CHAPTER 7

Cavities and castability analysis

7.1 Introduction

Feature recognition has been considered an important research area in computer-aided design and computer-aided manufacturing [43, 44, 24]. Informally, features are product's generic shapes or characteristics that are associated with properties, attributes, and engineering knowledge about the product [46, 47]. Manufacturing features are geometric structures of an object, such as holes or depressions, which have engineering meanings related to manufacturing operations. A hole of an object, for example, may have an engineering meaning "drilling" or "assembling site".

In applications such as manufacturing and molecular analysis, geometric structures such as cavities or docking sites are important. In manufacturing, features of a CAD model imply manufacturing information, which facilitates the process of analyzing manufacturability [44, 23].

A small hole or a depression on the boundary of an object, for example, restricts the set of directions for which this object is castable, because the portion of the cast in the hole or in the depression must be removed from the object without breaking the object. Most industrial parts such as engine rooms, telephone bodies, and small parts for car and aircraft have such features. This suggests a new approach to castability analysis: For given part removal directions, instead of examining the whole boundary of an object, we identify such features (holes and depressions) which play key roles in the preliminary decision process. If any such features contradicts the removal directions, we can stop and conclude that these directions are not feasible, or that the object needs additional cast parts. So identifying features not only facilitates the decision process and the automated design of a cast, but also greatly reduces the search space for feasible casting directions. When we search for the set of all feasible casting directions, features can greatly reduce the search space. A hole with the shape of cylinder in an object, for example, reduces the search space to a pair of two opposite directions parallel to the generator of the cylinder. Features, furthermore, can be used to minimize the number of casting parts (called *side cores*).

Based on the definitions of the reflex-free hull and cavities in Chapter 6, in this chapter we consider applications using cavities as a geometric feature in castability analysis. We assume that the cast (mould) consists of two parts and that these parts must be removed in opposite direction without damaging the parts or the object.

We present an algorithm which is useful for casting analysis. The algorithm partitions the faces of $\mathcal{P}$ into disjoint subsets, such that each subset must belong to one of the two mould parts. Furthermore, we prove that the bounding faces of a cavity belong to a single subset. By basing the algorithm on faces, we obtain a finite process. Our algorithm is an effective method to restrict the search space for feasible casting directions. In fact, we conjecture that this algorithm can be extended so that, in the end, for any two distinct subsets, there is a feasible casting direction in which the mould is removed from the corresponding faces in opposite directions.

7.2 Definitions and assumptions

Recall the process of iteratively filling plane-cavities of Section 6.5. We denote this process by $(\mathcal{P}, \sigma_k)$, where σ_k is any sequence of k plane-cavities. A connected component of $\mathcal{P}_k \setminus \mathcal{P}$ is a *cavity* of $(\mathcal{P}, \sigma_k)$. A face f of $\mathcal{P}$ *bounds* a cavity $\mathcal{F}$ of $(\mathcal{P}, \sigma_k)$ if f lies partially or completely on the boundary of $\mathcal{F}$. For the rest, we follow the definitions and the notations in Section 6.2.

We make one assumption for robustness: removal directions parallel to faces of $\mathcal{P}$ are not allowed. Thus, when the two parts of a mould for $\mathcal{P}$ meet along a *parting surface*, this parting surface meets the boundary of $\mathcal{P}$ along a closed curve called the *parting line* that consists of polyhedron edges. This is a practical consideration in mould design as well, since casting imperfections on the object may occur along the parting line, and if the parting line crosses a face then additional treatment and polishing may be required. For further information, see Section 4.1.

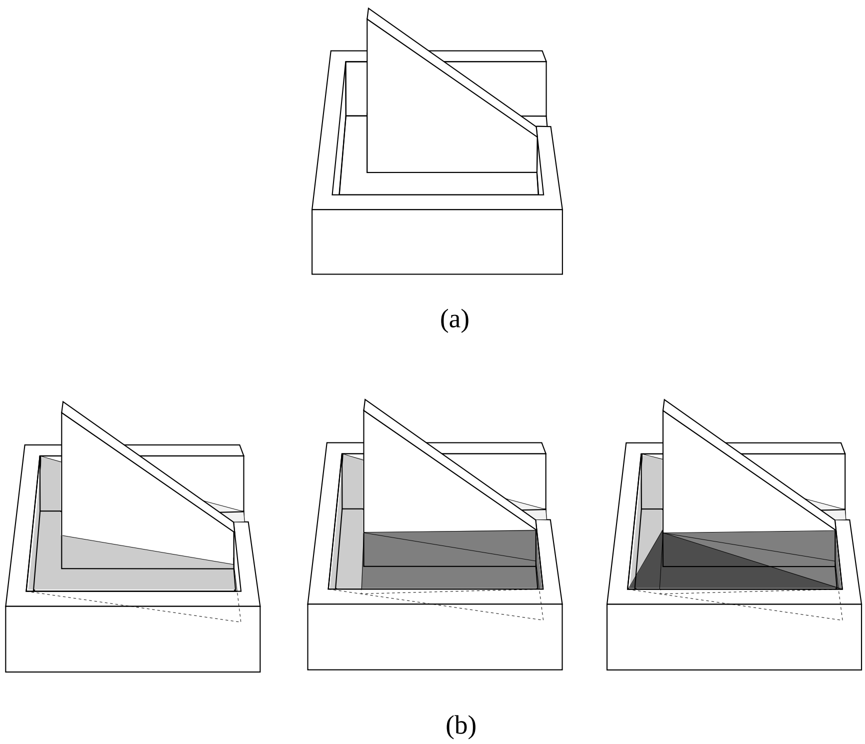

Figure 7.1: Cavities. (a) A container object, and (b) Filling process.

7.3 Theorems for coloring faces

In this section we show how to color the faces of $\mathcal{P}$ such that faces with the same color must appear in the same part of a two-part mould.

This algorithm is based on two geometric observations: Every polyhedron face belongs to one of the two mould parts of any cast, and the two polyhedron faces incident to a reflex edge must belong to the same mould part of any cast.

Lemma 33 *The bounding faces of a cavity of* $(\mathcal{P}, \sigma_k)$ *for any* σ_k *must belong to the same mould part of any cast.*

Proof: Given a cavity $\mathcal{F}$ of $(\mathcal{P}, \sigma_k)$, the boundary of $\mathrm{cl}(\mathcal{F})$ can be divided into two parts: one shared with $\mathcal{P}$ and one not shared with $\mathcal{P}$. We call the part not shared with $\mathcal{P}$ the *lid* of $\mathcal{F}$. It suffices to prove that no two points at the boundary between $\mathcal{P}$ and a cavity can be removed in opposite directions.

Assume that two distinct points on the bounding faces of $\mathcal{F}$ are removed in opposite directions. Inside the cavity the two cast parts meet along a common boundary surface. Then any line ℓ through the interior of the surface and parallel to the removal direction does not intersect $\mathcal{P}$. Since the reflex-free hull of $\mathcal{P}$ does not contain the intersection point between ℓ and the surface, $\mathcal{P}_k$ is not a subset of the reflex-free hull of $\mathcal{P}$, which contradicts Theorem 14. ◻

We need a few definitions in order to describe our algorithm. Let $\mathcal{S}$ be the sphere of directions. Given a face f, we denote by $cone(f)$ the set of directions on $\mathcal{S}$ that have positive projection on the normal of f. Note that if we translate f such that f passes through the center of $\mathcal{S}$, then $cone(f)$ is an open hemisphere defined by a plane through f. Symmetrically, we define $\overline{cone}(f)$ to be the set of directions on $\mathcal{S}$ that have negative projection on the normal of f. We denote the double cone $cone(f) \cup \overline{cone}(f)$ by $dcone(f)$. We generalize the notation to reflex edges. Given a reflex edge e with incident faces f and g, define $cone(e) = cone(f) \cap cone(g)$, $\overline{cone}(e) = \overline{cone}(f) \cap \overline{cone}(g)$, and $dcone(e) = cone(e) \cup \overline{cone}(e)$. Note that $cone(e)$ is the set of removal directions for f and g. (Recall that f and g must belong to the same mould part by assumption.)

Our algorithm works by assigning a color and a positive or negative sign to each face. Faces of the same color (regardless of the sign) form a *color group*. Given a color group G, we define its *cone of directions*, $cone(G)$, to be the common intersection of $cone(f)$ for all positive faces $f \in G$ and $\overline{cone}(g)$ for all negative faces $g \in G$. Symmetrically, $\overline{cone}(G)$ is the set of directions opposite to those in $cone(G)$. The double cone $dcone(G)$ is $cone(G) \cup \overline{cone}(G)$.

The algorithm consists of two phases. Initially, each face is assigned a positive sign and a distinct color, and therefore forms a color group by itself. In the first phase, we repeatedly recolor two groups G_1 and G_2 of faces by one common color if G_1 and G_2 meet along some reflex edge. In the second phase, we repeatedly recolor two groups G_1 and G_2 of faces by one common color if $dcone(G_1) \cap dcone(G_2)$ consists of exactly two connected components. We may also update the signs of faces in $G_1 \cup G_2$ and there are two cases: (1) $cone(G_1) \cap cone(G_2) \neq \emptyset$ and $cone(G_1) \cap \overline{cone}(G_2) = \emptyset$, (2) $cone(G_1) \cap cone(G_2) = \emptyset$ and $cone(G_1) \cap \overline{cone}(G_2) \neq \emptyset$. In case (1), we preserve the signs of all faces in $G_1 \cup G_2$. In case (2), we flip the sign of each face in G_2.

Lemma 34 *For any color group G, all positive faces in G must be removed in a common direction in $cone(G)$, and all negative faces in G in a common direction in $\overline{cone}(G)$, with respect to any mould.*

Proof: We prove this by induction. In the first phase, two faces f and g incident to a reflex edge must be removed in the same direction by our assumption that no face is parallel to a casting direction. In the second phase, suppose we decide to combine two color groups

G_1 and G_2. By induction assumption, all positive (resp. negative) faces in G_1 must be removed in a common direction in $cone(G_1)$ (resp. $\overline{cone}(G_1)$) with respect to any cast, and the same holds for G_2.

Suppose that $cone(G_1) \cap \overline{cone}(G_2) = \emptyset$. Then positive faces in G_1 cannot be removed in a direction in $\overline{cone}(G_2)$ and vice versa. Thus, positive faces in $G_1 \cup G_2$ must be removed in a common direction with respect to any cast, and this set of common directions is clearly $cone(G_1) \cap cone(G_2)$. A symmetric statement holds for negative faces in $G_1 \cup G_2$ and $\overline{cone}(G_1) \cap \overline{cone}(G_2)$.

Suppose that $cone(G_1) \cap cone(G_2) = \emptyset$. Then positive faces in G_1 cannot be removed in a direction in $cone(G_2)$ and vice versa. Thus, positive faces in G_1 and negative faces in G_2 must be removed in a common direction in $cone(G_1) \cap \overline{cone}(G_2)$. Since signs of faces in G_2 are flipped in merging, the lemma is satisfied. A symmetric statement holds for negative faces in G_1 and positive faces in G_2. ◻

Lemma 35 *Let $\mathcal{P}$ be a castable polyhedron. Let f and g be two bounding faces of a cavity of $(\mathcal{P}, \sigma_k)$ for some σ_k. Suppose that f and g belong to two different color groups G_1 and G_2 at some point during the coloring algorithm. If f and g have identical signs, then $cone(G_1) \cap cone(G_2)$ is nonempty. Otherwise, $cone(G_1) \cap \overline{cone}(G_2)$ is nonempty.*

Proof: Let $\mathcal{C}$ be a two-part mould for $\mathcal{P}$. By Lemma 33, f and g belong to the same part of $\mathcal{C}$. If f and g have identical sign, then Lemma 34 implies that faces in $G_1 \cup G_2$ of the same sign belong to the same part of $\mathcal{C}$. Thus, the removal direction of the positive faces in $G_1 \cup G_2$ belongs to $cone(G_1) \cap cone(G_2)$ which must then be nonempty. If f and g have opposite signs, then Lemma 34 implies that positive (resp. negative) faces in G_1 and negative (resp. positive) faces in G_2 belong to the same part of $\mathcal{C}$. Thus, the removal direction of positive faces in G_1 and negative faces in G_2 belongs to $cone(G_1) \cap \overline{cone}(G_2)$ which must then nonempty. ◻

We will prove that the bounding faces of a cavity of $(\mathcal{P}, \sigma_k)$ for any k will receive the same color. The proof is by induction on k. Since a face f will reside in different color groups G during the coloring algorithm, the set of removal directions for f changes as G changes. For ease of exposition, we disregard the group that f belongs to. Instead, we say that different cones of directions D are *associated with* f at different times during the coloring algorithm.

Furthermore, in making the inductive argument, we need to work with plane-cavity with respect to a nearly reflex vertex instead of a reflex vertex. Recall that v is a nearly reflex vertex if v is neither reflex nor flat, and all faces incident to v lie within a closed halfspace whose bounding plane passes through v and lies locally inside $\mathcal{P}$ at v.

Define the *star* of a vertex v, $St(v)$, to be the union of v and the interior of faces and edges incident to v. We use $\overline{St}(v)$ to denote the closure of $St(v)$. The *link* of v, denoted by $Lk(v)$,

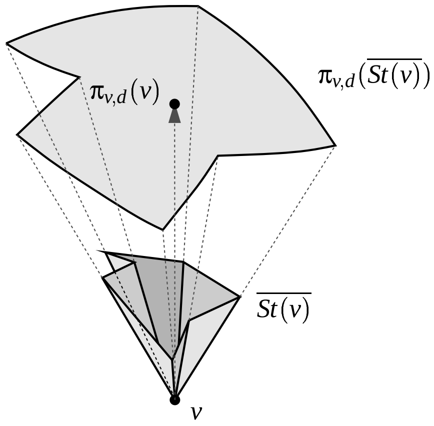

Figure 7.2: The closure, $\overline{St}(v)$, of the star of a reflex vertex v and the spherical polygon $\pi_{v,d}(\overline{St}(v))$ on the sphere, where d is the upward vertial direction.

is defined to be $\overline{St}(v) \setminus St(v)$. Let v be a reflex or nearly reflex vertex. Let d be any feasible casting direction from which v is visible. Note that all faces incident on v are visible from direction d. Put v at the center of the sphere of direction. First, we stretch $\overline{St}(v)$ radially away from v so that the stretched link of v lies on the sphere. This yields a polygon with curved boundary inside the sphere. Second, we project v in direction d onto the sphere. We use $\pi_{v,d}$ to denote the composite mapping from $\overline{St}(v)$ to the spherical polygon on the sphere. Note that $\pi_{v,d}(v)$ is in the kernel of $\pi_{v,d}(\overline{St}(v))$. Thus, $\pi_{v,d}(\overline{St}(v))$ can be triangulated into spherical triangles by drawing great circular arcs from $\pi_{v,d}(v)$ to vertices of $\pi_{v,d}(\overline{St}(v))$. Clearly, $\pi_{v,d}$ maps the circular arcs incident to $\pi_{v,d}(v)$ to edges incident to v, and the spherical triangles to triangles. Also, the angle at a vertex $\pi_{v,d}(x)$ of $\pi_{v,d}(\overline{St}(v))$ is exactly the exterior dihedral angle at the edge vx.

Lemma 36 *Let v be a reflex or nearly reflex vertex. Let va and vb be two reflex edges incident to v. Let $D_{va} \subseteq cone(va)$ and $D_{vb} \subseteq cone(vb)$ be two cones of directions associated with the two faces incident to va and vb, respectively. If $cone(va)$ and $cone(vb)$ lie on the same side of a great circle through $\pi_{v,d}(a)$ and $\pi_{v,d}(b)$, then $D_{va} \cap \overline{D_{vb}}$ is empty.*

Proof: Figure 7.3 illustrates the situation. It follows from the fact that $cone(vx)$ and $\overline{cone}(vx)$ for any reflex edge vx lie on opposite sides of any great circle that does not intersect $cone(vx)$ (such a great circle must pass through $\pi_{v,d}(x)$). $\square$

Lemma 37 *Let v be a reflex or nearly reflex vertex. Let d be a feasible removal direction from which v is visible. Let va, vb, and vx be three reflex edges such that $\pi_{v,d}(vx)$ lies in the smaller angle between $\pi_{v,d}(va)$ and $\pi_{v,d}(vb)$. Suppose that $cone(v\alpha)$ does not contain*

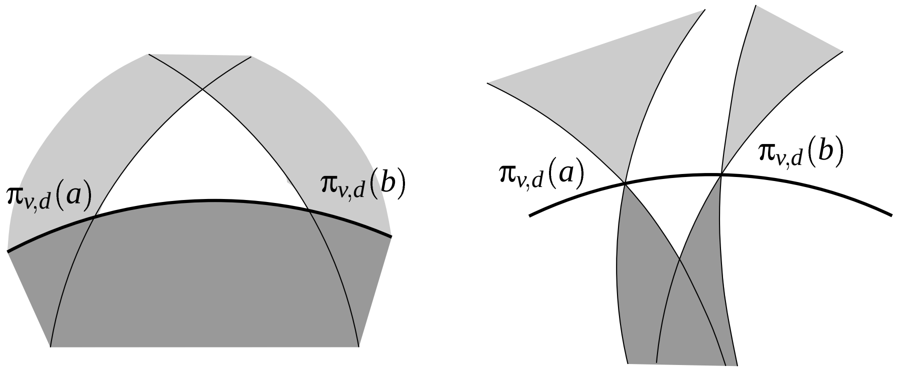

Figure 7.3: The bold curves are parts of great circles. The left picture shows the case where *va* and *vb* are incident to the same face. The right picture shows the case where they are not. The area of darker shade represents $cone(va)$ and $cone(vb)$. The area of lighter shade represents $\overline{cone}(va)$ and $\overline{cone}(vb)$. Since $cone(va)$ and $\overline{cone}(vb)$ lie opposite sides of the great circle, they cannot intersect. The same is true for $\overline{cone}(va)$ and $cone(vb)$.

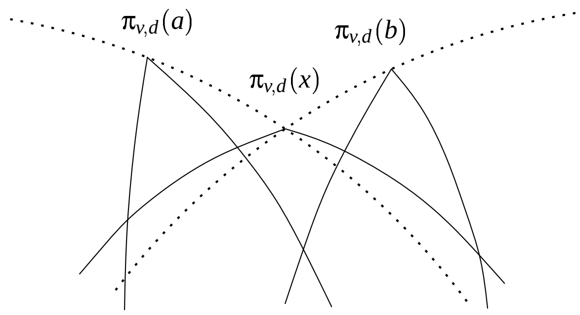

Figure 7.4: The dotted curves are great circles through $\pi_{v,d}(a)$ and $\pi_{v,d}(x)$, and $\pi_{v,d}(b)$ and $\pi_{v,d}(x)$.

$\pi_{v,d}(\beta)$ *for all* $\alpha, \beta \in \{a, b, x\}$. *Then* $D \cap \overline{D_x}$ *is empty, where* $D \subseteq cone(va) \cap cone(vb)$ *is a common cone of directions associated with the faces incident to va and vb, and* D_x *is the cone of directions associated with the faces incident to vx.*

Proof: Since $\pi_{v,d}(v)$ lies in the kernel of $\pi_{v,d}(St(v))$, $cone(vx)$ and $cone(va)$ cannot lie on opposite sides of the great circle through $\pi_{v,d}(a)$ and $\pi_{v,d}(x)$. The same holds for $cone(vx)$ and $cone(vb)$. If $cone(vx)$ and $cone(va)$ lies on the same side of the great circle through $\pi_{v,d}(a)$ and $\pi_{v,d}(x)$, then Lemma 36 implies that $D \cap \overline{D_x}$ is empty. We obtain the same conclusion if $cone(vx)$ and $cone(vb)$ lies on the same side of the great circle through $\pi_{v,d}(b)$ and $\pi_{v,d}(x)$. The remaining possibility is that $cone(vx)$ contains $cone(va) \cap cone(vb)$

which is a superset of D. See Figure 7.4. Thus, D cannot intersect $\overline{cone}(vx)$ which is a superset of $\overline{D_x}$. □

We are ready to prove that the bounding faces of a cavity of $(\mathcal{P}, \sigma_k)$ for any σ_k receive the same color and sign. We will restrict σ_k to simplify the analysis without loss of generality. Specifically, we want to refine σ_k to a sequence of *special* plane-cavities such that the filling of each special plane-cavity corresponds to sweeping a plane from a reflex or nearly reflex vertex until a vertex in the cavity is encountered. Given σ_k, we can refine it as follows: When we fill the plane-cavity $\sigma_k \setminus \sigma_{k-1}$ of $\mathcal{P}_{k-1}$ to produce $\mathcal{P}_k$, the plane-cavity can be decomposed into several steps. Take the plane H_k defining the plane-cavity $\sigma_k \setminus \sigma_{k-1}$. Sweep H_k towards the interior of $\sigma_k \setminus \sigma_{k-1}$ until it hits a vertex w of the plane-cavity. Record the volume swept over as one special plane-cavity. Continue the sweeping to the next vertex in the plane cavity and define another special plane-cavity. During the sweeping, we may need to split at the vertex encountered. In this case, we continue the sweeping of the different parts independently. Figure 7.5 shows an example. We repeat the above sweeping until no vertex in $\sigma_k \setminus \sigma_{k-1}$ remains. Now, we can think of the growing of $\mathcal{P}_{k-1}$ to $\mathcal{P}_k$ as filling the special plane-cavities obtained in reverse order.

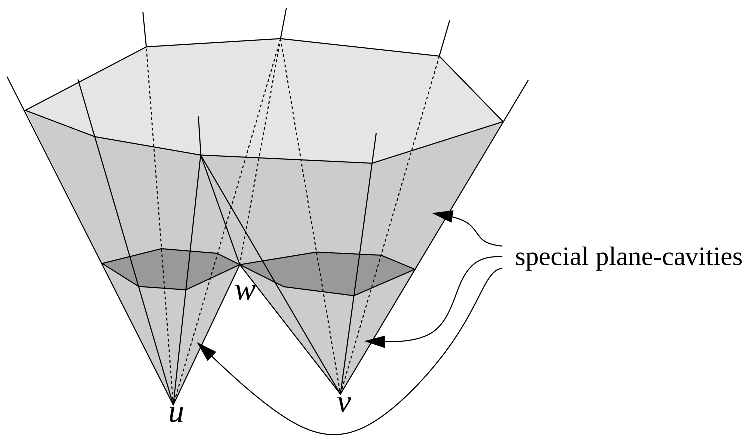

Figure 7.5: The topmost patch of lightest shade bounds a plane-cavity. This plane-cavity can be refined into a sequence of three special plane-cavities. The first two are the volume swept from u and v to w respectively. The third is the volume swept from the nearly reflex vertex w.

Theorem 15 *Let $\mathcal{P}$ be a castable polyhedron. The bounding faces of a cavity of $(\mathcal{P}, \sigma_k)$, for any sequence σ_k of special plane-cavities, receive the same color and sign.*

Proof: We prove this by induction on k. The basis step involves sweeping a plane from a reflex vertex v of $\mathcal{P}$ until a new vertex in the cavity is encountered. We show that the faces incident to v (i.e., faces swept over) will receive the same color and sign. Let d be a feasible removal direction from which v is visible. Without loss of generality, we can assume that each face f is assigned a positive sign if $cone(f)$ contains d.

Each extreme vertex of the spherical convex hull of $\pi_{v,d}(\overline{St}(v))$ is $\pi_{v,d}(x)$ for some reflex edge vx. Consider two neighbouring extreme vertices $\pi_{v,d}(a)$ and $\pi_{v,d}(b)$. Let D_{va} and D_{vb} be the cones of directions for the color groups containing the faces incident to va and vb, respectively. By Lemma 35, $D_{va} \cap D_{vb}$ is nonempty. By Lemma 36, $D_{va} \cap \overline{D_{vb}}$ is empty. Thus, the two color groups containing the faces incident to va and vb are eligible for merging. By applying this argument to every pair of neighbouring extreme vertices of the convex hull, we conclude that the faces incident to vx for all extreme vertices $\pi_{v,d}(x)$ will receive the same color.

Next, we argue that the faces that are incident to reflex edges between neighbouring extreme vertices $\pi_{v,d}(a)$ and $\pi_{v,d}(b)$ will also receive the same color. The portion of $Lk(v)$ between va and vb projects to a bay of $\pi_{v,d}(\overline{St}(v))$. Let D be the common cone of directions associated with the faces incident to va and vb.

First, pick out all reflex edges vx between va and vb such that $cone(vx)$ does not contain $\pi_{v,d}(a)$ and $\pi_{v,d}(b)$, and neither $cone(va)$ nor $cone(vb)$ contains $\pi_{v,d}(x)$. By Lemma 37, $D \cap \overline{D_x}$ is empty where D_x is the cone of directions associated with the faces incident to vx. By Lemma 35, $D \cap D_x$ is nonempty. Thus, the coloring algorithm will eventually assign the same color for faces incident to va, vb, and all such reflex edges vx.

There are four kinds of remaining reflex edges unaccounted for. The first two kinds include reflex edges vx such that $\pi_{v,d}(x)$ lies outside $cone(va) \cup cone(vb)$, and $cone(vx)$ contains either $\pi_{v,d}(a)$ or $\pi_{v,d}(b)$. The other two kinds include reflex edges vx such that $\pi_{v,d}(x)$ lies inside $cone(va) \cup cone(vb)$.

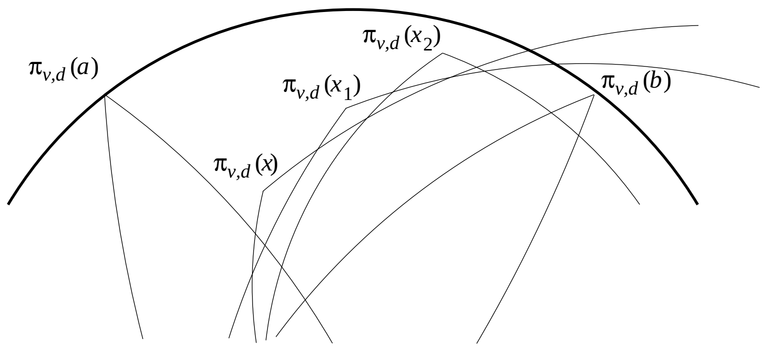

Figure 7.6: $cone(vx_k)$ does not contain $\pi_{v,d}(b)$

Suppose that $\pi_{v,d}(x)$ lies outside $cone(va) \cup cone(vb)$ and $cone(vx)$ contains $\pi_{v,d}(b)$. Then examine the reflex edge vx_1 after vx. Note that $\pi_{v,d}(x_1)$ lies outside $cone(va) \cup cone(vb)$, and $cone(vx_1)$ does not contain $\pi_{v,d}(a)$. If $cone(vx_1)$ also contains $\pi_{v,d}(b)$, then we examine the next reflex edge vx_2 and so on. Thus, we obtain a sequence $vx, vx_1, \cdots, vx_k$ such that $cone(vx_k)$ does not contain $\pi_{v,d}(b)$ in its interior as in Figure 7.6. So the faces incident to vx_k are in the same group for va and vb. By construction, Lemma 37 is applicable to va, vx_{k-1}, and vx_k. Thus, together with Lemma 35, we conclude that the faces incident to vx_{k-1} will receive the same color as those incident to va. Now, repeat the argument for va, vx_{k-2}, and vx_{k-1}, and so on. Eventually, the faces incident to $vx, vx_1, \cdots, vx_{k-1}$ will receive the same color as those incident to va and vb. Similar argument works for the case where $\pi_{v,d}(x)$ lies outside $cone(va) \cup cone(vb)$ and $cone(vx)$ contains $\pi_{v,d}(a)$. This takes care of the first two kinds of remaining reflex edges.

Take a successive pair of reflex edges vy_1 and vy_2 that we have already put in the same group for va and vb. Note that $cone(vy_1)$ does not contain $\pi_{v,d}(y_2)$ and $cone(vy_2)$ does not contain $\pi_{v,d}(y_1)$. We can apply the previous reasoning to color faces incident to each reflex edge vx between vy_1 and vy_2 such that $\pi_{v,d}(x)$ lies outside $cone(vy_1) \cup cone(vy_2)$. By repeating this overall argument, we will assign the same color to faces incident to reflex edges between va and vb as those incident to va and vb.

Now, all the edges between a successive pair of reflex edges vx and vy are convex. Thus, if f is a face incident to such a convex edge, then $cone(f)$ contains $cone(vx) \cap cone(vy)$ and hence $cone(f)$ contains the common cone of directions, say D, for all reflex edges between va and vb. Thus, $D \cap \overline{cone}(f)$ is empty. See Figure 7.7. So if D_f is the cone of directions associated with f, $D \cap \overline{D_f}$ is also empty as $\overline{D_f} \subseteq \overline{cone}(f)$. By Lemma 35, $D \cap D_f$ is nonempty and so f will also receive the same color as those faces incident to va and vb.

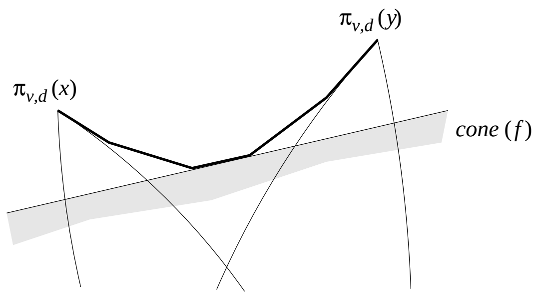

Figure 7.7: The bold polygonal chain is the projection of the link of v between vx and vy.

The above establishes the basis case of the induction. To proceed to the induction step, we need to associate cones of directions to the new faces introduced after filling a special plane-cavity from a reflex or nearly reflex vertex v. These new faces are not original polyhedron faces and we call them artificial faces. Let D be the intersection of cones of directions associated with (artificial or original) faces incident to v. We make D the cone of directions for all artificial faces introduced after filling this special plane-cavity. Then in the induction step, we will sweep a plane from a nearly reflex vertex w to fill another special plane-cavity. We can use the same argument for the basis case to show that the cones of directions associated with the (artificial or original) faces incident to w satisfy the conditions for receiving the same color. Since the cones of directions for artificial faces are derived inductively from intersection of cones of directions for original faces swept in the past, we conclude that the original faces incident to w will receive the same color as other original faces swept in the past. $\square$

7.4 An implementation for coloring faces

The first phase of the algorithm merges color groups at reflex edges. This can be easily done in $O(n)$ time by traversing the boundary of $\mathcal{P}$.

In phase 2, we merge groups G_1 and G_2 whenever $dcone(G_1) \cap dcone(G_2)$ consists of exactly two connected components. There are two cases: (1) $cone(G_1) \cap cone(G_2) \neq \emptyset$ and $cone(G_1) \cap \overline{cone}(G_2) = \emptyset$, (2) $cone(G_1) \cap cone(G_2) = \emptyset$ and $cone(G_1) \cap \overline{cone}(G_2) \neq \emptyset$. Recall that the cone of a group is the intersection of cones for each face in the group. In case (1), the condition could also be stated that there exists a direction with positive projection on all face normals in G_1 and G_2, and there is no direction with positive projection on the normals in G_1 that has negative projection on the normals in G_2. Similarly, the condition in case (2) could also be stated that there exists a direction with positive projection on the normals in G_1 that has negative projection on the normals in G_2, and there is no direction with positive projection on all face normals in G_1 and G_2.

To identify a mergeable pair, we build the arrangement of cones and their complements by building an arrangement of the n great circles that contribute to the current set of cones and their complements. Two spherical convex polygons A and B representing cones have one of four relationships: either their boundaries intersect, A is inside B, B is inside A, or they are disjoint. For a given cone A, all boundary intersections can be detected by walking the boundary of A in the arrangement. All cones including A can be found while building the arrangement by determining which cones include any chosen vertex of the boundary of A. Once all cones including cones are known, then the reverse relationship is also known, and the disjoint pairs are those that remain. If there are pairs that can be merged, at least one pair will be identified after $O(n^2)$ steps. Given a mergeable pair, it is not difficult to merge them in time $O(n^2)$. Thus, in $O(n^3)$ time we can color all faces. Since much of the above computation can be reused in subsequent steps, we suspect that this can be improved.

Theorem 16 *Given a castable polyhedron $\mathcal{P}$ of size n, the coloring algorithm assigns color and signs to faces of $\mathcal{P}$ in $O(n^3)$ time so that faces of the same color and sign belong to same mould part. Moreover, boundary faces of a cavity of $(\mathcal{P}, \sigma_k)$, for any sequence σ_k of plane-cavities, receive the same color and sign.*

Upon completion, the coloring algorithm will return several double cones of directions, and any feasible pair of opposite removal directions belongs to such a double cone. Afterwards, an eminently practical approach to identify a feasible removal direction is to select a random sample of directions from each cone and test the feasibility of these selected directions using the $O(n \log n)$-time algorithm of Chapter 3

7.5 Further applications to casting

Objects to be manufactured may also be non-castable. For example, a cube with a depression on each face is not castable using two mould parts. Such a problem is usually resolved by using *side-cores*. A side-core is an additional part. For the example of a cube with a depression on each face, we can introduce four side-cores to occupy the depression on each vertical face. The two main mould parts are in contact with the rest of the cube. During object ejection, the four side-cores are removed first, and the two main mould parts can then be removed without blockage.

There is an alternative way to define plane-cavities that facilitates the handling of side-cores. Sweep a plane from a reflex vertex of $\mathcal{P}$ until a saddle vertex (with respect to the sweeping direction) is encountered. We call the volume swept a *restricted plane-cavity* of $\mathcal{P}$. Consider the union of restricted plane-cavities of $\mathcal{P}$. Our techniques can be carried over to show that bounding faces of a connected component in the union must be removed in the same direction. Furthermore, our coloring algorithm will assign the same color and sign to such bounding faces. It is natural to assert that bounding faces of one such connected component should be occupied by a side-core or a main mould part. Thus, this helps us to identify where to use side-cores as well as their retraction directions.