

On monotone sequences of directed flips, triangulations of polyhedra, and structural properties of a directed flip graph

Hang Si

submitted: November 21, 2018

Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: hang.si@wias-berlin.de

No. 2554
Berlin 2018



2010 *Mathematics Subject Classification.* 65D18, 68U05, 65M50.

Key words and phrases. weighted Delaunay triangulations, non-regular triangulations, Lawson's flip algorithm, directed flips, monotone sequence, flip graph, Higher Stasheff-Tamari poset, redundant interior vertices, Schönhardt polyhedron, Steiner points.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: <http://www.wias-berlin.de/>

On monotone sequences of directed flips, triangulations of polyhedra, and structural properties of a directed flip graph

Hang Si

Abstract

This paper studied the geometric and combinatorial aspects of the classical Lawson's flip algorithm [21, 22]. Let \mathbf{A} be a finite point set in \mathbb{R}^2 and $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function which lifts the vertices of \mathbf{A} into \mathbb{R}^3 . Every flip in triangulations of \mathbf{A} can be assigned a direction [6, Definition 6.1.1]. A sequence of directed flips is monotone if all its flips follow the same direction. We first established a relatively obvious relation between monotone sequences of directed flips on triangulations of \mathbf{A} and triangulations of the lifted point set \mathbf{A}^ω in \mathbb{R}^3 . We then studied the structural properties of a directed flip graph (a poset) on the set of all triangulations of \mathbf{A} . We proved several general properties of this poset which clearly explain when Lawson's algorithm works and why it may fail in general. We further characterised the triangulations which cause failure of Lawson's algorithm, and showed that they must contain redundant interior vertices which are not removable by directed flips. A special case of this result in 3d has been shown in [19]. As an application, we described a simple algorithm to triangulate a special class of 3d non-convex polyhedra without using additional vertices. We prove sufficient conditions for the termination of this algorithm, and show it runs in $O(n^3)$ time, where n is the number of input vertices.

1 Introduction

The motivation of this paper is to study and understand the geometric meaning of the Lawson's flipping algorithm [21, 22]. At first, it is well-understood that this algorithm in \mathbb{R}^d acts like to add $d + 1$ simplices to the lower envelope of a $d + 1$ polyhedron. This makes the lower envelope of this polyhedron towards convex. However, it only guarantees the termination in the plane.

This paper takes another view of Lawson's flip algorithm by focusing on the sequence of flips produced by this algorithm. It turns out that this sequence allows a nice geometric interpretation. It corresponds to a triangulation of the volume of a $d + 1$ polyhedron. This relation has already been discovered for a long time. Sleater et al [33, Lemma 5] proved that a sequence of edge-flips converting one triangulation into another of a convex n -gon corresponds to a tetrahedralisation of a 3d polyhedron with these two triangulations as its boundary, see Figure 1. They used this fact to prove that the maximum flip distance between two triangulations of a convex n -gon is $2n - 10$, and it is tight for sufficiently large n [33, Theorem 2]¹.

All flips in the sequences produced by Lawson's algorithm have the same properties, such as (i) every flip corresponds to a unique tetrahedron inside a 3d polyhedron, and (ii) every flip in this sequence follows the same "direction". This means that this sequence is monotone with respect to this direction.

1.1 Our contributions and outline

The excellent book by De Lorea, Rambau, and Santos [6] gives a comprehensive exposition on the topics of this paper and laid down the fundamental mathematical concepts and definitions which are

¹Due to the isomorphism between the flip graph of $n + 2$ nodes and the binary tree rotation graph of n nodes [33, Lemma 1]. It is equivalent to the maximum rotation distance between two n -node binary trees, which is $2n - 6$ for sufficiently large n .

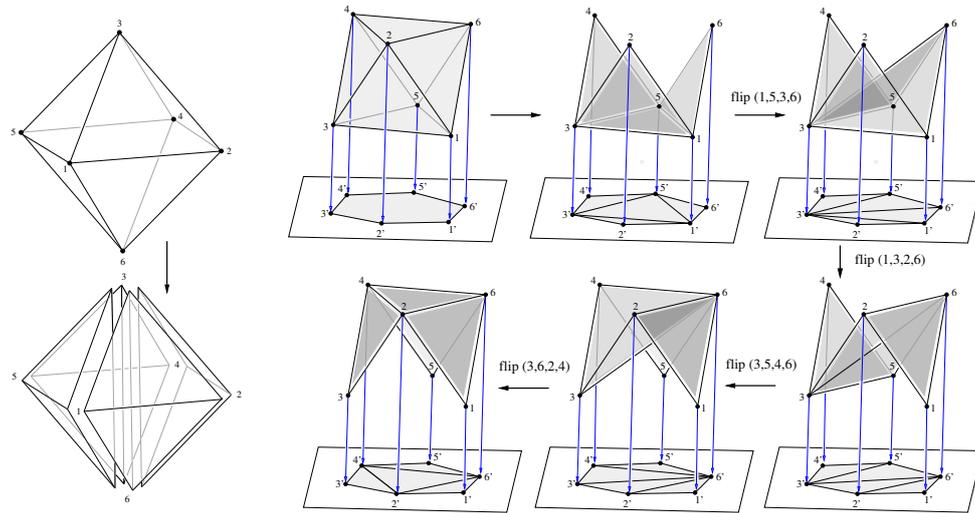


Figure 1: Left: the regular octahedron is decomposed into 4 tetrahedra. Right: a sequence of edge flips corresponds to this tetrahedralisation.

followed by this paper.

Section 2 gives the basic definitions together with an overview of the most related topics and previous works. In particular, the related theorems are introduced. This section is self-contained. It may be skipped for the first reading.

The precise definitions of directed flips, monotone sequence of direct flips, and the directed flip graph are given in section 3. We briefly introduce them here for presenting our results. Let \mathbf{A} be a finite point set in \mathbb{R}^d and let $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function which lifts every vertex in \mathbf{A} into a lifted point in \mathbb{R}^{d+1} . Let \mathbf{A}^ω in \mathbb{R}^{d+1} be the set of lifted points of \mathbf{A} . Then every flip in a triangulation of \mathbf{A} corresponds to a $d + 1$ -simplex whose vertices are in \mathbf{A}^ω . Let \mathcal{T}_1 and \mathcal{T}_2 be two triangulations of \mathbf{A} , and \mathcal{T}_2 is obtained from \mathcal{T}_1 by a flip. We call this flip *down-flip* if the upper faces of the $d + 1$ simplex are in \mathcal{T}_1 and lower faces are in \mathcal{T}_2 , otherwise, it is an *up-flip*. Figure 2 illustrates these two directed flips.

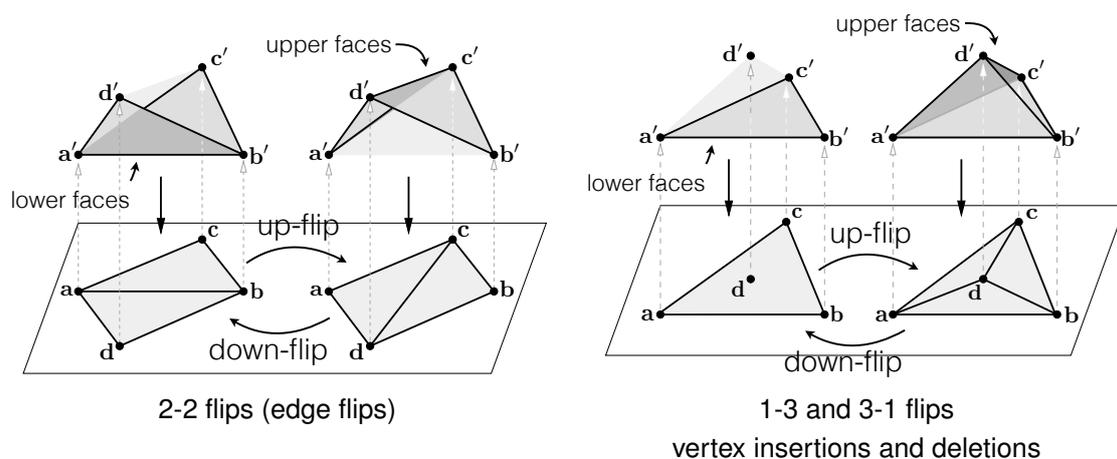


Figure 2: The two directed flips in the plane. In particular, an up-flip replaces the lower by the upper faces of the support of this flip, and a down-flip does the opposite way as an up-flip does.

By this definition, any flip has a direction. A sequence of directed flips is *monotone* if all flips in this

sequence have the same direction.

Section 4 and Section 5 are devoted to the proof of a relation between monotone sequences of directed flips in the triangulations of a point set in the plane and triangulations of a polyhedron in \mathbb{R}^3 .

- (1) Any monotone sequence of directed flips between two triangulations \mathcal{T}_1 and \mathcal{T}_2 of \mathbf{A} (or vice versa) corresponds to a 3d triangulation of a polyhedron whose boundary consists of these two triangulations (Theorem 5 and Corollary 6).
- (2) A 3d triangulation \mathcal{T}_{uv} of a polyhedron with vertices in \mathbf{A}^ω corresponds to a monotone sequence of directed flips between \mathcal{T}_u and \mathcal{T}_v (or vice versa) if it is acyclic with respect to a viewpoint in the direction of the directed flips (Theorem 7 and Corollary 8).

This relation is relatively obvious and has been partially proven and remarked in many previous works which will be reviewed in Section 2. We comment that (1) is a straightforward generalisation of many special cases. It is proven for the set of points of convex n -gons in [33, Lemma 5]². It is also proven for the set of triangulations of cyclic polytopes [6, Theorem 6.1.19]. It is implicitly proven in [6, Theorem 5.3.7] for the set of regular triangulations (see Theorem 3). However, we're not aware of a complete proof of this fact for an arbitrary point set. We prove this relation for an arbitrary finite point set in \mathbb{R}^2 which involves all types of elementary flips of \mathbb{R}^2 , i.e., edge flips (2-2 flips), vertex insertions (1-3 flips), and vertex deletions (3-1 flips). In particular, we show that non-regular triangulations may appear within the sequence. The proof of (1) is given in Section 4.

Fact (2) has been partially proven in the theory of Secondary polytopes, see e.g. [6, Theorem 5.3.7] as well as the set of triangulations of cyclic polytopes [6, Theorem 6.1.19]. To the best of our knowledge, this fact has not been established for an arbitrary finite point set. The proof of (2) is given in Section 5.

From the directed flips, one can construct a directed flip graph (a poset) on the set of all triangulations of \mathbf{A} , exactly the same as the construction of the first Higher Stasheff-Tamari (**HST**) poset from triangulations of point sets of cyclic polytopes [8, 25]. The structure properties of the **HST** has been well understood [6, Theorem 6.1.19]. However, the structure of this poset for an arbitrary finite point set is more complicated and interesting. Figure 3 shows a motivation example of this poset which is discussed in detail in Section 3.1.

The main result of this paper is to illustrate several general and special structural properties of this poset. Section 6 is devoted to this purpose.

In Section 6.1, a general structural theorem (Theorem 11) of this poset which includes various obvious properties is proven. These properties are easily derived from the relations proven in (1) and (2). In particular, we showed that this poset is in general not bounded, i.e., it has neither unique minimum nor unique maximum. On the other hand, the convex or concave property of the height function ω gives special property of this poset. In Section 6.2, we proved that such a poset will contains either unique minimum or unique maximum, but not both in general (Theorem 12). The existence of a unique minimum or maximum guarantees that the Lawson's flip algorithm will terminate when its flips are coincident with this direction.

We further characterised those triangulations which cause the failure of Lawson's algorithm. Figure 4 shows such as example. Theorem 16 showed that such triangulations must contains *redundant inte-*

²In their proof, only the case of edge-flips is considered, but it could be generalised to show the cases which involve vertex insertion and deletion as well.

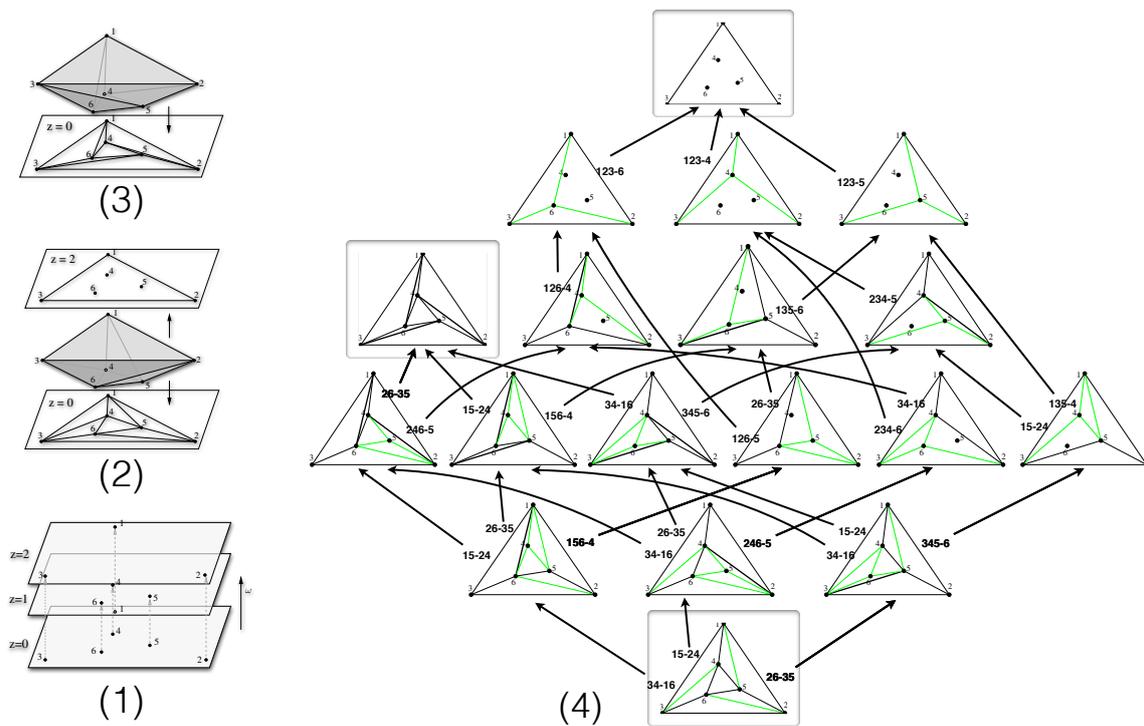


Figure 3: The poset of monotone sequence of directed flips of a point set of 6 points in the plane and a height function ω . (1) The lifted point set \mathbf{A}^ω is produced by the height function ω which sends the vertices 1, 2, 3 to the plane $z = 2$, and the vertices 4, 5, 6 to the plane $z = 1$. (2) The regular and farthest point regular triangulations of (\mathbf{A}, ω) . (3) A non-regular triangulation of \mathbf{A} . (4) The green edges are locally non-regular and flippable edges.

rior vertices (defined in Section 6.3) which cannot be removed by the corresponding directed flips. We proved this by showing that these vertices are connected by cycles of connected unflippable locally non-regular edges with respect to the corresponding directed flips. (The definition of unflippability is given in Section 2.5.) Interestingly, Joe [19] proved a special case in 3d triangulations, which showed that the failure of Lawson's algorithm in 3d is due to the existence of a cycle of connected unflippable locally non-Delaunay faces [19, Lemma 4, 5]. This results implies that there must exist cycles of unflippable locally non-Delaunay edges in such 3d triangulations. Moreover, these edges are redundant interior edges.

In Section 7, we showed an application of the obtained structural property theorem. We showed that a slightly modified Lawson's algorithm can be used to triangulate a special class of 3d non-convex polyhedra without using additional vertices. We prove sufficient conditions for the termination of this algorithm, and show it runs in $O(n^3)$ time, where n is the number of vertices of the polyhedron.

Finally, We present some open questions regarding the full structural properties of this poset and the set of all triangulations of \mathbf{A}^ω in Section 8.

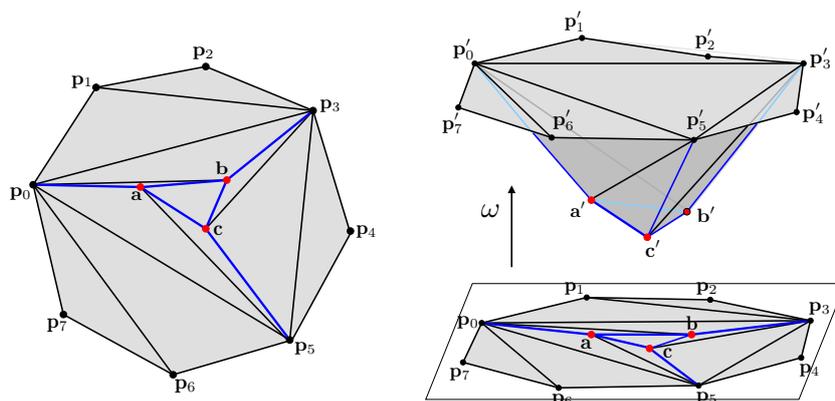


Figure 4: Redundant interior vertices (definition is given in Section 6.3) are shown in red. Left: A triangulation corresponds to an upper non-extreme node in this poset. The red vertices are lower interior vertices. The blue edges are locally non-regular with respect to up-flips. Moreover, all blue edges are unflippable. Right: The lifted triangulation are shown. The three edges: e_{ab} , e_{bc} , e_{ca} form a cycle of unflippable locally non-regular edges with respect to up-flips.

2 Definitions and Related Works

2.1 Triangulations of a point set

Let \mathbf{A} be a finite point set in \mathbb{R}^d . A *triangulation* of \mathbf{A} [6, Def 2.2.1] is a collection \mathcal{T} of i -dimensional simplices, where $i = -1, 0, \dots, d$, whose vertices are in \mathbf{A} , and it satisfies three conditions, which are:

- (1) All faces of simplices of \mathcal{T} are in \mathcal{T} . (Closure Property)
- (2) The intersection of any two simplices of \mathcal{T} is a (possibly empty) face of both. (Intersection Property)
- (3) The union of these simplices equals to the convex hull of \mathbf{A} , denoted as $\text{conv}(\mathbf{A})$. (Union Property)

In other words, a triangulation of \mathbf{A} is a geometrically realised simplicial complex whose underlying space is $\text{conv}(\mathbf{A})$. A triangulation of \mathbf{A} needs not contain all vertices of \mathbf{A} .

Figure 5 shows several triangulations of a two dimensional point set. In particular, the famous (*nearest point*) *Delaunay triangulation*, is the triangulation of \mathbf{A} such that no vertex of \mathbf{A} is inside the circumcircle of any triangle of it [7]. Likewise, the *farthest point Delaunay triangulation* is the triangulation of \mathbf{A} such that no vertex of \mathbf{A} is outside the circumcircle of every triangle of it, see e.g. [11].

2.2 Regular triangulations

Given a point set, there are many triangulations of it. A (nice) family of triangulations of a finite point set in \mathbb{R}^d are canonical projections of the envelopes of a convex $d + 1$ dimensional polytope. They are called *regular triangulations*, also known as *weighted Delaunay* [10], *Gale*, or *coherent* triangulations, see e.g., [37, 6].

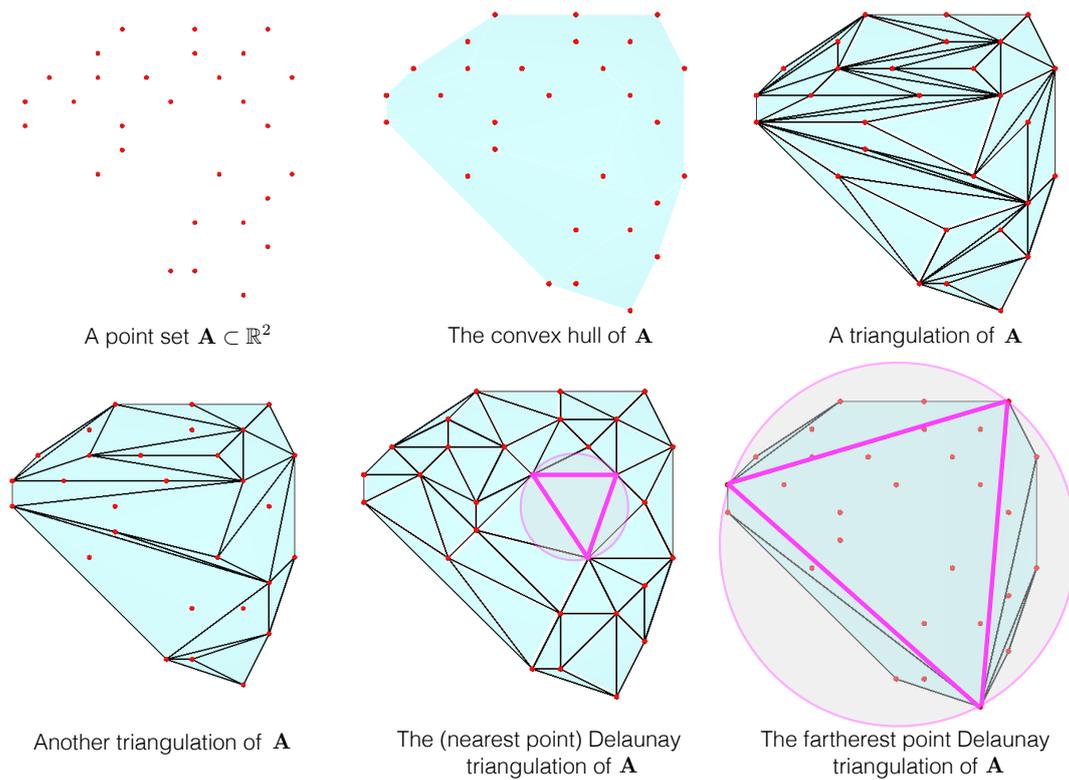


Figure 5: Triangulations of a point set \mathbf{A} in the plane.

Let \mathbf{A} be a finite point set in \mathbb{R}^d . Pick a *height function* $\omega : \mathbf{A} \rightarrow \mathbb{R}$. For a point $\mathbf{p} = (p_1, \dots, p_d) \in \mathbf{A}$, define its *lifted point* $\mathbf{p}' = (p_1, \dots, p_d, p_{d+1}) \in \mathbb{R}^{d+1}$, where $p_{d+1} = \omega(\mathbf{p})$ is called the *height* of \mathbf{p} . The point set $\mathbf{A}^\omega := \{\mathbf{p}' \mid \mathbf{p} \in \mathbf{A}\} \subseteq \mathbb{R}^{d+1}$ is called the *lifted point set* of \mathbf{A} .

The convex hull of \mathbf{A}^ω , $\text{conv}(\mathbf{A}^\omega)$, is a 3d polytope. A *lower face* of $\text{conv}(\mathbf{A}^\omega)$ is a facet of $\text{conv}(\mathbf{A}^\omega)$ that has a non-vertical supporting hyperplane with $\text{conv}(\mathbf{A}^\omega)$ lies entirely in the half-space which does not contain the point $(\mathbf{0}, -\infty)$. Simply saying, it is a face of $\text{conv}(\mathbf{A}^\omega)$ that is “visible from below” by placing the viewpoint at $(\mathbf{0}, -\infty)$. Likewise, an *upper face* of $\text{conv}(\mathbf{A}^\omega)$ is a facet of $\text{conv}(\mathbf{A}^\omega)$ that has a non-vertical supporting hyperplane with $\text{conv}(\mathbf{A}^\omega)$ lies entirely in the half-space which does not contain the point $(\mathbf{0}, +\infty)$. Simply saying, it is a face that is “visible from above” by placing the viewpoint at $(\mathbf{0}, +\infty)$, see Figure 6 for examples.

The *canonical projection* is a map $\pi : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d, (p_1, \dots, p_d, p_{d+1}) \rightarrow (p_1, \dots, p_d)$ which “deletes the last coordinate of a given point. Let us assume that ω is *sufficiently generic*, by which we mean that no $d + 1$ vertices of \mathbf{A} are lifted to lie in a non-vertical hyperplane in \mathbb{R}^{d+1} (This implies that \mathbf{A} is in general position). The canonical projection of the set of lower faces of $\text{conv}(\mathbf{A}^\omega)$ gives a unique triangulation of \mathbf{A} , which is called the *(nearest point) regular triangulation* of (\mathbf{A}, ω) . Likewise, the canonical projection of the set of upper faces of $\text{conv}(\mathbf{A}^\omega)$ is also a unique triangulation of \mathbf{A} , which is the *farthest point regular triangulation* of (\mathbf{A}, ω) . Figure 6 Left shows an example of these two triangulations of \mathbf{A} . Note that some points of \mathbf{A} may be contained in neither the regular nor farthest point regular triangulations of (\mathbf{A}, ω) . By varying the height function ω , we will obtain different regular triangulations of \mathbf{A} . In particular, the Delaunay and the farthest point Delaunay triangulations of \mathbf{A} are two special triangulations when $\omega = x^2 + y^2$, see Figure 6 Right.

Both the regular and farthest point regular triangulations of (\mathbf{A}, ω) in $\mathbb{R}^d \times \mathbb{R}$ can be constructed efficiently using incremental vertex insertion algorithms [4, 35, 20, 10]. In particular, Joe first proved

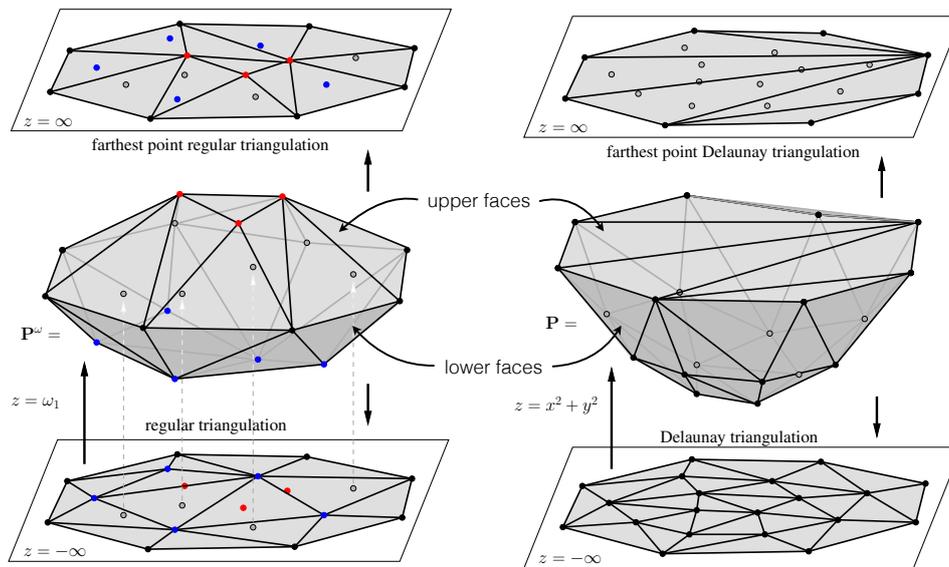


Figure 6: Left: the regular and farthest point regular triangulations of (\mathbf{A}, ω) in $\mathbb{R}^2 \times \mathbb{R}$. Right: the Delaunay and farthest point Delaunay triangulations of \mathbf{A} when $\omega = x^2 + y^2$.

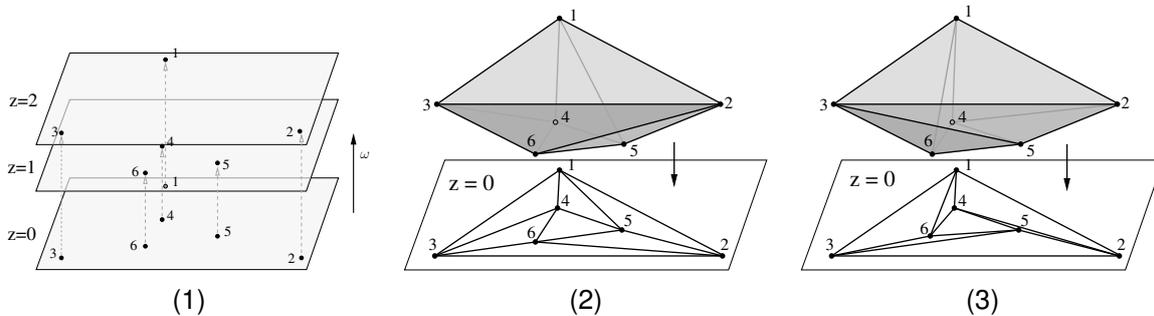


Figure 7: Triangulations of a (weighted) point set \mathbf{A} of 6 points in the plane. (1) The lifted point set \mathbf{A}^ω is produced by the height function ω which sends the vertices 1, 2, 3 to the plane $z = 2$, and the vertices 4, 5, 6 to the plane $z = 1$. (2) The regular triangulation of (\mathbf{A}, ω) . (3) A non-regular triangulation of \mathbf{A} .

an incremental flip algorithm to construct 3d Delaunay triangulations. The proof [20, Lemma 8, 9] is based on the Acyclic Theorem of Edelsbrunner [9] (see Section 2.4). Edelsbrunner and Shah [10] proved that this algorithm can construct regular triangulations in any dimensions.

2.3 Non-regular triangulations

A triangulation of \mathbf{A} which is not regular is called a *non-regular triangulation* of \mathbf{A} . Equivalently, a *non-regular triangulation* of \mathbf{A} is a triangulation of \mathbf{A} such that there exists no height function ω which can “lift it into the lower or upper envelope of the convex hull of \mathbf{A}^ω ”. Although the properties of regular triangulations are well studied, e.g. by the theory of secondary polytopes [13] (see Section 2.7), “not so much is known about non-regular triangulations”(quoted from [17]).

Figure 7 illustrates the simplest example of a non-regular triangulation of set of 6 vertices in the plane. In this example, the height function is chosen such that the outmost three vertices are on a higher plane than the three innermost vertices, see (1). The regular triangulation is shown in (2). The

triangulation in (3) is non-regular. In fact, no matter which height function ω is, the triangulation in (3) will correspond to neither the lower nor the upper faces of $\text{conv}(\mathbf{A}^\omega)$.

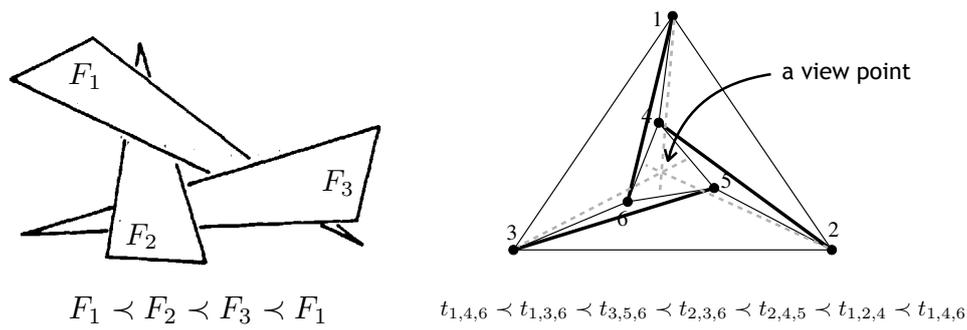


Figure 8: Cycles of simplifies in the in_front/behind relation. Left: the three triangles forms a cycle with respect to your eye position. Right: the three triangles form a cycle with respect to the fixed point.

Given a triangulation of a point set, to check whether it is regular or not is to find whether there exists a height function ω (a set of weights on the set of points). This is a linear program problem.

2.4 The Acyclic Theorem

Edelsbrunner proved an important property which owed by all regular triangulations [9]. Let $\mathbf{x} \in \mathbb{R}^d$ be an arbitrary but fixed viewpoint. A simplex σ lies *in front of* another simplex τ if there is a half-line starting at \mathbf{x} that first passes through σ and then through τ , denoted as $\sigma \prec_{\mathbf{x}} \tau$. The set of simplices in a triangulation together with $\prec_{\mathbf{x}}$ forms a relation. In general, this relation can have cycles, which are sequences $\tau_0 \prec_{\mathbf{x}} \tau_1, \dots, \tau_k \prec_{\mathbf{x}} \tau_0$, see Figure 8 Left.

Theorem 1 (Edelsbrunner [9]). *From any fixed viewpoint in \mathbb{R}^d , the in_front/behind relation for simplices of regular triangulations forms no cycle.*

This theorem, besides its many applications, provides a useful way to prove the non-regularity of a triangulation: *If a triangulation contains a cycle with respect to any viewpoint, then it is non-regular.* The cycle with respect to the non-regular triangulation in Figure 7 (3) is shown in Figure 8 Right.

However, this theorem does not imply that the acyclic property belongs to regular triangulations only. Indeed, there are non-regular triangulations which are also acyclic viewed from any viewpoint. This is first shown by Jaume and Rote [17].

2.5 Flips

All point sets in this subsection are assumed to be in *general position*, which we mean that no $d + 2$ points of a set of d -dimensional points lie on the same hyperplane.

Let \mathbf{T} be a set of $d + 2$ points in \mathbb{R}^d , the Radon's theorem [24] states that there exists a partition of \mathbf{T} into two subsets $\mathbf{U}, \mathbf{V} \subset \mathbf{T}$ such that $\text{conv}(\mathbf{U}) \cap \text{conv}(\mathbf{V}) \neq \emptyset$. This means every triangulation of \mathbf{T} either contains \mathbf{U} as a subset, or \mathbf{V} , but not both. It is shown by Lawson [23], that a set of $d + 2$ points can only have two triangulations. A *flip* is the operation that substitutes one triangulation of \mathbf{T} for the other.

In \mathbb{R}^2 , there are two Radon partitions of 4 points. This defines three types of flips, shown in Figure 9. We can distinguish these types of flips by the number of triangles in the triangulations before and after the flip. Hence there are *2-2 flips* (or *edge-flips*), *1-3 flips* (or *vertex-insertion*), and *3-1 flips* (or *vertex-deletion*), see Figure 9. Flips in \mathbb{R}^3 are classified accordingly.

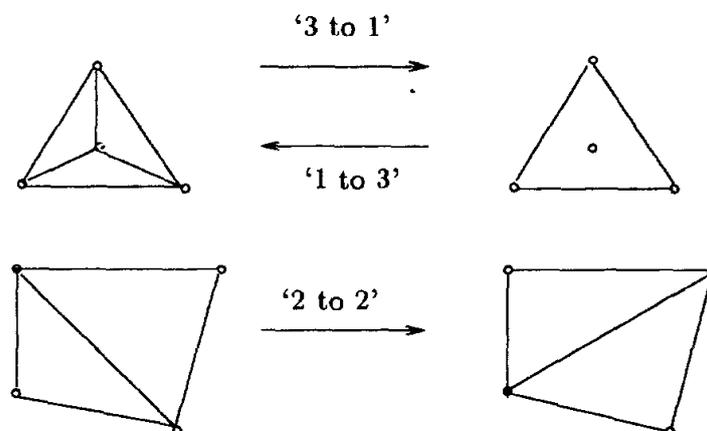


Figure 9: Flips in \mathbb{R}^2 .

Flippability Let \mathcal{K} be triangulation of a point set \mathbf{A} in \mathbb{R}^2 . Consider an edge $e_{ab} \in \mathcal{K}$. Assume it belongs to two triangles $f_{abc}, f_{abd} \in \mathcal{K}$. We say e_{ab} is *flippable* in \mathcal{K} if there exists a subcomplex $\mathcal{S} \subseteq \mathcal{K}$ such that the underlying space $|\mathcal{S}|$ is equal to the convex hull of $\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$. The edge e_{ab} is flippable implies that either $|\mathcal{S}| = f_{abc} \cup f_{abd}$ or there exists a third triangle, without loss of generality, assume \mathbf{a} is the reflex vertex, $f_{acd} \in \mathcal{K}$, such that $|\mathcal{S}| = f_{abc} \cup f_{abd} \cup f_{acd}$. In the former case, e_{ab} admits a 2-2 flip (an edge flip), and in the latter, it admits a 3-1 flip (vertex removal). We say that the convex hull of $\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$ is the *support* of this flip. Otherwise, e_{ab} is *unflippable*. This means, if we perform a flip on this edge, it will result a simplicial complex which is not a geometric triangulation according to our definition. Especially, the Intersection Property (IP) does not hold, see Figure 10 for examples. If an edge e_{ab} is unflippable, there must exist a vertex, either \mathbf{a} or \mathbf{b} such that it lies in the interior of the support $\text{conv}\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$ of this flip, and the degree of this vertex must larger than 3. We say that *the unflippability of this edge is due to this vertex*. For an example, in Figure 10, the edge e_{ef} is unflippable and the unflippability of this edge is due to \mathbf{e} .

2.6 Lawson's flip algorithm

A classical algorithm (due to Charles Lawson [22]) shows that one can transform any triangulation of \mathbf{A} into the Delaunay triangulation of \mathbf{A} by a sequence of edge flips.

Let \mathcal{K} be a triangulation of \mathbf{A} in \mathbb{R}^2 . An edge $e_{ab} \in \mathcal{K}$ is *locally Delaunay* if either (i) it is on the convex hull of \mathbf{A} , or (ii) it belongs two triangles, $f_{abc}, f_{abd} \in \mathcal{K}$, and \mathbf{d} does not lie inside of the circumcircle of f_{abc} . A locally Delaunay edge is not necessarily a Delaunay edge. The Delaunay Lemma [7] proved that if every edge of \mathcal{K} is locally Delaunay, then \mathcal{K} is the Delaunay triangulation of \mathbf{A} .

If an edge $e_{ab} \in \mathcal{K}$ is not locally Delaunay, then it must be flippable (an easy exercise). A 2-2 flip can be used to replace e_{ab} by the edge e_{cd} in \mathcal{K} such that e_{cd} must be locally Delaunay, see Figure 11. We can use edge flips as elementary operations to convert an arbitrary triangulation \mathcal{K} to the Delaunay triangulation.

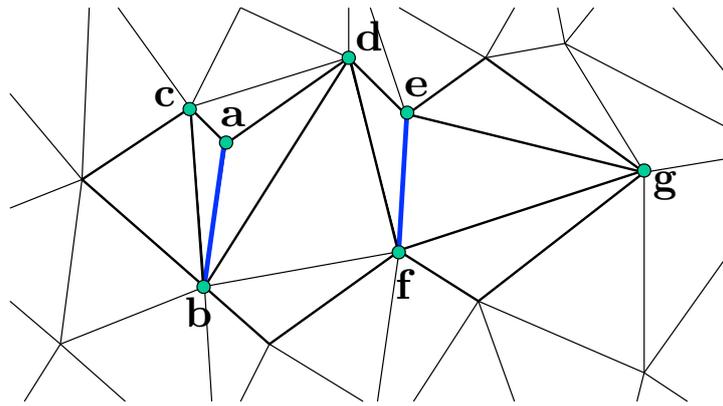


Figure 10: Flippability of edges in \mathbb{R}^2 . Edge e_{ab} is flippable (by a 3-1 flip) whose support of this flip is the convex hull of the set $\{a, b, c, d\}$, so do the edges e_{bd} and e_{df} (by 2-2 flips), while the edge e_{ef} is unflippable. Moreover, the unflippability is due to the vertex e .

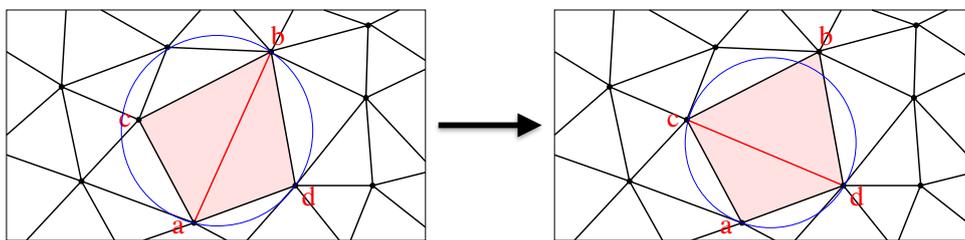


Figure 11: The edge flips from e_{ab} to e_{cd} replaces the locally non-Delaunay edge e_{ab} by the locally Delaunay edge e_{cd} .

The Lawson's algorithm uses a stack to maintain all edges which may be locally non-Delaunay. Initially, all edges of \mathcal{K} are pushed on the stack. It then pop these edges (one after one) and performs edge flips on those edges which are not locally Delaunay. After each flip, the stack will be updated by new edges on the convex hull of the support of this flip. The algorithm finishes when the stack is empty.

Lawson's algorithm can be understood as gluing a sequence of tetrahedra. A locally non-Delaunay edge e_{ab} corresponding a locally non-convex edge $e_{a'b'}$ in \mathbb{R}^3 . Flipping e_{ab} to e_{cd} is likely gluing a tetrahedron $t_{a'b'c'd'}$ from below to the triangles $f_{a'b'c'}$ and $f_{a'b'd'}$ in \mathbb{R}^3 . Once we glue $t_{a'b'c'd'}$ we cannot glue another tetrahedron right below $e_{a'b'}$. In other words, once we flip e_{ab} we cannot introduce e_{ab} again by some other flip. This implies that Lawson's algorithm will eventually terminate when all locally non-Delaunay edges are flipped. By the Delaunay lemma, the triangulation is Delaunay. Figure 12 illustrates the algorithm in both 2d and 3d. This also implies there are at most as many flips as there are edges connecting n points, namely $\binom{n}{2}$. Each flip takes constant time, hence the total running time is $O(n^2)$.

2.7 The undirected flip graph

The *undirected flip graph* of a finite point set \mathbf{A} is defined as follows: Each triangulation of \mathbf{A} is a node, and there is an edge from μ to ν if there is a flip that changes the triangulation μ to ν , see Figure 13 for an example.

Lawson's algorithm [22] shows that one can transform any triangulation of \mathbf{A} into the Delaunay tri-

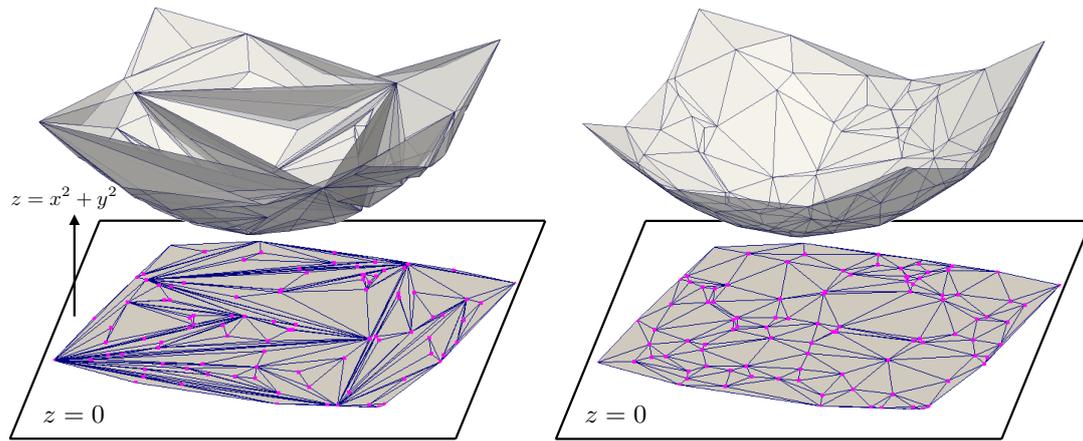


Figure 12: The lifted view of Lawson's flip algorithm which transforms a non-convex surface (left) into a convex one (right).

angulation of \mathbf{A} by a sequence of edge flips. Since one can always add or remove a vertex in a 2d triangulation by a sequence of 2-2, 1-3 and 3-1 flips. This shows that the (undirected) flip graph of any 2d point set is connected.

The undirected flip graph of a point set in three and higher dimensions are much more complicated. It was first proven by Joe [19] that Lawson's algorithm may fail in 3d. Moreover, the failure is due to the existence of a cycle of connected unflippable locally non-Delaunay faces [19, Lemma 4, 5].

A construction of a triangulation of a point set in \mathbb{R}^5 given by Santos [29, 28] shows that the undirected flip graph of a point set in \mathbb{R}^5 is disconnected. This leads to the following open problem.

Problem 2 (Open Problem). *Whether the undirected flip graph of a point set in \mathbb{R}^3 or \mathbb{R}^4 is connected or not.*

A nice result from the secondary polytope theory [13] shows that the undirected flip graph of the set of all regular triangulations of any point set in \mathbb{R}^d is connected. This implies the following very useful theorem which will be extensively used in this paper. A proof of this theorem is found in the book [6, Theorem 5.3.7]. The definition of "directed flips" is given later in Section 3.

Theorem 3. *Between every two regular triangulations \mathcal{T}_1 and \mathcal{T}_2 of \mathbf{A} , there exists a monotone sequence of directed flips which starts at \mathcal{T}_1 and ends at \mathcal{T}_2 or vice versa.*

Another interesting example is the vertex sets of *cyclic polytopes*, see e.g. [37, Chap 0] and [6, Section 6.1]. It has been shown that there is a "friendly" poset structure on the set of all triangulations of cyclic polytopes (introduced below). The Hasse diagram of this poset will be the graph of all triangulations. In particular, this graph is connected, see [8, 25] and [6, Section 6.1].

For the most recent and comprehensive exposition of undirected flip graphs, as well as the most latest results, we refer to the book [6].

2.8 Higher Stasheff-Tamari (HST) posets

Besides the undirected flip graph, there is another structure on a point set by placing a partial order on triangulations of this point set.

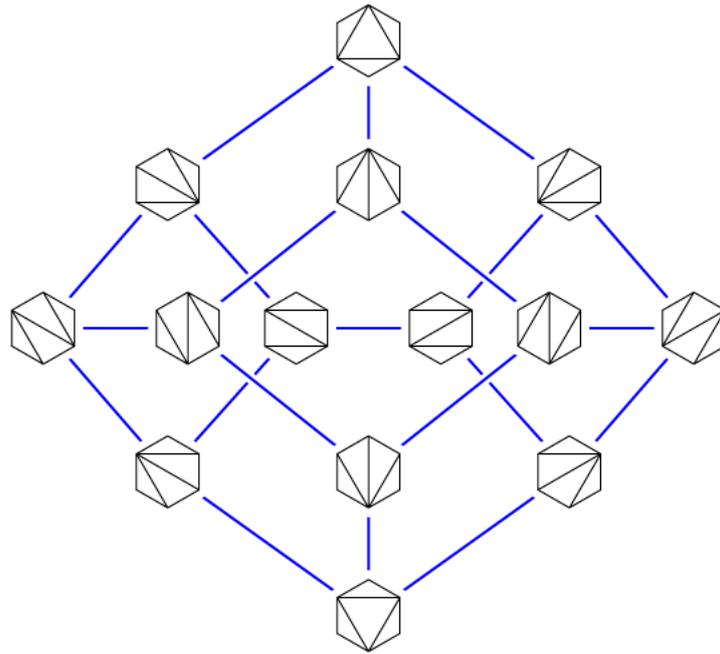


Figure 13: The (undirected) flip graph of a the vertex set of a hexagon ([12]).

Let \mathbf{C} be the vertex set of a *cyclic polytope* $\mathbf{C}(n, 2)$ in \mathbb{R}^2 (see e.g. [37, Chap 0]), where n is the number of points of \mathbf{C} . Let the four points $\mathbf{p}_i, \mathbf{p}_j, \mathbf{p}_k, \mathbf{p}_l \in \mathbf{A}$, where $i < j < k < l$, form a quadrilateral. A flip is called *up-flip* if it replaces the diagonal edge $[\mathbf{p}_i, \mathbf{p}_k]$ by $[\mathbf{p}_j, \mathbf{p}_l]$. It is a *down-flip* otherwise. Let $\mathcal{T}_1, \mathcal{T}_2$ be two triangulations of \mathbf{C} , we say $\mathcal{T}_1 \leq_1 \mathcal{T}_2$ when \mathcal{T}_2 can be produced from \mathcal{T}_1 by a (possibly empty) sequence of up-flips. Then the set of all triangulations of \mathbf{C} is a partially ordered set (poset), see Figure 14 Left. It is known as the *first Higher Stasheff-Tamari poset*, denoted as \mathbf{HST}_1 . By using the height of the characteristic sections [6, Def. 6.1.15] as the partial order, one obtain another poset structure known as the *second Higher Stasheff-Tamari poset*, denoted as \mathbf{HST}_2 , see Figure 14 Right.

These two poset link to many nice combinatorial and geometric objects. Its property has been well studied [8, 25]. The following theorem is about the structure properties of triangulations of cyclic polytopes proven in [6]. It is proven for cyclic polytope of any dimension. We state the two-dimension case only.

Theorem 4. [6, Theorem 6.1.19] *Let \mathbf{A} be the point set of $\mathbf{C}(n, 2)$. Let ω be a height function defined on the vertices of \mathbf{A} . Then*

- (1) *This poset is bounded. It has a unique minimum which is the regular triangulation of (\mathbf{A}, ω) . It has a unique maximum, which is the farthest-point regular triangulation of (\mathbf{A}, ω) .*
- (2) *For any triangulation of \mathbf{A}^ω there exists a maximum sequence of flips in the poset, i.e., all triangulations of \mathbf{A}^ω are coded in this poset.*
- (3) *All internal nodes of this poset are 2d regular triangulations of \mathbf{A} .*

It is shown that the Hasse diagram of this poset will be the graph of all triangulations. It is a lattice, which means, it has unique minimum and maximum elements and every two elements have a unique minimum upper bound. In particular, there is a bijective relation between the maximal chain (sequence)

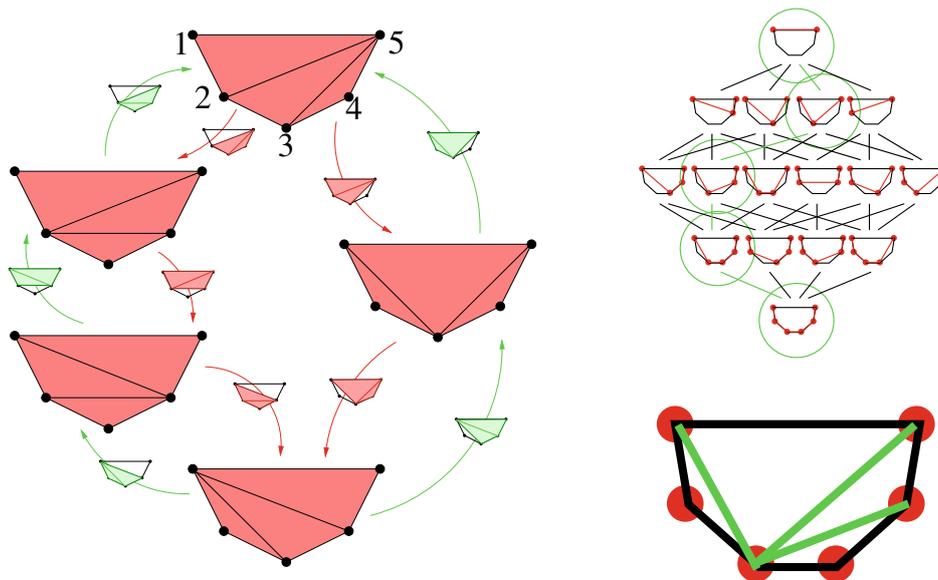


Figure 14: Examples of higher Stasheff-Tamari (HST) Posets for triangulations of cyclic polytopes (Courtesy of [6]). Left: HST_1 of the vertex set a 5-gon. It is the vertex set of the $2d$ cyclic polytope $C(5, 2)$. Right: HST_2 of the vertex set $C(6, 1)$, and a triangulation of $C(6, 2)$ resulted by a maximal chain.

of flips and all triangulations of the lifted cyclic polytope in \mathbb{R}^3 , see an example in Figure 14 Right. This is the key fact to show that the flip graph of the vertex set (up to \mathbb{R}^5) is connected.

2.9 Triangulating non-convex polyhedra

Decomposing a 3d polyhedron into a set of tetrahedra that forms a simplicial complex is a classical problem in computational geometry.

It is shown that this problem has many theoretical difficulties. There exists simple non-convex polyhedra whose interior cannot be triangulated with its own vertices. The smallest example is constructed by Schönhardt in 1928 [30], which is a twisted triangular prism of 6 vertices, see Figure 15 Left. Other construction of such polyhedra are reported in [1, 18, 5, 26, 3, 32]. Ruppert and Seidel [27] proved that the problem to determine whether a simple non-convex polyhedron can be triangulated without Steiner points is NP-complete.

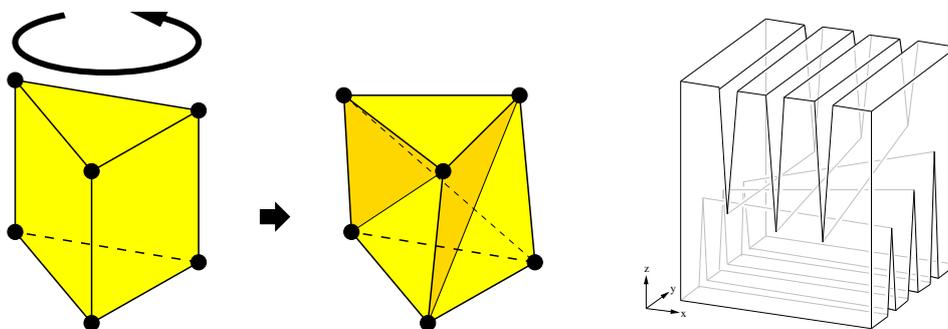


Figure 15: Left: The Schönhardt polyhedron. Right: A Chazelle polyhedron.

If addition vertices are allowed, so-called *Steiner points*, Chazelle [5] constructs a family of polyhedra with n vertices and proved that $\Omega(n^2)$ Steiner points are needed to triangulate them, see Figure 15 Right. Chazelle also proved that any simple polyhedron of n vertices can be triangulated into $O(n^2)$ tetrahedra [5]. In practice, a challenging question is: *Given any polyhedron, how to triangulate it with the number of Steiner points as small as possible?*

The problem is much difficult when the boundary of the polyhedra needs to be preserved, i.e., Steiner points, if they are needed, are only allowed in the interior of the polyhedra. There are algorithms based on heuristics, see e.g [15, 36, 14, 31], but none of them provides any bound on the number of interior Steiner points.

There are some interesting results on triangulating special classes simple and non-simple polyhedra without Steiner points. We refer to the work of Goodman and Pach [16], Bern [2], and Toussaint et al [34]. In particular, Toussaint et al [34] showed that certain classes of rectilinear (isothetic) simple polyhedra can always be triangulated in $O(n^2)$ time where n is the number of vertices in the polyhedron. They also showed that polyhedral slabs (even with holes) as well as subdivision slabs can always be triangulated in $O(n \log n)$ time. Furthermore, for simple polyhedral slabs $O(n)$ time suffices. They showed that polyhedra that are the union of three convex polyhedra can always be triangulated in $O(n^2)$ time. They remarked that: *one direction of future research is to find other nontrivial classes polyhedra which can be triangulated.*

3 Monotone Sequences of Directed Flips and the Directed Flip Graph

Let \mathbf{A} be a finite point set in \mathbb{R}^2 and let $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function which lifts every vertex $\mathbf{p} = (p_1, p_2) \in \mathbf{A}$ into a lifted point $\mathbf{p}' = (p_1, p_2, \omega(\mathbf{p})) \in \mathbb{R}^3$. Let $\mathbf{A}^\omega = \{\mathbf{p}' \mid \mathbf{p} \in \mathbf{A}\}$ in \mathbb{R}^3 be the set of lifted points of \mathbf{A} .

Let $\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\} \subset \mathbf{A}$ be the support of a flip (2-2 flip, 1-3 flip, or 3-1 flip) in a triangulation of \mathbf{A} . This flip corresponds to a tetrahedron $t_{\mathbf{a}'\mathbf{b}'\mathbf{c}'\mathbf{d}'}$, where $\{\mathbf{a}', \mathbf{b}', \mathbf{c}', \mathbf{d}'\} \subset \mathbf{A}^\omega$. (The definition of the support of a flip is given in Section 2.5.) Let \mathcal{T}_1 and \mathcal{T}_2 be two triangulations of \mathbf{A} differing by a flip. We say \mathcal{T}_1 is *decreasing with respect to the height function ω* if the upper faces of t appear in \mathcal{T}_1 and the lower faces of t appear in \mathcal{T}_2 . In this case, we call this flip a *down-flip*. Likewise, we say \mathcal{T}_1 is *increasing with respect to the height function ω* if it is the other way around. In this case, we call this flip an *up-flip*. The two types of directed flips are illustrated in Figure 2.

We say a sequence of directed flips *monotone* if all directed flips in this sequence are of the same type, i.e., all flips follow the same direction.

Similar to the construction of the first Higher Stasheff-Tamari (\mathbf{HST}_1) poset [25, 6]. We can construct a directed flip graph on the set of triangulations of a point set by using one of the two types of directed flips.

We first define a partial order on the set of triangulations of \mathbf{A} which has a height function ω defined on vertices of \mathbf{A} . Let \mathcal{T}_1 and \mathcal{T}_2 be two triangulations of \mathbf{A} . We say that $\mathcal{T}_1 \leq_1 \mathcal{T}_2$ if \mathcal{T}_2 can be produced from \mathcal{T}_1 by a (possibly empty) sequence of up-flips. Likewise, we say $\mathcal{T}_1 \geq_1 \mathcal{T}_2$ if \mathcal{T}_2 can be produced from \mathcal{T}_1 by a (possibly empty) sequence of down-flips.

The \leq_1 relation defines a poset on the set of triangulations of \mathbf{A} . If \mathbf{A} is the vertex of a cyclic polytope, then this poset is exactly \mathbf{HST}_1 . What would be a poset of an arbitrary point set look like?

3.1 A motivation example

Figure 3 shows the poset of the point set \mathbf{A} of 6 points with a given height function ω . It is the smallest example which contains a non-regular triangulation.

This poset is produced using the up-flips. All sequence of up-flips are starting from the regular triangulation (bottom) to the farthest point regular triangulation (top) of (\mathbf{A}, ω) . In each triangulation of \mathbf{A} , the green edges are locally non-regular with respect to up-flips and flippable. The (thin) black edges are also locally non-regular, but not flippable. The thick black edges are locally regular.

Each arrow in this graph represents an up-flip, which is either a 2-2 or a 3-1 flip. The numbers in front of each arrow, like 34-56 or 123-4, indicate the involved vertices which support this flip. For example, 34-56 means a 2-2 flip which replaces edge e_{34} by e_{56} , and 123-4 means a 3-1 flip which removes the vertex v_4 from the triangle f_{123} . The four vertices of each flip are the vertices of a tetrahedron. Note that in this example, there is no 3-1 flip.

From this graph, we observe some interesting properties of this poset.

- There is a unique minimum, which is the regular triangulation of (\mathbf{A}, ω) , denoted as \mathcal{T}_{reg} .
- There is no unique maximum. Instead, there are two top elements, which are the farthest point regular triangulation of (\mathbf{A}, ω) , denoted as \mathcal{T}_{far} , and the non-regular triangulation of \mathbf{A} , denoted as \mathcal{T}_{nonreg} .
- There is no path (up-flips) from \mathcal{T}_{nonreg} to \mathcal{T}_{far} .

This graph could partially explain the behaviour of Lawson's flip algorithm [22]:

- Any sequence of down-flips (which are reversing the arrows) will reach \mathcal{T}_{reg} . Even it starts from \mathcal{T}_{nonreg} . This explains that the Lawson's algorithm with respect to the down-flips must terminate at the unique minimum, the regular triangulation of (\mathbf{A}, ω) .
- However, it may fail to reach the \mathcal{T}_{far} with respect to up-flips. There exists a sequence which ends at \mathcal{T}_{nonreg} .

We call a sequence of up-flips *maximal* if it starts from the \mathcal{T}_{reg} and ends at the \mathcal{T}_{far} of (\mathbf{A}, ω) . From 3 (4) we could observe there are many maximal sequences. In the next section, we will prove that every maximal sequence gives a tetrahedralisation of \mathbf{A}^ω . It is not hard to observe that several maximal sequences may lead to the same tetrahedralisation. They contain the same set of flips but with have different orders. By a filtering of maximal sequences with the same set of flips, we can get all six tetrahedralisations of \mathbf{A}^ω , shown in Figure 16 Top.

Note that the lengths of the six tetrahedralisations are either 4 or 5. This is due to 3d flips, 2-3 or 3-2 flips, between these tetrahedralisations. It could be indeed observed from this poset (see Figure 16 Bottom):

- If two maximal sequences have the same length, then there is no 3d flip between them.
- If two maximal sequences have two elements in common, and their lengths are exactly different by one, then there is a 3d flip between their corresponding tetrahedralisations.

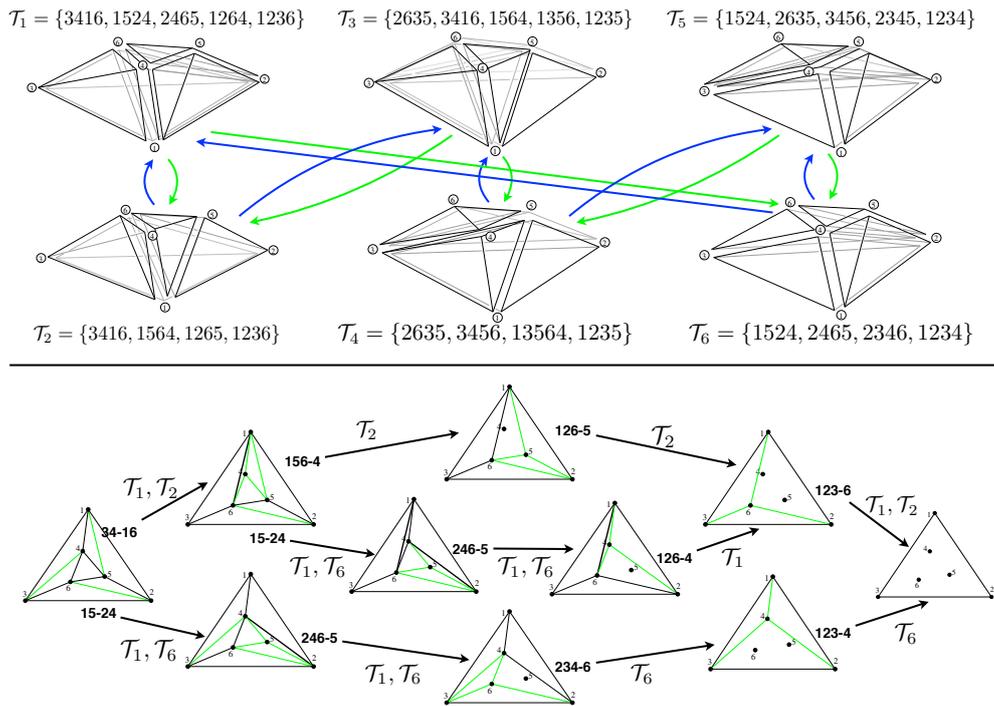


Figure 16: Top: the (indirected) flip graph of the six tetrahedralisations $\mathcal{T}_1, \dots, \mathcal{T}_6$ of \mathbf{A}^ω . A green arrow indicates a 3-2 flip and a blue arrow indicates a 2-3 flip, respectively. Bottom: The three sequences of flips corresponding to the tetrahedralisations $\mathcal{T}_1, \mathcal{T}_2$, and \mathcal{T}_6 are shown, respectively. The 3d flips between $\mathcal{T}_1, \mathcal{T}_2$ and $\mathcal{T}_1, \mathcal{T}_6$ can be observed.

By the above observation, we could “see” that the set of six tetrahedralisations is connected by 3d flips. In other words, the undirected flip graph of \mathbf{A}^ω is connected.

Finally, we discuss the existence of the non-regular triangulation \mathcal{T}_{nonreg} of \mathbf{A} . Since there is no path between \mathcal{T}_{nonreg} and \mathcal{T}_{far} , this shows that the 3d non-convex polyhedron with \mathcal{T}_{nonreg} and \mathcal{T}_{far} as its boundary is indecomposable with its own vertices. This links to the well-known Schönhardt polyhedron [30].

4 From Monotone Sequences of Directed Flips to Triangulations of Polyhedra

Recall a *triangulation* of a 3d point set is a 3d simplicial complex such that its underlying space is this convex hull of this point set. Let \mathbf{A} be a finite point set in \mathbb{R}^2 and let $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function defined on the vertices of \mathbf{A} . In this section, we prove the following theorem.

Theorem 5. *Any monotone sequence of directed flips which transforms from the regular triangulation into the farthest point regular triangulation of (\mathbf{A}, ω) (or vice versa) corresponds to a 3d triangulation of \mathbf{A}^ω .*

The theorem in one dimension lower is illustrated in Figure 14 Right. A nice example shows that a monotone sequence of up-flips between a set of 1-dimensional triangulations produces a triangulation

of 2-dimensional polygon.

4.1 Characteristic sections and lifted triangulations

We first introduce a very useful geometric connection between a triangulation in the plane and another one in \mathbb{R}^3 .

Let \mathcal{T} be a triangulation of \mathbf{A} in \mathbb{R}^2 , and let $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function defined on vertices of \mathbf{A} . We can define a piecewise-linear function by

$$g_{\omega, \mathcal{T}} : \text{conv}(\mathbf{A}) \rightarrow \mathbb{R}^3, \quad g_{\omega, \mathcal{T}}(\mathbf{p}) := \omega(\mathbf{p}), \quad \forall \mathbf{p} \in \mathbf{A},$$

and extended affinely on each simplex $\sigma \in \mathcal{T}$. This function is called the *characteristic section* of \mathcal{T} [6, Chap 5, Definition 5.2.12, page 229].

Geometrically, $g_{\omega, \mathcal{T}}$ is a piecewise flat surface in \mathbb{R}^3 whose canonical projection (which deletes the last coordinate of any point $\mathbf{x} \in g_{\omega, \mathcal{T}}$) is the planar domain $\text{conv}(\mathbf{A})$. Equivalently, it lifts the planar domain $\text{conv}(\mathbf{A})$ into a piecewise flat surface in \mathbb{R}^3 .

We define the *lifted triangulation* of \mathcal{T} via ω as a two-dimensional triangulation \mathcal{T}^ω embedded in \mathbb{R}^3 , such that:

- (1) The vertex set of \mathcal{T}^ω is the lifted vertex set of \mathcal{T} . It is a subset of \mathbf{A}^ω .
- (2) There is a bijection between the sets of simplices of \mathcal{T}^ω and \mathcal{T} , i.e., every simplex of \mathcal{T} is mapped uniquely to a simplex of \mathcal{T}^ω .
- (3) The union of all simplices of \mathcal{T}^ω is equal to the characteristic section $g_{\omega, \mathcal{T}}$.

It is easy to verify that the lifted triangulation \mathcal{T}^ω is a triangulation of the characteristic section $g_{\omega, \mathcal{T}}$. Figure 12 shows two planar triangulations and their corresponding lifted triangulations in \mathbb{R}^3 .

Every triangulation of \mathbf{A} corresponds to a lifted triangulation via ω . In particular, the regular and the farthest point regular triangulations of (\mathbf{A}, ω) correspond to the two extreme lifted triangulations of \mathbf{A} , denoted as \mathcal{R}^ω and \mathcal{F}^ω , which are the sets of lower and upper faces of $\text{conv}(\mathbf{A}^\omega)$, respectively. All other lifted triangulations of \mathbf{A} must lie between \mathcal{R}^ω and \mathcal{F}^ω .

We say that the characteristic section $g_{\omega, \mathcal{T}_1}$ *lies vertically below* $g_{\omega, \mathcal{T}_2}$ if for any two points $\mathbf{x}_1 = (x, y, z_1) \in g_{\omega, \mathcal{T}_1}$ and $\mathbf{x}_2 = (x, y, z_2) \in g_{\omega, \mathcal{T}_2}$, the inequality $z_1 \leq z_2$ holds. This means $g_{\omega, \mathcal{T}_1}$ is nowhere higher than $g_{\omega, \mathcal{T}_2}$.

With the concept of lifted triangulations, let us re-look at what a directed flip (up-flip or down-flip) in a triangulation \mathcal{T} of \mathbf{A} changes the characteristic sections in \mathbb{R}^3 . Let \mathcal{T}_1 be a triangulation of \mathbf{A} . An up-flip in \mathcal{T}_1 is equivalent to exchange the lower faces by upper faces of a tetrahedron $t_{\mathbf{a}'\mathbf{b}'\mathbf{c}'\mathbf{d}'}$ in $\text{conv}(\mathbf{A}^\omega)$, where $\text{conv}\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$ is the support of this flip. As a result, the lifted triangulation \mathcal{T}_1^ω is “deformed” into the lifted triangulation \mathcal{T}_2^ω which lies vertically above \mathcal{T}_1^ω .

This fact gives a 3d view of a monotone sequence of directed flips which transforms the regular into the farthest point regular triangulation of (\mathbf{A}, ω) . It “deforms” the \mathcal{R}^ω gradually into the \mathcal{F}^ω by passing through a sequence of lifted triangulations of \mathbf{A} ,

$$\mathcal{S} := \{\mathcal{R}^\omega = \mathcal{T}_0^\omega, \mathcal{T}_1^\omega \dots, \mathcal{T}_m^\omega = \mathcal{F}^\omega\}.$$

Every two adjacent lifted triangulations \mathcal{T}_1^ω and \mathcal{T}_2^ω in this sequence is connected by an up-flip in the planar triangulation. This sequence of lifted triangulations have the following properties:

- \mathcal{S} is ordered, i.e., $\mathcal{T}_i^\omega \leq_2 \mathcal{T}_{i+1}^\omega, \forall i = 0, \dots, m-1$.
- Let $\text{Vol}(\mathcal{R}^\omega, \mathcal{T}_i^\omega)$ denote the volume between \mathcal{R}^ω and \mathcal{T}_i^ω . Then

$$\text{Vol}(\mathcal{R}^\omega, \mathcal{T}_{i+1}^\omega) - \text{Vol}(\mathcal{R}^\omega, \mathcal{T}_i^\omega) = t_i,$$

where t_i is a tetrahedron in $\text{conv}(\mathbf{A}^\omega)$. The faces of t_i are exactly the triangles involved in the up-flip which transforms \mathcal{T}_i into \mathcal{T}_{i+1} . This shows that the set of lifted triangulation \mathcal{S} corresponds to a set of tetrahedra in $\text{conv}(\mathbf{A}^\omega)$,

$$T := \{t_0, t_1, \dots, t_{m-1}\}.$$

With the above definitions and facts, we can prove our theorem.

4.2 The proof

We prove the above theorem by showing that the set of tetrahedra T related to a monotone sequence of directed flips together with the faces of these tetrahedra fulfil all of the three properties of a triangulation, which are the (i) Closure Property, (ii) Intersection Property, and (iii) Union Property.

Proof. Let V, E, F be the set of all vertices, edges, and triangles of \mathcal{S} . Let $\mathcal{T}_\mathcal{S} = \{V, E, F, T\}$. We want to show that $\mathcal{T}_\mathcal{S}$ fulfils all the three properties of a triangulation, i.e., it is a triangulation of \mathbf{A}^ω .

First, every $t_i, i = 0, \dots, m-1$, connects two lifted triangulations \mathcal{T}_i^ω and \mathcal{T}_{i+1}^ω . The fact that the sequence of lifted triangulations sweeps through the volume of $\text{conv}(\mathbf{A}^\omega)$ implies that there is no hole left in the volume. Therefore $\cup T = \text{conv}(\mathbf{A}^\omega)$, i.e., the Union Property (iii) holds.

Second, it is obvious that the faces of every tetrahedron t_i are in F , and the edges of t_i are in E , and the vertices of t_i are in V . Moreover, E and V are the subsets of edges and vertices of faces of F . It is sufficient to show that every face in F belongs to at least one tetrahedron $t_i \in T$. By the general position assumption there is no identical triangle in \mathcal{R} and \mathcal{F} of (\mathbf{A}, ω) . Then every face $f \in F$ is either the lower or the upper face of at least one tetrahedron in T . These facts together ensure that the Closure Property (i) holds for $\mathcal{T}_\mathcal{S}$.

Last, the Intersection Property (ii) holds in $\mathcal{T}_\mathcal{S}$ by the fact that every two lifted triangulations are either disjoint or share at their common triangles in F .

This proves that $\mathcal{T}_\mathcal{S}$ is a triangulation of $\text{conv}(\mathbf{A}^\omega)$. □

Although the above theorem consider the biggest convex polytope whose boundaries are the \mathcal{R}^ω and \mathcal{F}^ω of \mathbf{A}^ω . Our proof shows that it holds for non-convex polyhedra as well, as long as the boundary of a polyhedron can be represented by two planar triangulations which are canonical projections of their lifted triangulations, respectively.

Corollary 6. *Let \mathcal{T}_u and \mathcal{T}_v be two triangulations of \mathbf{A} in the plane. Assume the lifted triangulation \mathcal{T}_u^ω lies strictly lower than \mathcal{T}_v^ω . Let P_{uv} be the 3d polyhedron whose boundary is $\mathcal{T}_u^\omega \cup \mathcal{T}_v^\omega$, then any monotone sequence of directed flips which transform \mathcal{T}_u into \mathcal{T}_v corresponds to a triangulation of P_{uv} .*

Based on this fact, we present a triangulation algorithm for a special class of 3d non-convex polyhedra in Section 7.

5 From Triangulations of Polyhedra to Monotone Sequences of Directed Flips

In this section, we consider the reverse of the Theorem 5, which is the following **question**: *given a triangulation of a 3d polyhedron, does it corresponds to a monotone sequence of directed flips?*

So far, it is only proven for a special case when \mathbf{A} is the vertex set of a cyclic polytope [6, Section 6.1.5]. The essential part of its proof is to show that between any two triangulations of \mathbf{A} in \mathbb{R}^d , there exists a triangulation of the polyhedron in \mathbb{R}^{d+1} bounded by these two lifted triangulations without new vertex. It is achieved by an explicit (very technical) construction of such triangulation in \mathbb{R}^{d+1} . However, there are difficulties to directly generalise this result to an arbitrary point set in \mathbb{R}^2 . The main difficulty is that there are 3d non-convex polyhedron which can not be tetrahedralised without additional vertices, such as the Schönhardt polyhedron [30].

Let \mathbf{A} be a finite point set in \mathbb{R}^2 , and ω be a height function defined on the vertices of \mathbf{A} . Let \mathcal{T}_u and \mathcal{T}_v be two triangulations of \mathbf{A} . Without loss of generality, assume the lifted triangulation \mathcal{T}_u^ω lies vertically below \mathcal{T}_v^ω , and assume that the characteristic sections of \mathcal{T}_u^ω and \mathcal{T}_v^ω only intersect at their boundary. Let P_{uv} be the 3d polyhedron whose lower faces are \mathcal{T}_u^ω and upper faces are \mathcal{T}_v^ω . Assume a triangulation \mathcal{T}_{uv} of P_{uv} is given. Moreover, the vertex set of \mathcal{T}_{uv} is exactly the vertex set of P_{uv} , i.e., \mathcal{T}_{uv} contains no additional vertex. We describe an algorithm to find a monotone sequence of directed flips between \mathcal{T}_u and \mathcal{T}_v .

5.1 A tetrahedra-driven flipping algorithm

Notations We will simultaneously work in two and three dimensions. We will relate a vertex in \mathbb{R}^3 and its canonical projection in a plane at the same time. Since the vertices in \mathcal{T}_{uv} are in one-to-one correspondence with the vertices in \mathcal{T}_u and \mathcal{T}_v , to simplify the notations, we do not distinguish them between \mathbb{R}^2 and \mathbb{R}^3 . For example, we denote t_{abcd} be a tetrahedron in \mathcal{T}_{uv} with vertices $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d} \in \mathbb{R}^3$. At the same time, we also denote f_{abc} a triangle in \mathcal{T}_u and \mathcal{T}_v with vertices $\mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^2$.

Removability and flippability We assume that the vertex set \mathbf{A} is in *general position* in the plane which means, no three vertices of \mathbf{A} lie on a common line. So the only flips we need are: 2-2, 1-3, and 3-1 flip.

The idea of our algorithm is to incrementally remove all tetrahedra of \mathcal{T}_{uv} , one at a time. Every tetrahedron in \mathcal{T}_{uv} corresponds a directed flip, which changes its lower and upper faces (or vice versa). However, not every tetrahedron is “removable” at any time. Let t_{abcd} be a tetrahedron in \mathcal{T}_{uv} . Let \mathcal{T} be a triangulation produced by the corresponding sequence of directed flips. We say that a face f_{abc} of t_{abcd} is *exposed* in \mathcal{T} if f_{abc} is a triangle in \mathcal{T} . We say that a tetrahedron $t_{abcd} \in \mathcal{T}_{uv}$ is *removable* (from \mathcal{T}_{uv}) if it in one of the following three cases.

- (1) One face of t_{abcd} is exposed in \mathcal{T} . Let the face be f_{abc} . If \mathbf{d} lies in the interior of the triangle f_{abc} , then t_{abcd} is *removable*. It corresponds to a 1-3 flip which inserts the vertex \mathbf{d} into \mathcal{T} , see Figure 17 (1).
- (2) Two faces of t_{abcd} are exposed in \mathcal{T} . Let the two faces be f_{abc} and f_{abd} , respectively. If the vertices $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ form a strictly convex quadrilateral in \mathcal{T} , then t_{abcd} is *removable*. It corresponds to a 2-2 flip which replaces edge e_{ab} by e_{cd} in \mathcal{T} , see Figure 17 (2).

- (3) Three faces of t_{abcd} are exposed in \mathcal{T} . Let the three faces be f_{abd} , f_{bcd} , and f_{cad} , then t_{abcd} is *removable*. It corresponds to a 3-1 flip which removes the vertex d from \mathcal{T} , see Figure 17 (3).

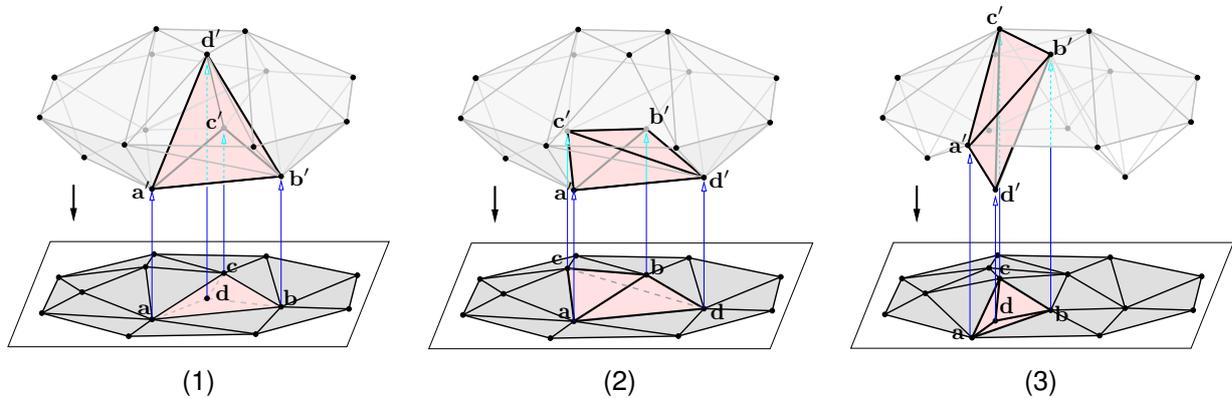


Figure 17: A removable tetrahedron t_{abcd} . (1) The face f_{abc} is exposed in \mathcal{T} , t_{abcd} corresponds to a 1-3 flip. (2) Two faces f_{abc} and f_{abd} are exposed in \mathcal{T} , t_{abcd} corresponds to a 2-2 flip. (3) Three faces f_{abd} , f_{bcd} , f_{cad} are exposed in \mathcal{T} , t_{abcd} corresponds to a 3-1 flip.

The cases when t_{abcd} is not removable will be discussed in the analysis of the algorithm in the next subsection.

The algorithm The tetrahedron-driven flip algorithm is described in Figure 18. It takes a tetrahedralisation \mathcal{T}_{uv} and a triangulation \mathcal{T} (\mathcal{T} is either \mathcal{T}_u or \mathcal{T}_v). This algorithm initialises two lists, \mathcal{Q} , which contains all tetrahedra in \mathcal{T}_{uv} , and \mathcal{L} , which is empty at the beginning. Then it runs into a loop until \mathcal{Q} is empty. At each iteration, it searches a removable tetrahedron $t_{abcd} \in \mathcal{Q}$. If such a tetrahedron exists, the triangulation \mathcal{T} is updated by performing the corresponding flip induced by t_{abcd} . \mathcal{Q} is reduced by one tetrahedron and it is added to \mathcal{L} .

```

Input:   $\mathcal{T}_{uv}$  and  $\mathcal{T}_u$ ;
Output:  $\mathcal{T}_v$  and  $\mathcal{L}$  (a sorted array of tetrahedra);
1  Let  $\mathcal{T} := \mathcal{T}_u$ ,  $\mathcal{Q} := \{t_{abcd} \mid t_{abcd} \in \mathcal{T}_{uv}\}$ ,  $\mathcal{L} := \emptyset$ ;
2  while  $\mathcal{Q} \neq \emptyset$  do
3    if a removable  $t_{abcd}$  exists then
4      Update  $\mathcal{Q} := \mathcal{Q} \setminus \{t_{abcd}\}$ ;
5      Perform the corresponding flip in  $\mathcal{T}$ ;
6      Update  $\mathcal{L} := \mathcal{L} \cup \{t_{abcd}\}$ ;
7    endif
8  endwhile

```

Figure 18: A tetrahedralisation-driven flip algorithm.

If this algorithm terminates, it outputs a list \mathcal{L} of sorted tetrahedra of \mathcal{T}_{uv} . Since each tetrahedron in \mathcal{L} corresponds to a directed flip, it is also a monotone sequence of directed flips which transforms \mathcal{T}_u into \mathcal{T}_v .

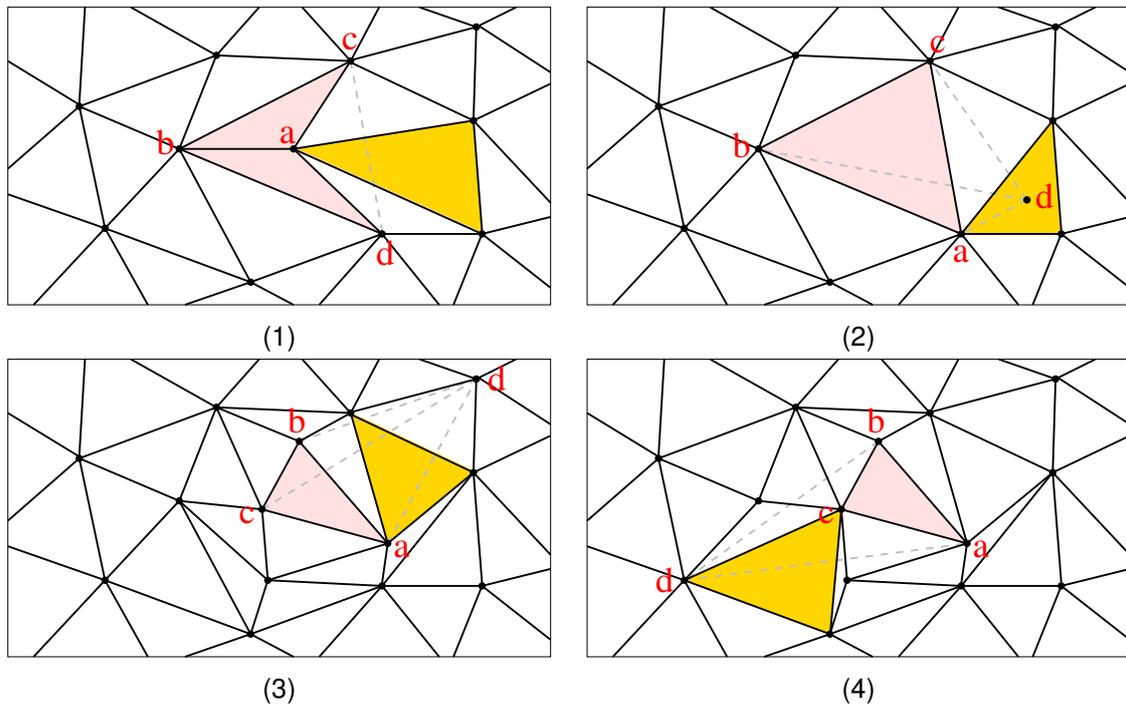


Figure 19: All cases when a tetrahedron t_{abcd} is not removable. In (1) neither a 2-2 flip nor a 3-1 flip is possible on edge e_{ab} . In (2), a 1-3 flip is not possible since the new vertex d does not inside the triangle f_{abc} . In (3) and (4) a 1-3 flip is not possible since the vertex d already exists in \mathcal{T} .

5.2 Analysis of termination

The crucial question in this algorithm is: does there exist a removable tetrahedron in \mathcal{T}_{uv} (line 3)? If all tetrahedra are not removable, then it gets stuck.

We will use the *in-front/behind* relation [9] to relate two tetrahedra in \mathbb{R}^3 . Without loss of generality, we assume that the lifted triangulation \mathcal{T}_u^ω lies vertically below \mathcal{T}_v^ω . We use \mathcal{T}_u as the input, so \mathcal{T}_v is the target (output). By placing our viewpoint at $\mathbf{x} = (0, 0, -\infty)$, i.e., we're looking from bottom to top. We say that a tetrahedron t_{pqrs} is *in-front* of a tetrahedron t_{abcd} , denoted as $t_{pqrs} \prec t_{abcd}$, if there exists a ray starting from \mathbf{x} that intersects first t_{pqrs} then t_{abcd} . It is the same to say that t_{abcd} is *behind* t_{pqrs} , denoted as $t_{abcd} \succ t_{pqrs}$.

Our algorithm initialises the planar triangulation $\mathcal{T} := \mathcal{T}_u$. At each time, this algorithm searches a removable tetrahedron $t_{abcd} \in \mathcal{T}_{uv}$. The search starts from an arbitrary tetrahedron $t_{abcd} \in \mathcal{T}_{uv}$ such that it has at least one exposed face $f_{abc} \in \mathcal{T}$. If t_{abcd} is removable, then the algorithm continues. Assume t_{abcd} is not removable. We show that it is always possible to find another tetrahedron $t_{pqrs} \in \mathcal{T}_{uv}$ such that it is in front of t_{abcd} , i.e., $t_{pqrs} \prec t_{abcd}$. There are following cases to be considered.

- (i) Two faces of t_{abcd} are exposed in \mathcal{T} . Without loss of generality, let the two faces are f_{abc} and f_{abd} which share the common edge $e_{ab} \in \mathcal{T}$. Since t_{abcd} is not removable, then a, b, c, d form a non-convex quadrilateral in the plane. Let a be the non-convex vertex of this quadrilateral. Then the face f_{cda} of t_{abcd} is missing in \mathcal{T} , see Figure 19 (1). Then there must exist another triangle $f_{pqr} \in \mathcal{T}$, such that $f_{pqr} \notin \{f_{abc}, f_{abd}\}$ and $f_{pqr} \cap f_{cda} \neq \emptyset$. Hence the tetrahedron $t_{pqrs} \prec t_{abcd}$.
- (ii) Only the face f_{abc} of t_{abcd} is exposed in \mathcal{T} , and $d \notin \mathcal{T}$. Since t_{abcd} is not removable, then d lies in a triangle $f_{pqr} \in \mathcal{T}$, such that $f_{pqr} \neq f_{abc}$, see Figure 19 (2). Hence the tetrahedron

$$t_{\mathbf{pqrs}} \prec t_{\mathbf{abcd}}.$$

(iii) Only the face $f_{\mathbf{abc}}$ of $t_{\mathbf{abcd}}$ is exposed in \mathcal{T} , and $\mathbf{d} \in \mathcal{T}$. Then we can show that the face $f_{\mathbf{abd}} \notin \mathcal{T}$. We distinguish two cases:

- (a) The edge $e_{\mathbf{cd}}$ intersects $e_{\mathbf{ab}}$, see Figure 19 (3). If $f_{\mathbf{abd}}$ exists, it must belong to another tetrahedron $t_{\mathbf{abde}}$, where $\mathbf{e} \neq \mathbf{c}$. Then either the edge $e_{\mathbf{ae}}$ or $e_{\mathbf{be}}$ must go through the interior of $t_{\mathbf{abcd}}$ – a contradiction that \mathcal{T}_{uv} is a tetrahedralisation. Then there must exist a face $f_{\mathbf{pqr}} \in \mathcal{T}$, and $f_{\mathbf{pqr}} \cap f_{\mathbf{cda}} \neq \emptyset$, see Figure 19 (3).
- (b) The edge $e_{\mathbf{cd}}$ does not intersect $e_{\mathbf{ab}}$, see Figure 19 (4). Then the vertex \mathbf{c} blocks the appearance of $f_{\mathbf{abd}}$. Therefore any face except $f_{\mathbf{abc}}$ in \mathcal{T} which contains \mathbf{c} as a vertex can be the face $f_{\mathbf{pqr}} \in \mathcal{T}$, and $f_{\mathbf{pqr}} \cap f_{\mathbf{cda}} \neq \emptyset$, see Figure 19 (4).

In both cases, there exists a tetrahedron $t_{\mathbf{pqrs}} \prec t_{\mathbf{abcd}}$.

From the above analysis, we see that starting from an arbitrary tetrahedron $t_{\mathbf{abcd}} \in \mathcal{T}_{uv}$, it is either a removable tetrahedron, or there exists another tetrahedron $t_{\mathbf{pqrs}} \in \mathcal{T}_{uv}$ and $t_{\mathbf{pqrs}} \prec t_{\mathbf{abcd}}$. This process can be repeated. This generates a sequence of tetrahedra, which are:

$$t_{\mathbf{abcd}} = t_{\mathbf{a_0b_0c_0d_0}} \succ t_{\mathbf{a_1b_1c_1d_1}} \succ \dots \succ t_{\mathbf{a_mb_mc_md_m}} = t_{\mathbf{pqrs}}$$

such that each tetrahedron $t_{\mathbf{a_ib_ic_id_i}} \in \mathcal{T}_{uv}$, $i = 0, \dots, m$, is not removable. Therefore this algorithm will terminate if the above sequence does not form a cycle, i.e., $t_{\mathbf{a_mb_mc_md_m}} = t_{\mathbf{a_0b_0c_0d_0}}$.

Theorem 7. *This algorithm terminates as long as \mathcal{T}_{uv} contains no cycle of tetrahedra from the viewpoint $\mathbf{x} = (0, 0, -\infty)$ (or $\mathbf{x} = (0, 0, +\infty)$).*

By the Acyclic Theorem [9], a regular triangulation contains no cycle of faces from any viewpoint. An immediate result is that this algorithm will terminate if \mathcal{T}_{uv} is a regular tetrahedralisation.

Our analysis shows that only cycles of tetrahedra from the given viewpoint \mathbf{x} will cause problem. More generally, we can consider a parallel projection of \mathcal{T}_{uv} through any line direction L (a one-dimensional affine subspace) in \mathbb{R}^3 . We will also get two planar triangulations $\mathcal{T}_{u,L}$ and $\mathcal{T}_{v,L}$ in two planes, one lies in front of \mathcal{T}_{uv} , and one lies behind \mathcal{T}_{uv} from a viewpoint on L . We can re-define the height function ω on vertices of $\mathcal{T}_{u,L}$ and $\mathcal{T}_{v,L}$ with respect to L as following: Let $\mathbf{p} = (p_x, p_y, p_z) \in \mathbb{R}^3$ be a vertex in $\mathcal{T}_{u,L}$ or $\mathcal{T}_{v,L}$, we project the vector (p_x, p_y, p_z) onto the line L and let the $\omega(\mathbf{p})$ be the length of the projected vector. Our algorithm will transform $\mathcal{T}_{u,L}$ into $\mathcal{T}_{v,L}$ or vice versa by a sequence of monotone flips as long as there is no cycle in \mathcal{T}_{uv} viewed by any viewpoint on L , an example is shown in Figure 21.

We thus have the following theorem about the termination of this algorithm.

Corollary 8. *This algorithm terminates as long as the 3d triangulation \mathcal{T}_{uv} of \mathbf{A}^ω in \mathbb{R}^3 contains no cycle of tetrahedra with respect to a viewpoint on the direction of the height function ω .*

We also see another nice fact of the 3d triangulations produced by monotone sequence of directed flips.

Corollary 9. *Let \mathcal{T} be a 3d triangulation of a polyhedron P in \mathbb{R}^3 . If \mathcal{T} corresponds to a monotone sequence of directed flips between two triangulations of a point set \mathbf{A} in \mathbb{R}^2 with a height function $\omega : \mathbf{A} \rightarrow \mathbb{R}$, then the set of tetrahedra of \mathcal{T} is acyclic with respect to a viewpoint on the direction of the height function ω .*

Note that the 3d triangulation \mathcal{T} in the above Corollary is not necessarily a regular triangulation.

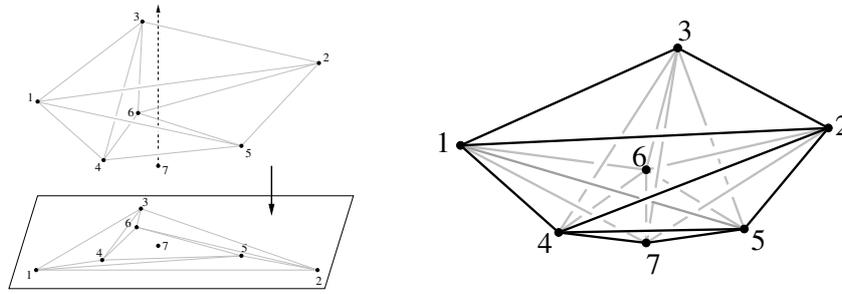


Figure 20: Left: The point set of this tetrahedralisation contains the point set of a Schönhardt polyhedron with an additional vertex lie outside this polyhedron and is visible by all vertices of it from the interior of this polyhedron. Right: the tetrahedralisation of the left point set consists of 10 tetrahedra.

5.3 A cycle consists of three tetrahedra

Our analysis of this algorithm shows, \mathcal{T}_{uv} may be non-regular and the algorithm still terminates as long as there is no cycle from the chosen viewpoint which produces \mathcal{T}_u and \mathcal{T}_v . It is thus interesting to study and understand such (non-regular) tetrahedralisations containing cycles (from some viewpoints).

The simplest non-regular tetrahedralisation contains 7 vertices. It is first appear in [19]. It is the simplest tetrahedralisation which will cause the 3d Lawson's flip algorithm fails [19]. A construction of this example is given in the book [6, Chap 3, Example 3.6.15, page 139].

The basic fact is that this tetrahedralisation contains the boundary triangulation of the Schönhardt polyhedron. It is necessary to contain an additional vertex which does not belong to the Schönhardt polyhedron. In particular, this additional vertex is visible by all vertices of the Schönhardt polyhedron through the bottom face, see Figure 20 Left. This property ensures that a tetrahedralisation containing this Schönhardt polyhedron exists. Figure 20 Right shows the the tetrahedralisation, denoted as \mathcal{T}_{uv} . There are 10 tetrahedra in \mathcal{T}_{uv} . Three tetrahedra are highlighted by different color. They form a cycle if the viewpoint is chosen along the z axis. This shows that \mathcal{T}_{uv} is a non-regular triangulation. However, these tetrahedra do not contain cycle if the viewpoint is chosen along the x axis, see Figure 21.

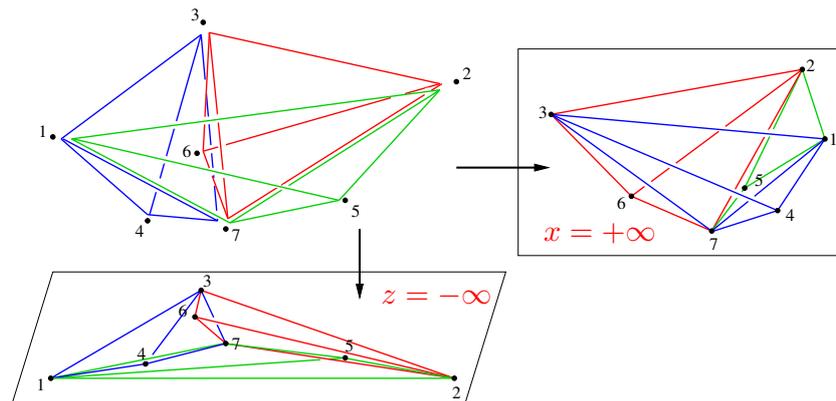


Figure 21: Two orthogonal projections of the three tetrahedra in Figure 20. In particular, the projection along the z -direction contains a cycle $t_{1374} \prec_z t_{1275} \prec_z t_{2376} \prec_z t_{1374}$, while along the x -direction does not.

Let $\mathcal{T}_{u,z}$ and $\mathcal{T}_{v,z}$ be the two traingulaitons obtained by the orthogonal projections of \mathcal{T}_{uv} along the z -axis. Likewise, let $\mathcal{T}_{u,x}$ and $\mathcal{T}_{v,x}$ be the two traingulaitons obtained by the orthogonal projections of

\mathcal{T}_{uv} along the x -axis, see Figure 22. Then the flip algorithm will fail to transform $\mathcal{T}_{u,z}$ to $\mathcal{T}_{v,z}$, or vice versa. However, it will succeed to transform $\mathcal{T}_{u,x}$ to $\mathcal{T}_{v,x}$, or vice versa.

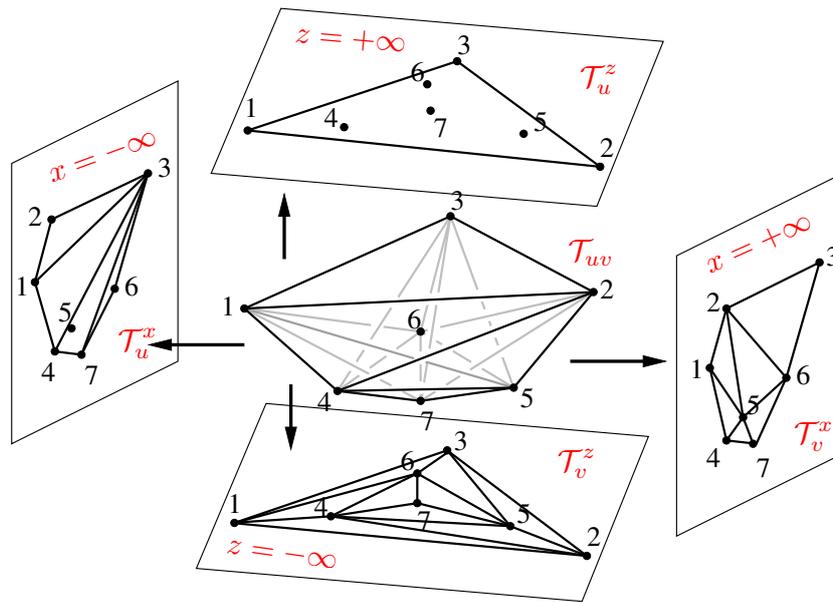


Figure 22: Triangulations obtained from different directions of the parallel projection of the same tetrahedralisation \mathcal{T}_{uv} .

From this example, we see that the existence of a viewpoint which “sees” no cycle of tetrahedra is a one of a key properties of a 3d triangulation of a point set. It is necessary to study the following question.

Question 10. *If \mathcal{T}_{uv} contains no interior point, is it always possible to find a viewpoint which produces two triangulations \mathcal{T}_u and \mathcal{T}_v such that the algorithm will succeed?*

6 Properties of the Directed Flip Graph

Since triangulations of \mathbf{A} can be ordered by the directed flips, we obtain a directed flip graph (poset) defined on the set of all triangulations of \mathbf{A} . As we have already shown in the motivation example in Section 3, it encodes many useful information of this point set. It is important to understand the structure of this poset. However, as we could observe that the size of this poset grows exponentially as the number of vertices grows. Figure 23 shows partially of this poset for the point set of 7 vertices for the point set shown in Figure 20.

6.1 General properties of this poset

In this section, some (obvious) properties of this poset are proven. Let \mathbf{A} be a finite point set in \mathbb{R}^2 and let $\omega : \mathbf{A} \rightarrow \mathbb{R}$ be a height function. Let $\mathcal{G}(\mathbf{A}, \omega)$ be the directed flip graph (poset) of all triangulations of \mathbf{A} ordered by one of the directed flips.

Definitions A *path* between two nodes in this poset represents a monotone sequence of directed flips between two triangulations of \mathbf{A} . Call a path *maximum* if it starts from the regular and ends at the

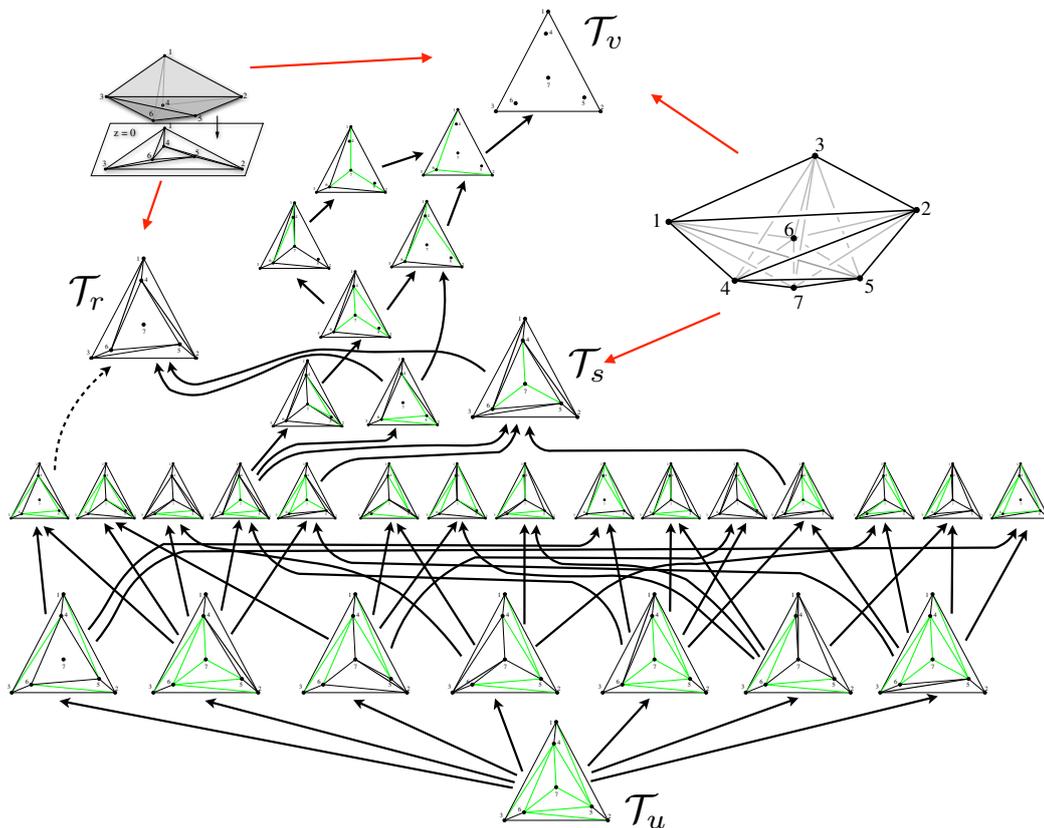


Figure 23: A subset of the directed flip graph of a set of 7 vertices.

farthest point regular triangulation of (\mathbf{A}, ω) , or vice versa. A node of this poset is called an *external node* if it can only be either the start or end of a path, otherwise, it is an *internal node*.

Theorem 11.

- (1) Every path between two nodes in this poset corresponds to a triangulation of the 3d polyhedron whose boundary is the union of the two triangulations.
- (2) Every maximum path in this poset corresponds to a triangulation of \mathbf{A}^ω . Moreover, this triangulation may be non-regular.
- (3) There exists non-regular triangulation of \mathbf{A}^ω which does not correspond to any maximum path in this poset.
- (4) This poset is in general not bounded.
- (5) Except the regular and the farthest point regular triangulations of (\mathbf{A}, ω) , all other external nodes in this poset are non-regular triangulations of \mathbf{A} .
- (6) Internal nodes of of this poset may be non-regular triangulations of \mathbf{A} .

Proof. Case (1) is exactly proven by the Corollary 6. Note that the 3d polyhedron is not necessarily convex.

The first part of (2) is proven by Theorem 5. We prove the second part of (2) by showing an example. Figure 22 shows such a 3d triangulation. First of all, there exists a monotone sequence of flips from

$\mathcal{T}_{u,x}$ to $\mathcal{T}_{v,x}$ which produces this triangulation. Hence a poset produced by viewpoints on the x -axis will contain this triangulation. However, the fact that there exists no monotone sequence of flips from $\mathcal{T}_{u,z}$ to $\mathcal{T}_{v,z}$ shows that it contains a cycle from viewpoints on the z -axis. By the Acyclic Theorem [9] it must be non-regular. The same example also proves (3).

Next we show that (2) and (3) together imply (4), i.e., this poset is not bounded. This can be proven by showing two examples (therefore there are at least two reasons). Both examples are shown in Figure 23. The first example is that there exists no path of up-flips (as well as down-flips) from the non-regular triangulation \mathcal{T}_r to the further point regular triangulation \mathcal{T}_v . In this case, there exists no 3d triangulation between these two triangulations. The second example is that there exists no path of up-flips (as well as down-flips) from the non-regular triangulation \mathcal{T}_s to the further point regular triangulation \mathcal{T}_v . In this case, although there exists a 3d triangulation between these two triangulations, but there is no monotone path between them.

To prove (5) we use the theory of secondary polytopes. By Theorem 3, between every two regular triangulations of \mathbf{A} there must exist a monotone sequence of directed flips. The negative of this Theorem shows that if there exists no monotone sequence of directed flips between two triangulations of \mathbf{A} , then at least one of the triangulations is non-regular. This shows that all external nodes of this poset, except the regular and farthest-point regular triangulations of (\mathbf{A}, ω) , are non-regular.

Finally, we construct an example to prove (6). The point set of this example contains 8 vertices, denoted as \mathbf{A} . The constructed tetrahedralisation \mathcal{T} has 15 tetrahedra. They are shown in Figure 24. It is necessary to have the 8th vertex which makes this example work. Figure 25 shows one of the maximum chains of the poset of this point set, in particular, it starts from the regular triangulation of \mathbf{A} to the farthest-point regular triangulation of \mathbf{A} . The sequence of flips are ordered by the tetrahedra as we have ordered in the list shown in the Figure 24.

One can show that in this chain, the internal nodes (3), (4), ..., (9) are non-regular triangulations. We use the negative of the Acyclic Theorem [9] to show that in each of these triangulations, there exists a cycle of triangles with respect to the in-front/behind relation viewed from the vertex 7. Figure 26 shows two of these cycles. \square

6.2 When ω is convex or concave

In this section we prove a special property of this poset when the height function ω is a convex or concave. From this property we can derive the exact conditions for when the Lawson's flip algorithm will terminate.

Theorem 12.

- (1) *If ω is convex, then the poset of down-flips has a unique minimum, which is the regular triangulation of (\mathbf{A}, ω) . However, it might have no unique maximum.*
- (2) *If ω is concave, then the poset of up-flips has a unique maximum, which is the farthest point regular triangulation of (\mathbf{A}, ω) . However, it might have no unique minimum.*

Proof. Note that (2) is just the reverse of (1). We only need to prove (1).

ω is convex implies that all lifted points of \mathbf{A} lies on the lower envelop of the convex hull of \mathbf{A}^ω in \mathbb{R}^3 . Then the vertex set of the regular triangulation of (\mathbf{A}, ω) is equal to \mathbf{A} . The vertex set of any other triangulation of \mathbf{A} must be either a subset of or equal to \mathbf{A} . This means that all monotone sequences

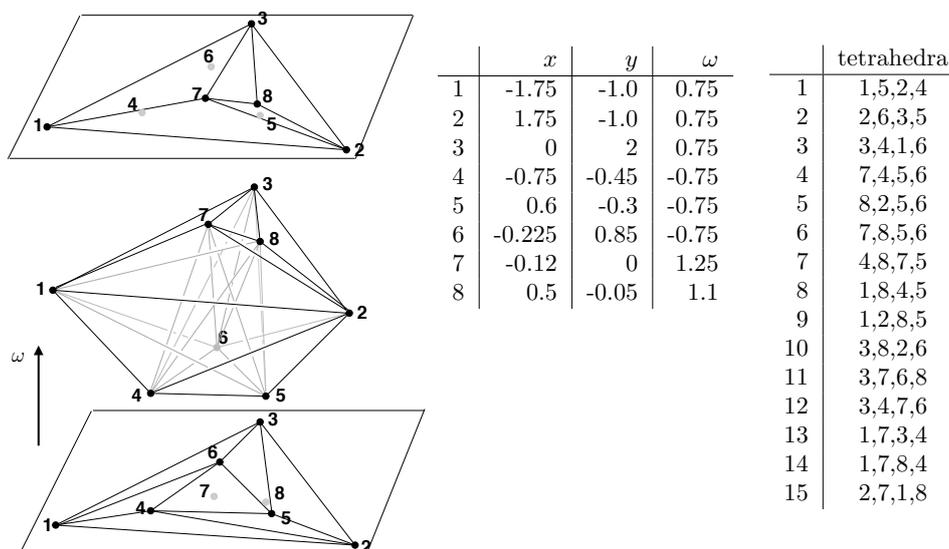


Figure 24: A point set \mathbf{A} of 8 points in \mathbb{R}^2 with a height function ω . Left: The regular and the farthest-point regular triangulation of (\mathbf{A}, ω) , and a tetrahedralisation \mathcal{T} of the convex hull of \mathbf{A}^ω . Right: Tables of coordinates of the 8 vertices of \mathbf{A} and the definition of ω , and the list of the 15 tetrahedra of \mathcal{T} .

of down-flips are either 2-2 flips (edge flips) or 1-3 flips (vertex insertions). There is no need of 3-1 flips (vertex deletions).

In the following, we show that between any triangulation \mathcal{T} of \mathbf{A} and the regular triangulation of (\mathbf{A}, ω) , denoted as \mathcal{R} , there always exists a sequence of down-flips. The proof is essentially the same as proving the termination of Lawson’s flip algorithm for producing the Delaunay triangulation of \mathbf{A} . In addition to using the 2-2 flips (edge flips), we need to use 1-3 flip (vertex insertions) as well.

Assume $\mathcal{T} \neq \mathcal{R}$. Call an edge $e_{ab} \in \mathcal{T}$ *locally non-regular with respect to down-flips* if it is shared by two triangles $f_{abc}, f_{abd} \in \mathcal{T}$, and the lifted point \mathbf{d}' lies vertically below the plane containing the lifted points \mathbf{a}' , \mathbf{b}' , \mathbf{c}' . In other words, the lifted edge $e_{a'b'}$ is locally non-convex between the lifted triangles $f_{a'b'c'}$ and $f_{a'b'd'}$ when viewed from below. Then e_{ab} must admit a 2-2 flip. Otherwise, at least one of the lifted points, either \mathbf{a}' or \mathbf{b}' does not lie on the lower convex hull of \mathbf{A}^ω which implies that ω is non-convex – a contradiction. Applying a 2-2 flip to replace e_{ab} by e_{cd} is the same as attaching a tetrahedron $t_{a'b'c'd'}$ below the edge $e_{a'b'}$. This means that the new edge $e_{cd} \in \mathcal{T}$ is locally regular. Whenever there exists a locally non-regular edge in \mathcal{T} we can flip it and replace it by a locally regular edge of \mathcal{T} , this process is irreversible, since it never creates locally non-regular edges.

Assume $\mathcal{T} \neq \mathcal{R}$ and all edges in \mathcal{T} are locally regular. Then there must exist a vertex in \mathcal{R} but not yet in \mathcal{T} . Let $\mathbf{p} \in \mathcal{R}$ and $\mathbf{p} \notin \mathcal{T}$. There must exist a triangle $f_{abc} \in \mathcal{T}$ such that $\mathbf{p} \in f_{abc}$. Assume no three points in \mathbf{A} are collinear, then f_{abc} and \mathbf{p} admit a 1-3 flip (vertex insertion) in \mathcal{T} . After \mathbf{p} is inserted, apply 2-2 flips if there are locally non-regular edges in \mathcal{T} .

The above process creates no cycles. This means that every triangulation of \mathbf{A} can be eventually transformed into the regular triangulation of (\mathbf{A}, ω) by a sequence of down-flips. This shows that this poset is bounded below, i.e., it has a unique minimum.

The fact that it might have no unique maximum in (1) can be proven by showing an example such that there exists no monotone sequence of up-flips between a triangulation of \mathbf{A} and the farthest point regular triangulation of \mathbf{A} . Such examples are shown in 23. There is no monotone sequence of flips between the triangulation \mathcal{T}_r and \mathcal{T}_v , and between \mathcal{T}_s and \mathcal{T}_v . □

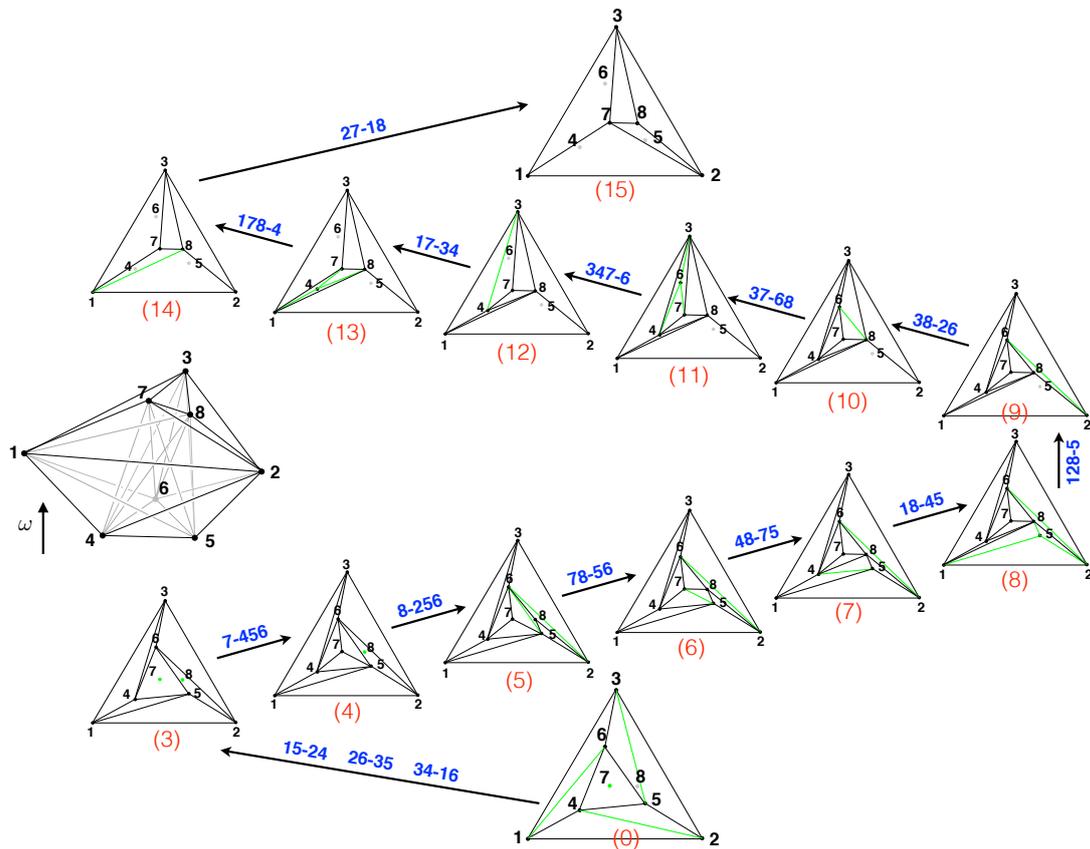


Figure 25: A maximum chain of flips from the regular to the farthest-point regular triangulation of (\mathbf{A}, ω) . The triangulations from (3) to (9) are non-regular triangulations of \mathbf{A} .

From the above theorem, we can state precise conditions when Lawson’s flip algorithm will terminate.

Corollary 13. *The Lawson’s flip algorithm terminates in one of the following cases.*

- (1) ω is convex and the directed flips are down-flips.
- (2) ω is concave and the directed flips are up-flips.

From the above theorem, we see that the condition that there is no 3-1 flip (vertex deletion) makes the poset bounded below or above. We remark that there are point sets which automatically satisfy this condition, such as the point set of a convex n -gon. In this case, no matter what property ω has, Lawson’s flip algorithm will terminate. This provides another proof of the Theorem 4 of the properties of point sets of cyclic polytopes in \mathbb{R}^2 .

6.3 Characterisation of external nodes of the poset

When ω is neither convex nor concave, starting from an arbitrary triangulation of \mathbf{A} , a monotone sequence of directed flips (up- or down-flips) might end at a non-regular triangulation of \mathbf{A} which is an external node of this poset. Besides the non-regular property, we show another property of these external nodes (triangulations).

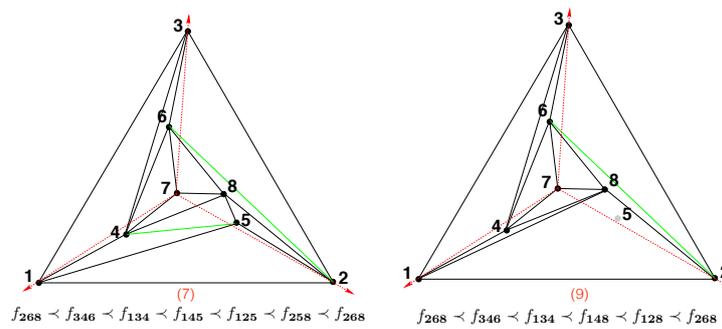


Figure 26: Two non-regular triangulations, (7) and (9). The cycles of triangles in each of them are shown. They are all viewed from the vertex 7.

Definitions We assume a given poset is produced by one of the two directed flips (up- and down-flips). We call the minimum and maximum nodes of this poset *lower extreme* and *upper extreme* nodes of this poset, respectively. External nodes which are not extreme nodes are called *non-extreme* nodes of this poset. We further distinguish two types of the non-extreme nodes. We call a non-extreme node *upper non-extreme* if there exists no directed flip in this triangulation towards the triangulation of the upper extreme node, and call it *lower non-extreme* if there exists no directed flip from this triangulation towards the triangulation of the lower extreme node. Figure 27 Left illustrates these definitions.

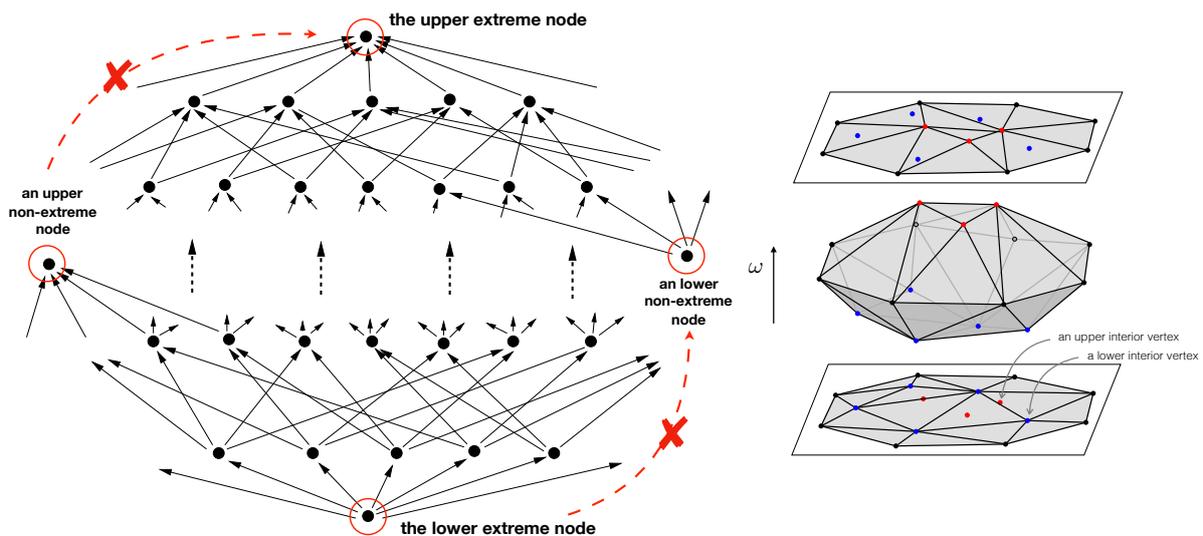


Figure 27: Left: Assume ω is neither convex nor concave. An example of a poset of triangulations of (\mathbf{A}, ω) . Extreme nodes, upper non-extreme nodes, and lower non-extreme nodes are shown. Right: The lower and upper interior vertices of a point set (\mathbf{A}, ω) .

A vertex of \mathbf{A} is an *interior vertex* if it is not on the convex hull of \mathbf{A} . We distinguish two types of the interior vertices of \mathbf{A} . We call an interior vertex $\mathbf{a} \in \mathbf{A}$ *upper interior* if the lifted point \mathbf{a}' appears in the triangulation of the upper extreme node, and call it *lower interior* if \mathbf{a}' appears in the triangulation of the lower extreme node, see Figure 27 Right for examples.

Redundant interior vertices By Theorem 12, when the directed flips are up-flips, if a triangulation contains no interior vertices of \mathbf{A} , then it must can be transformed into the farthest point regular triangulation of (\mathbf{A}, ω) since there is no 3-1 flip (vertex deletion) needed. Hence this triangulation

must not be an upper extreme node. We then can state the following property of all non-extreme nodes of this poset.

- (1) Any triangulation corresponds to an upper non-extreme node in this poset must contain lower interior vertices of \mathbf{A} .
- (2) Any triangulation corresponds to a lower non-extreme node in this poset must contain upper interior vertices of \mathbf{A} .

Any triangulation of a non-extreme node has no directed flip towards one of the two regular triangulations of (\mathbf{A}, ω) . We see the reason of the failure of Lawson's flip algorithm – there exist interior vertices in a triangulation which cannot be removed by the corresponding directed flips. These vertices, either lower or upper interior vertices, in (1) and (2) are called *redundant interior vertices*, see Figure 4 for examples.

In the following, we study the question: why a redundant interior vertex cannot be removed by any monotone sequence of directed flips?

Locally non-regular edges Recall that the definition of “locally non-regular edges with respect to down flips” was given in Section 6.2 (in the proof of Theorem 12). Below we give the definition of “locally non-regular edges with respect to up flips” which is the opposite of the former. Let \mathcal{T} be a triangulation of \mathbf{A} . We say an edge $e_{ab} \in \mathcal{T}$ is *locally non-regular with respect to up-flips* if the lifted vertex $\mathbf{d}' \in \mathbb{R}^3$ lies vertically above the plane passing through the lifted vertices $\mathbf{a}', \mathbf{b}', \mathbf{c}' \in \mathbb{R}^3$, where $f_{abc}, f_{abd} \in \mathcal{T}$ are the two triangles sharing at e_{ab} . In other words, the lifted edge $e_{a'b'}$ is locally non-convex when viewed from above (the viewpoint is placed at $(0, 0, +\infty)$). For examples, the blue edges in Figure 4 are all locally non-regular with respect to up-flips. The two versions of “locally non-regular edges” together make the clear definition of “locally non-regular edges with respect to the corresponding directed flips”.

In the rest of this section, we will simply say a “locally regular” or “locally non-regular” edge when the directed flips (up-flips or down-flips) are clear from the context.

The following two facts are due to the definition of non-extreme nodes. Recall that the definition of unflippability is given in Section 2.5.

Fact 14. *Let \mathcal{T} be a triangulation of \mathbf{A} and it corresponds to a non-extreme node of the poset. Then*

- (1) *All flippable edges in \mathcal{T} are locally regular.*
- (2) *All locally non-regular edges in \mathcal{T} are unflippable.*

We first prove an important lemma. Let $\mathbf{a} \in \mathcal{T}$ be a redundant interior vertex, the following lemma shows two special properties of \mathbf{a} .

Lemma 15. *Let \mathcal{T} be a triangulation of \mathbf{A} and it corresponds to a non-extreme node of the poset. Let $\mathbf{a} \in \mathbf{A}$ be a redundant interior vertex, then we show the following:*

- (i) *There exist at least two locally non-regular edges at \mathbf{a} .*

(ii) There exists at least one locally non-regular edge, e_{ap} , and the unflippability is due to p , moreover, p is a redundant interior vertex of A .

Proof. Without loss of generality, we assume that the poset is produced by up-flips, and \mathcal{T} corresponds to an upper non-extreme node of the poset, see Figure 4. Then $a \in \mathcal{T}$ is a lower interior vertex of A .

To shorten the notation, in the rest of this proof, all “locally regular” and “locally non-regular” are with respect to up-flips.

We first prove (i). Let the link vertices of a in clockwise direction be p_0, p_1, \dots, p_m . Figure 28 shows an example of a possible star of a . Since \mathcal{T} corresponds to a non-extreme node, then all flippable edge in star of a must be locally regular. Only unflippable edges in star of a might be locally non-regular. For examples, the edges ap_i, ap_j, ap_k in Figure 28. Without loss of generality, assume the edge e_{ap_0} is a locally non-regular edge. It is unflippable and the unflippability is due to a .

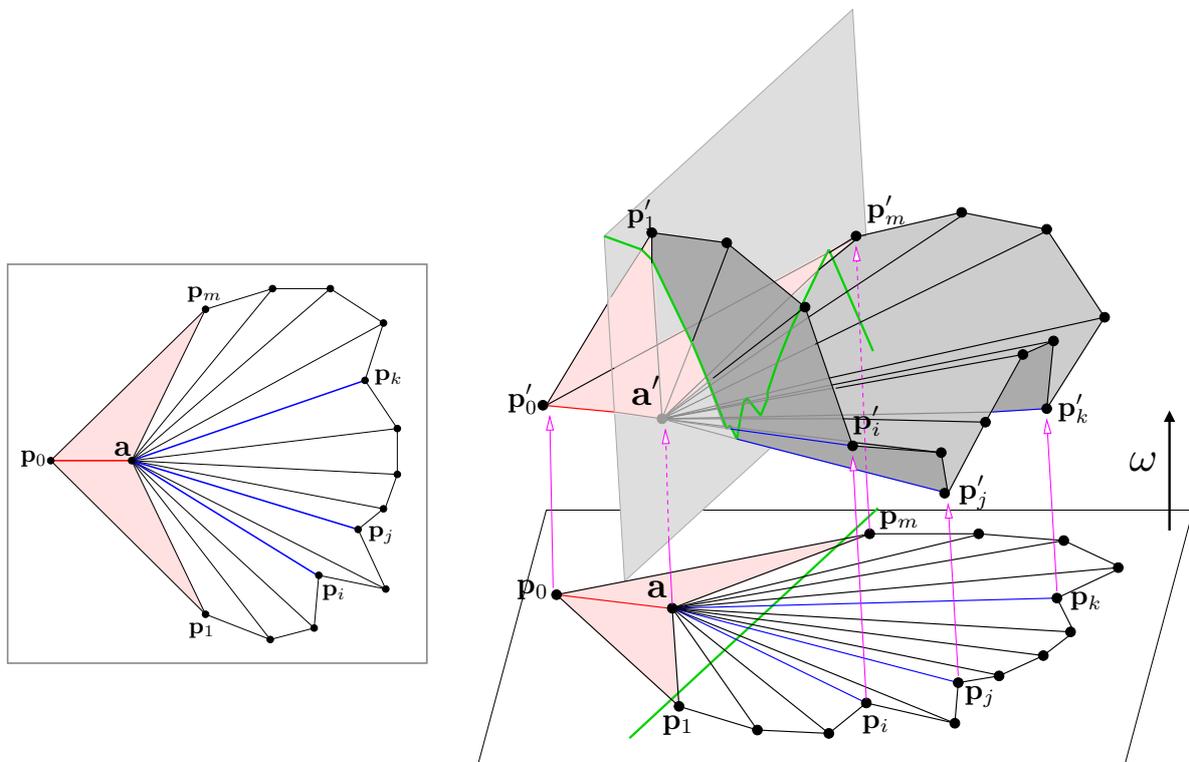


Figure 28: Left: The star of a redundant interior vertex $a \in \mathcal{T}$. Since \mathcal{T} is a triangulation corresponds to a non-extreme node, all flippable edges in this figure must be locally regular. Unflippable edges (shown in red and blue) might be locally non-regular. Right: A vertical plane (shown in grey) cuts the lifted triangles in the star of a redundant interior vertex $a \in \mathcal{T}$. The intersection curve (shown in green) is a piecewise linear line segment in \mathbb{R}^3 . It must contain at least a vertex which is locally convex. In this example, there are three such vertices on this curve, which corresponds to the intersections between the plane and the lifted edges $e_{a'p'_i}, e_{a'p'_j}$, and $e_{a'p'_k}$.

We search another locally non-regular edge from the clockwise direction of e_{ap_0} . Before we meet this edge, we will get a sequence of edges: $e_{ap_1}, e_{ap_2}, \dots, e_{ap_m}$, such that all edges are locally regular (with respect to up-flips). This means, p'_{i+1} lies vertically below the plane passing through the lifted vertices a', p'_i, p'_{i-1} , $i = 1, \dots, m - 1$. If we do not find a locally non-regular edge, this means that

p'_m lies vertically below the plane passing through the lifted points a', p'_0, p'_1 , this contradicts to the assumption that the edge e_{ap_0} is locally non-regular (with respect to up-flips). In other words, if we intersect the set of all lifted triangles in the star of a in \mathbb{R}^3 by a vertical plane, the intersection curve (which is a piecewise linear line segment) must not be concave (ideal for up-flips). It must contain at least a vertex which is locally convex, and the edge in the star of a which corresponds to this locally convex intersection point must be locally non-regular (with respect to up-flips). Figure 28 Right shows an example. Therefore, we must find at least one edge $e_{ap_i} \neq e_{ap_0}, 0 < i \leq m$, which is locally non-regular (with respect to up-flips) and it is unflippable. This proves (i).

Before we prove (ii). We show two basic facts. First, there can exist at least two locally non-regular edges whose unflippability are due to a . The other edge can be either ap_1 or ap_m , i.e., one of the adjacent edges of ap_0 . They are shown in Figure 29. Second we show that there are at most two locally non-regular edges at a . This can be proven by basic planar geometry. Without loss of generality, consider the case when ap_1 is locally non-regular, see Figure 29 Left. Then we have the following inequalities about the internal angles at a : $\alpha + \beta > 180^\circ, \beta + \gamma > 180^\circ$, then $\gamma + \delta < 180^\circ$ and $\alpha + \delta < 180^\circ$, which implies there exists no more unflippable edge at a whose unflippability can be caused by a .

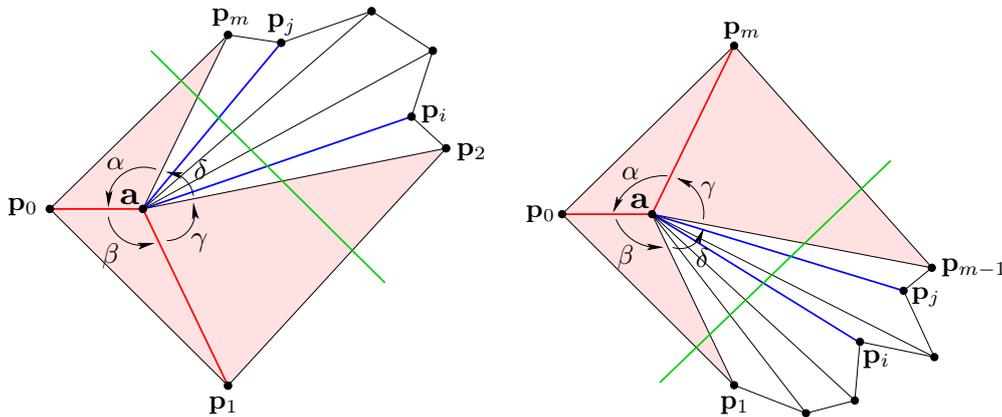


Figure 29: Two examples of the two special stars of a redundant interior vertex. There are two unflippable locally non-regular edges (red) at a . The green lines are the projections of the intersection curves in \mathbb{R}^3 between a vertical plane and the set of lifted triangles at a .

Now we prove (ii). First of all, if there exists more than two locally non-regular edges at a , then by the above facts, at least one edge whose unflippability is not due to a . If there exists only two locally non-regular edges at a , then we show that the unflippability of the other locally non-regular edge must not due to a . Consider the contrary, the unflippability of the other locally non-regular edge is also due to a , without loss of the generality, let this edge be ap_1 , see Figure 29 Left. Same as we prove (i), we must find another locally non-regular edge between edges ap_1 and ap_m . In other words, the cutting curve (shown in green in Figure 29) between a vertical plane and all lifted triangles, $f_{a'p'_1p'_2}, \dots, f_{a'p'_{m-1}p'_m}$ must contain at least one locally convex point. We arrive a contradiction. Therefore, in this case, the unflippability of the other edge must not due to a .

Finally, let e_{ap_j} be a locally non-regular edge at a whose unflippability is due to p_j . We show that p_j is a lower interior vertex of A . The fact that e_{ap_j} is locally non-regular with respect to up-flip shows that the lifted vertex p'_j lies above the plane passing through the lifted vertices a', p'_{j-1}, p'_{j+1} , where $f_{ap_jp_{j-1}}, f_{ap_jp_{j+1}} \in \mathcal{T}$. This shows that p_j must be a lower interior vertex of A . \square

The above lemma shows that in the triangulation \mathcal{T} of a non-extreme node, around a redundant interior vertex $\mathbf{a} \in \mathcal{T}$ there must exist another redundant interior vertex $\mathbf{p} \in \mathcal{T}$, such that the edge $e_{\mathbf{a}\mathbf{b}} \in \mathcal{T}$ is locally non-regular with respect to directed flip and unflippable. The same happens to \mathbf{p} . This will naturally generate a sequence of unflippable locally non-regular edges with respect to directed flips.

$$\mathcal{E} := \{e_{\mathbf{a}_0\mathbf{p}_0}, e_{\mathbf{p}_0\mathbf{p}_1}, e_{\mathbf{p}_1\mathbf{p}_2}, \dots, e_{\mathbf{p}_{m-1}\mathbf{p}_m}, e_{\mathbf{p}_m\mathbf{p}_{m+1}}\},$$

such that every edge $e_{\mathbf{a}_i\mathbf{p}_i} \in \mathcal{T}$ is locally non-regular with respect to directed flip and is unflippable.

Remark 1 Indeed, there might exist many sequences at a redundant interior vertex. This is due to the fact that there might exist more than one locally non-regular edges at this vertex whose unflippability are due to other vertices.

Since \mathcal{T} contains no directed flip at all, this means that at least one of such sequences form a cycle, i.e., $e_{\mathbf{p}_m\mathbf{p}_{m+1}} = e_{\mathbf{a}_0\mathbf{p}_0}$. For example, in Figure 4 the three edges $e_{\mathbf{a}\mathbf{b}}$, $e_{\mathbf{b}\mathbf{c}}$, and $e_{\mathbf{c}\mathbf{a}}$ form such a cycle.

Another easy outcome of the above lemma is that a cycle needs at least 3 edges. Hence there are at least 3 redundant interior vertices of \mathbf{A} in each triangulation.

We then see that any triangulation of a non-extreme external node of the poset must have the following property.

Theorem 16. *Any triangulation corresponds to a non-extreme node in this poset must contain at least 3 redundant interior vertices of \mathbf{A} . Moreover, at least 3 of these vertices are connected by cycles of unflippable locally non-regular edges with respect to the directed flips of this poset.*

Interestingly, Joe [19] proved a special case in 3d triangulations, which showed that the failure of Lawson's algorithm in 3d is due to the existence of a cycle of connected unflippable locally non-Delaunay faces [19, Lemma 4, 5]. This results implies that there must exist cycles of unflippable locally non-Delaunay edges in such 3d triangulations. Moreover, these edges are redundant interior edges.

Remark 2. From the above theorem, we see that cycles of edges which are unflippable locally non-regular with respect to directed flips are the cause of failure of Lawson's flip algorithm. In other words, if a triangulation is non-regular but it does not contain such cycles, Lawson's flip algorithm will terminate.

Remark 3. We only showed that there must exist such cycles in non-extreme triangulations. There are many interesting questions regarding the number of cycles, and the relation between these cycles. For example, do they link to each other?

7 Triangulating 3d Non-convex Polyhedra

Based on the known structural properties of the directed flip graph, this section describes an algorithm to triangulate a special class of 3d decomposable polyhedron without using Steiner points. Such polyhedra may be non-convex. This algorithm is simple since it only performs flips between triangulations of a planar point set.

7.1 The input polyhedra

Let \mathcal{P} be a family of 3d polyhedra such that any polyhedron $P \in \mathcal{P}$ satisfies the following conditions:

- (i) P is homeomorphic to a 3-ball.
- (ii) The vertices of P are in convex position.
- (iii) The vertex set of P can be represented by a planar point set \mathbf{A} with a height function $\omega : \mathbf{A} \rightarrow \mathbb{R}$, such that $\forall \mathbf{p} \in \mathbf{A}$, the lifted point \mathbf{p}' is a vertex of P .
- (iv) The boundary of $\partial P = \mathcal{T}_u^\omega \cup \mathcal{T}_v^\omega$, where \mathcal{T}_u and \mathcal{T}_v are two triangulations of \mathbf{A} .
- (v) ω is either convex or concave.

All convex polyhedra are in this family. However, only non-convex polyhedra which satisfy the (strong conditions) (iv) and (v) are in this family. The condition (iv) enforces that the orthogonal projection of the boundary of P are two planar triangulations of the same point set which share at the same convex hull of this point set. Figure 30 Left shows an example of such polyhedron P .

The condition (v) implies that either \mathcal{T}_u or \mathcal{T}_v has the property that it does not contain interior vertices of \mathbf{A} , see Figure 30 Left. With this property, the existence of a path from this triangulation to the two extreme triangulations \mathbf{A} is guaranteed (Theorem 12). They can be computed by the Lawson's flip algorithm in time $O(n^2)$.

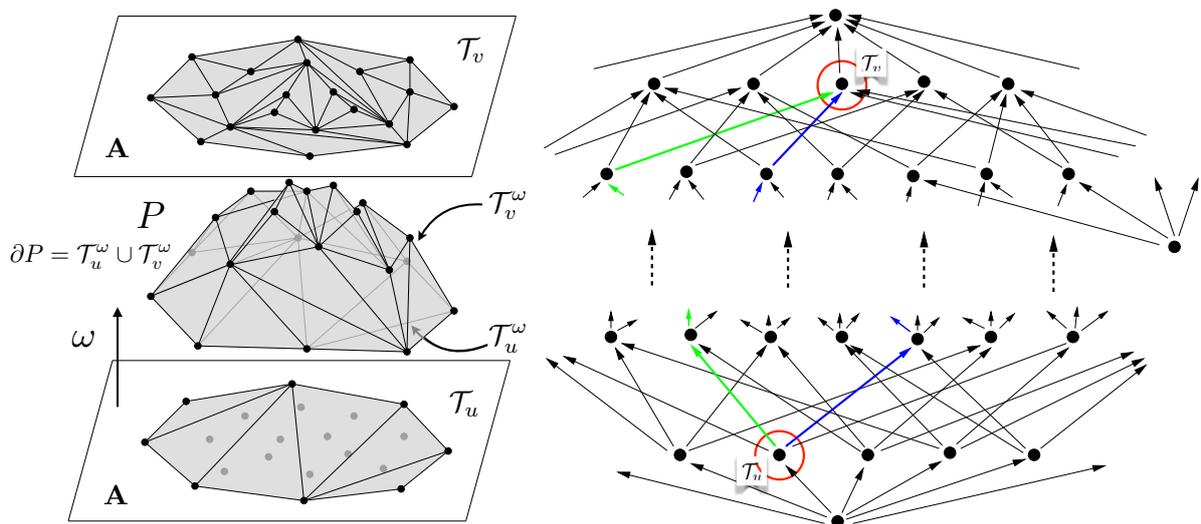


Figure 30: Left: An example of the input polyhedron. Right: Examples of existing paths from the source to the target triangulations.

Obviously, P is convex when \mathcal{T}_u and \mathcal{T}_v are just the two extreme triangulations of (\mathbf{A}, ω) . While if they are not (or one of them is not), then P must be non-convex.

7.2 The triangulation algorithm

Let P be a 3d polyhedron satisfies the above conditions. We describe an algorithm to triangulate P without using Steiner points.

The 3d triangulation problem is then transformed into finding a monotone sequence of directed flips between the source and target triangulations. It is then equivalent to find a path in the directed flip graph between these two nodes. Let us assume that the existence of a path between these two nodes in this graph is known. The question is how to find such a path?

The ideas of the algorithm Let us assume the source triangulation is the one which contain no interior vertices of \mathbf{A} , and the target triangulation is an internal nodes of this graph. If we follow an arbitrary path starting from the source triangulation, it is guaranteed that we will reach the extreme triangulation (in its direction) of \mathbf{A} , but we may by-pass the internal node corresponds to the target triangulation.

Here comes the main ideas of our algorithm: Assume there exists a directed flip. Before doing this flip in a triangulation \mathcal{T} of \mathbf{A} , we ensure that the corresponding tetrahedron of this flip does not intersect the target triangulation in its interior. Let us assume \mathcal{T}_v is the target triangulation, let $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d} \in \mathcal{T}$ are the support of this flip. We ensure the following condition:

$$\text{int}(t_{\mathbf{abcd}}) \cap |\mathcal{T}_v| = \emptyset,$$

where $\text{int}(t_{\mathbf{abcd}})$ means the interior of the tetrahedron $t_{\mathbf{abcd}}$, and $|\mathcal{T}_v|$ is the underlying space of \mathcal{T}_v . We call such flip is *conforming to the target triangulation*. If this flip does not satisfy the above condition, we skip this flip. We will show later, if we only do conforming flips, then we must end at the target triangulation.

The algorithm Without loss of generality, assume \mathcal{T}_u is the source triangulation and \mathcal{T}_v is the target. And assume that \mathcal{T}_u lies vertically below \mathcal{T}_v . Hence the directed flips we need are up-flips. We initialise a working triangulation $\mathcal{T} := \mathcal{T}_u$, and initialise a working list $L := \emptyset$ for returning the list of tetrahedra of the triangulation.

Let \mathcal{E} be the set of all locally non-regular edges with respect to up-flips in \mathcal{T} . For each edge $e_{\mathbf{ab}} \in \mathcal{E}$. if it is already an edge of \mathcal{T}_v , we simply remove it from \mathcal{E} . Otherwise, if it is flippable, and this flip is conforming to \mathcal{T}_v , then flip $e_{\mathbf{ab}}$, and update \mathcal{E} by all new locally non-regular edges with respect to up-flips in \mathcal{T} , add the tetrahedron $t_{\mathbf{abcd}}$ of this flip to L . This phase ends until \mathcal{E} is empty.

If $\mathcal{E} = \emptyset$ and $\mathcal{T} \neq \mathcal{T}_v$, check if there exists a vertex $\mathbf{a} \in \mathcal{T}_v$ and $\mathbf{a} \notin \mathcal{T}$. If such a vertex exists, there must exist a triangle $f_{\mathbf{bcd}} \in \mathcal{T}$ which contains \mathbf{a} , insert \mathbf{a} into \mathcal{T} by performing a 1-3 flip in \mathcal{T} . update \mathcal{E} by all new locally non-regular edges with respect to up-flips in \mathcal{T} , add the tetrahedron $t_{\mathbf{abcd}}$ of this flip to L .

7.3 Termination and running time

In this section, we prove the following theorem regarding the termination as well as running time of this algorithm.

Theorem 17. *If the target triangulation is a regular triangulation, then this algorithm terminates, and it runs in $O(n^3)$ time, where n is the number of vertices of the input polyhedron P .*

Proof. If the target triangulation is the unique extreme node, then every path starting from the source triangulation will lead to it. There is even no need to check conformity of each flip. It is just the Lawson's algorithm with edge flips and vertex insertions. In this case, the algorithm runs in $O(n^2)$ which is also the upper bound of the number of tetrahedra of a set of n vertices.

From now on we assume the target triangulation is an internal node of this poset. The condition imposed on the target triangulation means that it must not be a non-extreme node of the poset, which means that there exists a path from the source triangulation toward the target triangulation, this also means a 3d triangulation between the source and target triangulations exist.

Without loss of generality, assume the source triangulation \mathcal{T}_u lies vertically below \mathcal{T}_v , and the directed flips we need are up-flips. By performing only conforming up-flips, we ensure that every intermediate triangulation \mathcal{T} must lie vertically below \mathcal{T}_v , i.e., the lifted triangulation \mathcal{T}^ω never crosses the characteristic section of \mathcal{T}_v^ω . This shows that we will never by-pass the target node in the poset.

We still need to show that the target node must be reached. It is based on the following facts: there exists a 3d triangulation between the source and target triangulation. Our input ensures that the source triangulation contains no interior vertex. This implies we only need 2-2 flips (edge flips) and 1-3 flips (vertex insertions). This means that if the target triangulation is not reached, there always exists at least one up-flip in current triangulation. Therefore, the target triangulation must be reached.

Now we show the upper bound of the running time of this algorithm. The number of total tetrahedra of this triangulation is $O(n^2)$. At creation of each tetrahedron, we need to test the conformity against the target triangulation, which has $O(n)$ triangles. In the worst case, we will need to test t locally non-regular edges in order to find a valid conforming flip. If $t = O(n)$, then we get a naive upper bound $O(n^4)$. We can avoid this additional $O(n)$ tests at each creation of a tetrahedron by tagging all the tested locally non-regular edges which correspond to non-conforming flips. So we only need to test such flip once and will automatically avoid to test it again. Therefore the worst-case running time of this algorithm is $O(n^3)$. \square

Remark 1. We remark that the condition on source triangulations is sufficient but not necessary. This means that it is possible to start with a source triangulation which does contain interior vertices of \mathbf{A} . This will end at a node which is non-regular. As we showed in Theorem 11 (6) that non-regular triangulation might be also an interior node. This means, it still can find directed flips but this must be a 1-3 flip (vertex insertion), and the new vertex must be in the target triangulation.

Remark 2. The target triangulation may be non-regular triangulation as long as it does not correspond to a non-extreme node of this poset.

8 Discussions

In this paper, we studied monotone sequence of directed flips between triangulations of a finite point set (\mathbf{A}, ω) in $\mathbb{R}^2 \times \mathbb{R}$, where ω is a height function which lifts each vertex of \mathbf{A} into \mathbb{R}^3 . It is shown that such sequences correspond to 3d triangulations of the lifted point set \mathbf{A}^ω . We then studied the structural properties of the directed flip graph of the set of all triangulations of this point set. These properties clearly explained the general behaviour of the Lawson's flip algorithm. Based on this study, we proposed a simple algorithm to triangulation a special family of 3d non-convex polyhedra without using additional vertices.

There are many interesting questions and problems remain to be studied. We have mentioned some of them within the paper.

- (1) Question 10 in Section 5 asks: given a 3d triangulation \mathcal{T} of a point set \mathbf{V} in \mathbb{R}^3 . Assume \mathbf{V} contains no interior vertices. Is it always possible to find a viewpoint in \mathbb{R}^3 such that \mathcal{T} is acyclic with respect to this viewpoint?
- (2) In Section 6 we have characterised all external (extreme and non-extreme) nodes of this poset. It remains to understand what are properties of internal nodes? For examples, we would like to know where the set of internal nodes form a graded lattice?

- (3) Continue from (2), we would like to know more properties of triangulations corresponding to non-extreme nodes. Although it is proven there must exist cycles of unflippable locally non-regular edges, it remains unclear how different cycles relate to each other. For example, can they form a non-trivial link?
- (4) There is still a piece of information from the motivation example in Section 3.1 not be studied. In this example, we could easily “see” that two triangulations are connected by a 3d flip (2-3 or 3-2 flip), see Figure 16. How to characterise this 3d flips in the poset?
- (5) We have seen from many previous examples, e.g., Figure 3, that the existence of redundant interior vertices may cause the failure of the Lawson’s algorithm. It is then necessary to study the following problem: *How to remove an interior vertex by a monotone sequence of directed flips?* We comment that if we know how to do this, we could generalise our algorithm to triangulate a more general class of non-convex polyhedron whose vertices are on the convex hull.

The last question of this paper is how to generalise the result of this paper to higher dimensional point sets. The case of point sets in 3d is already a challenging question to consider.

References

- [1] F. Bagemihl. On indecomposable polyhedra. *The American Mathematical Monthly*, 55(7):411–413, 1948.
- [2] M. Bern. Compatible tetrahedralizations. In *Proc. 9th Annual ACM Symposium on Computational Geometry*, pages 281–288, 1993.
- [3] A. Bezdek and B. Carrigan. On nontriangulable polyhedra. *Beiträge zur Algebra und Geometrie (Contributions to Algebra and Geometry)*, 57(1):51–66, 2016.
- [4] A. Bowyer. Computing Dirichlet tessellations. *Comp. Journal*, 24(2):162–166, 1981.
- [5] B. Chazelle. Convex partition of polyhedra: a lower bound and worst-case optimal algorithm. *SIAM Journal on Computing*, 13(3):488–507, 1984.
- [6] J. A. De Loera, J. Rambau, and F. Santos. *Triangulations, Structures for Algorithms and Applications*, volume 25 of *Algorithms and Computation in Mathematics*. Springer Verlag Berlin Heidelberg, 2010.
- [7] B. N. Delaunay. Sur la sphère vide. *Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk*, 7:793–800, 1934.
- [8] P. Edelman and V. Reiner. The higher Stasheff-Tamari posets. *Mathematika*, 43:127–154, 1996.
- [9] H. Edelsbrunner. An acyclicity theorem for cell complex in d dimension. *Combinatorica*, 10(3):251–260, 1990.
- [10] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. *Algorithmica*, 15:223–241, 1996.
- [11] D. Eppstein. The farthest point delaunay triangulation minimizes angles. *Computational Geometry*, 1(3):143 – 148, 1992.

- [12] D. Eppstein, Marc van Kreveld, Elena Mumford, and Bettina Speckmann. Edges and switches, tunnels and bridges. *Computational Geometry*, 42(8):790 – 802, 2009. Special Issue on the 23rd European Workshop on Computational Geometry.
- [13] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Birkhäuser Boston, Boston, MA, 1994.
- [14] P.-L. George, H. Borouchaki, and E. Saltel. últimate robustness in meshing an arbitrary polyhedron. *International Journal for Numerical Methods in Engineering*, 58:1061–1089, 2003.
- [15] P.-L. George, F. Hecht, and E. Saltel. Automatic mesh generator with specified boundary. *Computer Methods in Applied Mechanics and Engineering*, 92:269–288, 1991.
- [16] J. Goodman and J. Pach. Cell decomposition of polytopes by bending. *Israel J. Mathematics*, 64:129–138, 1988.
- [17] R. Jaume and G. Rote. Recursively regular subdivisions and applications. *Journal of Computational Geometry*, 7(1), 2016.
- [18] B. Jessen. Orthogonal icosahedra. *Nordisk Mat. Tidskr*, 15:90–96, 1967.
- [19] B. Joe. Three-dimensional triangulations from local transformations. *SIAM Journal on Scientific and Statistical Computing*, 10(4):718–741, 1989.
- [20] B. Joe. Construction of three-dimensional Delaunay triangulations using local transformations. *Computer Aided Geometric Design*, 8:123–142, 1991.
- [21] C. L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3(4):365–372, 1972.
- [22] C. L. Lawson. Software for c^1 surface interpolation. *Mathematical Software III, Academic Press*, pages 164–191, 1977.
- [23] C. L. Lawson. Properties of n-dimensional triangulations. *Computer Aided Geometric Design*, 3(4):231 – 246, 1986.
- [24] J. Radon. Mengen konvexer Körper, die einen gemeinschaftlichen Punkt enthalten. *Math. Ann.*, 83:113–115, 1921.
- [25] J. Rambau. *Polyhedral Subdivisions and Projections of Polytopes*. PhD thesis, Fachbereich 3 Mathematik der Technischen Universität berlin, Berlin, Germany, Oktober 1996.
- [26] J. Rambau. On a generalization of Schönhardt’s polyhedron. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*, volume 52, pages 501–516. MSRI publications, 2005.
- [27] J. Ruppert and R. Seidel. On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete & Computational Geometry*, 7:227–253, 1992.
- [28] F. Santos. A point set whose space of triangulations is disconnected. *Amer. Math. Soc.*, 13:611–637, 2000.
- [29] F. Santos. Triangulations with very few geometric bistellar neighbors. *Discrete & Computational Geometry*, 23(1):15–33, 2000.

- [30] E. Schönhardt. Über die zerlegung von dreieckspolyedern in tetraeder. *Mathematische Annalen*, 98:309–312, 1928.
- [31] H. Si. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.*, 41(2):11:1–11:36, February 2015.
- [32] H. Si and N. Goerigk. Generalised Bagemihl polyhedra and a tight bound on the number of interior Steiner points. *Computer-Aided Design*, 103:92 – 102, 2018.
- [33] D. D. Sleator, W. P. Thurston, and R. E. Tarjan. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1:647–682, 1988.
- [34] G. T. Toussaint, C. Verbrugge, C. Wang, and B. Zhu. Tetrahedralization of simple and non-simple polyhedra. In *Proc. 5th Canadian Conference on Computational Geometry*, pages 24–29, 1993.
- [35] D. F. Watson. Computing the n -dimensional Delaunay tessellations with application to Voronoi polytopes. *Comput. Journal*, 24(2):167–172, 1981.
- [36] N. P. Weatherill and O. Hassan. Efficient three-dimensional Delaunay triangulation with automatic point creation and imposed boundary constraints. *International Journal for Numerical Methods in Engineering*, 37:2005–2039, 1994.
- [37] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1997.