Algorithmica (1989) 4: 3-32

Algorithmica

Primitives for the Manipulation of Three-Dimensional Subdivisions[1]

David P. Dobkin[2] and Michael J. Laszlo[3]

Abstract. Algorithms for manipulating three-dimensional cell complexes are seldom implemented due to the lack of a suitable data structure for representing them. Such a data structure is proposed here along with the primitive operations necessary to make it useful. Applications of the structure are also given.

Key Words. Three-dimensional cell complexes, Data structures, Computational geometry.

1. Introduction. A major impediment to the implementation of algorithms that manipulate three-dimensional cell complexes and subdivisions is the lack of a suitable data structure. What is needed is a data structure powerful enough to model such objects yet simple enough to allow their manipulation in well-defined ways. We focus attention here on the development of such a data structure. Our structure is analogous (though one dimension higher) to the winged-edge [Ba], [BHS], [EW] and quad-edge [GS] data structures which are widely accepted for modeling 2-manifolds. Just as these structures can be used to represent both planar polygonal cell complexes in R^2 and surfaces of polyhedra, our data structure can model polyhedral complexes in R^3 and surfaces of 4-polyhedra.

Our results can be viewed as similar to the work done by Guibas and Stolfi in deriving the quad-edge structure. Lifting the results one dimension higher increases the complexity of our data structure. They consider an edge as their atom, and consider the edge rings to which it belongs. We consider a polygon-edge pair as an atom, and consider the polygon ring and edge ring to which it belongs. The quad-edge atom could be considered to connect two vertices and two polygons. Similarly, our atom connects two vertices and two polyhedra. We simplify our structure by treating only complexes that are orientable, and whose cells do not puncture the interior of other cells.

There are numerous applications we envision for such a data structure. One application we consider is that of decomposing a polyhedron into tetrahedra [Wö]. We rederive and model one of his applications in our system. A second application we consider is the implementation of an algorithm for incrementally

[1] This work was supported in part by the National Science Foundation under Grant No. DCR85-05517.
[2] Department of Computer Science, Princeton University, Princeton, NJ 08544, USA.
[3] Department of Electrical Engineering and Computer Science, University of Illinois at Chicago, Chicago, ILL 60680, USA.

Received April 15, 1987; revised March 14, 1988. Communicated by Chee-Keng Yap.

computing the Delaunay triangulation of a three-dimensional point set [AB], [Bh]. Further applications are also possible. For one, our data structure provides an approach to an efficient divide-and-conquer algorithm for building three-dimensional Voronoi diagrams. A second possibility is for modeling the motion of a three-dimensional polyhedron through time, which can be viewed as a four-dimensional polyhedron (in x, y, z, t space) where hidden surface removal is done by projecting into x, y, t space and taking t cross-sections to determine individual scenes.

What we attempt to achieve in this paper is a blend between a derivation of the data structure and a small set of primitive operators for its manipulation, the development of macro operations from these primitives, and the use of these macros in the first two applications mentioned above. The results of this paper are implementable (see [La] for details).

2. Definitions and Prerequisites. In this section we define the class of objects to be manipulated by our data structure. It is assumed the reader is familiar with some basic concepts of point-set topology.

2.1. Basic Definitions. Where T is a topological space, a k-cell is a subspace of T whose interior is homeomorphic to R^k, and whose boundary is nonnull. In this paper we assume that $T = R^3$, though our results hold for more general T. We call a 0-cell a *vertex*, a 1-cell an *edge*, a 2-cell a *polygon* or a *facet*, and a 3-cell a *polyhedron*. Note that a cell may be unbounded; for instance, an edge can be a closed segment (bounded by two vertices) or a ray (bounded by one vertex).

A cell complex of T is a finite collection C of cells of T such that:

(i) the relative interiors of cells of C are pairwise disjoint,
(ii) for each cell $c \in C$, the boundary bd c of cell c is the union of elements of C,
(iii) if $c, d \in C$ and $c \cap d \neq \emptyset$, then $c \cap d$ is the union of elements of C.

We let $\mathcal{U}(C)$ be the union of the cells of C, and consider C a subdivision of $\mathcal{U}(C)$. An n-dimensional complex for which every k-cell is contained in (the boundary of) some n-cell is called an n-complex.

The combinatorial boundary of cell c of C, denoted ∂c, is defined to be the set of cells of C contained in bd c. Note that $\mathcal{U}(\partial c) = \text{bd } c$. The combinatorial boundary ∂C of complex C is defined as the set of cells of C contained in bd $\mathcal{U}(C)$. An open cell $d \subset c$ is said to be a face of c; if in addition $c \neq d$, then d is a proper face of c. If one of c or d is a proper face of the other, c and d are said to be *incident*. For instance, a polyhedron is incident with each vertex, edge, and facet that lies in its boundary. The star of a cell c, denoted *star* c, is the subset of C consisting of the cells of which c is a face.

Given n-complex C, by convention these exists one (null) $(n+1)$-cell of which every n-cell of C is a face; likewise there exists one (null)(-1)-cell which is a face of every vertex. Distinct k-cells c and d (for $0 \leq k \leq n$) are then said to be

adjacent if (i) there exists some $(k-1)$-cell of C that is a face of both c and d, and (ii) there exists some $(k+1)$-cell of C of which each of c and d is a face. For instance, two vertices connected by an edge are adjacent; in addition, two facets incident to the same polyhedron and the same edge are adjacent.

2.2. Space-Duality. The space-dual of a complex C of space T is a second complex C^* of T for which there exists a one-to-one mapping Ψ from C onto C^* such that:

(i) the image of a k-cell under Ψ is an $(n-k)$-cell, and
(ii) cells c and d are adjacent in C iff cells $\Psi(c)$ and $\Psi(d)$ are adjacent in C^*.

In particular, with respect to 3-complexes C and C^*, each vertex (edge) of one corresponds to a polyhedron (facet) of the other, and adjacency relations between cells are preserved. The space-dual of cell c, denoted c^*, is that cell which corresponds to c under Ψ.

The complex C^* space-dual to C is by no means unique. However, up to the topological property which we intend our data structure to represent—adjacency relations between cells—the numerous complexes that serve as space-dual to C in T are identical. For our puposes, C^* is well defined. Furthermore, $(C^*)^* = C$.

The 3-complexes treated in this paper are regarded as subdivisions of the three-dimensional sphere $S^3 = \{(x, y, z, w) \| x^2 + y^2 + z^2 + w^2 = 1\}$, which is R^3 plus a point at infinity. A subdivision C of the closed ball D^3 is obtained by omitting one 3-cell c. The space-dual C^*, a subdivision of R^3, is obtained by omitting the vertex c^* from S^3.

3. Traversal Functions. In this section we present the five traversal functions *Fnext*, *Enext*, *Spin*, *Clock*, and *Sdual*. We call these *traversal* functions because they provide the means of traversing or moving about the cells of a complex. The first two traversal functions are used to move from cell to adjacent cell. *Spin* and *Clock* are used to change a local sense of rotation, so *Fnext* and *Enext* know the direction in which each is to traverse. The function *Sdual* is used to move between a complex and its space-dual. Since edges and facets interchange in the space-dual, the roles of *Fnext* and *Enext* are interchanged in going between C and C^*.

3.1. Basic Traversal Functions. Let f be a facet of complex C. The combinatorial boundary of f contains a ring of edges $e^0 \cdots e^{n-1}$ where edges e^i and e^{i+1} are adjacent in C (addition modulo n). We call this ring, denoted $\mathscr{E}_f$, the edge-ring of facet f. $\mathscr{E}_f$ can be assigned either of two senses of rotation whereby we can distinguish between the two edges belonging to the ring that are adjacent to edge e^i. The particular sense of rotation assigned to $\mathscr{E}_f$ is called the *orientation* of facet f. We write $\mathscr{E}_f = (e^0 \cdots e^{n-1})$ to indicate the edge-ring with sense of rotation

such that, of the two edges e^{i-1} and e^{i+1} of $\mathscr{E}_f$ adjacent to e^i, e^{i-1} precedes and e^{i+1} follows, edge e^i.

Similarly we define the facet-ring of edge e, denoted $\mathscr{F}_e$, to be the ring of facets $\mathscr{F}_e = (f^0 \cdots f^{m-1})$ incident to e in C. Facets f^{i-1} and f^{i+1} of $\mathscr{F}_e$ are adjacent to facet f^i, and f^{i-1} precedes, while f^{i+1} follows, facet f^i.

The atomic unit on which queries are formulated is called a facet-edge pair. This is a pair consisting of a facet f and an edge e, such that f and e are incident. The edge component e of a is denoted e_a, and the facet component f of a is denoted f_a. The facet-edge pair a determines two rings in C, these being edge-ring $\mathscr{E}_{f_a}$ and facet-ring $\mathscr{F}_{e_a}$. There are four versions of a which derive from the two senses of rotation that each of its two rings can assume. Henceforth by facet-edge pair we mean one such version—each of the two rings determined by the facet-edge pair has a fixed sense of rotation. $\mathscr{E}_a$ denotes the edge-ring $\mathscr{E}_{f_a}$ with sense of notation determined by a; facet-ring $\mathscr{F}_a$ is similarly defined.

Given facet-edge pair a, it is useful to distinguish between the two vertices incident to e_a (its endpoints), and between the two polyhedra incident to f_a. To distinguish between the endpoints, observe that the orientation of $\mathscr{E}_a$ directs edge e_a in a natural way. We call that vertex which serves as the endpoint both to e_a and to the edge that precedes e_a in $\mathscr{E}_a$, the *origin* of e_a, denoted *aOrg*. Similarly, the *destination aDest* of e_a is that vertex incident both to e_a and to the edge that follows e_a in $\mathscr{E}_a$.

To distinguish between the polyhedra incident to f_a, we assume edge e_a possesses a sense of rotation, called the *spin* of a. When the edge is viewed from destination toward origin, its spin is left-handed (right-handed) if its sense of rotation appears clockwise (counterclockwise). Where the spin of a is left-handed (right-handed), we define H_a^+ to be that open half-space determined by facet f_a from which the orientation of $\mathscr{E}_a$ appears clockwise (counterclockwise). We then define the positive polyhedron of a, denoted *aPpos*, to be that polyhedron p of C incident to f_a for which points of the interior of p arbitrarily close to the relative interior of f_a lie in H_a^+. The negative *polyhedron aPneg* of a is the other polyhedron of C incident to f_a. Figure 1 illustrates some of the definitions presented so far in this section.

We are now able to define the traversal functions *Fnext*, *Enext*, *Spin*, and *Clock*. Each is applied to some facet-edge pair and returns a new facet-edge pair.

Fnext is defined by $a' = aFnext$ where $e_{a'} = e_a$ and facet $f_{a'}$ follows f_a in the facet-ring $\mathscr{F}_a$. The sense of rotation in the rings of $\mathscr{F}_{a'}$ is assigned so that $\mathscr{F}_{a'} = \mathscr{F}_a$ and $a'Org = aOrg$. In particular, a and a' have the same spin.

Enext is defined by $a' = aEnext$ where $f_{a'} = f_a$ and edge $e_{a'}$ follows edge e_a in the edge-ring $\mathscr{E}_a$. The rings of a' are directed so that $\mathscr{E}_{a'} = \mathscr{E}_a$ and $a'Ppos = aPpos$. Observe that a and a' necessarily have the same orientation.

Spin is defined by $a' = aSpin$ where a' and a are different versions of the same facet-edge pair—that is, $e_{a'} = e_a$ and $f_{a'} = f_a$—for which the sense of rotation of $\mathscr{F}_{a'}$ is opposite that of $\mathscr{F}_a$, and the senses of rotation of $\mathscr{E}_{a'}$ and $\mathscr{E}_a$ are the same.

Clock is defined by $a' = aClock$ where a' and a are different versions of the same facet-edge pair, for which the senses of rotation of $\mathscr{E}_{a'}$ and $\mathscr{F}_{a'}$ are opposite those of $\mathscr{E}_a$ and $\mathscr{F}_a$, respectively.

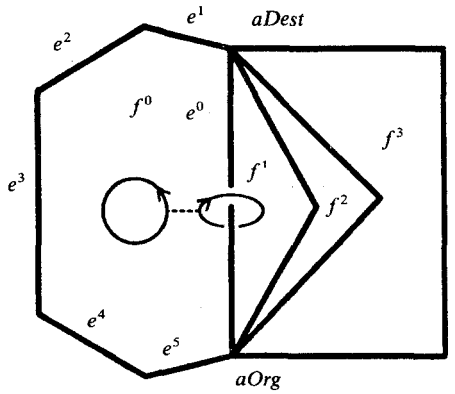

Fig. 1. We call this a handcuff diagram. It pictures a region of some complex. The "handcuff" represents facet-edge pair *a*. The placement and direction of its circular loop indicates the clocked facet component f_a, and its elliptical loop the space-oriented edge component e_a. In this example, $\mathcal{F}_a = (f^0 \cdots f^3)$ and $\mathcal{E}_a = (e^0 \cdots e^5)$, where $f_a = f^0$ and $e_a = e^0$. Polyhedron *aPpos* lies above the page and contains facets f^0 and f^1, while *aPneg* lies behind the page and contains f^0 and f^3.

Figure 2 illustrates these various traversal functions. Traversal functions *Spin* and *Clock* can be viewed as follows. Let *a* be a facet-edge pair with orientation and spin. The effect of *Spin* is to reverse spin. This reverses the sense of rotation in the facet-ring. The effect of *Clock* is to reverse orientation. This reverses the direction of edge e_a, as well as the sense of rotation in both facet- and edge-rings. Each of the four versions of a facet-edge pair has unique orientation and spin. Its orientation and spin are used as handles to manipulate the sense of rotation in its two rings.

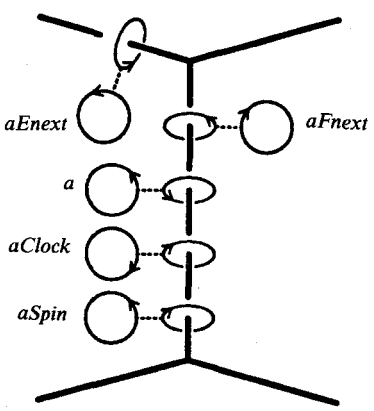

Fig. 2. This handcuff diagram illustrates the four traversal functions *Clock*, *Spin*, *Fnext*, and *Enext*. The region pictured is a winged-edge, consisting of five edges and (part of) two facets (to the left and right of the vertically drawn line). We assume these two facets to belong to a common polyhedron that lies behind the plane of the page, this being *aPpos* in this figure.

Traversal functions *Fnext* and *Enext* enable us to move from facet-edge pair to adjacent facet-edge pair. Facet-edge pairs a and b are said to be adjacent if either:

(i) f_a and f_b are adjacent in the facet-ring $\mathscr{F}_a$ $(=\mathscr{F}_b)$,
(ii) e_a and e_b are adjacent in the edge-ring $\mathscr{E}_a$ $(=\mathscr{E}_b)$.

The facet-edge pair a is adjacent to the facet-edge pairs a*Fnext*, a*Fnext*$^{-1}$, a*Enext*, and a*Enext*$^{-1}$.

The following relations hold among the traversal functions:

(A1) $aSpin^2 = a$.
(A2) $aClock^2 = a$.
(A3) $aSpinClock = aClockSpin$.
(A4) $aFnext^{-1} = aClockFnextClock$.
(A5) $aFnext^{-1} = aSpinFnextSpin$.
(A6) $aEnext^{-1} = aClockEnextClock$.
(A7) $aEnext^{-1} = aClockSpinEnextClockSpin$.
(A8) $aClockFnext^i \neq a$ for any i.
(A9) $aSpinEnext^i \neq a$ for any i.
(A10) $aClockEnext^i \neq a$ for any i.
(A11) $aSpinFnext^i \neq a$ for any i.
(A12) $a \in C$ iff $aFnext \in C$.
(A13) $a \in C$ iff $aClock \in C$.
(A14) $a \in C$ iff $aSpin \in C$.

3.2. Space-Duality. The traversal function *Sdual* is applied to a facet-edge pair a of complex C, and returns a second facet-edge pair *aSdual* belonging to C^*. The edge component of *aSdual* is $e_{aSdual} = f_a$, and its facet component is $f_{aSdual} = e_a$. In order to define the particular version of *aSdual*—that is, the sense of rotation of its two rings—we first extend the notion of space-duality to facet- and edge-rings.

Given edge-ring $\mathscr{E}_a = (e_a^0 \cdots e_a^{n-1})$ of C, its space-dual is the facet-ring $(\mathscr{E}_a)^* = (e_a^0 \cdots e_a^{n-1^*})$ of C^*. The space-dual of a facet-ring is similarly defined. The rings of *aSdual* are then assigned a sense of rotation such that

$$\mathscr{E}_{aSdual} = (\mathscr{F}_a)^* \quad \text{and} \quad \mathscr{F}_{aSdual} = (\mathscr{E}_a)^*.$$

The relation between a and *aSdual* can be grasped by imaging the two facet-edge pairs superimposed, edge e_a piercing facet f_{aSdual} orthogonally, and facet f_a pierced by edge e_{aSdual} orthogonally. Edge e_{aSdual} is directed from *aPneg* toward *aPpos*. Facet-ring $\mathscr{F}_{aSdual}$ moves from *aOrg* toward *aDest*, so *aDest* is the space-dual of *aSdualPpos*. Facet-edge pairs a and *aSdual* necessarily have the same orientation. This is depicted in Figure 3.

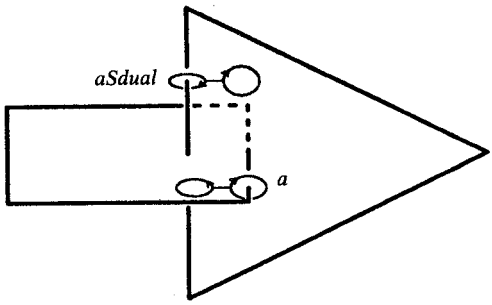

Fig. 3. This diagram depicts the relation between facet-edge pairs a and $aSdual$. Facet f_a is a square protruding from the page, and so appears foreshortened.

The following relations hold between *Sdual* and the other traversal functions:

(A15) $aSdual^2 = a$.
(A16) $aClockSdual = aSdualClock$.
(A17) $aSpinSdual = aSdualClockSpin$.
(A18) $aFnext = aSdualEnextSdual$.
(A19) $aEnext = aSdualFnextSdual$.
(A20) $a \in C$ iff $aSdual \in C^*$.

Relation (A16) indicates that changing the sense of rotation in both $\mathcal{F}_a$ and $\mathcal{E}_a$ corresponds to changing the sense of rotation in both $\mathcal{F}_{aSdual}$ and $\mathcal{E}_{aSdual}$. Relation (A17) implies that a change in the sense of rotation of $\mathcal{F}_a$ in C corresponds to a change in the sense of rotation of $\mathcal{E}_{aSdual}$ in C^*. Relation (A18) indicates that the rings $\mathcal{F}_a$ and $\mathcal{E}_{aSdual}$ rotate in the same direction. Relation (A19) defines *Enext* in terms of *Fnext* and *Sdual*, so maintaining the facet-rings in both C and C^*, as well as the correspondence between each cell and its space-dual, is sufficient to maintain the edge-rings in both complexes. The facet-edge data structure uses this fact.

4. The Facet-Edge Data Structure. In this section we present the facet-edge data structure. The scheme for representing a polyhedral subdivision and simultaneously its dual is described. The implementation of the facet-edge functions is presented. For the implementation, it is useful to first introduce a new operator called *Srot*.

4.1. Traversal Function Srot. Operator *Srot* (for Space ROTation) is defined by

$$aSrot = aSdualSpin = aSpinClockSdual.$$

Facet-edge pair *aSrot* is called the rotated version of a. Its edge component is directed from *aPneg* toward *aPpos*, and its spin is opposite the spin of a. Observe

the following relations:

$$\begin{aligned} aSrot^2 &= aSdualSpinSpinClockSdual \\ &= aClock, \\ aSrot^3 &= aClockSrot \\ &= aClockSpinClockSdual \\ &= aSpinSdual \\ &= aSrot^{-1}, \\ aSrot^4 &= a. \end{aligned}$$

Srot plays a significant role in the facet-edge structure. Given facet-edge pair a, the two facet-edge pairs of $\mathcal{F}_a$ adjacent to a are of the form $aSrot^0Fnext$ and $aSrot^2Fnext = aFnext^{-1}Clock$. The two facet-edge pairs of $\mathcal{E}_a$ adjacent to a are of the form $aSrot^1FnextSrot = aEnext^{-1}Clock$ and $aSrot^3FnextSrot = aEnext$.

By associating with a the facet-edge pairs $aSrot^rFnext$ for $r = 0, 1, 2, 3$, we can obtain the four facet-edge pairs adjacent to a (assuming we can move easily between a subdivision and its dual). By storing the $aSrot^rFnext$ in a *facet-edge node* associated with a, the four facet-edge pairs adjacent to a are available in constant time.

4.2. Implementation of the Traversal Functions. Polyhedral subdivision C (and simultaneously subdivision C^*) are represented by the facet-edge data structure. The facet-edge pairs (with spin and direction) comprising C and C^* may be partitioned into groups of eight. Where facet-edge pair a is an arbitrary member of some group, the facet-edge pairs of the group are of the form $aSrot^rSpin^s$ where $r \in \{0, 1, 2, 3\}$ and $s \in \{0, 1\}$. An arbitrary member $\bar{a}$ of each group is designated the canonical representative of the group.

A group is represented by a *facet-edge node* n, an array consisting of elements $n[0]$ through $n[3]$. Element $n[r]$ corresponds to the facet-edge pair $\bar{a}Srot^r$. The facet-edge pair $\bar{a}Srot^rSpin^s$ is represented by the triplet (n, r, s), where $r \in \{0, 1, 2, 3\}$ and $s \in \{0, 1\}$. Such a triplet is called a *facet-edge reference.* The facet-edge reference can be viewed as a pointer to the array element $n[r]$, plus a bit s indicating whether *Spin* is to be applied to the facet-edge pair $\bar{a}Srot^r$ which corresponds to $n[r]$.

Each element $n[r]$ of the facet-edge node contains two fields, *data* and *next.* Field *data* is used to hold application-dependent information corresponding to $\bar{a}Srot^r$ such as geometry, and need not concern us. Field *next* contains a facet-edge reference to $\bar{a}Srot^rFnext$. Given arbitrary facet-edge reference (n, r, s), the functions *Srot*, *Spin*, and *Fnext* are given by the formulas

$$\begin{aligned} (n, r, s)Srot &= (n, r+1+2s, s), \\ (n, r, s)Spin &= (n, r, s+1), \\ (n, r, s)Fnext &= (n[r+2s].next)Srot^{2s}Spin^s, \end{aligned}$$

where the r and s components are computed modulo 4 and 2, respectively.

Observe that in the third formula, we have

$$(n, r, 0)\,Fnext = (n[r].next)$$

and

$$\begin{aligned}(n, r, 1)\,Fnext &= (n[r+2].next)\,Srot^2\,Spin\\ &= (n, r+2, 0)\,FnextSrot^2\,Spin\\ &= (n, r+2, 0)\,FnextClockSpin\\ &= (n, r, 0)\,ClockFnextClockSpin\\ &= (n, r, 0)\,Fnext^{-1}\,Spin.\end{aligned}$$

The last equation indicates that moving forward around a facet-ring is the same as moving backward around the same ring given with a reverse sense of rotation.

The remaining traversal functions are defined in terms of *Srot*, *Spin*, and *Fnext* as follows. Observe that all traversal queries take constant time.

$$\begin{aligned}(n, r, s)\,Clock &= (n, r, s)\,Srot^2\\ &= (n, r+1+2s, s)\,Srot\\ &= (n, (r+1+2s)+1+2s, s)\\ &= (n, r+2, s).\\ (n, r, s)\,Fnext^{-1} &= (n, r, s)\,ClockFnextClock.\\ (n, r, s)\,Sdual &= (n, r, s)\,SrotSpin\\ &= (n, r+1+2s, s+1).\\ (n, r, s)\,Enext &= (n, r, s)\,SdualFnextSdual.\\ (n, r, s)\,Enext^{-1} &= (n, r, s)\,ClockEnextClock.\end{aligned}$$

5. Primitive Construction Operators. In this section we present the primitive construction operators *make_facet_edge*, *splice_facets*, and *splice_edges*. The first operator obtains and initializes a new facet-edge node, and returns a facet-edge reference to one of the eight facet-edge pairs represented by the node. Operators *splice_facets* and *splice_edges* are used to modify the facet- and edge-rings of a complex.

Two caveats accompany these operators. First, no class of complexes is closed with respect to these operators: their use does not guarantee that complexes are produced. Operator *make_facet_edge* does not, in fact, create a complex at all—edge e_a of the facet-edge pair a it returns is incident to facet f_a and to no other facet, and so does not belong to the boundary of a polyhedron. Furthermore, misuse of *splice_facets* or *splice_edges* can lead to nonsense objects. Second, these primitives are not easy to use in constructing complexes of complexity. The reader need not be vexed. In Section 6 we define higher-level operators in terms of these

primitives which make the task of construction quite feasible (if not also easy). Because of this, we make no attempt to protect the user who prefers to create objects from these most primitive operators.

To help describe these primitives, we introduce some notation for manipulating rings. The notation allows us to describe the manipulation of complexes in terms of the essentially one-dimensional manipulation of rings. Let $\Phi = (a_1 \cdots a_m)$ and $\Phi' = (a_{m+1} \cdots a_n)$ be two rings with all a_i distinct. Then $concat(\Phi, \Phi')$ represents the ring

$$concat(\Phi, \Phi') = (a_1 \cdots a_n).$$

The operation $split(\Phi, a_p)$ represents the pair of rings

$$split(\Phi, a_p) = ((a_1 \cdots a_{p-1}), (a_p \cdots a_m)),$$

where $0 < p \leq m+1$. Operations *first* and *second* are used to access the first and second rings of the pair $split(\Phi, a_p)$, respectively. Furthermore, rings Φ and Φ' are equivalent, denoted $\Phi \equiv \Phi'$, if they represent the same cycle of elements—that is, where $|\Phi| = |\Phi'| = n$, there exists an integer j such that, for each $1 \leq i \leq n$, the ith element of Φ is identical to the $(i+j)$th element of Φ', modulo n. By convention, $\mathscr{F}_a$ denotes the ring $\mathscr{F}_a = (aFnext^0 aFnext^{-1} \cdots aFnext^{n-1})$ where $|\mathscr{F}_a| = n$. $\mathscr{E}_a$ is similarly defined.

5.1. Make_facet_edge. Construction primitive *make_facet_edge* returns a facet-edge reference to a new (canonical) facet-edge $\bar{a}$. Relations (A) of Section 3 hold among the eight facet-edge pairs $\bar{a}Srot^r Spin^s$, where $r \in \{0, 1, 2, 3\}$ and $s \in \{0, 1\}$.

Primitive *make_facet_edge* is implemented as follows. Operation *make_facet_edge*( ) obtains a free node n. Element $n[r]$ is assigned the facet_edge reference $(n, r, 0)$ for $r \in \{0, 1, 2, 3\}$. In [La] it is shown that relations (A) hold over the eight versions of the facet-edge pair $\bar{a}$.

5.2. Splice_facets. The operation *splice_facets*(a, b) takes as arguments two facet-edge pairs, and returns no value. The operation affects the facet-rings $\mathscr{F}_a$ and $\mathscr{F}_b$ as follows:

(a) if the two rings are distinct, it combines them into one ring;
(b) if the rings are identical, it breaks the ring into two distinct rings.

The arguments determine where the facet-rings are to be cut and joined. In rings $\mathscr{F}_a$ and $\mathscr{F}_b$, the cuts occur immediately after facets f_a and f_b, respectively. If the two rings are distinct, the distinct edges e_a and e_b are coalesced into one edge, and the two rings combined at the cuts. If the two rings are identical, the edge e_a $(= e_b)$ is cleaved lengthwise into two new edges, and each serves as pivot to one of the two new facet-rings resulting from the cuts. The operator is illustrated in Figure 4.

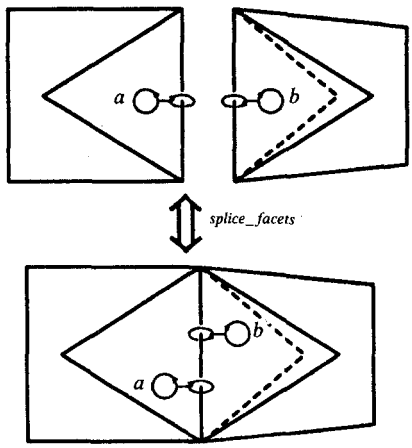

Fig. 4. This diagram illustrates the effect operator *splice_facets* has upon facet rings.

The operation can be viewed as a way of replacing certain facet-rings with others.

splice_facets(a, b)
{
if $(\mathcal{F}_a \equiv \mathcal{F}_b)$
replace $\mathcal{F}_a$ by the two rings *split*$(\mathcal{F}_{aFnext}, bFnext)$;
else
replace $\mathcal{F}_a$ and $\mathcal{F}_b$ by *concat*$(\mathcal{F}_{aFnext}, \mathcal{F}_{bFnext})$;
}

Operation *splice_facets*(a, b) is accomplished by interchanging the value of *aFnext* with *bFnext*. The operation affects the *Fnext* relation in complexes C_a and C_b (where by C_a we mean the complex to which cells f_a and e_a belong). Let $\overline{Fnext}$ denote the *Fnext* relation immediately after the operation is performed. Where $\alpha = aFnextClock$ and $\beta = bFnextClock$, relations *Fnext* (immediately before the operation) and $\overline{Fnext}$ are related as follows:

(B1) $a\overline{Fnext} = bFnext$.
(B2) $b\overline{Fnext} = aFnext$.
(B3) $\alpha\overline{Fnext} = \beta Fnext$.
(B4) $\beta\overline{Fnext} = \alpha Fnext$.
(B5) $aClockSpin\overline{Fnext} = \beta Spin$.
(B6) $bClockSpin\overline{Fnext} = \alpha Spin$.
(B7) $\alpha ClockSpin\overline{Fnext} = bSpin$.
(B8) $\beta ClockSpin\overline{Fnext} = aSpin$.
(B9) $\gamma\overline{Fnext} = \gamma Fnext$ for all other facet-edge pairs γ.

The implementation for *splice_facets*(a, b) is quite simple. In the following we assume the last four assignments are performed simultaneously; in practice, some temporary variables would be used when swapping values.

```
splice_facets(a, b)
  {
  assume (n, r, s) = a;
  assume (n', r', s') = b;
  (ν, ρ, σ) ← aFnextClock;
  (ν', ρ', σ') ← bFnextClock;
  n[r+2s].next ← bFnextClock^s Spin^s;
  n'[r'+2s'].next ← aFnextClock^s' Spin^s';
  ν[ρ+2σ].next ← bClock^(σ+1) Spin^σ;
  ν'[ρ'+2σ'].next ← aClock^(σ'+1) Spin^σ';
  }
```

Correctness of implementation is shown by proving that the procedure *splice_facets* results in the (B) relations. This is done formally in [La].

5.3. Splice_edges. The operation *splice_edges*(a, b) takes as arguments two facet-edge pairs, and returns no value. The operation modifies the edge-rings $\mathscr{E}_a$ and $\mathscr{E}_b$ as follows:

(a) if the two rings are distinct, it combines them into one ring;
(b) if the rings are equivalent, it breaks the ring into two rings.

As with *splice_facets*, the arguments to *splice_edges* determine where edge-rings are to be cut and joined. In rings $\mathscr{E}_a$ and $\mathscr{E}_b$, cuts occur immediately after edges e_a and e_b, respectively. Figure 5 illustrates the effect of *splice_edges*.

The operator can be seen as replacing certain edge-rings with others.

```
splice_edges(a, b)
  {
  if (E_a ≡ E_b)
    replace E_a by the two rings split(E_aEnext, bEnext);
  else
    replace E_a and E_b by concat(E_aEnext, E_bEnext);
  }
```

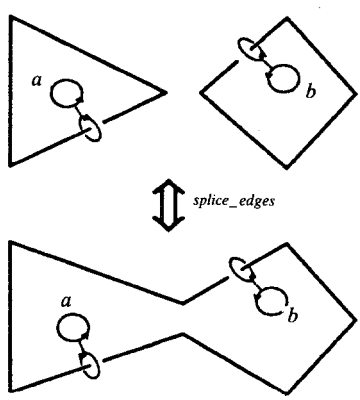

Fig. 5. This diagram illustrates the effect operator *splice-edges* has upon edge-rings.

Operation *splice_edges*(a, b) affects the *Enext* relation in complexes C_a and C_b (or equivalently, the *Fnext* relation in C_a^* and C_b^*). Let $\overline{Enext}$ denote the *Enext* relation immediately after the operation. Where $\alpha = aEnextClock$ and $\beta = bEnextClock$, *Enext* (before the operation) and $\overline{Enext}$ are related as follows:

(C1) $a\overline{Enext} = bEnext$.
(C2) $b\overline{Enext} = aEnext$.
(C3) $\alpha\overline{Enext} = \beta Enext$.
(C4) $\beta\overline{Enext} = \alpha Enext$.
(C5) $aSpin\overline{Enext} = \beta ClockSpin$.
(C6) $bSpin\overline{Enext} = \alpha ClockSpin$.
(C7) $\alpha Spin\overline{Enext} = bClockSpin$.
(C8) $\beta Spin\overline{Enext} = aClockSpin$.
(C9) $\gamma\overline{Enext} = \gamma Enext$ for all other facet-edge pairs γ.

As we might expect, an edge-ring of one complex can be modified by operating on the corresponding facet-ring of the dual complex. Indeed, *splice_edges* is implemented in terms of *splice_facets*.

```
splice_edges(a, b)
  {
  splice_facets(aSdual, bSdual);
  }
```

To show correctness of this implementation, it suffices to show that the (C) relations are satisfied by *splice_facets*$(aSdual, bSdual)$. Below we show this for (C2, 4, 6, 8); the remaining relations (except for (C9)) are shown similarly, whereas (C9) holds since *splice_facets* affects no more than four facet-edge mode fields.

Operation *splice_facets*$(aSdual, bSdual)$ establishes the following relations, where $\overline{Fnext}$ denotes the *Fnext* relation immediately after the operation:

(B2′) $bSdual\overline{Fnext} = aSdualFnext$.
(B4′) $bSdualFnextClock\overline{Fnext} = aSdualClock$.
(B6′) $bSdualSpinClock\overline{Fnext} = aSdualFnextClockSpin$.
(B8′) $bSdualFnextSpin\overline{Fnext} = aSdualSpin$.

Relation (Bi′), which derives from relation (Bi) of Section 5.2, is then used in showing (Ci) below.

(C2)
$$\begin{aligned} b\overline{Enext} &= bSdual\overline{Fnext}Sdual \\ &= aSdualFnextSdual \\ &= aEnext. \end{aligned}$$

(C4)
$$\begin{aligned} \beta\overline{Enext} &= bEnextClock\overline{Enext} \\ &= bSdualFnextSdualClockSdual\overline{Fnext}Sdual \\ &= bSdualFnextClock\overline{Fnext}Sdual \\ &= aSdualClockSdual \\ &= aClock \\ &= \alpha Enext. \end{aligned}$$

(C6) $$\begin{aligned} bSpin\overline{Enext} &= bSpinSdual\overline{Fnext}Sdual \\ &= bSdualSpinClock\overline{Fnext}Sdual \\ &= aSdualFnextClockSpinSdual \\ &= aSdualFnextSdualSpin \\ &= aEnextSpin \\ &= \alpha ClockSpin. \end{aligned}$$

(C8) $$\begin{aligned} \beta Spin\overline{Enext} &= bEnextClockSpin\overline{Enext} \\ &= bSdualFnextSdualClockSpinSdual\overline{Fnext}Sdual \\ &= bSdualFnextSpin\overline{Fnext}Sdual \\ &= aSdualSpinSdual \\ &= aClockSpin. \end{aligned}$$

6. Manipulating Individual Polyhedra. It is worthwhile to be able to manipulate the individual polyhedra of a complex that is represented by the facet-edge structure. First, we would permit the traversal of the combinatorial boundary of an arbitrary polyhedron, while ignoring the rest of the complex that contains the polyhedron. Such traversal should be accomplished using functions appropriate for moving around a two-dimensional subdivision—the facet-edge functions are too general to be appropriate. In this section we reduce each of a small set of such functions to the facet-edge functions. The set of functions we choose to work with are the edge (traversal) functions of the quad-edge structure.

Second, we would permit the construction of the facet-edge representation of a single (connected) polyhedron. The construction should be accomplished using operators appropriate to the task—the use of the facet-edge operators would be overkill. In this section we reduce the edge (construction) operators of the quad-edge structure to the facet-edge operators. Using these edge operators implemented in terms of the facet-edge operators, polyhedra can be built that are represented by the facet-edge structure. Each such polyhedron can be regarded as a primitive complex.

Third, we would permit two polyhedra to be glued together along a polygon of each. The complexes to which each belongs would thus be combined, or modified if they are one and the same. With the *meld* operator primitive complexes built with the edge operators can be combined to form nontrivial complexes.

6.1. Traversing the Boundary of a Polyhedron. In this subsection we concern ourselves with traversal in the combinatorial boundary ∂p of an arbitrary polyhedron p belonging to complex C or to C^*. We first briefly present as background the elements of the quad-edge structure that we will need. The presentation is intended as a reminder to the reader, and at places applies only to traversal of

orientable surfaces; the reader is encouraged to read [GS] if not already familiar with this seminal work. We then present an edge representation scheme whereby an edge $e \in \partial p$, viewed as a cell of the two-dimensional complex ∂p, can be represented in terms of the facet-edge structure that models C and C^*. The edge representation scheme is then used as a basis for describing each edge function in terms of its affect upon the facet-edge structure.

6.1.1. The Edge Functions. Where p is a polyhedron of complex C, the 2-complex $Q = \partial p$ is a subdivision of the sphere. Given edge $e \in Q$, the orientation (with respect to Q) and direction of e can be chosen independently, so there are four oriented, directed versions of e. We write $\hat{e}_p$ to denote any such version, or simply $\hat{e}$ when polyhedron p is known by context. Edge $e \in C$ is said to *underlie* edge $\hat{e}$ in Q.

The direction of $\hat{e} \in Q$ determines the edge's vertex of origin ($\hat{e}Org$) and vertex of destination ($\hat{e}Dest$), in the natural way. In addition, the orientation and direction of $\hat{e}$ together determine the edge's left polygon ($\hat{e}Left$) and right polygon ($\hat{e}Right$). Specifically, where Q is coherently oriented under the orientation of $\hat{e}$, $\hat{e}Left$ is that polygon of Q incident to $\hat{e}$ whose orientation agrees with $\hat{e}$'s direction; $\hat{e}Right$ is the other polygon of Q incident to $\hat{e}$. The orientation of the cells $\hat{e}Org$, $\hat{e}Dest$, $\hat{e}Left$, and $\hat{e}Right$ are taken by definition to agree with that of $\hat{e}$.

There are three primitive edge functions—*Flip*, *Sym*, and *Onext*—in terms of which the remaining edge functions of the quad-edge structure (except for *Dual*) are defined. The flipped version $\hat{e}Flip$ of edge $\hat{e}$ has orientation opposite that of $\hat{e}$, but the two edges have the same direction. The symmetric version $\hat{e}Sym$ of $\hat{e}$ has direction opposite that of $\hat{e}$, but the edges have the same orientation. Furthermore, considering the cycle of edges (in Q) incident to $\hat{e}Org$, we define $\hat{e}Onext$ to be that edge that immediately follows $\hat{e}$ in that cycle, where the direction of the cycle is induced by the orientation of $\hat{e}$.

The dual of $Q = \partial p$ is defined to be a 2-complex Q^* obtained from Q by interchanging vertices and polygons, and which preserves incidence relations. The dual of edge $\hat{e} \in Q$ is an oriented and directed edge $\hat{e}Dual \in Q^*$, for which:

(D1) $\hat{e}DualDual = \hat{e}$
(D2) $\hat{e}DualSym = \hat{e}SymDual$.
(D3) $\hat{e}DualFlip = \hat{e}FlipSymDual$.
(D4) $\hat{e}DualOnext = \hat{e}OnextSymDual$.

(This definition of dual is equivalent to that of [GS] where $\hat{e}Lnext$, the edge following $\hat{e}$ in $\hat{e}Left$, is defined by $\hat{e}Lnext = \hat{e}SymOnext^{-1}$.) *Dual* is extended to vertices and polygons by defining $(\hat{e}Org)Dual = \hat{e}DualLeft$ and $(\hat{e}Left)Dual = \hat{e}DualOrg$. *Dual* establishes a correspondence between the vertices (edges, polygons) of Q, and the polygons (edges, vertices) of Q^*.

Since $\hat{e}$ and $\hat{e}Dual$ have opposite orientation, it is convenient to define a rotated version $\hat{e}Rot$ of $\hat{e}$, given by

$$\hat{e}Rot = \hat{e}FlipDual = \hat{e}DualFlipSym.$$

Edge $\hat{e}Rot$ is the dual of $\hat{e}$, directed from $\hat{e}Right$ to $\hat{e}Left$, and oriented so that moving around $\hat{e}Right$ corresponds to moving around $\hat{e}RotOrg$.

To describe later how a subdivision may be modified, it is convenient to define $\hat{e}Org$ to be the ring of edges in Q incident to $\hat{e}$'s vertex of origin. More formally, $\hat{e}Org$ is the cycle under $Onext$ of e. Polygon $\hat{e}Left$ is defined in terms of the ring of edges in Q^* incident to the vertex dual to $\hat{e}$'s left polygon—$\hat{e}Left$ is defined to be the ring $\hat{e}OnextRotOrg$.

6.1.2. The Edge Representation Scheme. Let p be a polyhedron of complex C or C^*, where C and C^* are represented by the facet-edge structure. Where primal edge $\hat{e}_p \in \partial p$ and $d \in \{0, 1\}$, edge $\hat{e}_pDual^d$ is represented by the pair $\langle a, d\rangle$, called an *edge reference.* The first component is a reference to the facet-edge pair a, determined by the following:

(i) Edge e_a underlies $\hat{e}_p$.
(ii) Facet f_a underlies $\hat{e}_pLeft$.
(iii) The orientation of a coincides with the orientation of $\hat{e}_p$.
(iv) $aPpos = p$.

The second component d, called a *duality bit,* has value 0 (1) iff the edge being represented is primal (dual), and is identical to the exponent d of $\hat{e}_pDual^d$. The scheme is depicted in Figure 6.

Each edge $\hat{e}_pDual^d$, as p ranges over the polyhedra of $C \cup C^*$ and $d \in \{0, 1\}$, is uniquely and unambiguously represented by an edge reference. Given edge $\hat{e}_pDual^d$, conditions (i) and (ii) uniquely determine the components of a, then (iii) determines the orientation of a, then (iv) the orientation of a. On the other hand, given $\langle a, d\rangle$, consider first the edge that $\langle a, 0\rangle$ represents. Conditions (i) and (iv) together determine $\hat{e}_p$, then (iii) determines the orientation of $\hat{e}_p$, then (ii) the direction of $\hat{e}_p$. Finally, since $\hat{e}_p$ is unambiguously represented by $\langle a, 0\rangle$ and edge $\hat{e}_pDual$ is well defined, $\langle a, d\rangle$ unambiguously represents edge $\hat{e}_pDual^d$.

6.1.3. Implementation of the Edge Functions. One purpose of the edge representation scheme is to enable the traversal of the boundary of polyhedron p using the underlying facet-edge structure. Each edge function can be described

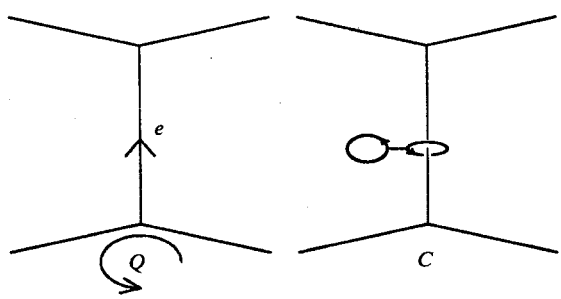

Fig. 6. This diagram illustrates the edge-reference scheme. The winged-edge corresponds to a region of $Q = \partial p$, where p lies behind the page. The facet-edge pair a depicted by the handcuff is such that edge $\hat{e}$ is given by $\langle a, 0\rangle$, and $\hat{e}Dual$ by $\langle a, 1\rangle$.

in terms of how it affects an edge reference. More precisely, for edge function O_p, there exists a sequence of facet-edge functions Op' for which

$$\langle a, d\rangle Op = \langle aOp', d'\rangle.$$

The following characterizes each edge operator in this fashion:

(E1) $\langle a, 0\rangle Flip = \langle aFnextSpin, 0\rangle$.
(E2) $\langle a, 0\rangle Sym = \langle aFnextClock, 0\rangle$.
(E3) $\langle a, 0\rangle Onext = \langle aEnext^{-1}FnextClock, 0\rangle$.
(E4) $\langle a, d\rangle Dual = \langle a, 1-d\rangle$.
(E5) $\langle a, 1\rangle Flip = \langle aSpinClock, 1\rangle$.
(E6) $\langle a, 1\rangle Sym = \langle aFnextClock, 1\rangle$.
(E7) $\langle a, 1\rangle Onext = \langle aEnext^{-1}, 1\rangle$.

The correctness of this scheme can be verified by showing that the edge operators so characterized possess the properties stated in Section 2.3 of [GS]. For instance, where $\hat{e}$ is represented by $\langle a, 0\rangle$, we have

$$\begin{aligned}\hat{e}Flip^2 &= \langle a, 0\rangle FlipFlip \\ &= \langle aFnextSpinFnextSpin, 0\rangle \\ &= \langle a, 0\rangle \\ &= \hat{e}.\end{aligned}$$

The derivation of (E1-7) is straightforward. Figure 7 pictorially motivates equations (E1-3). Equation (E4) follows from the edge representation scheme.

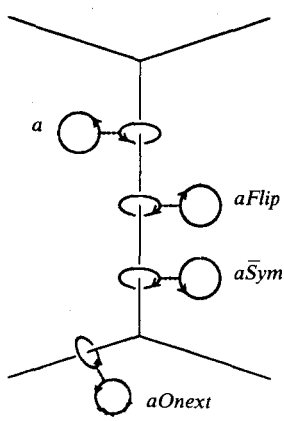

Fig. 7. This diagram illustrates the use of facet-edge pairs to represent directed, clocked edges in the boundary Q of a polyhedron. Each handcuff stands for facet-edge pair a where $\langle a, 0\rangle$ represents the directed, clocked edge $\hat{a}Op$ with which the handcuff is labeled.

Equations (E5–7) follow from the relation developed so far in this section. For instance, (E5) is derived as follows (see [La] for the remaining derivations):

$$\begin{aligned}
\langle a, 1\rangle \mathit{Flip} &= \langle a, 0\rangle \mathit{DualFlip} && \text{(E4)}\\
&= \langle a, 0\rangle \mathit{FlipSymDual} && \text{(D3)}\\
&= \langle a\mathit{FnextSpin}, 0\rangle \mathit{SymDual} && \text{(E1)}\\
&= \langle a\mathit{FnextSpinFnextClock}, 0\rangle \mathit{Dual} && \text{(E2)}\\
&= \langle a\mathit{SpinClock}, 0\rangle \mathit{Dual} && \text{(A1, A5)}\\
&= \langle a\mathit{SpinClock}, 1\rangle. && \text{(E4)}
\end{aligned}$$

6.2. Constructing a Polyhedron. A polyhedron can be characterized by its combinatorial boundary, this being a two-dimensional subdivision of the sphere. A facet-edge structure representing a single polyhedron is most easily created and modified by manipulating the polyhedron's boundary. We choose the edge operators of the quad-edge structure as the means of performing these manipulations. These edge operators handle (in particular) the class of *open subdivisions* of the sphere, of which the (closed) subdivisions may be regarded as a special case. The construction of a polyhedron involves using the edge (construction) operators to build open subdivisions incrementally until one is produced which coincides with the boundary of the target polyhedron. In this section we describe the effect each edge operator has upon the facet-edge structure by giving an implementation of each operator in terms of the facet-edge operators.

During the construction of polyhedron p, *open* subdivision Q is maintained under the edge-representation scheme. It is designated the *primal* subdivision, and its cells, and only its cells, may eventually belong to the target ∂p.

6.2.1. Open Subdivisions. An *open k-cell* (for our purposes) is an open subspace of the sphere S^2 homeomorphic to R^k. An *open complex* S of S^2 is a finite collection of open cells of S such that:

(i) the cells of S are pairwise disjoint,
(ii) for each cell $c \in S$, bd c is the union of elements of S, and
(iii) if $c, d \in S$ and cl $c \cap$ cl $d \neq \emptyset$, then cl $c \cap$ cl d is the union of elements of S.

Here cl c denotes the closure of cell c. An open complex whose union is S^2 is an *open subdivision* (of the sphere).

Let S be an open subdivision such that for each cell $c \in S$, the closure cl c is a (closed) cell. S corresponds to a (closed) subdivision of the sphere, obtained by replacing each cell $c \in S$ by its closure cl c. To build a polyhedron, the edge operators are applied successively to construct new elementary open subdivisions, and to combine and modify existing open subdivisions. The process proceeds until an open subdivision is produced that corresponds to the boundary of the target polyhedron.

6.2.2. Elementary Open Subdivisions of a Sphere. There are two elementary subdivisions of the sphere. The first consists of a single edge $\hat{e}$ that is not a loop, and is denoted S_e (subscript e stands for "edge"). Where $\hat{e} \in S_e$ is some arbitrary (but fixed) oriented and directed edge, we have $\hat{e}Org \neq \hat{e}Dest$ and $\hat{e}Left = \hat{e}Right$. The following properties hold in S_e.

(F1) $\hat{e}Onext = \hat{e}$.
(F2) $\hat{e}SymOnext = \hat{e}Sym$.
(F3) $\hat{e}FlipOnext = \hat{e}Flip$.
(F4) $\hat{e}FlipSymOnext = \hat{e}FlipSym$.

The other elementary open subdivision of the sphere consists of a single edge $\hat{e}'$ that is a loop, and is denoted S_l (subscript l stands for "loop"). S_l is dual to open subdivision S_e; there exists a version of $\hat{e}'$ for which edges $\hat{e}'$ and $\hat{e}Sdual$ represent identical open subdivisions. Since $\hat{e}'$ is a loop, we have $\hat{e}'Org = \hat{e}'Dest$ and $\hat{e}'Left \neq \hat{e}'Right$. Writing $\hat{e}Dual$ for $\hat{e}'$, the following properties hold in S_l:

(F5) $\hat{e}DualOnext = \hat{e}DualSym$.
(F6) $\hat{e}DualSymOnext = \hat{e}Dual$.
(F7) $\hat{e}DualFlipOnext = \hat{e}DualFlipSym$.
(F8) $\hat{e}DualFlipSymOnext = \hat{e}DualFlip$.

The operator *make_edge* builds a data structure representing both S_e and S_l, and returns an edge reference to one version of S_e's (nonloop) edge. Open subdivision S_e is primal. Its implementation is given as follows:

```
make_edge( )
  {
  a ← make_facet_edge( );
  b ← make_facet_edge( );
  splice_facets(a, b);
  splice_edges(a, bClock);
  return(⟨a, 0⟩);
  }
```

The operation $\hat{e} \leftarrow$ *make_edge*( ) obtains two new facet edge nodes. In one of the nodes it designates a facet-edge pair a for which $\langle a, 0\rangle$ represents $\hat{e}$, while in the other node it designates facet-edge pair b for which $\langle bSpin, 0\rangle$ represents $\hat{e}Flip$. Operation *splice_facets*(a, b) results in

(i) $aFnext = b$,
(ii) $bFnext = a$,

while operation *splice_edges*$(a, bClock)$ results in

(iii) $aEnext = bClock$,
(iv) $aEnext^{-1} = bClock$.

Edge $\hat{e}$ under the edge-representation scheme is depicted in Figure 8.

Relations (i)–(iv) ensure that $\hat{e} \leftarrow$ *make_edge*( ) does indeed build a facet-edge structure representing both S_e and S_l, and that edge $\hat{e}$ is represented by $\langle a, 0\rangle$.

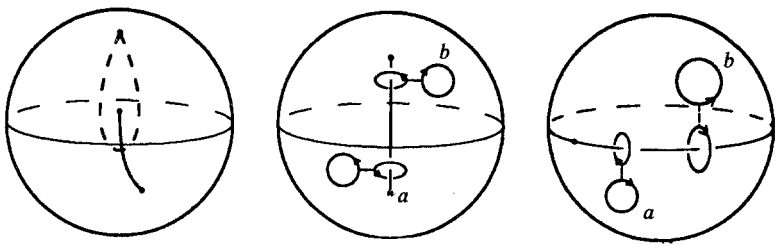

Fig. 8. This diagram depicts the open subdivisions under the edge-reference scheme. The left figure depicts edges $\hat{e} \in S_e$ and $\hat{e}' \in S_l$, superimposed on the same sphere to suggest how they are related. The center figure depicts the facet-edge representation for $\hat{e}$ where S_e is primal (constructed by *make_edge*), while the right figure depicts the facet-edge representation for loop $\hat{e}'$ where S_l is primal (constructed by *make_loop*).

This is verified by showing that relations (F1–8) are satisfied. For instance, (F1) is shown as follows (for other proofs refer to [La]):

$$\begin{aligned} \hat{e}Onext &= \langle a, 0\rangle Onext \\ &= \langle aEnext^{-1}FnextClock, 0\rangle \\ &= \langle bClockFnextClock, 0\rangle \qquad \text{(iv)} \\ &= \langle bFnext^{-1}, 0\rangle \\ &= \langle a, 0\rangle \qquad \text{(i)} \\ &= \hat{e}. \end{aligned}$$

The operator *make_loop* also builds a facet-edge structure representing both S_e and S_l, but *it* returns an edge-reference to that version of S_l's loop that corresponds to $\hat{e}Dual$ (where $\hat{e} \leftarrow make\_edge(\,)$). Open subdivision S_l is primal. Its implementation is given as follows:

```
make_loop( )
  {
  a ← make_facet_edge( );
  b ← make_facet_edge( );
  splice_facets(a, b);
  return(⟨a, 0⟩);
  }
```

The operation $\hat{e}Dual \leftarrow make\_loop(\,)$ obtains two new facet-edge nodes. In one node it designates facet-edge pair a for which $\langle a, 0\rangle$ represents edge $\hat{e}Dual$, and in the other node a facet-edge pair b for which $\langle bSpin, 0\rangle$ represents $\hat{e}DualFlip$. Operation $splice\_facets(a, b)$ results in

(i) $aFnext = b$,
(ii) $bFnext = a$,

while the absence of a call to *splice_edges* results in

(iii) $aEnext = a$,
(iv) $bEnext = b$.

Edge $\hat{e}Dual$ under the edge-representation scheme is depicted in Figure 8.

6.2.3. Modifying Open Subdivisions. The operator *splice* is used to modify open subdivisions. The operation $splice(\hat{a}, \hat{b})$ takes as arguments two edges $\hat{a}$ and $\hat{b}$, and returns no value. The operation affects the two rings $\hat{a}Org$ and $\hat{b}Org$ and, independently, the two rings $\hat{a}Left$ and $\hat{b}Left$. In each case, if the two rings are distinct, *splice* combines them into one ring; and if the two rings are identical, *splice* breaks it into two distinct rings. The arguments $\hat{a}$ and $\hat{b}$ determine where the rings will be cut and joined. For rings $\hat{a}Org$ and $\hat{b}Org$, the cuts occur immediately after $\hat{a}$ and $\hat{b}$; for rings $\hat{a}Left$ and $\hat{b}Left$, the cuts occur immediately after $\hat{a}OnextRot$ and $\hat{b}OnextRot$.

Operation $splice(\hat{a}, \hat{b})$ is performed by interchanging the values of $\hat{a}Onext$ with $\hat{b}Onext$, and $\hat{\alpha}Onext$ with $\hat{\beta}Onext$, where $\hat{\alpha} = \hat{a}OnextRot$ and $\hat{\beta} = \hat{b}OnextRot$. More formally, where $\overline{Onext}$ denotes the *Onext* relation immediately after the operation, $splice(\hat{a}, \hat{b})$ establishes the following relations between $\overline{Onext}$ and *Onext*:

(G1) $\hat{a}\overline{Onext} = \hat{b}Onext$.
(G2) $\hat{b}\,\overline{Onext} = \hat{a}Onext$.
(G3) $\hat{\alpha}\overline{Onext} = \hat{\beta}Onext$.
(G4) $\hat{\beta}\,\overline{Onext} = \hat{\alpha}Onext$.
(G5) $\hat{\gamma}\overline{Onext} = \hat{\gamma}Onext$ for all other edges $\hat{\gamma} \in Q \cup Q^*$.

Operation $splice(\hat{a}, \hat{b})$ is implemented in terms of the facet-edge operator *splice_edges* as follows:

```
splice(⟨a, d⟩, ⟨b, d⟩)
  {
  if (d = 0)
    splice_edges(aClockSpin, bClockSpin);
  else
    splice_edges(aEnext^-1, bEnext^-1);
  }
```

The duality bit d of the two arguments to *splice* are assumed to be identical—$splice(\hat{a}, \hat{b})$ is defined only if $\hat{a}$ and $\hat{b}$ are both primal, or both dual.

To show the correctness of the implementation, let $\overline{Onext}$ ($\overline{Enext}$) denote the *Onext* (*Enext*) relation immediately after $splice(\hat{a}, \hat{b})$, given primal edges $\hat{a}$ and $\hat{b}$ represented by $\langle a, 0\rangle$ and $\langle b, 0\rangle$. Operation *splice_edges*(*aClockSpin*, *bClockSpin*) establishes:

(i) $a\overline{Enext}^1 = bEnext^{-1}$.
(ii) $b\overline{Enext}^1 = aEnext^{-1}$.
(iii) $aEnext^{-1}ClockSpin\overline{Enext}^1 = bClockSpin$.
(iv) $bEnext^{-1}ClockSpin\overline{Enext}^1 = aClockSpin$.

Relations (i)–(iv), which follow from the (C) relations of Section 5.3, are used to show that values have been correctly swapped. To show (G1), we have

$$\begin{aligned}\hat{a}\overline{Onext} &= \langle a, 0\rangle \overline{Onext} \\ &= \langle a\overline{Enext}^{1}\, FnextClock, 0\rangle \\ &= \langle bEnext^{-1}\, FnextClock, 0\rangle \qquad \text{(i)} \\ &= \langle b, 0\rangle Onext \\ &= \hat{b}Onext.\end{aligned}$$

Similarly (ii) is used to show relation (F2). To show relation (F3), we have

$$\begin{aligned}\hat{\alpha}\overline{Onext} &= \langle a, 0\rangle OnextRot\overline{Onext} \\ &= \langle aEnext^{-1}\, FnextClockFnextSpin, 1\rangle \overline{Onext} \\ &= \langle aEnext^{-1}\, ClockSpin, 1\rangle \overline{Onext} \\ &= \langle aEnext^{-1}\, ClockSpin\overline{Enext}^{-1}, 1\rangle \\ &= \langle bClockSpin, 1\rangle \qquad \text{(iii)} \\ &= \langle bEnext^{-1}\, ClockSpinEnext^{-1}, 1\rangle \\ &= \langle bEnext^{-1}\, ClockSpin, 0\rangle DualOnext \\ &= \langle bEnext^{-1}\, FnextClockFnextSpin, 0\rangle DualOnext \\ &= \langle b, 0\rangle OnextFlipDualOnext \\ &= \hat{b}OnextRotOnext \\ &= \hat{\beta}Onext.\end{aligned}$$

Similarly (iv) is used to show relation (H4). Notice that

$$splice\_edges(aClockSpin, bClockSpin)$$

only modifies facet-rings of the complex dual to the complex to which $\hat{a}$ and $\hat{b}$ belong. Since each occurrence of *Fnext* in the derivations above apply only to facet-rings of the complex to which $\hat{a}$ and $\hat{b}$ belong, we have been free to assume (in the derivations) that *Fnext* has not been changed by *splice_edges*; no $\overline{Fnext}$ is necessary in the derivations.

We have shown correctness of an implementation for *splice* when its arguments are primal edges. Assume now that *splice* is passed two dual edges $\hat{a}$ and $\hat{b}$. To show correctness of implementation in this case, we note that operations $splice(\hat{a}, \hat{b})$ and $splice(\hat{\alpha}, \hat{\beta})$ are equivalent, where $\hat{\alpha} = \hat{a}OnextRot$ and $\hat{\beta} = \hat{b}OnextRot$. Since edges $\hat{\alpha}$ and $\hat{\beta}$ are primal, it is sufficient to show that $splice\_edges(\hat{\alpha}ClockSpin, \hat{\beta}ClockSpin)$—which implements $splice(\hat{\alpha}, \hat{\beta})$—and $splice\_edges(\hat{a}Enext^{-r}, \hat{b}Enext^{-1})$ are equivalent; this is done in [La].

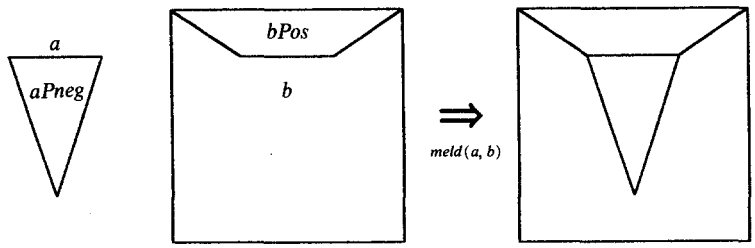

Fig. 9. This figure depicts a two-dimensional analogue of the effect of *meld*. Edges of the figure correspond to facets, and polygons to polyhedra. Note that $C_a \neq C_b$ in the figure, but this need not be the case.

6.3. Meld. The operator *meld* is used to glue a complex C_a to a second complex C_b. With its use one melds an n-sided polygon $f_a \in \partial C_a$ to an n-sided polygon $f_b \in C_b$, thereby locating complex C_a in the polyhedron *bPneg* of C_b. More formally, it establishes the topological relations for

$$\mathcal{U}(C_a) \subset bPneg \quad \text{and} \quad \mathcal{U}(C_a) \cap \mathcal{U}(\partial bPneg) = f_b.$$

A two-dimensional analogue of the situation is depicted in Figure 9.

Let a and b be facet-edge pairs, and that their spin and orientation are chosen so that C_a lies in *bPneg* and C_b lies in *aPpos*. Where $a_i = aEnext^i$ and $b_i = bEnext^i$, $meld(a, b)$ identifies polygon f_a with f_b, and edge e_{a_i} with e_{b_i} for $0 \leq i \leq n-1$. The operation first coalesces distinct edge rings $\mathcal{E}_a$ and $\mathcal{E}_b$, forming a "pillow" consisting of the edges of $\mathcal{E}_a$ $(=\mathcal{E}_b)$ and the polygons f_a and f_b, and then removes polygon f_a from the complex. The two polyhedra that end up incident to f_b are *bPpos* and *aPneg*.

The boundary of polyhedron *aPneg* is slightly changed to produce polyhedron p—facet f_a is replaced by f_b. In addition, polyhedra *aPpos* and *bPneg* are combined to form a new polyhedron q, the effect of locating C_a inside *bPneg*. We have

$$facets\_of(p) = facets\_of(aPneg) - f_a + f_b$$

and

$$facets\_of(q) = facets\_of(bPneg) \cup facets\_of(aPpos) - f_a - f_b.$$

To build the facet-rings of the "pillow" formed by coalescing $\mathcal{E}_a$ and $\mathcal{E}_b$, facet-rings of C_a and C_b are combined as follows:

for $i = 0, \ldots, n-1$
  if $\mathcal{F}_{a_i} \not\equiv \mathcal{F}_{b_i}$
  replace $\mathcal{F}_{a_i}$ and $\mathcal{F}_{b_i}$ by $concat(\mathcal{F}_{a_i Fnext}, \mathcal{F}_{b_i Fnext})$;

Facet f_a is removed as follows:

for $i = 0, \ldots, n-1$
  replace $\mathcal{F}_{a_i}$ by $first(split(\mathcal{F}_{a_i Fnext}, b_i Fnext))$;

```
meld(a, b)
  {
  firsta ← a;
  do {
    if (F_a ≠ F_b)
      splice_facets(a, bFnext^-1);
    splice_facets(a, aFnext^-1);
    a ← aEnext
    b ← bEnext
    }until a = firsta;
  }
```

Fig. 10. Procedure *meld.*

Operator *meld,* given in Figure 10, is implemented by a single loop in which the construction of the "pillow" and removal of f_a are interleaved. The necessary facet-ring manipulations are done with *splice_facets.* Note that the facet-edge pairs $a_i Srot^r Spin^s$ (where $0 \le i \le n-1$, $r \in \{0, 1, 2, 3\}$, and $s \in \{0, 1\}$) are effectively deleted from the data structure; the n facet-edge nodes representing these could be garbage-collected.

7. Decomposing a Polyhedron. The process of partitioning a polyhedron into simpler constituent polyhedra is called *decomposition.* One reason for decomposing a polyhedron $\bar{p}$ is that $\bar{p}$ may possess properties that preclude certain algorithms from being applied to it—for instance, it may be nonconvex, or possess cavities or handles. Sometimes the difficulty may be overcome by decomposing $\bar{p}$ into more amenable pieces, and then applying the algorithm to these [CD]. Alternatively, $\bar{p}$ may be well behaved but its volume might be too large to allow the efficient solution of equations via finite element methods, and further decomposition may be desired [JB].

There are various strategies for performing decomposition. We concern ourselves with an incremental strategy in which polyhedra are iteratively detached from the original polyhedron until nothing of the original remains. Each simpler piece split off from the original is not subject to further decomposition, and satisfies whatever "simplicity" criteria is required of the algorithm. Such an algorithm maintains a current polyhedron S (initially $\bar{p}$), and a current collection of constituent polyhedra C (initially $\varnothing$). The algorithm iteratively detaches a polyhedron p_i (in iteration i) from S and transfers it to C. The process stops when S represents a null polyhedron—collection C then represents the decomposition of $\bar{p}$. In this section we show how collection C assembled during the course of decomposition can be represented by the facet-edge structure. Each polyhedron detached from S is attached to C by meld operations. For simplicity, we assume that $\bar{p}$ is polyhedral in the sense we have been using the word (that is, having genus zero and no cavities), and that C always consists of zero or more ball complexes.

Wördenweber uses this incremental strategy in [Wö] to decompose a polyhedron into tetrahedra. He makes no attempt to assemble the pieces, but allows

the sequence of operations by which they were detached to represent the resulting decomposition. We refer the reader to [Wö] for a description of how the algorithm selects a tetrahedron to be detached from the current polyhedron S in each iteration. The actual removal of the tetrahedron from S is accomplished by one of the four operators *op0*, *op1*, *op2*, or *op3*. Each *opk* modifies the polyhedron S to reflect the removal of the tetrahedron. The index k of *opk* indicates how the tetrahedron t that *opk* is designed to detach is connected to the rest of S: k is the number of triangular facets that connect t to the rest of S, while $4-k$ is the number of t's exposed facets (that is, belonging to the boundary of S). In modifying S, *opk* removes each of the $4-k$ exposed facets of t, a consequence of detaching t from S. In general, some of these facets correspond to facets of C (in fact, to facets of ∂C); where f is such a facet, $\tilde{f}$ denotes that facet of C to which f corresponds. Facet $\tilde{f}$ will have been created when S was modified in some earlier iteration (say iteration j where $j<i$); the polyhedron p_j transferred to C at that time contains $\tilde{f}$ which was then cleaved from f. The rest of the $4-k$ exposed facets of t belong to the boundary of the original polyhedron $\bar{p}$, and so correspond to no facets of C. Let S' denote S immediately after being modified by *opk*. S' contains k new facets, the "connecting" facets of t; these facets are created by *opk*—again the consequence of deleting t from S—and replace the $4-k$ exposed facets removed by *opk*. Each of the k facets revealed by *opk* corresponds to an exposed facet of the tetrahedron that has been added to C. Figure 11 illustrates the effect each *opk* has upon S.

A decomposition algorithm employing the incremental strategy requires the use of an operator *op*, analogous to Wördenweber's *opk* operators, for transferring a polyhedron p_i from S to C. The operator must build a facet-edge representation for p_i, attach p_i to C (using calls to *meld*), and modify S (to reflect its loss of p_i). The operator's most formidable task is in determining exactly how p_i is to be attached to C—that is, in determining the arguments to each of its calls to *meld*. To guide *op* in attaching p_i to C, each facet $f \in S$ possesses a *link pointer* *link*(f) which references that facet $f \in \partial C$ to which f corresponds. The facets of S that belong to the boundary of the original polyhedron $\bar{p}$ do not correspond

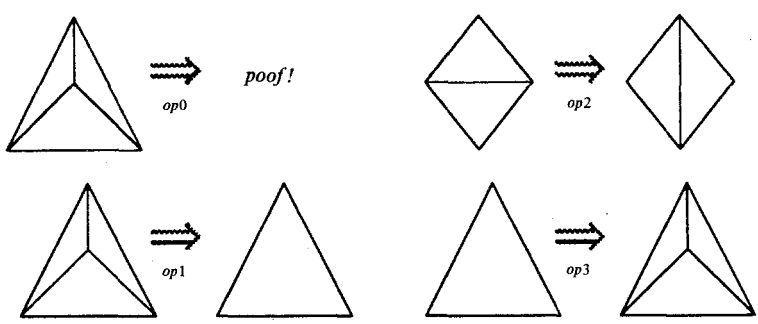

Fig. 11. This figure demonstrates how each *opk* locally modifies S to produce S'. Each drawing depicts a patch on the boundary of S or S'.

to any facet of C, and so have null link pointers. To elaborate, in iteration i, *op* performs the following steps in succession:

(i) constructs a facet-edge representation for polyhedron p_i, to be transferred from S to C,
(ii) attaches p_i to C, thereby forming C' (to serve as C in the next iteration),
(iii) modifies S to reflect its loss of p_i, thereby forming S' (to serve as S in the next iteration), and
(iv) updates the link pointers of S'.

We do not elaborate on step (i); it is performed using the quad-edge operators, whose implementation in terms of the facet-edge operators was given in Section 6. Assuming S is suitably represented—for concreteness we assume by the quad-edge structure—step (iii) also need not be treated. Presumably the description of p_i handed to *op* is adequate for *op* to perform steps (i) and (iii). Steps (ii) and (iv) do require elaboration. Henceforth denoting by p the polyhedron p_i constructed in step (i), we discuss in turn how we ascertain which facets of p are to be melded to C, how the link pointer is used to guide each meld operation, and how the link pointers are updated in S' to serve later iterations.

Consider the relation between p and S. That patch of p to be glued to C coincides with subcomplex $S_p = \bigcup \{star f | f \in S$ is removed by *op*$\}$. (Here *star* f is the complex consisting of the faces of cell f; in this case since f is a facet, it consists of f and the vertices and edges that bound f.) Subcomplex S_p is generally a patch of S, homeomorphic to a closed disk. (In the final iteration, however, S itself is transferred to C, in which case $S_p = S$.) We denote by $\varphi(c)$ that cell of p that coincides with cell $c \in S_p$. The mapping $\varphi : S_p \to p$ is an isomorphism, not generally onto. Consider next the relation between p and S'. That patch of p that lies in ∂C(after *op* has attached p to C) coincides with subcomplex $S'_p = \bigcup \{star f | f \in S'$ is created by *op*$\}$. (At the last iteration, however, $S' = S'_p = \varnothing$). We denote by $\varphi'(c)$ that cell of p that coincides with cell $c \in S'_p$; the isomorphism φ': $S'_p \to p$ is not onto. The patches $\varphi(S_p)$ and $\varphi'(S'_p)$ cover polyhedron p. Their intersection $\varphi(S_p) \cap \varphi'(S'_p)$ is a vertex-edge cycle in p, called the silhouette of p. These notions are depicted in Figure 12.

To attach p to C, for each facet $f \in S_p$, facet $\varphi(f) \in p$ is melded to facet $\tilde{f} \in C$. Each facet of S_p is obtained by treating the dual 2-complex S_p^* as a graph, and performing a search in S_p^*. Each vertex visited corresponds to a facet of S_p. The silhouette of p is used to restrict the search to S_p^*, prohibiting it from passing into the rest of S^*. Specifically, the search algorithm considers two vertices adjacent iff the edge that connects them is not dual to a silhouette edge of p. Pointer *link*(f) consists of two fields *edge* and *pair*, whose contents are as follows:

link(f).*edge*: an edge reference to $e \in S$ such that $eLeft = f$.
link(f).*pair*: an edge reference to *edge* $e' \in \partial C$ such that:
(a) $e'Left = \tilde{f}$,
(b) $e'Lnext^i = \tilde{e}Lnext^i$ for all i, and p^∞.

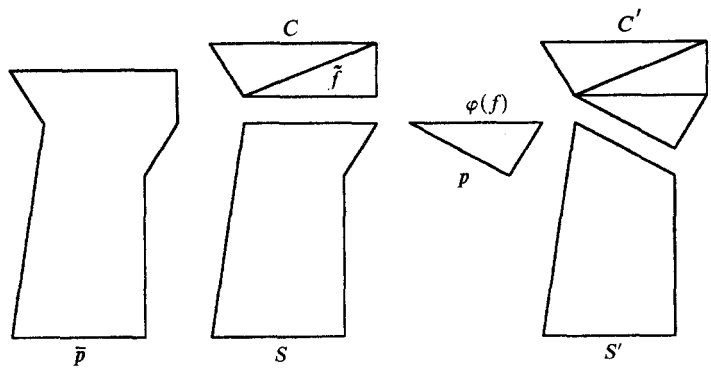

Fig. 12. This figure depicts a two-dimensional analogue of the effect of *op*. Each edge of the figure corresponds to a facet, and each polygon to a polyhedron.

When facet $f \in S_p$ is visited, facets $\varphi(f)$ and $\tilde{f}$ are melded by the following operation:

```
if link(f) ≠ ∅ {
  e ← link(f).edge;
  a ← that facet-edge pair a for which ⟨a, 0⟩ is φ(e);
  b ← link(f).pair;
  meld(a, b);
  }
```

Edge $\varphi(e) \in p$, required in the above block of code, is obtained by performing an identical graph search in p^*, coincident with the search in S_p^*.

Having attached p to C and modified S to produce S', the link pointers of S' must be updated. This involves setting the link pointer of each facet created by *op* (that is the facets of S'_p); the link fields of the other facets of S' are still correct. Much as before, we perform a graph search in S'^*_p and a coinciding search in p^*, using the silhouette of p to limit both searches. When we visit a vertex of S'^*_p, dual say to facet $f \in S'_p$, $link(f)$ is set by the following:

```
let e be an edge for which eLeft = f;
link(f).edge ← e;
link(f).pair ← a where φ(e) is represented by ⟨a, 0⟩;
```

The isomorphisms φ and φ' are each computed on the fly by performing identical searches in two distinct graphs. Each pair of searches must start at coinciding cells for each isomorphism to be correctly computed. To do this, we select some silhouette edge e—since e belongs to both S_p and S'_p, it can be used to compute the starting point for both pairs of searches. Let $e \in S$ be oriented and directed so the $eLeft \in S_p$, and let $\varphi(e) \in p$ be oriented and directed so that $\varphi(e)Org = \varphi(eOrg)$ and $\varphi(e)Left = \varphi(eLeft)$. The searches in S_p^* and p^* then begin at vertices *eLeftDual* and $\varphi(e)$*LeftDual*, respectively. To compute φ', we note that coinciding cells of S'_p and p have orientations that *disagree*: facets f and $\varphi(f)$ *appear* to have the same orientation when viewed for instance from a point beyond f but beneath $\varphi(f)$, say from the interior of a convex p. To determine

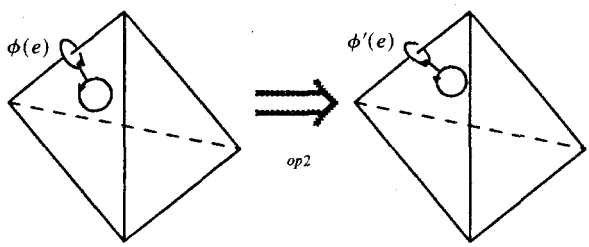

Fig. 13. This diagram depicts a tetrahedron t transferred by Wördenweber's *op2*. The tetrahedron is attached to C along the two facets behind the page, while the two facets in front of the page occur in ∂C.

the starting points for the searches in $S_p'^*$ and p^*, let edge e be oriented and directed as above. Facet $eLeft \in S$ is replaced by $eLeft \in S'$. The edge of S' coinciding with e is then $\varphi(e)Flip$, so the searches of $S_p'^*$ and p^* begin at the vertices *eLeftDual* and $\varphi(e)FlipLeftDual$, respectively. This is illustrated for Wördenweber's *op2* in Figure 13.

8. Incremental Construction of a Three-Dimensional Delaunay Triangulation. We describe how to build the Delaunay triangulation DT(S) of a set S on $n \geq 4$ points (called *sites*) of R^3, in general position. Since the facet-edge structure represents both a complex and its dual, the algorithm also serves to construct the Voronoi diagram of S. The strategy is first to construct some tetrahedron of DT(S)—called a *D-tetrahedron*—to serve as an initial current complex C. C is then grown by iteratively discovering, constructing, and melding a new D-tetrahedron to one or more triangular facets on the boundary of C, until it is known that $C = \mathrm{DT}(S)$. The algorithm is described in [AB], and under geometric inversion that maps S to a set of points S' on a three-dimensional hypersphere in R^4 [Br], corresponds to the gift-wrapping method of [CK] for building the convex hull of S'. The process of finding an initial and subsequent D-tetrahedra is described in [AB], so we describe this only briefly in the next paragraph, before presenting the entire algorithm.

Assume triangle f of DT(S) is on the boundary of complex C, and that the D-tetrahedron t incident to f is known. Operation *find_tetrahedron*(f, t) constructs the other D-tetrahedron t' adjacent to f (if it exists). Let $H_{f,t}$ denote that open half-space determined by aff f which does not contain t. The vertices that define t' are then the vertices of f together with site q, where $q \in H_{f,t}$ is that site for which the sphere determined by q and the vertices of f is of minimal radius. It is shown in [Bh] that the interior of this sphere contains no sites, hence t' is indeed a D-tetrahedron. If $S \cap H_{f,t}$ is empty, then f lies on the convex hull of S and t' does not exist.

An initial D-tetrahedron is found by first finding some triangular facet f on the convex hull of S by the method of [CK]. The D-tetrahedron adjacent to f is discovered using the strategy given above, where candidate sites q range over all sites (except for the three that determine f).

```
delaunay(S)
  {
  t ← an initial D-tetrahedron of S;
  𝓕 ← facet_edges_of(t);
  while (𝓕 ≠ ∅) {
    a ← some element of 𝓕;
    t ← find_tetrahedron(f_a, aPpos);
    if (t does not exist)
      𝓕 ← 𝓕 − {a};
    else {
      for each â ∈ facet_edges_of(t) {
        a ← 𝓕(â);
        if (a ≠ ∅) {
          𝓕 ← 𝓕 − {a};
          a ← align(â, aClock);
          meld(a, â);
          }
        else
          𝓕 ← 𝓕 ∪ {â};
        }
      }
    }
  }
```

Fig. 14. Procedure *delaunay*.

The algorithm *delaunay* of Figure 14 constructs the Delaunay triangulation DT(S) of a finite set of sites S of R^3, in general position. The algorithm initializes the current complex C to contain a D-tetrahedron, then iteratively melds D-tetrahedra to C until, for every facet of C, a D-tetrahedron has been sought on both sides of the facet.

Let F denote the set of facets for which a D-tetrahedron has been sought on exactly one side of the facet. F consists of those facets belonging to the boundary of the current complex C, less those facets that have been determined to lie on the convex hull of S. Dictionary $\mathscr{F}$ contains the triangles of F; more precisely, it contains one facet-edge reference to a for each triangle f_a of F. $\mathscr{F}(a)$ performs a look-up in dictionary $\mathscr{F}$, returning that element of $\mathscr{F}$ whose determining vertices are *aOrg*, *aEnextOrg*, and $aEnext^2Org$ if it exists, or $\emptyset$ if the dictionary contains no such element. A scheme for addressing the elements of $\mathscr{F}$ using (the indices of) the three vertices that determine its elements is easily concocted.

A tetrahedron t is represented by some facet-edge pair a such that $aPpos = t$. The set *facet_edges_of*(t) contains one facet-edge pair $\hat{a}$ for each of the four triangular facets of ∂t, where $\hat{a}Ppos = t$; the set is easily derived by traversal from that facet-edge pair a that represents t. Finally, *align*(a, b) denotes that facet-edge $bEnext^i$ for which $aOrg = bEnext^iOrg$; the algorithm ensures that some such i exists for each use of *align*.

9. Conclusion. The applications presented here but scratch the surface of the data structure's potential uses. Future research includes the development and

rederivation of applications that would markedly benefit from use of the structure. Two examples of these were mentioned in the Introduction: a divide-and-conquer algorithm for constructing three-dimensional Voronoi diagrams, and a scheme for modeling the motion of three-dimensional polyhedra. Future research also includes completely characterizing the class of complexes the data structure can model, and developing sets of construction operators with respect to which various classes of complexes are closed.

Acknowledgment. We would like to thank the anonymous referee for clarifying our exposition.

References

[AB] D. Avis and B. K. Bhattacharya, Algorithms for computing d-dimensional Voronoi diagrams and their duals, in *Advances in Computing Research*, vol. 1, F. P. Preparata, ed., JAI Press, Greenwich, CT, 1983, pp. 159–180.

[Ba] B. G. Baumgart, A polyhedron representation for computer vision, in *1975 National Computer Conference*, AFIPS Conference Proceedings, vol. 44, AFIPS Press, 1976, pp. 589–596.

[Bh] B. K. Bhattacharya, Application of computational geometry to pattern recognition problems, Tech. Rep. 82-3, Simon Fraser University, 1982.

[BHS] I. C. Braid, R. C. Hillyard, and I. A. Stroud, Stepwise construction of polyhedra in geometric modelling, in *Mathematical Methods in Computer Graphics and Design*, K. W. Brodlie, ed., Academic Press, London, 1980, pp. 123–141.

[Br] K. Q. Brown, Voronoi diagrams from convex hulls, *Inform. Process. Lett.*, **9**, 1979, 223–228.

[CD] B. Chazelle and D. P. Dobkin, Detection is easier than computation, *Proc. 12th ACM SIGACT Symposium*, Los Angeles, May 1980, pp. 146–153.

[CK] D. R. Chand and S. S. Kapur, An algorithm for convex polytopes, *J. Assoc. Comput. Mach.*, **17**(1), 1970, 77–86.

[EW] C. M. Eastman and K. Weiler, Geometric modeling using the Euler operators, Research Rep. 78, Institute of Physical Planning, Carnegie-Mellon University, February 1979.

[GS] L. Guibas and J. Stolfi, Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams, *ACM Trans. Graphics*, **4**(2), 1985, 75–123.

[JB] A. Jameson and T. Baker, Improvements to the aircraft Euler method, Paper AIAA-87-0452, AIAA 25th Aerospace Sciences Meeting, 1987.

[La] M. J. Laszlo, A data structure for manipulating three-dimensional subdivisions, Dissertation, Department of Computer Science, Princeton University, August 1987.

[Wö] B. Wördenweber, Volume-triangulation, C.A.D. Group, University of Cambridge, 1980.