Miscellaneous topological algorithms

Colm Ó Dúnlaing
Colum Watt
David Wilkins
Mathematics, Trinity College, Dublin 2, Ireland.
Chee-Keng Yap
Computer Science, Courant Institute,
251 Mercer Street, New York, NY 10012, U.S.A.[*]

Abstract

These notes survey some techniques of Computational Topology and include a few original results.

1 Introduction.

Computational Topology, i.e., the algorithmic study of topological methods, is developing rapidly [18, 33]. Topics covered in these two books range from algorithms based on Hopf's theorem to constructive methods of homological algebra. Algorithms for calculating in the homotopy groups and homology groups of surfaces have been reported at Computational Geometry conferences [38, 13, 14].

A strong motive in studying computational topology is 3-dimensional graphics and visualisation. Some work has been done on automating techniques of visualisation [16, 2].

In collecting the results surveyed here we have been most interested in the extension of 2-dimensional algorithmic problems such as graph planarity to 3 dimensions. Therefore most of the results focus on the natural extension of a *graph* to that of a *2-complex*, whether simplicial or more general. A particularly interesting problem is to decide whether a given 2-complex can be embedded in 3-dimensional space.

Apart from appearing much more difficult that the corresponding graph planarity problem, there appears to be a distinction between 'flat' and 'topological' embedding. The question of 'flat' embeddability has been dealt with satisfactorily in the work of Bokowski and others [6, 7, 8, 9]. Their methods are sketched very briefly at the end of these notes.

[*]Work of the first three authors was supported in part by the EC under ESPRIT BRAs 3075 (ALCOM) and 7141 (ALCOM II). Work by the fourth author was supported in part by the US NSF under grant #CCR-9402464. E-mail: odunlain@maths.tcd.ie, colum@maths.tcd.ie, dwilkins@maths.tcd.ie, yap@cs.nyu.edu. Some of this paper was presented at CG'92, the Workshop on Computational Geometry, Utrecht, in March 1992.

Another distinction is that while all (embedded) planar graphs are triangulable, the same is not true for embedded 2-complexes. Indeed, it is NP-complete to decide if a polyhedron can be tetrahedralised [31].

Apart from the work of Bokowski et al., the survey is mostly about topological embeddings. We begin with a review of a non-algorithmic result of Boyle's [10], namely, that an embedding of a 2-complex is usually determined topologically by the embedding of its 1-skeleton. This is included (Section 10) because it illustrates a technique (cancelling self-intersections from immersions) which is echoed in work of Shapiro [32] and elsewhere. Also we feel that the technique may yield some procedure to extend an embedding of the 1-skeleton to that of the whole complex, or decide when this is impossible. For this reason we also sketch (Section 5) some of Taniyama's work on invariants of graph embeddings in 3 dimensions [34, 35, 36, 37].

In Section 3 we sketch Whittlesey's methods of classifying 2-complexes up to homeomorphism. Our survey includes a satisfactory description of the complexity of the classification problem: it is polynomially equivalent to graph isomorphism. These results augment previous algorithmic work on the classification of surfaces [38].

Section 4 describes a homological invariant of n-complexes, a cocycle, whose vanishing is necessary for these complexes to be embeddable in $\mathbb{R}^{2n}$. This follows work of Shapiro's [32] and is related to work of Wu's [44]. For $n \neq 2$ the condition is also sufficient.

Sections 6 and 7 describe sufficient conditions for embeddability in certain cases.

Section 8 describes methods of Neuwirth's for investigating embeddability of n-complexes in manifolds of dimension $n+1$. We sketch Neuwirth's idea of 'in-groups' and 'out-groups' and give a necessary condition for embeddability. Section 9 gives a semi-algorithm of Neuwirth's for deciding if a 2-complex is the spine of a 3-manifold.

The survey concludes, as mentioned above, with a sketch of the methods of synthetic geometry [9].

Most of the work surveyed here is rather old; where the literature is more up-to-date our descriptions are more sketchy. The list of references includes papers that are of some interest but not necessarily cited in the paper.

It is interesting to note that methods described here give another three graph planarity algorithms: through obstructions [32, 39, 44], through out-groups [28], and through oriented matroids [6, 7].

Some of the results reported here are new: the complexity of homeomorphism of 2-complexes, and the embeddability criteria for some special cases (Sections 6 and 7).

2 Embeddings fixed by 1-skeletons.

In [10], Boyle considers embeddings of 2-complexes M in the 3-dimensional sphere S^3. S^3 can be considered in some sense as $\mathbb{R}^3$ closed in upon itself; it is possible to speak of 'piecewise-linear embeddings' in S^3.

He first gives a rather generalised definition of '2-d cell complex' M, then specifies certain restrictions on M, 'proper' and 'fine enough,' without which the theorem is false. His theorem says:

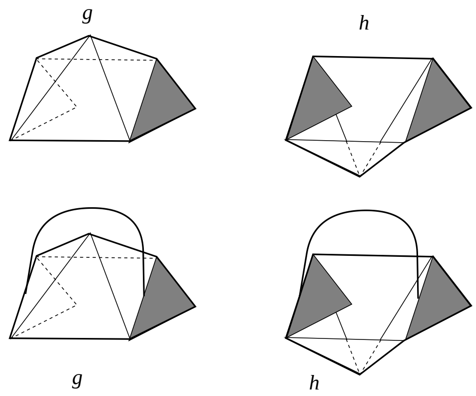

Figure 1: Theorem 2.1 fails when M is not fine enough.

THEOREM 2.1 *If M is a 2-d PL cell-complex, proper, and fine enough, and g and h are two PL embeddings of M in S^3, and they agree on the 1-skeleton of M, then g can be deformed into h by a PL ambient isotopy of S^3 which leaves the 1-skeleton fixed at all times.*

This means that to all intents and purposes g and h are topologically the same.

To avoid 'wild' embeddings, the embeddings are assumed to be piecewise linear. Any such embedding of the 2-sphere S^2 has the property that it separates S^3 into two open cells, of which it is the common boundary, and the closures of which are homeomorphic to closed 3-dimensional balls. The fact that the inside and outside are homeomorphic is significant, and the result does not carry over easily to embeddings in $\mathbb{R}^3$.

DEFINITION 2.2 *A* closed *(respectively, open) PL k-cell in $\mathbb{R}^n$ is the image of $[0,1]^k$ (respectively, $(0,1)^k$) under a 1-1 piecewise linear map; for $k=0$ it's a point (and identified with its interior in this case). A* 2-d PL cell-complex *M consists of a set of* 0-, 1-, *and* 2-*cells in some $\mathbb{R}^n$, such that the underlying space $|M|$ is partitioned by the cell interiors, and such that the boundary of each k-cell ($k=1,2$) is a union of $(k-1)$-cells.*

The underlying space, properly written $|M|$, may by abuse of notation be written as M.

The r-skeleton *$M^{(r)}$ is the union of cells of dimension $\leq r$. The 1-skeleton $M^{(1)}$ is a topological graph.*

DEFINITION 2.3 *A 2-d PL cell-complex M is* proper *if the intersection of any collection of 2-cells is always connected. It is* fine enough *if* (i) *$M \backslash f$ is connected for every closed 2-cell f;* (ii) *$M \backslash e$ is connected for every arc e which is the intersection of two 2-cells;* (iii) *$M \backslash \{v_1, v_2\}$ is connected for any pair of points v_1, v_2 on the boundary of f, for every 2-cell f.*

Figure 1 illustrates how the theorem fails if the complex M is not fine enough. In the first example, M is a closed disc expressed as the union of a square with two triangles; removing the square face disconnects M. The two exhibited embeddings are not mutually isotopic relative to $M^{(1)}$; Boyle gives as a reason the fact that if one 'joins together' the two pyramids to form an octohedron S, one of the triangle vertices u is inside S and one, v, is outside. The space $S^3 \backslash \{u,v\}$ is homotopic to S^2, so the second homology group $H_2(S^3 \backslash \{u,v\}) \equiv \mathbb{Z}$, and S is a generator for this group, hence is not null-homologous; hence there is no isotopy of $S^3 \backslash \{u,v\}$ fixing the boundary of the square and carrying g into h.

epsfysize=1.5in

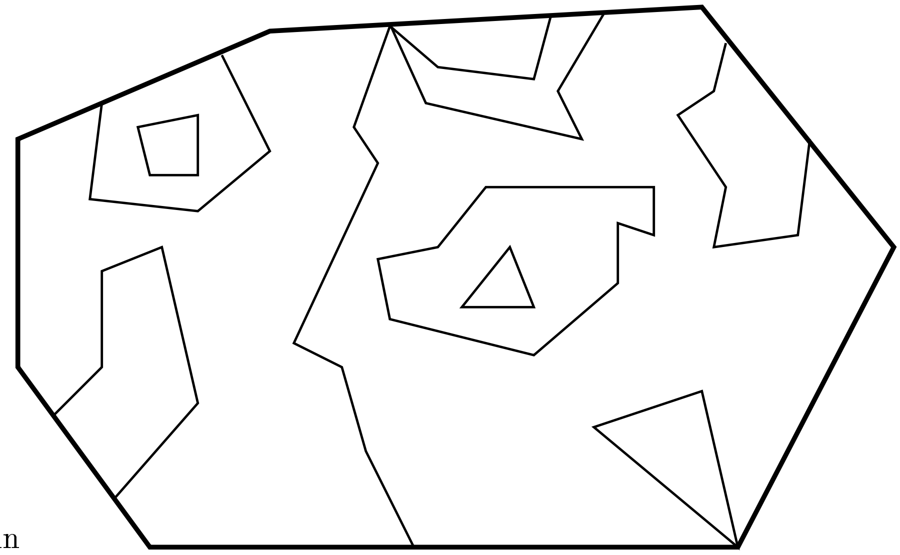

Figure 2: Intersection curves on the interior of f_h.

In the second illustrated example, requirement (i) of 'fine enough' is satisfied but not requirement (iii). The two embeddings are not isotopic relative to $M^{(1)}$. Boyle proves this as follows. The complex M is a disc with a handle attached to two points on its interior. In each embedding the handle can be continued into a closed loop L by joining its ends to the apex of the 'pyramid' shown. Consider also the (6-sided) boundary B of M, part of the 1-skeleton. In g, L and B are unlinked; in h they are linked. Hence the embeddings cannot be isotopic in S^3.

The fundamental technique of the paper is to take two different embeddings of M which agree on the 1-skeleton, 'disentangle' intersecting faces by continuous (PL) deformations, then for each face to 'merge together' its two images.

To prove the theorem, first some piecewise-linear adjustments are made to the embeddings to ensure that the embedded complexes intersect cleanly. Next the author shows, using M_g to denote the image of M under g, and f_h to denote the image of a face f of M under h, that M_g intersects f_h in the boundary and in curves of the following types:

(i) Simple closed curves interior to f_h; (ii) simple closed curves intersecting the boundary ∂f_h in a single point; and (iii) simple arc-segments meeting ∂f_h at their endpoints only.

(See Figure 2.) Note that since g and h agree on $M^{(1)}$, each of these curves (restricted to the interior of f_h) is interior to a unique face f_g of M_g (though not always the same face, of course).

Each such curve α separates the interior of f_h into two open regions; the closure of at least one of these regions is homeomorphic to a closed disc D_h. The author characterises some of the disc-like components as being 'innermost,' a term we shall leave undefined, and shows that if M is proper and f_h has some intersection curves, then there exists an innermost disc D_h. For every such disc D_h, arising from a curve α, there is a unique disc D_g which is the closure of a component of $f_g \backslash \alpha$ and such that $D_g \cup D_h$ is homeomorphic to S^2. (D_g might not be innermost on f_g.)

The author shows that if M is fine enough, and f_h has some intersection curves, then there exists an innermost disc D_h, with corresponding disc D_g, such that $M_g \backslash D_g$ is connected. Write

S for the PL 2-sphere $D_g \cup D_h$. It partitions S^3 into two open balls. Then $M_g \backslash D_g$ is entirely contained in one or other of these balls, which we shall designate as the *exterior* of S. Let B denote the interior: the closed complement of the exterior.[1] B does not intersect $M_g \backslash D_g$.

The author shows that there exists an isotopy of S^3, fixing $M_g^{(1)}$, whose effect on M_g is to 'move' B to a new placement B' outside S, which is strictly disjoint from B except where it intersects the 1-skeleton. The closed ball B' is sufficiently close to S to ensure that no new intersections with M_h are introduced. This isotopy brings M_g into a new position $M_{g'}$ with fewer intersection curves (α being elminated). Proceeding iteratively, the embedding M_g is brought by (ambient) isotopy to a new embedding g'' such that M_h and $M_{g''}$ are disjoint outside the 1-skeleton.

The process is finished by repeating the argument so by an ambient isotopy all the faces of $M_{g''}$ are made to coincide with the corresponding faces of M_h.

3 Classification of finite 2-complexes.

In Whittlesey's two papers [41] an algorithm is given to decide whether two 2-complexes are homeomorphic.[2] The first part of the analysis describes the essential features of a 2-complex as follows.

A *finite simplicial 2-complex* K is like a PL 2-complex with the following exceptions: (i) the 'piecewise linear' requirement is dropped; (ii) all the faces (i.e., 2-cells) are triangular, bounded by cycles of 3 1-cells. As before, we confuse K (a finite combinatorial object) with the topological space $|K|$.

We call the 0-cells vertices, the 1-cells edges, and the 2-cells faces, of K. It is convenient to re-interpret the faces as homeomorphic images of a 'standard 2-simplex,' such as the triangle $\Delta = \{(x,y,z) \in \mathbb{R}^3 : x,y,z \geq 0 \;\&\; x+y+z=1\}$. The midpoint of an edge is the image of the midpoint (1/2,1/2) of $[0,1]$, and the midpoint (barycentre) of a face is the image of the midpoint $(1/3, 1/3, 1/3)$ of Δ. If from K we construct a new simplicial complex by adding to the vertices of K all midpoints of its edges and faces, replacing each edge by two edges connected at its midpoint, and replacing each face by six edges and faces meeting at its midpoint as illustrated in Figure 3, then we obtain a simplicial complex K' called the (first) *barycentric subdivision* of K. Repeating the process with K' we obtain the second barycentric subdivision of K.

Given a point p in K, the *open star* $\mathrm{st}(p)$ is the union of the interiors of all cells containing p. This is an open neighbourhood of p in K.

First a description of the local topology of K: A point p is *regular* if K is surface-like near p, i.e., $\mathrm{st}(p)$ is homeomorphic to an open disc; it is *line-singular* if it has a neighbourhood like a 'book' with any number of pages except 2, and p is on the spine of the book; in this case p must be on an edge; it is a *conical point* (or 'pinch-point,' necessarily at a vertex) if the space locally resembles a union of open discs joined together at p; and otherwise it is a *node*. If p is not regular, then it is *singular*.

[1]If the embedding is in $\mathbb{R}^3$ rather than S^3, we do not have this freedom of choice, and the proof does not apply.

[2]This was Whittlesey's Ph.D. topic supervised by A.W. Tucker, who had thought the problem unsolvable, since it is undecidable of two 2-complexes whether their fundamental groups are isomorphic. See [1].

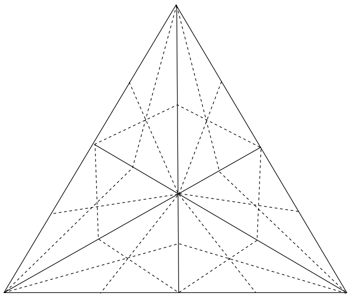

Figure 3: first and second barycentric subdivisions of a face

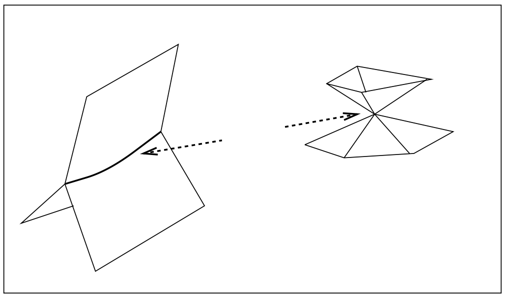

Figure 4: Line-singular points and conical point, with strip-neighbourhoods.

The set of all singular points of K is called the *singular graph* of K. It is contained in the 1-skeleton $K^{(1)}$.

In [41] the following is shown: take a narrow 'strip-neighbourhood' of the singular graph, so we have a union of narrow strips glued together along the singular graph. Let S be this neighbourhood; then $|K|\backslash S$ is a union of finitely many connected pieces M_j, where each M_j is a surface-with-boundary. The (1-dimensional) boundary ∂S is a union of closed loops along which the surfaces M_j are attached. Each such closed loop ρ_i projects down to a cycle σ_i in the singular graph[3]. Then the topological structure $|K|$ is characterised uniquely by

1. The singular graph in which each edge has been given an orientation.
2. The set of surfaces M_j, each given an orientation (where orientable).
3. The paths σ_i.
4. The attaching correspondences between each σ_i and the surface M_j whose boundary contains ρ_i.

Technically, the 'strip neighbourhood' S can be defined by proceeding to the second barycentric subdivision K'' of K and defining S to be the open star of the singular graph in K''. The open star is the union of interiors of cells which intersect the singular graph. The closed star is the union of cells which intersect the singular graph. The boundary of the closed star is called the *link* of the singular graph in K''. (It is the union of closed cells contained in the closed star which do not intersect the singular graph.) The closed loops ρ_i are the components of the link. See Figure 5

[3]Generally σ_i will be a non-simple cycle.

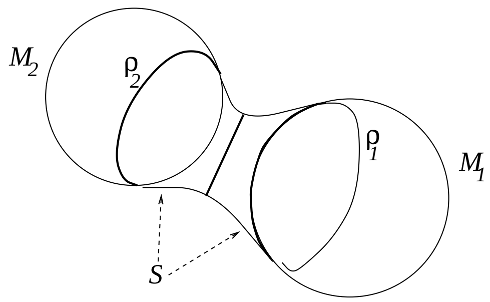

Figure 5: K homeomorphic to two spheres stitched together. The singular graph is a curve-segment, and its link is the two loops illustrated.

The whole structure can be revised by allowing the following possible reorientations:

- Reorienting any edge of the singular graph.
- Reorienting any orientable M_j with its incident line-paths σ_i.
- Reorienting any σ_i incident to any nonorientable surface M_j.

Having produced a topological description which can be derived from any 2-complex, Whittlesey shows how to reduce it to a normal form. This generalises Brahana and Tuckers' normal-form classification of triangulated compact surfaces [24].

We begin with a 2-complex K, formed by attaching the boundaries of closed discs D_i to its 1-skeleton by closed, not necessarily simple, paths σ_i. Let the edges of the 1-skeleton $K^{(1)}$ be oriented, yielding a digraph G. Let the vertices of G be labelled with uppercase letters and its directed edges with lowercase letters. Each directed cycle σ_i corresponds to a sequence $AaBb \dots A$ of letters or inverted letters labelling vertices and directed edges (if, for example, the edge $\{A, B\}$ is oriented as (B, A) in G then $a = d^{-1}$ where d is the letter labelling (B, A)). Now σ_i describes the attachment of a topological disc D_i, and the topology of $|K|$ depends essentially on these *words* $AaBb \dots A$. This Whittlesey calls a *word-system* W.

To avoid uninteresting details, we assume that every vertex and edge is incident to at least one face; so the list of vertices and edges is implicit in the word-system W.

If a is an edge-label, then say it is 'mentioned' wherever a or a^{-1} occurs in some word. An edge (label) is *regular* if it is mentioned exactly twice. The following alterations to a word-system W are shown to be admissible, since they produce homeomorphic complexes:

1. Edge composition: if a, b are labels (perhaps inverted) for edges with common vertex P and they only appear in W in the contexts aPb or $b^{-1}Pa^{-1}$ then the first can be replaced by a new edge-letter x and the second by x^{-1} uniformly throughout W.
2. Edge subdivision, inverse to this.
3. Word (or polygon) subdivision: a word $A\alpha P\mu Q\beta A$ can be replaced by two words $A\alpha PxQ\beta A$ and $P\mu Qx^{-1}P$ where x is a new letter.
4. Word composition, inverse to this, providing the edge label x is regular.

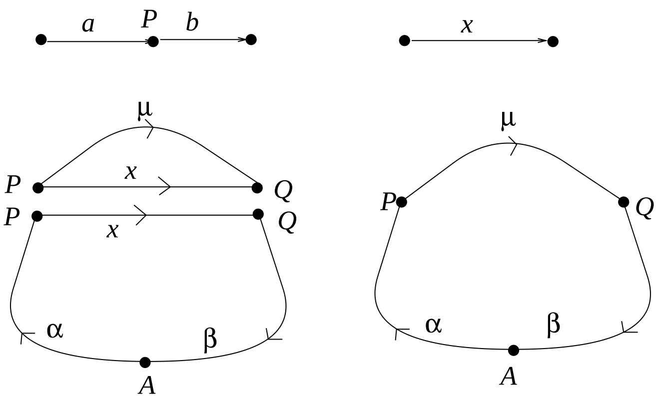

Figure 6: illustrating edge and polygon composition and subdivision.

See Figure 6. Two word-systems are *combinatorially equivalent* iff one can be converted into the other by a sequence of compositions and/or subdivisions, together with the following 'formal equivalences' on words, which clearly yield homeomorphic complexes: relabelling a vertex throughout, relabelling an edge throughout, replacing an edge-label by its inverse throughout, cyclically permuting a word, or replacing a word $AaBb\ldots A$ by its inverse $A\ldots b^{-1}Ba^{-1}A$.

It must be emphasised that polygon composition— combining two different words into one— is allowed only on *regular* edges.

Whittlesey introduces further 'derived rules' which can be applied to individual words:

- Incision: replacing a word $A\alpha\mathbf{Q}\beta A \to A\alpha\mathbf{PxQx^{-1}P}\beta A$ where x and P are new labels.
- Inverse to incision: if x is regular, $A\alpha\mathbf{PxQx^{-1}P}\beta A \to A\alpha\mathbf{Q}\beta A$.
- Circulation rule 1a: where x is regular, $A\alpha\mathbf{xQ}\mu\mathbf{T}\beta\mathbf{xQ}\gamma A \Leftrightarrow A\alpha\mathbf{xT}\beta\mathbf{xT}\mu^{\mathbf{-1}}\mathbf{Q}\gamma A$.
  $\ldots xQ\mu T\ldots xQ\ldots \Leftrightarrow \ldots xT\ldots xT\mu^{-1}Q\ldots$.
- Circulation rule 1b: where x is regular, $A\alpha\mathbf{T}\mu\mathbf{Px}\beta\mathbf{Px}\gamma A \Leftrightarrow A\alpha\mathbf{Tx}\beta\mathbf{P}\mu^{\mathbf{-1}}\mathbf{Tx}\gamma A$.
  $\ldots T\mu Px\ldots Px\ldots \Leftrightarrow \ldots Tx\ldots P\mu^{-1}Tx\ldots$.
- Circulation rule 2a: where x is regular, $A\alpha\mathbf{xQ}\mu\mathbf{T}\beta\mathbf{Qx^{-1}}\gamma A \Leftrightarrow A\alpha\mathbf{xT}\beta\mathbf{Q}\mu\mathbf{Tx^{-1}}\gamma A$.
  $\ldots xQ\mu T\ldots Qx^{-1}\ldots \Leftrightarrow \ldots xT\ldots Q\mu Tx^{-1}\ldots$.
- Circulation rule 2b: where x is regular, $A\alpha\mathbf{Q}\mu\mathbf{Tx}\beta\mathbf{x^{-1}T}\gamma A \Leftrightarrow A\alpha\mathbf{Qx}\beta\mathbf{x^{-1}Q}\mu\mathbf{T}\gamma A$.
  $\ldots Q\mu Tx\ldots x^{-1}T\ldots \Leftrightarrow \ldots Qx\ldots x^{-1}Q\mu T\ldots$.

These rules become more transparent if the vertex letters are omitted: incision becomes $\alpha\beta \Leftrightarrow \alpha xx^{-1}\beta$, where x is a regular edge or inverted edge; the circulation rules become: where x is a regular edge or inverted edge:

$$\alpha x\mu\beta x\gamma \Leftrightarrow \alpha x\beta x\mu^{-1}\gamma, \quad \alpha\mu x\beta x\gamma \Leftrightarrow \alpha x\beta\mu^{-1}x\gamma,$$

$$\alpha x\mu\beta x^{-1}\gamma \Leftrightarrow \alpha x\beta\mu x^{-1}\gamma, \quad \alpha\mu x\beta x^{-1}\gamma \Leftrightarrow \alpha x\beta x^{-1}\mu\gamma.$$

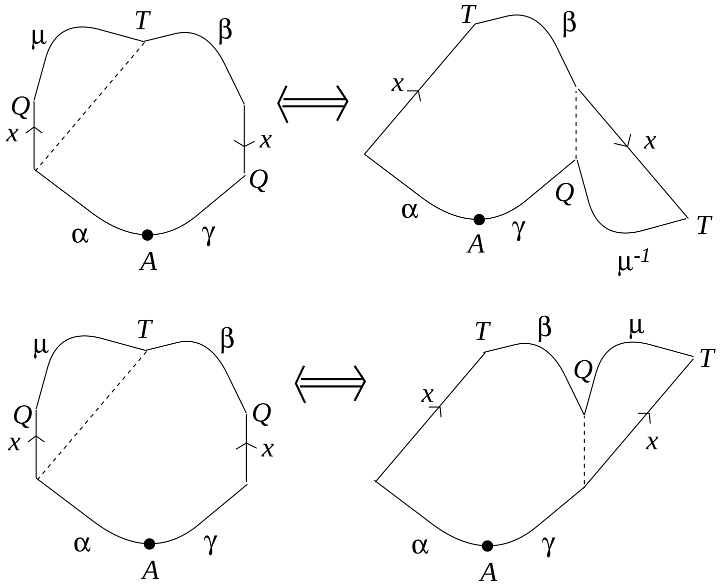

Figure 7: illustrating circulation rules 1a and 2a.

Figure 7 illustrates proofs for circulation rules 1a and 2a; the others can be shown similarly. Whittlesey shows how to apply these rules in a systematic way reducing W to a normal form, such that two 2-complexes are homeomorphic if and only their normal forms are formally equivalent.

First, one combines together polygons as much as possible: if a regular edge is mentioned in two words, combine them according to the polygon composition rule. Then one obtains a list of words in which no two words can be composed through regular edges.

A regular edge occurring twice in a word is called a 'crosscap.' Collect all the crosscaps as follows: given a word (perhaps cyclically permuted) $AxQ\mu AxQ\gamma A$. This matches the left-hand side of circulation rule (1a), and it can be replaced by $AxAxA\mu^{-1}Q\gamma A$. Replace x by yTz, getting $AyTzAyTzA\mu^{-1}Q\gamma A$. Cyclically permute: $TzAyTzA\mu^{-1}AyT$, and apply circulation rule (1a) to get $TzTzTy^{-1}A\mu^{-1}AyT$. Or: $TzTzTy^{-1}A\ldots AyT$. If the ellipsis contains some more crosscaps, they can be brought to the left of the y^{-1} as follows:

Within the word $y^{-1}A\ldots Ay$, if a crosscap $\ldots v\ldots v$ exists, bring the occurrences together as done for x already, yielding $y^{-1}A\ldots MvMvM\ldots Ay$. This matches the pattern $xQ\mu T\beta\ldots Qx^{-1}$, with $x = y^{-1}$, $Q = A$, μ the first ellipsis, $T = M$, $\beta = vMvM\ldots$

Applying circulation rule (2a), we get $y^{-1}MvMvM\ldots My$.

Now, $Ty^{-1}MvMv$ matches the pattern $T\mu Px\beta Px$, so it can be replaced by $TvMyTv$ using circulation rule (1b). Again, $vMyTvM$ can be replaced by $vTvTy^{-1}M$ using circulation rule (1a). Combining together all these transformations we reduce the original word to

$$TzTzTvTvTy^{-1}M\ldots MyT.$$

If further crosscaps exist in the ellipsis, they can be moved to the front in the same way, and we eventually arrive at a word of the form

$$Tz_1Tz_1T\ldots Tz_qTz_qTy^{-1}\ldots yT,$$

where all the q crosscaps are brought to the beginning of the word— there are no further crosscaps in the ellipsis.

Having collected all the crosscaps, we next collect all the 'handles.' A handle is a subsequence

$$\ldots a \ldots b \ldots a^{-1} \ldots b^{-1} \ldots$$

where a and b are regular edges (or inverted regular edges).

The handle can be brought to the front of a word (or a subword!) as follows: beginning with

$$Ay^{-1} \ldots Ta \ldots b \ldots a^{-1}T \ldots b^{-1} \ldots$$

apply circulation rule (2b) to move a to the left:

$$Aa \ldots bU \ldots a^{-1}Ay^{-1} \ldots Ub^{-1} \ldots$$

then apply circulation rule (2a) (reversed):

$$AaV \ldots bAy^{-1} \ldots Va^{-1}Ab^{-1} \ldots$$

Then apply rule (2a), again reversed:

$$AaAy^{-1} \ldots WbAa^{-1}Ab^{-1}W \ldots$$

and finally apply rule (2b) to get

$$AaAbAa^{-1}Ab^{-1}Ay^{-1} \ldots$$

Since the circulation rules used here do not invert symbols, we have not introduced any crosscaps not already there. The process can be repeated until all handles are brought to the front.

If this is applied after first collecting the crosscaps, we obtain a word in the form

$$Tz_1Tz_1T \ldots Tz_qTz_qTa_1Tb_1Ta_1^{-1}Tb_1^{-1}T \ldots Ta_sTb_sTa_s^{-1}Tb_s^{-1}Ty^{-1} \ldots yT$$

where the final ellipsis contains no crosscaps and no handles. Next, if a word contains crosscaps, then all handles can be replaced by pairs of crosscaps, as follows.

$$zz(aba^{-1})b^{-1} \Leftrightarrow zab^{-1}(a^{-1}z)b^{-1} \Leftrightarrow za(z^{-1})ab^{-1}b^{-1} \Leftrightarrow zzaab^{-1}b^{-1}.$$

Thus we can reduce any word to the form $\alpha y^{-1}\beta yA$ where β contains no handles or crosscaps (note: if necessary, there being no crosscaps, y can be introduced by incision) and α is of the form A, $Az_1Az_1A \ldots Az_qAz_qA$ ($q > 0$), or $Aa_1Ab_1Aa_1^{-1}Ab_1^{-1}A \ldots Aa_sAb_sAa_s^{-1}Ab_s^{-1}$ ($s > 0$).

A 'cuff' is a regular edge t such that t and t^{-1} both occur in the above-mentioned suffix β. By hypothesis, two such pairs don't interlace, since otherwise they would make a handle. If there exists a cuff, choose an innermost cuff, and bring it in front of β as follows:

$$y^{-1}\rho t\sigma t^{-1}\tau y \Leftrightarrow y^{-1}t\sigma t^{-1}\rho\tau y \Leftrightarrow t\sigma t^{-1}y^{-1}\rho\tau y.$$

In this way one 'separates the cuffs,' continuing until there is no nested pair $\ldots r \ldots s \ldots s^{-1} \ldots r^{-1} \ldots$.

This reduces our word to the form $\alpha r_1 \beta_1 r_1^{-1} \ldots A r_t \beta_t r_t^{-1} A$, where α is a sequence of crosscaps or handles, and no β_j contains any regular edges. Some of these β_j are just vertices, and the cuff becomes a removable incision: remove all incisions. (Note: an incision rPr^{-1} cannot be removed if P is mentioned elsewhere.)

Within a cuff, the word β_j can be permuted cyclically using the circulation rules (2a).

Then, within each β_j, compose edges as long as possible (a cyclic permutation of the word may allow further composition). This edge composition is on the singular graph; the singular graph may end up with self-loops PeP.

This process, applied to all the words of the system, yields a canonical form for the system.

A word with normal form A is just a 2-sphere.

We now consider a list of *trivial isomorphisms* of canonical word-systems. Note first that if a canonical word has crosscaps and cuffs, any cuff can be inverted; if zz is the rightmost crosscap:

$$zz \ldots r_j \beta_j r_j^{-1} \ldots \Leftrightarrow \ldots zz r_j \beta_j r_j^{-1} \ldots \Leftrightarrow \ldots z r_j \beta_j^{-1} r_j^{-1} z \ldots \Leftrightarrow$$
$$\ldots zz r_j \beta_j^{-1} r_j^{-1} \ldots \Leftrightarrow zz \ldots r_j \beta_j^{-1} r_j^{-1} \ldots$$

We can now summarise the trivial isomorphisms on a canonical word-system:

1. Permute the words in the system.
2. Permute the order of cuffs in any word (circulation rule (2b)).
3. Cyclically permute any word β_j in a cuff.
4. Reorient any edge.
5. Invert any word in the system.
6. Invert any cuff in a 'non-orientable' word, i.e., a word with crosscaps.
7. Relabel any vertex or edge.

This concludes the description. Whittlesey asserts the following theorem:

THEOREM 3.1 *Two complexes (in which every edge and vertex is incident to at least one face) are homeomorphic iff their word-systems are combinatorially equivalent iff their canonical forms are trivially isomorphic.*

We turn to consideration of the complexity of the homeomorphism problem. As a first step,

THEOREM 3.2 *Graph isomorphism is efficiently reducible to homeomorphism of simplicial 2-complexes.*

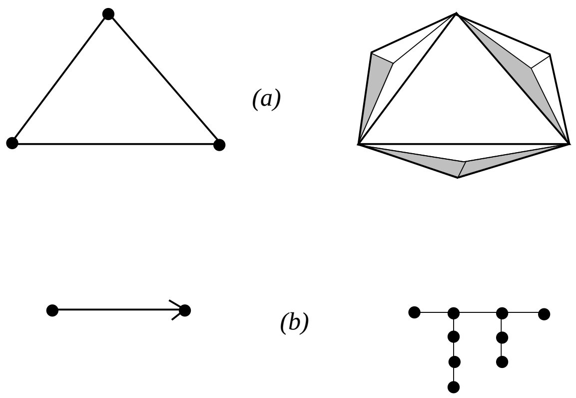

Figure 8: (a) Graph G and corresponding complex K; (b) Directed edge and corresponding 9-node undirected graph.

PROOF (SKETCH). Let G be a graph. We construct a simplicial 2-complex resembling G as follows: for each edge $\{u, v\}$ of G there is a tetrahedron on which two vertices are identified with u and v. The other two vertices are 'internal' to the structure for $\{u, v\}$. Let K be the associated complex. See Figure 8 (a). Clearly, to each vertex u of G incident to more than one edge there corresponds an isolated conical point of K. The other vertices of K (except for isolated point components) are regular. If G has n nodes and m edges then K has $n + 2m$ nodes, $6m$ edges, and $4m$ faces. Let S be the set of isolated conical points of K.

Suppose for brevity's sake that G is connected, and G' is another connected graph with corresponding simplicial 2-complex K', with set S' of isolated conical points. Clearly if G and G' are isomorphic then K and K' are homeomorphic. Conversely, if $f: K \to K'$ is a homeomorphism, it carries S onto S', and every connected component of $X = |K| \backslash S$ onto a connected component of $Y = |K'| \backslash S'$.

These connected components are either (a) a point, in which case G and G' are trivial; (b) a sphere, in which case G and G' are both single edges; or (c) spheres with one or two points deleted. These cases are mutually exclusive, so we consider only case (c), which applies when S and S' are nonempty.

Each edge e of G corresponds to a unique component of X, hence, through f, to a unique edge $g(e)$ of G'. If two edges of G have a node in common, then so do the corresponding edges of G', and it follows that the bijection g defines an isomorphism of the two graphs. ∎

We shall next argue that homeomorphism of simplicial 2-complexes, implemented through Whittlesey's reduction method, can be reduced efficiently to graph isomorphism. We begin with the observation (see [23]) that isomorphism of directed graphs can be reduced efficiently to graph isomorphism, by the construction illustrated in Figure 8 (b).

A second observation is that Whittlesey's reduction procedure can be implemented in polynomial time. The procedure involves composition, a few incisions, and circulation rules; only incision increases the length of the words in the system, and there is at most one per word; and the circulation rules can be applied (to words stored in an array, say) in polynomial time. Therefore it is enough to reduce trivial isomorphism of canonical word-systems to directed graph isomorphism.

Let W be a canonical word-system. Among the trivial isomorphisms one can invert any word. To allow for this, for every word w in W add a new word w' obtained by inverting w and cyclically permuting so that the crosscaps or handles come at the beginning. We assume that W has been augmented in this way.

We construct a directed graph D as follows. D has 5 'levels,' where level 0 comprises the nodes of out-degree zero, and a node in level $r+1$ has an adjacent node in level r but none at any lower level.

To each vertex- and edge-letter in W there corresponds a node at level 0. Let us imagine these nodes as being labelled with the corresponding letters; the labelling is not part of the construction.

Level-0 nodes labelled with vertex labels have unique predecessors at level 1, also imagined labelled with the same vertex labels; level-0 nodes labelled with edge labels a have two predecessors at level 1, labelled with a and a^{-1} respectively. Thus for each vertex- or (inverted) edge-letter in W there corresponds a unique node at level 1.

For each word w in the system there is a set of $|w|$ nodes at level 2, each linked to the level-1 node with the appropriate label. So far, all the linkages have been 'vertical.' At this and the next levels we introduce horizontal linkages reflecting the sequencing of each word, and its cuffs, and within its cuffs.

Consider a fixed word w of the system. Write it in the form $\alpha A\beta$ where α represents the handles or crosscaps (mutually exclusive) and β represents its cuffs. Factorise β as $b_1B_1\beta_1b_1^{-1}A\ldots b_sB_s\beta_sb_s^{-1}A$, a product of its s 'cuffs.'

For each subword β_j, connect the corresponding nodes at level 2 by a single directed cycle (if the cuff has no crosscaps) or a bidirected cycle (if the cuff has crosscaps). This reflects the trivial isomorphisms in which a cuff can be cyclically permuted or inverted.

Again corresponding to the word w there are $|\alpha|+1+5s$ nodes at level 3, which we can imagine as being labelled with a word $\alpha Ab_1B_1c_1b_1^{-1}Ab_2B_2c_2\ldots$

There are vertical links from each, except those corresponding to a c_j (c_j for the j-th cuff in w), to the corresponding nodes at level 2.

For the j-th cuff there are horizontal links connecting the nodes for $b_j, B_j, c_j, b_j^{-1}, A$ in that order, in level 3. There are vertical links connecting c_j with those nodes of β_j with vertex labels. This reflects the possibility of cyclically permuting β_j by an even-phase cycle.

Finally, at level 4, corresponding to the word w there are $|\alpha|+2$ nodes, if w has cuffs, or $|\alpha|+1$ nodes if w has no cuffs. We imagine the nodes being labelled either αAc or αA. There are horizontal links connecting them in that order; there are vertical links connecting c (if it exists) with all the s nodes c_j at level 3, corresponding to the $s \geq 0$ cuffs of w, and connecting the other nodes at level 4 with the corresponding nodes labelled with αA at level 3. See Figure 9.

The graph D has the property that if another directed graph D' represents a canonical word-system W' (augmented in the same way, by inverting words) then the graphs are isomorphic iff the systems are trivially isomorphic.

Summarising,

THEOREM 3.3 *Homeomorphism of simplicial 2-complexes as a decision problem is polynomially equivalent to graph isomorphism.* ∎

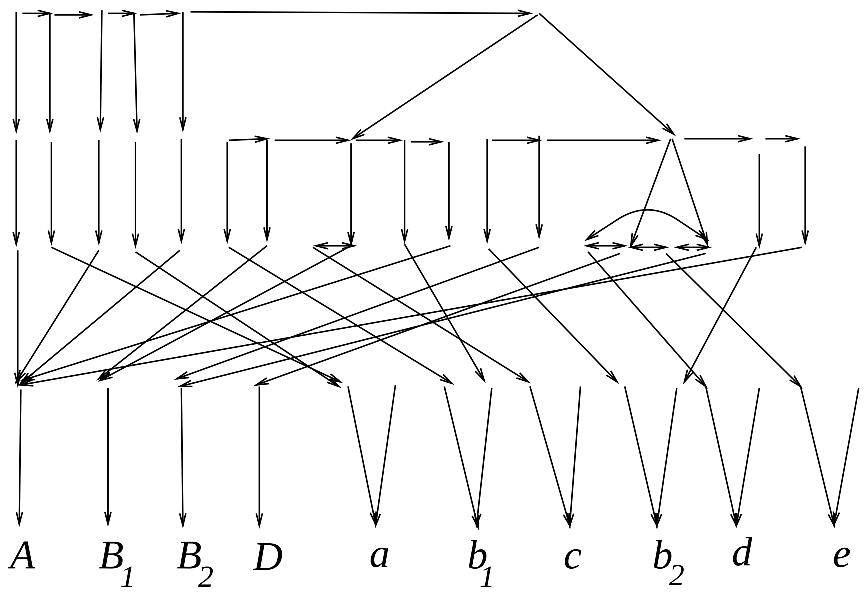

Figure 9: directed graph for the single word $AaAaAb_1B_1cB_1b_1^{-1}Ab_2B_2dDeB_2b_2^{-1}A$.

4 Obstructions to embedding complexes.

This section discusses Shapiro's paper [32] on obstructions to embedding complexes.

4.1 Intersection Numbers

Let σ^p and τ^q be oriented simplices such that $p+q=r$. Let $f\colon\sigma\to\mathbb{R}^r$ and $g\colon\tau\to\mathbb{R}^r$ be continuous maps such that $f(\partial\sigma)\cap g(\tau)$ and $g(\partial\tau)\cap f(\sigma)$ are both empty. Then it is possible to choose an integer which in some sense 'counts' the number of intersections of $f(\sigma^\circ)$ with $g(\tau^\circ)$. This is done by choosing transverse smooth maps $f_1\colon\sigma\to\mathbb{R}^r$ and $g\colon\tau\to\mathbb{R}^r$ which are 'close' to f and g respectively and for which $f_1(\partial\sigma)\cap g_1(\tau)$ and $g_1(\partial\tau)\cap f_1(\sigma)$ are both empty. Then the number of points $(x,y)\in\sigma\times\tau$ for which $f_1(x)=g_1(y)$ is finite and to each such point we assign the value 1 if the sum of the orientations of $f_1(\sigma)$ and $g_1(\tau)$ at this point is equal to the (standard) orientation of $\mathbb{R}^r$ or -1 otherwise.

DEFINITION 4.1 *The intersection number $I(f,g)$ of f and g is then equal to the sum of these integers. If $c_p=\sum_{i=1}^s a_if_i$ and $c_q=\sum_{j=1}^t b_jg_j$ are p- and q-chains respectively with the property that $I(f_i,g_j)$ is defined for each pair (i,j), then the intersection number $I(c_p,c_q)$ is defined by*

$$I(c_p,c_q)=\sum_{1\le i\le s,1\le j\le t} a_ib_jI(f_i,g_j).$$

With this definition one can verify that intersection numbers have the following properties:

(i) $I(f,g)=(-1)^{pq}I(g,f)$.

(ii) $I(f\circ\alpha,g\circ\beta)=(-1)^{(\mathrm{sign}(\alpha))(\mathrm{sign}(\beta))}I(f,g)$, where $\alpha\colon\sigma_1\to\sigma$ and $\beta\colon\tau_1\to\tau$ are simplicial isomorphisms.

(iii) There is some $\eta > 0$ such that if f_1 and g_1 are η-displacements of f and g respectively, then $I(f_1, g_1)$ is defined and $I(f, g) = I(f_1, g_1)$.

(iv) suppose that $f_1 : \sigma \to \mathbb{R}^r$ and $g_1 : \tau \to \mathbb{R}^r$ are oriented singular p- and q-simplices respectively in $\mathbb{R}^r$, and that $I(f_1, g_1)$ is defined. Suppose also that F is a homotopy between f and f_1 and G is a homotopy between g and g_1 with the property that $I(F_t, G_t)$ is defined for each $t \in [0, 1]$. Then $I(f, g) = I(f_1, g_1)$.

(v) Let $f : \sigma^p \to \mathbb{R}^r$ and $g : \tau^q \to \mathbb{R}^r$ be oriented singular simplices in $\mathbb{R}^r$ with $p + q = r + 1$. If the intersection number $I(f|_\alpha, g|_\beta)$ is defined for each pair of faces α and β of σ and τ respectively with $\dim(\alpha) + \dim(\beta) = r$, then

$$I(f, \partial g) = (-1)^p I(\partial f, g).$$

4.2 The Primary Obstruction

Let K be a finite simplicial complex. Assume that the simplices of K have been oriented and give the product cells in $K \times K$ the product orientation. Denote by K^* the subcomplex of the cell complex $K \times K$ consisting of those product cells $\sigma \times \tau$ for which σ and τ are simplices in K having no vertices in common. Let $T : K^* \to K^*$ be defined by $T(x, y) = (y, x)$ for all $(x, y) \in K^*$. Then T is a cellular automorphism and induces a cochain omplex isomorphism $T^\# : C^*(K^*) \to C^*(K^*)$. For each positive integer r let A_r^* denote the subcomplex of $C^*(K^*)$ given by

$$A_r^i = \{c \in C^i(K^*) : T^\# c = (-1)^r c\}$$

It is easy to verify that the sequence of groups $\{A_r^i : i = 0, 1, 2, \ldots\}$ together with the restricted coboundary operator δ does indeed define a subcomplex of $C^*(K^*)$.

Consider now the question of measuring how close a map $f : K \to \mathbb{R}^r$ comes to being an embedding. One way to do this would be to count the points (x, y) for which $f(x) = f(y)$ (provided there are only a finite number of such points). A cruder version of this would be to count the number of points in $f(\sigma) \cap f(\tau)$ for each pair of open simplices σ and τ in K. Now if σ and τ have no vertices in common, $\dim(\sigma) + \dim(\tau) = r$, and both $f(\partial\sigma) \cap f(\tau)$ and $f(\partial\tau) \cap f(\sigma)$ are both empty, then the intersection number $I(f|_\sigma, f|_\tau)$ achieves this for us. This is a motivation for the following definitions.

DEFINITION 4.2 *Let $\mathcal{F}(r)$ denote the set of continuous maps $f : K \to \mathbb{R}^r$ which have the property that $f(\sigma) \cap f(\tau) = \emptyset$ whenever σ and τ are disjoint closed simplices in K for which* $\dim(\sigma) + \dim(\tau) = r$. *For $f \in \mathcal{F}(r)$ define a cochain $m_f^r \in C^r(K^*)$ by*

$$< m_f^r, \sigma \times \tau >= (-1)^q I(f|_\sigma, f|_\tau)$$

for each pair of disjoint closed oriented simplices σ and τ in K with $q = \dim(\tau)$ and $\dim(\sigma) + q = r$.

Using the properties of intersection numbers listed above, it is straightforward to verify that m_f^r is in fact a cocycle in A_r^r. Thus it represents a cohomology class in the cohomology group $H^r(A_r^*)$. In addition it can be shown that this cohomology class is independent of the particular map $f \in \mathcal{F}(r)$ used to define it. This cohomology class is denoted $m^r(K)$ and is in fact a topological invariant of the space $|K|$ called the *primary obstruction* to embedding K in $\mathbb{R}^r$.

In the statement and proof of the main theorem of Shapiro's paper [32], the author uses a restricted subset of the set of maps $\mathcal{F}(r)$ which he calls 'general maps.' The following definition [39] is stronger than his.

DEFINITION 4.3 *A* general map $f: K \to \mathbb{R}^r$ *is a continuous map with the following properties:*

(1) *For each open simplex σ in K, $f|\sigma$ is a smooth embedding.*

(2) *If σ and τ are open simplices in K and $\dim(\sigma) + \dim(\tau) \leq r$, then $f(\sigma)$ intersects $f(\tau)$ transversally in $\mathbb{R}^r$.*

(3) *If σ^p and τ^q are open simplices in K and $p + q = r$, then no point of $f(\sigma) \cap f(\tau)$ lies in the image of any other open simplex.*

(4) *For each vertex v of K, there is a neighbourhood U of v in $|K|$ such that for each $x \in U$, $f^{-1}f(x) = \{x\}$. (If $\dim(K) < r$ it is sufficient that f be injective on U. Then using (2), one can shrink U so that $f^{-1}f(x) = \{x\}$ for each x in this new neighbourhood.)*

(5) *For each closed simplex σ in K, the restriction $f|_\sigma$ is smooth except possibly at the vertices of σ.*

It is straightforward to deduce from the first 3 properties that if $f: K \to \mathbb{R}^r$ is a general map, then for any pair of simplices σ^p and τ^q in K with $p + q = r$, $f(\sigma) \cap f(\tau)$ is a finite set of points. Therefore the algebraic intersection number, denoted $I(f|_{\sigma^p}, f|_{\tau^q})$, is well-defined and given by the sum of the orientation numbers of $f(\sigma)$ and $f(\tau)$ at their points of intersection.

Now if $f: K \to \mathbb{R}^r$ is a general map with the property that $f(\sigma) \cap f(\tau)$ is empty whenever σ and τ are open simplices of K with $\dim(\sigma) + \dim(\tau) = r$ then it is clear that m^r_f vanishes and hence $m^r(K)$ is zero. In his paper, Shapiro proves the following converse of this observation.

THEOREM 4.4 *If $\dim(K) \leq r - 3$, the vanishing of the primary obstruction $m^r(K)$ is a necessary and sufficient condition for the existence of a general map $f: |K| \to \mathbb{R}^r$ which has the property that for each pair of open simplices σ^p and τ^q in K, with $p+q = r$, $f(\sigma^p) \cap f(\tau^q) = \emptyset$.*

Remark: If $n = \dim(K) \geq 3$, then a map $f: K \to \mathbb{R}^{2n}$ obtained by applying the above theorem is an injective map and hence, by compactness of $|K|$, it is an embedding. If we examine the cases $\dim(K) = 1$ or 2 then the methods used in the above proof are not applicable. However, in the case $\dim(K) = 1$, it is straightforward to show that $m^2(K_5)$ and $m^2(K_{3,3})$ are both nonzero. Using the functorial properties of these obstructions and Kuratowski's result on the planarity of graphs, one can show that a 1-dimensional simplicial complex is planar if and only if the primary obstruction to embedding it in $\mathbb{R}^2$ vanishes. (This latter assertion is also valid if all the calculations are carried out mod(2). Wu showed this independently of Shapiro's work [44].)

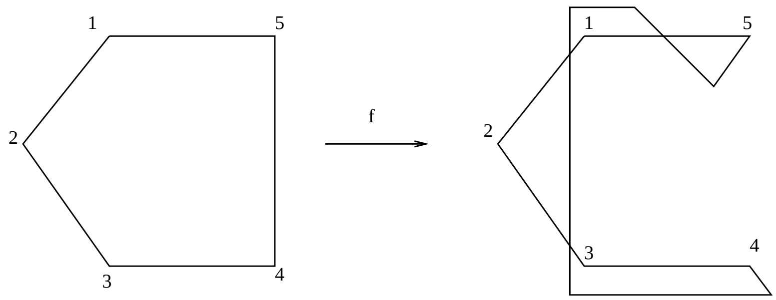

Figure 10: An example

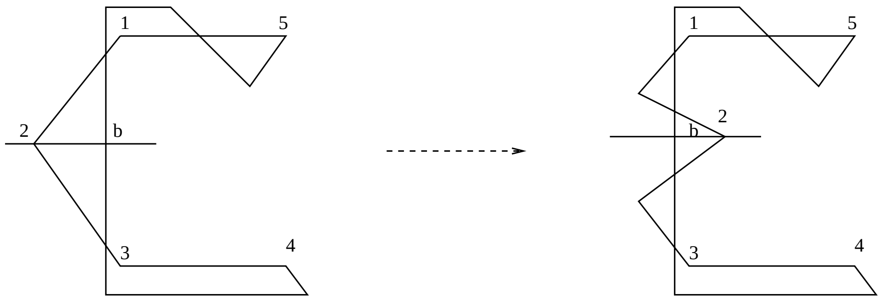

Figure 11: The first alteration

Although the methods in the proof of the above theorem are not always applicable to the problem of embedding a 1-complex in $\mathbb{R}^2$, they can be illustrated by a simple example, that of embedding a pentagon in the plane starting with the map indicated in diagram 10. We know that this complex is planar and therefore its primary obstruction vanishes.

All the maps encountered will satisfy the additional requirement of being 'radial at all of its vertices'.

DEFINITION 4.5 *A map $f\colon |K| \to \mathbb{R}^r$ is called radial near the vertex $v \in K$ if there is a neighbourhood U of v in $|K|$ such that for each $x \in U$ the restriction of f to the line segment $[v, x]$ is linear.*

For the map f indicated, a straightforward calculation shows that

$$m^r_f = -b(23, 45) + b(12, 45) = \delta b(2, 45)$$

The first step is to alter f to obtain a map g for which m^r_g is zero. (In the general case we start with a general map from K to $\mathbb{R}^r$ which is radial at the vertices, and alter it to obtain a general map g which is also radial at the vertices and for which $m^r_g = 0$.) To do this we choose points a and b interior to the simplices 2 and 45 respectively and which do not lie in the image of any other simplices. We join these points by an open arc which doesn't meet the image of the complex any where else (this isn't always possible in low dimensions). The point a is pushed along this arc past b and this deformation is extended to a small neighbourhood of a as indicated in diagram 11 (in this case our extension is radial at a— if a is not a 0-dimensional simplex then the extension is smooth in a neighbourhood of a).

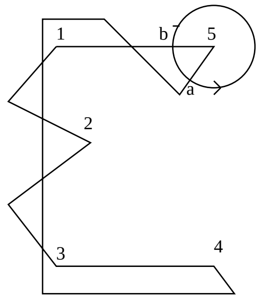

Figure 12: alteration of g.

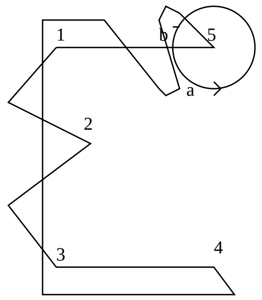

Figure 13:

It is straightforward to show that $m_f^r - m_g^r = \delta b(2, 45)$. (In the general case we need to be able to produce a negative sign if necessary; this is done by adding a twist to deformation, for which the extra room of higher dimensions is needed.)

Next we alter g so that for each pair of open simplices σ^p and τ^q in K with $p + q = r$, the algebraic intersection number $I(g|_\sigma, g|_\tau)$ is zero. In our example the only nonzero intersection number occurs for the pair 15 and 45. First choose a small circle about 5 inside which g is radial. Let a and b be the two points in which this circle meets $g(K)$. (In the general case we choose a small $(r-1)$-sphere about one of the vertices. This sphere meets the image of K in two disjoint submanifolds and we choose a in one of these and b in the other.) Join a to b by the open arc indicated in diagram 12.

Push a past b along this arc and extend this deformation radially inside the circle and smoothly in a small neighbourhood of a outside the circle (see diagram 13).

This increases the intersection number $I(g|_\sigma, g|_\tau)$ by 1 so that all intersection numbers now vanish. (In the general case we need the extra dimensions to allow us to choose whether these intersection numbers are altered by 1 or -1.) Now for each pair of intersecting simplices of K we can pair off each intersection giving rise to a $+1$ with one giving rise to a -1. For each such pair we insert a disc as indicated by the shaded regions in diagram 14.

(In general we need the extra space of higher dimensions to be able to construct these discs.) For each such disc we deform the image of K as indicated by the arrows to eliminate the corresponding pair of intersections (see diagram 15)

When this has been carried out for each disc, we obtain the embedding whose image is shown in diagram 16.

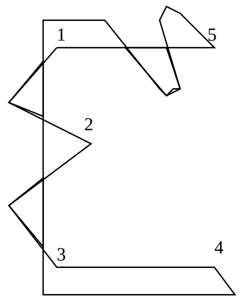

Figure 14:

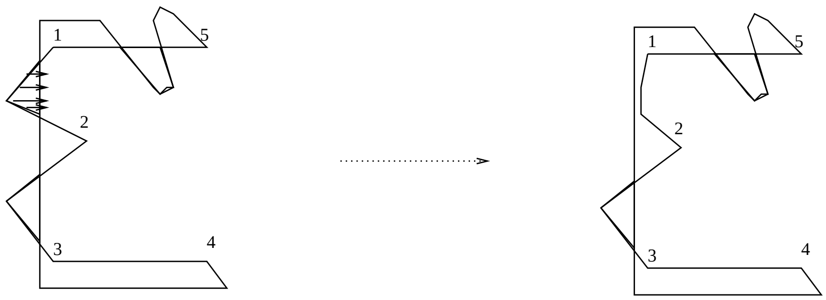

Figure 15:

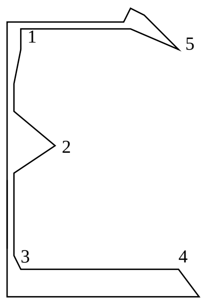

Figure 16:

4.3 Outline of Proof of the Embedding Theorem

Sketch of Proof. Assume the first obstruction for embedding $|K|$ in $\mathbb{R}^r$ vanishes. Choose a general map $f: |K| \to \mathbb{R}^r$ which is radial near each vertex of K. Then the obstruction cocycle m_f^r represents the zero cohomology class. The first step is to replace f by a general map g (also radial near the vertices) for which $m_g^r = 0$. The basic construction to achieve this is to alter f in such a way that m_f^r is modified by adding a coboundary of the form $\pm\delta b(\sigma, \tau)$ where the dimensions of σ and τ add to $r-1$. To do this join a fixed point a in the interior of σ to a fixed point b in the interior of τ by a path which meets $f(|K|)$ nowhere else and which is nowhere tangent to the image of a simplex of K. (Both a and b lie in the interior of this path). Push a past b along this curve and extend this to a deformation of a small neighbourhood of a in $|K|$. This has the required effect (with the alternate sign achievable by putting a twist in this deformation).

The process used to obtain the map g may have introduced intersections between the images of certain pairs of simplices which share exactly one vertex and whose dimensions add to r. Our next step is to introduce more intersections of this type (if necessary) so that the algebraic intersection number for each such pair of simplices is zero. If (σ, τ) is one such pair let v denote the shared vertex. Choose a small neighbourhood U of v on which g is radial and let S be a sphere centred at $g(v)$ and contained in $g(U)$. Then $A = g(\sigma) \cap S$ and $B = g(\tau) \cap S$ are smooth submanifolds of S. Fix $a \in A$ and $b \in B$ and join a to b by a curve in S which meets A only once and at a (an interior point of the curve) and similarly meets B only once and at b (also interior to the curve). We choose this curve to be transverse to A and B. Push a along the curve past b and extend this smoothly to a deformation of a small neighbourhood of a in A (the deformed A will be disjoint from B). This is extended radially inside S and smoothly in a small region of $g(\sigma)$ outside S. The result is a new general map, radial at the vertices, with exactly one intersection more than g had: this new intersection occurs between the images of σ and τ. We can control whether this adds or subtracts 1 from the old intersection number of σ and τ. Hence in a finite number of steps we can obtain a general map h, radial at the vertices and for which the intersection number of any pair of simplices (whose dimensions add to r) is zero.

To conclude, for each pair (σ^p, τ^q) of simplices of K with $p + q = r$ we pair off each $+1$ intersection with a -1 intersection and eliminate them together. The model for this process is the circular arcs P and Q centred at $(-\frac{1}{2}, 0)$ and $(\frac{1}{2}, 0)$ respectively, each of radius 1 and each contained in the closed unit disc. It is clear that we may push P across Q in the negative x direction (leaving it unchanged near the boundary of the disc) and hence eliminate the two intersection points of P and Q. This is carried over to our problem by embedding the disc in $\mathbb{R}^r$ so that

(i) P is carried into $h(\sigma)$ and Q into $h(\tau)$

(ii) the image of the disc contains no intersection points other than the image of $P \cap Q$

(iii) no other point of the disc maps into a point on the image of a closed $(r-3)$ simplex of K, and

(iv) the image of the disc is nowhere tangent to the image of an $(r-3)$-simplex in K.

The intersections of P and Q are mapped to a pair of intersections of $h(\sigma)$ and $h(\tau)$ having opposite signs. We carry over the deformation of P in the disc to deform the corresponding arc in σ and extend this smoothly to a small neighbourhood of this arc in $|K|$. The result is that we eliminate the pair of intersections of $h(\sigma)$ and $h(\tau)$. We carry this procedure out for each such pair of intersections and hence obtain our required map.

4.4 A word on the 4-dimensional case

In 4 dimensions the vanishing of the first obstruction is a necessary condition for the embedding of a 2-complex. However the method of proof that shows the vanishing of the first obstruction is sufficient when the codimension is at least 3 is only partly applicable: we can achieve a map f for which the intersection numbers $I(f|_{\sigma^p}, f|_{\tau^q})$ vanish whenever $p+q=4$ (so $p=q=2$ for a 2-complex). However, when we try to position a disc in order to eliminate a pair of intersections between two 2-simplices, we cannot guarantee that it can be kept clear of the rest of the image of K. Consequently we may introduce new intersections when we try to eliminate old ones by means of such a disc. However if the first obstruction to embedding a 2-complex in $\mathbb{R}^4$ does vanish then it is possible to add a finite number of handles to the complex and hence obtain a new complex which can be embedded in $\mathbb{R}^4$. Each such handle is added to a single face of the old complex so we don't get handles straddling two faces of our old complex.

4.5 Second Obstruction

The first obstruction was obtained by showing that the embedding problem implied the existence of a nonzero section of a certain vector bundle over $|\bar{K}^*|$. For an n-complex K, the space $|\bar{K}^*|$ is 2n-dimensional and if we are examining the problem of embedding in $\mathbb{R}^r$ the vector bundle will have rank r. The first obstruction then tells if we can construct a nonzero section over the r-skeleton of $|\bar{K}^*|$. If this is possible the second obstruction then tells if such a section can be extended to a nonzero section over the $r+1$-skeleton of $|\bar{K}^*|$. Assuming the second obstruction vanishes, as in the case of the first obstruction we still have to do a lot of work to produce a map of the required type. A rough idea of the procedure follows. We start with a map in $\mathcal{F}(r)$ for which the images of any p- and q-cells in $\mathbb{R}^r$ are disjoint whenever $p+q=r$. Then the images of a p'-cell and a q'-cell (where $p'+q'=r+1$) will typically intersect in a finite number of disjoint simple closed curves. The model for eliminating such an intersection curve is called a *deformation cell* and is described as follows. Consider the closed unit ball B in $\mathbb{R}^3$. The spherical caps $P=\{x : |x-(-\frac{1}{2},0,0)|=1\}$ and $Q=\{x : |x-(\frac{1}{2},0,0)|=1\}$ intersect in a circle S and it is clear that we may push P (say) past Q in B in the negative x_1 direction and hence eliminate this intersection curve. We fit such a deformation cell such that S coincides with an intersection circle of one p'- and one q'-cell, P lies entirely in the p'-cell and Q in the q'-cell and the rest of the deformation cell is disjoint from the image of our simplicial complex. By extending the deformation described in B to a small neighbourhood of the deformation cell in $\mathbb{R}^r$ we eliminate one intersection circle. Repeating this process, we eventually get to the stage where the images of any pair of simplices of K will be disjoint provided their dimension sum is no greater than $r+1$.

5 Embedding graphs in $\mathbb{R}^3$.

In several recent papers [34, 35, 36, 37], Taniyama has analysed embeddings of graphs in $\mathbb{R}^3$. Two graph embeddings may be considered equivalent under a range of criteria ranging from ambient isotopy to $\mathbb{Z}_2$-homology, and he has classified these in (partial) order of increasing strength.

For the weakest two, homology and $\mathbb{Z}_2$-homology, Taniyama describes a cohomology invariant ascribed to Wu [42] which is enough to classify embeddings. There are points in common with Shapiro's work described in these notes.

Let f be an embedding of a graph G in $\mathbb{R}^3$. For any space X let $C_2(X)$ be the set of ordered pairs of distinct elements of X (equivalent to X^* in the previous section), and let σ be the involution which swaps coordinates: $\sigma(x, y) = (y, x)$.

Let $A_p(X, \sigma)$ be the $\mathbb{Z}$-module of singular p-chains a in $C_2(X)$ such that $\sigma a = -a$. In other words, if $a = \sum \alpha_j \sigma_j$ where each σ_j is a continuous map of the standard p-simplex into X, then $\sigma a = \sum \alpha_j \sigma \circ \sigma_j$.

Then $\{A_p(X, \sigma) \colon p \in \mathbb{N}\}$ forms a chain complex (with usual boundary operator). Let $H^2(C_2(X), \sigma)$ be the second cohomology group of this complex.

Similarly, let $\sigma \colon S^2 \to S^2$ be the antipodal map ($x \mapsto -x$); define $A_p(S^2, \sigma)$ as those singular p-chains a such that $\sigma a = -a$; this defines a chain complex; let $H^2(S^2, \sigma)$ be its second cohomology group.

An embedding $f \colon G \to \mathbb{R}^3$ defines a map $f^2 \colon C_2(G) \to C_2(\mathbb{R}^3)$, $(x, y) \mapsto (f(x), f(y))$, with the property that $\sigma \circ f^2 = f^2 \circ \sigma$. Hence there is a map $(f^2)^{\#} \colon H^2(\mathbb{R}^3, sigma) \to H^2(G, \sigma)$.

PROPOSITION 5.1 [42, 37] *The map*

$$(x, y) \mapsto \frac{x - y}{|x - y|}$$

is an equivariant homotopy equivalence between $C_2(\mathbb{R}^3, \sigma)$ and (S^2, σ). ∎

The word 'equivariant' means 'commutes with σ,' implying that the groups $A_p(\mathbb{R}^3, \sigma)$ and $A_p(S^2, \sigma)$ are isomorphic. Hence $H^2(\mathbb{R}^3, \sigma)$ is isomorphic to $H^2(S^2, \sigma)$.

PROPOSITION 5.2 [42, 37] $H^2(S^2, \sigma) \equiv \mathbb{Z}$. ∎

The composition of these isomorphisms yields Wu's invariant $\mathcal{L}(f)$. Let τ be a fixed generator for $H^2(S^2, \sigma)$ and i the isomorphism with $H^2(\mathbb{R}^3, \sigma)$. Then

$$\mathcal{L}(f) = (f^2)^{\#}(i\tau).$$

Thus $\mathcal{L}(f)$ is a class in $H^2(G, \sigma)$. It can be calculated explicitly (if f is piecewise linear), by methods similar to those of [39]. The significance of the invariant is:

PROPOSITION 5.3 (Main Theorem of [37].) *Two (piecewise linear) embeddings f and g of a graph G are homologous iff $\mathcal{L}(f) = \mathcal{L}(g)$.* ∎

(The constructions and definitions for $\mathbb{Z}_2$-homology are very similar.) In [35] Taniyama introduces another invariant of 3-dimensional graph embeddings, which is an invariant under the rather stronger equivalence of 'link homotopy.' It is defined in terms of knot polynomials.

6 A simplicial 2-complex with planar 1-skeleton is embeddable in $\mathbb{R}^3$.

THEOREM 6.1 *If the 1-skeleton of a 2-dimensional simplicial complex is planar, then that complex can be embedded (linearly) in 3-space.*

Proof of Theorem 6.1 will be given below, following a few lemmas. The curve parametrised by (t, t^2, t^3) is called a *G-curve* by Shapiro [32], or sometimes the 'moment curve' in $\mathbb{R}^3$ [11]. Given a sequence of five points on the curve, in vertical order, three of them labelled a and two of them labelled b, the sequence 'interlocks' if it has the pattern *ababa*.

LEMMA 6.2 *A triangle and a line segment with distinct vertices along the curve $t \mapsto (t, t^2, t^3)$ will intersect if and only if the vertices 'interlock' along the curve.*

PROOF. (This sketch-proof is fleshed out in [39].) Let the triangle vertices be labelled a and the segment vertices be labelled b. Suppose that the triangle and the line segment do not intersect. Then convexity implies the existence of a hyperplane separating them. The points at which the curve crosses the hyperplane would be roots of a cubic polynomial. There can be at most three such points, hence the vertices of the line segment cannot interlock with the vertices of the triangle in order along the curve: because if the sequence were *ababa* then the curve would intersect the hyperplane at least four times, once between each consecutive pair.

Conversely suppose that the vertices do not interlock. Now no two line segments with vertices along the curve will intersect. (Otherwise the curve would contain four coplanar points, which is impossible.) The fact that the vertices do not interlock implies that one of the vertices of the triangle can be moved along the curve (like a bead along a string) to bring it arbitrarily close to another vertex of the triangle. For example, if the sequence were *abaab* then the second and third triangle vertices could be made arbitrarily close. In this way the triangle can be made to lie within an arbitrarily small neighbourhood of a line segment, and this line segment will not intersect the given line segment, so the given triangle and line segment cannot intersect. ∎

PROOF OF THEOREM 6.1. Suppose that the 1-skeleton G of the complex is a planar graph, i.e. G can be embedded in the plane. Choose such an embedding. The embedding induces a cyclic ordering (anticlockwise, say) of the edges around each vertex.

For simplicity, and with negligible loss of generality, assume that G is connected. In order to get (linear) embeddability of the 2-complex itself into 3-space, it suffices, by the above lemma, to find an ordering of its vertices such that the vertices of a triangle of the complex do not interlock with the vertices of any edge in this ordering.

Apply depth-first search [26] to G, but with this restriction: when a non-root vertex v is first visited, process the edges incident to v in anticlockwise order around v, beginning with the edge following (v, u) in anticlockwise order, where u is the parent of v in the depth-first-search tree. Assign preorder numbers to the vertices as they are visited.

It is enough to show that if A, B, C, D, E are five vertices, where ACE is a triangle (i.e., AC, CE, AE are all edges of G) and BD is an edge then the preorder ranking $A < B < C < D < E$ is impossible. Because then the preorder numbering will avoid the interlocking discussed above.

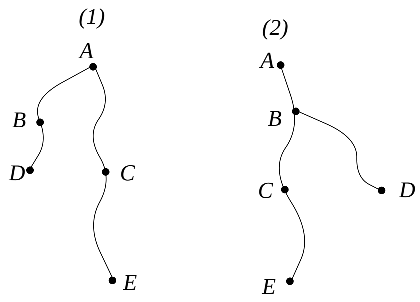

Figure 17: Theorem 6.1, cases (1) and (2).

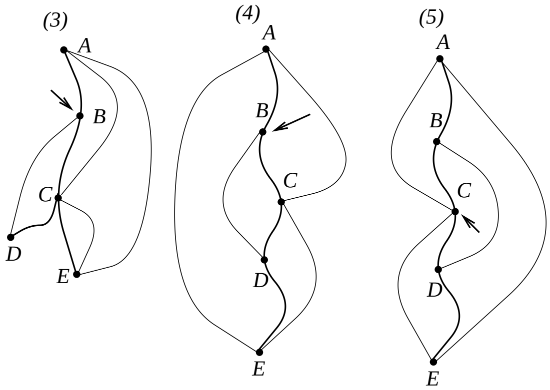

Figure 18: Theorem 6.1, cases 3, 4, and 5. The arrows show where the anticlockwise rule has been broken.

So suppose that $A < B < C < D < E$ in preorder, and the four given edges are edges of G: we argue a contradiction. Depth-first search trees have the property that if x is a vertex of G and X is the set of vertices visited earlier than x (i.e., with lower preorder rank), then the descendants of x are all vertices connected to x in $G \backslash X$. Also, the descendants of x have contiguous preorder ranks.

It follows here that, since AC, BD, and CE are edges of G and $A < B < C < D < E$, that B is an ancestor of D and C an ancestor of E, and they are all descendants of A (because E is and all the others are ranked between A and E).

Case (1): B is not an ancestor of C. Then $B < D < C$.

Case (2): B is an ancestor of C, but D is not a descendant of C. Then $C < E < D$. See Figure 17.

The other three cases use the anticlockwise ordering of edges visited in depth-first search. See Figure 18.

Case (3): B is an ancestor of C and D a descendant of C, but E is not a descendant of D. Assuming the anticlockwise condition applies at C, D is to the left of the $C \to E$ branch of the depth-first search tree (this ordering is meaningful because the cycle $ABCEA$ exists). Then the back-edge DB is to the left of the tree path $B \to C$, and the anticlockwise condition is violated at B.

Cases (4) and (5): A, B, C, D, E all lie along a branch of the depth-first search tree. Relative to the simple cycle $ABCDEA$, the edges AC, CE, and BD form 'bridges' [26], and BD

interlaces with the other two. This means that either (4) BD is to the left of the branch $ABCDE$ and the other two edges to its right, or (5) vice-versa. In case (4), the anticlockwise condition is violated at B and in case (5) at C. This concludes the proof. ∎

7 Embeddability for another restricted class of complexes

We give a necessary homological condition for embeddability of a simplicial complex K in $\mathbb{R}^3$.

Our calculations will invoke the following isomorphism (17, 23.18]): assuming that $H_{q-1}(X)$ is finitely generated,

$$H^q(X) \equiv F_q \oplus T_{q-1}, \tag{1}$$

where T_r is the torsion submodule of $H_r(X)$ and F_r its free part (the quotient H_r/T_r).

In particular, since H_0 is torsion-free, $H^1(X)$ is torsion-free.

LEMMA 7.1 *If K is embeddable in $\mathbb{R}^3$ then its first homology group is (torsion-) free.*

PROOF. We can replace $\mathbb{R}^3$ with its compactification S^3. Let us assume that the embedding in S^3 is sufficiently well-behaved to make Alexander duality work with the singular cohomology groups. For instance, we could assume the embedding is piecewise linear. Let K' be a nice open neighbourhood of the embedding of K, so that K' is homeomorphic to $K \times (0,1)$, and let E be the (closed) complement of K'. Alexander's duality theorem ([17, 27.5]) implies that

$$H^q(E) \equiv H_{3-q}(S^3, K').$$

From the exact homology sequence,

$$\ldots \to H_2(S^3) \to H_2(S^3, K') \to H_1(K') \to H_1(S^3) \to \ldots$$

we get an isomorphism

$$H_2(S^3, K') \equiv H_1(K') \equiv H_1(K).$$

By Alexander Duality,

$$H_{3-q}(S^3, K') \equiv H^q(S^3 \backslash K'),$$

and when $q = 1$ the latter group is torsion-free; hence $H_1(K)$ is torsion-free. ∎

SOME EXPLICIT CALCULATIONS FOR A 'SHELL.' By 'shell' we mean a 2-complex formed by gluing the boundary S^1 the closed unit ball B to the 1-skeleton S^1 by a continuous map f. For example, if $f(\theta) = 2\theta$ the space is the projective plane.

Let X be such a space, and choose an excisive couple A, B in X as the images under f of the open ball of radius $2/3$ and the half-open annulus of inner radius $1/3$ and outer radius 1 respectively. The intersection is an open annulus.

Part of the Meyer-Vietoris sequence (for reduced homology) can be transcribed:

$$\ldots H_2(A) \oplus H_2(B) \to H_2(X) \to H_1(A \cap B) \to H_1(A) \oplus H_1(B) \to H_1(X) \to H_0^{\#}(A \cap B) \ldots$$

This is

$$0 \to H_2(X) \to \mathbb{Z} \to 0 \oplus \mathbb{Z} \to H_1(X) \to 0$$

where the third isomorphism maps i to $(0, -\mu i)$, where μ is the degree of f. Hence $H_1(X) \equiv \mathbb{Z}/(\mu\mathbb{Z})$, and this is torsion-free iff $\mu = \pm 1$ or $\mu = 0$.

If $\mu \neq 0$ then the third homomorphism is injective, so $H_2(X) \equiv 0$. Otherwise the third homomorphism is trivial and $H_2(X) \equiv \mathbb{Z}$. Let us resume the calculations involving Alexander duality and various sequences; write K for X, and identify K with its embedded copy in S^3. We assume that $H_1(K)$ is torsion-free.

By Alexander duality,

$$H^2(K) \equiv H_1(S^3, S^3 \backslash K)$$

and by the formula (1), assuming $H_1(K)$ is finitely generated, which it is,

$$H^2(K) \equiv H_2(K) \oplus T_1,$$

hence

$$H_2(K) \equiv H_1(S^3, S^3 \backslash K).$$

By the exact homology sequence for reduced groups,

$$\ldots \to H_1(S^3) \to H_1(S^3, S^3 \backslash K) \to H_0^{\#}(S^3 \backslash K) \to 0$$

we get an isomorphism between the middle two groups; hence

$$H_2(K) \equiv H_0^{\#}(S^3 \backslash K).$$

Consequently, if $\mu = \pm 1$, $H_0^{\#}$ is trivial and the embedded K does not enclose space; otherwise, when $\mu = 0$, the groups are isomorphic to $\mathbb{Z}$ and the embedded K encloses space. We can summarise all of this in the following:

THEOREM 7.2 *Let U be the unit disc in $\mathbb{R}^2$, S^1 its boundary. For the case where the complex K is the adjunction space of a map $f: \partial U \to S^1$, a necessary condition for K to be embeddable is that f have degree 0 or ± 1. If the degree is 0, then the embedded K encloses space; otherwise it doesn't.* ∎

It turns out that the condition given is also sufficient for this very restricted kind of space. That is, if f has degree 0, then there exists a (rather uninteresting) embedding of K in $\mathbb{R}^3$. The embedding amounts to choosing an arc on the 2-sphere S^2 and gluing the arc to a ring in 3-space.

The construction is as follows. Given the adjunction map $f: S^1 \to S^1$, it can be 'lifted' to a map $\ell: [0, 1] \to \mathbb{R}$ where $\ell(0) = 0$ and $\ell(1)$ is the degree of f, assumed to be zero. (Explicitly, $f(e^{2\pi it}) = e^{2\pi i\ell(t)}$.) Choose t_0 so that $\ell(t_0)$ is minimal, and define $h(t) = \ell((t + t_0) \bmod 1) - \ell(t_0)$. Because f has degree 0, h is also a continuous map, which we can view as another lifting of f — only the 'phase' has been altered — with the property that $h(t)$ is always nonnegative and $h(0) = h(1) = 0$. Let $[0, t_m]$ be the range of h and construct an embedded copy of the adjunction space of U with $[0, t_m]$ under the continuous map $e^{2\pi it} \mapsto h(t)$. Call this space Z. To embed Z, identify $[0, t_m]$ with a line J in the xy-plane; identify the *graph* of the map h with a closed curve-segment in the plane $z = 1$; complete this graph to a Jordan curve in the same plane; connect points on this Jordan curve to J by line-segments as follows: if a point is on the copy of the graph, with parameter t, connect it to the point of J corresponding to

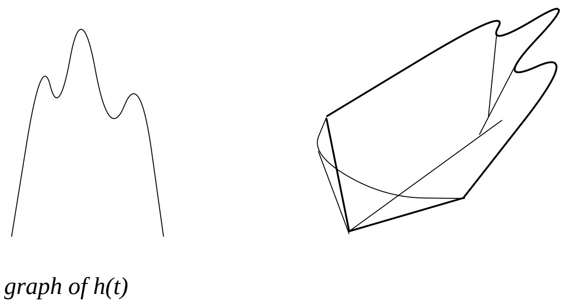

Figure 19: Embedding the adjunction space Z.

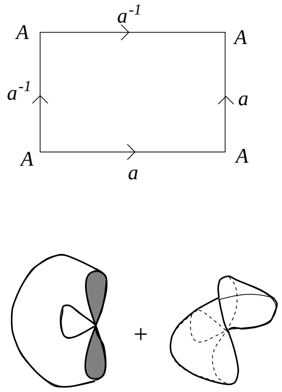

Figure 20: Alternative embedding.

$h(t)$; otherwise connect it to the point of J corresponding to 0. See Figure 19. Then the union of the surface swept out by these line-segments, together with the inside of the Jordan curve in the plane $z = 1$, is the embedding of Z. Clearly it is homeomorphic to a 'pleated sphere.' Finally, to obtain the embedding we desire, wrap this pleated sphere around a suitable copy of S^1 in $\mathbb{R}^3$ so that a point $h(t)$ on J maps to the point corresponding to $e^{2\pi i h(t)}$ on this circle.

In all such embeddings, the interior of the embedded complex is homeomorphic to an open ball. This is not the only possibility: Figure 20 illustrates another embedding, where the map f can be defined:

$$f(e^{i\theta}) = \begin{cases} e^{2i\theta} & \text{if } \theta < \pi, \\ e^{-2i\theta} & \text{otherwise.} \end{cases}$$

For the illustrated embedding, the inside of the embedded complex is homeomorphic to the inside of a torus.

8 In-groups, out-groups, and embeddings.

In [27] Neuwirth associates with an embedding of an n-complex K in an oriented $(n+1)$-manifold M a group called the *in-group* and written $G(K, M)$.

Specifically, suppose that M is an oriented compact triangulated $(n+1)$-manifold, so its structure is defined by an $(n+1)$-complex. Let K be an n-subcomplex of its n-skeleton. Fix an orientation of M and apply arbitrary orientations to each n- and $(n-1)$-dimensional face

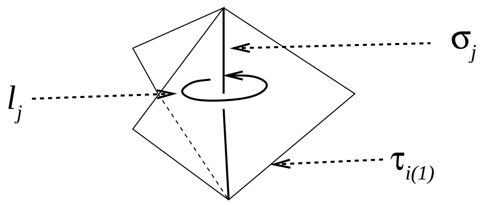

Figure 21: Oriented loop around an $(n-1)$-face.

in K. Let τ_i and σ_j respectively denote these oriented faces. Each face σ_j has codimension 2 in M, so one can form a small oriented loop ℓ_j going around it. Then ℓ_j intersects some of the faces τ_i in cyclic order $(\tau_{i(1)}, \ldots, \tau_{i(m)})$, say, and the orientations of $\tau_{i(r)}$ and ℓ_j define associated intersection numbers $\epsilon_{i(r)} = \pm 1$. In the group F freely generated by the n-faces $\tau_1, \ldots, \tau_p$ of M, let w_j be the word

$$\tau_{i(1)}^{\epsilon_{i(1)}} \cdots \tau_{i(m)}^{\epsilon_{i(m)}}.$$

The in-group $G(K, M)$ is defined as the group with generators τ_i and relators w_j.

Suppose that $M \backslash K$ has components $C_1 \ldots C_t$. Form an arc-connected topological space Y by joining these components by arcs to a centre point v (this is a formal construction: the arcs only join to Y at their endpoints). Form X by gluing K back onto Y. Let $G = \pi_1(X)$ be the fundamental group of X with basepoint v, and let H be the normal subgroup of G generated by all loops disjoint from K. The following is Theorem 4 of [27].

THEOREM 8.1 *$G(K, M)$ is isomorphic to the quotient group G/H.* ∎

Each in-group $G(K, M)$ is defined uniquely by the cyclic ordering of n-faces τ_i around $(n-1)$-faces σ_j. Each cyclic ordering determines a group, called an *out-group* of K. When the dimension n of K is 1, i.e., K is a 1-complex or graph, the out-groups provide detailed information about surface embeddings. Let ω be a fixed cyclic ordering of the incident edges around each vertex of K. There is an orientable surface $N_K(\omega)$ defined by ω, such that K is embedded in this surface and its complement consists of open discs [15, 45, 28]. Let $\pi_\omega(K)$ be the out-group of K defined by the ordering ω. The following is Theorem III.1 of [28].

PROPOSITION 8.2 *When $n = 1$, $\pi_\omega(K) \equiv F(r) * \pi_1(N_K(\omega))$, where $r+1$ is the number of connected components of $N_K(\omega)\backslash K$, and F(r) is the free group on r generators.* ∎

This gives us a criterion for planarity. If ω is an ordering compatible with a planar embedding of K, then $N_K(\omega) = S^2$, and has trivial fundamental group. Hence,

COROLLARY 8.3 *K is planar iff one of its out-groups is free.* ∎

Figure 22 (copied from [27]) illustrates the idea when K is a figure-8. The cyclic orderings around the vertices of degree 2 (there is only one) simply reduce the number of generators of the out-group: for example, the ordering around σ_1 makes the letters τ_1 and τ_2 equal in every out-group. Hence the number of generators can be reduced to two, τ_2 and τ_4, say, and the out-groups are defined by the 6 possible cyclic orderings of edges around the vertex σ_3.

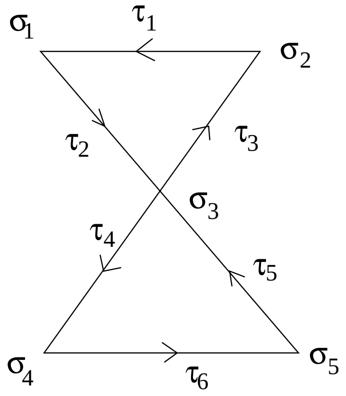

Figure 22: graph example from [27].

Let us interpret the intersection numbers as yielding +1 for an out-edge and -1 for an in-edge. Thus the cyclic ordering 2453 yields the relator $\tau_2^{-1}\tau_4\tau_4^{-1}\tau_2$ which collapses to the identity: thus this ordering produces the out-group $F(2) = \mathbb{Z} * \mathbb{Z}$. This is as expected, since it is the ordering illustrated, which defines a planar embedding whose complement has 3 components.

The ordering 2435 yields the relator $\tau_2^{-1}\tau_4\tau_2\tau_4^{-1}$, which can be interpreted as $\tau_2^{-1}\tau_4 = \tau_4\tau_2^{-1}$, so the two letters commute and the group is $\mathbb{Z} \times \mathbb{Z}$, which shows that the ordering is not compatible with a planar embedding. Similarly, the ordering 2543 yields $\mathbb{Z} \times \mathbb{Z}$— incompatible with an embedding— and the other three orderings yield $F(2)$, and they all define planar embeddings. This concludes the discussion of Figure 22.

When $n > 1$, we get a necessary condition for embeddability in an oriented $(n+1)$-manifold M:

PROPOSITION 8.4 ([27], Theorem 6. *If $n > 1$, a necessary condition for embeddability of K in M where the complement of K has $(r + 1)$ components is that at least one out-group of K should be a quotient of $\pi_1(M) * F(r)$.* ∎

9 Expansion of 2-complexes to 3-manifolds.

Neuwirth [29] gives an algorithm to determine whether a restricted kind of 2-cell-complex K can be expanded to an oriented 3-manifold M by attaching a closed 3-cell. In this case, K is called a 'spine' of M. It can be assumed that M is 'built up' by taking a single vertex or 0-cell, to which one attaches n 1-cells, forming a generalised figure-8, to which one attaches n 2-cells (discs). These together form the 2-skeleton of (this description of) M. Finally, M is obtained by attaching a closed 3-cell (ball) to K. (A lower-dimensional analogy is constructing a torus by attaching a rectangle to a figure-8. The idea is indicated in Figure 23). The complex K corresponds in a direct way to a group presentation: the author obtains therefore a semi-algorithm to determine whether a finitely presented group is the fundamental group of an oriented 3-manifold.

The ideas are roughly as follows. Within the imagined orientable 3-manifold M, take a thin tubular neighbourhood of the 1-skeleton, form the intersection of the boundary of this neighbourhood with a small sphere around the 0-cell, obtaining disjoint circles, and intersect

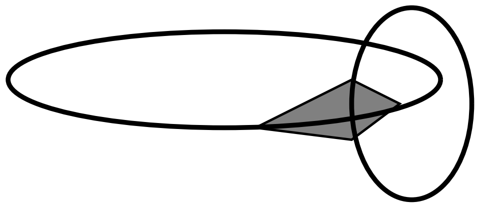

Figure 23: 1-skeleton supporting a torus.

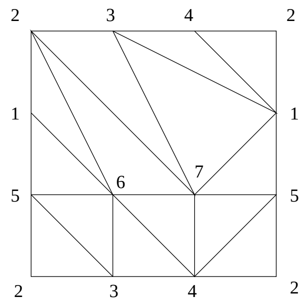

Figure 24: Möbius torus.

these circles with the 2-skeleton, obtaining $2n$ points. Let E be this set of points. The layout of K can be used to define three distinct permutations A, B, C on the set E. Then

THEOREM 9.1 *Suppose* (i) *that A and C together generate a transitive permutation group on E. Then* (ii) *the group generated by AC and BC is also transitive, and* (iii) $|AC| - |A| + |C| = 2$.[4] ∎

More significantly, the possible permutations A, B and C can be listed from the combinatorial structure of K, and if they satisfy (i)–(iii) then M can be constructed. Generalised versions of these results are described when (i) does not hold.

10 Computational synthetic geometry.

This field has been actively pursued in recent years. The book [9] summarises the principal tools employed, which are algebraic. A typical computational problem in this field (to which a positive answer was given in 1949) was to determine whether *Möbius's torus*, which is a triangulation of a torus with 7 vertices, could be realised by a polyhedron in 3 dimensions [7].

The method is explained in [6]. Broadly speaking, it is as follows: an arrangement of the 7 vertices in space so that the 14 triangles can be filled in (flat) without intersecting any other edge, corresponds to a certain assignment of signs ± 1 to the set of all ordered 4-tuples of points. (The assignment χ of signs to the 35 tuples is called a 'simplicial chirotope.') The chirotope χ satisfies certain restrictions: more specifically, for every triple a, b, c forming a triangle and edge u, v, there are exactly two patterns of signs excluded for the five ordered 4-tuples from the set a, b, c, u, v. For each such chirotope, its *realisation* is a mapping of vertices to points $(1, x, y, z)$

[4] $|A|$ is the number of cycles in the cycle decomposition of A.

in $\mathbb{R}^4$ such that the sign of the determinant of any 4 of the points matches the chirotope sign. This in turn defines a realisation of the Möbius torus as a polyhedron in $\mathbb{R}^3$.

For example, a study of the paper [7] yields the following coordinatisation of the Möbius torus:

$$(0,1,0),(8,0,-1),(8,0,1),(0,-1,0),(3.8,-0.33,0.33),(1,0,0),(3.8,0.33,-0.33)$$

11 References.

1. D. Albers and G. Alexanderson (1985). *Mathematical people — profiles and interviews.* Birkhäuser.

2. V. Akman, A. Arslan, and W. Franklin (1991). Implementing an electronic topological picturebook. *Proc. 13th IMACS World Congress, Dublin,* 1564–1565.

3. T. Banchoff (1976). Self linking numbers of space polygons. *Indiana University Mathematics Journal* **25:12**, 1171–1188.

4. T. Banchoff (1990). *Beyond the third dimension: geometry, computer graphics, and higher dimensions.* Scientific American Library, New York.

5. U. Betke and P. Gritzmann (1982). A combinatorial condition for the existence of polyhedral 2-manifolds. *Israel J. Maths* **42:4**, 297–299.

6. J. Bokowski and U. Brehm (1985). A new polyhedron of genus 3 with 10 vertices. *Colloquia Math. Soc. János Bolyai*, Siofok.

7. J. Bokowski and A. Eggert (1986). All realizations of Möbius' torus with 7 vertices. Preprint 1009, FB Mathematik, Technische Hochschule Darmstadt.

8. J. Bokowski, A. Guedes de Oliveira, and J. Richter (1988). Algebraic varieties characterizing matroids and oriented matroids. Technical report.

9. J. Bokowski and B. Sturmfels (1989). *Computational Synthetic Geometry.* Springer Lecture notes in Mathematics 1355.

10. J. Boyle (1987). Embeddings of 2-dimensional cell complexes in S^3 determined by their 1-skeleton. *Topology and its applications* **25**, 285–299.

11. A. Brøndsted (1983). *An Introduction to Convex Polytopes.* Springer Graduate Texts in Mathematics **90**.

12. H. Crapo and J. Ryan (1986). Spatial realizations of linear scenes. *Structural Topology* **13**, 33–68.

13. T. K. Dey (1993). An optimal algorithm for detecting null-homologous cycles on 2-manifolds. Presented at the 5th Canadian Conference on Computational Geometry, Waterloo, Canada, August 1993.

14. T. K. Dey and H. Schipper (1994). A new technique to compute the polygonal schema for 2-manifolds with application to null-homotopy detection. *Proc. 10th Annual Symposium on Computational Geometry.* (To appear, *Discrete and Computational Geometry.*)

15. J. Edmonds (1960). A combinatorial representation for polyhedral surfaces. *A.M.S. Notices* **7**, 646.

16. G. Francis (1987). *A topological picturebook.* Springer.

17. M. Greenberg (1967). *Lectures on algebraic topology.* Benjamin, New York.

18. C. Hayat-Legrand and F. Sergeraert (eds) (1990). *Algorithmique, topologie et géométrie algébriques.* (Astérisque **192**.)

19. N. Habegger (1984). Obstructions to embedding disks II (a proof of a conjecture of Hudson). *Topology and its applications*, 123–130.

20. N. Habegger (1984). Embedding up to homotopy type — the first obstruction. *Topology and its applications*, 131–143.

21. J. Hudson (1963). Knotted tori. *Topology* **2**, 11–22.

22. J. Hudson (1963). A non-embedding theorem. *Topology* **2**, 123-128.

23. J. Köbler, U. Schöning, and J. Torán (1993). *The graph isomorphism problem: its structural complexity.* Progress in theoretical computer science, Birkhäuser.

24. S. Lefschetz (1949). *Introduction to Topology.* Princeton University Press.

25. G. Markowsky and M. Wesley (1980). Fleshing out wire frames. *IBM J. Research and Development* **24:5**, 582–597.

26. K. Mehlhorn (1984). *Data structures and algorithms* **2**: *Graph Algorithms and NP-completeness.* Springer.

27. L. Neuwirth (1965). In-groups, coverings, and embeddings. *Proc. Cambridge Phil. Soc* **61**, 647–656.

28. L. Neuwirth (1966). Imbedding in low dimensions. *Illinois J. Math.* **10:3**, 470–478.

29. L. Neuwirth (1968). An algorithm for the construction of 3-manifolds from 2-complexes. *Proc. Camb. Phil. Soc* **64**, 603–613.

30. M. Paterson and A. Razborov (1988). The set of minimal braids is co-NP-complete. Technical report, Computer Science Department, University of Warwick.

31. J. Ruppert and R. Seidel (1992). On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete and Computational Geometry* **7**, 227–253.

32. A. Shapiro (1957). Obstructions to the imbedding of a complex in a Euclidean space. I: the first obstruction. *Annals of Mathematics* **66:2**, 256–269.

33. M. Tangora (1989). *Computers in Geometry and Topology,* (edited Tangora), Dekker Lecture notes in Pure and App. Maths.

34. K. Taniyama (1994). On embeddings of a graph into $\mathbb{R}^3$. In *Contemporary Mathematics* **164:** *geometric topology*, 239–246, ed. C. Gordon, Y. Moriah, and B. Wajnryb. AMS, Providence, RI.

35. K. Taniyama (1994). Link homotopy invariants of graphs in $\mathbb{R}^3$. *Revista matematica de la Univ. Complutense de Madrid* **7:1**, 129–144.

36. K. Taniyama (1994). Cobordism, homotopy, and homology of graphs in $\mathbb{R}^3$. *Topology* **33:3**, 509–523.

37. K. Taniyama (1994). Homology classfication of spatial embeddings of a graph. Preprint, Dept. of Maths, College of Arts and Sciences, Tokyo Women's Christian University. To appear, *Topology and its applications.*

38. G. Vegter and C-K Yap (1990). Computational complexity of combinatorial surfaces. *Prof. 6th ACM Symposium on Computational Geometry*, 102–111.

39. C. Watt (1991). *Embedding simplicial complexes in Euclidean spaces.* M.Sc. Dissertation, Mathematics Department, Trinity College, Dublin.

40. M. Wesley and G. Markowsky (1981). Fleshing out projections. *IBM J. Research and Development* **25:6**, 934–954.

41. E. Whittlesey (1960). Finite surfaces: a study of finite 2-complexes, I: local structure, II: the canonical form. *Mathematics Magazine* **34,1:** 11–22, **2:** 67–80.

42. Wu Wen-Jun (1960). On the isotopy of complexes in a Euclidean space I. *Science Sinica* **9**, 21–46.

43. Wu Wen-Jun (1965). *A theory of imbedding, immersion, and isotopy of polytopes in a euclidean space.* Science Press, Peking, 1965.

44. Wu Wen-Jun (1985/86). On the planar imbedding of linear graphs. *J. Sys. Sci. and Math. Scis.* **5:4,** 290-302, and **6:1** 23–35.

45. J. Youngs (1963). Minimal imbeddings and the genus of a graph. *J. Math. Mech.* **12**.