Some Algorithmic Aspects of PL Embeddings

Lecture notes

Francis Lazarus

2020

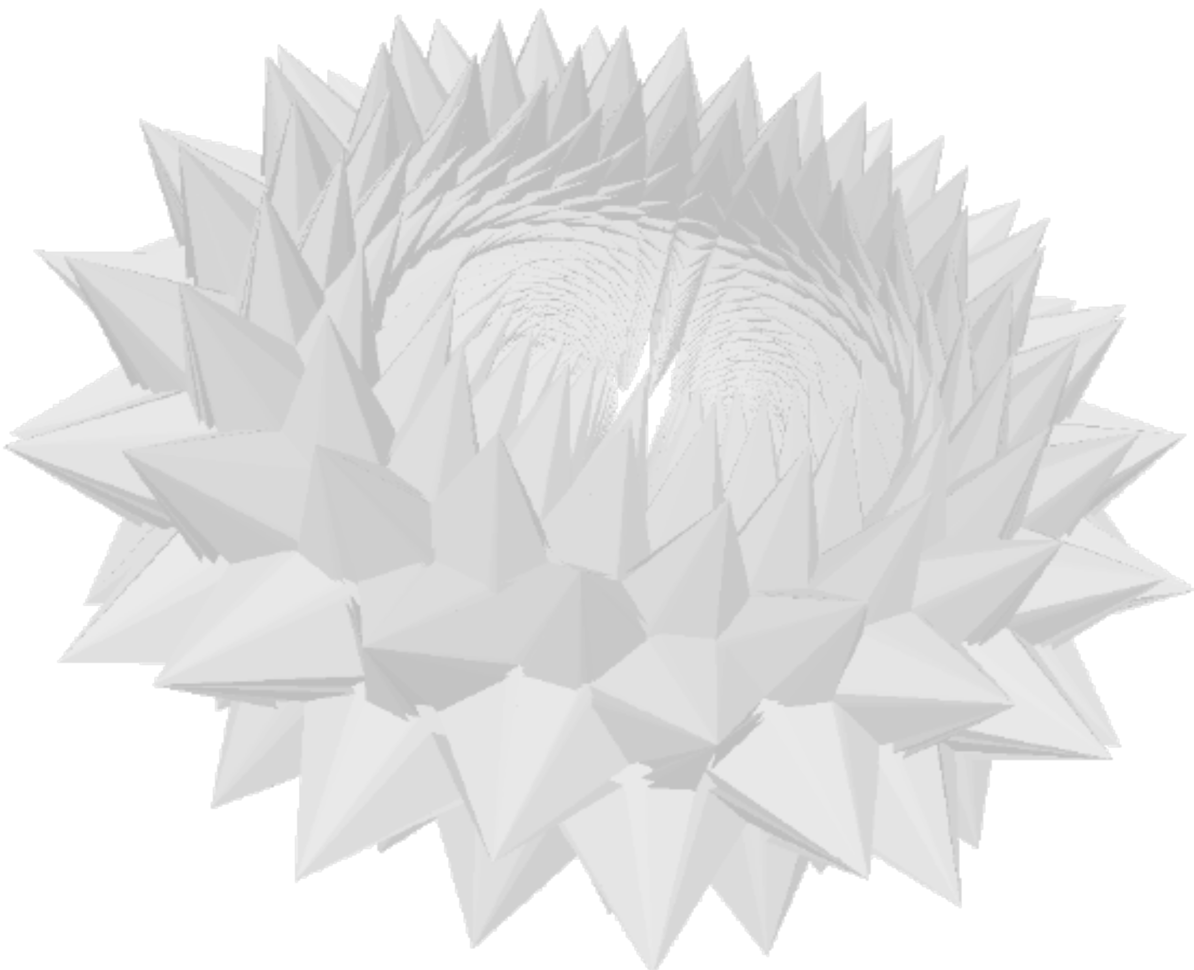

Contents

1

Isometric PL Embedding of Surfaces

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.

1.1 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 1.1.1. 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.

Exercise 1.1.2. 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 1.1.3. 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 1.1.4. 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.

[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.

1.2 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 1.2.1 (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 1.2.2. 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 1.2.3. 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.

1.3 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.

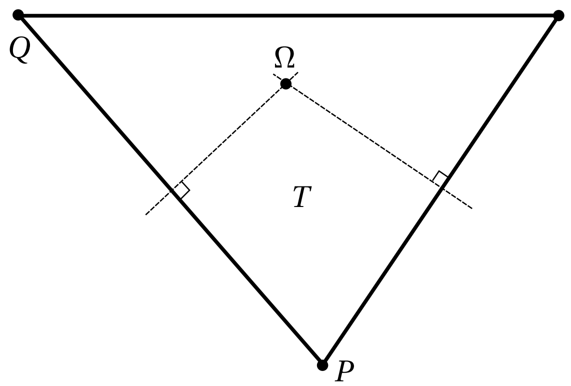

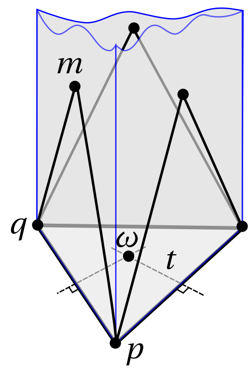

Lemma 1.3.1. *The above embedding of the sides of T extends to a PL isometric embedding of T lying inside the prism above t. Moreover, refining this isometric embedding we can enforce that its C^0 distance to the linear embedding $T \to t$ is arbitrarily small.*

PROOF. Let ω' the point vertically above ω such that $|p\omega'| = |P\Omega|$. Refer to Figure 1.2 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.

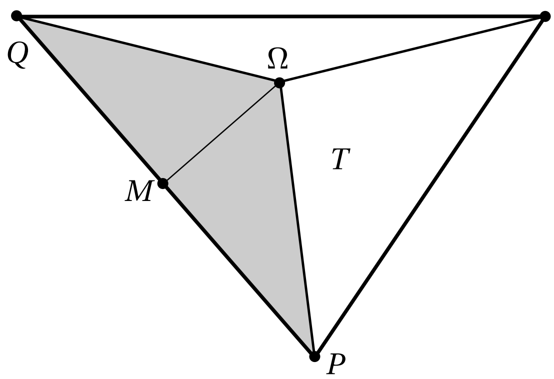

We show below how to fold the subtriangle $[PQ\Omega]$ in the prism above t so that its boundary fits the broken line $pmq\omega'p$. Similar constructions apply to the other two subtriangles so that putting together the three constructions we obtain the desired embedding of T. The second part of the lemma will follow after subdividing both T and t uniformly into sufficiently small triangles as on Figure 1.1. We can indeed

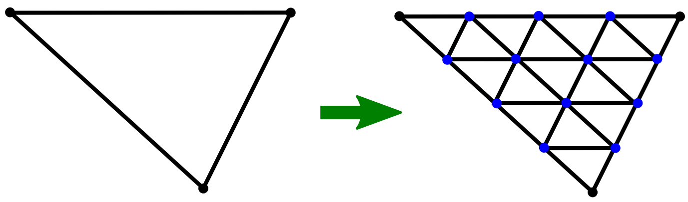

Figure 1.1: Uniform subdivision of a triangle. The vertices of the subdivision have barycentric coordinates $(i/n, j/n, k/n)$ for $i, j, k \in \mathbb{N}$ and $i + j + k = n$ for some fixed n.

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 1.2.

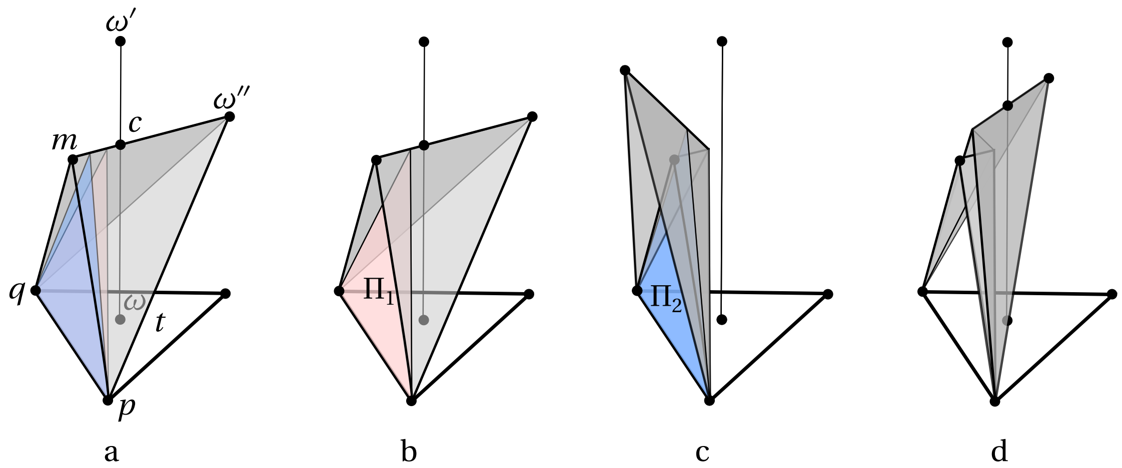

Figure 1.2: a: fold of $[PQ\Omega]$ along its height $M\Omega$ and the two planes Π_1 and Π_2. b: the plane Π_1. c: the plane Π_2 and reflection in Π_1. d: reflection in Π_2.

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. $\square$

Exercise 1.3.2. 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).

1.4 Proof of the Burago and Zalgaller theorem

We shall assume once for all that the surface S in Theorem 1.2.1 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 1.2.1. 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 1.1 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 1.3.
- 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 1.3 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 1.2.1 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 1.2.1 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 1.4.1. 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 1.4.2 (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 1.4.3. Prove that $f_{\varphi,\psi,\lambda}$ is indeed conformal with the claimed conformal factor.

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 $C_{2\pi}$

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

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.

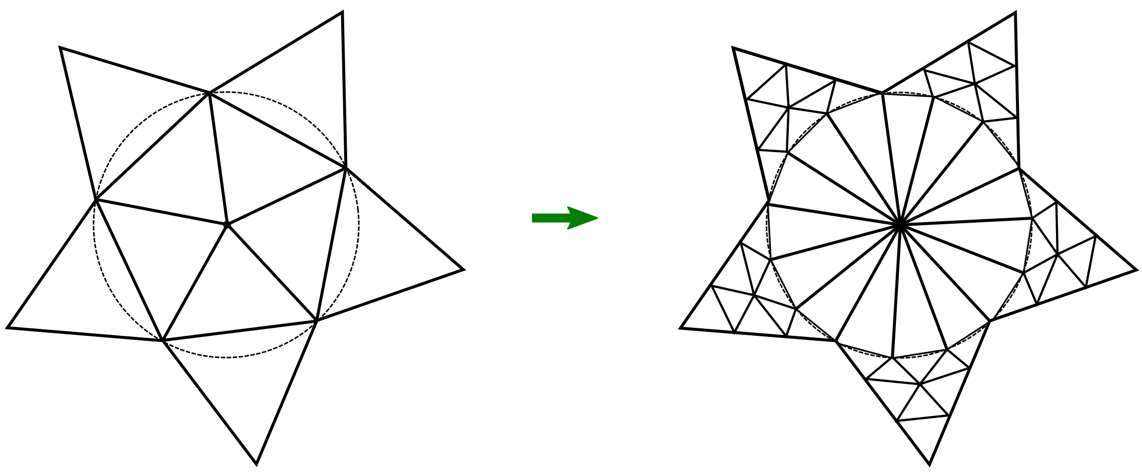

Replacing $N_v(5)$ by $N_v(15)$ allows to extend the uniform subdivision of the boundary of the cogged disk $N'_v(5)$.

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 1.4.2 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 1.4.2 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 1.5 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 1.3 we can find uniform conditions on the degree of similarity and on the contraction factor that allow to apply the basic construction of Section 1.3. 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 1.3 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 1.3 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.

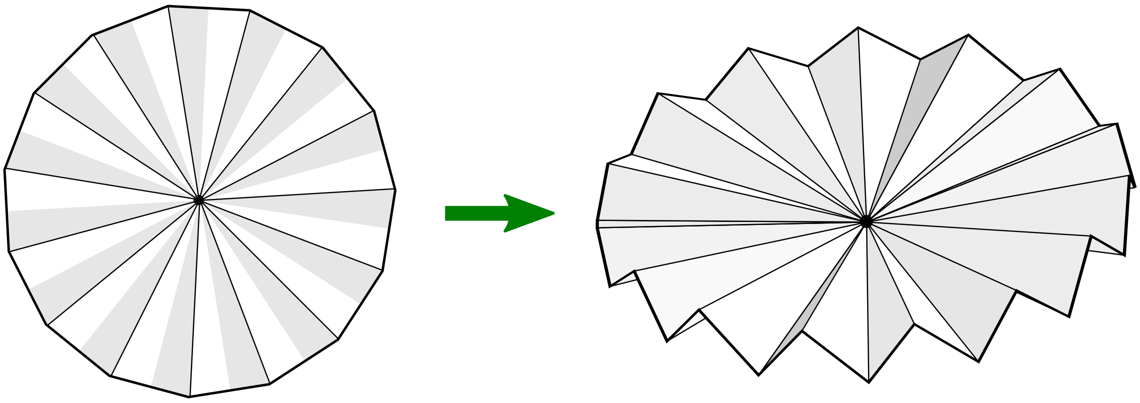

Beware that the left disk is not flat at its center!

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 1.3 has been applied to each triangle of a PL approximation (Figure 1.5, left) of this Hopf (conformal) torus to obtain the PL isometric embedding of Figure 1.5, right.

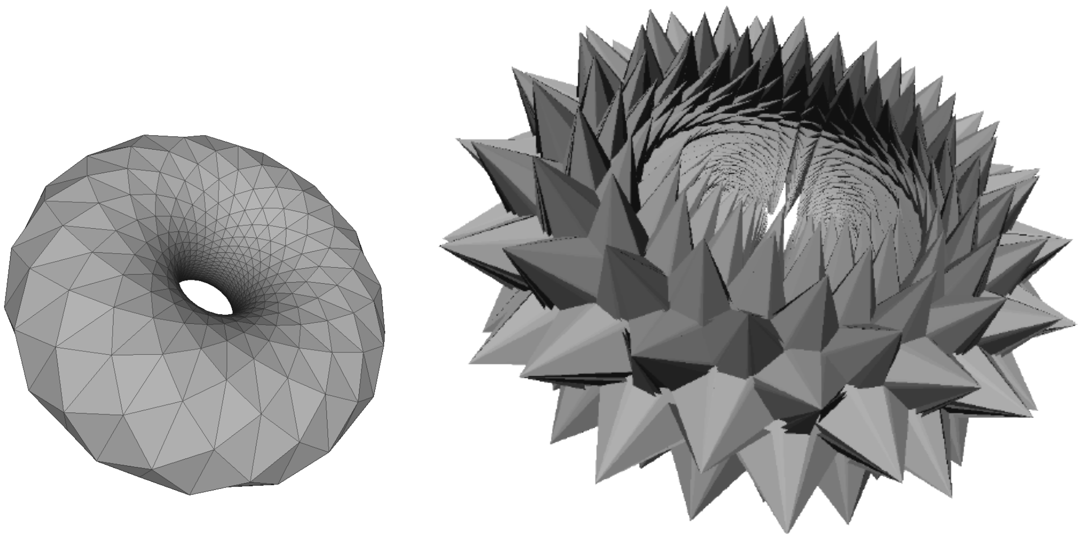

Figure 1.5: Left, a short PL embedding of the square fat torus. Right, The resulting PL isometric embedding of the square flat torus computed by Florent Tallerie. The triangulation is composed of 170,040 triangles.

1.5 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 1.5.1. 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 1.5.2 (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 1.5.1 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 1.6. 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 1.6. 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.

PROOF. Consider a disk of radius ρ with center c on the middle segment of $e = pq$ and crossing the angle bisector at p (with angle ϕ). Reducing ρ if necessary we may assume that the disk is tangent to the bisector at a point x. The triangle $[pxc]$ has thus a right angle at x so that $\rho = |cx| = |pc|\sin(\phi/2) \ge (h/3+\rho)\sin(\theta/2) > r$. Choosing $\rho \le r$ we ensure that the disk does not intersect the bisector. This is true for all three sides of t, hence none of the disks intersect. □

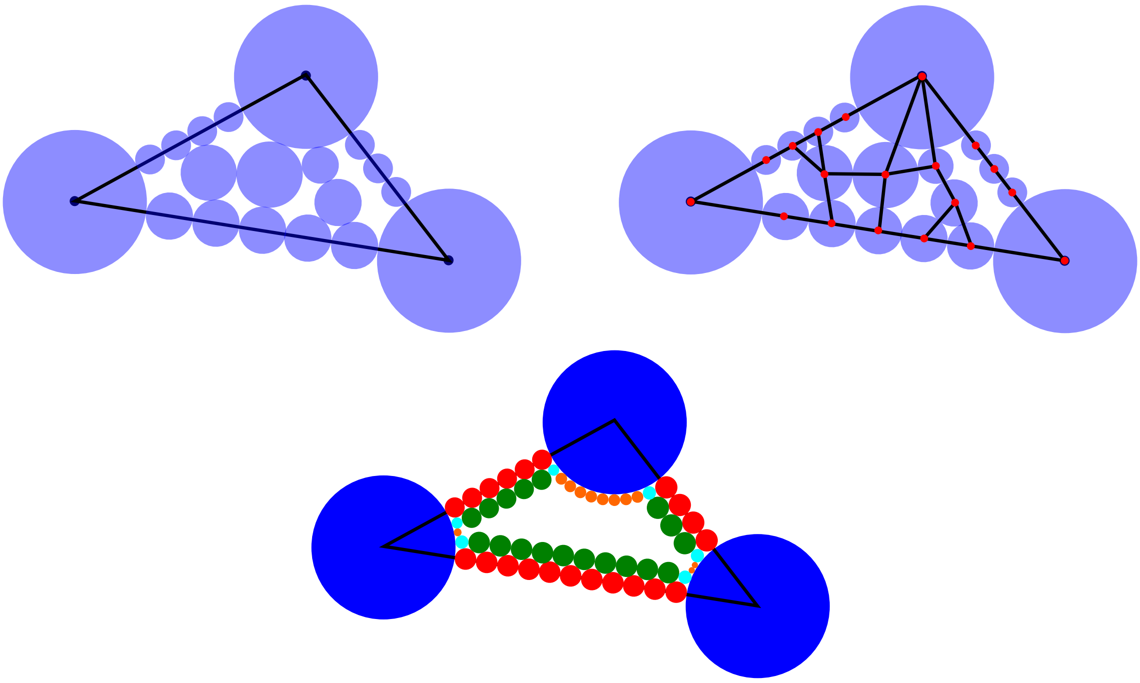

Figure 1.6: Upper left: a packing of disks covering the edges. Upper right: the corresponding contact graph. Lower middle: the vertex-disks (dark blue), the edge-disks (red), and the three types of disks in the circular row: 1 (green), 2 (light blue) and 3 (orange).

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 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.

   PROOF. Same proof as above, noting that each disk of type 1 is included in a larger disk of radius $\rho(1+\sqrt{3})$ centered at the tangent point of two disks on the first row. □

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 1.6.

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 1.7. The graph has

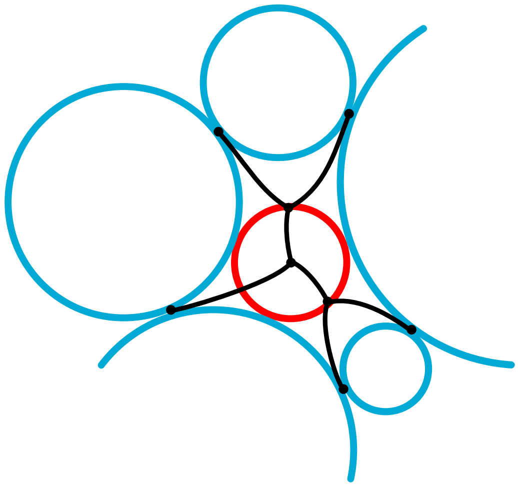

Figure 1.7: The medial axis of a circular sequence of 5 disks. Adding the middle disk (red) splits the face of the contact graph into smaller faces.

one leaf vertex per contact point of touching disks in $\mathscr{D}$. Every vertex 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.

[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.

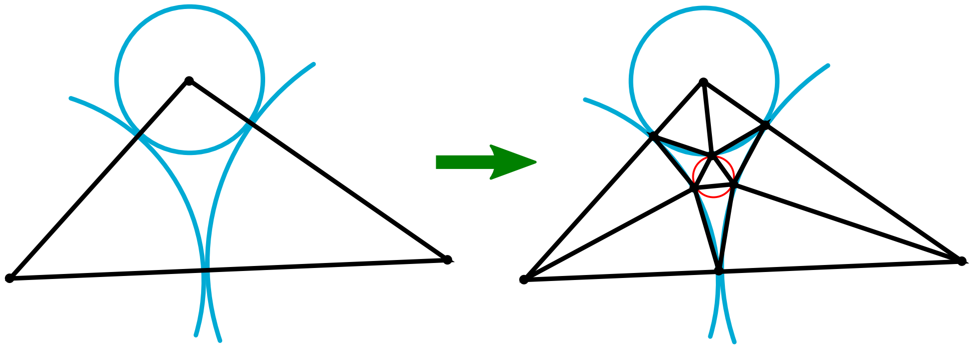

The fact that the faces of the contact graph incident to the edges of t are triangles thus implies that they are subdivided into acute triangles, whence the second part of the lemma. □

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 1.5.3 (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 1.5.2. Then, inside every right triangle of $\mathscr{T}'$ we flip the interior subdividing edge parallel to its hypotenuse.

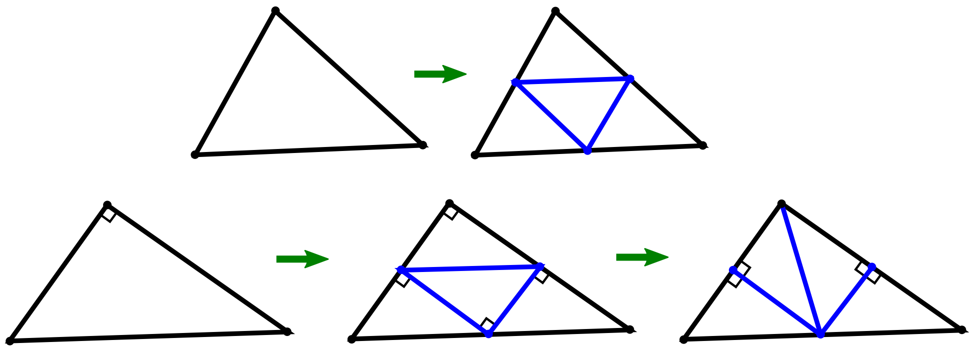

This replaces two right subtriangles by two other congruent right subtriangles. Let $\mathscr{T}''$ be the resulting triangulation. We also denote by $\mathscr{M}$ the set of midpoints introduced in $\mathscr{T}'$ (or equivalently in $\mathscr{T}''$) and by V'' the set of vertices of $\mathscr{T}''$.

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.

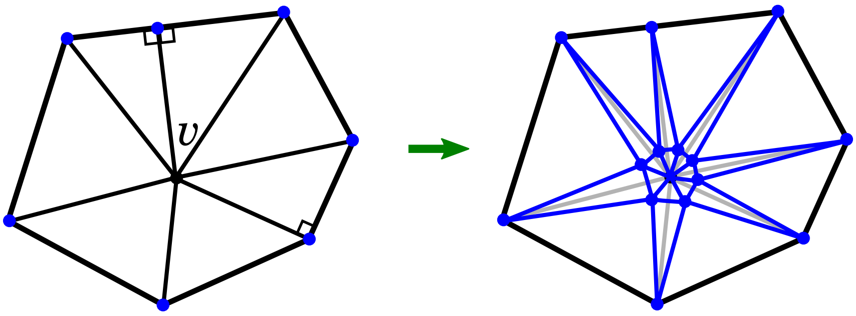

If the central "wheel" replacing v is small enough all the triangles in the new star subdivision will be acute. Since the vertices in $V'' \setminus \mathscr{M}$ are pairwise non-adjacent, their open stars are pairwise non-intersecting and we can perform a similar re-triangulation in every star independently.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

1.6 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 1.2.1 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 1.6.1 (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 1.5.3 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, $|\mathscr{T}_2| = 12|\mathscr{T}|$.

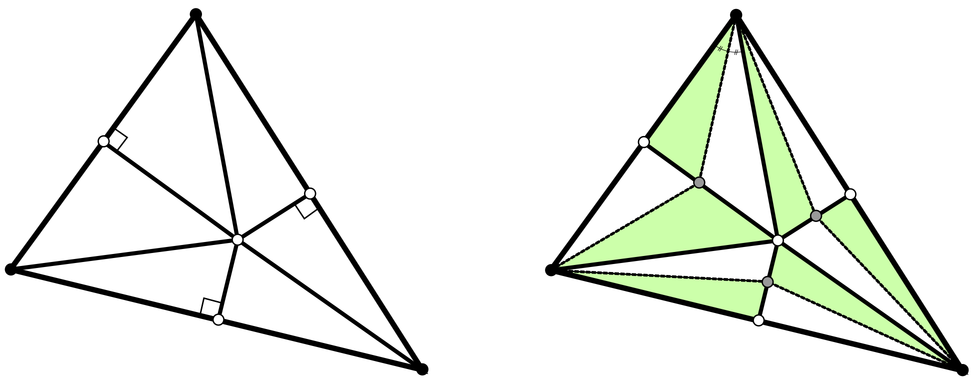

Left, first subdivision. Right, each subtriangle is further split along a bisector resulting in a triangulation $\mathscr{T}_2$. Every triangle of $\mathscr{T}_2$ has one (black) vertex in V_0, one (white) vertex in V_1 and one (grey) vertex in V_2.

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 1.6.2 (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 1.6.3 (Brehm, 1981)**.** *Let $\{p_1, \dots, p_n\}$ be a set of n points contained in a convex polygon P in the plane. Then, any nonexpanding map $\{p_1, \dots, p_n\} \to \mathbb{E}^2$ 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 $i = 2, \dots, n$. 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 $\mathscr{T}$ of P such that h is an isometry on each triangle t of $\mathscr{T}$. 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 $\Omega = \cup_{t \in \mathscr{T}} \Omega \cap t$ 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 $a p_1 b C$ as a fan)

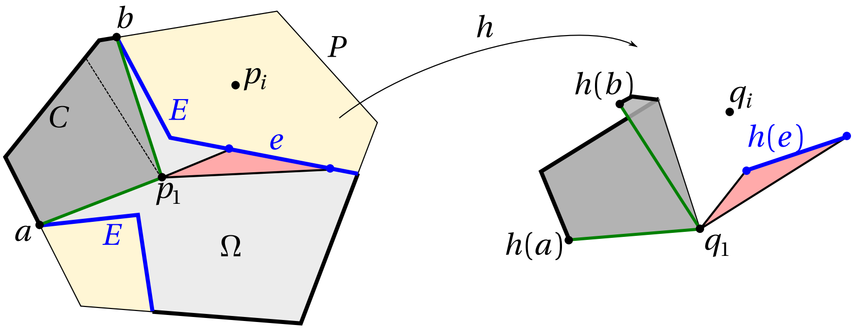

The map f thus defined is clearly continuous and piecewise distance preserving. Moreover, we have $f(p_1) = q_1$ and $p_i \in P \setminus \Omega$ for $i \geq 2$, so that $f(p_i) = h(p_i) = q_i$ and f is indeed an extension of $p_i \mapsto q_i$. □

PROOF OF THEOREM 1.6.2. 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.

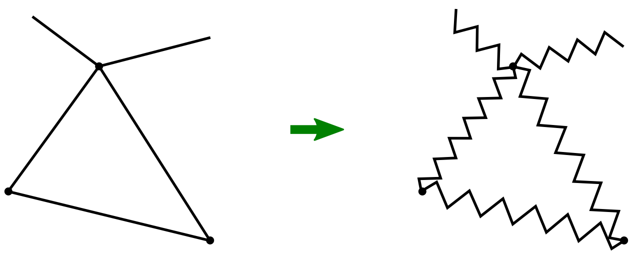

Denote by $\mathcal{T}^1$ the 1-skeleton of $\mathcal{T}$, which is the union of its edges. We would like f to be nonexpanding on $\mathcal{T}^1$. This is true for the restriction of f to each edge individually but might become false in general. To overcome this problem we first "reparametrize" $\mathcal{T}^1$

by contracting a small neighborhood of each vertex in $\mathscr{T}^1$ 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 $\varphi:\mathscr{T}^1\to\mathscr{T}^1$ 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 $|f(x)f(y)|\le d_S(x,y)$ for all $x,y\in\mathscr{T}^1$: either $d_S(x,y)>\delta$ and then

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

or $d_S(x,y)\le\delta$ 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 $|f(x)f(y)|=|f(x)f(y')|\le d_S(x,y')=d_S(x,y)$ 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 $\mathscr{T}^1$, which is also an approximation of h on $\mathscr{T}^1$.

It remains to invoke the extension theorem 1.6.3 for each triangle t of $\mathscr{T}$. Consider a subdivision of ∂t such that f is linear on each segment of this subdivision. Let $p_1,\dots,p_n$ be the vertices of the subdivision. The extension theorem applied to $\{p_1,\dots,p_n\}$, the restriction of f to the p_i's and $P:=t$ provides the desired piecewise distance preserving map. Moreover, if the triangles of $\mathscr{T}$ are small enough, applying a uniform subdivision if necessary, then f and h will be C^0 close on the whole surface S. $\square$

2

Embedding in Euclidean spaces: the double dimension case

Contents

Historically, the notion of manifold, say at the time of Gauss (1777-1855), was thought *extrinsically* as a subspace of some Euclidean space. Starting with Riemann (1854) this notion evolved toward the more abstract *intrinsic* definition of a space which is locally Euclidean. It results from the famous Whitney embedding theorems that the intrinsic and extrinsic point of view are indeed the same. The weak Whitney embedding theorem (1936) claims that every n-manifold embeds in $\mathbb{R}^{2n+1}$, while the strong version reduces the dimension of the target space from $2n+1$ to $2n$. However, those theorems do not say anything for embedding more complicated spaces. In this lecture we look at this question from the algorithmic point of view. For this, we need to describe in a combinatorial way the spaces we are interested in.

2.1 Topological prerequisites

2.1.1 Complexes

A natural way to describe spaces is to express them as assembly of elementary pieces. In practice, the pieces are **cells**, i.e. subspaces homeomorphic to balls of various dimensions. Using cells in place of more complicated building blocks greatly simplifies the computation of topological invariants such as homotopy or homology groups. An assembly of cells is called a **complex**. Depending on the shape of the cell, we obtain different categories of complexes with suitable notions of morphisms. We thus have among others, simplicial, cubic, polyhedral, delta, or cellular (CW) complexes. The most general complexes are the cellular ones. Their definition is not entirely combinatorial as it relies on the notion of attaching maps which are continuous maps sending the boundary of a cell to cells of lower dimensions. It is not the purpose of these notes to give a formal definition of all the kinds of complexes. We will essentially stick to finite simplicial complexes and finite one dimensional cellular complexes.

Graphs: One dimensional cellular complexes are also called *graphs*. Their zero and one dimensional cells are called **vertices** and **edges**, respectively. A graph is thus a set of vertices connected by edges. Its topological type, up to homeomorphism, is described combinatorially by two sets, one for the vertices and one for the edges, and a map associating each edge to a pair of possibly identical vertices, called its **endpoints**. An edge whose endpoints coincide is a **loop** edge. If distinct edges share the same endpoints, they form a **multiple** edge. A graph without loop and multiple edge is said **simple**. A simple graph is thus another name for a one dimensional simplicial complex.

Simplicial complexes: Their cells are simplices. A k dimensional simplex, or **k-simplex**, is the convex hull of $k+1$ affinely independent points $p_0,\dots,p_k$ in some $\mathbb{R}^d$ and is denoted by $[p_0,\dots,p_k]$. The empty set is also considered as a simplex[1] with dimension -1. The convex hull of any subset of the p_i's is a **face** of the k-simplex and is itself a simplex of dimension at most k. A geometric **simplicial complex** K is a collection of simplices in some $\mathbb{R}^d$ such that (1) any face of a simplex in K is in K, (2) the intersection of any two simplices in K is a common face of the two simplices. The **dimension** of K is the maximum dimension of its simplices. The union of the simplices of K is denoted by $|K|$ and indifferently called the **underlying set**, the **polyhedron**, the **carrier**, or the **total space** of K. Any simplex $\sigma \in K$ (formally its carrier $|\sigma|$) is closed in $|K|$. Its interior, as a cell, is denoted by $\mathring{\sigma}$. By the above property (2), $|K|$ is the disjoint union of the interior of its simplices. In other words, every point in $|K|$ belongs to the interior of exactly one **supporting** simplex. A **subdivision** of K is any simplicial complex L such that $|K| = |L|$ and such that every simplex of L is contained in a simplex of K. Two complexes are **isomorphic** if there is a one-to-one correspondence between their simplices that preserves dimension and commutes with faces: A face of a simplex corresponds to a face of the corresponding simplex. The complexes are said **PL homeomorphic** when they have isomorphic subdivisions.

[1]This assumption facilitates the definition of the join operation. See Section 2.3.1.

A simplicial complex can be described combinatorially by an **abstract simplicial complex**. This is a collection of finite subsets of a ground set with the hereditary property: Any subset of a subset in the collection is itself in the collection. The subsets in the collection are its abstract simplices. A **simplicial map** $f : K \to L$ between abstract simplicial complexes is a map between their vertex sets that sends simplices to simplices: $\sigma \in K \implies f(\sigma) \in L$. For geometric simplicial complexes simplicial maps extend uniquely to continuous maps by affine interpolation over each simplex of the value at its vertices. Up to homeomorphism, we can realize an abstract simplicial complex by gluing along common faces realizations of its simplices in some Euclidean space. Alternatively, the realization of an abstract simplicial complex A with ground set V can be defined as the subset of $[0,1]^V$, with the induced topology, of all points $(t_v)_{v\in V}$ such that $\{v : t_v > 0\} \in A$ and $\sum_{v\in V} t_v = 1$.

The **barycentric subdivision**, sd K, of a geometric simplicial complex is obtained by subdividing its simplices recursively by dimension order: The edges are replaced by starring their two boundary points from their barycenter, then the triangles are replaced by starring their already subdivided edges from their barycenter, and so on. This process is repeated, each time replacing a simplex by a cone with apex its barycenter over its already subdivided boundary. Each simplex of the resulting barycentric subdivision is the convex hull of the barycenters of an increasing sequence of simplices of K. The abstract simplicial complex associated to sd K has thus $K \setminus \{\emptyset\}$ itself for ground set and its nonempty simplices have the form $(\sigma_0, \sigma_1, \dots, \sigma_{k-1})$ where $\sigma_0 \subset \cdots \subset \sigma_k$ is a strictly increasing sequence of nonempty simplices of K.

2.1.2 Embeddings

A topological embedding is just a map inducing a homeomorphism onto its image (endowed with the induced topology of the target space). For a compact space, in particular for a finite complex, an embedding is just a continuous injective map. For a simplicial complex K, we may consider more constrained kinds of embeddings. A **linear mapping**[2] of K is a map $f : |K| \hookrightarrow \mathbb{R}^d$ whose restriction to each simplex of K is affine. In other words, f sends simplices in K to geometric simplices in $\mathbb{R}^d$, and is entirely determined by the image of the vertices of K. The mapping is **piecewise linear**, or **PL**, if K has a subdivision K' such that f is a linear mapping of K'. A **linear embedding**[3] of K is a linear mapping which is also an embedding, and similarly for a PL embedding. The three notions of embeddings (topological, PL and linear) are increasingly restrictive in the sense that K may have a topological embedding but no PL embedding into $\mathbb{R}^d$, while K may have a PL embedding but no linear embedding into $\mathbb{R}^d$. For more details on this, see Section 2 and Appendix C in [MTW11] or the notes of Section 5.1 in [Mat08]. From a computational perspective, we will be mainly interested in PL and linear embeddings.

The weak Whitney theorem has a simple extension to complexes.

---

[2] It is common practice in PL topology to use the term *linear* where the term *affine* would be more appropriate.

[3] Linear embeddings are also called geometric embeddings.

Proposition 2.1.1. *Any finite simplicial complex of dimension n embeds linearly into $\mathbb{R}^{2n+1}$.*

PROOF. Define a linear mapping f of the n dimensional complex K into $\mathbb{R}^{2n+1}$ by mapping the vertices of K to points in general position in $\mathbb{R}^{2n+1}$, *i.e.*, such that no hyperplane contains more than $2n+1$ points. One may for instance choose the points on the *moment curve* $t \mapsto (t, t^2, \dots, t^{2n+1})$. We claim that f is an embedding. This is clearly the case when restricted to any simplex of K: The simplex has at most $n+1$ vertices which are sent to affinely independent points by the general position assumption. To see that f is injective we just need to prove that distinct simplices have their interior sent to disjoint sets. So, let $\sigma = [v_1, \dots, v_k]$ and $\tau = [w_1, \dots, w_\ell]$ be two distinct simplices of K. Since $k + \ell \leq 2n+2$, the general position assumption implies that the image points $f(v_1), \dots, f(v_k), f(w_1), \dots, f(w_\ell)$ span a simplex of dimension $k+\ell-1$ and that $f(\sigma), f(\tau)$ are two distinct faces of this simplex. It follows that $f(\mathring{\sigma})$ and $f(\mathring{\tau})$ are indeed disjoint. As already observed, injectivity implies embedding for finite simplicial complexes. □

Exercise 2.1.2. Prove that any set of points on the moment curve $t \mapsto (t, t^2, \dots, t^d)$ in $\mathbb{R}^d$ is in general position, *i.e.*, that no hyperplane contains more than d of the points.

In view of the Proposition, the question of whether an n-dimensional complex embeds into $\mathbb{R}^d$ is only interesting for $d \leq 2n$. In these notes we will focus on the case $d = 2n$. There is indeed a nice invariant that leads to practical algorithms in this case. In the next section we consider the case $n = 1$, which amounts to decide if a graph is planar.

2.2 Graph embedding

The graph planarity problem has received much attention in the computer science community, culminating with the linear time algorithm of Hopcroft and Tarjan [HT74]. It happens that topological, PL and linear embeddability are equivalent for embedding graphs into the plane, so that any planar graph may be drawn with straight lines for the edges. See the lecture notes [LdM17] for more details. The most striking result concerning graph planarity is probably the Kuratowski's criterion in terms of forbidden graphs. Recall that the **complete graph** K_5 is obtained by connecting five vertices in all possible ways, while the **complete bipartite graph** $K_{3,3}$ is obtained by connecting each of three independent (*i.e.*, pairwise non-connected) vertices to each of three other independent vertices.

Theorem 2.2.1 (Kuratowski, 1929)**.** *A graph is planar if and only if it does not contain a subdivision of K_5 or $K_{3,3}$ as a subgraph.*

See [LdM17] for a proof. We shall refer to this theorem but use a different path to derive a planarity criterion due to van Kampen (1932) that is more amenable to a generalization to higher dimensions. We follow the presentation of Wu [Wu85].

Let $G = (V, E)$ be a graph with vertex set V and edge set E. Consider the set C_2 of unordered pairs of disjoint edges. We denote by $\sigma \times \tau$ such an ordered pair, where $\sigma, \tau \in E$ do not share any vertex. (Remark that $\sigma \times \tau = \tau \times \sigma$.) Intuitively, $\sigma \times \tau$ is a 2-dimensional rectangular cell. We denote by C^2 the vector space $\mathbb{Z}_2^{C_2}$ (we write $\mathbb{Z}_2$ for $\mathbb{Z}/2\mathbb{Z}$), viewing vectors as maps $C_2 \to \mathbb{Z}_2$. Similarly, we consider the set C_1 of pairs $(v, \sigma) \in V \times E$ such that v is not an endpoint of σ and the vector space C^1 of maps $C_1 \to \mathbb{Z}_2$. We also write $v \times \sigma$ for (v, σ). The **coboundary** operator is the morphism $\delta : C^1 \to C^2$ defined for any $c \in C^1$ by[4]

$$\delta c(\sigma \times \tau) = c(v_1 \times \tau) + c(v_2 \times \tau) + c(w_1 \times \sigma) + c(w_2 \times \sigma)$$

where $\sigma \times \tau \in C_2$, and v_1, v_2 (resp. w_1, w_2) are the endpoints of σ (resp. τ). Let $(\sigma \times \tau)^* \in C^2$ take value 1 at $\sigma \times \tau$ and 0 elsewhere. Similarly, let $(v \times \tau)^* \in C^1$ take value 1 at $v \times \tau$ and 0 elsewhere. Then, the above formula amounts to define the coboundary on the canonical basis of C^1 by

$$\delta(v \times \tau)^* = \sum_{v \in \sigma} (\sigma \times \tau)^* \tag{2.1}$$

It appears that the quotient $C^2/\operatorname{Im}\delta$ is a topological invariant[5].

Lemma 2.2.2. *PL homeomorphic graphs have isomorphic quotient groups $C^2/\operatorname{Im}\delta$.*

PROOF. Since any subdivision of a graph can be obtained by repeatedly splitting edges, it is enough to prove the lemma for a graph G' obtained by splitting an edge $e \in E$ of a graph $G = (V, E)$. Let v be the new vertex splitting e and let e_1, e_2 be the resulting edges in G'. We denote with a prime the groups or maps related to G'. Hence, $\delta' : C'^1 \to C'^2$ is the coboundary operator for G'. We view edges in $E \setminus \{e\}$ as edges of G' as well as edges of G. Define the morphisms $s_1 : C'^1 \to C^1$ and $s_2 : C'^2 \to C^2$ by

$$s_1(c)(u \times \tau) = \begin{cases} c(u \times e_1) + c(u \times e_2) & \text{if } \tau = e, \\ c(u \times \tau) & \text{otherwise} \end{cases} \quad \text{for } c \in C'^1, u \in V, \tau \in E, u \notin \tau$$

and

$$s_2(d)(\sigma \times \tau) = \begin{cases} d(\sigma \times e_1) + d(\sigma \times e_2) & \text{if } \tau = e, \\ d(\sigma \times \tau) & \text{if } \sigma, \tau \neq e \end{cases} \quad \text{for } d \in C'^2, \sigma, \tau \in E, \sigma \cap \tau = \emptyset$$

It is easily checked that s_1 and s_2 are onto and satisfy $\delta s_1 = s_2 \delta'$. The proof is left as an exercise. It follows that $s_2(\operatorname{Im}\delta') \subset \operatorname{Im}\delta$ and that s_2 induces an epimorphism $s_2^* : C'^2/\operatorname{Im}\delta' \to C^2/\operatorname{Im}\delta$. It remains to see that s_2^* is injective. So, suppose that $s_2^*(d + \operatorname{Im}\delta') = 0$, i.e. that $s_2(d) \in \operatorname{Im}\delta$. We have $s_2(d) = \delta c$ for some $c \in C^1$. By surjectivity of s_1, $c = s_1 c'$ for some $c' \in C'^1$, so that $s_2(d) = \delta s1(c') = s_2(\delta' c')$, or equivalently, $s_2(d - \delta' c') = 0$. Now, it is easily seen that this implies $d - \delta' c' = \sum_\sigma \alpha_\sigma \delta'(v \times \sigma)^*$ for some coefficients $\alpha_\sigma \in \mathbb{Z}_2$ (see (2.1)). In other words, $\ker s2 \subset \operatorname{Im}\delta'$. We conclude that $d \in \operatorname{Im}\delta'$ as desired. □

For an element $c \in C^2$, we denote by $[c]_2$ its coset in $C^2/\operatorname{Im}\delta$.

[4] It is common practice to write δc for $\delta(c)$.

[5] This topological invariant is the second equivariant cohomology group of the deleted product of G. Equivalently, this is the second (ordinary) cohomology group of the same deleted product quotiented by the action that exchanges coordinates in the (deleted) product.

2.2.1 The mod 2 van Kampen obstruction

Two paths in the plane are said in **general position** if each one avoids the endpoints of the other one, except at common endpoints, and if they otherwise cross transversally at their finitely many intersection points. An immersion into the plane of $G=(V,E)$ is said in general position if the image of its edges are pairwise in general position. We now associate to any PL immersion $f:|G|\to\mathbb{R}^2$ in general position the element $c_f\in C^2$ given by

$$c_f(\sigma\times\tau)=|f(\sigma)\cap f(\tau)| \mod 2, \qquad \text{for } \sigma,\tau\in E, \sigma\cap\tau=\emptyset$$

Lemma 2.2.3. *$[c_f]_2$ is independent of f.*

We give two proofs, a short proof by picture, and a longer formal one.

PROOF BY PICTURE. Every two general position immersions are related by a sequence of isotopies of $\mathbb{R}^2$ and of local moves as on Figure 2.1. An isotopy or any of the I-IV

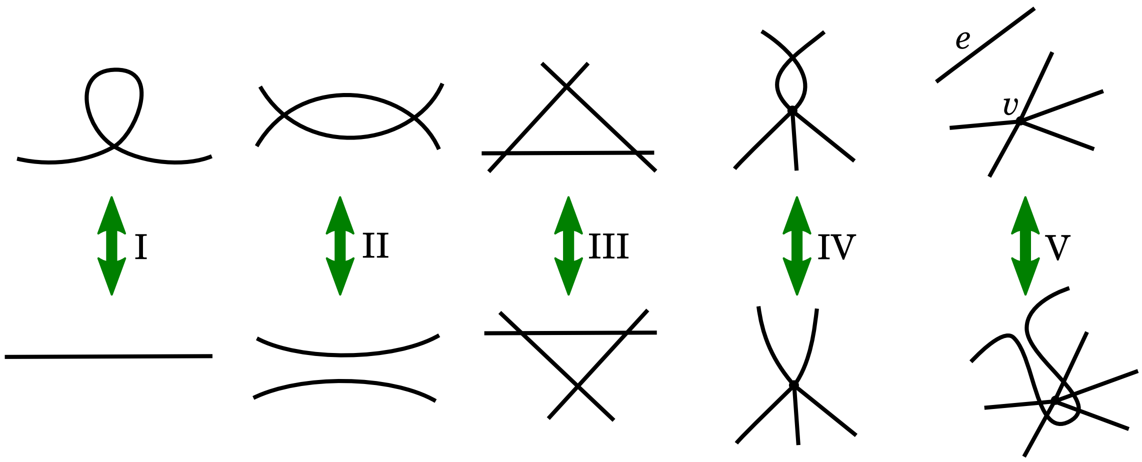

Figure 2.1: The first three moves I, II, III are known as (shadows of) Reidemeister moves. The IV move amounts to a transposition in the edge order around a vertex, while the V move is referred to as a **finger move** or an (e,v)**-move**.

moves leaves c_f unchanged while an (e,v)-move results in an additional term $\delta(v\times e)^*$ in c_f. In any case, $[c_f]_2$ is preserved. □

Exercise 2.2.4. Figure 2.1 actually applies to smooth curves. Can you adapt the proof and find a list of moves specific to the PL category?

A formal proof would require showing that the five moves in Figure 2.1 are the only required moves to transform an immersion into another one. (See Exercise 2.2.4 for the PL case.) We give below a more combinatorial proof due to Wu [Wu85]. We first give a simple relation between winding number and intersection number. Recall that the **winding number** $w(\gamma,p)$ of a plane closed curve γ with respect to a point $p\notin\gamma$ is the total number of times γ travels counterclockwise around[6] p.

[6]Formally, identifying the plane with $\mathbb{C}$, $w(\gamma,p)=\frac{1}{2\pi i}\oint_\gamma\frac{\mathrm{d}z}{z-p}=\frac{1}{2\pi i}\int_\alpha^\beta\frac{\gamma'(t)}{\gamma(t)-p}dt$, where $\gamma:[\alpha,\beta]\to\mathbb{C}$.

Lemma 2.2.5. *Let $w_2(\cdot,\cdot) = w(\cdot,\cdot) \mod 2$ be the mod 2 winding number. For any path π with endpoints p, q in general position with respect to a closed curve γ:*

$$w_2(\gamma, p) - w_2(\gamma, q) = |\gamma \cap \pi| \mod 2$$

where $|\gamma \cap \pi|$ counts the number of intersections between γ and π.

PROOF. We prove the lemma when γ and π are PL curves. The case of continuous curves follows by PL approximation. It is well-known that $w(\gamma, p)$ is the algebraic number of intersections of γ with a ray originating from p. In particular, all rays with origin p have the same algebraic number of intersections with γ. As we move from p towards q along π, aligning the rays from p and from q we see that the winding number changes exactly as we traverse γ and the change is ± 1 depending on the orientation of γ and π at the intersection point. The lemma follows. □

PROOF OF LEMMA 2.2.3, WU'S VERSION. Let $f, g : G \to \mathbb{R}^2$ be two immersions of G in general position.

- Let $e = [p, q]$ be an edge of G. We first consider the case where f and g coincide on $G - e$ (the graph G with the interior of edge e removed). For every edges σ, τ distinct from e, we obviously have $c_f(\sigma \times \tau) = c_g(\sigma \times \tau)$ since both values only depends on the embedding of σ and τ. Let $C_e := f(e) \cdot g(e)^{-1}$ be the closed curve formed by concatenating $f(e)$ with the path $g(e)$ traversed in the opposite direction. Consider the cochain

$$c = \sum_v w_2(C_e, f(v))(v \times e)^*$$

where the sum runs over all vertices of G not incident to e, *i.e.*, distinct from p and q. We compute, writing $\partial \sigma = s - r$

$$\begin{aligned} c_f(\sigma \times e) - c_g(\sigma \times e) &\equiv |f(\sigma) \cap f(e)| - |g(\sigma) \cap g(e)| \mod 2 \\ &\equiv |f(\sigma) \cap C_e| \mod 2 \qquad (\text{since} f(\sigma) = g(\sigma)) \\ &= w_2(C_e, f(s)) - w_2(C_e, f(r)) \qquad (\text{by Lemma 2.2.5}) \end{aligned}$$

On the other hand, we compute

$$\begin{aligned} \delta c(\sigma \times e) &= c(\partial \sigma \times e) + c(\sigma \times \partial e) \\ &= w_2(C_e, f(s)) - w_2(C_e, f(r)) \end{aligned}$$

For disjoint edges σ, τ both distinct from e we trivially have $c_f(\sigma \times \tau) - c_g(\sigma \times \tau) = c(\partial \sigma \times \tau) + c(\sigma \times \partial \tau) = 0$. It follows that $c_f - c_g = \delta c$, or equivalently that $[c_f]_2 = [c_g]_2$.

- We now consider the case where f and g only agree on the vertices of G. Let $e_1, \ldots, e_m$ be the edges of G. We define immersions f_i that agree with g on $e_1, \ldots, e_i$ and with f on the remaining edges. Putting $f_0 = f$, we have by the preceding paragraph, that $[c_{f_{i-1}}]_2 = [c_{f_i}]_2$ for $i = 1, \ldots, m$. It follows that $[c_f]_2 = [c_g]_2$.

- We finally consider the case of arbitrary f and g in general position. Denote by $v_1, \dots, v_n$ the vertices of G. Using an induction on the number of vertices one can construct a PL homeomorphism H of the plane that sends $f(v_i)$ to $g(v_i)$. On the one hand, we have $[c_f]_2 = [c_{H\circ f}]_2$ and on the other hand $[c_{H\circ f}]_2 = [c_g]_2$ by the preceding paragraph[7]. We conclude that $[c_f]_2 = [c_g]_2$ in the general case.

□

In view of Lemma 2.2.3, we denote by[8] $\kappa_2(G)$ the value of $[c_f]_2$ computed from any immersion f in general position. It appears that $\kappa_2(G)$ only depends on the topology of $|G|$ and not on its cellular decomposition. In the proof of Lemma 2.2.2, we introduced an edge splitting isomorphism $s_2^*: C'^2/\mathrm{Im}\,\delta' \to C^2/\mathrm{Im}\,\delta$. By composing such morphisms, we obtain a natural subdivision isomorphism s^* between the quotients $C^2/\mathrm{Im}\,\delta$ computed for a subdivision of a graph or the graph itself. Likewise, if H is a subgraph of G, there is a natural inclusion morphism ι^* between the quotient groups for G and H. The topological invariance of $\kappa_2(G)$ is formalized in the following easy lemma whose proof is left to the reader.

Lemma 2.2.6. *If $\iota: H \hookrightarrow G$ is a cellular inclusion, we have $\iota^*(\kappa_2(G)) = \kappa_2(H)$. Similarly, if G' a subdivision of G and s is the corresponding subdivision operator, we have $s^*(\kappa_2(G')) = \kappa_2(G)$.*

The topological invariant $\kappa_2(G)$ of G is called the mod 2 **van Kampen obstruction**. Note that $c_f = 0$ if f is an embedding. Hence, $\kappa_2(G) = 0$ whenever G is planar. It thus follows from the next lemma that the Kuratowski forbidden graphs K_5 and $K_{3,3}$ are non-planar.

Lemma 2.2.7. *$\kappa_2(K_5)$ and $\kappa_2(K_{3,3})$ are each nonzero.*

PROOF. Compute κ_2 using your preferred embeddings of K_5 and $K_{3,3}$. Can you draw them with a single crossing? □

We are now ready to state that the van Kampen obstruction is a good invariant to test graph embeddability in the plane.

Theorem 2.2.8. *A graph is planar if and only if its mod 2 van Kampen obstruction cancels.*

PROOF. We already observed that the condition is necessary. Suppose that a graph G satisfies $\kappa_2(G) = 0$. By the preceding lemmas 2.2.6 and 2.2.7, G cannot contain a subdivision of K_5 or $K_{3,3}$. It ensues from Kuratowski's theorem that G is planar. □

This theorem is known as the (strong) Hanani-Tutte theorem in graph theory and is expressed as follows: any (generic) immersion of a non-planar graph contains two disjoint edges whose images cross oddly.

[7] This argument found by Axel Péneau simplifies the proof of Wu.
[8] The subscript 2 is used to emphasize that we consider mod 2 cohomology.

2.3 The van Kampen-Flores Theorem

The mod 2 van Kampen obstruction constructed for a graph G as in the previous section can be interpreted as a certain equivariant cohomology class of the **deleted product** $G \times_\Delta G$ of G. This deleted product is composed of all the products of disjoint cells (vertex or edge) of G and has the same equivariant homotopy type (it is even an equivariant deform retract) as the topological deleted product $|G| \times |G| \setminus \Delta$, where $\Delta = \{(x,x) \in |G| \times |G|\}$ is the diagonal of $|G| \times |G|$. Here, by *equivariant* we refer to invariance with respect to some action of $\mathbb{Z}_2$ on the deleted product[9]. The mod 2 van Kampen obstruction for graphs can be generalized to complexes of dimension $n > 1$ using integer instead of $\mathbb{Z}_2$ coefficients and its non-vanishing is indeed an obstruction to embedding in $\mathbb{R}^{2n}$. For $n > 2$ this obstruction also provides a sufficient condition for embeddability in $\mathbb{R}^{2n}$. However, this is not the case for $n = 2$ as Freedman, Krushkal and Teichner [FKT94] constructed a relatively simple simplicial complex of dimension 2 whose van Kampen obstruction vanishes but that cannot be embedded in $\mathbb{R}^4$. In this section we look at a slightly different approach based on the deleted join rather than the deleted product. It leads to the van Kampen (1932) – Flores (1933) theorem that for every dimension n the n-skeleton of the $(2n+2)$-simplex does not embed into $\mathbb{R}^{2n}$. We follow the exposition of de Longueville [dL13, Ch. 4].

2.3.1 Join operations

The join

The **join** $X * Y$ of two topological spaces X and Y is the quotient $X \times Y \times I/\sim$ where $I = [0,1]$ is the unit interval and the equivalence classes of $\sim$ are of the form $\{x\} \times Y \times \{0\}$, $X \times \{y\} \times \{1\}$ and are otherwise singletons. Intuitively, $X * Y$ is the "cube" $X \times Y \times I$ where we have collapsed the face $X \times Y \times \{0\}$ to X and the face $X \times Y \times \{1\}$ to Y.

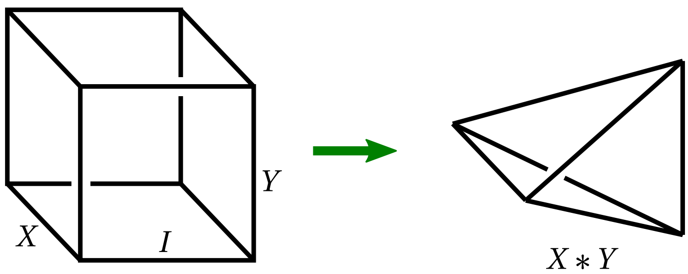

Suppose that X and Y are subspaces of some Euclidean space, and that X and Y are contained in respective affine subspaces that are affinely independent, meaning that the union of affinely independent pointsets, one in each subspace, is itself independent. Then, $X * Y$ is homeomorphic to the union of all line segments connecting points of X to points of Y. The points of this **geometric join** are convex combinations of the form $(1-t)x + ty$ with $(x,y,t) \in X \times Y \times I$. The formal combination $(1-t)x \oplus ty$ can also be used to describe points of the topological join if we consider that $0.x \oplus 1.y = 0 \oplus y$ is independent of x and $1.x \oplus 0.y = x \oplus 0$ is independent of y.

[9]It is also possible to quotient the deleted product by this action and to consider the usual cohomology on the quotient space.

When $Y = X$, beware that $x \oplus 0$ and $0 \oplus x$ represent points in disjoint copies of X. Formally, one should consider two distinct copies $X \times \{1\}$ and $X \times \{2\}$ of X and write $(x,1) \oplus 0$ and $0 \oplus (x,2)$. We however drop the second component for concision.

The join $\sigma * \tau$ of two simplices is a simplex of dimension $\dim\sigma + \dim\tau + 1$. If σ, τ are geometric simplices with affinely independent vertices, the vertices of their geometric join is the union of their vertices. In particular, a simplex with vertices $\ldots, p_i, \ldots$ may be written as $*_i p_i$. Considering abstract simplices as subsets of a ground set, the join operation of simplices thus corresponds to the union of subsets. More generally, the **join of two simplicial complexes** K, L, either geometric or abstract, is the simplicial complex

$$K * L = \{\sigma * \tau \mid \sigma \in K, \tau \in L\}$$

When $K = L$, we insist on the fact that the empty set is considered as a simplex in K and that for all $\sigma \in K$, the simplices $\sigma * \emptyset$ and $\emptyset * \sigma$ are distinct in $K * K$.

Exercise 2.3.1. Prove that the geometric join of two embedded subspaces is indeed homeomorphic to their topological join. Deduce that for simplicial complexes K and L the carrier of their join $|K * L|$ is homeomorphic to the join of their carriers $|K| * |L|$

Note that simplicial complexes behave well with respect to the join operation. This is less true for the product operation as it is not so immediate to obtain a simplicial decomposition of the product of two simplicial complexes.

The deleted join

The **deleted join** of a simplicial complexes K is the subcomplex of $K * K$ defined as

$$K *_\Delta K = \{\sigma * \tau \mid \sigma, \tau \in K, \sigma \cap \tau = \emptyset\}$$

More generally, if K and L are subcomplexes of a same complex we set

$$K *_\Delta L = \{\sigma * \tau \mid \sigma \in K, \tau \in L, \sigma \cap \tau = \emptyset\}$$

The deleted join can also be defined for a topological space X. Using the formal convex combination notation, we define

$$X *_\Delta X := X * X \setminus \{\frac{1}{2}x \oplus \frac{1}{2}x \mid x \in X\}$$

The simplicial and topological deleted join are closely related.

Proposition 2.3.2. *For any simplicial complex K, the space $|K *_\Delta K|$ is an equivariant deform retract of $|K| *_\Delta |K|$. In particular, both spaces have the same homotopy type.*

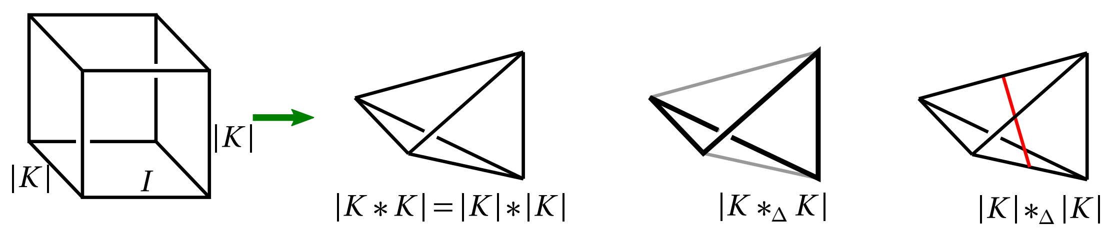

PROOF. Denote by $\rho : |K| *_\Delta |K| \to |K *_\Delta K|$ the retraction we are looking for. Geometrically, if $\Delta' := |K| * |K| \setminus |K| *_\Delta |K| = \{\frac{1}{2}x \oplus \frac{1}{2}x \mid x \in |K|\}$ is the *diagonal* of $|K| * |K|$, we shall define ρ so that it retracts every conic slice of the form $p * \Delta'$ with $p \in |K *_\Delta K|$ to p. The next figure illustrates the case where K is a 1-simplex.

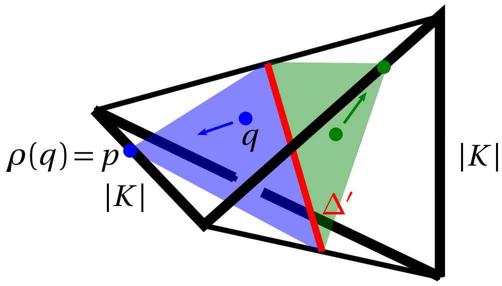

More formally, denote by Supp(x) the supporting simplex of a point $x \in |K|$. We have

$$|K *_\Delta K| = \{(1-t)x \oplus t\,y \mid x, y \in |K| \text{ and } \mathrm{Supp}(x) \cap \mathrm{Supp}(y) = \emptyset\}$$

We let $\rho((1-t)x \oplus t\,y) := (1-t')x' \oplus t\,y'$ where x', y', t' are defined as follows. Denote by V the set of vertices of K and by $(x_v)_{v\in V}$ and $(y_v)_{v\in V}$ the respective barycentric coordinates of x and y as defined in Section 2.1.1. We let x' and y' be the points with respective barycentric coordinates $(x'_v)_{v\in V}$ and $(y'_v)_{v\in V}$ satisfying

$$(1-t')x'_v = \max\{(1-t)x_v - t\,y_v, 0\}/S \quad \text{and} \quad t'y'_v = \max\{t\,y_v - (1-t)x_v, 0\}/S$$
$$\text{with } S = S_x + S_y, \quad S_x = \sum_{v\in V} x'_v, \quad S_y = \sum_{v\in V} y'_v \quad \text{and } t' = S_y/S$$

Note that the division by S is well-defined. Indeed, $S = 0$ implies $(1-t)x_v = t\,y_v$ for all $v \in V$, whence by summing over V, $t = 1/2$ and $x = y$. In turn, $(1-t)x \oplus t\,y = \frac{1}{2}x \oplus \frac{1}{2}x$ cannot be a point of $|K| *_\Delta |K|$, and S does not cancel on $|K| *_\Delta |K|$. Since x' and y' have disjoint support we have $(1-t)x' \oplus t\,y' \in |K *_\Delta K|$ as desired. Also, when x and y have disjoint support, we have $x' = x$, $y' = y$ and $t' = t$. It follows that ρ is the identity over $|K *_\Delta K|$. Moreover, the linear interpolation between ρ and the identity on $|K| *_\Delta |K|$ is a well-defined equivariant map at every interpolating parameter. (Refer to the next section for the notion of equivariance.) This concludes the proof of the lemma. □

2.3.2 The $\mathbb{Z}_2$-index

A **$\mathbb{Z}_2$-space** (X, α) is a space X together with an action of $\mathbb{Z}_2$ on it. Such a $\mathbb{Z}_2$-action is determined by the action of 1 which must be a continuous involution $\alpha : X \to X$. We may speak of the $\mathbb{Z}_2$-space X, omitting the involution when the $\mathbb{Z}_2$-action is implicitly clear. The $\mathbb{Z}_2$-action, or $\mathbb{Z}_2$-space, is **free** if α has no fixed point. The most important example of free $\mathbb{Z}_2$-space is given by the antipodality acting on the Euclidean sphere $\mathbb{S}^d$. Another basic examples are provided by squaring a space, as in $X \times X$ or $X * X$, and exchanging coordinates for the $\mathbb{Z}_2$-action. To be specific for $X * X$, this action is given by $(1-t)x \oplus t\,y \mapsto t\,y \oplus (1-t)x$, or equivalently by $(x, y, t) \mapsto (y, x, 1-t)$, recalling that $X * X$ is a quotient of $X \times Y \times I$. These actions are not free, but become free if we restrict the squared space to its deleted product or its deleted join.

The $\mathbb{Z}_2$-spaces form a category whose morphisms are called **$\mathbb{Z}_2$-maps** or **equivariant maps**. A continuous map $f:(X,\alpha)\to(Y,\beta)$ is equivariant if it commutes with the $\mathbb{Z}_2$-actions, *i.e.*, if the diagram

$$\begin{tikzcd} X \arrow[r, "f"] \arrow[d, "\alpha"] & Y \arrow[d, "\beta"'] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

is commutative. $\mathbb{Z}_2$-spaces have their simplicial counterpart where we ask that the spaces are simplicial complexes and the involved maps are simplicial. The above sphere example has a simplicial version. Consider the barycentric subdivision $\mathrm{sd}(\partial\sigma^{d+1})$ of the boundary of a $(d+1)$-simplex[10] $\sigma^{d+1}=2^{[d+2]}$. The vertices of $\mathrm{sd}(\partial\sigma^{d+1})$ are thus the proper subsets of $[d+2]$. Consider the **antipodal simplicial map** α_s on $\mathrm{sd}(\partial\sigma^{d+1})$ sending such a subset to its complement in $[d+2]$. Then $(\mathrm{sd}(\partial\sigma^{d+1}),\alpha_s)$ is a simplicial $\mathbb{Z}_2$-complex $\mathbb{Z}_2$ homeomorphic to $\mathbb{S}^d$ endowed with the antipodality.

Let us write $X\preceq_{\mathbb{Z}^2} Y$ if there exists an equivariant map between the $\mathbb{Z}_2$-spaces X and Y. It is easily seen that $\preceq_{\mathbb{Z}^2}$ is a reflexive and transitive relation on $\mathbb{Z}_2$-spaces. Define the **$\mathbb{Z}_2$-index**, $\mathrm{Ind}(X)$, of a $\mathbb{Z}_2$-space X as the minimum d such that $X\preceq_{\mathbb{Z}^2}\mathbb{S}^d$, *i.e.*, such that there exists an equivariant map $X\to\mathbb{S}^d$. We put $\mathrm{Ind}(X)=\infty$ is no such d exists. The transitivity of $\preceq_{\mathbb{Z}^2}$ directly implies that the $\mathbb{Z}_2$-index is non-decreasing for this relation.

Exercise 2.3.3. Show that any non free $\mathbb{Z}_2$-space is an upper bound for $\preceq_{\mathbb{Z}^2}$ and that its $\mathbb{Z}_2$-index is infinite.

Proposition 2.3.4. $\mathrm{Ind}(\mathbb{S}^d)=d$.

PROOF. We obviously have $\mathrm{Ind}(\mathbb{S}^d)\le d$ by reflexivity of $\preceq_{\mathbb{Z}^2}$. The other direction $\mathrm{Ind}(\mathbb{S}^d)\ge d$ is a direct consequence of the Borsuk-Ulam theorem. Indeed, one of the classical formulation of this theorem says that every continuous map $\mathbb{S}^d\to\mathbb{R}^d$ must send a pair of antipodal points to the same point. The existence of an equivariant map $\mathbb{S}^d\to\mathbb{S}^n$ with $n<d$ would however provide a map $\mathbb{S}^d\to\mathbb{S}^n\hookrightarrow\mathbb{R}^d$ without this property. □

Lemma 2.3.5. $\mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d)\le d$

PROOF. We just need to exhibit a continuous equivariant map $\mathbb{R}^d *_\Delta \mathbb{R}^d\to\mathbb{S}^d$. The map $\mathbb{R}^d\times\mathbb{R}^d\times I\to\mathbb{R}^{d+1}$, $(x,y,t)\mapsto(1-2t,(1-t)x-ty)$ is constant on each fiber of $\mathbb{R}^d\times\mathbb{R}^d\times I\to\mathbb{R}^d*\mathbb{R}^d$ and thus quotients to a map $\mathbb{R}^d*\mathbb{R}^d\to\mathbb{R}^{d+1}$. Moreover, the norm of this map never cancels on $\mathbb{R}^d *_\Delta \mathbb{R}^d\subset\mathbb{R}^d*\mathbb{R}^d$, so that the map

$$\mathbb{R}^d *_\Delta \mathbb{R}^d\to\mathbb{S}^d,\qquad (1-t)x\oplus ty\mapsto\frac{(1-2t,(1-t)x-ty)}{\|(1-2t,(1-t)x-ty)\|}$$

is well-defined. We easily check that it is equivariant. □

Exercise 2.3.11 in the next section asks you to strengthen Lemma 2.3.5 to show that $\mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d)=d$.

[10] As usual we write $[n]$ for $\{1,\dots,n\}$.

2.3.3 An obstruction to embedding

Suppose that $f : |K| \hookrightarrow \mathbb{R}^d$ is an embedding of a simplicial complex K. Then, the $\mathbb{Z}_2$-map $f * f : |K| * |K| \to \mathbb{R}^d * \mathbb{R}^d$, $(1-t)x \oplus ty \mapsto (1-t)f(x) \oplus tf(y)$ restricts to a $\mathbb{Z}_2$-map $|K *_\Delta K| \to \mathbb{R}^d *_\Delta \mathbb{R}^d$. It ensues that $\mathrm{Ind}(|K *_\Delta K|) \leq \mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d)$. In view of Lemma 2.3.5, we have

Proposition 2.3.6. *If* $\mathrm{Ind}(|K *_\Delta K|) > d$ *then* K *has no embedding in* $\mathbb{R}^d$.

In fact, the $\mathbb{Z}_2$-map $f * f : |K *_\Delta K| \to \mathbb{R}^d *_\Delta \mathbb{R}^d$ is well-defined as soon as $f(x) \neq f(y)$ for every $x, y \in |K|$ with disjoint support. It follows that the condition $\mathrm{Ind}(|K *_\Delta K|) > d$ implies that for every map $|K| \to \mathbb{R}^d$ there are two disjoint simplices whose images intersect.

We will now apply Proposition 2.3.6 to to prove that some d-dimensional complexes cannot be embedded into $\mathbb{R}^d$. Before that we introduce yet another simple operation on complexes.

The combinatorial Alexander dual. Let K be a proper subcomplex of the $n-1$ dimensional simplex $2^{[n]}$. (Note that every complex is a subcomplex of the simplex over its vertices.) Since K is a proper subcomplex it must be included in the boundary of the $(n-1)$-simplex. This boundary can be identified with a sphere and the intuition behind the Alexander dual is to take the complement of the antipodal image of K on the sphere. More precisely, the **Alexander dual** of K with respect to $2^{[n]}$ is the proper subcomplex of $2^{[n]}$ defined by

$$K^A = \{\sigma \in 2^{[n]} \mid [n] \setminus \sigma \notin K\}$$

In other words, K^A is composed of the simplices whose complements are not in K. The following exercise makes the above intuition more concrete. Here, by a subcomplex induced by a subset W of vertices we mean the set of simplices whose vertices fall in W.

Exercise 2.3.7. Let $V := K \setminus \{\emptyset\} \subset 2^{[n]}$ denote the set of vertices of $\mathrm{sd}\, K$. Show that $\mathrm{sd}\, K^A$ is the subcomplex of $\mathrm{sd}(\partial 2^{[n]})$ induced by the complement of $\alpha_s(V)$, where α_s is the antipodal simplicial map on $\mathrm{sd}(\partial 2^{[n]})$ sending a vertex $\sigma \in V$ to its complement $[n] \setminus \sigma$.

The proof of the following lemma is immediate from the definitions.

Lemma 2.3.8. *Let* $K \subset 2^{[2d+3]}$ *be the* d*-skeleton of the* $(2d+2)$*-simplex. Then* $K^A = K$.

Bier spheres. Given a proper subcomplex K of $2^{[n]}$, the **Bier sphere** of K with respect to n is

$$\mathrm{Bier}_n(K) = K *_\Delta K^A$$

Quite surprisingly the topology of the Bier sphere is independent of K. To see this we first subdivide $\mathrm{Bier}_n(K)$ using a subdivision process specific to subcomplexes of the join of complexes. Given the simplicial complexes K and L, the **shore subdivision** of a subcomplex $J \subset K * L$ is given by

$$\mathrm{ssd}\, J = \bigcup_{\sigma * \tau \in J} \mathrm{sd}\, \sigma * \mathrm{sd}\, \tau$$

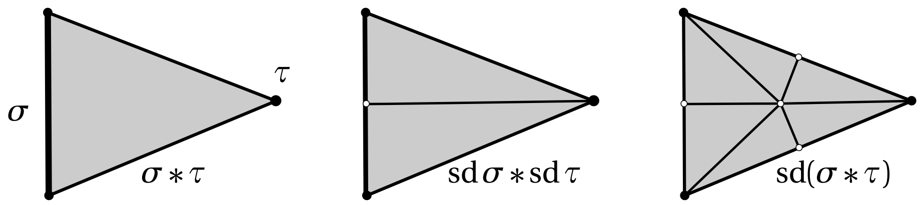

Comparison between the shore and barycentric subdivision of the 2-simplex expressed as the join of an edge and a vertex.

Being a subdivision, the shore of J has a carrier homeomorphic to $|J|$.

Proposition 2.3.9. $\mathrm{ssd}(\mathrm{Bier}_n(K))$ *is isomorphic to* $\mathrm{sd}(\partial 2^{[n]})$.

PROOF. For a simplex $\sigma \in \mathrm{sd}(\partial 2^{[n]})$, we write $\cup\sigma \subset [n]$ for the union of its vertices, viewed as proper subsets of $[n]$. Recall that α_s is the antipodal simplicial map on $\mathrm{sd}(\partial 2^{[n]})$. We have

$$\mathrm{ssd}(\mathrm{Bier}_n(K)) = \bigcup_{\sigma*\tau\in\mathrm{Bier}_n(K)} \mathrm{sd}\,\sigma * \mathrm{sd}\,\tau = \bigcup_{\substack{\sigma\in K,\,\tau\in K^A,\\ \sigma\cap\tau=\emptyset}} \mathrm{sd}\,\sigma * \mathrm{sd}\,\tau = \bigcup_{\substack{s\in\mathrm{sd}\,K,\ \alpha_s(\cup t)\notin K,\\ \cup s\subset\alpha_s(\cup t)}} s * t$$

To see the last equality, first note that $s \in \mathrm{sd}\,\sigma$ with $\sigma \in K$ is equivalent to $s \in \mathrm{sd}\,K$. Similarly, $t \in \mathrm{sd}\,\tau$ with $\tau \in K^A$ is equivalent to $t \in \mathrm{sd}\,K^A$. In turn, writing $\tau_0 \subset \cdots \subset \tau_\ell$ for the vertices of t, this means $[n]\setminus\tau_\ell \notin K$, *i.e.*, $\alpha_s(\cup t) \notin K$. Finally, writing $\sigma_0 \subset \cdots \subset \sigma_k$ for the vertices of s, the conditions $s \in \mathrm{sd}\,\sigma$, $t \in \mathrm{sd}\,\tau$ becomes $\sigma_k \in \sigma$ and $\tau_\ell \in \tau$. It follows that the condition $\sigma\cap\tau = \emptyset$ reduces to $\sigma_k \cap \tau_\ell = \emptyset$ which in turn can be written $\cup s \subset \alpha_s(\cup t)$.

We now consider the simplicial map $\varphi\colon \mathrm{ssd}(\mathrm{Bier}_n(K)) \to \mathrm{sd}(\partial 2^{[n]})$ sending the simplex $s * t$ to the simplex $s*\alpha_s(t)$. Equivalently, φ sends a vertex of the form $\sigma_0*\emptyset$ to itself and of the form $\emptyset*\tau_0$ to $\emptyset*\alpha_s(\tau_0)$. This map is well-defined since the condition $\cup s \subset \alpha_s(\cup t)$ implies that the vertices of $s*\alpha_s(t)$ form an increasing sequence of subsets of $[n]$, hence a simplex in $\mathrm{sd}(\partial 2^{[n]})$. φ is injective and it remains to see that it is surjective. For this, consider a simplex σ of $\mathrm{sd}(\partial 2^{[n]})$ with vertices $\sigma_0 \subset \cdots \subset \sigma_m$. Let k be the minimum index such that $\sigma_k \notin K$. Then, defining s as the simplex with vertices $\sigma_0 \subset \cdots \subset \sigma_{k-1}$ and defining t as the simplex with the remaining vertices of σ, we see that $\sigma = \varphi(s*\alpha_s(t))$. □

We are now ready to prove that

Theorem 2.3.10 (van Kampen - Flores)**.** *The d-skeleton of the $(2d+2)$-simplex does not embed in $\mathbb{R}^{2d}$.*

PROOF. Let $K \subset 2^{[2d+3]}$ be the d-skeleton of the $(2d+2)$-simplex. By Lemma 2.3.8 and Proposition 2.3.9, $\mathrm{Bier}_{2d+3}(K) = K *_\Delta K$ is isomorphic to the $(2d+1)$-sphere. From Proposition 2.3.4, we have $\mathrm{Ind}(K *_\Delta K) = 2d+1$ and we conclude by invoking Proposition 2.3.6. □

Exercise 2.3.11. Consider[11] the d-simplex $\sigma = 2^{[d+1]}$ as a subcomplex of $2^{[d+2]}$. What is the Alexander dual of σ? Deduce that $\mathrm{Ind}(\sigma *_\Delta \sigma) = d$. Conclude that $\mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d) \geq d$.

[11]This exercise was suggested by Axel Péneau.

Exercise 2.3.12. Consider the simplicial complex represented in the figure below.

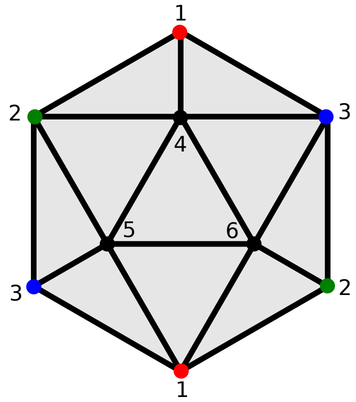

It is composed of 6 vertices and 10 triangles forming a disk, the boundary edges of which should be identified according to the boundary vertex numbering. This complex is topologically a projective plane. Mimic the proof of the van Kampen - Flores theorem to prove that the projective plane cannot be embedded into $\mathbb{R}^3$.

3

Embedding Linearly

Contents

We already saw that every m-dimensional complex embeds linearly into $\mathbb{R}^{2m+1}$. What about the existence of linear embeddings into $\mathbb{R}^d$ with $d \leq 2m$? It turns out that independently of obstruction theories, like Whitney or van Kampen obstructions, this question is decidable. We first look at a combinatorial approach based on the notion of chirotope for a point configuration.

3.1 Affine point configurations

Given a set $\{p_1, \ldots, p_n\}$ of n points in $\mathbb{R}^d$, its **chirotope** is the map $\{1, \ldots, n\}^{d+1} \to \{-1, 0, 1\}$ defined by

$$(i_0, \ldots, i_d) \mapsto \operatorname{sign}(\det \begin{pmatrix} 1 & \cdots & 1 \\ p_{i_0} & \cdots & p_{i_d} \end{pmatrix}) \tag{3.1}$$

In other words, the chirotope returns for every $(d+1)$-tuple of points the orientation of the d-simplex defined by those points. Here, the orientation is assumed to be zero if the points are affinely dependent. Intuitively, the chirotope records the relative positions of the points in a point configuration. For instance, it is easily seen that the chirotope determines the combinatorial structure of the convex hull of a point configuration. Not every map $\{1,\dots,n\}^{d+1} \to \{-1,0,1\}$ can be the chirotope of a point configuration. The map has to satisfy certain conditions related to the Grassman-Plücker relations and to Radon partitions.

3.1.1 Grassman-Plücker relations

Let $V = (v_1,\dots,v_n)$ be a family of n vectors in $\mathbb{R}^d$. As usual, put $[n] := \{1,\dots,n\}$. For a sequence $I = (i_1,\dots,i_d)$ of d indices in $[n]$, we denote by

$$m_I := \det(v_{i_1},\dots,v_{i_d})$$

the determinant with respect to the canonical basis of $\mathbb{R}^d$ of the d-tuple of vectors of V indexed by I. The *minors* $(m_I)_{I\in\binom{[n]}{d}}$, where $\binom{[n]}{d} \subset [n]^d$ denotes the set of all increasing sequences of d indices in $[n]$, are the **homogeneous Plücker coordinates** associated to V.

Theorem 3.1.1. *The homogeneous Plücker coordinates associated to V satisfy the* ***Grassman-Plücker relations****:*

$$\forall I \in \binom{[n]}{d+1}, \forall J \in \binom{[n]}{d-1} : \sum_{s=0}^{d} (-1)^s m_{I-i_s} m_{J+i_s} = 0 \tag{3.2}$$

where $I - i_s$ is obtained by deleting i_s in $I = (i_0, i_1, \dots, i_{d-1})$, and $J + i_s$ is obtained by appending i_s at the end of J. Note that $J + i_s$ is not necessarily increasing and that m_{J+i_s} cancels whenever $i_s \in J$.

Proof. For $J = (j_0, j_1, \dots, j_{d-2-1})$ fixed, consider the $(d+1)$-linear map $f : (\mathbb{R}^d)^{n+1} \to \mathbb{R}$ given by

$$(u_0,\dots,u_d) \mapsto \sum_{s=0}^{d} (-1)^s \det(u_0,\dots,\widehat{u_s},\dots,u_d)\det(v_{j_0},\dots,v_{j_{d-2}},u_s)$$

We easily check that f is alternating. However, an alternating $(d+1)$-linear map over a d dimensional space must be zero. In particular, $f(v_{i_0},\dots,v_{i_d}) = 0$, which is precisely Equation (3.2). □

Exercise 3.1.2. Prove that the map f in the above proof is indeed alternating.

Cultural note: The homogeneous Plücker coordinates provide an embedding of the **Grassmannian** $\mathbf{Gr}(d,\mathbb{R}^n)$ into the projective space $\mathbb{P}\left(\bigwedge^d \mathbb{R}^n\right)$ of the d-fold exterior product $\bigwedge^d \mathbb{R}^n$ of $\mathbb{R}^n$. Let us briefly explain why. Recall that $\mathbf{Gr}(d,\mathbb{R}^n)$ is the set of

d dimensional subspaces of $\mathbb{R}^n$. The exterior (or Grassmann) algebra $\bigwedge\mathbb{R}^n$ can be defined as the quotient of the tensor algebra $\bigotimes\mathbb{R}^n$ by the two-sided ideal generated by the tensor products $\{v\otimes v\}_{v\in\mathbb{R}^n}$. The exterior product $\wedge$ thus induced by the tensor product is antisymmetric as can be seen by expanding $(x+y)\otimes(x+y)$. As a vector space, the d-fold exterior product $\bigwedge^d\mathbb{R}^n$ has a basis composed of the d-vectors $e_I = e_{i_1}\wedge\cdots\wedge e_{i_d}$, where $I=(i_1,\dots,i_d)\in\binom{[n]}{d}$ and $(e_1,\dots,e_n)$ is the canonical basis of $\mathbb{R}^n$. Consider a family $W=(w_1,\dots,w_d)$ of d vectors in $\mathbb{R}^n$ as a $d\times n$ matrix whose columns are the components of the v_i expressed in the canonical basis. Viewing the transpose matrix W^t as a family $V=(v_1,\dots,v_n)$ of n vectors in $\mathbb{R}^d$, we compute $w_1\wedge\cdots\wedge w_d=\sum_I m_I e_I$, where the m_I are the homogeneous Plücker coordinates associated to V. One can show that two families of d independent vectors have proportional wedge products if and only if they span the same vector space, whence the claimed embedding. In fact, the Grassman-Plücker relations (3.2) are a necessary and sufficient condition on the m_I to come from a wedge product of d vectors, the so-called *decomposable* d-vectors. The Grassmannian $\mathbf{Gr}(d,\mathbb{R}^n)$ is thus embedded in $\mathbb{P}\left(\bigwedge^d\mathbb{R}^n\right)$ as a projective algebraic variety determined by quadratic equations.

3.1.2 Radon partitions

Let $\mathscr{P}$ be a set of points in $\mathbb{R}^d$. Any partition $\mathscr{P}=\mathscr{P}'\cup\mathscr{P}''$ such that the convex hulls $\operatorname{Conv}\mathscr{P}'$ and $\operatorname{Conv}\mathscr{P}''$ have a nonempty intersection is called a **Radon partition** of $\mathscr{P}$. Recall that $\mathscr{P}$ is in **general position** if no affine hyperplane contains more than d points of $\mathscr{P}$.

Lemma 3.1.3. *Let $\mathscr{P}=\mathscr{P}'\cup\mathscr{P}''$ be a partition of a set of $d+1$ points in general position in $\mathbb{R}^d$. Then $\operatorname{Conv}\mathscr{P}'$ and $\operatorname{Conv}\mathscr{P}''$ are disjoint.*

Note that the lemma just says that two faces of a d-simplex with disjoint vertex sets are indeed disjoint.

PROOF. Write $\mathscr{P}=\{p_i\}_{i\in I}$, $\mathscr{P}'=\{p_i\}_{i\in I'}$ and $\mathscr{P}''=\{p_i\}_{k\in I''}$ with $I=I'\cup I''$. Suppose by way of contradiction that $\operatorname{Conv}\mathscr{P}'\cap\operatorname{Conv}\mathscr{P}''$ contains a point p. Then we can write p as two convex combinations $\sum_{i\in I'}\alpha_i p_i$ and $\sum_{i\in I''}\alpha_i p_i$ with $\sum_{i\in I'}\alpha_i=\sum_{i\in I''}\alpha_i=1$. It follows that $\sum_{i\in I'}\alpha_i p_i-\sum_{i\in I''}\alpha_i p_i=0$. This provides an affine dependency between the points of $\mathscr{P}$ in contradiction with the general position assumption. $\square$

Theorem 3.1.4 (Radon, 1921)**.** *Any set of $d+2$ points in $\mathbb{R}^d$ admits a Radon partition. Moreover, if the $d+2$ points are in general position any two of them are in the same part if and only if they are separated by the hyperplane spanned by the remaining d points. In particular, the Radon partition is unique.*

PROOF. Any $d+2$ points, say $\mathscr{P}=\{p_1,p_2,\dots,p_{d+2}\}$, must be affinely dependent in $\mathbb{R}^d$. We can thus find real numbers $\alpha_1,\alpha_2,\dots,\alpha_{d+2}$, not all zero, such that $\sum_{i=1}^{d+2}\alpha_i=0$ and $\sum_{i=1}^{d+2}\alpha_i p_i=0$. Let $I_+ := \{i\in[d+2]\mid \alpha_i\geq 0\}$ and $I_- := \{i\in[d+2]\mid\alpha_i<0\}$. Then, $\sum_{I_+}\alpha_i=\sum_{I_-}-\alpha_i$ and denoting the common sum by A we derive the two convex combinations $\sum_{I_+}(\alpha_i/A)p_i = \sum_{I_-}(-\alpha_i/A)p_i$. It follows that $\mathcal{P} = \{p_i\}_{i\in I_+} \cup \{p_i\}_{i\in I_-}$ is a Radon partition of $\{p_i\}_{i\in[d+2]}$.

Suppose that $\mathcal{P}$ is in general position and consider a Radon partition $\mathcal{P} = \mathcal{P}' \cup \mathcal{P}''$. Let $p, q \in \mathcal{P}$ and let H be the affine hull of the remaining points $\mathcal{P} \setminus \{p, q\}$. By general position, H is a hyperplane that does not contain p nor q. By Lemma 3.1.3 applied to $\mathcal{P} \setminus \{p, q\}$ in H, the convex hulls Conv $(\mathcal{P}' \setminus \{p, q\})$ and Conv $(\mathcal{P}'' \setminus \{p, q\})$ are disjoint. If p and q are in a same part, then they must be separated by H. Otherwise, Conv $\mathcal{P}'$ and Conv $\mathcal{P}''$ would also be disjoint, contradicting that $\mathcal{P}' \cup \mathcal{P}''$ is a Radon partition. Conversely, if $p \in \mathcal{P}'$ and $q \in \mathcal{P}''$, then they must lie on the same side of H since otherwise Conv $\mathcal{P}'$ and Conv $\mathcal{P}''$ would be disjoint, again contradicting that $\mathcal{P}' \cup \mathcal{P}''$ is a Radon partition. □

Corollary 3.1.5. *Let $\mathcal{P}$ be a set of $d+2$ points in $\mathbb{R}^d$, not all on a same hyperplane. There exists a hyperplane H that contains d of the points in $\mathcal{P}$ and such that the two remaining points are on the same side of H, i.e. contained in the same component of $\mathbb{R}^d \setminus H$.*

PROOF. By induction on the dimension d. The base case $d = 1$ is trivial and left to the reader. If $d > 1$, first suppose that $\mathcal{P}$ is in general position. By the previous theorem, $\mathcal{P}$ has a (unique) Radon partition. Choose one point in each part and take for H the affine hull of the remaining points. Then H has the required properties by the same previous theorem.

If $\mathcal{P}$ is not in general position, there must be a hyperplane K that contains a subset $\mathcal{Q}$ of $d+1$ points of $\mathcal{P}$. Let p be the remaining point in $\mathcal{P} \setminus \mathcal{Q}$. Note that the points in $\mathcal{Q}$ cannot lie on a same $(d-1)$-plane. For otherwise, $\mathcal{P}$ would be contained in a hyperplane. By induction applied to $\mathcal{Q}$ in K, there is a $(d-1)$-plane L in K that contains $d-1$ of the points in $\mathcal{Q}$ such that the two remaining points of $\mathcal{Q}$ are on the same side of L. Taking for H the affine hull of $L \cup \{p\}$, we obtain a hyperplane with the required properties. □

3.1.3 From chirotopes to oriented matroids

Let $\mathcal{P} = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$. Recall that its chirotope χ returns for every $(d+1)$-tuple $I \in [n]^{d+1}$ the orientation $\chi(I) \in \{-1, 0, 1\}$ of the d-simplex spanned by the vertices of $\mathcal{P}$ indexed by I. The fact that $\mathcal{P}$ is a subset of a d dimensional affine space imposes some relations between the signs of its chirotope.

Theorem 3.1.6. *The chirotope χ of a set of n points in $\mathbb{R}^d$ is alternating and satisfies the following conditions.*

- **C-GP:** *For all $I = (i_0, i_1, \ldots, i_{d+1-1}) \in \binom{[n]}{d+2}$ and $J \in \binom{[n]}{d}$ the set of signs*

$$\{(-1)^s \chi(I - i_s)\chi(J + i_s)\}_{s=0,\ldots,d+1}$$

  either contains $\{-1, 1\}$, or is reduced to $\{0\}$.

- **C-R:** *For all* $I = (i_0, i_1, \dots, i_{d+1-1}) \in \binom{[n]}{d+2}$ *the set of signs*

$$\{(-1)^s \chi(I - i_s)\}_{s=0,\dots,d+1}$$

either contains $\{-1, 1\}$, *or is reduced to* $\{0\}$.

Proof. Let $\mathscr{P} = \{p_1, \dots, p_n\}$ be a set of n points in $\mathbb{R}^d$. The definition of their chirotope as the sign of a determinant shows that it is indeed alternating. Put $v_i = \begin{pmatrix} 1 \\ p_i \end{pmatrix} \in \mathbb{R}^{d+1}$ and let $V = (v_1, \dots, v_n)$. From the very definitions we see that the chirotope χ of $\mathscr{P}$ coincides with the signs of the homogeneous Plücker coordinates $(m_K)_{K \in \binom{[n]}{d+1}}$ of V:

$$\forall K \in \binom{[n]}{d+1} : \chi(K) = \operatorname{sign}(m_K)$$

The Grassman-Plücker relations (3.2) in Theorem 3.1.1 implies that either all terms in $\sum_{s=0}^{d+1} (-1)^s m_{I-i_s} m_{J+i_s}$ are zero or two terms are non-zero with opposite signs. Condition **C-GP** follows.

For Condition **C-R**, we first remark that when $\mathscr{P}_I = \{p_{i_0}, \dots, p_{i_{d+1}}\}$ is contained in a hyperplane, then the chirotope cancels on all $(d+1)$-tuples of indices in I and thus satisfies **C-R**. Otherwise, we may apply Corollary 3.1.5 to find two points p_{i_j}, p_{i_k} in $\mathscr{P}_I$ such that

$$\det(v_{i_j}, v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}}) = \det(v_{i_k}, v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}}) \tag{3.3}$$

and this quantity is nonzero. By the alternating property of the determinant we have

$$\det(v_{i_j}, v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}}) = (-1)^j \det(v_{i_0}, \dots, v_{i_j}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}})$$

and

$$\det(v_{i_k}, v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}}) = (-1)^{k-1} \det(v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, v_{i_k}, \dots, v_{i_{d+1}})$$

reporting in (3.3), we get that

$$(-1)^j \det(v_{i_0}, \dots, v_{i_j}, \dots, \widehat{v_{i_k}}, \dots, v_{i_{d+1}}) = -(-1)^k \det(v_{i_0}, \dots, \widehat{v_{i_j}}, \dots, v_{i_k}, \dots, v_y\, i_{d+1})$$

It ensues that $(-1)^j \chi(I - i_j)$ and $(-1)^k \chi(I - i_k)$ have opposite signs and are both nonzero, so that **C-R** holds in all cases. □

The pair $([n], \chi)$, where $\chi : [n]^{d+1} \to \{-1, 0, 1\}$ is an alternating map satisfying the condition of Theorem 3.1.6, is called an **affine oriented matroid** of rank $d+1$. χ is the chirotope of this oriented matroid. Any set of points in $\mathbb{R}^d$ whose chirotope coincides with χ is a *realization* of χ.

3.2 Linear embeddings and immersions

Recall that a linear mapping of a simplicial complex K into $\mathbb{R}^d$ is entirely determined by the image of the vertices of K. It is an embedding if it induces an injective map $|K| \hookrightarrow \mathbb{R}^d$. For an **immersion** we only require that this map is locally injective, which amounts to ask that the restriction of the map to the star of each vertex is injective. Here, the **star** of a vertex of K is the subcomplex comprising all the simplices containing that vertex and all their faces.

Lemma 3.2.1. *A linear map $f : K \to \mathbb{R}^d$ is an embedding if and only if every pair of disjoint simplices in K is sent to disjoint simplices in $\mathbb{R}^d$. It is an immersion if and only if the previous condition holds locally, i.e., f sends every pair of disjoint simplices in the star of a vertex to disjoint simplices in $\mathbb{R}^d$.*

Proof. The conditions in the lemma are trivially necessary. Suppose that the condition for f to be an embedding holds. We first claim that the restriction of f to each simplex $[v_0,\dots,v_k] \in K$ is injective. Otherwise, $f(v_0),\dots,f(v_k)$ must span a flat (affine subspace) of dimension at most $k-1$. By Radon's theorem 3.1.4 we can partition the $f(v_i)$ in two subsets whose convex hulls intersect. The corresponding subsets of v_i define two disjoint faces of $[v_0,\dots,v_k]$ whose images have a common intersection. This is however in contradiction with the embedding condition in the lemma.

Now, by way of contradiction, consider two points $x \neq y$ in $|K|$ such that $f(x) = f(y)$. Let $\sigma, \tau \in K$ be the supporting simplices of x and y, respectively. By the previous claim and the embedding condition, σ and τ must have a common face different from both σ and τ. Let $\{u_i\}_{i\in I}$ be the vertices of that face, and let $\{v_j\}_{j\in J}$ and $\{w_k\}_{k\in K}$ be the remaining vertices of σ and τ, respectively. We have $x = \sum_I \alpha_i u_i + \sum_J \beta_j v_j$ and $y = \sum_I \alpha'_i u_i + \sum_K \gamma_k w_k$ for some positive coefficients $\alpha_i, \beta_j, \alpha'_i, \gamma_k$ with $\sum_I \alpha_i + \sum_J \beta_j = \sum_I \alpha'_i + \sum_K \gamma_k = 1$. Set $I_+ = \{i \in I \mid \alpha_i > \alpha'_i\}$ and $I_- = \{i \in I \mid \alpha_i < \alpha'_i\}$. We deduce from $f(x) = f(y)$ that $\sum_{I_+}(\alpha_i - \alpha'_i) f(u_i) + \sum_J \beta_j f(v_j) = \sum_{I_-}(\alpha'_i - \alpha_i) f(u_i) + \sum_K \gamma_k f(w_k)$. Remarking that $\sum_{I_+}(\alpha_i - \alpha'_i) + \sum_J \beta_j = \sum_{I_-}(\alpha'_i - \alpha_i) + \sum_K \gamma_k$ and denoting by A the common positive sum, we obtain after dividing by A two convex combinations of $\{u_i\}_{i\in I_+} \cup \{v_j\}_{j\in J}$ on one side and of $\{u_i\}_{i\in I_-} \cup \{w_k\}_{k\in K}$ on the other side whose image by f coincide. This again contradicts the embedding condition. It follows that the linear extension of f is indeed injective. The second part of the lemma is proved similarly, working separately in the star of each vertex. □

Lemma 3.2.2. *If a simplicial complex K has a linear embedding into $\mathbb{R}^d$, then it has a linear embedding sending the vertices to a pointset in general position in $\mathbb{R}^d$. The same holds, replacing embedding by immersion. Moreover, one may enforce that the image vertices have rational coordinates.*

Proof. By the previous lemma, being an embedding or an immersion is ensured by a finite set of open conditions, namely the existence of a separating hyperplane for the images of pairs of disjoint simplices. It ensues that any sufficiently small perturbation of the vertex images preserves the property of being an embedding or an immersion. In particular, one may require that the image vertices are in general position and that all their coordinates are rational. □

3.2.1 A certificate of non-embeddability

Suppose that a simplicial complex K has a linear embedding $f : K \to \mathbb{R}^d$. Let $V = \{v_i\}_{i\in I}$ be the vertices of K. By Lemma 3.2.2, we can assume that $f(V)$ is in general position. In other words, the chirotope $\chi : I^{d+1} \to \{-1,0,1\}$ of $f(V)$ does not cancel on $\binom{I}{d+1}$. An oriented matroid with such a chirotope is said **uniform**. We shall also say that the chirotope itself is uniform. Lemma 3.2.1 provides a simple criterion for f to be an embedding. This criterion turns out to be encoded in the chirotope of $f(V)$ as stated in the next Corollary 3.2.4.

Lemma 3.2.3. *Let σ, τ be two intersecting simplices in $\mathbb{R}^d$ such that $\dim\sigma + \dim\tau > d$. Then, we can find a face of σ and a face of τ that intersect and whose dimensions add up to exactly d.*

PROOF. We first make two simple observations.

1. Let H be a flat intersecting a set S in a Euclidean space. Then, the boundary points of $H \cap S$ in H (it is all of $H \cap S$ if its interior in H is empty) are contained in the boundary of S.

2. If two sets intersect in a Euclidean space, then one of the two intersects the boundary of the other one.

Let $k = \dim\sigma$ and $\ell = \dim\tau$. We prove the lemma by induction on $k+\ell$. Denote by H the intersection of the affine hulls of σ and τ. Then, $\sigma \cap H$ and $\tau \cap H$ are two intersecting convexes in H. If one of them, say $\sigma \cap H$, has empty interior in H, then by observation (1) applied in the affine hull of σ, it is included in the boundary of σ. It follows that a proper face σ' of σ intersects τ. Replacing σ' by a larger face of σ if necessary, we may assume that $\dim\sigma + \ell > \dim\sigma' + \ell \geq d$. We can thus invoke the induction to conclude. If both $\sigma \cap H$ and $\tau \cap H$ have nonempty interior in H, then their intersection contains a boundary point of one of them, say $\sigma \cap H$, by observation (2). By observation (1) this boundary point is also in the boundary of σ and we may conclude as in the previous case. □

Corollary 3.2.4. *Let K be a simplicial complex of dimension at most d with vertex set $[n]$. Consider a map $f : [n] \to \mathbb{R}^d$ such that $f([n])$ is in general position and denote its chirotope by $\chi : [n]^{d+1} \to \{-1,0,1\}$. Then f linearly extends to an embedding $f : |K| \to \mathbb{R}^d$ if and only if the following condition is satisfied.*

- **C-E:** *for all $I \in \binom{[n]}{d+2}$, the subsets*

$$I_+ := \{i \in I \mid (-1)^i \chi(I-i)\} = 1 \quad \textit{and} \quad I_- := \{i \in I \mid (-1)^i \chi(I-i) = -1\}$$

  are not the vertex sets of a pair of simplices in K.

A similar condition **C-I** *characterizes immersions, where we only ask that I^+, I^- are not the vertex sets of a pair of simplices in the star of some vertex in K.*

PROOF. From Radon's theorem 3.1.4 and looking at the proof of Condition **C-R** in Theorem 3.1.6, it is easily seen that $f(I_+)Uf(I_-)$ defines the unique Radon partition of $f(I)$. In particular, Conv $f(I_+)$ and Conv $f(I_-)$ intersect. Condition **C-E** is thus necessary for the extension of f to be an embedding. Conversely, assume that **C-E** holds. Consider two disjoint simplices $\sigma, \tau \in K$. If $\dim\sigma + \dim\tau < d$ then $f(\sigma)$ and $f(\tau)$ are disjoint by the general position hypothesis. If $\dim\sigma + \dim\tau \geq d$ we also claim that $f(\sigma)$ and $f(\tau)$ are disjoint. Otherwise, by Lemma 3.2.3 we can assume that $\dim\sigma+\dim\tau = d$. Let I be the concatenation of the vertices of σ and τ. Then, $I \in \binom{[n]}{d+2}$ and the uniqueness of the Radon partition for $f(I)$ implies that $\{I_+, I_-\} = \{\sigma, \tau\}$. This would however be in contradiction with condition **C-E**. It follows that every pair of disjoint simplices in K is sent by f to disjoint simplices in R^d. Lemma 3.2.1 implies that f indeed defines an embedding. A similar proof holds for Condition **C-I** on immersions. □

The previous theorem, together with Lemma 3.2.2 and Theorem 3.1.6 have the following consequence. If K has a linear embedding in $\mathbb{R}^d$, then there should exist a uniform chirotope *admissible* for the embedding of K, *i.e.*, satisfying conditions **C-GP**, **C-R** and **C-E**. The existence of an admissible chirotope is purely combinatorial and only depends on d and K. It can thus be checked by a computer. If no admissible chirotope is found then we can claim that K has no linear embedding in $\mathbb{R}^d$. A brute force algorithm would try all maps $\binom{[n]}{d+1} \to \{-1, 1\}$ to see if one satisfies conditions **C-GP**, **C-R** and **C-E**. The number of possible maps, $2^{\binom{n}{d+1}}$, is already far too large, not to mention the tests for conditions **C-GP**, **C-R** and **C-E**, to be tractable in practice, except for very small complexes.

3.2.2 Linear embedding of surfaces

A finite simplicial surface is a simplicial complex S whose carrier $|S|$ is a compact two dimensional manifold. Equivalently, every simplex of S should be a face of a triangle in $|S|$ and every edge should be a face of at most two triangles. One says that S *triangulates* $|S|$, or is a *triangulation* of $|S|$. Recall that every simplicial surface embeds linearly in $\mathbb{R}^5$. It follows from their classification that all orientable surfaces can be obtained from the connected sum of a sphere, possibly with boundary, with a certain number of tori. In particular, all orientable surfaces have a topological embedding into $\mathbb{R}^3$. In fact, the method of Burago and Zalgaller described in the first lecture shows that all orientable surfaces have a PL embedding in $\mathbb{R}^3$. The answer becomes less trivial if one asks for the linear embedding into $\mathbb{R}^3$ of a specific triangulation of a surface. Until a counterexample was found in 2000, it was not known whether all simplicial surfaces could be linearly embedded in 3-space. Here are some known facts.

- It follows from a celebrated theorem of Steinitz (1922) that all triangulations of a sphere have a linear embedding into $\mathbb{R}^3$. In fact, each such triangulation is the boundary complex of a convex polyhedron in $\mathbb{R}^3$. See [Zie95, Chap. 4] for a proof.

- Archdeacon et al. [ABEM07] proved that all triangulations of the torus can be linearly embedded into $\mathbb{R}^3$. In particular, the toroidal triangulation with the smallest number of vertices, the so-called **Möbius torus**, has many linear embeddings. The 1-skeleton of this triangulation is the complete graph K_7 on 7 vertices. The first known linear embedding of the Möbius torus, due to Császár (1949), is shown Figure 3.1. Bokowski and Eggert [BE91] have listed all the 72

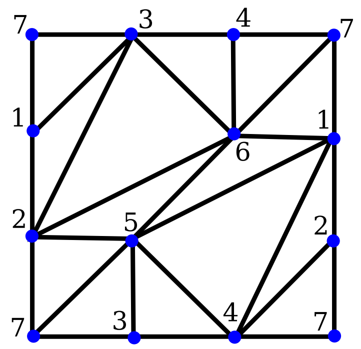

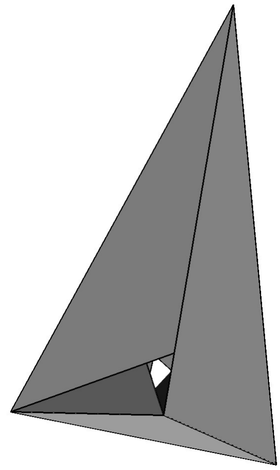

Figure 3.1: Left, layout of the Möbius torus. Right, Császár's linear embedding. The vertex coordinates are, in order :
$(3,-3,0),(-3,3,0),(-3,-3,1),(3,3,1),(-1,-2,3),(1,2,3),(0,0,15)$

admissible uniform chirotopes of the Möbius torus (up to an automorphism of the triangulation) and they were able to exhibit realizations for each of them.

- For higher genus, there exists simplicial orientable surfaces without any linear embedding in $\mathbb{R}^3$. All the results in this direction were proved with the help of a computer to check that some specific triangulation had no admissible chirotope. For instance, Altshuler et al. [ABS96] proved that a certain simplicial surface of genus 6 with 12 vertices has no admissible chirotope. Using a more efficient heuristic to explore the set of chirotopes Schewe [Sch10] proved that none of the 59 genus 6 triangulations with 12 vertices has an admissible chirotope. He proved a similar result for a triangulation of genus 5 with one triangle removed. As a consequence, any triangulation obtained from a connected sum along this triangle cannot be realized into $\mathbb{R}^3$. Similar nonrealizability results were obtained only asking for immersions rather than embeddings.

3.3 Deciding linear embeddability

The preceding approach, based on chirotopes, does not always allow to decide when a simplicial complex K is linearly embeddable in some $\mathbb{E}^d$. Even if K has an admissible chirotope, we still have to exhibit an actual embedding, or prove that no such embedding exists in order to conclude. The conditions for this existence happens to be dictated by a set of polynomials inequalities. Indeed, assuming that K has an admissible chirotope χ all what we need to find is a set of points in $\mathbb{E}^d$, one for each vertex of K, such that the corresponding chirotope is equal to χ. Now, the chirotope of the set of points is given by sign conditions on determinants (see (3.1)) which are polynomials in the coordinates of the points.

In fact, it is not necessary to know in advance an admissible chirotope to express that K has a linear embedding. By Lemma 3.2.1, it is equivalent to look for a set of points $\{p_1, p_2, \dots, p_n\}$, corresponding to the vertices $i \in [n]$ of K, such that every pair of disjoint simplices $([i_0, i_1, \dots, i_{k-1}], [j_0, j_1, \dots, j_{\ell-1}])$ in K is sent to non-intersecting simplices in $\mathbb{E}^d$. This condition can be rephrased as the existence of a hyperplane separating $[p_{i_0}, \dots, p_{i_k}]$ and $[p_{j_0}, \dots, p_{j_\ell}]$. In other words, there should exist coefficients $c_0, c_1, \dots, c_{d-1}$ such that the hyperplane equation $c_0 + \sum_{i=1}^d c_i x_i$ evaluates positively on $p_{i_0}, \dots, p_{i_k}$ and negatively on $p_{j_0}, \dots, p_{j_\ell}$. Hence, by introducing new variables c_i, we are again reduced to the satisfiability of a set of polynomials inequalities.

A subset of $\mathbb{R}^d$ defined by polynomials inequalities is said **real semi-algebraic**. Deciding linear embeddability thus reduces to decide whether a real semi-algebraic set is nonempty. Decision problems that reduce (in a sense to be defined) to the (non)vacuity of a real semi-algebraic[1] set are known as decision problems for the **existential theory of the reals**. The existential theory of the reals thus defines a complexity class that turns out to lie somewhere between the classes **NP** and **PSPACE**. In particular, the existential theory of the reals is decidable. In order to make sense out of these claims we need to recall some basic definitions from the theory of computation.

3.3.1 Turing machines and complexity

This section is intended to be a crash introduction to computational complexity. The following notes are greatly inspired by Avi Widgerson [Wig06].

Turing machines

The most popular model of computation was introduced by Alan Turing in 1936. It was proved equivalent to other notions of computation such as recursive functions or λ-calculus. Formally, a Turing machine is a triple $(\mathscr{A}, \mathscr{Q}, \mathscr{T})$, where $\mathscr{A}$ is a finite alphabet including a special **blank** character denoted by $\emptyset$, $\mathscr{Q}$ is a finite set of **states**, and[2] $\mathscr{T} \subset \mathscr{A} \times \mathscr{Q} \times \mathscr{A} \times \mathscr{Q} \times \{R, L\}$ is a **transition table** specifying how the machine operates on **configurations**. Those are words of the form $uqv \in \mathscr{A}^* \times \mathscr{Q} \times \mathscr{A}^*$, where $\mathscr{A}^*$ denotes the set of words (*i.e.*, finite sequences) over $\mathscr{A}$. Intuitively, the machine can be represented by a linear **tape** composed of a bi-infinite sequence of **cells** that each contains one alphabet symbol, and by a read/write head pointing to one cell and containing the machine state. Configuration uqv then corresponds to a tape marked with the word uv and otherwise with blanks and whose read/write head points to the first letter in v (the empty word is interpreted as a blank). Transition $aqbpD \in \mathscr{T}$ applies to any configuration uqv such that a is the first letter in v. It transforms uqv replacing a with b, the state q by p, and moves the head one step to the left or right according to whether D equals L or R, respectively.

[1] Here, we are only interested in systems of polynomials with integer (equivalently, rational) coefficients.

[2] In these notes, we use the symbol $\subset$ to indicate the subset relation, not necessarily proper.

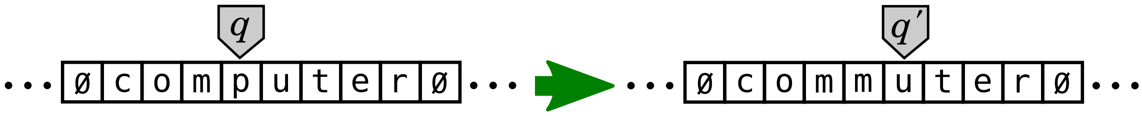

Figure 3.2: Illustration of the transition $\mathtt{p}q\mathtt{m}q'R$ applied to configuration $\mathtt{com}q\mathtt{puter}$ on a Turing machine operating on the Latin alphabet.

A Turing machine is **deterministic** if at most one transition applies to a given configuration: $aqbpD \in \mathscr{T}$ and $aqb'p'D' \in \mathscr{T}$ implies $b' = b$, $p' = p$ and $D' = D$. The machine is **halting** in a given configuration when no transition applies. Usually, a Turing machine has two special halting states interpreted as *accepting* and *rejecting*. As opposed to a deterministic machine, a nondeterministic Turing machine may lead to several computations starting from a same configuration.

Complexity classes

In computer science a **decision problem** refers to a subset of words over a fixed alphabet $\mathscr{A}$. Words in the subset are the *YES instances* of the problem. Intuitively, the YES instances correspond to the encoding of objects – such as numbers, graphs, or Boolean formulas – satisfying a certain property. For instance, one may consider the problem of primality testing where the YES instances are the binary encoding of prime integers over the alphabet $\mathscr{A} = \{0, 1\}$. In full generality, a decision problem can be any subset $I \subset \mathscr{A}^*$. Such a subset is also called a **language**. A Turing machine is said to solve or decide[3] problem I if given any word $w \in \mathscr{A}^*$ as input, *i.e.*, starting with a configuration of the form $q_0 w$, where q_0 is a chosen initial state, it halts in the accepting state whenever $w \in I$ and halts in the rejecting state otherwise. An *algorithm* for problem I is just another name for a Turing machine solving I. The **time complexity** of the computation on input w is the number of transitions needed to reach a halting state. The **space complexity** is the maximum length of a configuration during the computation.

Polynomial and exponential classes. An algorithm has **polynomial time complexity** if for every $n \in \mathbb{N}$ and every input of length n the computation on this input has time complexity at most $p(n)$, where p is a polynomial that only depends on the algorithm. The set of problems admitting algorithms of polynomial time complexity is denoted by **P**. Replacing $p(n)$ by $2^{p(n)}$ we obtain the class **EXP** of problems with exponential time complexity. Analogously, the set of problems solved by Turing machines whose space complexity is polynomial is denoted by **PSPACE**. It is believed, but not known, that **EXP** $\not\subset$ **PSPACE**.

Exercise 3.3.1. Show that **PSPACE** $\subset$ **EXP**.

The class NP. The acronym **NP** stands for the class of *nondeterministic polynomial time* algorithms. A problem I is in **NP** if there is a nondeterministic Turing machine such that (1) given any $w \in I$ as input at least one computation leads to an accepting state in polynomial time and (2) no computation leads to an accepting state whenever

[3] or, referring to the language terminology, to *recognize*.

$w \notin I$. Case (2) leaves the possibility that the machine runs forever, but computations that take more than polynomial time may be discarded without affecting the functionality of the machine, so that we can always assume that the computation takes polynomial time in both cases (1) and (2). However, the two cases are highly asymmetric since a computation leading to a rejecting state does not say anything about the input. There is another useful definition of the class **NP** in terms of efficiently verifiable certificate. A problem I is in **NP** if there is a deterministic Turing machine with polynomial time complexity, the *verifier*, such that (a) for every $w \in I$ there exists $c \in \mathscr{A}^*$ so that the verifier accepts wc in polynomial time and (b) if $w \notin I$ the verifier rejects wc whichever c we choose. Hence, c acts as a **certificate**, or efficiently verifiable proof for being a YES instance.

Theorem 3.3.2. *The two definitions of the class* **NP** *by means of nondeterministic machines or in terms of certificates and deterministic verifiers are equivalent.*

PROOF. Suppose that a language I is recognized by a nondeterministic machine M in polynomial time. An input word w determines a directed rooted tree of computations where each node corresponds to a configuration of M and the children of a configuration node correspond to the various transitions that apply to that configuration. The degree of a node is bounded by a constant, namely the size of the transition table of M. A computation path in this tree is easily encoded as the list ℓ of branching choices at the nodes along the path. By assumption, ℓ has polynomial size and may serve as a certificate. We can define a verifier V that takes the concatenation $w\ell$ (with some predefined separator) as input and essentially simulates the computation of M on w guided by ℓ. The successive branching choices in ℓ allow V to maintain the current configuration of M determined by those choices. The main task of the verifier is thus to check that each branching choice corresponds to an actual transition of M that applies to the current configuration. Clearly, V operates in polynomial time and $w \in I$ if and only if we can choose ℓ so that the simulation leads to an accepting state of M. We have thus proved that I is in **NP** according to the second definition.

Conversely, suppose that every word in I has a certificate verifiable by a polynomial time Turing machine V. We define a nondeterministic machine M operating in two stages. In the first stage, M guesses a certificate with polynomial length. In the second stage, M simulates V deterministically on the input word concatenated with the guessed certificate. The nondeterminism of M is thus concentrated in the first stage. It is easily seen that M recognize I as a member of **NP** in the sense of the first definition. □

Exercise 3.3.3. Show that **NP** $\subset$ **PSPACE**.

Reduction and completeness

The notion of reduction allows to compare the difficulty of different problems. Given two problems $I, J \subset \mathscr{A}^*$, we say that I **reduces** (in polynomial time) to J, written $I \leq J$, if there is a function $r : \mathscr{A}^* \to \mathscr{A}^*$, computable by a Turing machine with polynomial time complexity, such that $I = r^{-1}(J)$. In other words, r transforms YES and NO instances of the first problem to, respectively, YES and NO instances of the second problem[4]. Hence, $I \le J$ and $J \in \mathbf{P}$ implies $I \in \mathbf{P}$. This is obviously true replacing **P** by any other larger complexity class. If I reduces to J and J to K, it is easily seen I reduces to K. The reduction relation is thus a preorder (*i.e.*, a reflexive and transitive relation). Any problem which is an upper bound for a complexity class C is said **C-hard**. It is said **C-complete** if it furthermore belongs to **C**. A **C**-complete problem is thus a hardest representative in **C**. It is a priori not clear whether a complexity class has complete problems.

Exercise 3.3.4. Show that every non-trivial problem (proper subset of $\mathscr{A}^*$) in **P** is **P**-complete.

It turns out that the class **NP** has complete problems, among which the satisfiability problem. A Boolean formula is a logical expression over Boolean variables connected by the usual $\wedge, \vee, \neg$ operators. A formula is *satisfiable* if there is an assignment of its variables that makes the formula evaluate to true. The problem **SAT** is the set of satisfiable formulas encoded, say, over the alphabet $\{0, 1, \wedge, \vee, \neg, (,)\}$.

Theorem 3.3.5 (Cook'71 - Levin'73)**.** *SAT is* **NP***-complete.*

PROOF. Any truth assignment of a formula in SAT provides a certificate that is easily checkable in polynomial time. It follows that SAT $\in \mathbf{NP}$. It remains to show that every problem $I \in \mathbf{NP}$ reduces to SAT. Let $M = (\mathscr{A}, \mathscr{Q}, \mathscr{T})$ be a nondeterministic machine solving I in polynomial time. For every instance w, we need to construct a formula Φ_w so that $w \in I$ if and only if Φ_w is satisfiable.

Number the cells of the tape once for all from left to right so that at the initial step the tape contains $w = w_1 w_2 \dots w_n$ with cell 1 containing w_1. By assumption on M, the number of computation steps given w as input is bounded by $p(n)$ for some polynomial p, where $n := |w|$ is the length of w. By convention, we consider that M stays in the same configuration once in a halting state. This way we can assume that the number of computation steps is exactly $p(n)$. It follows that the head of M can only point to a cell with index in the range $J := [-p(n), p(n)]$. In particular, cells with index outside this range must contain the empty symbol. The whole computation is thus entirely described by the content of the jth cell at the ith step (configuration) of the computation, with $1 \le i \le p(n)$ and $j \in J$, and the sequence of $p(n)$ states and head positions during the computation. In accordance with this description, we introduce Boolean variables $C_{i,j,s}, Q_{i,q}, H_{i,j}$ with $1 \le i \le p(n)$, $j \in J$, $s \in \mathscr{A}$ and $q \in \mathscr{Q}$. The variable $C_{i,j,s}$ is intended to be true whenever the jth cell at the ith step contains s and false otherwise. Similarly, $Q_{i,q}$ and $H_{i,j}$ are intended to be true exactly when M is in state q at step i with the head pointing to the jth cell.

We next consider the following Boolean formulas. We recall that $A \implies B$ is a shorthand for $\neg A \vee B$.

[4]This type of reduction is called *many-one*, or *Karp reduction*. *Polynomial-time Turing reduction*, also known as *Cook reduction*, is another common notion of reduction, where I reduces to J if I can be solved in polynomial time by a Turing machine with an oracle for J, meaning that the machine is allowed to call a subroutine for problem J at anytime during the computation, in constant time for each call.

- $\phi_{i,j} = \bigvee_{s\in\mathscr{A}} (C_{i,j,s} \wedge (\bigwedge_{t\neq s} \neg C_{i,j,t}))$ expresses that the jth cell at the ith step takes one and only one value.
- $\phi_i = \big(\bigvee_{q\in\mathscr{Q}} (Q_{i,q} \wedge (\bigwedge_{r\neq q} \neg Q_{i,r}))\big) \wedge \big(\bigvee_{j\in J} (H_{i,j} \wedge (\bigwedge_{k\neq j} \neg H_{i,k}))\big)$ expresses that the state and head position each take exactly one value at the ith step.
- $\phi_b = \bigwedge_{1\leq j\leq n} C_{1,j,w_j} \wedge \bigwedge_{j\notin[1,n]} C_{1,j,\emptyset} \wedge Q_{1,q_0} \wedge H_{1,1}$ expresses that the initial tape contains the input w and that M is in the initial state q_0 with the head pointing to the first symbol of w.
- $\phi_e = Q_{p(n),q_a}$, where q_a is the accepting state, expresses that M accepts w.
- $\psi_i = \bigwedge_{\substack{j\in J\\ s\neq t}} \big((C_{i,j,s} \wedge C_{i+1,j,t}) \implies H_{i,j}\big)$ expresses that only the cell pointed by the head may change from step i to $i+1$.
- $\psi_{i,j,q,s} = (Q_{i,q} \wedge H_{i,j} \wedge C_{i,j,s}) \implies \bigvee_{s\,q\,t\,r\,D\in\mathscr{T}} (Q_{i+1,r} \wedge H_{i+1,j+D} \wedge C_{i+1,j,t})$ expresses that when M is in state q at step i with the head pointing to the jth cell containing s, only the relevant transitions may apply. Here, $j+D$ is $j-1$ or $j+1$ depending on whether $D = L$ or $D = R$.

We finally set $\Phi_w = \bigwedge_{i,j} \phi_{i,j} \wedge \phi_b \wedge \phi_e \wedge \bigwedge_i \psi_i \wedge \bigwedge_{i,j,q,s} \psi_{i,j,q,s}$. To conclude, it remains to notice that the description of the formula Φ_w can be computed in polynomial time (with respect to n) and that Φ_w is satisfiable if and only if M recognizes w, *i.e.* $w \in I$. $\square$

3.3.2 Existential theory of the reals

We are now ready to characterize the complexity of the linear embedding problem. Given as input an abstract simplicial complex K and a dimension d, the problem is to decide if K has a linear embedding into $\mathbb{R}^d$. As we shall see this problem can be reduced in polynomial time to test the non-emptiness of a semi-algebraic set defined by polynomials with integer coefficients.

Semi-algebraic set. An *atomic formula* may have one of two forms $\{p = 0\}$ or $\{p > 0\}$, where p is a polynomial in a finite number of variables, with integer coefficients. A predicate $\Phi(X_1,\dots,X_d)$ in the language of fields with integer coefficients is a Boolean predicate applied to atomic formulas using the free variables $X_1,\dots,X_d$. In other words, $\Phi(X_1,\dots,X_d)$ can be obtained recursively from atomic formulas using the logical connectors $\wedge, \vee$ and $\neg$. A **semi-algebraic set** over the integers is any set of the form

$$\{x = (x_1,\dots,x_d) \in \mathbb{R}^d \mid \Phi(x)\}$$

with Φ a predicate as above. An **existential formula** is a proposition of the form

$$\exists x \in \mathbb{R}^d \mid \Phi(x)$$

Deciding the falsity or truth of an existential formula is thus the same as deciding if a semi-algebraic set is empty or not. The set of problems that reduces in polynomial time to deciding the status of existential formulas has been gathered under the name of **existential theory of the reals**. This complexity class is denoted by $\exists\mathbb{R}$.

Lemma 3.3.6. **NP** $\subset \exists\mathbb{R}$.

PROOF. By Theorem 3.3.5 of Cook and Levin, it is enough to prove that SAT reduces to $\exists\mathbb{R}$. Let $\Phi(X)$ be a Boolean formula with variables $X = (X_1,\dots,X_d)$. Using the distributivity rules of negation over disjunction and conjunction we can assume that the negations in Φ may only apply to the atomic variables X_i. Let $Y = (Y_1,\dots,Y_d)$ be free variables, we recursively define a polynomial $P_\Phi(Y)$ using the formulas: $P_{X_i}(Y) = Y_i$, $P_{\neg X_i}(Y) = 1 - Y_i$, $P_{\Phi_1\wedge\Phi_2} = P_{\Phi_1}(Y)\times P_{\Phi_2}(Y)$ and $P_{\Phi_1\vee\Phi_2} = P_{\Phi_1}(Y) + P_{\Phi_2}(Y)$. We now consider the existential formula defined by the conjunction of the following predicates.

- $Y_i^2 - Y_i = 0$, $i = 1\dots d$.
- $P_\Phi(Y) > 0$.

Noting that $Y_i^2 - Y_i = 0$ implies $Y_i \in \{0,1\}$, it is easily checked that $\Phi(X)$ can be satisfied if and only if the above existential formula defines a nonempty semi-algebraic set. Moreover, a description of the existential formula can be obtained in time proportional to the length of the description of Φ, thus providing the required reduction. $\square$

A much more challenging task is to provide an upper bound on the complexity of $\exists\mathbb{R}$. The first approach to decide the vacuity of a system of polynomials (in)equations used the cylindrical decomposition of Collins (1975). This cylindrical decomposition includes a decomposition of $\mathbb{R}^k$, where k is the number of variables in the system, into semi-algebraic cells such that each polynomial in the system has a constant sign $-, 0$ or $+$ over each cell. Hence, the system has at least one solution if we can find one cell in the decomposition such that the sign of each polynomial agrees with the corresponding (in)equality in the system. The best known computation of such an adapted decomposition takes time $O(sd^{2^k})$ where s is the number of polynomials in the system and d is their maximal degree. The cylindrical decomposition approach thus leads to a doubly exponential time algorithm. It was eventually shown that $\exists\mathbb{R}$ could be solved using polynomial space only [Can88, Ren88].

Theorem 3.3.7 (Canny'88)**.** $\exists\mathbb{R} \subset$ **PSPACE**

The proof of this result is far beyond the purpose of this lecture. Describing all the details takes a whole thick book [BPR06]. There are excellent surveys [Bas14, RRSED00] that can serve as introductory lectures. An important step is to decide the (non)emptiness of a real algebraic set defined by a system $\mathscr{S}$ of polynomial equations. The main idea is to augment $\mathscr{S}$ with other polynomial conditions so that the new system has only a finite number of solutions, a so-called *zero-dimensional system*, with at least one solution in each (semi-algebraically) connected component defined by $\mathscr{S}$. Those solutions can even be returned implicitly using *rational univariate* *representations*. This is done by searching for the critical points of a given functional (e.g. the squared distance to a fixed point) over the algebraic set. For this method to work it is required that the critical points are non-singular and that the components are bounded. One way to enforce these conditions is to use symbolic perturbations. They are obtained by introducing new variables playing the role of infinitesimals, replacing equations of the form $P = 0$ by $P = \varepsilon$, for ε an infinitesimal. Other modifications may be introduced to take care of the unbounded components leading to a new system of polynomial equations whose coefficients are now polynomials in the infitesimals. After solving the modified system, it remains to substitute zero for the infinitesimals to obtain real solutions. A huge amount of techniques from real algebraic geometry are necessary, such as the use of resultants, root counting, Gröbner basis computations, etc. In the end, it can be proved that the emptiness of a semi-algebraic set defined by a system of polynomial (in)equations can be decided using polynomial space, in terms of the size of the encoding of the system. Only a few implementations seems to exist and are hardly able to deal with more than a dozen variables with polynomials of relatively low degree.

Exercise 3.3.8. Show that the emptiness of a semi-algebraic set defined by polynomial (in)equations can be reduced to the emptiness of an algebraic set defined by polynomial equations. Show that you can furthermore impose that the algebraic set is defined by a single polynomial equation.

Linear embeddability belongs to $\exists\mathbb{R}$

In the introduction to Section 3.3 we already observed that the embeddability of a simplicial complex K could be reduced to the satisfiability of a set of polynomial inequalities. We still need to check that this reduction takes polynomial time. Recall that we have to encode the conditions that pairs of disjoint simplices are sent to non-intersecting simplices in $\mathbb{R}^d$. The transcription into polynomials of those conditions for each pair of simplices just claims the existence of a separating hyperplane and clearly takes polynomial time. There still remains the potential problem that the number of simplices, hence the number of polynomials conditions, is very large compared to the encoding of K. A reasonable encoding should indeed only records the **maximal simplices** of K — those that are not a face of larger simplices — the other simplices being implicitly encoded as faces of the maximal ones. For instance, if $|K|$ is an m-dimensional simplex, its total number of faces is 2^{m+1} while its encoding is essentially the single set $[m+1]$. Nonetheless, since $m \le d$ is an obvious condition for embeddability in $\mathbb{R}^d$, we are led to conclude that

Theorem 3.3.9. *The linear embedding problem into $\mathbb{R}^d$ is in $\exists\mathbb{R}$ for any fixed dimension d.*

The question raised by the potentially large number of polynomial conditions can be dealt with at the expense of getting larger polynomials. We can indeed replace the conditions in Lemma 3.2.1 by a smaller number of conditions. To see this, we first make a simple observation.

Lemma 3.3.10. *Let σ, τ be two simplices in $\mathbb{R}^d$ intersecting along a common face. There exists a hyperplane intersecting each of σ, τ along their common face and otherwise separating them.*

PROOF. Let $\{u_i\}_{i\in I}$ be the vertices of the common face, and let $\{v_j\}_{j\in J}$ and $\{w_\ell\}_{\ell\in L}$ be the remaining vertices of σ and τ, respectively. Let $u_i^{\sigma,\varepsilon} := (1-\varepsilon)u_i + \varepsilon v_{j_0}$ for some fixed $j_0 \in J$. Likewise, let $u_i^{\tau,\varepsilon} := (1-\varepsilon)u_i + \varepsilon w_{\ell_0}$ for some $\ell_0 \in L$. For $0 < \varepsilon < 1$, $\sigma_\varepsilon := \mathrm{Conv}(\{u_i^{\sigma,\varepsilon}\}_{i\in I} \cup \{v_j\}_{j\in J})$ and $\tau_\varepsilon := \mathrm{Conv}(\{u_i^{\tau,\varepsilon}\}_{i\in I} \cup \{w_\ell\}_{\ell\in L})$ are disjoint compact convexes, hence separated by a hyperplane H_ε defined by a unit normal vector ν_ε and a point u_ε, say between $u_1^{\sigma,\varepsilon}$ and $u_1^{\tau,\varepsilon}$. As ε tends to zero, ν_ε and u_ε converge toward a vector ν_0 and a point u_0 defining a hyperplane H_0. It is easily seen that H_0 has the required property. $\square$

With some abuse of terminology we still call the hyperplane as in Lemma 3.3.10 a *separating hyperplane* for (σ, τ).

Corollary 3.3.11. *The embedding conditions in Lemma 3.2.1 can be replaced by the following: (1) the vertices of each maximal simplex of K are sent to affinely independent points in $\mathbb{R}^d$ and (2) for every pair of distinct maximal simplices of K, there exists a separating hyperplane in the sense of the previous lemma.*

PROOF. Condition (1) is trivially necessary for any linear embedding. Lemma 3.3.10 implies that condition (2) is also necessary. Conversely, suppose that a linear map $f: K \to \mathbb{R}^d$ satisfies (1) and (2). Let σ, τ be to disjoint simplices of K. σ and τ are faces of two maximal simplices, say σ' and τ' respectively. On the one hand, if $\sigma' = \tau'$, condition (1) implies that σ, τ are sent to disjoint faces of a non-degenerate simplex in $\mathbb{R}^d$. On the other hand, if $\sigma' \neq \tau'$, condition (2) implies the existence of a separating hyperplane for (σ', τ') providing a separating hyperplane for (σ, τ). In any case, σ and τ are sent to non-intersecting simplices in $\mathbb{R}^d$, showing that f is an embedding by Lemma 3.2.1. $\square$

By Corollary 3.3.11 we just need a number of polynomial conditions that is quadratic in the number of maximal simplices. Beware, though, that condition (1) is expressed by the non-cancellation of determinants that may contain up to $d!$ terms. The potential benefit of this approach in terms of the number of polynomials should thus be balanced with the increase in the size of the polynomials.

Bibliography

[ABEM07] Dan Archdeacon, C. Paul Bonnington, and Joanna A. Ellis-Monaghan. How to exhibit toroidal maps in space. *Discrete & Computational Geometry*, 38(3):573–594, 2007.

[ABS96] Amos Altshuler, Jürgen Bokowski, and Peter Schuchert. Neighborly 2-manifolds with 12 vertices. *journal of combinatorial theory, Series A*, 75(1):148–162, 1996.

[Bas14] Saugata Basu. Algorithms in real algebraic geometry: a survey. 2014.

[BE91] Jürgen Bokowski and Anselm Eggert. All realizations of Möbius' torus with 7 vertices. *Structural Topology 1991 núm 17*, 1991.

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

[BPR06] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer, 2006.

[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.

[Can88] John Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 460–467. ACM, 1988.

[dL13] Mark de Longueville. *A Course in Topological Combinatorics*. Universitext. Springer, 2013.

[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.

[FKT94] Michael H Freedman, Vyacheslav S Krushkal, and Peter Teichner. van Kampen's embedding obstruction is incomplete for 2-complexes in $\mathbb{R}^4$. *Math. Res. Lett*, 1(2):167–176, 1994.

[HT74] J. Hopcroft and R. Tarjan. Efficient planarity testing. *Journal of the ACM*, 21(4):549–568, 1974.

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

[LdM17] Francis Lazarus and Arnaud de Mesmay. Planar graphs (lecture notes). pdf, 2017.

[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.

[Mat08] J. Matoušek. *Using the Borsuk-Ulam theorem.* Universitext. Springer Verlag, 2008.

[MTW11] Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. *Journal of the European Mathematical Society*, 13(2):259–295, 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.

[Ren88] James Renegar. A faster PSPACE algorithm for deciding the existential theory of the reals. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 291–295. IEEE Computer Society, 1988.

[RRSED00] Fabrice Rouillier, Marie-Françoise Roy, and Mohab Safey El Din. Finding at least one point in each connected component of a real algebraic set defined by a single equation. *Journal of Complexity*, 16(4):716–750, 2000.

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

[Sch10] Lars Schewe. Nonrealizable minimal vertex triangulations of surfaces: showing nonrealizability using oriented matroids and satisfiability solvers. *Discrete & Computational Geometry*, 43(2):289–302, 2010.

[Wig06] Avi Wigderson. P, NP and mathematics – a computational complexity perspective. In *Proc. of the 2006 International Congress of Mathematicians*, volume 3, 2006.

[Wu85] Wen-Jun Wu. On the linear embedding of linear graphs. *J. Sys. Sci. & Math. Sci.*, 5(4):290–302, 1985. Translated from Chinese, 1973.

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

[Zie95] Günter M. Ziegler. *Lectures on polytopes.* Number 152 in Graduate texts in mathematics. Springer, rev. first edition, 1995.