

2. DIAGRAMS OF CELLULAR SURFACE MAPS

Objektyp: **Chapter**

Zeitschrift: **L'Enseignement Mathématique**

Band (Jahr): **43 (1997)**

Heft 3-4: **L'ENSEIGNEMENT MATHÉMATIQUE**

PDF erstellt am: **26.09.2024**

Nutzungsbedingungen

Die ETH-Bibliothek ist Anbieterin der digitalisierten Zeitschriften. Sie besitzt keine Urheberrechte an den Inhalten der Zeitschriften. Die Rechte liegen in der Regel bei den Herausgebern.

Die auf der Plattform e-periodica veröffentlichten Dokumente stehen für nicht-kommerzielle Zwecke in Lehre und Forschung sowie für die private Nutzung frei zur Verfügung. Einzelne Dateien oder Ausdrucke aus diesem Angebot können zusammen mit diesen Nutzungsbedingungen und den korrekten Herkunftsbezeichnungen weitergegeben werden.

Das Veröffentlichen von Bildern in Print- und Online-Publikationen ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. Die systematische Speicherung von Teilen des elektronischen Angebots auf anderen Servern bedarf ebenfalls des schriftlichen Einverständnisses der Rechteinhaber.

Haftungsausschluss

Alle Angaben erfolgen ohne Gewähr für Vollständigkeit oder Richtigkeit. Es wird keine Haftung übernommen für Schäden durch die Verwendung von Informationen aus diesem Online-Angebot oder durch das Fehlen von Informationen. Dies gilt auch für Inhalte Dritter, die über dieses Angebot zugänglich sind.

conjugate of itself or its inverse. In Section 5 we indicate how the algorithm can be described in terms of formal manipulations of presentations of surface groupoids, by describing a trivial example which illustrates the algorithm.

2. DIAGRAMS OF CELLULAR SURFACE MAPS

In this section we introduce the setting in which we shall work, and describe the connection with group theory.

2.1. DEFINITIONS. By a *two-dimensional CW-complex* X we shall mean a combinatorial CW-complex of dimension at most two, in which each cell has a preferred orientation. Formally we have the following situation.

As a set, X is the disjoint union of three sets V , E , F , whose elements are called the *vertices*, *edges*, and *faces*, of X , respectively.

There are given maps ι , τ , from E to V , and, for each edge e , the vertices ιe , τe are called the *initial* and *terminal* vertices of e , respectively. If $\iota e = \tau e$, we say that e is a *loop*. For each vertex v we understand that $\iota v = v = \tau v$.

We write $E^{\pm 1}$ for the Cartesian product $E \times \{1, -1\}$, and for any $(e, \epsilon) \in E^{\pm 1}$ we write e^ϵ for (e, ϵ) . We identify $e = e^1$. We use the same conventions for the faces. For a vertex v , we understand that $v^1 = v = v^{-1}$.

For $e \in E$, we define $\iota(e^{-1}) = \tau e$, and $\tau(e^{-1}) = \iota e$.

Each face f of X has an associated *boundary cycle* which is a finite alternating sequence

$$(2.1) \quad \partial f = v_0, e_1^{\epsilon_1}, v_1, \dots, v_{n-1}, e_n^{\epsilon_n}, v_n,$$

where $n \geq 0$, the v_i are vertices, $v_n = v_0$, the e_i are edges, each ϵ_i is ± 1 , and $\iota(e_i^{\epsilon_i}) = v_{i-1}$, $\tau(e_i^{\epsilon_i}) = v_i$. We define

$$\partial f^{-1} = v_n, e_n^{-\epsilon_n}, v_{n-1}, \dots, v_1, e_1^{-\epsilon_1}, v_0.$$

It is thus implicit that we are assigning to each closed two-cell a polygonal structure, and a distinguished vertex where the boundary cycle begins and ends. Notice that we are allowing vertices of valence one, so a boundary cycle need not be reduced.

A one-dimensional CW-complex, that is, a two-dimensional CW-complex with no faces, will be called a *graph*.

2.2. DEFINITIONS.

Let X be a two-dimensional CW-complex.

The edges and vertices of X form a graph, denoted $X^{(1)}$ and called the *one-skeleton* of X . We say that X is *connected* if $X^{(1)}$ is a connected graph. The free groupoid on $X^{(1)}$ will be denoted $\pi X^{(1)}$.

Let e be an edge of X . A vertex v is said to be *incident* to e if $\iota e = v$ or $\tau e = v$, and in the former (resp. latter) case we call e^{-1} (resp. e^1) an edge with a *distinguished* incidence to v .

Let f be a face of X , and suppose that ∂f is as in (2.1). The distinguished vertex $v_0 = v_n$ will be denoted $\text{vert}(f)$. There is associated an element of $\pi X^{(1)}$, denoted $\omega(f)$, which is $\text{vert}(f)$ if $n = 0$, and is the product $e_1^{\epsilon_1} \cdots e_n^{\epsilon_n}$ if $n \geq 1$. A vertex v is said to be *incident* to f if some v_i equals v , and we then call the pair (f, i) a face with a *distinguished* incidence to v . An edge e is said to be *incident* to f if some e_i equals e , and we then call the pair (f, i) a face with a *distinguished* incidence to e . For $1 \leq i \leq n$, we say that $e_{i-1}^{\epsilon_{i-1}}$ and $e_i^{-\epsilon_i}$ are *adjacent* in f , where the subscripts are interpreted modulo n , that is, 0 is interpreted as n .

By the *groupoid* of X , denoted πX , we mean the groupoid obtained from $\pi X^{(1)}$ by imposing the relation $\omega(f) = \text{vert}(f)$ for each face f of X . For any vertex v of X the *fundamental group of X at v* , denoted $\pi(X, v)$, is the subgroup(oid) of πX consisting of all elements with initial and terminal vertex v . If X is connected then changing the choice of v gives an isomorphic group, and there is specified an isomorphism which is unique up to conjugacy.

We say that X is a (closed) *CW-surface* if it is a finite, connected, two-dimensional CW-complex such that for each edge e there are exactly two faces with a distinguished incidence to e , and for each vertex v the edges with a distinguished incidence to v form a single (non-empty) equivalence class under the equivalence relation generated by the relation of being adjacent in some face.

The former condition, on edges, implies that the edges with a distinguished incidence to v form cycles under the relation of being adjacent in some face, and the latter condition, on oriented edges, then requires that there be exactly one cycle at v , called the *edge cycle around v* .

2.3. EXAMPLES.

(i) A simplicial complex structure on a surface yields a CW-surface.

(ii) Any surface group presentation $\langle S \mid r \rangle$ has an *associated* CW-surface X with one vertex, denoted v , with edge set S , and with one two-cell, denoted

f , and the boundary cycle of f is the sequence in $S^{\pm 1}$ determined by r . Here $\pi X^{(1)} = \pi(X^{(1)}, v)$, $\pi X = \pi(X, v)$, and there are natural identifications $\pi X^{(1)} = \langle S \mid \rangle$ and $\pi X = \langle S \mid r \rangle$.

2.4. DEFINITIONS.

Let $X = (V, E, F)$ be a CW-surface.

The *dual surface* $X^* = (F^*, E^*, V^*)$ of X is defined to be any CW-surface constructed as follows. Let V^* , E^* , F^* be copies of V , E , F respectively, with bijective correspondence denoted by $*$. Then X^* has F^* , E^* and V^* as vertex set, edge set, and face set, respectively. For any $e \in E$, there are two different faces with a distinguished incidence to e . If we denote these by (f, i) , (f', i') , with f, f' in F , then in X^* , the edge e^* is incident to the vertices f^* , f'^* . For any $v \in V$, the elements of X with a distinguished incidence to v are cyclically ordered, and this cyclically ordered set is called the *face-and-edge cycle around v* ; by considering alternate terms we get the *edge cycle around v* and the *face cycle around v* . Applying $*$ to the face-and-edge cycle around v gives a cyclic sequence which is taken to be the boundary cycle of v^* , once a distinguished vertex is chosen.

We say that X is *oriented* if each edge e occurs with opposite signs in the two faces with a distinguished incidence to e , and, in this event, we can use the signs to orient the dual surface X^* consistently.

We say that X is *orientable* if we can obtain an oriented CW-surface by replacing some faces with their inverses; otherwise X is *unorientable*.

In the remainder of this section and the next, all paragraphs which are devoted to the unorientable case are marked with a Maltese cross (\boxtimes), and, by skipping these, the reader interested primarily in orientable surfaces can follow the discussion for that case.

Consider a loop e in X , let v be the vertex incident to e , and consider a face f incident to e . Here two vertices in the boundary cycle of f are equal, which results in f getting attached to itself at a point.

If this attachment is performed without twisting, we say that e is an *orientable* or *two-sided* loop. If X is orientable than clearly all loops are orientable.

\boxtimes If this attachment is performed with a twist, we say that e is an *unorientable*, or *one-sided* loop in X . This can be expressed in a more combinatorial manner by saying that e is unorientable if the boundary cycle of $f^{\pm 1}$, viewed cyclically, contains a subsequence e', v, e, v, e'' , and the sequence e', e^{-1}, e, e''^{-1} of four distinct edges with a distinguished incidence to v is

not in the correct order, with respect to the cyclic ordering by face-adjacency. Here the CW-complex resulting from collapsing e to a vertex is a CW-surface.

A useful way to codify a groupoid presentation of πX is to write $\langle E \mid \omega(\partial f) (f \in F) \rangle$, so $\pi X = \langle E \mid R \rangle$, where R is the set of words in $E^{\pm 1}$ determined by the boundary cycles, one word for each face. There is no need to specify the vertices, since they correspond to equivalence classes in $E^{\pm 1}$ under the equivalence relation generated by face-adjacency.

We can form a new CW-surface Y from X by successively erasing edges incident to two distinct faces (so melding two faces into one) until only one face f is left. The set E' of erased edges then corresponds to a maximal subtree in the one-skeleton of the dual complex of X . Here X and Y both have the same vertex set, and $Y^{(1)}$ can be viewed as the complement of E' in $X^{(1)}$, and πY is a subgroupoid of πX . One can even choose a retraction of πX onto πY by choosing a suitable image in $\pi Y^{(1)}$ of each erased edge. Notice that $\omega(\partial f)$ is an element of the free group $\pi(Y^{(1)}, v)$, where $v = \text{vert}(f)$, and there is an isomorphism $\pi(X, v) \simeq \pi(Y, v)$. Hence we have a homomorphism from a free group $\pi(Y^{(1)}, v)$ onto $\pi(Y, v) \simeq \pi(X, v)$, and the kernel is the normal subgroup generated by $\omega(\partial f)$.

Frequently we will want to alter the choice of E' by exchanging an edge b for some edge y of Y , such that b divides the face f into two faces f_1, f_2 , each of which has a single occurrence of y in the boundary cycle. Either of these subfaces can be used to choose an element of the free groupoid $\pi Y^{(1)}$ which gets equated to b in the groupoid πX . We now have a new Y' with $Y'^{(1)} = Y^{(1)} \cup \{y\} - \{b\}$, and a map $Y'^{(1)} \rightarrow \pi Y^{(1)}$ which induces an isomorphism of free groupoids $\pi Y \simeq \pi Y'$. The single face f' of Y' is obtained by glueing together f_1 and f_2 along the two copies of b . It is straightforward to check that the isomorphism $\pi Y \simeq \pi Y'$ carries $\omega(f)$ to a conjugate of $\omega(f')$ or its inverse. The situation is amply illustrated in Figure 2.1, which depicts a labelled CW-subcomplex formed from two faces which are adjacent in two ways, so there are two ways to choose edges to be erased. In general, the symbol denoting an oriented path in $X^{(1)}$ is placed on the right of the path, and similarly for edges. Here, if we erase b , we get a face with clockwise boundary cycle $\bar{z}\bar{y}\bar{x}\bar{q}\bar{p}xyzrs$, where overlines denote inverses. But if we erase y , we get a face with clockwise boundary cycle $\bar{c}\bar{b}\bar{a}\bar{p}\bar{q}abcsr$. Algebraically, erasing b corresponds to using one of the small faces as a relation to eliminate b by identifying $b = \bar{a}qxyz\bar{s}\bar{c}$, and then $\bar{c}\bar{b}\bar{a}\bar{p}\bar{q}abcsr = \bar{c}(cs\bar{z}\bar{y}\bar{x}\bar{q}\bar{a})\bar{a}\bar{p}\bar{q}\bar{a}(\bar{a}qxyz\bar{s}\bar{c})csr = s(\bar{z}\bar{y}\bar{x}\bar{q}\bar{p}xyzrs)\bar{s}$. Choosing whether to erase b or y affects the choice of free group mapping onto the

surface group, but, as we have just seen, the free groups are related via an isomorphism which respects the given relators, up to conjugacy and inverse.

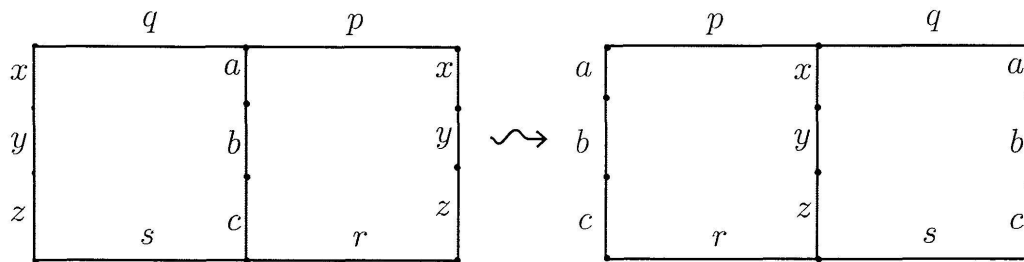


FIGURE 2.1

Changing edges to be erased

It usually happens that we are given a basis S of the free group $\pi(Y^{(1)}, v)$, and an expression of $\omega(\partial f)$ as a word r in $S^{\pm 1}$, so we get a presentation $\pi(X, v) = \langle S \mid r \rangle$. One standard method of choosing a basis S of $\pi(Y^{(1)}, v)$ is to choose a maximal subtree Y_0 of $Y^{(1)}$, and associate a free generator to each edge of $Y^{(1)} - Y_0$ in the natural way. This choice of S ensures that the above presentation is a surface group presentation. An alternative construction is to collapse the edges E_0 of Y_0 to get a new CW-surface Z with one vertex and one face, such that πZ is isomorphic to $\pi(X, v)$. Algebraically, in passing from the groupoid presentation $\pi X = \langle E \mid R \rangle$ to the groupoid presentation $\pi Y = \langle E - E' \mid \omega(\partial f) \rangle$, we successively use the erased edges to meld pairs of relators, and then annihilate the elements of E_0 to get a surface group presentation of πZ .

We will be interested in the situation where we are given a presentation to start with.

2.5. REMARKS. Let G be a surface group, and let $\langle S \mid r \rangle$ be a surface group presentation of G .

In this article we will be applying homotopy equivalences to a CW-surface X with fundamental group G , and we wish to ensure that the given presentation is always recoverable. Some of the difficulties arise from the choices involved. The choice of base point v affects the data only up to conjugacy. The choice of set of edges E' determining a maximal tree in the one-skeleton of the dual surface affects the data up to isomorphism of the covering free group $\pi(X^{(1)} - E', v)$, and we have seen that the isomorphism respects the relators up to conjugacy and inverse. Thus if S is associated to a basis of one of the free groups in such a way that r corresponds to

a conjugate of the relator or its inverse, then S is associated to a basis of each of the free groups in such a way that r corresponds to a conjugate of the relator or its inverse. We want to ensure that each homotopy equivalence specifies an isomorphism of covering free groups so as to respect relators in this way.

We start with the CW-surface associated to the presentation $\langle S \mid r \rangle$, and apply homotopy equivalences using four operations called subdivision, erasing, collapsing and expanding.

Subdivision of edges and faces changes the covering free group by a simple isomorphism which preserves relators.

Erasing a set of edges E'' which determine a subtree of the one-skeleton of the dual complex changes, by a simple isomorphism which preserves relators, the covering free group corresponding to a choice of E' containing E'' .

Collapsing, in the cases of interest to us, concerns the three elementary operations of collapsing to a vertex an edge which is not a loop, collapsing to an edge a two-edged face which is not a sphere, and collapsing a one-edged face to a vertex. If we want to collapse an edge which is not a loop, we first adjust the choice of E' to ensure that it does not contain the edge to be collapsed. It is then straightforward to check that, for each of the three elementary collapsing operations, the covering free group changes by a simple isomorphism which preserves relators.

Expanding is the reverse of collapsing, and changes the covering free group by a simple isomorphism which preserves relators.

At any stage we can lose the base vertex, and prior to its disappearance we have to change the covering free group by conjugating by a chosen path to a new base vertex.

2.6. DEFINITIONS. Let X_1 and X_2 be two-dimensional CW-complexes. A map of sets $\beta: X_1 \rightarrow X_2$ is said to be *cellular* if the following are satisfied:

If v is a vertex of X_1 , then $\beta(v)$ is a vertex of X_2 .

If e is an edge of X_1 , then $\beta(e)$ is a vertex or an edge of X_2 , and $\iota\beta(e) = \beta\iota(e)$, $\tau\beta(e) = \beta\tau(e)$.

If f is a face of X_1 , exactly one of the following holds:

$\beta(f)$ is a vertex v , and all the terms of $\beta(\partial f)$ are v ;

$\beta(f)$ is an edge e , one of the terms of $\beta(\partial f)$ is e , one is e^{-1} , and the rest are vertices;

$\beta(f)$ is a face, every edge incident to f is mapped to an edge, and $\partial\beta(f) = \beta(\partial f)$; here, with ∂f as in (2.1), we understand that

$$(2.2) \quad \beta(\partial f) = \beta(v_0), \beta(e_1)^{\epsilon_1}, \beta(v_1), \dots, \beta(v_{n-1}), \beta(e_n)^{\epsilon_n}, \beta(v_n).$$

This definition of cellular map is more restrictive than the usual definition, but it does include simplicial maps, with suitably chosen orientations of simplices, so we do not lose any homotopy classes of maps.

2.7. CONSTRUCTION. Let $\beta: X_1 \rightarrow X_2$ be a cellular map of two-dimensional CW-complexes.

Then β induces a cellular map on the one-skeletons $\beta^{(1)}: X_1^{(1)} \rightarrow X_2^{(1)}$, and this determines a groupoid homomorphism $\pi(\beta^{(1)}): \pi X_1^{(1)} \rightarrow \pi X_2^{(1)}$. The latter then induces a groupoid homomorphism $\pi(\beta): \pi X_1 \rightarrow \pi X_2$. If we specify a vertex v of X_1 , then we obtain a group homomorphism $\pi(\beta, v): \pi(X_1, v) \rightarrow \pi(X_2, \beta(v))$.

The *diagram associated to β* consists of the CW-complex X_1 together with a labelling of its cells, which labels each cell with its image cell in X_2 . For any cells c_1 of X_1 , c_2 of X_2 , if $c_2 = \beta(c_1)$ we say that c_1 is a c_2 -cell. With our definition of cellular map, there are three types of labelling that a face of X_1 can have, namely, we can have a v -face, an e -face, or an f -face, and these can be depicted as in Figure 2.2.

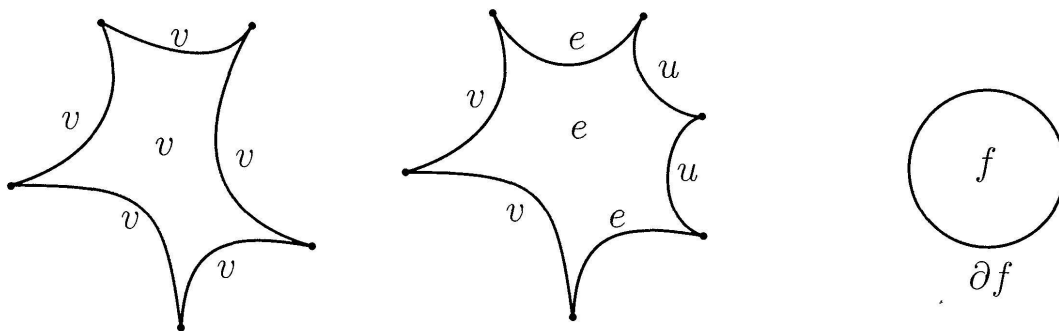


FIGURE 2.2

The three types of labelled cell

We remark that the labelled regions in Figure 2.2 are precisely the types of regions used by Ol'shanskii [15]. These diagrams, which represent elementary concepts in topology, can be viewed as Lyndon-van Kampen diagrams in which trivial relators play a larger part than usual. Here we have a groupoid setting which allows various vertices, rather than just the one vertex allowed when considering groups.

In terms of diagrams, β is a *covering* if for each vertex v of X_1 the cycle of labels (in X_2) of faces (in X_1) with a distinguished incidence to v is precisely the cycle of faces (in X_2) with a distinguished incidence to the label $\beta(v)$ of v . Also β is a *branched covering* if for each edge e of X_1 the labels (in X_2) of the two faces (in X_1) with a distinguished incidence to e are precisely the two faces (in X_2) with a distinguished incidence to the label $\beta(e)$ of e ; here deleting all vertices leaves a covering. We will abuse notation and say that β is a *pinching* if X_1 is obtained from X_2 by slicing open non-loop edges and inserting punctured projective planes and punctured tori with cell structures and labels as depicted in Figure 3.1. Here β acts by collapsing these subsurfaces to edges, so is homotopic to a pinching, and our abuse of notation is reasonable.

Our main activity will be to apply operations to these diagrams. Let us note one which will be used frequently.

2.8. CONSTRUCTION (type: subdivision). Suppose we are given two-dimensional CW-complexes X_1, X_2 , a face f of X_1 , and a cellular map $\beta: X_1 - \{f\} \rightarrow X_2$, such that $\pi(\beta^{(1)})(\omega(f))$ is trivial in $\pi X_2^{(1)}$, that is, $\pi(\beta^{(1)})(\omega(f)) = \beta(\text{vert}(f))$.

We wish to subdivide f to obtain a refinement X'_1 of X_1 , and an extension $\beta': X'_1 \rightarrow X_2$ of β .

Let ∂f be as in (2.1), and let $\beta(\partial f)$ denote the sequence in (2.2). Since $\pi(\beta^{(1)})(\omega(f))$ is trivial in $\pi X_2^{(1)}$, either every term of (2.2) is a vertex v , or some subsequence of (2.2) has the form $e^\epsilon, v, v, \dots, v, e^{-\epsilon}$. In the former case we can extend β to X_1 by labelling f as v . In the latter case, we can subdivide f into two faces by adding an edge with label the vertex $u = \iota(e^\epsilon)$ slicing off a piece of f with label e , as in Figure 2.3; here we use \emptyset to indicate a region with boundary label which determines a trivial element of the free groupoid.

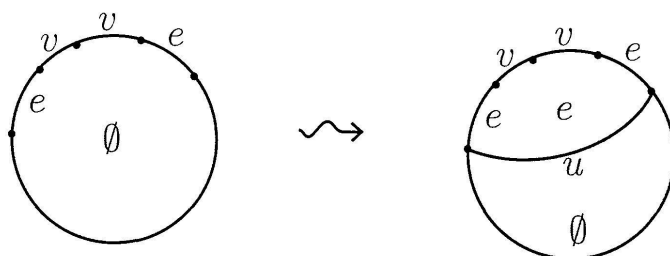


FIGURE 2.3
Subdividing

We can continue in this way, and by induction on the length of ∂f , we obtain a subdivision of f which allows an extension of β .

It is convenient to mention a more complicated operation at this stage.

2.9. CONSTRUCTION (type: subdivision). Suppose we are given two-dimensional CW-complexes X_1, X_2 , and a cellular map $\beta^{(1)}: X_1^{(1)} \rightarrow X_2^{(1)}$ of the one-skeletons, such that, for each face f of X_1 , $\pi(\beta^{(1)})(\omega(f))$ is trivial in πX_2 , that is, there exists $d \geq 0$, and elements w_i of $\pi X_2^{(1)}$, and faces f_i of X_2 such that, in the free groupoid $\pi X_2^{(1)}$,

$$(2.3) \quad \pi(\beta^{(1)})(\omega(f)) = \prod_{i=1}^d w_i^{-1} \omega(f_i)^{\epsilon_i} w_i.$$

Essentially as in the previous construction, we wish to subdivide each face of X_1 to obtain a refinement X'_1 of X_1 , and an extension $\beta': X'_1 \rightarrow X_2$ of $\beta^{(1)}$.

Let f be a face of X_1 , and suppose (2.3) holds. We first subdivide f into $d + 1$ two-cells by drawing in d balloons-on-sticks, as in Figure 2.4, which are subdivided and labelled in such a way that, if the boundary cycle is read clockwise, the labelling of the i th stick, starting at the basepoint, gives the word w_i , and the labelling of the boundary of the i th balloon, starting at the attaching point and reading counter-clockwise, is $\partial f_i^{\epsilon_i}$. We label the i th balloon f_i , and orient it in the manner dictated by the label on the boundary.

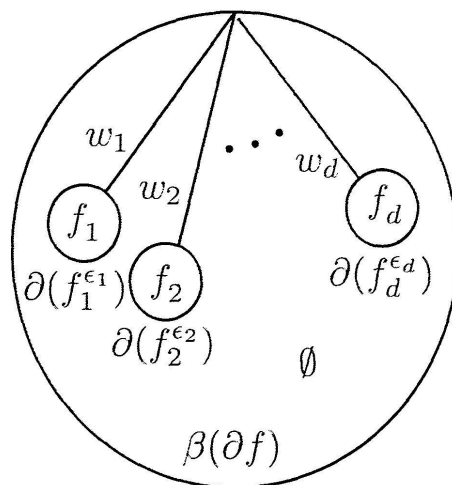


FIGURE 2.4

Subdividing for a relation

In the subdivided f there remains a single two-cell f' which is not labelled. Here, (2.3) implies that

$$\pi(\beta^{(1)})(\omega(f')) = \pi(\beta^{(1)})(\omega(f)) \prod_{i=1}^d w_i^{-1} \omega(f_i)^{-\epsilon_i} w_i = \beta^{(1)}(\text{vert}(f'))$$

in $\pi X_2^{(1)}$. By Construction 2.8, we can subdivide f' , and extend $\beta^{(1)}$. Thus we find we can subdivide each face of X_1 to obtain a refinement X'_1 of X_1 , and an extension $\beta': X'_1 \rightarrow X_2$ of $\beta^{(1)}$.

We conclude this section by recalling how one constructs a cellular map of CW-surfaces from a homomorphism of surface groups.

2.10. CONSTRUCTION (type: subdivision). Let $\alpha: G_1 \rightarrow G_2$ be a homomorphism of surface groups.

Let $\langle S_1 \mid r_1 \rangle$, $\langle S_2 \mid r_2 \rangle$ be surface group presentations of G_1 , G_2 , respectively, and choose a lifting $A: \langle S_1 \mid \rangle \rightarrow \langle S_2 \mid \rangle$ of α . Thus A is a homomorphism of free groups such that $A(r_1)$ lies in the normal subgroup generated by r_2 , and the resulting homomorphism $\langle S_1 \mid r_1 \rangle \rightarrow \langle S_2 \mid r_2 \rangle$ is α .

Let $X_1 = (w, S_1, g)$ and $X_2 = (v, S_2, f)$ denote the CW-surfaces associated to the presentations $\langle S_1 \mid r_1 \rangle$ and $\langle S_2 \mid r_2 \rangle$, respectively.

We want to subdivide X_1 to obtain a CW-surface X'_1 , and a cellular map $\beta': X'_1 \rightarrow X_2$.

We begin by subdividing the one-skeleton $X_1^{(1)}$, to get a graph Γ , as follows. For each $s \in S_1$, $A(s)$ is a word in S_2 , possibly empty, and hence $A(s)$ is either 1, or is a non-empty reduced word $e_1^{\epsilon_1} \cdots e_n^{\epsilon_n}$ in S_2 , that is, $n \geq 1$, each e_i lies in S_2 , and each ϵ_i is 1 or -1 . If $A(s) = 1$ we label s with the vertex v ; in the second case, we add $n - 1$ new vertices to subdivide s into n edges, denoted $s_1^{\epsilon_1}, \dots, s_n^{\epsilon_n}$, and label each s_i with e_i having the same orientation. Doing this for each element of S_1 gives us a labelled graph Γ , with the labels coming from $X_2^{(1)}$. Notice that the two-cell of the subdivided X_1 has as boundary cycle the subdivided r_1 , and the labelling gives a word in $\pi(X_2^{(1)}, v)$ which corresponds to $A(r_1)$. Since this word equals v in $\pi(X_2, v)$, we can use Construction 2.9 to further subdivide X_1 and obtain a CW-surface X'_1 and a cellular map $\beta': X'_1 \rightarrow X_2$. Moreover $\pi(\beta'^{(1)}, w): \pi(X_1'^{(1)}, w) \rightarrow \pi(X_2^{(1)}, v)$ can naturally be identified with $A: \langle S_1 \mid \rangle \rightarrow \langle S_2 \mid \rangle$, and $\pi(\beta', w): \pi(X'_1, w) \rightarrow \pi(X_2, v)$ can naturally be identified with $\alpha: \langle S_1 \mid r_1 \rangle \rightarrow \langle S_2 \mid r_2 \rangle$. Thus we have a cellular map of CW-surfaces which realizes α .

Notice that the cellular map is constructed from an equation of the form (1.1). We can apply the algorithm of the next section to this map, to get a new cellular map, from which we can extract a new equation of the form (1.1), without altering the given presentations, since at each step we can choose isomorphisms of the covering free groups which respect the relator up to conjugacy and inverse. Each element of S_1 will be transformed into a path in a labelled one-skeleton without changing the homotopy class in the surface underlying X_1 ; this amounts to choosing a new labelling for each element of S_1 , which, in turn, amounts to choosing a new lifting at the free group level. This whole process will then give non-trivial group-theoretical information, although not so much as in the topological situation.

3. THE ALGORITHM

Throughout this section let $\beta: X_1 \rightarrow X_2$ be a cellular map of CW-surfaces.

Let V , E , and F denote the sets of vertices, edges, and faces, respectively, of X_2 . We then have a diagram with V -faces, E -faces, and F -faces, as depicted in Figure 2.2.

The aim of this section is to alter β by composing it with various cellular homotopy equivalences of X_1 and X_2 (based on the operations of contracting, expanding, erasing, and subdividing), until we arrive at the minimum possible number of F -faces. These alterations of β can be viewed as homotopies, since one is free to imagine that there is a surface X underlying X_1 that has lines inscribed on it, and that these lines can be deformed continuously. Abusing notation then, we will say that the altered forms of β are *homotopic* to β .

3.1. CONSTRUCTION (type: subdivision). If X_2 has a loop e , we subdivide e by adding a new vertex v , and, in X_1 , subdivide each e -edge, and each e -face, by adding a new v -vertex, and a new v -edge, respectively.

By our definition of CW-surface, X_2 has an edge. Thus we have the following.

3.2. CONDITION. *There is at least one edge in X_2 , but there are no loops. Hence, in X_1 , no E -edge is a loop, or equivalently, all loops are V -loops.*