



The complexity of finding small triangulations of convex 3-polytopes

Alexander Below,^a Jesús A. De Loera,^{b,*}¹
and Jürgen Richter-Gebert^{a,2}

^a *Institut für Theoretische Informatik, ETH-Zürich, Switzerland*

^b *Department of Mathematics, University of California, Davis, CA 95616, USA*

Received 29 June 2000

Abstract

We present a polynomial transformation from the satisfiability problem to the problem of finding a triangulation of a convex 3-polytope that uses a small number of tetrahedra. We also discuss other complexity results related to optimal triangulations.

© 2003 Elsevier Inc. All rights reserved.

1. Introduction

A *triangulation* of a d -dimensional convex polytope P is a set of d -simplices whose union is the polytope, their vertices are extreme points of P , and any two simplices in it intersect in a common (possibly empty) face. The *size* of a triangulation is the number of its full-dimensional simplices. In this paper we discuss the computational complexity of finding small size triangulations of a convex polytope. We discuss in particular the case of *minimal triangulations*, i.e., those with smallest possible size.

This geometric minimization problem arises in several contexts. For example, minimal triangulations of the d -cube have been extensively studied (see references in [11,15,16,19,25]) due to connections with the simplicial approximation of fixed points of

* Corresponding author.

E-mail addresses: below@inf.ethz.ch (A. Below), deloera@math.ucdavis.edu (J.A. De Loera), richter@ma.tum.de (J. Richter-Gebert).

¹ The author was partially supported by NSF grant DMS-0073815 and is also grateful to ETH-Informatik for its hospitality.

² Current affiliation: Technische Universität München, Zentrum Mathematik.

continuous maps [29]. Optimal size triangulations appear also in the polyhedral techniques in Algebraic Geometry [28]. Understanding minimal triangulations of convex polytopes is related to the problem of characterizing the f -vectors of triangulations of balls and polytopes (see open problems in [7]). In fact, the study of minimal triangulations of topological balls also received attention due to its connections to data structures, in the calculation of rotation distance of binary trees [27].

The computational geometry literature has several papers interested in finding triangulations of optimal size [4,13]. In 1992 Bern and Eppstein asked whether there is a polynomial time algorithm to compute a minimal triangulation of a 3-dimensional convex polytope (open problem 12 in [6, Section 3.2]). Our main result shows that, under the hypothesis $P \neq NP$, such an algorithm cannot exist:

Theorem 1.1. *Given a convex 3-polytope P and a positive integer K , deciding whether P has a triangulation of size K or less is an NP-complete problem.*

We give a transformation to the Satisfiability (SAT) problem (cf. [14]): given an instance S of C logical clauses in V boolean variables, is there a truth assignment to the variables such that all clauses are simultaneously satisfied? We will give a number K and construct the vertices of a convex 3-polytope, a *logical polytope*, of size polynomial in C and V (polynomial size pertains to the binary encoding length), which has a triangulation of size at most K if and only if there is a satisfying truth assignment. In fact, we can restrict our discussion to the special case of the SAT problem where each variable appears in three clauses, two of the times negated (see [14, p. 259]).

It is straightforward to derive some interesting conclusions from the main theorem. First, consider the logical polytopes we have constructed. If we could find the size s_{\min} for their minimal triangulations in polynomial time, either $s_{\min} \leq K$, in which case we have indeed a triangulation smaller or equal to K , or $K < s_{\min}$, in which case we can be sure there is no triangulation for the logical polytope of size K or less. This proves it must be NP-hard to find a minimal triangulation in dimension 3. Second, notice that the hardness result can be extended for polytopes of any fixed dimension that are given by its vertices. Simply note that by taking a pyramid over the logical polytopes we create a 4-dimensional logical polytope whose possible sizes of triangulations are those of the original 3-dimensional logical polytope. Thus, it is at least as hard to find the size of smallest triangulations in four dimensions as it is in three. Repeating the pyramid construction enough times, adding a new dimension each time, we have the following corollaries (the second result was obtained in [21] via a direct transformation to 3-SAT):

Corollary 1.2. (1) *Let d be a fixed positive integer with $d \geq 3$. Finding a minimal-size triangulation of a convex d -polytope is NP-hard.*

(2) *Let d be a fixed positive integer with $d \geq 4$. Finding a minimal-size triangulation of the boundary of a convex d -polytope is NP-hard.*

We should remark that since these results are given in fixed dimension it is unnecessary to specify how the polytope is presented (by its facets inequalities or by its vertices) as the

transformation to the vertex representation, or vice versa, takes polynomial time when the dimension is fixed [10].

It is useful to remind the reader of a few known results about the sizes of triangulations of polytopes and the relevant literature before we discuss the details. It is known that the sizes of triangulations for a d -dimensional polytope with n vertices lie between $n - d$ and $f_d(\partial C(n + 1, d + 1)) - d - 1$, where $f_d(\partial C(n + 1, d + 1))$ is the number of d -dimensional facets of a $(d + 1)$ -dimensional cyclic polytope with $n + 1$ vertices [22]. In particular, for 3-polytopes the possible number of tetrahedra ranges from $n - 3$ to $\binom{n}{2} - 2n + 3$. Both bounds are known to be tight for three dimensions [13]. It is also known that the size of a minimal triangulation of a convex 3-polytope must lie between $n - 3$ and $2n - 10$, when $n > 12$ [13]. That the upper bound is tight was proved in [27] using hyperbolic geometry. It is worth noticing at this point that the size of the constant K we construct in the proof of Theorem 1.1 satisfies the inequalities $n - 3 < K < 2n$. More will be said about this in the final section. Now we discuss the general structure and main ideas of the proof of Theorem 1.1.

Two elementary properties of triangulations will be useful to reach our goal:

- (1) Every boundary triangular facet F of a polytope is contained in exactly one tetrahedron of a triangulation. The fourth vertex of that tetrahedron is said to *triangulate* F .
- (2) Simplices of a triangulation cannot intersect in their relative interiors.

We will primarily see this behavior in triangles being pierced by an edge of the triangulation, a so-called *bad intersection*. Our proof combines techniques presented in the articles [5,23].

Consider a long vertex-edge chain on the boundary of a polytope such that the adjacent triangular faces all meet in two points a and b (see Fig. 1).

Lemma 1.3. *Let P be a convex 3-polytope such that the triangles (a, q_i, q_{i+1}) and (b, q_i, q_{i+1}) for $i = 0, \dots, m$ are among its facets, with the additional restriction that $\text{conv}\{a, b\} \cap \text{conv}\{q_0, \dots, q_{m+1}\} = \emptyset$. Let n be the number of vertices of P .*

Then, for each triangulation of P that does not use the (interior) edge (a, b) the number of tetrahedra is at least $n + m - 3$.

The proof of this lemma, which can be found in [5, Section 2], involves an inductive argument for counting interior edges of a triangulation and the use of Euler's relations for




Fig. 1. The vertex-edge chain of Lemma 1.3.

2-spheres and 3-balls that yield a simple formula between the number of interior edges and the number of tetrahedra of a triangulation.

For us, the quantity $n + m - 3$ will be a relatively large number, forcing that a triangulation not using (a, b) cannot be small. When using the edge (a, b) we can triangulate $P_{\text{chain}} = \text{conv}\{a, b, q_0, \dots, q_{m+1}\}$ using the $m + 1$ tetrahedra (a, b, q_i, q_{i+1}) for $i = 0, \dots, m$. Call Q the (non-convex) polytope we get after cutting all these tetrahedra out of P . Let n_Q denote the number of vertices of Q . Note that $n = n_Q + m$. Suppose the number of tetrahedra in any triangulation of Q (if there is one at all) is bounded above by some number t . Then we can bound the size of a minimal triangulation of P using (a, b) by $t + m$. Note that t does not depend on the length m of the vertex-edge chain. Hence, by choosing m large enough (leaving Q as it is) makes

$$\text{size of minimal triangulations} \leq t + m < n_Q + 2m - 3 = n + m - 3,$$

and any close-to-minimal triangulation will *have to use* (a, b) . This argument still holds when we have many vertex-edge chains of the same length m present in other parts of the boundary of the polytope P . If m is large enough, a small triangulation is forced to use the edges (a, b) of *all* these vertex-edge chains.

We also use the famous non-convex *Schönhardt polytope* [6,17,18,20,23,24]. Roughly speaking, a Schönhardt polytope can be obtained by “twisting” the top face of a triangular prism in a clockwise direction (see Fig. 2). The three quadrangular sides are then broken up and “bent in,” thus creating the non-convex (*reflex*) edges (B_i, A_{i+1}) that we call *diagonals*. The resulting polytope is non-convex and we distinguish the two triangular faces without reflex edges: *the bottom* (A_1, A_2, A_3) , and the *top or skylight* (B_1, B_2, B_3) . Whenever dealing with vertices of a Schönhardt polytope, abusing the notation, by an index $i + 1$ we mean $(i \bmod 3) + 1$. For example, $3 + 1$ gives 1.

The reader can easily verify that Schönhardt polytopes cannot be triangulated using only its six vertices. Imagine the Schönhardt polytope is glued along its bottom face to a bigger polytope. Again, the resulting non-convex polytope can only be triangulated if its top triangle (the skylight) is *visible* by another vertex (we will rigorously define visible




Fig. 2. A Schönhardt polytope.




Fig. 3. The visibility cone and an invisible vertex.




Fig. 4. Cupola.

later, but it does correspond to the intuition that every point of the skylight can be joined to the vertex by a straight line segment. We will show that the triangular cone defined by the planes containing the faces (B_i, B_{i+1}, A_{i+1}) contains exactly the points that can view the skylight. For this reason we will call this cone the *visibility cone*.

Now we convexify the Schönhardt polytope by attaching three circular vertex-edge chains opposite to the concavities. Thus, we create a convex polytope that satisfies the hypothesis of Lemma 1.3, and that we will call a *cupola*, see Fig. 4. The cupola is usually glued along its bottom face to a bigger convex polytope and obtain a convex polytope P . We can combine what we know about vertex-edge chains and about Schönhardt polytopes. Namely, in order to have a small triangulation of P , the three diagonals of the Schönhardt polytope inside the cupola have to be used. But then, the vertex triangulating the skylight of the cupola must not be obstructed from seeing the skylight by the diagonals. Hence, the vertex has to lie in the visibility cone of the cupola.

In [23] Ruppert and Seidel used SAT to prove that it is NP-complete to decide whether a non-convex polyhedron admits a triangulation. Their constructions used Schönhardt polytopes, and in particular their visibility cones, to do the transformation. In our case, because we need convexity, we glue cupolas, instead of Schönhardt polytopes. They are glued to a bigger *frame* polytope along their bottom faces. Similar to [23], we have variable cupolas and clause cupolas. The visibility cones of the variable cupolas contain only two *truth-setting* vertices, one for *false* and one for *true*. The visibility cones of the clause cupolas contain as many *literal* vertices as there are literals in the logical clause. Each

variable must choose between a “true” or “false” value. Inside each clause at least one variable will be chosen to be true (to satisfy the clause). We model these logical choices by the geometric choices of which vertex in the visibility cone of a (variable/clause) cupola is used to triangulate the skylight. In addition, our polytope satisfies some *blocking* conditions: the tetrahedron spanned by the top face of a clause cupola and a literal vertex coming from a negated variable X_i will improperly intersect the tetrahedron spanned by the top face of the cupola of variable X_i and the truth-setting vertex corresponding to *true*. In this way the choices made for the truth values of the variables and for the literals satisfying the clauses will be *consistent*. We will call our polytope the *logical polytope* because it comes from a logical formula.

For the polynomial transformation (from SAT) we need to give an algorithm to compute the coordinates of the vertices of the logical polytope. The binary encoding length of the polytope, as well as the runtime of the algorithm, have to be polynomial in the encoding length of the SAT instance. Each step of the construction will be polynomial; this is a delicate point in the formalism of our argument. We apply a sequence of these constructions (polynomially many). The coordinates of the vertices of the polytope are potentially singly-exponential, but their binary encoding length is guaranteed to be polynomial.

Elementary steps of the construction include operations such as taking the join of two or three points, intersecting planes and lines, putting points on polynomial curves, etc. The coordinates of the resulting construction elements are therefore polynomials in coordinates of the input elements. On the other hand, we will have requirements on the positions of the points with respect to some planes or other points on lines etc. All these conditions can be formulated as strict polynomial inequalities in coordinates of the construction elements. An essential element of our construction is that our systems of strict polynomial inequalities will depend on *one single parameter* ϵ . All these polynomial inequalities are satisfied at $\epsilon = 0$, but an additional requirement for us is $\epsilon > 0$. The following lemma describes a polynomial algorithm to find a number ϵ_0 such that all $0 < \epsilon \leq \epsilon_0$ solve the inequality system.

Lemma 1.4. (1) Suppose $p(\epsilon) = a_d\epsilon^d + \dots + a_1\epsilon + a_0$ is a polynomial with $p(0) > 0$. Let

$$\epsilon_0(p) := \min\left(1, \frac{a_0}{2(|a_1| + \dots + |a_d|)}\right).$$

Then for $0 \leq \epsilon \leq \epsilon_0(p)$ we have $p(\epsilon) > 0$.

Hence, the construction of ϵ_0 can be done in time polynomial in the encoding length of the coefficients of p , and ϵ_0 has polynomial encoding length.

(2) p_1, \dots, p_l are univariate polynomials such that $p_1(0) > 0, \dots, p_l(0) > 0$ then there is a rational number, $\epsilon_0 > 0$, such that $p_1(\epsilon) > 0, \dots, p_l(\epsilon) > 0$ for all $0 < \epsilon \leq \epsilon_0$. Moreover, the encoding length of ϵ_0 is bounded by a polynomial in the encoding length of the coefficients of p_1, \dots, p_l .

Proof. For $0 \leq \epsilon \leq 1$ we have that $a_i\epsilon^i \geq -|a_i|\epsilon$. The reason is that for $a_i \geq 0$, $a_i\epsilon^i \geq 0 \geq -|a_i|\epsilon$, and for $a_i < 0$, $a_i\epsilon^i > a_i\epsilon = -|a_i|\epsilon$. Hence, for $0 \leq \epsilon \leq \epsilon_0(p)$

$$p(\epsilon) \geq \sum_{i=1}^d -|a_i|\epsilon + a_0 > -\sum_{i=1}^d |a_i| \frac{a_0}{2 \sum_{i=1}^d |a_i|} + a_0 > 0.$$

For the second part, take the value $\epsilon_0(p_1, \dots, p_r) := \min(\epsilon_0(p_1), \dots, \epsilon_0(p_r))$. Now all the conditions are simultaneously satisfied. \square

Of course, in general the real solutions of a multivariate system of inequalities coming from geometric requirements may be empty, but our steps of construction reduce everything to sequentially solving easy univariate systems of inequalities.

Here is the organization of our paper. In Section 2 we discuss useful properties of Schönhardt polytopes and of cupolas. Later in the section we explain how to construct and glue cupolas that have a prescribed visibility cone and how to construct visibility cones that fit our purposes. The polynomial transformation to SAT is presented in Section 3. From a given SAT instance we construct a frame polytope to which we then glue the cupolas. In the final Section 4, as a complement of our main theorem, we present a family of polytopes (the so-called stacked polytopes) for which the decision problem of Theorem 1.1 can be solved in polynomial time. We end the paper with a few open questions.

2. Basic building blocks

We recall the notion of *beyond a face* (see [30, p. 78]): A point p is *beyond* a face F of a polytope P if it (strictly) violates all inequalities defining facets of P containing F , but it strictly satisfies all other inequalities that define other facets of P . The polytope $P_{\text{beyond } F}$ is the (closure of the) set of all points beyond F . We denote by $P \setminus F$ the polyhedron defined by all facet-defining inequalities that do not hold with equality for all points in F . This is exactly $P \cup P_{\text{beyond } F}$. In our constructions we will often put one or more points beyond some face, and then take the convex hull. This will only destroy the facets containing this face, and introduce new ones containing the new points. We will say we *attach* one polytope P to another Q *along* a facets F_P of P and F_Q of Q if $P \subseteq Q_{\text{beyond } F_Q}$ and $Q \subseteq P_{\text{beyond } F_P}$. It is important to observe that the convex hull of their union contains both the face lattices of P and Q without, of course, F_P and F_Q .

2.1. The Schönhardt polytope

Let us turn to a well-known example of a non-convex non-triangulable polytope, the so-called Schönhardt polytope (named after its first occurrence in [24]. See also [20]). For the notion of non-convex polytope and what it means to triangulate them we refer to [9].

Definition 2.1. A *Schönhardt* polytope (Fig. 2) is a non-convex polytope with six vertices $A_1, A_2, A_3, B_1, B_2,$ and B_3 and facets $(A_1, A_2, A_3), (B_1, B_2, B_3), (A_1, B_1, A_2), (B_1, A_2, B_2), (A_2, B_2, A_3), (B_2, A_3, B_3), (A_3, B_3, B_1),$ and (B_3, B_1, A_1) . At exactly the edges $(B_1, A_2), (B_2, A_3), (B_3, A_1)$ the corresponding facets are to span an interior angle greater than π (the edges are said to be *reflex*). These edges are called the *diagonals* of the Schönhardt polytope. The top face (B_1, B_2, B_3) is called the *skylight* of the Schönhardt polytope.

Six points are said to be in *Schönhardt position* if they are the vertices of a Schönhardt polytope. We say that the skylight is *visible* from a point x (or x is *able to see* the skylight,

or x is a *viewpoint* of the skylight) if the tetrahedron spanned by x and the skylight does not intersect any of the diagonals in their relative interior. The *visibility cone* of the Schönhardt polytope is the triangular cone bounded by the planes $B_1B_2A_2$, $B_2B_3A_3$, and $B_3B_1A_1$. See Fig. 3.

The use of the word “skylight” is motivated by the idea that the skylight triangle is a glass window and light comes through it illuminating the interior of the Schönhardt polytope defining a cone of light. It is obvious that this non-convex polytope cannot be triangulated (without adding new points): The fourth point of the tetrahedron containing the skylight must be one of A_1 , A_2 , or A_3 , but the diagonals “obstruct the view” of the skylight from these vertices.

It is our intention to patch the sides of the Schönhardt polytope with vertex-edge chains in order to convexify it (and then glue it to a frame polytope). According to Lemma 1.3, a small triangulation of this convex polytope must necessarily contain the diagonals. In this case, the fourth point of the tetrahedron containing the skylight also has to be able to see the skylight. We will show where to place the vertex-edge chains in order for them not to be visible from the skylight. Hence, the triangulating vertex has to lie beyond (A_1, A_2, A_3) , we will show that it has to lie in the visibility cone.

Lemma 2.2. *Let $A_1, A_2, A_3, B_1, B_2, B_3$ be six points in Schönhardt position. We denote by $C_{A,B}$ the convex hull of the six points. Then*

- (1) *All orientations of simplices spanned by four of these six points are determined up to one global sign change. As a consequence, the six points are in convex position, and their convex hull $C_{A,B}$ is an octahedron that has (A_1, A_2, A_3) and (B_1, B_2, B_3) as facets and it has edges (A_i, B_{i+1}) ($i = 1, 2, 3$).*
- (2) *There are no points that can see the skylight (B_1, B_2, B_3) and, at the same time, (i) are beyond either of the edges (A_i, B_{i+1}) of $C_{A,B}$, and (ii) are on the side of the plane $B_1A_2B_3$ opposite to B_2 or similarly for the analogous planes $B_1A_3B_2$, $B_2A_1B_3$ and the points B_3, B_1 , respectively.*
- (3) *The visible points beyond the facet (A_1, A_2, A_3) of $C_{A,B}$ are exactly the points that are also in the visibility cone of the Schönhardt polytope.*

In what follows we will use the language of oriented matroids. For the theory of oriented matroids we refer to [8,30]. Here we only sketch the necessary definitions and how they are related to the notion of visibility. The orientation of a simplex (x_1, x_2, x_3, x_4) , is defined as

$$[x_1, x_2, x_3, x_4] = \text{sign det} \begin{pmatrix} x_1 & x_2 & x_3 & x_4 \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

All such orientations make up the *chirotope* of an oriented matroid (see [8, p. 123]).

Given the oriented matroid of points x_1, \dots, x_n in d -space, its circuits are functions $C : \{x_1, \dots, x_n\} \mapsto \{+, -, 0\}$ that correspond to so-called minimal Radon partitions. This means that the convex hulls of $C^+ = \{x_i \mid C(x_i) = +\}$ and $C^- = \{x_i \mid C(x_i) = -\}$ intersect

in their relative interiors, and C^+ and C^- are minimal at that. It is easy to check that the function

$$C(x) = \begin{cases} (-1)^i \cdot \overbrace{[x_1, \dots, x_{d+1}]}^{\text{omit } x_i} & \text{if } x \in \{x_1, \dots, x_{d+1}\}, \\ 0 & \text{otherwise,} \end{cases}$$

defines a circuit if it is not identical 0. In fact, all circuits can be obtained this way. We will compute circuits to use an argument of the following form: x does not see the skylight if and only if there is a circuit such that the positive part is one of the diagonals and negative part is the set containing x and a subset of vertices of the skylight. Since then the tetrahedron spanned by x and the skylight is pierced by the diagonal.

Important tools to compute simplex orientations are the Grassmann–Plücker relations (see [8, Section 2.4]): For points a, b, x_1, \dots, x_4 they state that the set of signs

$$\{[a, b, x_1, x_2] \cdot [a, b, x_3, x_4], -[a, b, x_1, x_3] \cdot [a, b, x_2, x_4], [a, b, x_1, x_4] \cdot [a, b, x_2, x_3]\}$$

is either identical 0 or contains both a + and a -. The typical use of the Grassmann–Plücker relations is to deduce one orientation when the others are known. We can read the orientations of some of the different tetrahedra from two-dimensional projections (drawings) of the point configurations as in Fig. 2. We use a left-handed rule system, i.e., we decide whether the triangle (x_1, x_2, x_3) is oriented counterclockwise (+) or not (-), also if x_4 is on our side of the plane spanned by $x_1, x_2,$ and x_3 (+) or not (-), and multiply these two signs to obtain the orientation $[x_1, x_2, x_3, x_4]$.

Proof of Lemma 2.2. (1) In a Schönhardt polytope, the simplices (A_1, A_2, A_3, B_1) and (A_1, A_2, A_3, B_2) have the same orientation since edges (A_1, A_2) and (A_2, A_3) are both incident to facet (A_1, A_2, A_3) and they are both non-reflex edges.

By the above argument, going around the boundary of a Schönhardt polytope, keeping in mind which edges are reflex, we can determine the orientation of 12 simplices up to one global sign change (there are 12 edges). But there are $\binom{6}{4} = 15$ simplices formed by the vertices of the Schönhardt polytope. The remaining three simplices are $(A_1, A_2, B_2, B_3), (A_2, A_3, B_1, B_3), (A_1, A_3, B_1, B_2)$. The signs are determined by the following Grassmann–Plücker relations: for (A_1, A_2, B_2, B_3) take $a = A_1, b = A_2, x_1 = A_3, x_2 = B_1, x_3 = B_2, x_4 = B_3$ (the other two by circular index shift). Then

$$\{- \cdot ?, - \cdot - \cdot \cdot +, - \cdot - \cdot -\} \supseteq \{+, -\}$$

the equation forces $[A_1, A_2, B_2, B_3] = +$. From the chirotope information it is easy to check that all vertices are in convex position (see description of how to read the facets of the convex hull from the chirotope in [8, Chapter 3]), and that their convex hull $C_{A,B}$ is indeed an octahedron.

(2) We will show that if a point x lies beyond A_1B_2 of $C_{A,B}$, on the side of $B_1A_2B_3$ opposite to B_2 , then (B_1, A_2) and the triangle (B_2, B_3, x) form a minimal Radon partition in the set of vertices $A_1, A_2, A_3, B_1, B_2, B_3,$ and x , hence have an interior point in common. This means x cannot see the skylight. For this, we compute the following orientations:

$$\begin{aligned}
 -[B_1, B_2, B_3, x] &= +, & \text{since } (B_1, B_2, B_3) \text{ is a facet of } C_{A,B} \setminus (A_1, B_2), \\
 +[A_2, B_2, B_3, x] &= +, & \text{since } (A_2, B_2, B_3) \text{ is a facet of } C_{A,B} \setminus (A_1, B_2), \\
 -[A_2, B_1, B_3, x] &= -, & \text{from the assumption on } x, \\
 +[A_2, B_1, B_2, x] &= -, & \text{from the Grassmann–Plücker relation below,} \\
 -[A_2, B_1, B_2, B_3] &= -, & \text{from part (1).}
 \end{aligned}$$

The necessary Grassmann–Plücker relation is the one with $a = B_1$, $b = B_2$, $x_1 = B_3$, $x_2 = A_1$, $x_3 = A_2$, and $x_4 = x$ such that

$$\{- \cdot ?, - \cdot - \cdot -, - \cdot +\} \supseteq \{+, -\}$$

forces $[B_1, B_2, A_2, x] = -$.

(3) If x is in the visibility cone V , then it is, by part (2) of this lemma, on the same side as B_3 with respect to the plane $B_1A_2B_2$. Hence, A_2 is on opposite side of B_3 with respect to the plane B_1B_2x . Therefore, the relative interior of the convex hull of B_1 and A_2 lies strictly on one side of the plane B_1B_2x , and the tetrahedron (B_1, B_2, B_3, x) on the other side of this plane. Therefore, those two point sets cannot have points in common. By symmetry it follows that the other two diagonals do not obstruct any point of V from seeing the skylight either.

Assume now that a point x is beyond face (A_1, A_2, A_3) , but outside of V , i.e., for instance on the A_1 side of the plane $B_1B_2A_2$. We claim that the pair $\{B_1, A_2\}, \{B_2, B_3, x\}$ forms a circuit in the oriented matroid of the point configuration of the vertices of $C_{A,B}$ and x . This means that the triangle (B_2, B_3, x) is pierced by the diagonal (B_1, A_2) in the relative interior, hence x is not visible.

$$\begin{aligned}
 -[B_2, B_3, A_2, x] &= -, & \text{since } (B_2, B_3, A_1) \text{ is a facet of } C_{A,B} \setminus (A_1, A_2, A_3), \\
 +[B_1, B_3, A_2, x] &= +, & \text{from the Grassmann–Plücker relations below,} \\
 -[B_1, B_2, A_2, x] &= +, & \text{from the assumption on } x, \\
 +[B_1, B_2, B_3, x] &= -, & \text{since } (B_1, B_2, B_3) \text{ is a facet of } C_{A,B} \setminus (A_1, A_2, A_3), \\
 -[B_1, B_2, B_3, A_2] &= +, & \text{from part (1).}
 \end{aligned}$$

In this case, we have to apply the Grassmann–Plücker relations twice to get $[B_1, B_3, A_2, x] = +$. First, we deduce $[A_1, A_2, B_3, x] = -$ from the Grassmann–Plücker relation with $a = A_1$, $b = A_2$, $x_1 = A_3$, $x_2 = x$, $x_3 = B_2$, $x_4 = B_3$:

$$\{- \cdot -, - \cdot + \cdot ?, + \cdot +\} \supseteq \{+, -\}.$$

Now we use this orientation to formulate $a = A_2$, $b = B_3$, $x_1 = A_1$, $x_2 = B_1$, $x_3 = B_2$, $x_4 = x$:

$$\{+ \cdot -, - \cdot + \cdot ?, - \cdot +\} \supseteq \{+, -\}$$

in order to get the desired $[A_2, B_3, B_1, x] = -$. \square




Fig. 5. A cupola as part of a larger convex polytope P .

2.2. The cupola

Definition 2.3. A polytope C is called an m -cupola (or for short a cupola) if it has the following properties:

- (1) the vertices of C are $A_1, A_2, A_3, B_1, B_2, B_3$, and $q_k^{i,j}$ ($k = 0, \dots, m + 1, (i, j) \in \{(1, 2), (2, 3), (3, 1)\}$), where $q_0^{i,j} = A_i$ and $q_{m+1}^{i,j} = B_j$ are identified;
- (2) the vertices $A_1, A_2, A_3, B_1, B_2, B_3$ are in Schönhardt position, and (A_1, A_2, A_3) (the bottom facet) and (B_1, B_2, B_3) (the skylight) are facets of C ;
- (3) the other facets are $(B_i, q_k^{i,j}, q_{k+1}^{i,j})$ and $(A_j, q_k^{i,j}, q_{k+1}^{i,j})$ for $k = 0, \dots, m + 1, (i, j) \in \{(1, 2), (2, 3), (3, 1)\}$;
- (4) the vertices $q_k^{1,2}$ ($k = 1, \dots, m$) lie on the side of the plane $B_1 A_2 B_3$ opposite to B_2 . Similar conditions must hold for $q_k^{2,3}$ and $q_k^{3,1}$.

Proposition 2.4. Let C be an m -cupola which is part of a larger polytope P (i.e., the set $Q = P - C$ is a convex polytope and Q and C share the common facet (A_1, A_2, A_3)). Let n be the number of vertices of P and n' be the number of vertices of Q .

If T is a triangulation of P with the property that the fourth point of the tetrahedron containing the skylight of C is not in the visibility cone of C , then there are at least $n + m - 3 = n' + 4m$ tetrahedra in the triangulation.

Proof. If the vertex triangulating the skylight of C is a vertex on a vertex-edge chain of C , then it does not see the skylight by Definition 2.3(4) and Lemma 2.2(2). If it is in Q instead, then it has to be beyond the face (A_1, A_2, A_3) of C . Hence, by Lemma 2.2(3) it cannot see the skylight either. Therefore, the triangulation T does not use one of the diagonals. By Lemma 1.3 the number of tetrahedra is at least $n + m - 3$. Since by construction $n = n' + 3(m + 1)$, the number of tetrahedra is also at least $n' + 4m$. \square

This means that any “badly” triangulated cupola gives rise to a high number of tetrahedra in the triangulation. The way we are going to use this is to construct many cupolas and apply the proposition to all of them.

2.3. Constructing a cupola from a visibility cone

In this subsection we will show that cupolas can be attached to any face of a frame polytope using intermediate polytopes and that the visibility cone can be prescribed. The following theorem does not have the full strength we need for the construction. In Section 3, we will use a slightly stronger version which we will present at the end of this section. However, this theorem captures the main ideas used to construct a cupola.

Theorem 2.5 (Cupola construction from a given visibility cone). *Let F be a facet of a 3-polytope P , and V be a triangular cone such that $F \cap V$ is a triangle in the relative interior of F , and m be a positive integer. Then there is an m -cupola C beyond F of P such that P is beyond (A_1, A_2, A_3) of C and such that V is the visibility cone of C . Moreover, the input length of C is polynomial in the input lengths of P , V , and m .*

Before we come to the proof, we will exhibit a necessary condition of the visibility cone V of a cupola C and the facet the cupola is being glued upon. It will imply that we cannot directly attach a cupola to a face (as in [23]), but we have to construct an intermediate polytope first.

Lemma 2.6. *Let $A_1, A_2, A_3, B_1, B_2, B_3$ be vertices in Schönhardt position. Define l_1 to be the intersection line of planes $B_3B_1A_1$ and $B_1B_2A_2$, lines l_2 and l_3 are defined accordingly (Fig. 6, note that they contain the extreme rays of V). The lines l_1, l_2 , and l_3 intersect the relative interior of the bottom face (A_1, A_2, A_3) of a cupola C . The intersection points D_1, D_2 , and D_3 are forced to have the following collinearities: $A_1D_1D_2$, $A_2D_2D_3$, and $A_3D_3D_1$.*

Proof. l_1 enters the Schönhardt polytope S in point B_1 , runs along facet (A_1, B_1, B_3) until it reaches the edge (A_1, B_3) where it goes into the interior of S . Then the relative interior of (A_1, A_2, A_3) contains the point D_1 . In this way, D_1, D_2, A_2 are all on the planes $A_1A_2A_3$ and $B_1B_2A_2$. \square




Fig. 6. Collinearity condition in the base triangle of a cupola.

Proof of Theorem 2.5. We proceed in three steps. The lines l_1, l_2, l_3 are defined as in Lemma 2.6.

The bottom triangle (A_1, A_2, A_3) . We will now construct an intermediate polytope beyond F which will have a triangular facet (A_1, A_2, A_3) which is (1) parallel to F , and which is (2) intersected by the cone V in a triangle (D_1, D_2, D_3) in the relative interior such that (3) the collinearity condition from Lemma 2.6 holds.

To do this, we place a plane H parallel to and slightly above F such that the intersection points D_i of H and l_i ($i = 1, 2, 3$). Also H has to be so close to F that the l_i do not cross between H and F . By prolonging the line segment D_3D_1 slightly beyond D_1 (staying in $P \setminus F$) we obtain point A_1 , analogously construct A_2 and A_3 (Fig. 7). Taking the convex hull of F and the points $A_1, A_2,$ and A_3 gives then the intermediate polytope, whose face (A_1, A_2, A_3) has the collinearity condition. These constructions are polynomially constructible in the sense of Lemma 1.4.

The frame of the cupola. As in the construction of the bottom facet (A_1, A_2, A_3) , we place a plane H' parallel and slightly above this facet. The intersection of H' and V is the triangle (B_1, B_2, B_3) (B_1 is on the same extreme ray of V as D_1 and so on). See Fig. 8.

It is clear from the construction that triangles (D_1, D_2, D_3) , (A_1, A_2, A_3) , and (B_1, B_2, B_3) are parallel and all oriented the same way. Therefore, it is not hard to check that the points $A_1, A_2, A_3, B_1, B_2,$ and B_3 are vertices of a Schönhardt polytope whose




Fig. 7. Building the intermediate polytope for the cupola.




Fig. 8. Building the frame of a cupola.




Fig. 9. Construction of the vertex-edge chain.

visibility cone is V . Polynomiality of this part of the construction follows from Lemma 1.4 as well.

Attaching the vertex-edge chains. Now that the frame of a cupola is done, i.e., the vertices A_1, \dots, B_3 are in Schönhardt position, it remains to patch the key structures of Lemma 1.3, the vertex-edge chains $q_i^{j,k}$ ($i = 1, \dots, m, (j, k) \in \{(1, 2), (2, 3), (3, 1)\}$), to the sides of the frame $\text{conv}(P \cup \{A_1, \dots, B_3\})$.

Given triangular faces (a, q_0, q_{m+1}) and (b, q_0, q_{m+1}) of a convex polytope P and a plane G which (strictly) separates points q_0 and q_{m+1} . We claim that we can construct points q_1, \dots, q_m beyond the edge (q_0, q_{m+1}) of P such that the convex hull of $P \cup \{q_i\}$ has the properties of Lemma 1.3 and such that the points q_1, \dots, q_m lie on the same side of G as q_0 . Moreover, the input length of the constructed points is polynomially bounded in the input length of P and G .

By applying our claim three times, we will conclude our proof. The vertices $q_i^{j,j+1}$ are placed beyond edge (A_j, B_{j+1}) , vertices B_j and A_{j+1} take the roles of a and b , G is the plane spanned by B_j, A_{j+1} , and B_{j+2} . It is easy to check that this is exactly what we want for Lemma 1.3 and for the cupola conditions.

Now we prove the claim. We will put the points q_i ($i = 1, \dots, m$) on a parabola segment, beyond the edge (q_0, q_{m+1}) . Let H be a plane containing q_0 and q_{m+1} which also intersects the interior of P . This plane has the property that it contains points beyond edge (q_0, q_{m+1}) . It is constructible in polynomial time. Let v be the sum of the two normal vectors of planes aq_0q_{m+1} and bq_0q_{m+1} , and H the plane containing q_0 and q_{m+1} parallel to v .

Let now D be the intersection point of G and (q_0, q_{m+1}) . Let w be a vector of direction of the intersection line of G and H , such that starting at D it is pointing out of P . Now let $E_\epsilon = D + \epsilon w$ for $\epsilon > 0$ to be specified later. For small ϵ , E_ϵ is beyond (q_0, q_{m+1}) . Hence, the parabola defined according to Lemma 2.7, stated and proved below, by $p(0) = q_0$, $p(1/2) = E_\epsilon$, and $p(1) = q_{m+1}$ lies entirely in H , and for arguments between 0 and 1 passes just beyond (q_0, q_{m+1}) . Let $q_i = p(i/(4m))$ for $i = 1, \dots, m$. For small ϵ all those points are beyond (q_0, q_{m+1}) and on the same side of G as q_0 (polynomial conditions, use Lemma 1.4). Also, they are in convex position such that the convex hull of $P \cup \{q_1, \dots, q_m\}$ has exactly the required face lattice. \square

Lemma 2.7. *Let p_0, p_1, p_2 be three non-collinear points in \mathbb{R}^3 and t_0, t_1, t_2 be three distinct real numbers. Then there is a unique curve $p: \mathbb{R} \rightarrow \mathbb{R}^3$ such that $p_0 = p(t_0)$, $p_1 = p(t_1)$, and $p_2 = p(t_2)$ which is quadratic in every coordinate. Furthermore, all points on $p(t)$ are in the plane spanned by p_0, p_1 , and p_2 , and they are in convex position. Also a plane containing $p(r)$ and $p(l)$ for some $r \neq l$ which does not contain all of p has all points between l and r on one of its sides and all other points on the other side.*

Proof. Since p_0, p_1, p_2 have to be on the t_0, t_1, t_2 positions of the curve

$$p(t) = \begin{pmatrix} a_x + b_x t + c_x t^2 \\ a_y + b_y t + c_y t^2 \\ a_z + b_z t + c_z t^2 \end{pmatrix},$$

we have the condition

$$\begin{pmatrix} a_x & b_x & c_x \\ a_y & b_y & c_y \\ a_z & b_z & c_z \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ t_0 & t_1 & t_2 \\ t_0^2 & t_1^2 & t_2^2 \end{pmatrix} = \begin{pmatrix} \vdots & \vdots & \vdots \\ p_0 & p_1 & p_2 \\ \vdots & \vdots & \vdots \end{pmatrix}.$$

By the non-singularity of the Vandermonde matrices, there is a unique solution to a, b, c , given the p_i and t_i .

The curves which are quadratic in every coordinate are linear transforms of the moment curve $m(t) = (1, t, t^2)$. This curve lies entirely in the $x = 1$ plane, is convex, and has the condition that it intersected by each plane at most twice (or it is in this plane). All these properties are invariant under linear transformations. \square

Proposition 2.4 stated that we get a large triangulation if we triangulate the skylight of a cupola by a vertex outside the visibility cone. Now we want to estimate how much smaller a triangulation is if we use a vertex v in the visibility cone instead. We give a relatively small triangulation of the cupola and of the space between the bottom face (A_1, A_2, A_3) of the cupola and the triangular face F of P with the help of the vertex v .

Proposition 2.8. *Let F be triangular face of a polytope P , and C an m -cupola attached to it according to Lemma 2.5. Let v be a vertex of P in the visibility cone of C . Then there is a triangulation of $\text{conv}(\{v\}, F, C)$ with at most $3m + 16$ tetrahedra.*

Proof. First of all, we triangulate along the vertex-edge chains using the tetrahedra $(B_i, A_{i+1}, q_k^{i,i+1}, q_{k+1}^{i,i+1})$ for $i = 1, 2, 3$, and $k = 0, \dots, m$.

After removing these tetrahedra, we are left with the union of the Schönhardt polytope on the vertices $A_1, A_2, A_3, B_1, B_2, B_3$, and the convex polytope $\text{conv}(\{v\}, F, (A_1, A_2, A_3))$. This is a non-convex polytope with all edges, except the diagonals, being convex (easy conclusion from Lemma 2.2 and the construction). Since the specified vertex v is inside the visibility cone, it sees all facets of this polytope, except the three facets it is incident to, from the *interior*. In particular, we can form tetrahedra of all these facets and v and none of them intersect badly. They are 7 tetrahedra for the facets of the Schönhardt polytope (since we do not count the bottom face) and at most 6 for the rest (the convex hull of F and

(A_1, A_2, A_3) has—by a planar graph argument—at most $2 \cdot 6 - 4 = 8$ facets, subtracting 2 for F and (A_1, A_2, A_3) gives 6). \square

It is this $3m$ in contrast to the $4m$ in Proposition 2.4 which makes this way of triangulating optimal for large m . We give more details on the use of these propositions in Section 3 when we actually construct a polytope with many cupolas.

2.4. Constructing a visibility cone

In order to use the cupola as a basic building block, we need to have a visibility cone that contains a specified set of vertices and intersects the relative interior of some face. Once we have that we can construct the cupola as described in the previous section. The set will consist of all vertices lying in a specified plane.

Lemma 2.9. *Let H be a plane which intersects the relative interior of some face F of a polytope P , and let $S = \{v_1, \dots, v_s\}$ be the set of vertices of P lying in H , not including the vertices of F . Let $S' = \{w_1, \dots, w_{s'}\}$ be a set of points in $\text{relint}(F) \cap H$. It is possible to construct a triangular cone V which intersects F in a triangle that lies in the relative interior of F and V contains S and S' in its interior and no other vertex of P .*

The reader may not see at this point the purpose of the set S' , but we will justify it at the end of this section.

Proof. $P \cap H$ is a polygon. Without loss of generality, $F \cap H$ is horizontal and situated on the top of the polygon $P \cap H$ (see Fig. 10). Let l be the line connecting the leftmost point of S' and leftmost vertex of S (the one encountered first when walking around $P \cap H$ counterclockwise, starting at $F \cap H$). Analogously, let r be the line connecting M and the rightmost vertex of S .

The area between l and r (in H) is already a cone containing S and no other vertices of P . We will perturb it in a way that the other conditions are satisfied as well.

First, shift l and r parallelly outwards, guaranteeing that they still intersect $F \cap H$ in its relative interior (easy open conditions); we obtain l' and r' . Also, let f' be a line in H parallel to F just outside P , i.e., such that l' and r' intersect f' in the same order as $F \cap H$ (again using Lemma 1.4).




Fig. 10. Construction of l and r , then l' and r' (viewed in H).




Fig. 11. Rotated hyperplanes, viewed by their intersections with F .

Now, we will rotate H about l' and r' and f' , getting three planes bounding the desired triangular cone. Let H be oriented in some way, and $a_H x \geq b_H$ be its defining inequality. Let v be some point which lies on the positive side of H . Let $G_{l'}$ be the plane through l' and v . By construction, all vertices in S lie on the same side of $G_{l'}$, so we can orient it such that S is on its positive side. Let $a_{l'} x \geq b_{l'}$ be its defining inequality. Perform the same construction for r' and f' obtaining $G_{r'}$ and $G_{f'}$, also orienting them in a way that v is on their respective positive sides. Let $G_{l'}^\epsilon$ be the plane defined by $(a_H + \epsilon a_{l'})x \geq b_H + \epsilon b_{l'}$. This plane contains l' and for small ϵ it is very close to H . Hence, it is the *rotation of H about l' in the direction of plane $G_{l'}$* . Also let $G_{r'}^\epsilon$ be defined by $(a_H + \epsilon a_{r'})x \geq b_H + \epsilon b_{r'}$, and $G_{f'}^\epsilon$ be defined by $(-a_H + \epsilon a_{f'})x \geq -b_H + \epsilon b_{f'}$.

Obviously, all points in S and in S' are on the positive sides of the planes $G_{l'}^\epsilon$, $G_{r'}^\epsilon$, and $G_{f'}^\epsilon$. For small $\epsilon > 0$, these planes do not “sweep” over vertices of P which are not in S , and it is easy to see that in this case, there are no vertices of P that satisfy all three new inequalities. Also for small ϵ , the points in F satisfying all three inequalities define a triangle in the relative interior of F with endpoints $G_{l'}^\epsilon \cap G_{r'}^\epsilon \cap F$, $G_{r'}^\epsilon \cap G_{f'}^\epsilon \cap F$, and $G_{f'}^\epsilon \cap G_{l'}^\epsilon \cap F$. Hence, the set of all points satisfying the three inequalities is a triangular cone V with the desired properties. The conditions on ϵ are open polynomial conditions according to Lemma 1.4. \square

This lemma can be used to build *one* cupola over the facet F . However, there might be problems if we keep on constructing around the polytope, like adding more cupolas over other facets of P . The visibility cone we just constructed might “catch” points we construct later. But these constructions all happen *beyond* facets of P , so we can use the following lemma to construct all cupolas one after the other without their visibility cones catching extra vertices.

Lemma 2.10. *Let H_1, \dots, H_n hyperplanes, intersecting facets F_1, \dots, F_n of a polytope P with the restriction that $F_i \cap H_j = \emptyset$ for all $i \neq j$. Then $P_{\text{beyond } F_i} \cap H_j = \emptyset$ for all $i \neq j$.*

Proof. Assume there is a point u in $P_{\text{beyond } F_i} \cap H_j$ ($i \neq j$). Then this point also lies in $(P \setminus F_i) \cap H_j$, but on the non-positive side of F_i . Let v be a point in $F_j \cap H_j$, Then v is

also in $(P \setminus F_i) \cap H_j$ (since $F_j \subseteq P \subseteq P \setminus F_i$), but on the positive side of F_i . Hence, there must be a point w on the line segment $[u, v]$ which is on the hyperplane containing F_i . The whole segment lies in $P \setminus F_i$, hence every point on it has to satisfy all of P 's defining inequalities except that of F_i . So w lies in the facet F_i . But it also lies in H_j (the whole line segment does), which contradicts the assumption $F_i \cap H_j = \emptyset$. \square

In Section 3.1 we will need an additional condition: given a set of lines in the plane H (of Lemma 2.9) that pierce the face F , we want to be sure that these lines also pierce the skylight of the constructed cupola. (This condition will play an important role when we want to force so-called blocking conditions, see Section 3: at some point two tetrahedra spanned by two skylights and two respective visible vertices v and v' will have to intersect in their interiors. This is already guaranteed if the corresponding lines g and g' intersect inside the polytope.)

The next theorem specifies the way in which we will use all the preceding lemmas in our construction in Section 3.

Theorem 2.11 (Full-strength cupola construction). *Let H_i be planes that intersect facets F_i of a polytope P in their relative interiors such that $F_i \cap H_j = \emptyset$ for all $i \neq j$. Let $S_i = \{v_1^i, \dots, v_{s_i}^i\} := (\text{vert}(P) \cap H_i) \setminus F_i$, and $L_i = \{g_1^i, \dots, g_{s_i}^i\}$ sets of lines. Assume further that each of the lines g_j^i pierces the relative interior of F_i and is incident to some v_k^i .*

Then we can sequentially construct all cupolas C_i beyond the faces F_i such that in the resulting polytope their visibility cones contain S_i and no other vertices. In addition, the skylight of the cupola over each F_i is pierced by the lines in L_i .

Proof. The theorem follows from the ideas in Lemmas 2.5 and 2.9. In the construction of the visibility cone over facet F_i , we invoke Lemma 2.9 with the polytope $P \cup \bigcup_{j \neq i} P_{\text{beyond } F_j}$. The set S'_i is, of course, $\{l \cap F_i \mid l \in L_i\}$. The cupola construction was such that the cupolas over F_j were always *beyond* the facet F_j , so the constructed visibility cones contain no vertices of the other cupolas. In order to have the lines in L_i pierce the skylight of cupola i we have to alter the construction of the cupola in Lemma 2.5: when we put the planes parallel to F , we do it in such a way that the triangles (A_1, A_2, A_3) and then (B_1, B_2, B_3) are pierced by these lines. These are both open conditions on the distance of the planes to F . \square

3. The transformation from SAT

It is our intention to model the well-known satisfiability problem (SAT) using the visibility cones of cupola polytopes. Just as Ruppert and Seidel did in [23], from now on we will restrict our attention to *special SAT instances: each variable appears exactly three times, twice unnegated and once negated*. This is not really necessary, but simplifies explanations. We will argue in Section 3.1 why this restriction is valid. The formula

$$f = (X_1 \vee \neg X_2 \vee X_3 \vee \neg X_4) \wedge (\neg X_1 \vee X_2 \vee \neg X_3 \vee X_4) \wedge (X_1 \vee X_2 \vee X_3 \vee X_4)$$

is such a special SAT formula. The figures in this section will correspond to this particular instance.

In Section 3.1 we will introduce the formal definition of the family of *logical polytopes* $\mathcal{P}_{f,m}$ associated to a given logical formula f and a natural number m . This number m will denote the length of the vertex-edge chains of the cupolas we use. In order to define logical polytopes, we will specify the face lattice of a *frame polytope*, reveal where the m -cupolas are glued, and then give more conditions on the positions of the vertices relative to each other. To achieve our goal of proving Theorem 1.1 we need two lemmas. The first lemma ensures that logical polytopes exist and that we can construct them in polynomial time (see the proof in Section 3.3). The second lemma assures that, among all logical polytopes of a fixed logical formula, the size m of the vertex chains in cupolas can be adjusted to (1) be polynomial in the number of logical variables and clauses, and (2) to be large enough to guarantee the equivalence between logical satisfiability of the logical formula and small triangulations of the logical polytope. We will prove this second lemma in Section 3.2.

Lemma 3.1. *There is a polynomial algorithm that, given any positive integer m and a logical formula f containing C clauses and V variables, produces a logical polytope $P \in \mathcal{P}_{f,m}$ with m vertices on each vertex-edge chain. The number of vertices of P is bounded by a polynomial in m , C , and V . Also, the coordinates of the vertices have binary encoding length polynomial in m , C , and V .*

Lemma 3.2. *Let f be a logical formula containing C clauses and V variables. There exists a polynomial $m(C, V)$ with integer coefficients such that for $m = m(C, V)$ and for any logical polytope $P \in \mathcal{P}_{f,m}$ the following is true: P admits a triangulation with $\leq K = n + m - 4$ tetrahedra if and only if there is a satisfying truth assignment to the variables of the formula f .*

Finally, using these two properties, we are ready for the proof of the main result.

Proof of Theorem 1.1. The problem is clearly in NP: checking whether a collection of tetrahedra is indeed a triangulation of the polytope P needs only a polynomial number of calculations. Every pair of tetrahedra is checked for proper intersection (in a common face or not at all), and the sum of the volumes equals the volume of P (computable for instance by the Delaunay triangulation of the polytope). Also the size of triangulations of a given polytope is bounded by a polynomial in n of degree two (this follows from the well-known upper bound theorem, for details see [22]).

By Lemma 3.1, from a given logical formula f on V logical variables and C clauses, we can construct a logical polytope $P \in \mathcal{P}_{f,m(C,V)}$ of encoding length polynomial in V and C . Hence, by Lemma 3.2 there is a polynomial transformation that establishes the polynomial equivalence of a solution for the SAT problem and the existence of small triangulations of P (small means of size less than K as given in the statement). This completes the proof. \square

3.1. Logical polytopes

First of all, we want to argue why the restriction to the *special SAT formulas* (every variable appears exactly twice unnegated and once negated) is NP-complete: the SAT problem remains NP-complete even for instances where each variable or its negation appear at most three times (see references in [14, p. 259]). In addition, a change of variables can be used to change a non-negated variable into a negated variable if necessary. Also note that if a variable appears only negated or only positive the variable and the clauses that contain it can be discarded. Finally, if a variable appears exactly once positive and exactly once negated then it can be eliminated by combining the two clauses that contain the two variables into one.

Now, we want to define the family of logical polytopes $\mathcal{P}_{f,m}$ for a given logical formula f and a given positive integer number m . We start by describing its face lattice. To prevent a possible confusion we remark that our vertices will be labeled by the letters c_j when they are related to logical clauses, x_j when they are related to logical variables, and z_j when the vertex is auxiliary. Points always have subscripts and/or superscripts thus should not be confused with their coordinate-entries (x, y, z) .

In a logical polytope there will be an m -cupola for each clause and one for each variable and its negation. The cupolas will be glued to a *frame polytope* which resembles a wedge. Look carefully at Fig. 12 for an example of the overall structure.

Figure 13 gives a view of the lower hull of the frame polytope. The sharp part of the wedge consists of $2C + 1$ vertices (where C is the number of clauses) c_0, \dots, c_{2C} . We call this part of the frame polytope the *spine*. We attach the *clause cupola* associated with clause i to the triangle $(c_{2i}, c_{2i+1}, c_{2i+2})$ (shaded in the picture).

On top of this wedge structure we will put a series of *roofs*. They are triangular prisms, spanned by the two triangles (z_T^i, z_F^i, z_A^i) and (z_L^i, z_R^i, z_B^i) , one for every variable X_i of the logical formula. The *variable cupolas* will be attached to the triangular facet (z_L^i, z_R^i, z_B^i) , the *back gables* (the triangular faces are shaded in Fig. 14).

The variable cupola of variable X_i is such that its visibility cone contains exactly the front vertices z_T^i and z_F^i . We will use these cupolas to read from a small triangulation of the polytope the logical value of variables according with the following rule: if the truth-setting vertex z_T^i associated to the i th logical variable is used to triangulate the skylight of the




Fig. 12. Sketch of a logical polytope.




Fig. 13. The spine of the wedge: here the clause cupolas are attached.




Fig. 14. The roofs, back gables shaded.




Fig. 15. A roof, back gable shaded, z -coordinate superelevated.

cupola for variable i , then we set $X_i = true$. If the truth-setting vertex used to triangulate the skylight of the cupola for variable i is instead z_F^i then $X_i = false$.

Beyond the quadrilateral face containing z_T^i we will place the *literal vertices* x_1^i and x_2^i which corresponds to the positive occurrences of X_i in the logical formula. Beyond the other quadrilateral face we will place the other literal vertex x_3^i which correspond to the negated occurrence of X_i . These vertices are in the visibility cones of the three cupolas of the clause where variable X_i or its negation appears.

We list the five conditions on logical polytopes which are necessary for the transformation to work in both ways, i.e., a small triangulation yields a satisfying truth assignment for our logical formula and vice versa.

Definition 3.3. For a given logical formula, the family $\mathcal{P}_{f,m}$ of *logical polytopes* is the set of all three-dimensional polytopes P that satisfy the following conditions:

- (1) (*Convexity and face lattice*) The logical polytope must be convex, and the face lattice is as we just described it. In particular, several m -cupolas are part of the polytope, one for each clause and variable in f .
- (2) (*Visibility*) The literal vertices x_1^i , x_2^i , and $\overline{x_3^i}$ are vertices in the visibility cone associated to their respective clause m -cupolas, but of no other clause visibility cone. The vertices z_T^i , z_F^i are the only vertices in the visibility cones of the i th variable m -cupola.
- (3) (*Blocking*) This constraint ensures that the assignment of true or false values for variables is done consistently, i.e., the positive (negative) literals can be used to make their clauses true if and only if the variable is set true (false).
 Concretely, the tetrahedron spanned by z_F^i and the skylight of the m -cupola of variable X_i intersects the interior of the tetrahedron spanned by x_1^i (by x_2^i) and the skylight of the clause m -cupola corresponding to x_1^i (to x_2^i). Also the tetrahedron spanned by z_T^i and the skylight of the m -cupola of variable X_i intersects the interior of the tetrahedron spanned by x_3^i and the skylight of the clause m -cupola corresponding to it. See Fig. 16 for an example.
- (4) (*Non-blocking*) Using the vertex z_T^i to triangulate the interior of the i th variable m -cupola should not prevent the non-negated literal vertices from seeing their associated m -cupolas. Concretely, if j is the clause corresponding to the literal vertex x_1^i , then tetrahedra $(z_T^i, z_L^i, z_R^i, z_B^i)$ and $(x_1^i, c_{2j-2}, c_{2j-1}, c_{2j})$ do not intersect at all. The canonical analogue shall hold for x_2^i and x_3^i (for x_3^i replace z_T by z_F).
- (5) (*Sweeping*) Because we intend to follow the same triangulation procedure which was proposed by Ruppert and Seidel [23], and which we will explain in Section 3.2, we will need that




Fig. 16. Blocking for consistent logical values.

- (a) the variable x_1^i is to the “left” (negative x direction) of the planes $c_{2k-1}c_{2k}z_F^i$, $c_{2k}c_{2k+1}z_F^i$, and $c_{2k-1}c_{2k+1}z_F^i$ for $0 \leq k \leq C - 1$.
- (b) x_2^i is to the “left” of the planes $c_{2k-1}c_{2k}x_1^i$, $c_{2k}c_{2k+1}x_1^i$, and $c_{2k-1}c_{2k+1}x_1^i$ for $0 \leq k \leq C - 1$.
- (c) x_3^i is to the “left” of the planes $c_{2k-1}c_{2k}z_F^i$, $c_{2k}c_{2k+1}z_F^i$, and $c_{2k-1}c_{2k+1}z_F^i$ for $0 \leq k \leq C - 1$.
- (d) z_T^i is to the “left” of the planes $c_{2k-1}c_{2k}x_2^i$, $c_{2k}c_{2k+1}x_2^i$, $c_{2k-1}c_{2k+1}x_2^i$, $c_{2k-1}c_{2k}x_3^i$, $c_{2k}c_{2k+1}x_3^i$, and $c_{2k-1}c_{2k+1}x_3^i$ for $0 \leq k \leq C - 1$.

3.2. Using the logical polytope

Proof of Lemma 3.2. If a triangulation T of the polytope has $\leq n + m - 4$ tetrahedra, then by Proposition 2.4 the skylight of each cupola is triangulated by a vertex in the visibility cone of the cupola. In particular, one of z_F^i and z_T^i is chosen to triangulate the cupola corresponding to variable X_i for each i . We claim that assigning to X_i the truth value according to this choice ($z_F \mapsto \text{false}$, $z_T \mapsto \text{true}$) satisfies all clauses of the formula.

Each clause cupola skylight is triangulated by one of the literal vertices, say clause j by the positive literal vertex x_1^i (or x_2^i). By the blocking conditions, it cannot be the case that the variable skylight of X_i is triangulated by z_F^i since these tetrahedra would intersect badly. So we had set X_i to *true*. Having x_1^i (or x_2^i) in clause j ’s visibility cone meant that variable X_i appears unnegated in this clause. If the skylight of clause cupola j is triangulated by x_3^i , by the same argument X_i was set to false, and clause j satisfied by the literal $\neg X_i$. Hence, all clauses are satisfied.

Now we need to prove the converse. If there is a *true–false* assignment that satisfies all logical clauses we must find a triangulation that has no more than K tetrahedra. For that we construct a “small” triangulation. There are three different kinds of tetrahedra: the ones triangulating the cupolas, the ones triangulating the roofs, and the ones triangulating of the rest of the wedge. We know how to triangulate a cupola if we know a vertex in its visibility cone (see the proof of Proposition 2.8). For the rest we will now follow a *sweeping procedure* which was first described by Ruppert and Seidel [23].

The sweeping triangulation proceeds by triangulating “slices” that correspond to the different variables X_1 to X_V , i.e., from right to left. The variable roofs are arranged sequentially for exactly this purpose. A slice is roughly speaking the part of the tetrahedra between a roof and vertices of the spine. After the i th step of the process the partial triangulation will have triangulated the first i slices. The part of the boundary of the partial triangulation that is inside the logical polytope will form a triangulated disk. We will call it the *interface* following the convention of Ruppert and Seidel. It contains the following triangles:

$$(z_T^i, c_{2C}, z_L^i) \quad \text{and} \quad \begin{cases} (z_T^i, c_{2j-2}, c_{2j}) \\ \text{if clause } j \text{ is satisfied by one of the first } i \text{ variables, or} \\ (z_T^i, c_{2j-2}, c_{2j-1}) \quad \text{and} \quad (z_T^i, c_{2j-1}, c_{2j}) \\ \text{otherwise,} \end{cases}$$

for all $j = 1, \dots, C$. Before the first step, the partial triangulation is empty. After the last




Fig. 17. The interface after step 2.




Fig. 18. Removing the tetrahedra spanned by z_T and the shaded triangles.

step the partial triangulation will cover the whole logical polytope. In general, the vertices of the i th roof will see all triangles of the interface and will be used as apexes to form new tetrahedra to add to the current partial triangulation. This way the interface will slowly move from right to left.

Now we describe in detail the triangulation step for the i th variable X_i . Since we are only concerned with roof vertices in roof i , we will drop all superscripts. The triangulation step depends on whether X_i is set *true* or *false* in the satisfying assignment. Let us consider first the case $X_i = \text{true}$.

The point z_T is used to triangulate the interior of the variable cupola associated to X_i according to Proposition 2.8. From z_T we also form six tetrahedra with the following triangles: (z_L, \bar{x}_3, z_B) , (\bar{x}_3, z_B, z_A) , (z_B, z_A, x_2) , (z_B, x_2, z_R) , (z_A, x_1, x_2) , and (x_1, z_A, z_F) .

Now we come to the part of the triangulation which gave the sweeping procedure its name. We form the tetrahedra between x_1 and the current interface triangles. This is possible by part (a) of condition (5). We also use the tetrahedron (x_1, z_T, c_0, z_F) . This is illustrated in the transition from (a) to (b) in Fig. 19. The interface advances to x_1 , i.e., if (z_F, c_j, c_k) was an interface triangle before, now (x_1, c_j, c_k) is an interface triangle. Also (z_F, c_{2C}, z_R) is replaced by the triangle (x_1, c_{2C}, z_R) .

Since X_i is set to *true* we can use x_1 to triangulate its clause cupola according to Proposition 2.8. We only do this if the clause cupola has not been previously triangulated using an other literal vertex. Condition (2) ensures that x_1 is in the visibility cone of the clause cupola coming from the clause that contains the unnegated literal X_i . Furthermore, condition (4) ensures that we can actually perform this triangulation of the clause cupola




Fig. 19. The sweep.




Fig. 20. The sweep for $X_i = false$.

without badly intersecting the tetrahedra of the variable cupola. In Fig. 19(c) we see that if x_1 is used to triangulate clause j 's cupola, then the interface triangle (x_1, c_{2j-2}, c_{2j}) is replaced by the two triangles $(x_1, c_{2j-2}, c_{2j-1})$ and (x_1, c_{2j-1}, c_{2j}) .

We repeat this procedure with x_2 , i.e., form tetrahedra with x_2 and the current interface triangles, and then use x_2 to triangulate its clause cupola if necessary (Fig. 19(d)). This is possible by part (b) of condition (5). We continue by forming tetrahedra using z_T as apex (Fig. 19(e), possible by condition (5), part (d)). At last, we will include the triangle (c_{2C}, z_L, z_B) . All these triangles are visible by part (d) of condition (6). After all these tetrahedra are added the interface is ready for the next variable.

Let us now consider the triangulation step in the case X_i is set to be *false*: we use the vertex z_F to triangulate the variable cupola as well as seven faces of the roof (see Fig. 20): (z_T, \bar{x}_3, z_A) , (\bar{x}_3, z_A, z_B) , (\bar{x}_3, z_L, z_B) , (z_B, z_A, x_2) , (z_B, x_2, z_R) , (z_A, x_2, x_1) , (x_2, x_1, z_R) . The reader can see that on the roof we are leaving only the vertex \bar{x}_3 . Next the tetrahedron (z_F, z_L, z_R, c_{2C}) is cut out. Hereby, the interface triangle (z_F, z_R, c_{2C}) is replaced by (z_F, z_L, c_{2C}) (Fig. 20(c)). Then \bar{x}_3 will be used as apex with the triangles in the interface. If the negated literal \bar{X}_i is used to satisfy its clause j , the j th clause cupola is triangulated by \bar{x}_3 . The interface advances as in the *true*-case. Then z_T can be used to form tetrahedra with the triangles in the interface. In the end the interface is again ready for the next variable.

How many tetrahedra can such a triangulation have? Triangulating all cupolas with a vertex in their visibility cones yields at most $(3m + 16)(C + V)$ tetrahedra (Proposition 2.8). In one step of the sweeping triangulation the tops of the roofs are each triangulated using six or seven tetrahedra (if the variable is unnegated or negated, respectively). Furthermore, the interface is triangulated by some vertices three times (in the positive case by x_1^i , by x_2^i , and by z_T^i) or two times (in the negative case by $\overline{x_3^i}$ and by z_T^i). The interface contains in each step between C and $2C$ triangles. Eventually, in either case there is one more tetrahedron (see above). An upper bound for the size of this triangulation is therefore

$$\begin{aligned} \#T &\leq (3m + 16)(C + V) + 7V + 3CV + 1 \\ &= m(3C + 3V) + \underbrace{16C + 23V + 3CV + 1}_{p_T(V, C)}. \end{aligned}$$

What is the number of the vertices of the logical polytope in terms of the number of clauses and variables? We have V logical variables and C clauses in the SAT instance. We have m interior points each of the vertex-edge chains we added (later we will determine the value of m as a polynomial function of V and C). We observe that we have $3m + 6$ vertices in each cupola, hence we have $(3m + 6)(V + C)$ for all cupolas. We have in each roof nine vertices, two of them are shared with the subsequent roof except for the last roof. Hence, the total number of vertices in roofs is $7V + 2$. We have left only the $2C + 1$ vertices along the spine. In conclusion, the number of vertices of P is

$$n = (3m + 6)(V + C) + 7V + 2 + 2C + 1 = m(3C + 3V) + \underbrace{8C + 13V + 3}_{p_n(C, V)}.$$

We had said before that a “bad” triangulation (where at least one cupola skylight is triangulated by a vertex not lying in its visibility cone) has at least $n + m - 3 = m(3C + 3V + 1) + p_n(C, V)$ tetrahedra. On the other hand, a “good” triangulation has at most $m(3C + 3V) + p_T(C, V)$ tetrahedra. We can now set $m = m(C, V) = p_T(C, V) - p_n(C, V) + 1$. Then, if a good triangulation exists, its size is smaller than or equal to $K = n + m - 4$, and if not, all triangulations are larger than K . \square

3.3. Constructing a logical polytope

Proof of Lemma 3.1. The construction will be carried out in five stages. By the time we end the construction all five requirements of the definition of logical polytopes must be satisfied, but three of the conditions will not be met until the last stage.

- (1) Give coordinates of the basic wedge, with rectangular faces on top for each variable.
- (2) Attach the roofs for each variable, giving preliminary coordinates for the literal vertices and preliminary coordinates for the points on the lower edge (the spine of the wedge).
- (3) Perturb the literal vertices to their final positions.
- (4) Perturb the vertices on the spine of the wedge.
- (5) Attaching the variable cupolas following the procedures of Section 2.

In every step we will build a construction element (a point, a line, or a plane) whose coordinates are polynomials in the construction elements up to that particular moment. Hence, the encoding length of each new construction element is bounded by a linear function of the encoding length of the construction so far. The number of construction steps is polynomially bounded in C and V . Hence, the encoding length of the whole construction is also polynomially bounded in C and V . Note however, that the coordinates themselves will in general be exponentially large.

Instead of writing explicit (and highly cumbersome) coordinates for the construction elements, we rely on Lemma 1.4 to ensure that such coordinates can be found if one has really the desire to see a particular logical polytope. A key property of stages (2)–(4) in the construction is that the geometric conditions we want to determine a finite collection of strict polynomial inequalities *in a single variable*. Then, by Lemma 1.4, we know there is an appropriate polynomial size solution. In subsequent stages of the construction similar new systems, for other independent parameters, will be solved, preserving what we had so far, but building up new properties.

Stage (1). The basic wedge. Consider the triangular prism which is the convex hull of the six points $c_0 = (0, 0, 0)$, $c_{2C} = (0, 1, 0)$, $z_T^V = (0, 0, 1)$, $z_F^1 = (1, 0, 1)$, $z_L^V = (0, 1, 1)$, and $z_R^1 = (1, 1, 1)$. See Fig. 21(a). In order to obtain a convex structure on the top of the wedge, we consider the function $f(x) = x(1 - x) + 1$. The vertices of each roof boundary (that is z_T^i and z_F^i as well as z_R^i and z_L^i) will lie on the surface $z = f(x)$. More specifically, $z_F^i = z_T^{i+1} = (i/V, 0, f(i/V))$ and $z_R^i = z_L^{i+1} = (i/V, 1, f(i/V))$ for $i = 0, \dots, n$. By the concavity of f , the points are indeed in convex position and their convex hull, the *wedge* has the desired face lattice (see Fig. 21(b)).

So far none of the conditions we want are satisfied (not even partially).

Stage (2). The roofs. We will first attach the points z_A^i and z_B^i to the quadrilateral face $(z_L^i, z_R^i, z_T^i, z_F^i)$. Then we give preliminary coordinates to the literal vertices and to the vertices on the spine.

Let $z_A^i = 1/2 \cdot (z_T^i + z_F^i) + (0, 1/3, t_{\text{roof}})$ and $z_B^i = 1/2 \cdot (z_T^i + z_F^i) + (0, 2/3, t_{\text{roof}})$ where t_{roof} is a non-negative parameter that is called the *roof height*. That is the points have the same x coordinate as the midpoint between z_T^i and z_F^i , y coordinate $1/3$ and $2/3$,




Fig. 21. Construction of the wedge.

respectively, and height t_{roof} over the face $(z_T^i, z_F^i, z_L^i, z_R^i)$. We want to choose t_{roof} in a way that z_A^i and z_B^i are beyond the facet $(z_T^i, z_F^i, z_L^i, z_R^i)$ (see Fig. 21(c)). We can easily achieve this by the technique presented in Lemma 1.4: The only possibly concave edges are the (z_T^i, z_L^i) . One restriction is therefore that all determinants $\det(z_T^i, z_L^i, z_A^{i-1}, z_A^i)$ have to be positive. These are finitely many open quadratic conditions on t_{roof} . For $t_{\text{roof}} = 0$ the points z_A^i and z_B^i are inside the facets $(z_T^i, z_F^i, z_L^i, z_R^i)$, hence the edges in question are trivially convex. We will get more polynomial constraints on t_{roof} below and then solve all simultaneously to find the suitable roof height.

The spine of the wedge is still a line. We now put preliminary points c_0, \dots, c_{2C} on this line. Let

$$u(j) = \frac{1}{2} \frac{j}{2C}$$

and $c_j = (0, u(j), 0)$ for $j = 0, \dots, 2C - 1$, and $c_{2C} = (0, 1, 0)$ (see Fig. 22). As an auxiliary point, let b_l be the barycenter of the points c_{2l-2}, c_{2l-1} , and c_{2l} ($l = 1, \dots, C$). At this moment, this point $b_l = c_{2l-1}$. Later, as we perturb the spine vertices b_l will move accordingly, always $b_l = 1/3(c_{2l-2} + c_{2l-1} + c_{2l})$.

Now we want to give initial positions to the literal vertices. Say variable X_i occurs unnegated in clauses l_1 and l_2 and negated in l_3 . Note that l_j depend on the variable we are considering. For instance, in our example logical formula on p. 151, for variable X_1 , $l_1 = 1, l_2 = 3$, and $l_3 = 2$. But for variable X_2 , $l_1 = 2, l_2 = 3$, and $l_3 = 1$.

The preliminary literal vertex x_1^i is the intersection of the $y = u(2l_1 - 1)$ plane with the line connecting z_F^i and z_B^i . We do the same for the other positive occurrence of X_i and obtain the preliminary x_2^i . For the negative occurrence of X_i , we take the line connecting z_T^i and z_B^i , intersect it with the $y = u(2l_3 - 1)$ plane, and obtain the preliminary x_3^i . We




Fig. 22. Preliminary coordinates for the spine vertices.




Fig. 23. Construction of the literal vertices in the X_i slice of the wedge.




Fig. 24. Construction of H^i and g^i .

join the preliminary x_1^i and b_{l_1} by a line d_1^i (this line lies in the $y = u(2l_1 - 1)$ plane). Do the analogue process for x_2^i and x_3^i , obtaining d_2^i and d_3^i . Later we will move the vertices x_1^i, x_2^i, x_3^i along their respective lines d_1^i, d_2^i, d_3^i a little out of polytope in order to turn them into extreme points. The lines d_j^i will also be used for blocking conditions.

Let H^i be the plane that contains z_T^i and z_F^i and the midpoint of the edge (z_L^i, z_B^i) (Fig. 24). The only vertices above H^i are $x_1^i, x_2^i, x_3^i, z_A^i$, and z_B^i , and the only vertices on H^i are z_T^i and z_F^i . This follows from the convexity of the current polytope.

Let g_1^i (g_2^i) be the line in the plane H^i which is incident to z_F^i and intersects the line d_1^i (d_2^i). Note that this intersection point lies in the segment (x_1^i, c_{2l_1-1}) (the line segment (x_2^i, c_{2l_2-1})), thus in the interior of the constructed polytope. Analogously, let g_3^i be the line in the plane H^i which is incident to z_T^i and intersects the line segment (x_3^i, c_{2l_3-1}) . It can be verified that if the roof height is small (z_L^i, z_R^i, z_B^i) is pierced by the g_j^i in its relative interior. This is another strict polynomial inequality in t_{roof} . It will be the planes H^i and lines g_j^i ($i = 1, \dots, V$) from which we make the visibility cones for the cupolas of variables X_i according to Theorem 2.11.

It is important to note right now that the non-blocking conditions are satisfied for this special position of the vertices. We do not want the tetrahedron $(z_T^i, z_L^i, z_R^i, z_B^i)$ and the triangle $(x_1^i, c_{2l_1-2}, c_{2l_1})$ to intersect. From this we get strict polynomial inequalities on t_{roof} . They are satisfied for $t_{\text{roof}} = 0$ since the y coordinates of the spine vertices c_l are smaller than $1/2$. A suitable value of t_{roof} can be found solving the univariate inequality system we accumulated in our discussion (Lemma 1.4). It is easy to check that the sweeping conditions are also satisfied for the preliminary position of the points x_1^i, x_2^i, x_3^i . So far we have met two of the five required conditions to have a logical polytope.

Stage (3). Literal vertices. Now we put the final x_j^i ($j = 1, 2, 3$) a little outward on line d_j^i (Fig. 23). A little for x_1^i and x_2^i means that the positive literal vertices lie in a plane parallel to the face $(z_R^i, z_B^i, z_A^i, z_F^i)$ very close to it. We treat x_3^i similarly. If the three literal vertices are moved a sufficiently small distance t_{literal} the face lattice of what we get after taking the convex hull is as Fig. 15(a) in all roofs. See also the Schlegel diagram in Fig. 15(b).

By construction H_i contains z_F^i and z_T^i , and the $y = u(2j - 1)$ planes contain all literal vertices corresponding to clause j . This will become important for the visibility

conditions (see stage (5)). Also, for small t_{literal} the non-blocking and sweeping conditions are satisfied.

Although we do not have the blocking condition yet auxiliary lines can be set up. As above, let l_1, l_2, l_3 be the clauses to which the literal vertex $x_1^i, x_2^i, \overline{x_3^i}$ belong. We made sure that the line segments (c_{2l_1-1}, x_1^i) and (z_F^i, z_B^i) intersect in their respective relative interiors. Hence, by the construction of line g_1^i , it is also pierced by (x_1^i, c_{2l_1-1}) between z_F and the face (z_L^i, z_R^i, z_B^i) . (Analogously, (c_{2l_2-1}, x_2^i) and $(z_F^i, g_2^i \cap (z_L^i, z_R^i, z_B^i))$ as well as $(c_{2l_3-1}, \overline{x_3^i})$ and $(z_T^i, g_3^i \cap (z_L^i, z_R^i, z_B^i))$ intersect in their relative interiors.) Later on this intersection will evolve into the real blocking conditions using Theorem 2.11.

Stage (4). Perturbing the vertices on the spine of the wedge. We now perturb the points c_j on the spine of the wedge. Every even-indexed c_{2l} is changed to lie on a parabola, and for the moment the odd-indexed vertices c_{2l-1} are changed to lie on the line connecting c_{2l-2} and c_{2l} . The y coordinates of all points stay the same

$$c_{2l} = \left(\frac{1}{2}(y-1)^2 \cdot t_{\text{even}}, y, (y-1)^2 \cdot t_{\text{even}} \right).$$

Note that by the $1/2$ in the x coordinate, the points are moved *into* the polytope. The changes (parameter t_{even}) must be small enough that the convex hull now has the desired appearance (Fig. 25) and the non-blocking conditions and the sweeping conditions are still satisfied. Once more we appeal to Lemma 1.4. The polynomial inequalities are now on the variable t_{even} and the sweeping and non-blocking were satisfied at $t_{\text{even}} = 0$. The reader should note that while the constructed vertices in the roofs do not change coordinates, dependent construction elements like the lines d_j^i (connecting x_j^i and c_{2l_j-1}) and g_j^i (lying in H^i and intersecting d_j^i) change when the spine vertices move. However, the parameter t_{even} has to be small enough that the preliminary blocking conditions are still met: g_j^i still pierce the facet (z_L^i, z_R^i, z_B^i) in its relative interior, and g_j^i and d_j^i intersect in the interior of the polytope.

Now we move the odd points c_{2l-1} beyond the face $G_l = (c_{2l-2}, c_{2l}, z_T^0)$: to this end, we choose a point p_l beyond G_l and move to $c_{2l-1} + t_{\text{odd}}(p_l - c_{2l-1})$. Such a point p_l is easily found by taking a normal to G_l through its barycenter and moving outwards while staying beyond the face (note that this involves again Lemma 1.4, see the definition of *beyond*). The parameter t_{odd} is chosen small enough: convexity and the correctness of the face lattice are easily achieved. Also the sweeping conditions are valid for slight moves. Keeping t_{odd} small also guarantees the non-blocking conditions: the tetrahedron $(x_1^i, c_{2l_1-2}, c_{2l_1-1}, c_{2l_1})$




Fig. 25. Perturbation of the vertices on the spine.

is only slightly bigger than just the triangle $(x_1^i, c_{2l_1-2}, c_{2l_1})$ which did not intersect the tetrahedron $(z_T^i, z_L^i, z_R^i, z_B^i)$ (x_2^i and $\overline{x_3^i}$).

For the blocking conditions, let X_i be the j th logical variable in clause l . Note that now the line d_j^i intersects the triangle $(c_{2l-2}, c_{2l-1}, c_{2l})$ in its relative interior. The lines g_j^i are updated as the lines d_j^i move. Since t_{odd} is small, g_j^i still pierces the facet (z_L^i, z_R^i, z_B^i) in its relative interior, and g_j^i and d_j^i intersect in the interior of the polytope. Note that d_j^i is still in the $y = u(2l - 1)$ plane because the y coordinates of the spine vertices were conserved.

Stage (5). Attaching the cupolas. It remains to construct all the cupolas. Over the facets (z_L^i, z_R^i, z_B^i) ($i = 1, \dots, V$) we construct cupolas using the planes H_i and sets of lines $\{g_1^i, g_2^i, g_3^i\}$, and over the facets $(c_{2l-2}, c_{2l-1}, c_{2l})$ ($i = 1, \dots, C$) we construct the clause cupolas using the $y = g(2l - 1)$ planes and the sets of lines $\{d_j^i \mid X_i$'s j th occurrence is in clause $l\}$. We invoke Theorem 2.11 and get the final polytope. By this construction, it is convex, has the correct face lattice, and the visibility conditions are satisfied.

The reader will recall that g_j^i and d_j^i intersect in the interior of the polytope. Say again variable X_i occurs unnegated in clauses l_1 and l_2 and negated in l_3 . By Theorem 2.11 g_j^i pierces the skylight of the cupola corresponding to variable X_i and d_j^i pierces the skylight corresponding to its clause l_j . Hence, the tetrahedron spanned by z_F^i and the variable X_i 's skylight together with the tetrahedron spanned by x_1^i (x_2^i) and clause l_1 's skylight (l_2 's skylight) intersect in their interiors. Analogously, the tetrahedron spanned by z_T^i and the variable X_i 's skylight and the tetrahedron spanned by $\overline{x_3^i}$ and clause l_3 's skylight intersect in their interiors. These are exactly the blocking conditions.

All other conditions concerned only points we constructed before, so they are still satisfied. The final polytope is therefore a logical polytope. \square

4. Final remarks and conclusions

It is worth noticing at this point that the size of the constant K we constructed in Section 3.2 satisfies in fact $n - 3 < K < 2n$. Now we discuss an interesting justification of why the lower bound is strict.

We say that a convex polytope is *stacked* if it has a triangulation whose dual graph is a tree (the dual graph of a simplicial complex is the graph that has one vertex for each maximal-dimensional simplex and two vertices are connected precisely when the corresponding simplices are adjacent via a common facet). The reader should be aware that in the literature the terminology stacked polytope is often restricted to simplicial polytopes. Here, of course, we use it allowing that the stacking of simplices may give coplanarities. For example, any 3-cube or triangular prism is a stacked polytope under our definition.

It turns out that a convex d -polytope P with n vertices has a triangulation of size $n - d$ precisely when P is a stacked polytope (see [22]). It is natural to ask which polytopes are stacked. If it were NP-hard to recognize stacked polytopes then this would provide another proof that the problem of finding minimal triangulations is also in the same class. However, we can prove:

Theorem 4.1. (1) For a convex d -polytope P , presented by its n vertices, there is a $O(n^3d^3)$ algorithm to decide whether P is stacked (i.e., P has a triangulation with $n - d$ maximal simplices). The algorithm uses only the 1-skeleton of the polytope. The size of a minimal triangulation of a stacked polytope does not depend on the particular coordinatization, but only depends on its face lattice.

(2) A convex 3-dimensional polytope P is stacked if and only if its graph does not contain as a minor the graph of an octahedron or a pentagonal prism.

Proof. We need some definitions that have been introduced earlier in the graph theory literature [1,2,12]. We say a graph G is k -decomposable if G has $k + 1$ or fewer vertices or there is a subset of vertices S of G with at most k vertices such that (i) S is a cut, i.e., $G - S$ is disconnected, and (ii) each of the connected components of $G - S$ has the property that when the vertices of S are added back together with the complete graph on those vertices, the resulting graph is again k -decomposable.

It was shown in [1, Theorem 2.7] that the class of k -decomposable graphs is the same as the class of partial k -trees: a graph is a k -tree if it can be reduced to the complete graph K_k , by a finite sequence of removals of degree k vertices with completely connected neighbors (i.e., neighbors of the vertex induce a complete graph K_k). A partial k -tree is simply an edge-subgraph of a k -tree.

Now we claim that a d -dimensional convex polytope P is stacked, if and only if its 1-skeleton is a partial d -tree. Here is the proof: the “only if” implication is clear from the definition of stacked polytope. We can prove the “if” implication by induction on n . The theorem is trivial if $n = d + 1$ because then P is a simplex and its graph is a complete graph, thus is a d -tree. Assume then $n > d + 1$ and that the result is true for polytopes with fewer than n vertices. Remember that if 1-skeleton $G(P)$ is a partial d -tree then it is d -decomposable. Thus, there is a cut S of cardinality at most d . The set S must have in fact cardinality d because $G(P)$ is d -connected by Balinski’s theorem. If one has a vertex-cutset S of cardinality d in $G(P)$, then the hyperplane $H(S)$ spanned by S intersects $G(P)$ only in the vertices of $G(P) \cap S$ and in no edges (otherwise S is not a cut). In conclusion, $H(S) \cap P$ is a $(d - 1)$ -simplex and because $G(P)$ is d -decomposable we can apply induction hypothesis to prove the polytopes $H(S)^+ \cap P$ and $H(S)^- \cap P$ are d -decomposable, and thus they are partial d -trees with fewer vertices than P , so both polytopes are stacked. Finally, note that their stacked triangulations match well at the common boundary simplicial facet $H(S) \cap P$, proving that P is stacked.

There is an algorithm that, for fixed values of k , decides whether a given graph is a partial k -tree in running time $O(n)$ [3]. Computing the 1-skeleton of a k -polytope P can be computed with some more effort. Two vertices v_1 and v_2 are adjacent in P if and only if there is $a \in R^d$ and $b \in R$ such that $a^T v_1 = a^T v_2 = b$ and $a^T p < b$ for all p in $P - \{v_1, v_2\}$. That is, there is a halfspace which supports P at exactly two points v_1, v_2 . Clearly, $O(n^2)$ of these LP calculations suffice. From the well-known Ellipsoid method each such LP can be solved in $O(nd^3)$ operations. This, together with the above claim about detecting partial k -trees, completes the proof of the first part of the theorem.

For the second part we observe that partial k -trees form a minor closed family. This means that the set of partial k -trees is closed under taking edge-deletion or edge-contraction operations. The famous results of [26] imply that they can be characterized by a finite set

of forbidden minors. El-Mallah and Colbourn [12] proved that a graph is a planar 3-tree if and only if it has no minor isomorphic to the graph of an octahedron or a pentagonal prism. This fact together with our claim complete the proof of the second part. \square

Final remarks.

- (1) The “coning” triangulation proposed in [13] provides an algorithm which is polynomial on the number of vertices and gives a 2-approximation of the minimal triangulation as it produces a triangulation of size less than or equal to $2n - 7$.
- (2) Given a 3-dimensional convex polytope P , a proper subset S of tetrahedra with vertices in $\text{vertices}(P)$, and a positive integer K . Deciding whether there is a triangulation of P that uses K simplices from S can also be proved to be an NP-hard problem using the constructions we explained.
- (3) It is interesting to note that the constructions presented in [5] prove also that *covering* a convex 3-polytope with tetrahedra can be done with fewer pieces than triangulating. A cover is a collection of simplices whose union is the whole polytope, but the elements can intersect in their interiors. *Question*: what is the computational complexity of finding minimal simplicial covers?
- (4) *Question*: what is the complexity of deciding whether a triangulation of the boundary of convex non-simplicial 3-polytope extends to a triangulation of the whole polytope without adding new interior points? If NP-hard, this could be used to establish another proof our results. The curious reader can easily prove that already for a triangular prism not all triangulations of the boundary extend to a triangulation of the whole polytope. Even more interesting. The triangulations of the boundary of a 3-cube extend or not depending on the coordinates of its vertices.

Acknowledgments

We thank M. Bern, K. Fukuda, J. O’Rourke, B. Sturmfels, Emo Welzl, and anonymous referees for their comments and suggestions.

References

- [1] S. Arnborg, A. Proskurowski, Characterization and recognition of partial 3-trees, *SIAM J. Algebraic Discrete Meth.* 7 (2) (1986) 305–314.
- [2] S. Arnborg, D.G. Corneil, A. Proskurowski, Complexity of finding embeddings in a k -tree, *SIAM J. Algebraic Discrete Meth.* 8 (2) (1987) 277–284.
- [3] S. Arnborg, A. Proskurowski, Linear time algorithms for NP-hard problems restricted to partial k -trees, *Discrete Appl. Math.* 23 (1989) 11–24.
- [4] D. Avis, H. ElGindy, Triangulating point sets in space, *Discrete Comput. Geom.* 2 (1987) 99–111.
- [5] A. Below, U. Brehm, J.A. De Loera, J. Richter-Gebert, Minimal simplicial dissections and triangulations of convex 3-polytopes, *Discrete Comput. Geom.* 24 (2000) 35–48.
- [6] M. Bern, D. Eppstein, Mesh generation and optimal triangulation, in: D.Z. Du, F.K. Hwang (Eds.), *Computing in Euclidean Geometry*, World Scientific, 1992, see also second edition appeared in 1995.

- [7] L. Billera, A. Björner, Face numbers of polytopes and complexes, in: J.E. Goodman, J. O'Rourke (Eds.), *Handbook of Discrete and Computational Geometry*, CRC Press, New York, 1997, pp. 271–290.
- [8] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, G.M. Ziegler, *Oriented Matroids*, Cambridge Univ. Press, Cambridge, 1992.
- [9] B. Chazelle, L. Palios, Triangulating a nonconvex polytope, *Discrete Comput. Geom.* 5 (1990) 505–526.
- [10] B. Chazelle, An optimal convex hull algorithm in any fixed dimension, *Discrete Comput. Geom.* 10 (1993) 377–409.
- [11] R.W. Cottle, Minimal triangulation of the 4-cube, *Discrete Math.* 40 (1982) 25–29.
- [12] E. El-Mallah, C.J. Colbourn, On two dual classes of planar graphs, *Discrete Math.* 80 (1990) 21–40.
- [13] H. Edelsbrunner, F.P. Preparata, D.B. West, Tetrahedrizing point sets in three dimensions, *J. Symbolic Comput.* 10 (1990) 335–347.
- [14] M.R. Garey, S.J. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.
- [15] M. Haiman, A simple and relatively efficient triangulation of the n -cube, *Discrete Comput. Geom.* 6 (1991) 287–289.
- [16] R.B. Hughes, M.R. Anderson, Simplexity of the cube, *Discrete Math.* 158 (1996) 99–150.
- [17] C.W. Lee, Subdivisions and triangulations of polytopes, in: J.E. Goodman, J. O'Rourke (Eds.), *Handbook of Discrete and Computational Geometry*, CRC Press, New York, 1997, pp. 271–290.
- [18] C.W. Lee, Regular triangulations of convex polytopes, in: P. Gritzmann, B. Sturmfels (Eds.), *Applied Geometry and Discrete Mathematics—The Victor Klee Festschrift*, in: DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 4, 1991, pp. 443–456.
- [19] T.H. Marshall, Volume formulae for regular hyperbolic cubes, *Conform. Geom. Dynam.* 2 (1998) 25–28.
- [20] J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford Univ. Press, New York/Oxford, 1987.
- [21] J. Richter-Gebert, Finding small triangulations of polytope boundaries is hard, *Discrete Comput. Geom.* 24 (2000) 503–518.
- [22] G.L. Rothschild, E.G. Straus, On triangulations of the convex hull of n points, *Combinatorica* 5 (1985) 167–179.
- [23] J. Ruppert, R. Seidel, On the difficulty of triangulating three-dimensional non-convex polyhedra, *Discrete Comput. Geom.* 7 (1992) 227–253.
- [24] E. Schönhardt, Über die Zerlegung von Dreieckspolyedern in Tetraeder, *Math. Ann.* 98 (1928) 309–312.
- [25] W.D. Smith, A lower bound for the simplexity of the n -cube via hyperbolic volumes, *Combinatorics of Polytopes*, *European J. Combin.* 21 (1) (2000) 131–137.
- [26] N. Robertson, P.D. Seymour, Graph minors—a survey, in: I. Anderson (Ed.), *Surveys in Combinatorics*, Cambridge Univ. Press, 1985.
- [27] D.D. Sleator, R.E. Tarjan, W.P. Thurston, Rotation distance, triangulations, and hyperbolic geometry, *J. Amer. Math. Soc.* 1 (1988) 647–681.
- [28] B. Sturmfels, Polynomial equations and convex polytopes, *Amer. Math. Monthly* 105 (1998) 907–922.
- [29] M.J. Todd, *The Computation of Fixed Points and Applications*, in: *Lecture Notes in Econom. and Math. Systems*, vol. 124, Springer-Verlag, Berlin, 1976.
- [30] G.M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, New York, 1994.