Homeomorphism of 2-complexes is equivalent to graph isomorphism.

Colm Ó Dúnlaing
Colum Watt
David Wilkins
Mathematics, Trinity College, Dublin 2, Ireland. *

February 22, 1999

Abstract

Whittlesey [12,13,14] gave a criterion which decides when two finite 2-dimensional complexes are homeomorphic. We show that graph isomorphism can be reduced efficiently to 2-complex homeomorphism, and that Whittlesey's criterion can be reduced efficiently to graph isomorphism. Therefore graph isomorphism and 2-complex homeomorphism are polynomial-time equivalent.

1 Introduction

This paper is concerned with a problem in Computational Topology: for recent surveys of the field, see [3, 10]

A *combinatorial surface* is a topological space representable as a finite union of triangular regions, joined along edges. Examples include a sphere, a torus, a Klein bottle, and so on. More generally one can have a surface with boundary — points on the boundary resemble points on the boundary of a closed disc.

It is possible to construct many examples of such surfaces by taking a convex polygonal region and 'gluing' its edges together in various ways. Examples are given in Figure 1.

The process is summarised by labelling the sides of the polygon with letters like a or inverted letters like a^{-1}. Letters (inverted or otherwise) occur at most twice. Each letter corresponds to an oriented loop or curve-segment to which the side is 'glued' by an order-preserving or reversing map depending on whether the letter is inverted. As a result, when a letter and its inverse occur, the sides are joined without a twist, as in a cylinder $abcb^{-1}$; otherwise a twist is involved, as in a Möbius band $abcb$.

When every side is glued to another side we have a compact surface without boundary. By reducing to a canonical form, it turns out that every such surface is homeomorphic either to a

*This work was supported in part by the EU under ESPRIT BRAs 3075 (ALCOM) and 7141 (ALCOM II).
E-mail: odunlain@maths.tcd.ie, colum@maths.tcd.ie, dwilkins@maths.tcd.ie

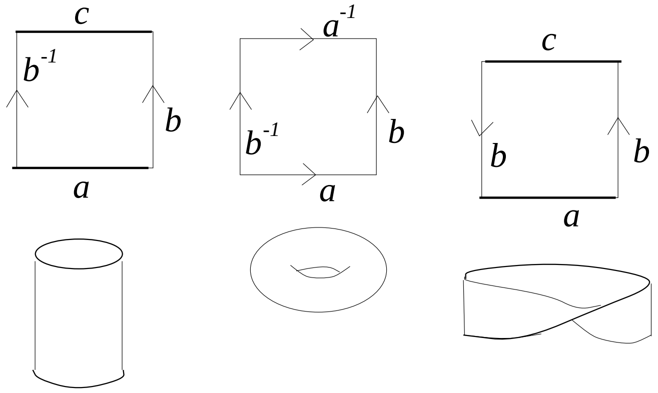

Figure 1: a bounded cylinder, a torus, and a Möbius band.

sphere, or to m tori joined together, or to n projective planes joined together. It can be shown that no two surfaces in this list are homeomorphic.

The canonical form is described in [6], due to Brahana with simplifications by Tucker. The paper [11] shows how to reduce efficiently to canonical form. Whittlesey's work [12,13,14] generalises the classification of compact surfaces.

This paper is based mainly on Whittlesey's work. It is concerned with classifying 2-complexes up to homeomorphism. The surfaces (with boundary) illustrated in Figure 1, constructed by 'gluing together' edges of a square, are simple examples of 2-complexes.

This paper is ultimately concerned with the computational complexity of homeomorphism testing for 2-complexes. We prove that it is polynomial-time equivalent to graph isomorphism testing, which gives as accurate an estimate of its complexity as present methods allow: graph isomorphism is a well-studied problem, obviously in the class NP, but currently not known to be NP-complete. See [5].

This paper presents quite a self-contained description of Whittlesey's method which, together with its relation to graph isomorphism, should provide a starting-point for implementation studies. The present description is an improved and extended version of results first presented in [8].

Apart from finding normal forms for surfaces, there has been related work about the homotopy and homology groups for surfaces [2,4,11]. It should be noted that homotopy group calculations such as in [4] cannot be generalised to 2-complexes. This is because, given any finite presentation of a group G (see [7]), one can construct a 2-complex whose fundamental group is G. If G has unsolvable word-problem [7] then it is undecidable if a loop in that complex is null-homotopic.

The classification of 2-complexes had been suggested as a topic by Whittlesey's Ph. D. supervisor Alan Tucker. Actually, Tucker doubted that homeomorphism testing for 2-complexes could be solved, because related questions about fundamental groups were unsolvable. Fortunately he forgot to mention his doubts to Whittlesey. An account is given in [1].

2 Cell complexes

(2.1) Graphs. A graph G is conventionally defined as a pair (V, E) where V is a finite set of 'vertices' and E a set of 'edges,' where each edge is an unordered pair $\{u, v\}$ of distinct vertices in V. The edge $e = \{u, v\}$ is *incident* to u and v, and one also says that u and v are *incident* to e.

A directed graph is similar, except that the edges are ordered pairs (u, v) of distinct vertices. The *inverse* e^{-1} of a directed edge e is (v, u), where $e = (u, v)$.

Paths. A *path* from u to v in a (directed) graph G is a sequence

(2.2) $$v_0, e_1, v_1, e_2, \dots e_k, v_k$$

where $v_0 = u$, $v_k = v$, and for $1 \leq j \leq k$, e_j is the (directed) edge from v_{j-1} to v_j. Of course, each v_i must be a vertex and each e_j an edge of G. The *length* of the path is k. (Note: the length of the path is the number of occurrences of edges. The same edge may be 'traversed' several times.)

Obviously the sequence in Formula 2.2 contains redundant data. The definition is unusual, and intended to apply to generalised kinds of graph, such as when there are several edges joining the same two vertices (2.5). Otherwise the edges could be omitted from the description.

(2.3) Simplicial 2-complexes. The notion of 'graph' can be extended to higher dimensions. A *simplicial 2-complex* is a triple (V, E, F) where (V, E) is a graph and F consists of 'triangular' *faces*. Each face is an unordered triple $\{u, v, w\}$ of distinct vertices, and it is *incident* to the three edges $\{u, v\}$, $\{v, w\}$, and $\{u, w\}$, which must all belong to E. Likewise we can have an oriented simplicial 2-complex, with directed edges and oriented faces. An oriented face is a *cyclically* ordered triple (u, v, w). Its incident edges are (u, v), (v, w), and (w, u), and they must belong to E. The same face can be written as (v, w, u) or (w, u, v). The '1-skeleton' of an (oriented) simplicial complex (V, E, F) is the (directed) graph (V, E).

(2.4) Giving graphs and simplicial complexes a topological structure. Topology can be introduced by identifying every edge with a line-segment and every face with a triangular plane region. Faces incident to the same edge are considered glued together along their common edge, and edges incident to the same vertex are considered glued together at that vertex. Parts are glued together by declaring them equivalent and forming a quotient topology. To define such a topology it is not necessary to embed the complex in Euclidean space (see [9]).

(2.5) 1-complexes. Whittlesey's algorithm applies to more general complexes. His notion of a graph is more directly topological: the edges are viewed as curve-segments 'attached' to the vertices. The conventional notion of a graph is insufficient to describe these graphs, since several edges can be incident to the same two vertices, and edges can have both ends incident to one vertex. So, 'graph' means '1-complex,' which means, roughly, a multigraph (allows multiple edges) with self-loops: edges incident to just one vertex. 'Oriented graph' means 'directed multigraph with self-loops.' A graph in the conventional sense Whittlesey calls a polygonal graph.

Whittlesey represents vertices by capital letters $A, B, \dots$ and edges by small letters $a, b, \dots$. We say 'e is an edge from A to B,' not '$e = (A, B)$,' since there could be several such edges.

The idea of path is generalised beyond Formula 2.2 in that an edge can be traversed in the opposite direction. Thus a path from A to H is a sequence

$$AsBt\dots zH$$

where s is an edge from A to B, or the inverse of an edge from B to A, and so on. If $A = H$ the path is *closed* (i.e., a cycle). There are trivial closed paths, containing just a single vertex.

(2.6) 2-complexes. Whittlesey's definition of 2-complex is likewise more general than a simplicial complex. A complex is 'built up' in three stages. First, a finite set of vertices ('0-cells') is chosen.

Second, a finite set of directed edges ('1-cells') whose ends are attached to these vertices. This defines the *1-skeleton* of the complex, which is a 1-complex or oriented graph in its generalised sense.

The 1-skeleton of a complex K is denoted $K^{(1)}$. The 1-skeleton is not just a graph; it is a 1-complex, with a topological structure.

Third, a finite set of discs ('2-cells') whose boundaries are attached to the 1-skeleton by continuous maps. For definiteness, the discs are all copies of the unit disc $\{(x,y) \in \mathbb{R}^2 : x^2+y^2 \leq 1\}$. Its boundary is the unit circle S^1.

Each disc is attached to $K^{(1)}$ along a closed path in $K^{(1)}$. Given such a closed path, an 'attaching map' f from the disc boundary S^1 is defined as follows: suppose the closed path has length k. If $k = 0$ the path is trivial and the attaching map f takes the entire boundary onto the path's single vertex.

Otherwise, separate S^1 into k circular segments. Taking the k segment endpoints in anticlockwise order around S^1, f maps them to the path vertices in the order given. The map f maps each segment monotonically onto the corresponding edge of the path, in the direction given. (If $k \geq 3$ it may be preferable to identify the disc with a regular k-gon and its interior.)

Obviously there is freedom of choice in this construction, but different choices yield homeomorphic complexes. So a 2-complex K can be described by listing its vertex labels A, 1-cells in the form AeB, say, (so e joins A to B), and its 2-cells by the attaching paths. This is a finite description. The attaching maps are implicit: they are unique up to homeomorphism.

Given a 2-complex K, we can assume that attaching maps, as just described, have been fixed. The topological space defined by these attaching maps will also be denoted K. This is 'an abuse of language.'

Whittlesey calls these structures *cell complexes*. They are more general than simplicial complexes and less general than CW complexes (in which the attaching maps can be 'pathological') [9].

(2.7) Barycentric subdivision. Whittlesey's algorithm is most conveniently described in terms of 2-complexes rather than the more basic simplicial complexes. However, it is easy to convert a 2-complex into a homeomorphic simplicial complex by *barycentric subdivision.*[1]

Subdividing a (directed) 1-cell. Let e be an edge from A to B. Introduce a new vertex C (C is meant to be a point along the edge e) and two new edges, c from A to C and d from C to B. The edge e is discarded. Boundary paths must also be modified by replacing occurrences of AeB by $AcCdB$ and $Be^{-1}A$ by $Bd^{-1}Cc^{-1}A$. It is easy to show that the new complex is homeomorphic to the original.

[1] 'Barycentric' indicates a centre of gravity, and refers to the new vertices introduced.

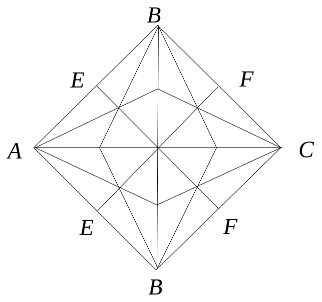

Figure 2: 3 subdivisions yield simplicial complex homeomorphic to S^2.

Subdividing a 2-cell. Suppose that in a complex K, a disc D is attached along a closed path of length k by an attaching map f. Let P be the closed path to which f attaches D. Assume first that $k \geq 2$.

Modify K by first adding a new vertex C corresponding to the centre of D. For every occurrence of a vertex Q along the path (except the last, which duplicates the first), add a new edge from C to Q. Call the edge e_{CQ}, though this is ambiguous since Q could occur more than once. For every subsequence $J\ell L$ of the path P, add a new 2-cell with attaching path

$$Ce_{CJ}J\ell Le_{CL}^{-1}C.$$

Finally delete the 2-cell attached along P.

Each of the new 2-cells corresponds to a piece D_i of the disc D. Their union is homeomorphic to the 2-cell they replace.

If $k = 1$ the construction as described produces a homeomorphic complex as follows: there is just one new cell with attaching path $Ce_{CJ}Je_{JJ}Je_{CJ}^{-1}C$ replacing a 'closed loop attachment' $Je_{JJ}J$. Both cells are homeomorphic to a closed disc.

If $k = 0$ the construction is much the same. The entire disc boundary is mapped onto a single vertex A. The new attaching path is $Ce_{CA}Ae_{CA}^{-1}C$. The 2-cell (before subdivision) is homeomorphic to the 2-sphere S^2, and the new 2-cell is homeomorphic to S^2 for different reasons.

(2.8) (Barycentric) subdivision of a 2-complex is done in two steps. First, subdivide all the 1-cells. Then subdivide all the 2-cells.

Consider the case of a complex K with one vertex, no edges, and one 2-cell attached to the vertex. This is homeomorphic to S^2. A barycentric subdivision produces a complex with two vertices, one edge, and one 2-cell attached by a path of length 2. Repeat, getting a complex with 4 vertices, 6 edges, and 4 2-cells attached by paths of length 3. This resembles a simplicial complex, but is not, because different 1-cells have the same endpoints, and different 2-cells have the same bounding vertices. One further subdivision yields a complex with 14 vertices, 36 edges, and 24 faces. This is a simplicial complex (see Figure 2). In general,

(2.9) Proposition *If K is a 2-complex then at most three barycentric subdivisions yields a simplicial 2-complex.* ∎

Whittlesey uses barycentric subdivision for other purposes not discussed here.

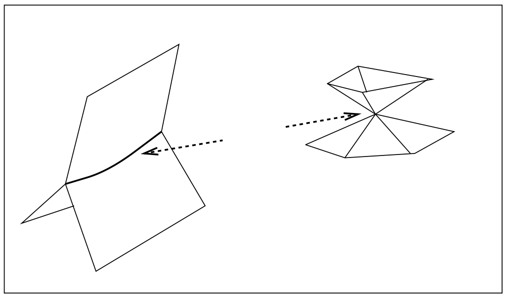

Figure 3: Line-singular points and conical point, with neighbourhoods.

3 Local structure of a 2-complex.

This section presents in outline Whittlesey's analysis of the topology of a 2-complex.

(3.1) Regular and singular points. Let K be a 2-complex. A point p is *regular* if K is surface-like near p, i.e., it has a neighbourhood 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*.

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

(3.2) Topological invariants for a 2-complex. Whittlesey [13] describes in detail the topology of the complex K in terms of the singular graph and the way in which the 2-cells are attached.

One can view the complex as a singular graph with spaces M_j attached where each M_j is a surface with boundary rather than just a disc. The boundary of each M_j is a union of cycles. Each cycle is given an attaching map defined by a closed path in the singular graph.

(3.3) Admissible changes to a presentation. In [12] Whittlesey states that the topology of a complex K can be presented by listing the singular graph together with these spaces M_j and the attaching maps. Certain changes can be made. For example, if M_j is orientable then it can be given the opposite orientation. If M_j is not orientable, then any of the attaching cycles can be inverted. Except for such changes, the topology of K is characterised fully by this presentation.

An outline proof of this result is given in [12].

4 Reducing graph isomorphism to 2-complex homeomorphism

(4.1) Graph isomorphism. Let $G = (V, E)$ and $G' = (V', E')$ be (directed) graphs. An isomorphism between G and G' is a bijective map $g: V \to V'$ such that for any $u, v \in V$, the (directed) edge from u to v is in E iff the (directed) edge from $g(u)$ to $g(v)$ is in E'.

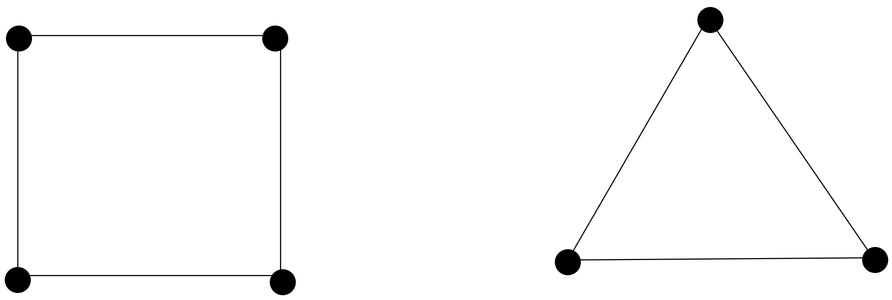

Figure 4: graphs homeomorphic but not isomorphic.

The main result of this paper is that graph isomorphism is polynomial-time equivalent to homeomorphism of 2-complexes. One direction is easy to prove — Lemma 4.2 below. In proving the Lemma, it is necessary to construct 2-complexes whose topology fully determines the vertices and edges of G and G' respectively. Comparing G and G' topologically is not enough: a 3-cycle is homeomorphic to a 4-cycle, but not isomorphic (Figure 4).

(4.2) Lemma *Graph isomorphism is efficiently reducible to homeomorphism of 2-complexes.*

Proof. Let G be a graph with n vertices and m edges. We construct a 2-complex (actually, a simplicial 2-complex) K resembling G as follows. It has $n + 2m$ 0-cells (vertices) of which the first n correspond directly to the vertices of G.

For each edge $e = \{u, v\}$ of G there is a 2-dimensional subcomplex $T(e)$, the boundary of a tetrahedron, on which u and v are incident 0-cells. $T(e)$ is, of course, homeomorphic to the 2-sphere S^2. The other two vertices are 'internal' to the structure for $\{u, v\}$. See Figure 5. If G has n vertices and m edges then K has $n + 2m$ 0-cells, $6m$ 1-cells, and $4m$ 2-cells.

An edge e of G is incident to a vertex u if and only if the subcomplex $T(e)$ contains the 0-cell u.

The singular graph (3.1) S of K consists of isolated conical points, corresponding to those vertices of G which are incident to more than one edge. Let $X = K \backslash S$ (K being viewed as a topological space). The components C of X correspond to the edges of G, i.e., each 'edge tetrahedron' $T(e)$ is the closure $\bar{C}$ in K of a unique component C of X.

Given another graph G', write K' for the corresponding complex, S' for its singular graph, and X' for $K' \backslash S'$.

Given an isomorphism g from G to G' it is easy to define a homeomorphism (indeed, a 2-complex isomorphism) f from K to K'. Every vertex u of G is a 0-cell of K: let $f(u) = g(u)$. For every edge $e = \{u, v\}$ of G, let $e' = \{g(u), g(v)\}$. Then extend f so it carries the tetrahedron $T(e)$ onto $T(e')$. We omit the few remaining details.

Conversely, suppose $f \colon K \to K'$ is a homeomorphism. It takes S bijectively onto S' and X onto X'.

The components of X are spheres with up to two points deleted. We construct a map g from the vertices of G to those of G' as follows. On S, g is defined to coincide with f. It is extended to the vertices not in S as follows.

If a component C of X is a sphere minus one point, then it contains a unique vertex v (which is incident to exactly one edge in G). Now, $f(v)$ need not be a vertex of G'. However, the image $f(C)$ contains a unique vertex v' of G': let $g(v) = v'$.

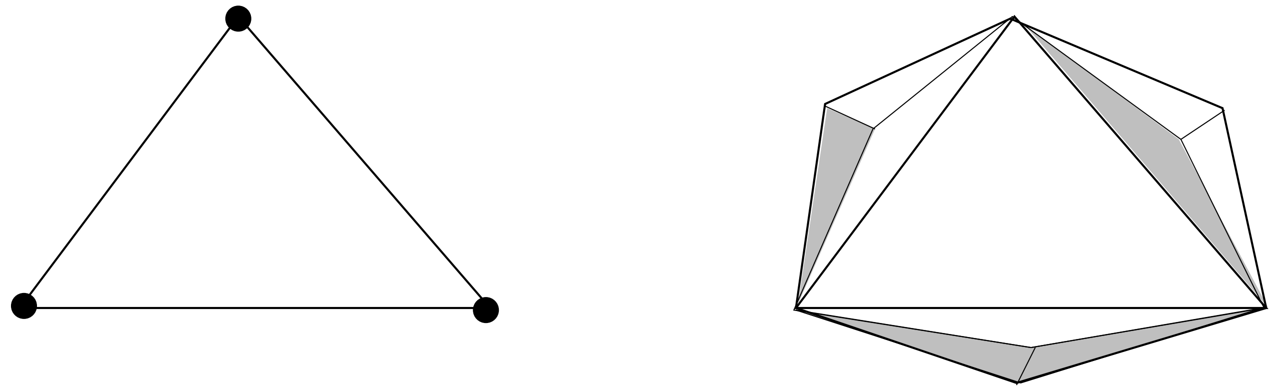

Figure 5: Graph G and corresponding complex K.

If a component C of X is a sphere with no point removed, it must coincide with an 'edge tetrahedron' $T(e)$, and its image $f(C)$ is another edge tetrahedron $T(e')$ in K'. Denote by u and v the vertices incident to e (e is the only edge incident to these vertices). Let $g(u)$ be one of the vertices incident to e' and $g(v)$ be the other.

This finishes the construction, and the map g is bijective. Every component C of X is contained in a unique edge-tetrahedron $T(e)$, and its image is contained in a unique $T(e') = f(T(e))$. By construction, if u is incident to e in G then $g(u)$ is incident to e' in G' and vice-versa. It follows that g is an isomorphism. ∎

5 Word systems and normal forms

We return to Whittlesey's treatment of 2-complex homeomorphism. Having produced a topological description which can be derived from any 2-complex (3.3), Whittlesey shows how to reduce it to a normal form. This generalises Brahana and Tuckers' normal-form classification of combinatorial surfaces [6].

(5.1) Word systems. 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 (generalised) directed graph 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 $AaBb\ldots A$ of letters or inverted letters labelling vertices and directed edges. Now σ_i describes the attachment of a topological disc D_i, and the topology of $|K|$ depends essentially on these *words* $AaBb\ldots A$. This Whittlesey calls a *word-system* W.

Words (and subwords) will be denoted by lowercase Greek letters $\alpha, \beta, \ldots$.

(5.2) Simplifying assumption: no isolated edges or vertices. To avoid inessential details, Whittlesey assumes 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.

Regular edges. An edge (label) a is *regular* if a occurs twice in the word-system, or a^{-1} occurs twice, or a and a^{-1} each occur once.

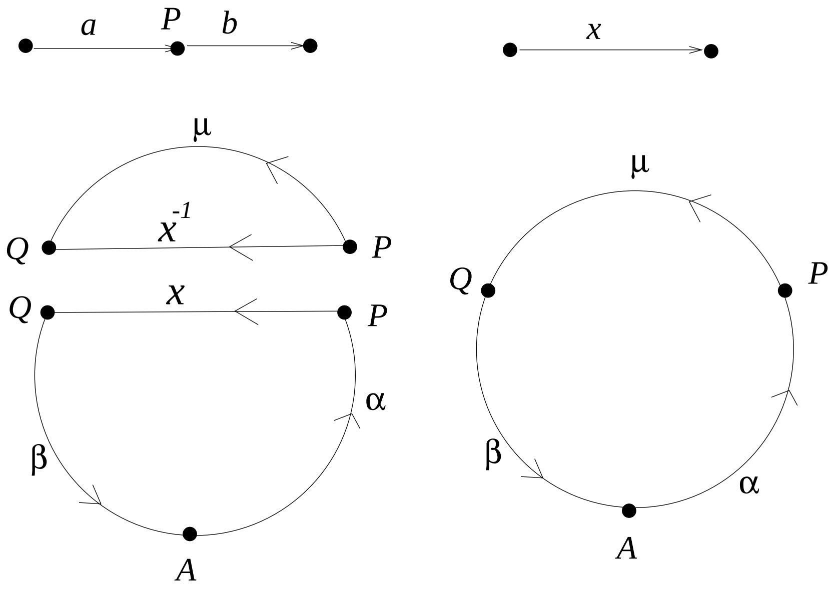

Figure 6: illustrating edge and polygon composition and subdivision.

(5.3) Admissible changes to a word-system. The following changes to a word-system W are shown to be admissible, since they produce homeomorphic complexes:

Edge composition: if a, b are labels (perhaps inverted) for edges with common vertex P and a, P, and b are only mentioned together in W, as $AaPbB$ or $Bb^{-1}Pa^{-1}A$, then $AaPbB$ can be replaced everywhere by a new edge-letter AxB and $Bb^{-1}Pa^{-1}A$ by $Bx^{-1}A$.

Edge subdivision, inverse to edge composition.

Word subdivision (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.

Word composition, inverse to this, providing the edge label x is regular.

It must be emphasised that polygon composition— combining two different words into one— is allowed only on *regular* edges. See Figure 6.

(5.4) Relabelling and cyclic permutation. Obviously a word system can be changed by changing a vertex label A to P throughout, if P does not already occur; or by changing an edge label a to b (and a^{-1} to b^{-1}) throughout, granted that neither b nor b^{-1} are already present. Also an edge label e may be swapped with its inverse e^{-1} throughout. A word may also be *inverted* and *permuted cyclically.* So

$$AaBb\ldots cA$$

may be replaced by its inverse

$$Ac^{-1}\ldots b^{-1}Ba^{-1}A$$

and a word $A\alpha B\beta A$ may be cyclically permuted, yielding

$$B\beta A\alpha B.$$

(5.5) Combinatorial equivalence. Two word-systems are *combinatorially equivalent* iff one can be converted into the other by a sequence of compositions and/or subdivisions, together

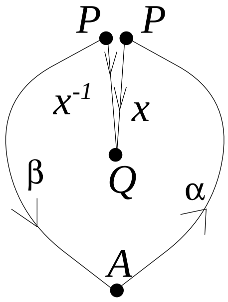

Figure 7: incision.

with the following 'formal equivalences' on words: 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 $A\ldots b^{-1}Ba^{-1}A$.

Whittlesey introduces further 'derived rules' which can be applied to individual words. These operations can be reduced to sequences of divisions and compositions, and relabelling. Again it is necessary that certain edges be regular.

Incision. Replace a word $A\alpha P\beta A$ by $A\alpha PxQx^{-1}P\beta A$ where x and Q are new labels. The inverse operation is allowed when x is regular and Q not mentioned elsewhere: $A\alpha PxQx^{-1}P\beta A$ can be replaced by $A\alpha P\beta$.

The rule can be justified by an appeal to Figure 7.

(5.6) Omitting vertex labels. Vertex labels are required for a full description of a complex, because of the existence of pinch-points. However, we shall often omit them, to make the formulae less cluttered.

The following 'circulation rules' can be justified in terms of divisions, compositions and relabelling. As a more direct explanation, consider Figure 8. The first pair of diagrams illustrates a transformation $\alpha\beta x\gamma\delta x \Rightarrow \alpha y\delta\beta^{-1}y\gamma^{-1}$. This amounts to cutting a Möbius band at two different places (x and y). In the case where β is empty we get rule 1a, and where γ is empty we get circulation rule 1b. (The diagram doesn't really allow for them to be empty, and y is relabelled as x, but the modifications needed are simple.)

Circulation Rule 1a. When x is a regular edge-label (or inverted label), replace $\alpha x\gamma\delta x$ by $\alpha x\delta x\gamma^{-1}$, or vice-versa.

Circulation Rule 1b. When x is regular, replace $\alpha\beta x\delta x$ by $\alpha x\delta\beta^{-1}x$, or vice-versa.

The second transformation illustrated in Figure 8 amounts to cutting a cylinder in different places. The second pair of circulation rules could be considered as special cases of this transformation.

Circulation rule 2a. When x is regular, replace $\alpha x\gamma\delta x^{-1}$ by $\alpha x\delta\gamma x^{-1}$, or vice-versa.

Circulation rule 2b. When x is regular, replace $\alpha\beta x\delta x^{-1}$ by $\alpha x\delta x^{-1}\beta$, or vice-versa.

(5.7) In the circulation rules, the edge x or its inverse occurs at the end of the word. This is mainly to simplify Figure 8 and can be achieved by cyclic permutation. However, the rules can easily be generalised: since the rules preserve a prefix α, they would preserve a prefix $\epsilon\alpha$, and hence, through cyclic permutation, they would preserve a suffix ϵ. Hence the rules can be considered to apply to subwords within the context $\alpha\ldots\epsilon$.

(5.8) Whittlesey shows how to apply these rules in a systematic way reducing a word-system

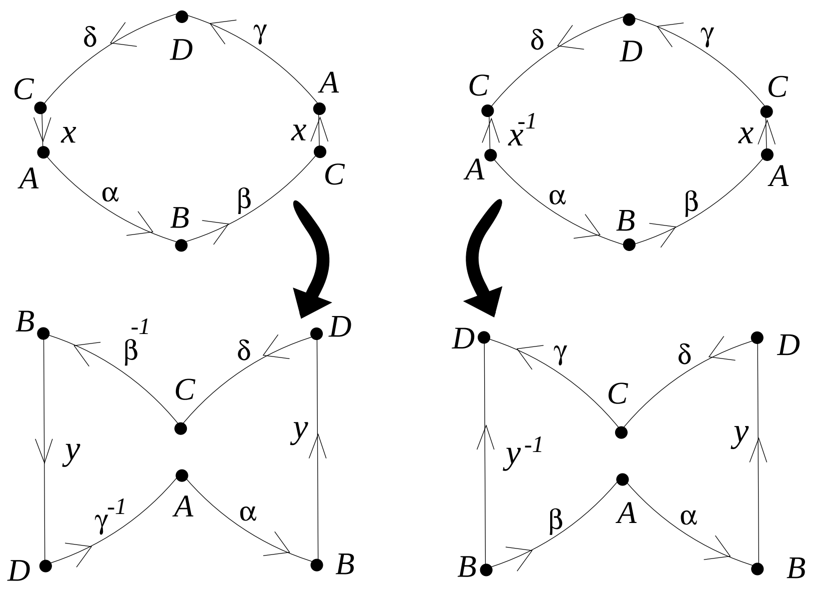

Figure 8: illustrating circulation rules.

to a normal form, such that two 2-complexes are homeomorphic if and only their normal forms are 'trivially isomorphic' (5.20).

First, one combines together polygons (words) 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.

(5.9) Collecting crosscaps. The next step is to collect crosscaps. A crosscap is where a regular edge or inverted edge x occurs twice in a word:

$$\alpha\beta x\gamma x\delta$$

It is possible to bring both occurrences of x directly after α:

$$\alpha\beta x\gamma x\delta \Rightarrow \alpha x\gamma\beta^{-1}x\delta \Rightarrow \alpha xx\beta\gamma^{-1}\delta,$$

using circulation rule 1b and then 1a. If another crosscap y exists to the right of xx, the same process can be repeated, yielding $\alpha xxyy\ldots$ (Moving a crosscap may eliminate others by inverting edge-labels.) In this way, all the crosscaps can be collected at the beginning of the word.

(5.10) Collecting handles. A handle is an occurrence of the subsequence

$$\ldots x\ldots y\ldots x^{-1}\ldots y^{-1}\ldots$$

where x and y are regular edges (or their inverses). The subsequence may be brought together as follows.

$$\alpha x\beta y\gamma x^{-1}\delta y^{-1}\epsilon \Rightarrow \alpha xy\gamma\beta x^{-1}\delta y^{-1}\epsilon$$

by applying rule 2a to the subword $x\ldots x^{-1}$. Then apply the same rule to $y\ldots y^{-1}$ and to $x\ldots x^{-1}$ again:

$$\Rightarrow \alpha xy\delta\gamma\beta x^{-1}y^{-1}\epsilon \Rightarrow \alpha x\delta\gamma\beta yx^{-1}y^{-1}\epsilon$$

Finally apply rule 2b to move $\delta\gamma\beta$ to the right:

$$\Rightarrow \alpha x y x^{-1} y^{-1} \delta\gamma\beta\epsilon.$$

This operation preserves the prefix α, so it can be repeated to collect all handles at the front of the word. The process does not (unlike when collecting crosscaps) invert edges.

(5.11) Moving crosscaps. A crosscap of the form xx can be moved anywhere without otherwise changing the word, by applying two circulation rules: for example,

$$\theta x x \Rightarrow x\theta^{-1} x \Rightarrow x x \theta.$$

(5.12) Cancelling handles with crosscaps. Suppose that a handle exists to the right of a crosscap:

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

First bring the handle together as above, and then move the crosscap up to it, yielding

$$\alpha x x a b a^{-1} b^{-1} \beta.$$

Using circulation rules of type 1 we get (the context $\alpha \ldots \beta$ is omitted)

$$xxaba^{-1}b^{-1} \Rightarrow xb^{-1}a^{-1}xa^{-1}b^{-1} \Rightarrow xb^{-1}b^{-1}ax^{-1}a \Rightarrow xb^{-1}b^{-1}xaa \Rightarrow xxbbaa$$

and hence the original word can be converted to the form $\alpha xxbbaa\beta$. Repeat this as often as necessary to replace every handle by a pair of crosscaps.

(5.13) Classification of surfaces. In this way we can reduce any word to the form $\alpha\beta$ where α is a product of crosscaps or a product of handles and β contains neither crosscaps nor handles. When the complex represents a surface, all edges are regular, so every occurrence of x in β is matched by an occurrence of x^{-1} (since β contains no crosscaps). Also there are no pinch-points (nor any other kind of singular point).

If β contains a subsequence $\ldots a\gamma b\delta a^{-1} \ldots$, then γ or δ must contain b^{-1} since otherwise β would contain a handle. In other words, in β edges occur in inverse pairs, and occurrences are nested. We can choose an innermost pair, which must be adjacent: aa^{-1}. This pair can be cancelled (reverse incision). Repeating as often as necessary, β becomes empty, and we conclude

(5.14) Proposition *Every (connected finite) 2-complex representing a surface is homeomorphic to a sphere or a product of crosscaps or of handles.* ∎

Vegter and Yap [11] have described an efficient implementation of the above reduction procedure: a surface (represented by a single word) can be reduced to a product of crosscaps or handles in time $O(n\log(n))$ where n is the number of edges.

Here we consider a reduction procedure for general 2-complexes, not necessarily representing surfaces.

(5.15) Moving handles. We have already seen that a crosscap xx can be moved past any subword. The same applies to handles using circulation rules of type 2, as follows.

$$aba^{-1}b^{-1}\theta \Rightarrow b^{-1}\theta aba^{-1} \Rightarrow b^{-1}a\theta ba^{-1} \Rightarrow a^{-1}b^{-1}a\theta b \Rightarrow \theta a^{-1}b^{-1}ab.$$

Now there is a handle after θ. It can be converted to the original form using the transformations

$$a^{-1}b^{-1}ab \Rightarrow b^{-1}aba^{-1} \Rightarrow aba^{-1}b^{-1}.$$

(5.16) By means of the transformations discussed above, any word can be converted into the form $\alpha\beta$ where α is either a product of crosscaps or a product of handles, and β contains no crosscaps nor handles. It is convenient to convert this to $\alpha y\beta y^{-1}$, as follows.

Incision yields $\alpha y^{-1}y\beta$ which can be transformed to $y^{-1}\alpha y\beta$, because α is a (possibly empty) product of crosscaps or handles past which y^{-1} can be moved. Cyclic permutation yields the desired form

$$\alpha y\beta y^{-1}.$$

Separating cuffs. 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. Therefore the cuffs are 'properly nested' in the sense of 5.13: i.e., $s\ldots t\ldots t^{-1}\ldots s^{-1}$ or $s\ldots s^{-1}\ldots t\ldots t^{-1}$ may occur, but not $s\ldots t\ldots s^{-1}\ldots t^{-1}$.

If there exists a cuff, choose an innermost cuff, and bring it in front of y:

$$y\gamma t\delta t^{-1}\epsilon y^{-1} \Rightarrow yt\delta t^{-1}\epsilon\gamma y^{-1} \Rightarrow t\delta t^{-1}y\epsilon\gamma y^{-1}$$

using rules 2a and 2b. Repeating this process on $\epsilon\gamma$, we eventually convert $y\beta y^{-1}$ to a product of separated cuffs, without nesting. Each cuff is of the form $t\delta t^{-1}$. If δ is empty this *could* be a removable incision. One must bear in mind that vertex labels have been omitted for simplicity: 'δ is empty' means that the cuff has the form tPt^{-1}. Only if the vertex P is not mentioned elsewhere can the incision be removed. Having removed all removable incisions, one then can compose the edges within each cuff, wherever possible. By circulation rule 2a the word δ can be permuted cyclically within the cuff $t\delta t^{-1}$, which can allow further edge compositions.

(5.17) Definition *A* normal (or canonical) form *for a system of words is obtained by composing words as much as possible by regular edges, then reducing each word to the form $\alpha\beta$ where α is a product of crosscaps or handles and β a product of cuffs (not containing other cuffs), and composing edges within each cuff as much as possible. The resulting word-system is called* canonical.

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

(5.18) Inverting cuffs. If a canonical word has crosscaps and cuffs, any cuff can be inverted. Let zz be the rightmost crosscap. It can be moved next to the cuff. Then apply rules of type 2b and 1a.

$$zzt\delta t^{-1} \Rightarrow zt\delta t^{-1}z \Rightarrow zzt\delta^{-1}t^{-1}$$

The cuff is thus inverted.

(5.19) Simultaneously inverting the cuffs. If a word is in normal form $\alpha\beta$, all the cuffs in β can be inverted simultaneously, whether or not it has crosscaps. If the word contains crosscaps we can invert each cuff separately, so assume that α is a product (perhaps empty) of handles.

Certainly the entire word $\alpha\beta$ can be inverted, and then cyclic permutation yields $\alpha^{-1}\beta^{-1}$. The word β^{-1} is a product of cuffs, which can be permuted arbitrarily (circulation rule 2b), so it is enough to show that α^{-1} can be replaced by α. The word α^{-1} is also a product of handles. It would appear that one can just relabel the edges in α^{-1}, since they are not mentioned elsewhere.

However, one needs to consider the vertices incident to α. In fact, all edges in α are incident to just one vertex A. To see this, consider the full description of a single handle, including vertices: $AaBbCa^{-1}Db^{-1}E$. Comparing the endpoints of edges, $A = D$, $B = C$, $C = D$, and $B = E$. Hence they are identical. Similarly all the edges in the product α are incident to a single vertex A; likewise for α^{-1}, so relabelling the edges does indeed convert it to α.

(5.20) Trivial isomorphisms. The following list Whittlesey calls the *trivial isomorphisms* of a canonical word-system. They yield homeomorphic complexes.

1. Permute the words in the system.
2. Permute the order of cuffs in any word (circulation rule 2b).
3. Cyclically permute any word δ in a cuff (rule 2a).
4. Invert any cuff in a 'non-orientable' word, i.e., a word with crosscaps (5.18).
5. Simultaneously invert all cuffs in a word without crosscaps (5.19).
6. Relabel any vertex or edge.
7. Reorient any edge.

The rest of this paper is based on the following proposition of Whittlesey's, stated in [14]. It is a corollary of the topological characterisation quoted above in 5.3.

(5.21) Proposition *Two complexes (in which every edge and vertex is incident to at least one face) are homeomorphic iff their word-systems are combinatorially equivalent (5.5) iff their canonical forms are trivially isomorphic.* ∎

6 Reducing homeomorphism to graph isomorphism

We shall next argue that homeomorphism of 2-complexes, implemented through Whittlesey's reduction method, can be reduced efficiently to graph isomorphism.

(6.1) Homeomorphism of 1-complexes. In order to describe complexes completely using a word-system, we assumed that every vertex and every edge was incident to at least one face (5.2). This assumption cannot be retained if one needs a complete reduction to graph isomorphism, so we first consider reducing 1-complex homeomorphism to graph isomorphism.

An *isomorphism* between 1-complexes G and G' is a map f which takes vertices bijectively onto vertices and edges bijectively onto edges, and preserves incidence. In other words, e is incident to A iff $f(e)$ is incident to $f(A)$. Isomorphic 1-complexes are obviously homeomorphic.

(6.2) Canonical form for 1-complexes. (See [13].) Given a 1-complex G, a vertex P is regular if it is incident to exactly two edges, a and b: it can be eliminated by composing the two edges. Repeating this process as long as possible produces a canonical form in which the only vertices of degree 2 (i.e., with 2 edge-incidences) are those incident to single closed loops. We say 'canonical complex' for 'complex in canonical form.'

(6.3) Lemma *If two canonical 1-complexes are homeomorphic then they are isomorphic.*

Proof. We shall show this by arguments similar to those in Section 4.

Let $h: G \to G'$ be a homeomorphism. Partition the vertices of G into four sets T, L, I, and S, denoting 'terminal,' 'loops,' 'isolated,' and 'singular' respectively. They contain the vertices of respective degree 1, 2, 0, and ≥ 3. (If both ends of an edge are incident to a vertex then the edge contributes 2 to the degree.) The vertices of G' are similarly partitioned into T', L', I', and S'.

The homeomorphism h clearly takes S bijectively onto S' and I onto I'. Therefore we can begin by defining $f(u) = h(u)$ for all u in $S \cup I$.

Consider the components of $G \backslash S$ (viewed as a topological space). They are of five kinds: (i) open line-segments (both ends in S), (ii) half-open line-segments (one end in S and one in T), (iii) closed line-segments (both ends in T), (iv) loops (incident to L) and (v) isolated vertices. Of these five kinds of component, no two are homeomorphic.

The homeomorphism takes components bijectively onto components of the same kind and the closures of components bijectively onto the closures of components. Each component (except (v)) contains a unique edge of G or G', so this gives a bijective map f from edges of G to those of G'. Every closed-loop component is incident to a unique vertex in L or L': this allows us to extend f to L.

Every vertex in T is incident to a unique edge, containing a unique component, which is of type (ii) or (iii), and h takes its endpoint or endpoints to a vertex in T. We can extend f to T so that whenever $u \in T$ is incident to e then $f(u)$ is incident to $f(e)$. (There are two ways of defining f when e coincides with a component of type (iii).)

Since the closure of a component is mapped onto the closure of a component, f preserves incidences and is an isomorphism from G onto G'. ∎

We can now reduce 1-complex homeomorphism to graph isomorphism. Using the above lemma we want to reduce isomorphism of (canonical) 1-complexes to graph isomorphism. (This is not the same thing, since 1-complexes are more general than graphs.)

(6.4) We shall use the fact (see [5]) that isomorphism of directed graphs can be reduced efficiently to graph isomorphism, by the construction illustrated in Figure 9.

Given a 1-complex G we shall construct a directed graph D such that two 1-complexes are isomorphic iff the directed graphs are. The construction is counter-intuitive in the sense that

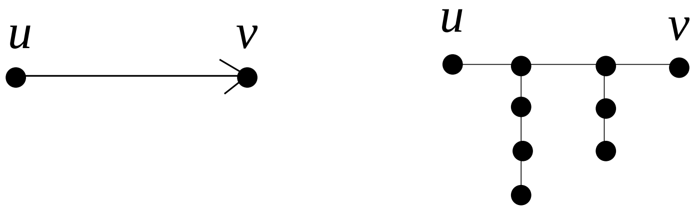

Figure 9: a directed edge simulated by undirected edges.

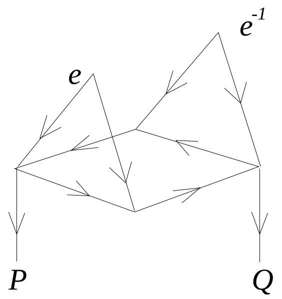

Figure 10: illustrating an edge gadget.

G and D are quite different in appearance (Figure 11). It has features in common with the reduction of word systems to directed graphs, discussed later.

Sometimes vertices of a directed graph will be called 'nodes,' to emphasise that the notion of a directed graph is combinatorial rather than geometric.

(6.5) Edge gadgets. For each inverse pair e, e^{-1} of edge labels, D has an 'edge gadget,' a certain directed subgraph. If we isolate the gadget from the rest of D, it has two source nodes, that is, nodes of indegree zero, which may conveniently be labelled e and e^{-1}. It has one or two sink nodes, nodes of outdegree zero, corresponding to the vertices P and Q (which could be the same) incident to e. They may conveniently be labelled P and Q. The sources are connected to the sinks by the small directed graph illustrated in Figure 10.

It is easy to show that any two such gadgets are isomorphic iff they have the same number (1 or 2) of sink nodes. An isomorphism must carry the source nodes to source nodes and sink nodes to sink nodes. There are two isomorphisms, each determined fully by where it maps the source nodes, or the sink nodes if there are two.

(6.6) The directed graph D is as follows. It has sink nodes corresponding to the vertices, of which the isolated vertices are also source nodes. The nodes corresponding to non-isolated vertices (if any) are all descendants of a single source node s. The source nodes are all at level 0. All the neighbours of s are at level 1. There is one for each edge of G. They are independent in the sense that no edges connect nodes at level 1.

Each level-1 node has two neighbours at level 2. Suppose the node represents an (undirected) edge joining P and Q. Its level-2 neighbours are the source nodes for an edge gadget with sink nodes P and Q. These edge gadgets furnish three levels, so the non-isolated sink nodes are all at level 4. See Figure 11.

(6.7) Remarks. Since the edge gadgets are embedded in a larger structure, their 'sources' have nonzero indegree. Conceivably edge gadgets could be matched against subgraphs which

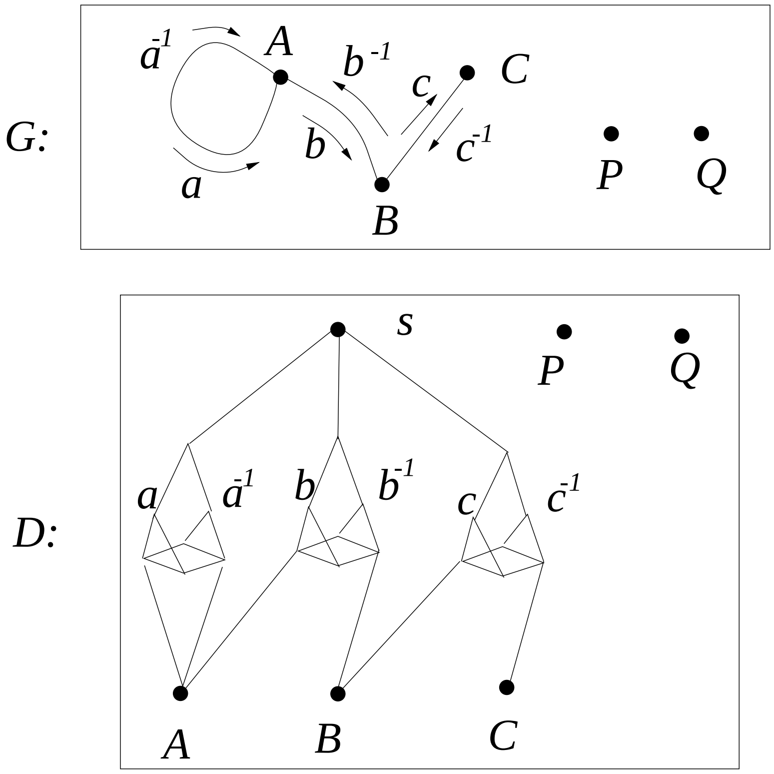

Figure 11: 1-complex G encoded as a directed graph D.

are not edge gadgets. It is the partition into levels which ensures that edge gadgets can be matched only with edge gadgets. The same idea is very important when we deal with word systems.

Two sources of an edge-gadget represent an inverse pair of directed edges, which are symmetrically connected to the same level-1 node. Of course, our construction can easily be made asymmetric to represent oriented 1-complexes. In the word-system construction discussed later, the connections will not be symmetric.

(6.8) Lemma *Given 1-complexes G and G', with associated directed graphs D and D' respectively,* (i) *G and G' are isomorphic iff D and D' are isomorphic, and* (ii) *If the 1-complexes are in canonical form then they are homeomorphic iff the directed graphs are isomorphic.*

Sketch of the proof. (i, 'only if') Let $f\colon G \to G'$ be an isomorphism. Define $g\colon D \to D'$ as follows. On the sink nodes of D, corresponding to the vertices of G, g imitates f. It takes the distinguished source s of D (if it exists: i.e., if G possesses edges), to the corresponding source of D'.

Let e be any edge of G, with incident vertices P and Q (perhaps the same). There is a unique corresponding level-1 node u in D, and the edge-gadget attached to it has sink node or nodes p and q corresponding to P and Q. Likewise D' possesses a level-1 node u' attached to an edge-gadget representing $f(e)$, whose sink nodes correspond to $f(P)$ and $f(Q)$. These must be $g(p)$ and $g(q)$ by construction. The two edge gadgets are of the same type (i.e., with one or two sinks), so g may be extended to take one isomorphically onto the other.

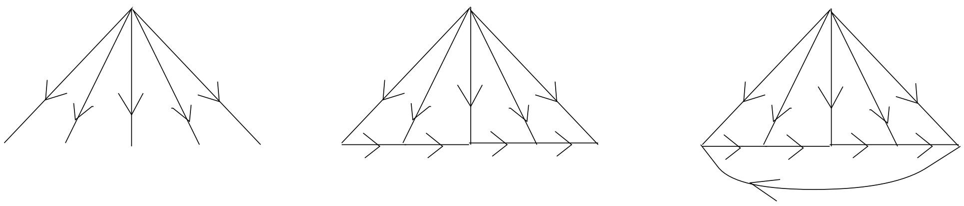

Figure 12: gadgets for a set, a list, and a directed cycle.

In this way a bijection g is formed between the nodes of G and those of G', and by its construction g is an isomorphism.

(i, 'if') Conversely, let g be an isomorphism from D to D'. It preserves levels. Construct a map f from vertices of G (sink nodes of D) and edges of G (level-1 nodes of D) to vertices and edges of G' in the obvious way: f is an isomorphism.

(ii) follows from Lemma 6.3. ∎

(6.9) Canonical form in polynomial time. 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 obviously be applied in polynomial time. Actually, Whittlesey's reduction procedure is just a modest extension of the Brahana-Tucker procedure for compact surfaces, which can be implemented in $O(n \log(n))$ time [11].

Therefore, to (efficiently) reduce homeomorphism of 2-complexes satisfying Whittlesey's simplifying assumption (5.2) to graph isomorphism, it is enough to (efficiently) reduce trivial isomorphism of canonical word-systems to directed graph isomorphism.

(6.10) Gadgets for sets, lists, and cycles. Next we see how to convert a word-system in normal form into a directed graph, such that two systems are trivially isomorphic (5.20) iff the corresponding directed graphs are isomorphic. The ideas are much the same as those leading to Lemma 6.8. The directed graph D constructed has a unique source node s and its sink nodes correspond to the vertex-labels of the system. Its nodes are separated into levels based on distance from the source node s: an isomorphism between two such graphs preserves levels. This is important since different 'gadgets' are introduced in the first few levels.

We make use of simple gadgets to represent a set, a list, and a directed cycle. Each gadget has a single source node, and the 'sink' nodes (or more properly, 'output' nodes, which have nonzero out-degree in the list and cycle gadgets) are all at the next level. Obviously the sink nodes for a set gadget can be permuted arbitrarily, those for a cycle gadget can be permuted cyclically, and those for a list cannot be permuted at all. The gadgets are illustrated in Figure 12.

(6.11) Multiple occurrences of objects. Where the definitions mention, for example, a set of objects then the directed graph D contains a gadget representing a set, whose 'sinks' are the objects, but if they mention a list of *occurrences* of objects, D contains a gadget whose 'sinks' are attached by directed edges to the objects in question. In other words, when occurrences are mentioned there is an extra level of indirection.

(6.12) Hierarchical construction of D. The directed graph D has a single source node s and its sink nodes correspond to the vertex-labels of the word system. It has gadgets for all the edge-labels in the system. The 'sinks' from these edge gadgets are connected to the appropriate sink nodes of D.

D is constructed with the list of trivial isomorphisms (5.20) in mind. The last two items in (5.20) are almost irrelevant, since they concern vertex and edge labellings, which are transparent in the graph representation. However to aid the description we retain the vertex- and edge-labels, bearing in mind that they do not figure in D.

D is assembled from various 'gadgets' based on the following hierarchy of definitions.

1. A word system is a *set* of words.

2. A word is an *ordered pair* (i.e, list) (α, β) where α is a *list* of crosscaps or handles and β is discussed in the next two items. A crosscap is an ordered pair of edge occurrences and a handle is an ordered quadruple of edge and inverted-edge occurrences. A cuff $t\delta t^{-1}$ is represented by an ordered triple (t, δ, t^{-1}) where δ is a cyclically ordered list of edge occurrences. Its *inverse* cuff is (t, δ^{-1}, t^{-1}).

   If α is empty, it should properly be a single vertex, and is represented as a list containing a single vertex occurrence. Similarly if β is empty.

3. When α is a nonempty product of crosscaps, β is a set of inverse pairs of cuffs, unordered: i.e., $\{(t, \delta, t^{-1}), (t, \delta^{-1}, t^{-1})\}$.

4. Otherwise, α is a product of handles, perhaps empty. Then β is represented as an unordered pair which could be denoted $\{\beta, \beta^{-1}\}$, where β and its inverse are represented as *sets* of cuffs.

(6.13) Lemma *If two word-systems are trivially isomorphic, then the directed graphs constructed are isomorphic.*

Proof. Let S and S' be two trivially isomorphic word-systems. Without loss of generality they share the same vertex- and edge-labels and differ only by one of the transformations listed in 5.20. Let D and D' be the corresponding directed graphs. The vertex- and edge-labels implicit in the construction are the same for both.

If S' is obtained by permuting the words in S, then D and D' are isomorphic since at levels 0 and 1 a *set* of words is represented.

If the cuffs are permuted in any word then the graphs are isomorphic since the cuffs are together represented either as an unordered pair of sets or a set of unordered pairs. The representation is independent of the order of cuffs within a word.

If in a cuff $t\gamma t^{-1}$ in S, the subword γ is cyclically permuted, then the graphs are the same since a cuff is encoded as (t, δ, t^{-1}) where δ is a *cyclically* ordered list of edge occurrences, representing γ.

If a cuff $t\delta t^{-1}$ in S is inverted in a word with crosscaps, then the graphs are the same, since the cuff and its counterpart $t\delta^{-1}t^{-1}$ in S' are both encoded as $\{(t, \delta, t^{-1}), (t, \delta^{-1}, t^{-1})\}$.

If all cuffs are inverted in a word without crosscaps in S, then graphs are the same since the product β in S and its counterpart β^{-1} in S' are both encoded as $\{\beta, \beta^{-1}\}$. ∎

(6.14) Levels apply locally. In the construction of the directed graph D there is one node at level 0, and those at level 1 correspond to the words of the system, and so on. Beyond this the 'global' level is ignored, though the idea of level 'locally' within a gadget (isolated from the rest of D) is necessary. Similarly, tracing backwards from the sink nodes only a few levels count.

(6.15) Lemma *If two word-systems S and S' yield isomorphic directed graphs D and D', then they are trivially isomorphic.*

Proof. Let f be an isomorphism between D and D'. The only sinks, nodes of out-degree zero, are those corresponding to vertex-labels of the word-systems. The isomorphism carries the sinks of D onto the sinks of D'. This allows us to alter the vertex labels in S' to match those of S.

Going back one level from these vertices we reach the 'sink' vertices from edge gadgets and vertex occurrences representing empty words α and β. These are easily distinguished and the isomorphism carries edge-gadgets onto edge-gadgets. This allows edge-labels to be altered and/or inverted so S' may now be assumed to have the same edge- and vertex-labels as S.

Considering Level 1, the isomorphism takes a node representing a word $\alpha\beta$ to one representing $\alpha'\beta'$. The first part α is either empty or a product of crosscaps or handles, and because of the list representation, it must be identical to α'. So α is a nonempty product of crosscaps iff α' is. This determines how β and β' should be represented.

Suppose that α contains crosscaps, so β and β' are represented as sets of pairs of cuffs. The isomorphism therefore takes a pair

$$\{(t, \delta, t^{-1}), (t, \delta^{-1}, t^{-1})\}$$

onto another pair

$$\{(t, \delta', t^{-1}), (t, {\delta'}^{-1}, t^{-1})\}$$

where $t\delta t^{-1}$ and $t\delta' t^{-1}$ are the actual subwords in S and S'. By construction δ' must be a cyclic permutation of δ or of δ^{-1}, and it is obtainable from δ by trivial isomorphisms.

Finally suppose that crosscaps are not involved, so the representation of β is denoted $\{\beta, \beta^{-1}\}$, β in its turn denoting a set of cuffs; similarly for β^{-1} and β'. The isomorphism takes the gadget for β onto that for β' or ${\beta'}^{-1}$. In the latter case can invert β (a trivial isomorphism on S) to bring about the first case. It can therefore be assumed that the isomorphism takes the gadget for β onto that for β'.

Now if the isomorphism takes a cuff gadget (t, δ, t^{-1}) in D to a gadget $(t', \delta', {t'}^{-1})$ in D', then $t = t'$ and δ' is a cyclic permutation of δ. Thus $t\delta' t^{-1}$ can be obtained from $t\delta t^{-1}$ by a trivial isomorphism. Therefore β' can be obtained from β by trivial isomorphisms.

In this way S can be transformed into S' by a sequence of trivial isomorphisms. ▮

(6.16) The full construction. Let K be a 2-complex. Let S be the subcomplex of 2-cells with all incident faces and vertices, and let G be the subcomplex containing isolated vertices and edges not incident to any face. We know how to construct directed graphs for S (a word-system) and G (a 1-complex) such that the homeomorphism problem for each is equivalent to isomorphism for the directed graphs.

(6.17) To get the general result, we create directed graphs C and D for S and G separately, and we make isomorphic copies if necessary to ensure that these graphs have disjoint vertex-sets. Let s be the source for C (if C is empty s should be an isolated vertex), and t the 'special' source for D. We then form an ordered pair (C, D) by putting the sources (s, t) in a list gadget. This allows C or D to be empty — the corresponding nodes in the list gadget being left with no attachments. D may have other sources (isolated vertices), which remain in the construction but are not reachable from t.

Let V be the 0-cells of K. Connect the sink vertices in C (if it is nonempty) to the vertices in V they represent; similarly for D. Let $C + D$ denote the resulting structure. The following lemma is an easy corollary of Lemmas 6.8, 6.13, and 6.15.

(6.18) Lemma *Let K and K' be 2-complexes with directed graphs $C + D$ and $C' + D'$ constructed as described in 6.17, from subcomplexes S, G, S', and G'. Let $I = S \cap G$; it contain only 0-cells.*

Then $C + D$ and $C' + D'$ are isomorphic iff the complexes S and S' and the complexes G and G' are simultaneously homeomorphic by maps which agree on I. In other words, K and K' are homeomorphic iff $C + D$ is isomorphic to $C' + D'$. ∎

Of course, all the constructions described can be implemented in polynomial time. We have therefore proved (Lemmas 4.2, 6.18)

(6.19) Theorem *Homeomorphism of 2-complexes is polynomially equivalent to graph isomorphism.* ∎

7 Acknowledgement

The authors wish to thank two anonymous referees for their careful reading of this paper and a related paper submitted previously.

8 References

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

2. T. K. Dey (1997). Optimal algorithms to detect null-homologous cycles on 2-manifolds. *IJCGA* **7:3**, 167–174.

3. T. K. Dey, H. Edelsbrunner, and S. Guha (1999). Computational Topology. In *Advances in Discrete and Computational Geometry (Contemporary mathematics* **223***)*, ed. B. Chazelle, J. E. Goodman, and R. Pollack, American Mathematical Society, 109–143.

4. T. K. Dey and H. Schipper (1994). A new technique to compute polygonal schema for 2-manifolds with application to null-homotopy detection. *Proc. 10th Annual Symposium on Computational Geometry.* Also *Discrete Computational Geometry,* **14**, (1995), 93–110.

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

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

7. R. Lyndon and P. Schupp (1977). *Combinatorial group theory.* Springer.

8. C. Ó Dúnlaing, C. Watt, D. Wilkins, and C-K Yap (1995). Miscellaneous topological algorithms. Report ALCOM-II-430 (1995) and TCDMATH 97-01 (1997).

9. E. Spanier (1966). *Algebraic topology.* McGraw Hill.

10. G. Vegter (1997). Computational Topology. In *Handbook of Discrete and Computational Geometry,* ed. J. E. Goodman and J. O'Rourke, CRC Press, 517–536.

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

12. E. Whittlesey (1958). Classification of finite 2-complexes. *Proc. AMS* **x:6**, 841–845.

13. E. Whittlesey (1960). Finite surfaces: a study of finite 2-complexes I: local structure. *Mathematics Magazine* **34:1**, 11–22.

14. E. Whittlesey (1960). Finite surfaces: a study of finite 2-complexes II: the canonical form. *Mathematics Magazine* **34:2**, 67–80.