho'} \\ \varpi f_1^*\langle\cdot,\cdot\rangle_{\mathbb{E}^3} + (1-\varpi)\alpha\mu & \text{on } \cup_{v\in V} B_{v,\rho'} \setminus B_{v,\rho'/2} \end{cases}$$

where ϖ is a smooth plateau function interpolating between 1 on the boundary of $B_{v,\rho'/2}$ and 0 on the boundary of $B_{v,\rho'}$. Note that f_1 is already conformal with respect to μ on $B_{v,\rho'} \setminus B_{v,\rho'/2}$ so that μ and μ' are conformal. Note also that f_1 is short with respect to $\alpha\mu$ on $B_{v,\rho'} \setminus B_{v,\rho'/2}$ so that f_1 remains short with respect to μ' in the interior of $S \setminus \cup_{v\in V} B_{v,\rho'/2}$ while being isometric on its boundary. We can now apply Theorem 4.2 to μ' and f_1 on $S \setminus \cup_{v\in V} B_{v,\rho'/2}$ to obtain an almost isometric immersion f_2 approximating f_1. In other words, $f_2^*\langle\cdot,\cdot\rangle_{\mathbb{E}^3} \approx \mu'$ and $f_2 \approx f_1$. Moreover f_2 and f_1 coincide on the boundary of $B_{v,\rho'/2}$. We extend f_2 to S by setting $f_2 = f_1$ on the disks $B_{v,\rho'/2}$. The map f_2 is C^2, short and almost conformal with respect to μ except on $\cup_{v\in V_-} B_{v,\rho'/2}$.

Next, we compute an acute triangulation $\mathscr{T}$ of $S \setminus \cup_{v\in V_-} N'_v$ as described in Section 5 so that $\mathscr{T}$ together with the triangulations of the N'_v define an acute triangulation of S. The triangles in $\mathscr{T}$ being in finite number admit a smaller and a larger angle. As noted at the end of Section 3 we can find uniform conditions on the degree of similarity and on the contraction factor that allow to apply the basic construction of Section 3. Recall that around each $v \in V_+$ the modified map f_2 is a standard conformal map. In particular its conformal factor tends to zero at v. Moreover, the default of conformality of f_2 outside the $B_{v,\rho'/2}$ can be quantified. Hence, we can subdivide $\mathscr{T}$ uniformly to get a sufficiently fine triangulation for which the PL approximation of f_2 with respect to $\mathscr{T}$ sends

- adjacent triangles to almost coplanar triangles, and
- each triangle in $S \setminus \cup_{v\in V_-} N'_v$ to a triangle that is either sufficiently similar or sufficiently smaller so that the basic construction of Section 3 can be applied.

It remains to extend this subdivision to the neighborhoods N'_v as described above. Let $\mathscr{T}'$ be the resulting triangulation. We finally apply the basic construction of Section 3 to each triangle of $\mathscr{T}'$ except those in the neighborhoods of the form $N_v(\ell k) \subset N'_v$, for $v \in V_-$. In those neighborhoods we use a simpler construction. The ℓk long isosceles triangles inside each $N_v(\ell k)$ are further split along their longest median and linearly embedded into a radially crimped surface above the plane tangent to $f(v)$ as shown on the figure below.

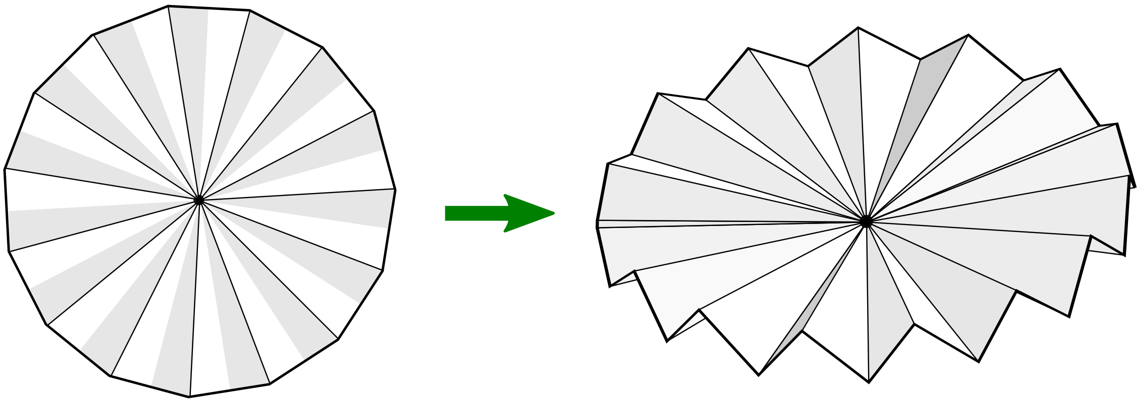

Beware that the left disk is not flat at its center!

If ℓk is large enough the boundary of $N_v(\ell k)$ is embedded almost perpendicularly to the tangent plane at $f(v)$ and can be glued with the rest of the construction. □

We end this section with a picture of a PL isometric embedding of the square flat torus approximating a short Hopf torus [Pin85]. The basic construction of Section 3 has been applied to each triangle of a PL approximation (Figure 5, left) of this Hopf (conformal) torus to obtain the PL isometric embedding of Figure 5, right.

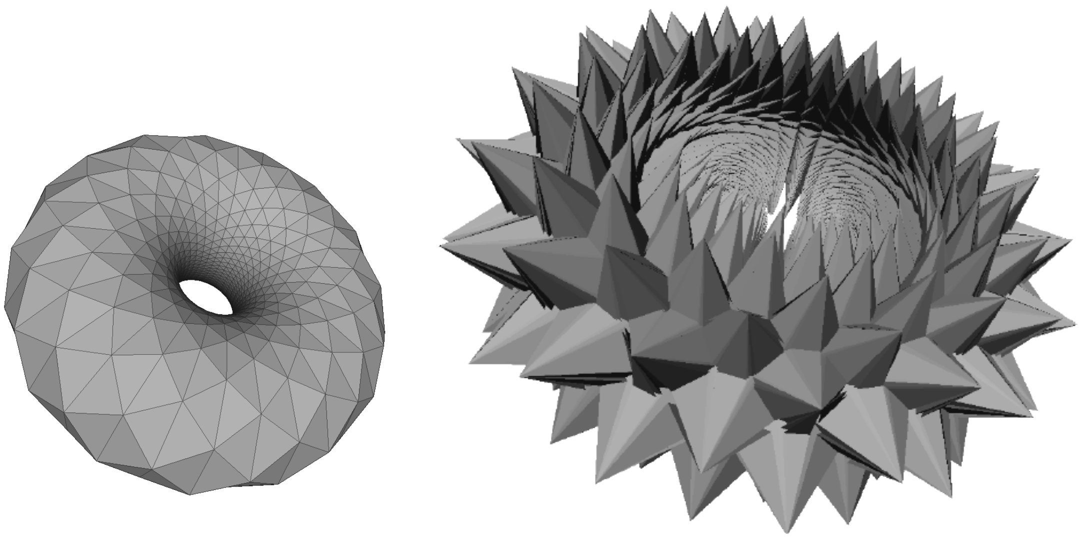

Figure 5: Left, a short PL embedding of the square fat torus. Right, The resulting PL isometric embedding of the square flat torus computed by Florent Tallerie. The triangulation is composed of 170,040 triangles.

5 Existence of acute triangulations

An **acute triangulation** of a polyhedral surface S is a simplicial triangulation such that every triangle is flat and acute in S. In particular, if S is already triangulated it might be desirable to subdivide this triangulation into an acute one. The existence of such acute triangulations and refinements has a long history, starting with Burago and Zalgaller in 1960. See [Zam13] for a comprehensive account on the subject. The existence proof of Burago and Zalgaller (only available in Russian) was recently simplified by Saraf [Sar09] and by Maheara [Mae11]. Saraf constructs a **non-obtuse triangulation** where the angles within the triangles are at most $\pi/2$. A non-obtuse triangulation may thus contain right angle triangles.

Exercise 5.1. Check that a triangle can always be subdivided into at most two non-obtuse triangles and prove that it can always be subdivided into at most 7 acute triangles.

Theorem 5.2 (Saraf'09, Maheara'11)**.** *Every triangulation $\mathscr{T}$ of a polyhedral surface can be subdivided into a non-obtuse triangulation. Moreover, we can impose that the triangles of the subdivision with at least one vertex which is a vertex of $\mathscr{T}$ or interior to an edge of $\mathscr{T}$ are acute.*

PROOF (SKETCH). Exercise 5.1 provides a seemingly short proof by subdividing each triangle into 7 acute triangles. However, the subdivisions of an edge induced by the subdivision of the two adjacent triangles have no reason to agree so that the resulting subdivision might not be simplicial. We thus need a more clever construction.

Let $\mathscr{T}$ be a triangulation of a polyhedral surface S. The main argument for the construction of an acute triangulation is to first cover the edges of $\mathscr{T}$ with a set of non-overlapping disks centered along the edges. Then, inside each triangle t, the disks covering its edges are completed into a packing, i.e. into a set of touching disks with disjoint interiors. Connecting the centers of touching disks with line segments we obtain a **contact graph** that induces a subdivision of t into polygons. See Figure 6. The packing can be chosen so that every polygon has at most four sides. Moreover, it is possible to subdivide such polygons into non-obtuse triangles subdividing each side in two by introducing the tangency point of the disks centered at its endpoints. In particular, the edges of t will be subdivided exactly at the center and contact points of the covering disks, thus matching the subdivision induced by the other adjacent triangle (for a boundary edge there is no matching to check). This provides the required non-obtuse triangulation of S as in the first part of the lemma.

In details, we let θ be the smallest angle in the triangles of $\mathscr{T}$, and we let h be the shortest altitude of any triangle of $\mathscr{T}$. Put $r = \frac{h}{9}\sin\frac{\theta}{2}$. Consider an edge e of length ℓ in $\mathscr{T}$. We place two disks of radius $R := h/3$ centered at the endpoints of e and cover the remaining middle segment with $k_e := \lceil(\ell/2 - R)/r\rceil$ equally spaced disks of radius $r_e := (\ell/2 - R)/k_e$. The disks placed at the vertices are said of *vertex type*, and the other disks are said of *edge type*. See Figure 6. We easily check that $3r/5 \le r_e \le r$. We cover similarly all the edges of $\mathscr{T}$. It is easily checked that the disks of type vertex and edge have pairwise disjoint interior. Consider a triangle t. The disks covering its edges form a circular row of packed disks. We partially extend this packing with a second row composed of three types of disks.

1. for every pair of touching edge-disks we place a disk of the same radius tangent to the two edge-disks.

2. for every pair of touching disks, one of which a vertex-disk and the other one an edge-disk, we place a disk tangent to both disks in the pair and to a disk of

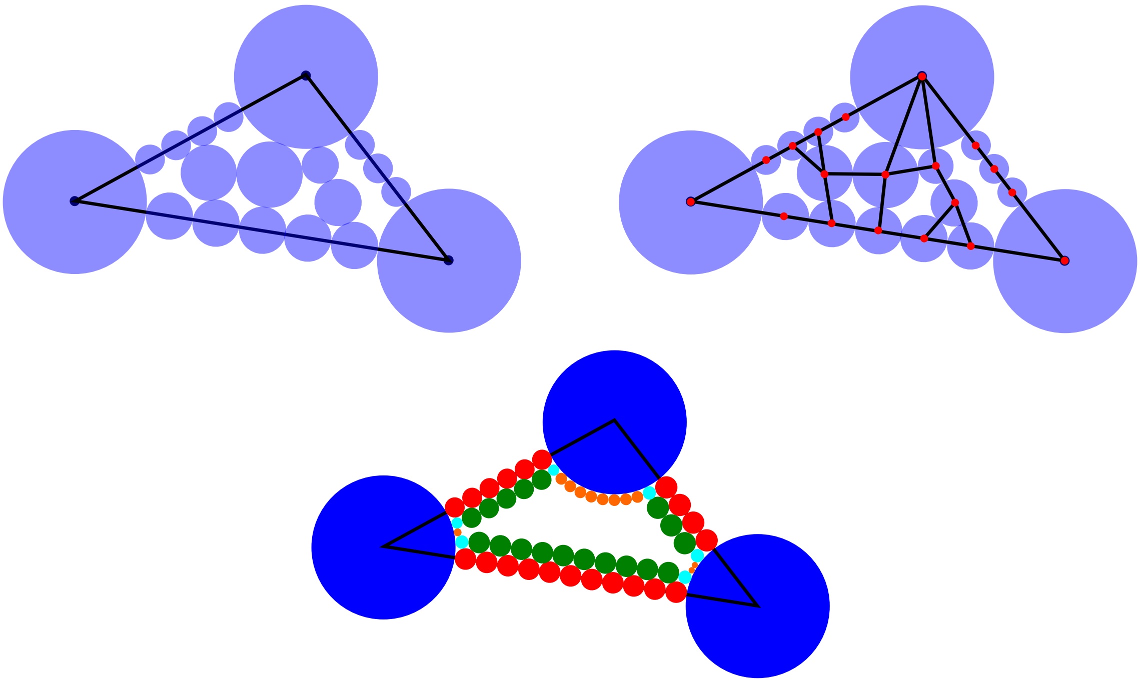

Figure 6: Upper left: a packing of disks covering the edges. Upper right: the corresponding contact graph. Lower middle: the vertex-disks (dark blue), the edge-disks (red), and the three types of disks in the circular row: 1 (green), 2 (light blue) and 3 (orange).

type 1. The above choice of R and r is such that the two first types of disks now form three disjoint sequences of touching disks – one per edge of t.

3. we finally pack disks of radius at most r tangent to the vertex-disks in order to connect these three sequences into a single circular sequence of tangent disk. See Figure 6.

The reason for this second row of disks is to enforce that the faces of the contact graph incident to the edges of t are triangles. We now extend inside t the packing formed by this row and the disks covering the edges of t.

Claim 1. The second row of disks can be extended towards its interior to form a packing whose contact graph has faces (excluding the exterior one) with at most four sides.

The proof, due to Bern et al. [BMR95, lem. 1] is by induction on the number of sides of a face. Initially the contact graph has a single face corresponding to the second row of disks. Consider the *medial axis* of the collection of disks $\mathscr{D}$ defining a face. This is the set of centers of all inclusion-wise maximal disks contained in the piecewise circular polygon bounding $\mathscr{D}$. It is a finite connected graph comprising arcs of hyperbolas possibly degenerated into line segments[3] as illustrated on figure 7.
For two disks $B_{c,r}, B_{c',r'}$ their medial axis is the set of x s.t. $|cx| - r = |c'x| - r'$ or equivalently, $|cx| - |c'x| = r - r'$ which defines an hyperbola.
The graph has one leaf vertex per contact point of touching disks in $\mathscr{D}$. Every vertex

[3] This medial axis is in fact part of the 1-skeleton of the Apollonius diagram of $\mathscr{D}$, also called the additively weighted Voronoi diagram.

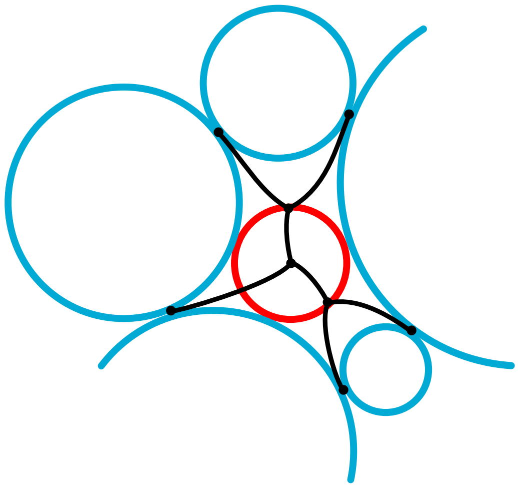

Figure 7: The medial axis of a circular sequence of 5 disks. Adding the middle disk (red) splits the face of the contact graph into smaller faces.

of the graph with degree d is the center of a maximal disk tangent do d disks in $\mathscr{D}$. If $|\mathscr{D}| > 4$, either the graph has a vertex of degree at least 4, or it contains vertices of degree 3 only and one of those is adjacent to two non-leaf vertices. In both cases adding the maximal disk centered at this vertex splits the contact graph into polygons of size less than $|\mathscr{D}|$. This ends the proof of the claim.

We now have a packing including the vertex and edge-disks, the above second row of disks and its extension. By construction, the faces of its contact graph incident to the edges of t are triangles and thanks to claim 1 we can extend the packing so that the remaining faces have at most four sides. It remains to prove that each of those faces, triangle or quadrilateral, can be subdivided into non-obtuse triangles so that the subdivisions agree on the face boundaries. More specifically, we show that a contact graph reduced to a triangle or a quadrilateral has a non-obtuse triangulation where the contact points of the disks defining the graph are the only vertices inserted along its edges (but the triangulation may contain other interior vertices). For the quadrilateral case one can obtain a triangulation into at most 56 non-obtuse triangles. The construction is rather tedious and described in [BMR95, lem. 4-7]. For the triangle case Maheara [Mae11] gives a subdivision into 10 *acute* triangles as shown below.

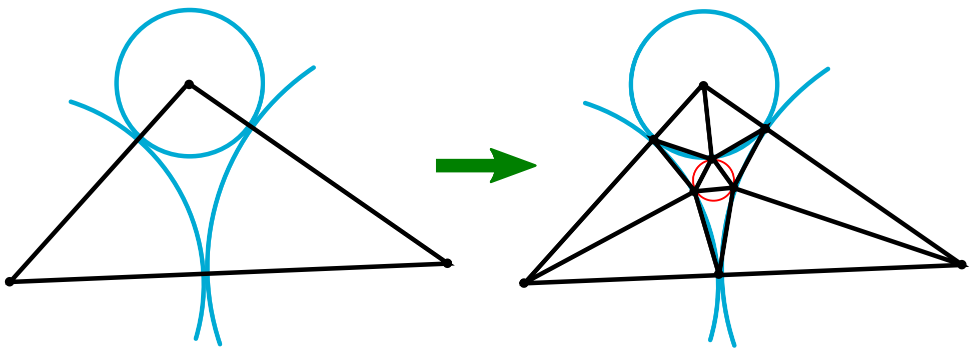

The fact that the faces of the contact graph incident to the edges of t are triangles thus implies that they are subdivided into acute triangles, whence the second part of the lemma. □

Maheara is able to bound the size of the non-obtuse triangulation by $3952\frac{\ell_{max} n}{h\theta}$ where ℓ_{max} is the maximum length of an edge of $\mathscr{T}$, n its number of triangles, and h, θ are defined as in the proof.

Corollary 5.3 (Maheara'11)**.** *Every triangulation of a polyhedral surface can be subdivided into an acute triangulation.*

Proof. Let $\mathscr{T}$ be a triangulation of a polyhedral surface S. From the preceding theorem there exists a subdivision $\mathscr{T}'$ into non-obtuse triangles such that the subdivision triangles incident to an edge of $\mathscr{T}$ are acute. We subdivide uniformly each triangle in $\mathscr{T}'$ by splitting every edge at its midpoint, connecting the three midpoints in each face. Each triangle in $\mathscr{T}'$ is thus subdivided into 4 similar triangles so that $\mathscr{T}'$ satisfies the properties in Theorem 5.2. Then, inside every right triangle of $\mathscr{T}'$ we flip the interior subdividing edge parallel to its hypotenuse.

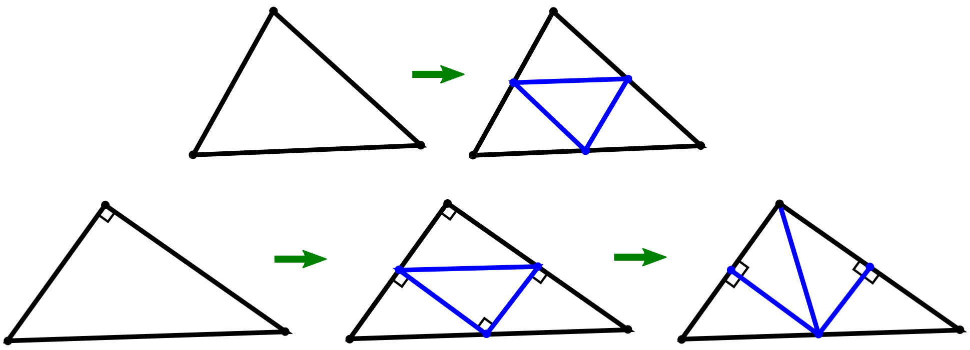

This replaces two right subtriangles by two other congruent right subtriangles. Let $\mathscr{T}''$ be the resulting triangulation. We also denote by $\mathscr{M}$ the set of midpoints introduced in $\mathscr{T}'$ (or equivalently in $\mathscr{T}''$) and by V'' the set of vertices of $\mathscr{T}''$.

The edge flipping operation implies that a vertex standing at the right corner of some right triangle must belong to $\mathscr{M}$. Remark that no such vertex is adjacent to a vertex of the original edges of $\mathscr{T}$ since all their incident triangles are acute. Consider a vertex $v \in V'' \setminus \mathscr{M}$. In particular, all its incident angles are acute. If v is incident to some right triangle replace the subdivision inside its star as described on the next figure.

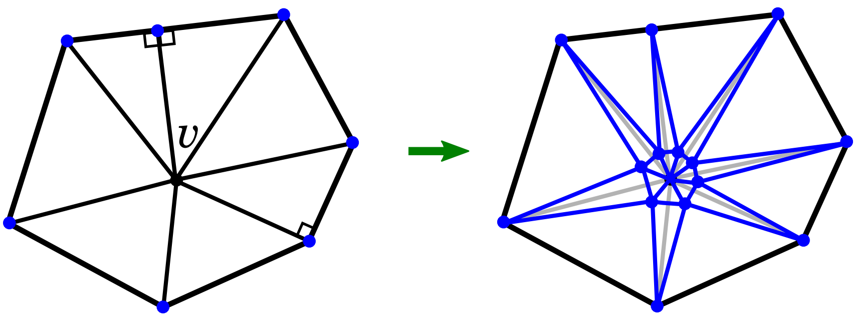

If the central "wheel" replacing v is small enough all the triangles in the new star subdivision will be acute. Since the vertices in $V'' \setminus \mathscr{M}$ are pairwise non-adjacent, their open stars are pairwise non-intersecting and we can perform a similar re-triangulation in every star independently. Note that these local modifications do not affect the edges of $\mathscr{T}''$ subdividing the original edges of $\mathscr{T}$ by the above remark. $\square$

6 Equidimensional piecewise distance preserving maps

For a polyhedral surface S, recall that $f : S \to \mathbb{E}^d$ is piecewise distance preserving if S admits a triangulation such that the restriction of f to any triangle is isometric. Theorem 2.1 of Burago and Zalgaller asserts the existence of piecewise distance preserving map when $d = 3$. Surprisingly, the result remains true for $d = 2$. I partly follow the notes of Petrunin and Yashinski [PY16].

Theorem 6.1 (Zalgaller)**.** *Every polyhedral surface S admits a piecewise distance preserving map into $\mathbb{E}^2$.*

PROOF. The proof is actually very simple once we know the existence of acute triangulations. By Corollary 5.3 we may assume that S comes equipped with an acute triangulation $\mathscr{T}$. Let V_0 be the set of vertices of $\mathscr{T}$. Subdivide each triangle t of $\mathscr{T}$ into 12 subtriangles as follows. In a first step split every edge at its midpoint and replace t by 6 triangles, starring its boundary at the circumcenter of t. Note that t being acute contains its circumcenter in its interior. Let V_1 be the set of vertices introduced in this step, comprising the edge midpoints and triangle circumcenters. Finally split each subtriangle along the angle bisector incident to its vertex in V_0, splitting the opposite edge accordingly. Denote by V_2 the set of vertices thus introduced and let $\mathscr{T}_2$ be the triangulation finally obtained. Hence, $|\mathscr{T}_2| = 12|\mathscr{T}|$.

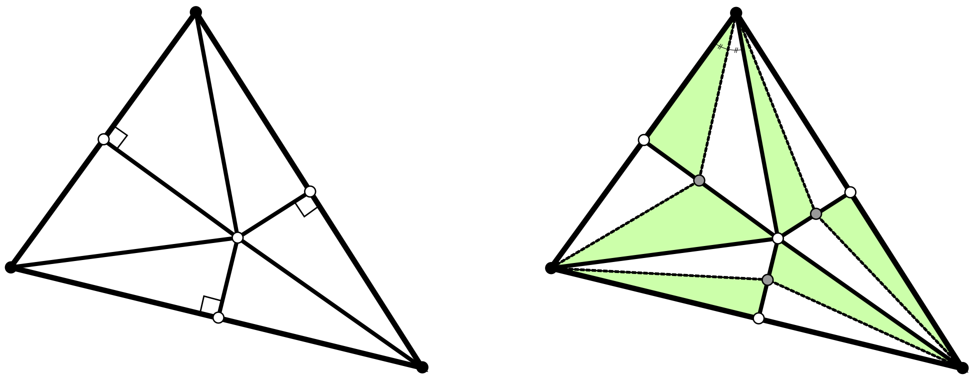

Left, first subdivision. Right, each subtriangle is further split along a bisector resulting in a triangulation $\mathscr{T}_2$. Every triangle of $\mathscr{T}_2$ has one (black) vertex in V_0, one (white) vertex in V_1 and one (grey) vertex in V_2.

We now define $f : S \to \mathbb{R}^2$, sending $\mathscr{T}_2$ linearly into $\mathbb{R}^2$ as follows. Let $[v_0, v_1, v_2] \in \mathscr{T}_2$ with $v_i \in V_i$, $i = 0, 1, 2$. Set $f(v_0) = (0,0) \in \mathbb{R}^2$ and $f(v_1) = f(v_0) + |v_0 v_1| e_1$, where (e_1, e_2) is the canonical basis of $\mathbb{R}^2$. Define $f(v_2)$ in the upper halfplane $\{x_2 > 0\}$ so that $[f(v_0), f(v_1), f(v_2)]$ is isometric to $[v_0, v_1, v_2]$. It is a simple matter to check that the image of a vertex is independent of the incident triangle chosen to define its image. The resulting linear extension f is clearly piecewise distance preserving. Note that in the above figure the restriction of f to green triangles is orientation preserving while its restriction to the white triangles is orientation reversing (or vice-versa). □

The preceding theorem has a stronger form which is the analog of the theorem of Burago and Zalgaller in dimension 2.

Theorem 6.2 (Akopyan, 2007)**.** *Let S be a polyhedral surface. Every nonexpanding PL map $S \to \mathbb{E}^2$ can be approximated by a piecewise distance preserving map, where the C^0 distance to the apprimation can be chosen arbitrarily small.*

The proof relies on an extension theorem of independent interest.

Theorem 6.3 (Brehm, 1981)**.** *Let $\{p_1, \ldots, p_n\}$ be a set of n points contained in a convex polygon P in the plane. Then, any nonexpanding map $\{p_1, \ldots, p_n\} \to \mathbb{E}^2$ extends to a piecewise distance preserving map $f : P \to \mathbb{E}^2$.*

PROOF. Denote by q_i the image of p_i by the nonexpanding map. The proof is by induction on the number n of points. The base case $n = 1$ is trivially solved by taking for f the plane translation of vector $p_1 q_1$. For $n > 1$ the induction hypothesis provides a piecewise distance preserving map $h : P \to \mathbb{E}^2$ such that $h(p_i) = q_i$ for $i = 2, \ldots, n$. We may assume $h(p_1) \neq q_1$ for otherwise we can set $f = h$. Consider the set

$$\Omega = \{x \in P \mid |p_1 x| < |q_1 h(x)|\}$$

Note that $p_1 \in \Omega$.

Claim. Ω is the interior (relative to P) of a star-shaped polygon with respect to p_1.

PROOF OF THE CLAIM. • Ω is star-shaped: if $x \in \Omega$ then for every y on the segment $[p_1, x]$ we have

$$|p_1 y| = |p_1 x| - |x y| < |q_1 h(x)| - |h(x) h(y)| \leq |q_1 h(y)|$$

Hence $y \in \Omega$ as desired. Here, we used the simple fact that the distance preserving map h is nonexpanding.

• Ω is the interior of a polygon: consider a triangulation $\mathcal{T}$ of P such that h is an isometry on each triangle t of $\mathcal{T}$. Denote by ι the extension of the isometry $h|_t$ to the plane. Then, the condition $|p_1 x| < |q_1 h(x)|$ can be written $|p_1 x| < |\iota^{-1}(q_1) x|$ on t. Hence, $\Omega \cap t$ is the intersection of t with the open halfplane containing p_1 and delimited by the bisector of the segment $[p_1, \iota^{-1}(q_1)]$. It follows that $\Omega = \cup_{t \in \mathcal{T}} \Omega \cap t$ has indeed a polygonal shape. □

Intersecting the boundary $\partial\Omega$ of Ω with $\mathcal{T}$, we may assume that h is an isometry on each segment of $\partial\Omega$. Let E be the set of segments of $\partial\Omega$ that are not contained in ∂P. For each segment $e \in E$ we have by continuity of h that $|p_1 x| = |q_1 h(x)|$ for $x \in e$. We now define f by parts as follows.

- We set $f = h$ on $P \setminus \Omega$.
- For $e \in E$ we define f on the triangle $p_1 * e$ (the cone with apex p_1 over e) as the isometry sending p_1 to q_1 and e to $h(e)$.

- The remaining part $\Omega \setminus \cup_{e\in E} p_1 * e$ is composed of disjoint open convex polygons with closure of the form $p_1 * C$ where C is a subpath of ∂P. Denote by a and b the endpoints of C. The partial definition of f already maps $[p_1, a]$ and $[p_1, b]$ to $[q_1, h(a)]$ and $[q_1, h(b)]$ respectively. Recalling that $|h(a)h(b)| \leq |ab|$, it is an exercise to extend f inside $p_1 * C$ in a piecewise distance preserving manner (hint: fold the polygon ap_1bC as a fan)

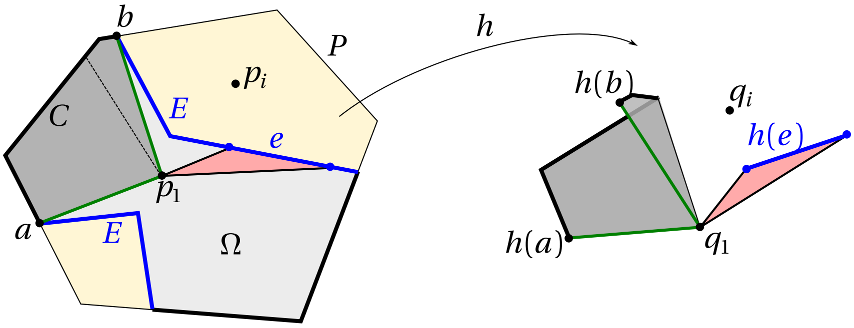

The map f thus defined is clearly continuous and piecewise distance preserving. Moreover, we have $f(p_1) = q_1$ and $p_i \in P \setminus \Omega$ for $i \geq 2$, so that $f(p_i) = h(p_i) = q_i$ and f is indeed an extension of $p_i \mapsto q_i$. $\square$

PROOF OF THEOREM 6.2. We first suppose that $h : S \to \mathbb{R}^2$ is a short PL map with Lipschitz constant $C < 1$. Let $\mathcal{T}$ be a triangulation such that h is linear on each triangle of $\mathcal{T}$. Denote by f the piecewise distance preserving map approximating h that we are looking for. We define f on the edges of $\mathcal{T}$. If e is such an edge we let $f(e)$ result from a corrugation process applied to $h(e)$: we simply replace the segment $h(e)$ by a polygonal curve with the same extremities and the same length as e but with a saw-tooth profile. The larger is the number of teeth the closer is f to h along e.

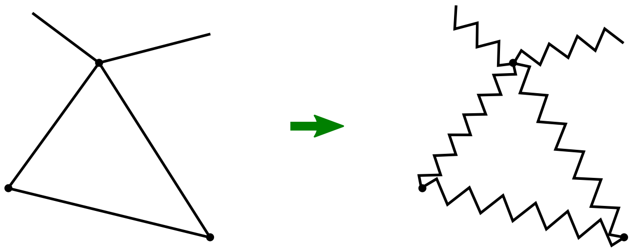

Denote by $\mathcal{T}^1$ the 1-skeleton of $\mathcal{T}$, which is the union of its edges. We would like f to be nonexpanding on $\mathcal{T}^1$. This is true for the restriction of f to each edge individually but might become false in general. To overcome this problem we first “reparametrize” $\mathcal{T}^1$ by contracting a small neighborhood of each vertex in $\mathcal{T}^1$ and by expanding linearly the remaining part of each edge to the whole edge: If $[p, q]$ is an edge, this parametrization smashes small subsegments $[p, p']$ and $[q', q]$ of a fixed length δ to p and q respectively and stretches $[p', q']$ to $[p, q]$. Denote by $\varphi : \mathcal{T}^1 \to \mathcal{T}^1$ the resulting parametrization. If δ is small enough, $h \circ \varphi$ remains short, say with Lipschitz constant $C' < 1$. We now apply the above corrugation process to $h \circ \varphi$ using the *same* corrugations for all the subsegments of length δ that are incident (hence contracted) to a same vertex. Hence, if $[p, p']$ and $[p, q']$ are two such segments, their image by f should coincide. It is now easy to check that choosing the corrugations so that f and $h' := h \circ \varphi$ are at C^0

distance $(1-C')\delta/2$, we have $|f(x)f(y)| \leq d_S(x,y)$ for all $x, y \in \mathcal{T}^1$: either $d_S(x,y) > \delta$ and then

$$|f(x)f(y)| \leq |f(x)h'(x)| + |h'(x)h'(y)| + |h'(y)f(y)| \leq |h'(x)h'(y)| + (1-C')\delta \leq C' d_S(x,y),$$

or $d_S(x,y) \leq \delta$ so that x, y are close to a same vertex v and belong to segments smashed to v by φ. Considering y' on the same segment as x and at the same distance to v as y we conclude that $|f(x)f(y)| = |f(x)f(y')| \leq d_S(x,y') = d_S(x,y)$ by construction.

When h is just nonexpanding rather than short, we replace h by Ch for some $C < 1$ arbitrarily close to 1 to obtain an approximation of Ch on $\mathcal{T}^1$, which is also an approximation of h on $\mathcal{T}^1$.

It remains to invoke the extension theorem 6.3 for each triangle t of $\mathcal{T}$. Consider a subdivision of ∂t such that f is linear on each segment of this subdivision. Let $p_1, \ldots, p_n$ be the vertices of the subdivision. The extension theorem applied to $\{p_1, \ldots, p_n\}$, the restriction of f to the p_i's and $P := t$ provides the desired piecewise distance preserving map. Moreover, if the triangles of $\mathcal{T}$ are small enough, applying a uniform subdivision if necessary, then f and h will be C^0 close on the whole surface S. □

References

[BMR95] Marshall Bern, Scott Mitchell, and Jim Ruppert. Linear-size nonobtuse triangulation of polygons. *Discrete & Computational Geometry*, 14(4):411–428, 1995.

[BZ95] Yuriy Dmitrievich Burago and Viktor Abramovich Zalgaller. Isometric piecewise-linear immersions of two-dimensional manifolds with polyhedral metrics into $\mathbb{R}^3$. *Algebra i analiz*, 7(3):76–95, 1995. Transl. in St Petersburg Math. J. (7)3:369–385.

[DM68] P. H. Doyle and D. A. Moran. A short proof that compact 2-manifolds can be triangulated. *Inventiones Mathematicae*, 5(2):160–162, 1968.

[Kui55] Nicolaas Kuiper. On C^1-isometric imbeddings. *Indag. Math.*, 17:545–556, 1955.

[LP15] Nina Lebedeva and Anton Petrunin. Local characterization of polyhedral spaces. *Geometriae Dedicata*, 179(1):161–168, 2015.

[Mae11] Hiroshi Maehara. On a proper acute triangulation of a polyhedral surface. *Discrete Mathematics*, 311(17):1903–1909, 2011.

[Pak06] Igor Pak. Inflating polyhedral surfaces. 2006.

[Pak08] Igor Pak. Inflating the cube without stretching. *The American Mathematical Monthly*, 115(5):443–445, 2008.

[Pin85] Ulrich Pinkall. Hopf tori in S^3. *Inventiones mathematicae*, 81(2):379–386, 1985.

[PY16] Anton Petrunin and Allan Yashinski. Piecewise distance preserving maps. *St. Petersburg Mathematical Journal*, 27(1):155–175, 2016.

[Rad25] Tibor Radó. über den Begriff der Riemannschen Fläche. *Acta Litt. Sci. Szeged*, 2:101–121, 1925.

[Sar09] Shubhangi Saraf. Acute and nonobtuse triangulations of polyhedral surfaces. *European Journal of Combinatorics*, 30(4):833–840, 2009.

[Zam13] Carol T. Zamfirescu. Survey of two-dimensional acute triangulations. *Discrete Mathematics*, 313(1):35–49, 2013.