MINIMAL SIMPLICIAL DISSECTIONS AND TRIANGULATIONS OF CONVEX 3-POLYTOPES.*

Alexander Below
Institut für Theoretische
Informatik, ETH-Zürich
below@inf.ethz.ch

Ulrich Brehm
Institut für Geometrie,
Technische Univ. Dresden
brehm@math.tu-dresden.de

Jesús A. De Loera
Dept of Mathematics,
Univ. of California-Davis
deloera@math.ucdavis.edu

Jürgen Richter-Gebert
Institut für Theoretische
Informatik, ETH-Zürich
richter@inf.ethz.ch

July 6, 1999

Abstract

This paper addresses three questions related to minimal triangulations of a 3-dimensional convex polytope P.

- Can the minimal number of tetrahedra in a triangulation be decreased if one allows the use of interior points of P as vertices?

- Can a dissection of P use fewer tetrahedra than a triangulation?

- Does the size of a minimal triangulation depend on the concrete realization of P?

The main result of the this paper is that all these questions have an affirmative answer. Even stronger the gaps of size produced by allowing interior vertices or by using dissections may be linear in the number of points.

1 Introduction.

Let $\mathcal{A}$ be a point configuration in R^d with its convex hull $conv(\mathcal{A})$ having dimension d. A set of d-simplices with vertices in $\mathcal{A}$ is a *dissection* of $\mathcal{A}$ if no pair

---

of simplices has an interior point in common and their union equals $conv(\mathcal{A})$. A dissection is a *triangulation* of $\mathcal{A}$ if in addition any pair of simplices intersects at a common face. Following the standard use in the literature (see for instance [13]), we say that a triangulation (dissection) of a convex polytope P is a triangulation (dissection) of the vertices of P. Notice that in this definition no additional points are allowed. In contrast to this, we will be interested on triangulations of point sets with additional interior points besides those of P and will called them *triangulations of P with interior points*. The *size* of a dissection is the number of d–simplices it contains. We call a triangulation or a dissection *minimal/maximal* when its size is the smallest/largest among all the triangulations or respectively all dissections. Note that dissections do not necessarily form simplicial complexes.

In this paper we solve two questions about triangulations and dissections, and present several relevant consequences:

First, in Section 3 we investigate the behavior of sizes of triangulations under the addition of new interior points. Independently Böhm [5] and Gritzmann and Klee [10] raised the issue that it is conceivable that by using auxiliary interior points the size of a triangulation of the new set of points could be smaller than a minimal triangulation of the original set. As pointed out in section 8.4 of [10], this question is relevant in the study of complexity classes of basic problems in computational convexity. In this article we exhibit a family of 3-dimensional polytopes that indeed have this property. This behavior is reminiscent of the fact that adding points can also reduce the size of a Delaunay triangulation [3, 6].

Second, we show in Section 3 that dissections can require fewer simplices than triangulations. We show that a linear difference between the size of a minimal triangulation and a minimal dissection is indeed possible. The question whether this could happen had been raised several times [5, 7, 12, 10]. The solution of the questions depends on two simple geometric lemmas that will be useful throughout the paper. They are proved in Section 2. The main results of Section 3 can be stated as follows:

Theorem 1.1 *There is a family of simplicial convex 3–polytopes P_m with the following properties:*

1. *all triangulations of P_m are larger than a triangulation with one interior point. This gap can be linear.*

2. *all triangulations of P_m are larger than the minimal dissections of P_m. This gap can be linear.*

Moreover, the approximate shape of the polytope can be prescribed.

Theorem 1.2   *1. Given three numbers $h_0 \geq 1$, $k \geq 1$, $\varepsilon > 0$ and a convex body $K \subseteq R^3$ there is a simplicial convex 3–polytope P with $\delta(K, P) < \varepsilon$*

(Hausdorff distance) such that every triangulation of P has at least k tetra-hedra more than a triangulation of P with h_0 suitably chosen additional interior points.

2. *For a 3–polytope P, if adding n' new interior points allows for a smaller triangulation than a minimal triangulation of P, then $n' < e_i$ where e_i is the number of interior edges in a minimal triangulation of P. In particular, it is impossible to have such behavior for 3–polytopes whose minimal triangulations use one or no interior edges.*

A famous open question in computational geometry asks for the computational complexity of finding minimal triangulations of convex polyhedra [2]. A related problem is whether the smallest size of a triangulation of a convex polytope is a determined from the face lattice or if it can change with the coordinates (see [8]). We present the negative answer in Section 4.

Theorem 1.3 *The minimal size of a triangulation of a convex 3–polytope is not an invariant of the face lattice. There is a simplicial 3-polytope with 10 vertices for which the minimal number of tetrahedra possible in a triangulation depends on its coordinates. The example is smallest possible in dimension and number of vertices.*

2   Two useful lemmas.

The following elementary formula shows that many of the arguments about the number of tetrahedra in a triangulation can be reduced to study the number of interior edges. Similar formulas have been written already many times in the literature (see for instance, Lemma 2.1 in [9] or [16]). The proof is a simple application of Euler's formula for triangulated 2-spheres and 3-balls and we omit the easy details.

Lemma 2.1 *Let P be a 3-polytope with n boundary vertices and n' interior auxiliary vertices. For a triangulation T of P that uses the n' interior points the number of tetrahedra in T, denoted #T, is related to the number of interior edges e_i of T by the formula:*

$$\#T = e_i + n - 3 - n'.$$

We now present a lemma that showcases the main non-trivial effect we will use in all the examples. The lemma shows that a certain substructure in the face lattice of a polytope forces certain interior edges to appear in triangulations of sufficiently small size.

Lemma 2.2 *Let P be a convex 3-polytope, that contains the following collection of triangular facets: (a, q_i, q_{i+1}), and (b, q_i, q_{i+1}) for $i = 1, \ldots, m-1$ (see Figure*
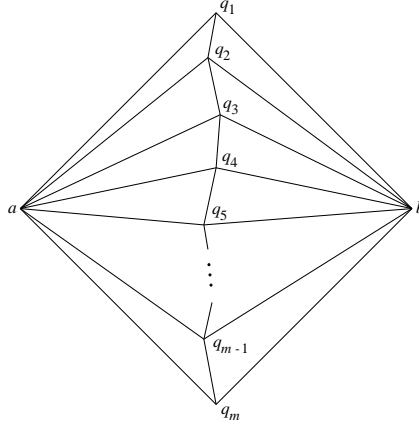
Figure 1: The key structure of Lemma 2.2

.

1), with the additional restriction that $conv\{q_1, q_m\}$ goes through the interior of P.

Then for each triangulation of P that does not use the edge ab the number of interior edges e_i is at least $m - 2$.

Proof: Since $conv\{q_1, q_m\}$ is in the interior of P, we obtain the following simple fact: For all $|i - j| \geq 2$, if $q_i q_j$ is an edge of a triangulation, it will be also an interior edge. The proof of the lemma proceeds by induction on m. The property is satisfied for $m = 3$. We can assume that all vertices q_i, with $2 \leq i \leq m - 1$, are incident to at least one interior edge of the triangulation T. Here are the reasons: (1) A q_i untouched by an interior edge belongs to the tetrahedra $\sigma_{i,a} = (a, q_{i-1}, q_i, q_{i+1})$ and $\sigma_{i,b} = (b, q_{i-1}, q_i, q_{i+1})$. This is because the triangle (a, q_i, q_{i+1}) is in some simplex, and if the fourth point is some other vertex besides q_{i-1} or b we have an interior edge touching q_i. Furthermore the fourth point cannot be b since in this case the edge ab would be present. (2) By chopping off these two tetrahedra together with the vertex q_i (i.e. considering the convex hull of all of P's vertices except q_i) we can apply induction to guarantee that the remaining triangulation $T \setminus \sigma_{i,a}, \sigma_{i,b}$ has at least $m - 3$ interior edges. Together with the edge $q_{i-1} q_{i+1}$ they amount for $m - 2$ interior edges in T.

Following the above assumption, we set up a one-to-one map from the set $\{q_2, \ldots, q_{m-1}\}$ to a subset of the interior edges that touch them: The vertices q_i come along a polygonal curve in a canonical order which is reflected by their indices. We mark and orient the interior edges $q_i v$ that touch a vertex q_i as follows: If $v \notin \{q_1, \ldots, q_m\}$, we call the edge $q_i v$ *special*, otherwise we *orient* it from smaller to larger index. For the vertices q_i with special edges incident to them, we map q_i to one of those. If a vertex q_i has no special edges, but has

4

outgoing interior edges, we map it to the outgoing edge $q_i q_k$ with the smallest index k. We are left with the case of those q_i that have only incoming interior edges incident to q_i. Consider the triangle $a q_i q_{i+1}$. It has to be in some tetrahedron of T whose fourth point is bound to be a q_{j_a} with $j_a < i$. Likewise $b q_i q_{i+1}$ is in a tetrahedron with fourth point q_{j_b} with $j_b < i$. If both $j_a = j_b = i-1$, there can be no interior edges incident to q_i (see above), a contradiction. Let j be any of j_a, j_b such that $j < i-1$. Map q_i to $q_j q_{i+1}$.

We claim that the given map is one-to-one. If some vertex q_i maps to the special edge $q_j v$, then necessarily $i = j$. There are potentially two vertices that can be mapped to an interior edge $q_j q_k$ with $j < k$: q_j when $q_j q_k$ is the chosen outgoing edge of q_j and q_{k-1}, in case q_{k-1} has only incoming edges. In the latter case one of the tetrahedra (a, q_j, q_{k-1}, q_k) and (b, q_j, q_{k-1}, q_k) has to be in the triangulation, and q_j will be mapped to the smaller indexed edge $q_j q_{k-1}$. This is an interior edge since $j < k-2$, so q_j cannot also be mapped to $q_j q_k$. The injectivity of the map is proven. $\square$

Roughly speaking the main use of Lemma 2.2 is the possibility to force the occurrence of many interior edges when other interior edges are absent. Our solutions of the open questions rely on a combination of this effect with a famous example, that appears in many textbooks of discrete and computational geometry, Schönhardt's example of a non-triangulable non-convex polyhedron (see [13, 14, 15, 17, 18]). The example, obtained from "twisting" a triangular prism, has the property that cannot be triangulated nor dissected unless we add a new point and then the triangulation is unique. A variation of the same example is a triangulation of the boundary of a triangular prism that chooses boundary edges in a "cyclic" manner. Such boundary triangulation cannot be extended to a triangulation of the whole convex hull.

3 Adding points or dissecting can reduce the size.

Proof: (of Theorem 1.1) Consider a triangular prism with bottom triangle (A_1, A_2, A_3) and top triangle (B_1, B_2, B_3) where both triangles are identical, parallel, equilateral, and vertex A_i lies directly below the vertex B_i. Denote by $n_{i,j}$ the inner normal to the quadrilateral facet (A_i, A_j, B_i, B_j) (see the figure below for a description). In each of its quadrilateral facets we will put m points following a diagonal $A_i B_j$ but slightly lifted to be in a parabola. The points we add will have coordinates along three parabolic curves $C_{1,2}$, $C_{2,3}$, and $C_{3,1}$. For the definition of these curves let $p(\lambda) = (\lambda - 1/2)^2 - 1/4$, and $\epsilon > 0$ a very small number.

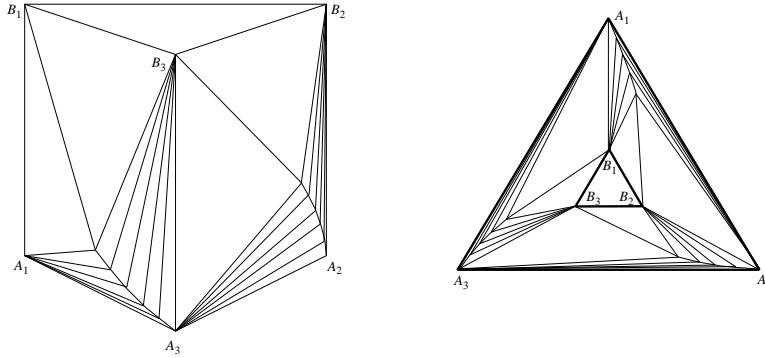$$C_{1,2} = \lambda A_1 + (1-\lambda)B_2 + \epsilon p(\lambda) n_{1,2},$$

Figure 2: The polytope P_m and its Schlegel diagram.

$$C_{2,3} = \lambda A_2 + (1 - \lambda)B_3 + \epsilon p(\lambda)n_{2,3},$$

$$C_{3,1} = \lambda A_3 + (1 - \lambda)B_1 + \epsilon p(\lambda)n_{3,1}.$$

Note that the curve $C_{i,j}$ joins the vertices A_i and B_j. The pattern of the curves follows cyclically arranged diagonals as in the case of Schönhardt's polyhedron. Take m points along each of the curves $C_{1,2}$, $C_{2,3}$, $C_{3,1}$ that we denote by $q_1(C_{i,j}), q_2(C_{i,j}), \ldots, q_m(C_{i,j})$. We will take those points such that all the triangles $(q_k(C_{1,2}), B_2, B_3)$ are pierced by the edge $B_1 A_2$, the triangles $(q_k(C_{2,3}), B_3, B_1)$ are pierced by the edge $B_2 A_3$, and similarly the triangles $(q_k(C_{3,1}), B_1, B_2)$ are pierced by the edge $B_3 A_1$. These *piercing conditions* are easily achieved by accumulating the points low enough toward the triangle (A_1, A_2, A_3). We observe that the sequences of points produce, when taking the convex hull, $2(m + 1)$ triangular facets, $m + 1$ on each side of the diagonal. This is exactly the situation of Lemma 2.2. The resulting 3-polytope, containing $3m + 6$ vertices, will be denoted by P_m. The polytope P_m is the union of the original prism with three "caps" placed on top of the quadrilateral facets. We show in Figure 2 the case for $m = 5$.

Now we construct a triangulation T_v of P_m that uses an auxiliary interior point v within the prism $conv(A_1, A_2, A_3, B_1, B_2, B_3)$. At each cap of the polytope P_m we triangulate in a "stacked" way, namely by choosing the tetrahedra of the form $(q_k(C_{i,j}), q_{k+1}(C_{i,j}), B_i, A_j)$. For each curve $C_{i,j}$ we get $m + 1$ tetrahedra. Complete the triangulation by coning from v to the eight triangles that triangulate the boundary of $conv(A_1, A_2, A_3, B_1, B_2, B_3)$. We have a triangulation, using the interior point v, with $3m + 11$ tetrahedra.

We now construct a dissection D of P_m. We do not use any interior points. Triangulate the three caps as before in a stacked way, then for the interior of the
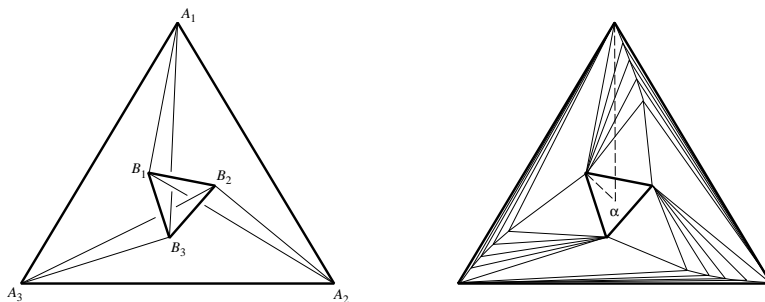
Figure 3: A twisted version of P_m, based on Schönhardt's example.

prism triangulate arbitrarily. All triangulations of the prism are combinatorially the same and have three simplices. The dissection D has size $3m + 6$.

From Lemma 2.2 we have that triangulations that do not use all the edges $B_1 A_2$, $B_2 A_3$, or $B_3 A_1$ will have at least m interior edges (note, we apply the lemma with $m + 2$ points). In this case Lemma 2.1 implies that the number of tetrahedra will be at least $4m + 3$. On the other hand, we claim that it is impossible to have a triangulation that uses those three edges simultaneously. The reason is the triangular face (B_1, B_2, B_3), must belong to a certain tetrahedron, and if the fourth point is along one of the curves $C_{i,j}$ then, by the piercing conditions, there is a bad intersection with one of the edges $B_1 A_2$, $B_2 A_3$, or $B_3 A_1$. Hence, the only hope is that the fourth point is another of the points A_l, but this is impossible in a triangulation as two edges will intersect improperly. In conclusion, any triangulation of P_m that does not use interior points is larger than the triangulation T_v and larger than the dissection D for $m > 8$.□

Before we come to the proof of Theorem 1.2, we present some observations on a twisted version of polytope P_m. We will patch several copies of this new polytope onto simplicial polytopes when constructing the examples of Theorem 1.2. We start with a triangular prism such that the bottom triangle (A_1, A_2, A_3) and the top triangle (B_1, B_2, B_3) are equilateral, and their barycenters are vertically on top of each other. The bottom triangle should be somewhat larger than the top triangle. We twist the top triangle by an angle of α degrees creating a Schönhardt style non-convex polyhedron. Then we patch the sides again with copies of caps used in Lemma 2.2. We call the resulting convex polytope $P_{m,\alpha}$. Therefore the original example P_m of Theorem 1.1 is isomorphic to $P_{m,0}$ (see Figure 3). The copies will have again the points $q_k(C_{i,j})$ close to the bottom triangle to guarantee the piercing conditions as in the proof of Theorem 1.1. The number ϵ should be small enough for the points $q_k(C_{i,j})$ to lie vertically above the bottom triangle. By *vertically above (below)* we mean that a line going vertically down (up) from any of those points pierces (A_1, A_2, A_3).

When the angle α is greater than zero, we are in the situation that the

dissection D we described becomes impossible, but one can still use an interior point v to create a triangulation T_v with an interior point, smaller than any triangulation of the polytope. The proof is identical as the one we discussed for Theorem 1.1.

An important difference from the example of Theorem 1.1 is that now the set of possible interior points v to construct T_v is restricted by the twisting angle α: We call the *set of visibility* of $P_{m,\alpha}$ all points p which are below the plane spanned by the top triangle and satisfy the following condition: The tetrahedron spanned by B_1, B_2, B_3, and p does not intersect the edges B_1A_2, B_2A_3, or B_3A_1. It is easy to check that the set of visibility is the union of four sets (see Figure 3): The tetrahedron S_0 bounded by the planes $B_1B_2B_3$, $B_1A_1B_3$, $B_2A_2B_1$, and $B_3A_3B_2$; the cone S_1, with apex B_1 lying between the planes $B_1B_2B_3$, $B_1A_1B_3$, $B_1A_3B_2$; as well as the analogous cones S_2 and S_3 with apexes B_2 and B_3. The reason we want the piercing conditions is exactly to avoid this set of visibility.

By symmetry, the three planes $B_1A_1B_3$, $B_2A_2B_1$, and $B_3A_3B_2$ intersect in a point p which lies on the vertical line connecting the barycenters of the bottom and top faces. If the twisting angle $\alpha = \pi/3$, then p is the intersection of edges B_1A_2, B_2A_3, and B_3A_1. By continuity of a rotation, if $\alpha = \pi/3 - \delta$ for small $\delta > 0$, the point p is still above the bottom triangle. Therefore S_0 has no points vertically below (A_1, A_2, A_3). Note that the sets S_1, S_2, S_3 do not contain points vertically below the bottom triangle.

In conclusion, for a suitable twisting angle α the polytope $P_{m,\alpha}$ will have four properties: (1) it contains only points in and vertically above its bottom face, (2) there are no points of visibility vertically below the bottom triangle, (3) inside $P_{m,\alpha}$ there is a point which can "see" the top and (by an analogous argument) the bottom triangular faces, and (4) a new auxiliary interior point v used in a smaller triangulation T_v must lie strictly in the set of visibility of $P_{m,\alpha}$.

Imagine now a polytope Q, with facet (A_1, A_2, A_3), lying vertically below this facet. Then $Q \cup P_{m,\alpha}$ is a convex polytope. There is a triangulation of $Q \cup P_{m,\alpha}$ using an interior point in $P_{m,\alpha}$ which does not use A_1B_3, A_2B_1, A_3B_2, but there is no such triangulation if we do not allow extra interior points within $P_{m,\alpha}$. Via a projective transformation we can transform any polytope Q' with a specified triangular facet (A_1', A_2', A_3') to such a polytope Q. We can even arrange for the preimage of $P_{m,\alpha}$ under this projective transformation to be as "flat" as we want it to be. Hence we can patch arbitrarily flat isomorphic copies of $P_{m,\alpha}$ to any triangular facet of Q' still obtaining a convex polytope.

Proof: (of Theorem 1.2 part (1)) Approximate the convex body K with a simplicial 3–polytope, with at least h_0 facets, Q_0 that is close to it in the Hausdorff distance (see [11] for references on approximation by polytopes). On h_0 of the triangular facets of Q_0 we stack polytopes $Y_0, Y_1, \ldots, Y_{h_0}$ where each Y_i is an isomorphic copy of a polytope $P_{m,\alpha}$ and $Q_1 = Q_0 \cup Y_1 \cup Y_2 \cup \ldots Y_{h_0}$ is a convex polytope still ϵ-close to K. This stacking procedure is possible by the above

discussion.

If one triangulates Y_i with an interior point that lies in its visibility set one finds a very small triangulation, but the auxiliary point is forced to lie in the visibility sets by construction of each copy Y_i. Moreover, the visibility sets of different Y_i, Y_j are disjoint from each other and from Q_0. The rest of Q_1 is the original polytope Q_0 that triangulate minimally. Thus if we add h_0 points to the interior of Q_1 we can produce a triangulation that has at least k fewer tetrahedra than the minimal triangulation of Q_1 without new points, since for large enough m we can surpass any value k we are given.□

Proof: (of Theorem 1.2 part (2)) We use the equation $\#T = e_i + n - 3 - n'$. from Lemma 2.1 where $\#T$ denotes the number of tetrahedra in T, n (n') equals the number of boundary (interior) vertices, and $n' = 0$ before we add interior points.

If using the n' interior auxiliary points reduces the size of the triangulation, then we get the inequality $e_i^* + n - 3 - n' < e_i + n - 3$, where e_i^* is the number of interior edges in the new triangulation T' that uses interior vertices. Notice that since every interior point is in at least 4 interior edges, $4n' \leq 2e_i^*$. In conclusion $n' < e_i$. The rest of the statement follows immediately. □

Note that if in the triangulation T_v of Theorem 1.1 we take the new point v not in the interior of P_m, but rather on the relative interior of (A_1, A_2, A_3), then there is a bigger difference between the size of T_v and any triangulation of P_m because there are at least three fewer interior edges. Hence we conclude that triangulating a convex polytope using new auxiliary boundary points can result on smaller triangulations as well. It is natural to ask what happens when we add points in the exterior since we have looked at adding the points at the interior or the boundary of the polytope. The following proposition summarizes the answer:

Proposition 3.1 *There are polytopes whose minimal triangulations are larger than the minimal triangulations of the new polytopes obtained by adding a new exterior point, this gap can be linear in fixed dimension three or even grow exponentially if we allow the dimension to grow.*

Example 3.2 Take an m-antiprism with vertices $\{(\cos(2i\pi/m), \sin(2i\pi/m), 0), (\cos((2i + 1)\pi/m), \sin((2i + 1)\pi/m), 1) : i \in \{0, \ldots, m - 1\}$. It is a polytope in R^3 having $2m$ vertices. A minimal triangulation has cardinality $3m - 5$ (see [8]). Adding a new vertex $(0, 0, M)$, M large enough, makes the polytope stacked, having minimal triangulation of cardinality $2m - 2$ (see Figure 4).

Take a d-dimensional crosspolytope $(d \geq 4)$ with vertices $\{e_i, -de_i : i \in \{1, \ldots, d\}\}$. By calculating the volume of the d-simplices of a crosspolytope with vertices $\{e_i, -e_i : i \in \{1, \ldots, d\}\}$, it can be seen that all triangulations have cardinality 2^{d-1}. Adding a vertex $(d/(2d - 1) - \epsilon, \ldots, d/(2d - 1) - \epsilon)$ with $\epsilon > 0$ small makes the polytope stacked, having minimal triangulation of cardinality $d+1$. The same effect can actually be achieved using $0/1$ coordinates.
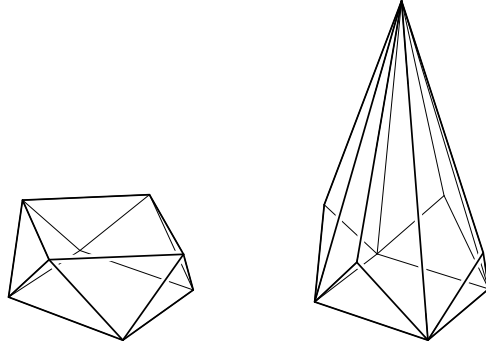
9

Figure 4: 4-antiprism (left) and the stacked polytope with a new vertex (right)

4   Minimal triangulations and coordinate changes.

In [8] the question was raised whether the minimal triangulation could be an invariant of the face lattice. We present the smallest example with variable size of minimal triangulation. This implies that any algorithm to compute minimal triangulations of 3–polytopes must take into account the coordinates of the polytope, not only its face lattice.

Proof: (of Theorem 1.3) The example is given by the 3–polytope with 10 vertices shown in Figure 5. The polytope is made by superimposing two copies of the combinatorial structure used in Lemma 2.2. If the points A, B, C, D are coplanar then one of the edges AB or CD cannot be in a triangulation at the same time. Hence by lemmas 2.1 and 2.2, any triangulation has at least 10 simplices.

On the other hand, one can arrange the z-coordinates of A, B, C and D in such a way that the edge AB is above the edge CD. This way it has a unique minimal triangulation with 9 simplices (A, C, D, H), (A, B, C, D), (A, B, E, F), (A, B, D, E), (C, D, I, J), (A, B, C, G), (C, D, H, I), (A, B, F, G), (B, C, D, J). The triangulation above is possible since the tetrahedron (A, B, C, D) exists and does not intersect the other simplices.

The fact that this example is minimal in dimension and number of vertices is a consequence of the following lemmas:

Lemma 4.1 *If a triangulation T of a 3-polytope P uses at most one interior edge, then for any polytope P' with the same face lattice, the set of simplices T' corresponding to T is still a triangulation.*

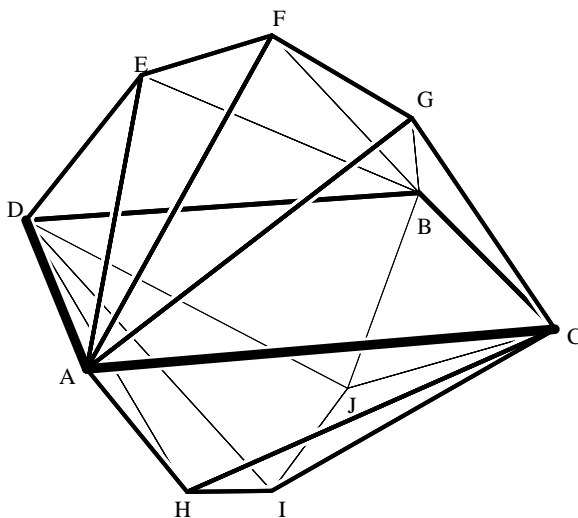*Proof:* We have to show two things: First, that T' still defines a simplicial

10

Figure 5: a view of the polyhedron with variable minimal triangulation.

complex, i.e., the tetrahedra of T' are all full-dimensional and intersect properly. Second, that all points of P' belong to one of the tetrahedra of T'.

Changing the coordinatization of P to the coordinates of P' while maintaining the face lattice gives us a bijection f of the vertices of P and of P'. In our proof we refer to the vertices of P' by primed letters, e.g., $f(v) = v'$ for v a vertex of P. The bijection f can be canonically extended to map T to T', or, for that matter, (abstract) faces of simplices in T to (abstract) faces of simplices in T'. Note that interior edges in P are always mapped to interior edges in P', and that edges on the boundary will stay on the boundary. For the proof of the lemma we will use that the vertices in P' lie in convex position and that the simplices in T' use at most one interior edge.

We prove now that the tetrahedra of T' are all full-dimensional. Any simplex in T' uses at most one interior edge. It has therefore vertices v_1, v_2 and v_3 which are connected by edges which lie on the boundary. If the fourth point had collapsed into a coplanarity with the other three points, it is forced, by $v_1 v_2$, $v_2 v_3$, $v_1 v_3$ being boundary edges, to lie in $conv(v_1, v_2, v_3)$. This contradicts its being in convex position. Hence all simplices of T' are full-dimensional.

Suppose two simplices in T' intersect improperly; that is, not in a common face. Such a bad intersection occurs if a pair of faces of T' forms a new minimal Radon partition (their relative interiors intersect) or in oriented matroid terminology a circuit [19]. In three dimensions the only five possible circuits are presented in Figure 6. Obviously, the cases (A), (B), and (C) change the face lattice when they occur because a vertex becomes an interior point.
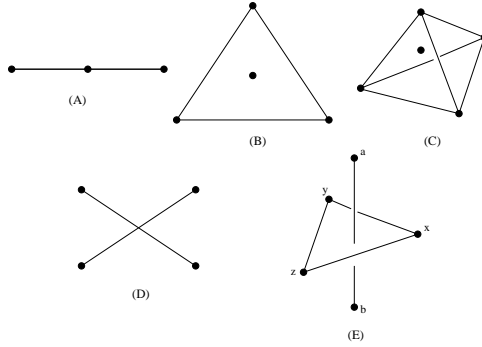
11

Figure 6: Possible bad intersections of faces of tetrahedra.

Consider now the pair of edges related to case (D). At least one of them is exterior since we only have one interior edge inside T. Their intersection point is therefore on the boundary of P'. Hence both edges have to be exterior edges lying in a common two-face of P'. This implies that their preimages in P already intersected improperly which is impossible since T was a triangulation. Case (D) is forbidden.

For case (E) note that if the edge $a'b'$ intersects the triangle $(x'y'z')$ in P', then $a'b'$ is interior; leaving the other three edges $x'y'$, $x'z'$, $y'z'$ to be exterior. This is also true for their preimages. This forces that, for the old coordinates when $relint(ab)$ and $relint(xyz)$ were disjoint, the vertices a, b were in the same side of the hyperplane spanned by x, y, and z; otherwise one of the edges xy, xz, yz is also interior. Finally, the vertices x, y, z form a cut of the graph of the polytope because their edges are all in the boundary. But then the face lattice was changed when going from the old to the new coordinates because the cut leaves the points a, b on the same side of the hyperplane but later they lie on opposite sides.

We have shown that the simplices in T' form a simplicial complex. Finally, we show that every point of P' belongs to one simplex of T'. Denote by $|T'|$ the point set $\bigcup_{\sigma' \in T'} \sigma$ (similarly for $|T|$). We can assume the bijection f is extended to a one-to-one point map $f : |T| = P \mapsto |T'|$, and it can be made continuous, e.g., by taking the canonical affine map on each simplex given by the values of the vertices of that simplex. Note that the boundary of P is mapped bijectively to the boundary of P'. If there were points in P' which are not any simplex of T', they have to be in the interior of P'. Hence, the set of all those points equals $int(P') \setminus |T'|$ which is an open set. Actually, so is $int(P') \cap |T'|$ since the boundary of $|T'|$ is by continuity and bijectivity of f the boundary of P'. This would imply to have a partition of $int(P')$ into two open sets which contradicts the connectivity of $int(P')$ in a topological sense. This concludes the proof of the lemma. $\square$

Lemma 4.2 *For convex 3–polytopes with up to nine vertices, the smallest possible number of tetrahedra in a triangulation is determined by the face lattice.*

Proof: First we will give necessary bounds for the number of vertices n and number of tetrahedra t in a minimal triangulation that is not present in every coordinatization.

From Lemma 2.1 we know that a triangulation of a polytope with less than $n-1$ tetrahedra has at most one interior edge. By the previous lemma, those triangulations are invariant under coordinate changes, hence $t \geq n-1$.

A triangulation that exists for all coordinatizations can be made as follows: select a vertex v, triangulate all facets of P that do not contain v. The tetrahedra of the triangulation of P are those formed by v and a triangle in the boundary not in a facet containing v. This procedure is called coning from vertex v or *pulling* triangulations, see [13] or [4, 14]. The number of tetrahedra in this triangulation is the number of triangles in the triangulated boundary minus the number of triangles in the faces that touch v. The first number is $2n-4$ by Euler's formula. The second number is the degree of v in the graph G_{face} consisting of the vertices of P and all edges between vertices lying in a common face. This graph contains the graph G' of any triangulation of the boundary. By Euler's formula, G' has $3n-6$ edges, hence the maximal degree in G' (and therefore in G_{face} is at least $\lceil 2 \cdot \frac{3n-6}{n} \rceil$.

We showed that pulling from a vertex v with maximal degree in G_{face} gives an invariant triangulation with $s(n) := 2n-4 - \lceil 2 \cdot \frac{3n-6}{n} \rceil$ tetrahedra. A non-invariant minimal triangulation must have $t < s(n)$ tetrahedra. Solving for the equations $n-1 \leq t < s(n)$ for $n = 4, \ldots, 9$ gives only the solution $n = 9$, $t = 8$.

Consider a convex polytope with nine vertices we show that either no triangulation of size eight is possible or it allows for an invariant triangulation of size eight or less. We can dispose of the case when P has a vertex v of degree three. The reason is such a vertex can be chopped off, the convex hull of the remaining eight points can be triangulated using an invariant triangulation of size no more than $s(8) \leq 7$ by our previous discussion. These at most seven tetrahedra plus the tetrahedron containing v give a size ≤ 8 invariant triangulation of P.

In the graph G_{face} there could be a vertex with degree more than five, but then our pulling procedure yields an invariant triangulation of at most eight simplices. Hence the only remaining case is that the triangulated planar graph of the triangulated boundary of P has only vertices of degree four or five. If n_i denotes the number of vertices of degree i, we have that twice the number of edges is $2 \cdot (3 \cdot 9 - 6) = 42 = 5n_5 + 4n_4$ and $n_5 + n_4 = 9$. The only non-negative integer solution is $n_5 = 6$ and $n_4 = 3$. From these degree restrictions and an easy case analysis, it can be seen that the only triangulated planar graph for the boundary is given in Figure 7:

We claim that all triangulations, that extend this triangulation of the boundary of P, will have 3 or more interior edges and then have nine or more tetrahedra by Lemma 2.1. Assume this is not true. Every boundary triangle is in
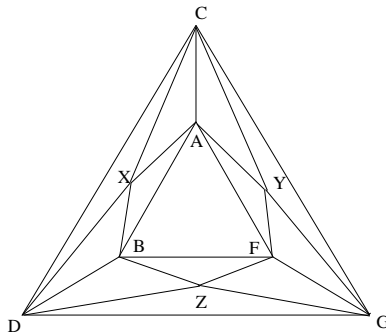
13

Figure 7: Unique triangulated planar graph with $n_5 = 6$ and $n_4 = 3$

a tetrahedron. The fourth point of the tetrahedron is in an adjacent triangle, otherwise instantly three interior edges are produced. But the tetrahedra that cover the triangles (ABX),(AFY), and (BFZ) will still induce three distinct interior edges. $\square$

References

[1] **Avis D. and ElGindy H.** *Triangulating point sets in space*, Discrete Comput. Geom. 2, (1987), 99-111.

[2] **Bern M. and Eppstein D.** *Mesh generation and optimal triangulations*, in *Computing in Euclidean Geometry* (D.Z Du and F.K Hwang eds.), World Scientific Publishing Co. 1992.

[3] **Bern M., Eppstein D. and Gilbert, J.**, *Provably good mesh generation.* J. Comput. Syst. Sci. 48, No.3, 1994, 384-409.

[4] **Billera L., Filliman P., and Sturmfels B.** *Constructions and complexity of secondary polytopes*, Adv. in Math., 83, 1990, 155-179.

[5] **Böhm J.**, private communication.

[6] **Chazelle B. Edelsbrunner H., Guibas L.J., Hershberger J.E., Seidel, R.and Sharir M.**, *Selecting heavily covered points.*, SIAM J. Comput. 23, No.6, 1994, 1138-1151.

[7] **Croft H.T. , Falconer K.J., and Guy R.K.** *Unsolved Problems in Geometry*, Springer-Verlag, New York, 1991.

[8] **De Loera J.A.** *Computing minimal and maximal triangulations of convex polytopes*, manuscript available at http://www.inf.ethz.ch/personal/deloera

14

[9] **Edelsbrunner H., Preparata F.P., and West D.B.** *Tetrahedrizing point sets in three dimensions*, Journal of Symbolic Computation, 10, 1990, 335-347.

[10] **Gritzmann P. and Klee V.** *On the complexity of some basic problems in computational convexity* POLYTOPES: abstract, Convex and Computational 374-466, (T. Bisztriczky et al. eds), Kluwer Academic Publishers, Netherlands, 1994.

[11] **Grüber P.** *Aspects of approximation of convex bodies* in Handbook of Convex Geometry, (P. Grüber and J. Wills eds.), volume A, 319-346, North-Holland, Amsterdam, 1993.

[12] **Hughes, R.B. and Anderson, M.R.** *Simplexity of the cube.* Discrete Math. 158, No.1-3, 1996, 99-150.

[13] **Lee C.W.** *Subdivisions and triangulations of polytopes* in Handbook of Discrete and Computational Geometry, 271-290, (Goodman J.E. and O'Rourke J. eds.) , CRC Press, New York, 1997.

[14] **Lee C.W.** *Regular triangulations of convex polytopes,* Applied Geometry and Discrete Mathematics- The Victor Klee Festschrift (P. Gritzmann and B. Sturmfels eds.) DIMACS series in Discrete Math. and Theoretical Comp. Science, 4, 1991, 443-456.

[15] **O'Rourke J.** *Art gallery theorems and algorithms.* Oxford University Press, New York, 1987.

[16] **Richter-Gebert, J.** *Finding small triangulations of polytope boundaries is hard*, manuscript available at http://www.inf.ethz.ch/personal/richter

[17] **Ruppert J. and Seidel R.** *On the difficulty of triangulating three-dimensional non-convex polyhedra*, Discrete Comput. Geom., 7 (1992), 227–253.

[18] **Schönhardt E.** *Über die Zerlegung von Dreieckspolyedern in Tetraeder*, Math. Ann., 98 (1928), 309–312.

[19] **Ziegler G.M.** *Lectures on Polytopes*, Springer-Verlag, New York, 1994.