Embedding in Euclidean spaces: the double dimension case

Francis Lazarus

2020

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.

1 Topological prerequisites

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 $p_0, \ldots, p_k$ in some $\mathbb{R}^d$ and is denoted by $[p_0, \ldots, p_k]$. 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.

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

[1]This assumption facilitates the definition of the join operation. See Section 3.1.

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**, $\mathrm{sd}\, K$, 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 $\mathrm{sd}\, K$ has thus $K \setminus \{\emptyset\}$ itself for ground set and its nonempty simplices have the form $(\sigma_0, \dots, \sigma_k)$ where $\sigma_0 \subset \dots \subset \sigma_k$ is a strictly increasing sequence of nonempty simplices of K.

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.

Proposition 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

[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.

assumption. To see that f is injective we just need to prove that distinct simplices have their interior sent to disjoint sets. So, let $\sigma = [v_1, \ldots, v_k]$ and $\tau = [w_1, \ldots, w_\ell]$ be two distinct simplices of K. Since $k + \ell \le 2n + 2$, the general position assumption implies that the image points $f(v_1), \ldots, f(v_k), f(w_1), \ldots, f(w_\ell)$ 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 1.2. Prove that any set of points on the moment curve $t \mapsto (t, t^2, \ldots, t^d)$ 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 $d \le 2n$. 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 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 2.1 (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]

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

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

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

$$\delta(v\times\tau)^* = \sum_{v\in\sigma}(\sigma\times\tau)^* \tag{1}$$

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

Lemma 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

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

and

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

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 (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 $C^2/\operatorname{Im}\delta$.

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.

[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.

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 2.3. *$[c_f]_2$ 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 1. An isotopy or any of the I-IV moves

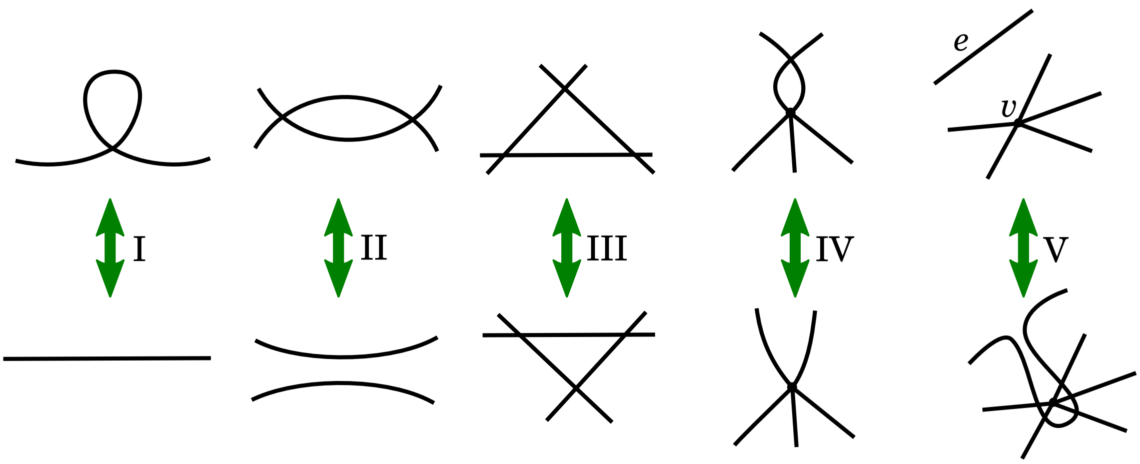

Figure 1: The first three moves I, II, III are known as (shadows of) Reidemeister moves. The IV move amounts to a transposition in the edge order around a vertex, while the V move is referred to as a **finger move** or an **(e,v)-move**.

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 2.4. Figure 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 1 are the only required moves to transform an immersion into another one. (See Exercise 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.

Lemma 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.

[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 $\gamma:[\alpha,\beta]\to\mathbb{C}$.

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. $\square$

PROOF OF LEMMA 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 $\partial \sigma = s - r$
$$\begin{aligned} c_f(\sigma \times e) - c_g(\sigma \times e) &\equiv |f(\sigma) \cap f(e)| - |g(\sigma) \cap g(e)| \mod 2 \\ &\equiv |f(\sigma) \cap C_e| \mod 2 \qquad (\text{since} f(\sigma) = g(\sigma)) \\ &= w_2(C_e, f(s)) - w_2(C_e, f(r)) \qquad (\text{by Lemma 2.5}) \end{aligned}$$
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 $e_1, \dots, e_m$ be the edges of G. We define immersions f_i that agree with g on $e_1, \dots, e_i$ 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 $i = 1, \dots, m$. 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.

[7] This argument found by Axel Péneau simplifies the proof of Wu.

□

In view of Lemma 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, 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 2.6. *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 2.7. *$\kappa_2(K_5)$ 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 2.8. *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 2.6 and 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.

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

[8] The subscript 2 is used to emphasize that we consider mod 2 cohomology.

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].

3.1 Join operations

3.1.1 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.

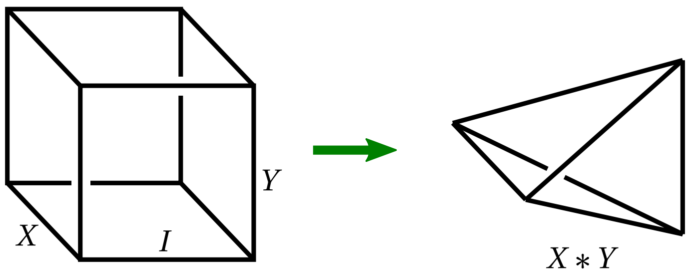

Suppose that X and Y are subspaces of some Euclidean space, and that X and Y are contained in respective affine subspaces that are affinely independent, meaning that the union of affinely independent pointsets, one in each subspace, is itself independent. Then, $X*Y$ is homeomorphic to the union of all 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. 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

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

geometric join is the union of their vertices. In particular, a simplex with vertices $\dots, p_i, \dots$ 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 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.

3.1.2 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 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.*

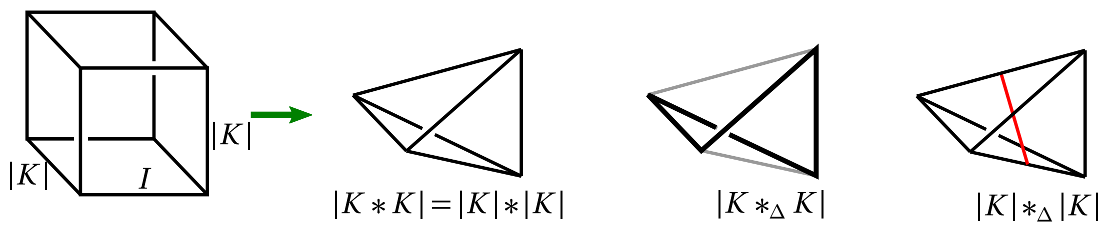

PROOF. Denote by $\rho : |K| *_\Delta |K| \to |K *_\Delta K|$ the retraction we are looking for. Geometrically, if $\Delta' := |K| * |K| \setminus |K| *_\Delta |K| = \{\frac{1}{2} x \oplus \frac{1}{2} x \mid x \in |K|\}$ is the *diagonal* of $|K| * |K|$, we shall define ρ so that it retracts every conic 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.

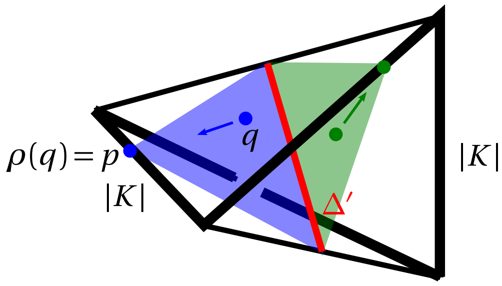

More formally, denote by $\mathrm{Supp}(x)$ the supporting simplex of a point $x \in |K|$. We have

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

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 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. $\square$

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 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 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. $\square$

Lemma 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. $\square$

Exercise 3.11 in the next section asks you to strengthen Lemma 3.5 to show that $\mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d) = d$.

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

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 $\mathrm{Ind}(|K *_\Delta K|) \le \mathrm{Ind}(\mathbb{R}^d *_\Delta \mathbb{R}^d)$. In view of Lemma 3.5, we have

Proposition 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 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 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 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

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

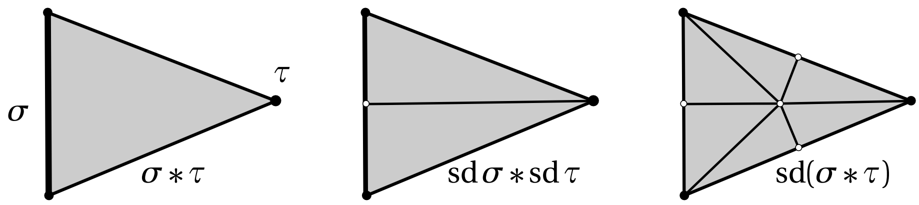

Comparison between the shore and barycentric subdivision of the 2-simplex expressed as the join of an edge and a vertex.

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

Proposition 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 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 3.8 and Proposition 3.9, $\mathrm{Bier}_{2d+3}(K) = K *_\Delta K$ is isomorphic to the $(2d+1)$-sphere. From Proposition 3.4, we have $\mathrm{Ind}(K *_\Delta K) = 2d+1$ and we conclude by invoking Proposition 3.6. □

Exercise 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 3.12. Consider the simplicial complex represented in the figure below.

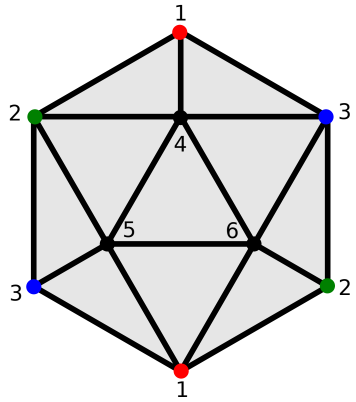

It is composed of 6 vertices and 10 triangles forming a disk, the boundary edges of which should be identified according to the boundary vertex numbering. This complex is topologically a projective plane. Mimic the proof of the van Kampen - Flores theorem to prove that the projective plane cannot be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

References

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

[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.

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

[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.

[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.