Flipping

`Topics:` Radon's theorem, flipside of a simplex, transformability, incremental construction, correctness, number of flips.

Radon's theorem. The goal of this section is to generalize the idea of edge flipping to three dimensions. We begin with Radon's theorem, which is a result on $d+2$ points in d dimensions. The case $d = 2$ is related to edge flipping in the plane.

THEOREM 3. Every set S of $d+2$ points in $\mathbb{R}^d$ has a partition $S = A \dot\cup B$ with $\operatorname{conv} A \cap \operatorname{conv} B \neq \emptyset$.

PROOF. Since there are more than $d+1$ points they are affinely dependent. Hence there are coefficients λ_i not all zero with $\sum \lambda_i p_i = 0$ and $\sum \lambda_i = 0$. Let I be the set of indices i with $\lambda_i > 0$, and let J contain all other indices. Note that $c = \sum_{i \in I} \lambda_i = -\sum_{j \in J} \lambda_j > 0$, and also

$$x \;=\; \frac{1}{c} \cdot \sum_{i \in I} \lambda_i p_i \;=\; -\frac{1}{c} \cdot \sum_{j \in J} \lambda_j p_j .$$

Let A be the set of points p_i with $i \in I$ and let B contain all other points. Point x is a convex combination of the points in A as well as of the points in B. Equivalently, $x \in \operatorname{conv} A \cap \operatorname{conv} B$. □

A $(d+1)$-dimensional simplex has $d+2$ vertices and a face for every subset of the vertices. If we project its boundary complex to $\mathbb{R}^d$ we get a simplex for every subset of at most $d+1$ vertices. By Radon's theorem, at least two of these simplices have an improper intersection. This intersection comes from projecting the two sides of the simplex boundary on top of each other.

Flipside of a simplex. Consider the case $d = 2$. The projection of a 3-simplex (tetrahedron) to $\mathbb{R}^2$ is either a convex quadrangle or a triangle. In the former case the two diagonals cross, and in the latter case the first vertex lies in the triangle spanned by the other three vertices, see Figure 7. The direction of projection defines an *upper* and a *lower* side of the tetrahedron boundary, and the two sides meet along the *silhouette*. Let $\alpha = \operatorname{conv} A$ and $\beta = \operatorname{conv} B$ be the two faces whose projections have an improper intersection. They lie on

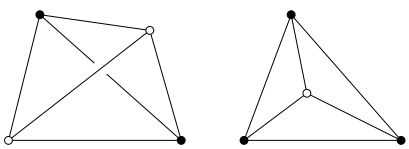

Figure 7: The two generic projections of a tetrahedron to the plane.

opposite sides, and we assume that α belongs to the upper and β to the lower side. The quadrangle case defines an edge flip, which replaces the projection of the upper by the projection of the lower side, or vice versa. We also call this a 2-2 *flip* because it replaces 2 old by 2 new triangles. The triangle case defines a new type of flip, which we refer to as a 1-3 or a 3-1 *flip* depending on whether a new vertex is added or an old vertex is removed.

How do these considerations generalize to the case $d = 3$? The projection of a 4-simplex to $\mathbb{R}^3$ is either a double pyramid and a tetrahedron, see Figure 8. In the

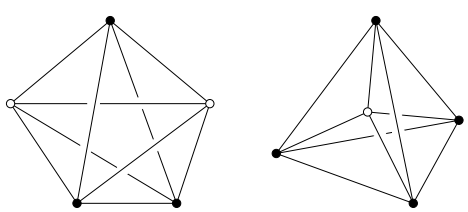

Figure 8: The two generic projections of a 4-simplex to 3-dimensional space.

double pyramid case, α is an edge and β is a triangle. There are three tetrahedra that share α and they form the upper side of the 4-simplex. The remaining two tetrahedra share β and form the lower side. The 3-2 *flip* replaces the projection of the upper side by the projection of the lower side, and the 2-3 *flip* does it the other way round. In the tetrahedron case, α is the first

vertex and β is the tetrahedron spanned by the other four vertices. The 1-4 *flip* adds α effectively replacing β be four tetrahedra, and the 4-1 *flip* removes α.

Transformability. In using flips to construct a Delaunay complex in $\mathbb{R}^3$ we encounter cases where we would like to flip but we cannot. This happens only for 2-3 flips. Let $abcd$ and $bcde$ share the triangle bcd. If the edge ae crosses bcd we can replace $abcd, bcde$ by $baec, caed, daeb$, which is a 2-3 flip. However, if the edge ae misses bcd, as illustrated in Figure 9 where ae passes behind bd, we cannot add ae because it might cross other triangles in the current complex. In this case the union of the two tetrahedra is non-convex. Assume without loss of generality that bd is the non-convex edge. There are two cases. If bd belongs to only

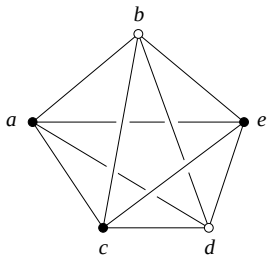

Figure 9: The edge ae does not pass through the triangle bcd but rather behind the edge bd.

three tetrahedra then the third one is $abde$, and we can replace $abdc, cbde, ebda$ by $bace, aced$. This is a 3-2 flip. However, if bd belongs to four or more tetrahedra then we are stuck and cannot remove the triangle bcd. This is the *non-transformable* case.

The occurrence of non-transformable cases does not imply that all hope is lost. It might still be possible to flip elsewhere in a way that resolves non-transformable cases by changing their local neighborhood.

Incremental construction. A special case where we can show that non-transformable configurations can always be resolved is where a single point p is added to the Delaunay complex D of a finite set $S \subseteq \mathbb{R}^3$. The algorithm is based on a generalization of Lemma 4 in Chapter I from the plane to space. Let K be a 3-complex with vertex set S and underlying space the convex hull of S. A triangle $bcd \in K$ is *locally delone* if it belongs to only one tetrahedron or it belongs to two tetrahedra, $abcd$ and $bcde$, and e lies inside the circumsphere of $abcd$. The proof of Lemma 4 below is similar to that of Lemma 4 in Chapter I and is therefore omitted.

LEMMA 4. If every triangle in K is locally delone then K is the Delaunay complex of S.

The first step of the algorithm finds the tetrahedron $abcd \in D$ that contains p. We use a 1-4 flip to split $abcd$ into $abcp, abdp, acdp, bcdp$. The second step applies 2-3 and 3-2 flips to remove triangles that are not locally delone. Assume we have a stack which stores all triangles that are not locally delone for future processing. Initially, the stack contains the four triangles of $abcd$. The iteration continues until the stack is empty.

```
while stack is non-empty do
  pop bcd off the stack;
  if bcd ∈ D and bcd is not locally delone
              and bcd is transformable then
    apply a 2-3 or 3-2 flip to remove bcd;
    push all new triangles on the stack
  endif
endwhile.
```

We prove shortly that eventually all triangles will be locally delone, which implies the algorithm halts with the Delaunay complex of $S \cup \{p\}$. The running time is constant per flip, provided we have an efficient pointer structure representing D. This data structure supports walking around in constant time per step, and it permits local changes in constant time per change.

Correctness. Consider the complex D at some moment in time after adding p when it is not yet the Delaunay complex of $S \cup \{p\}$. It suffices to show D has at least one triangle that is not locally delone but transformable. Let bcd be a triangle in D shared by two tetrahedra $abcd$ and $bcde$. If $p \notin \{a, b, c, d, e\}$ then bcd is locally delone because both $abcd, bcde$ belong to the Delaunay complex of S. This is true right after adding p and it remains true inductively. Indeed, the induction hypothesis permits only flips that destroy exactly one old tetrahedron from the Delaunay complex of S and introduce only tetrahedra that have p as one of their vertices. All new triangles are therefore either in the star or in the link of p. Instead of using the stack for all triangles that are not locally delone it suffices to store only those in the link of p. For each such triangle bcd we have one new tetrahedron in the star, $bcdp$, and one old tetrahedron outside the star, $abcd$, if it exists.

We still have to worry about the case where all link triangles that are not locally delone are non-transformable. To prove that this is not possible we let L be the set of tetrahedra outside the star that have at least one triangle in the link of p. They form a spiky

sphere around p, not unlike the spiky circle in Figure 10. Let $L' \subseteq L$ contain all tetrahedra whose triangles

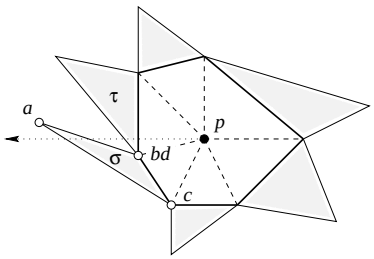

Figure 10: The bold edges belong to the link of p and the shaded triangles belong to L.

in the link are not locally delone, and note that $L \neq \emptyset$. For each tetrahedron $\tau \in L$ consider the power $\pi_\tau(p)$ of p from the circumsphere of τ, and let $\sigma = abcd$ be the tetrahedron that minimizes this distance function. We have $\pi_\sigma(p) < 0$ for else bcd would be locally delone and so would all other link triangles. We argue that σ is also transformable. To get a contradiction assume it is not and let bd be the non-convex edge of the union of $abcd$ and $bcdp$. Let $abdx$ be the tetrahedron on the other side of abd. We have $x \neq p$ for else bcd would be transformable. Consider a half-line that starts at p and passes through an interior point of abc. After crossing the link, the half-line goes through a tetrahedron of L, maybe $abdx$, before it encounters $\sigma = abcd$, see Figure 10. Outside the link we have a genuine Delaunay complex, namely a portion of the Delaunay complex of S. For tetrahedra in this complex the power distance of p from their circumspheres increases along the half-line, which contradicts the minimality assumption in the choice of σ. This finally proves the correctness of the incremental flip algorithm for constructing Delaunay complexes in $\mathbb{R}^3$.

Number of flips. Recall that we interpreted the construction of the Delaunay triangulation by flips in $\mathbb{R}^2$ as gluing tetrahedra in $\mathbb{R}^3$. Similarly, we interpret the construction of the Delaunay complex by flips in $\mathbb{R}^3$ as gluing 4-simplices in $\mathbb{R}^4$. Each flip either introduces an edge or it removes one. Once an edge is removed it cannot be introduced again. This implies that the total number of flips is less than $2\binom{n}{2} < n^2$. Modulo implementation details we thus have an algorithm that constructs the Delaunay complex of n points in $\mathbb{R}^3$ in $O(n^2)$ time. The size of the final Delaunay complex is therefore at most some constant times n^2. There are sets $S \subseteq \mathbb{R}^3$ with $\Omega(n^2)$ Delaunay simplices. However, for point sets that occur in practice, such as dense packings in simple domains, the number of Delaunay simplices is typically less than a constant times n, see also the open problem on average size Delaunay complexes.

Bibliographic notes. Theorem 3 is a byproduct of Johann Radon's attempt [5] to prove a theorem by Eduard Helly [2]. The two theorems are equivalent and form a cornerstone of modern convex geometry. Helly's theorem says that if a finite collection of convex sets in $\mathbb{R}^d$ has the property that every subcollection of $d+1$ sets has a non-empty common intersection then the entire collection has a non-empty intersection.

Algorithms that construct a Delaunay complex in $\mathbb{R}^3$ through flips have first been considered by Barry Joe. In [3] he gives an example where the non-transformable cases form a deadlock situation and flipping does not lead to the Delaunay complex. In [4] he shows that flipping succeeds if the points are added one at a time. The proof of Joe's result in this section is taken from [1] where the same is shown for weighted Delaunay complexes in $\mathbb{R}^d$.

[1] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica* **15** (1996), 223–241.

[2] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Deutsch. Math.-Verein.* **32** (1923), 175–176.

[3] B. Joe. Three-dimensional triangulations from local transformations. *SIAM J. Sci. Statist. Comput.* **10** (1989), 718–741.

[4] B. Joe. Construction of three-dimensional Delaunay triangulations from local transformations. *Comput. Aided Geom. Design* **8** (1991), 123–142.

[5] J. Radon. Mengen konvexer Körper, die einen gemeinschaftlichen Punkt enthalten. *Math. Ann.* **83** (1921), 113–115.