

Finding Small Triangulations of Polytope Boundaries Is Hard

J. Richter-Gebert

Institute for Theoretical Computer Science, ETH Zentrum,
CH-8092 Zürich, Switzerland
richter@inf.ethz.ch

Abstract. We prove that it is NP-hard to decide whether a polyhedral 3-ball can be triangulated with k simplices. The construction also implies that it is difficult to find the minimal triangulation of such a 3-ball. A lifting argument is used to transfer the result also to triangulations of boundaries of 4-polytopes.

The proof is constructive and translates a variant of the 3-SAT problem into an instance of a concrete polyhedral 3-ball for which it is difficult to find a minimal triangulation.

1. Introduction

Polytopal complexes are objects that arise from joining polytopes in a face-to-face manner. We here study triangulations of polytopal complexes. We will prove that in general it is computationally hard to calculate a triangulation of a complex that uses a minimal number of simplices. Formally, a polytopal complex \mathcal{C} is a finite collection of polytopes such that

- (i) the empty polytope is in \mathcal{C} ,
- (ii) if $P \in \mathcal{C}$, then all faces of P are also in \mathcal{C} ,
- (iii) the intersection $P \cap Q$ of two polytopes $P, Q \in \mathcal{C}$ is a face of both P and Q .

The *dimension* of \mathcal{C} is the largest dimension of a polytope in \mathcal{C} . A *triangulation* of \mathcal{C} is a polytopal complex \mathcal{T} that has the following properties:

- (i) the vertices (i.e., 0-faces) of \mathcal{C} are the vertices of \mathcal{T} ,
- (ii) all polytopes in \mathcal{T} are simplices,
- (iii) the intersection $P \cap Q$ of a polytope $P \in \mathcal{C}$ and a polytope $Q \in \mathcal{T}$ is a face of Q (i.e., \mathcal{T} is a refinement of \mathcal{C}).

This paper investigates the algorithmic complexity of finding triangulations of three-dimensional polyhedral complexes that use a minimal number of tetrahedra. Besides

its intrinsic interest this problem is relevant for questions in optimization (where one is interested in triangulations that minimize certain properties), in computational geometry (for instance, for mesh generation and finite elements methods), or even in algebraic geometry (where polytopes and subdivision arise from exponent vectors of polynomials). Originally, the research that led to the result presented in this paper was motivated by a different but related problem: *How difficult is it to calculate a minimal triangulation of a polytope?* For recent progress in this area consider the paper of de Loera [1].

Here we study two classes of polyhedral complexes of particular interest: polytopal 3-balls (these are polytopal complexes that arise as subdivisions of convex 3-polytopes) and boundaries of 4-polytopes. In both cases we prove that it is NP-hard to decide whether a given polyhedral complex can be triangulated using a fixed number k of tetrahedra. The number k that occurs in our proof is $k = \sum_{i=1}^n (\text{mintriang}(P_i))$, where P_1, \dots, P_n are the 3-polytopes (cells) in \mathcal{C} and $\text{mintriang}(P_i)$ is the minimal number of tetrahedra needed for a triangulation of P_i . This number k is the simplest lower bound on the number of tetrahedra needed for a triangulation of \mathcal{C} . Thus we can sharpen our main results to the following form:

It is NP-hard to decide whether a polytopal 3-ball (resp. the boundary of a 4-polytope) can be triangulated by using the minimal possible number of tetrahedra for each cell.

It is obvious that a similar statement about triangulations of 2-balls (resp. the boundary of 3-polytopes) does not hold. This is the case since here every triangulation of the maximal (i.e., two-dimensional) cells is already a triangulation of the entire polytopal complex. The reason for this is that in two-dimensional polyhedral complexes the triangulations of the individual cells do not interfere with each other.

For three-dimensional polyhedral complexes the situation is different. The triangulations of two adjacent cells that share a 2-face F must be *compatible*: they must induce identical triangulations on F . This “flow of information” is strong enough that we can embed a certain variant of the 3-SAT problem (that is known to be NP-hard) in a triangulation problem of three-dimensional complexes. As usual in cases of embedding a variant of 3-SAT into a geometric problem, this “implementation process” can be subdivided into several distinct (more or less standard) tasks (see, for instance, [3]–[6]):

- Construct a *frame*, i.e., a suitably rigid geometric (sub)structure that serves as a frame of reference for the rest of the construction.
- Find *switches*, these are small building blocks that realize a certain geometric property in two (or more) different ways.
- Find a method to produce *negated copies* of the switches.
- Find *gates*, these are small building blocks that provide some kind of logical connection between “input stages.”
- Find a way to *connect* switches and gates.

Our proof follows these five steps. We consider cell complexes that arise in polytopal 3-balls. All essential constructions are carried out on that level. Later we identify these cell complexes as subcomplexes of the boundary of a suitable 4-polytope.

2. The Construction

2.1. A Variant of 3-SAT

Let $X = (x_1, x_2, \dots, x_n)$ be boolean variables. The *literals* over X are the variables in X together with their negations $\neg x_1, \neg x_2, \dots, \neg x_n$. A three-clause is a triple of literals over X . Our construction is based on the following variant of the 3-SAT problem which is known to be NP-complete (compare [2]).

Problem 2.1 (NOT-ALL-EQUAL-3SAT). Given boolean variables x_1, \dots, x_n and a set S of m three-clauses. Is there an assignment of boolean values for the elements of X such that each clause contains at least one *false* and one *true* literal? Such an assignment is called admissible.

We describe a procedure that associates to each instance S of the above problem a 3-ball $K(S)$ and a number $k(S)$ such that $K(S)$ can be triangulated with $k(S)$ tetrahedra if and only if there is an admissible assignment for S .

2.2. Triangulations of 3-Polytopes

The idea of our construction is based on the following simple observation on 3-polytopes.

Lemma 2.1. *Let P be a 3-polytope with v vertices and let T be a triangulation of P that has e^i edges in the interior of P . Then the number of tetrahedra in T is $v + e^i - 3$.*

Proof. The proof follows from a simple (and standard) application of the Euler formula for polytopes. Let t be the number of tetrahedra in T , let e^b be the number of edges T that lie on the boundary of P , and let f^i and f^b be the number of 2-faces of T that lie in the interior (resp. on the boundary) of P . Since every tetrahedron has four faces and the interior faces of the triangulation are adjacent to two tetrahedra we have $4t = f^b + 2f^i$. On the other hand T induces a triangulation on the boundary and we get $f^b = 2v - 4$. Merging the last two equations we get $f^i = 2t - v + 2$. Euler's formula for 3-polytopes applied to the boundary of P states $v - e^b + f^b = 2$. Plugging in the last two equations in Euler's formula for 3-balls (i.e., $v - e^b - e^i + f^b + f^i = t + 1$) directly implies the desired result $v + e^i - 3 = t$. □

In particular this implies that any triangulation of a 3-polytope with v vertices needs at least $v - 3$ tetrahedra. If we are interested in a triangulation of a three-dimensional polytopal complex formed by 3-polytopes P_1, P_2, \dots, P_n , then the number of tetrahedra needed for this triangulation is at least

$$\sum_{i=1}^n v_i - 3n,$$

where v_i is the number of vertices of P_i .

We construct classes of such polytopal complexes where it is NP-hard to decide whether there is a triangulation that actually achieves this bound. The crucial obstruction that we use comes from the following fact: Triangulations of 3-polytopes P and Q that are adjacent along a 2-face F are not independent from each other. They have to induce identical triangulations on F . Later we show that our construction can be embedded as a substructure of a suitable 3-ball and also as a substructure of the boundary of a suitable 4-polytope.

The only 3-polytopes that are used (as cells) in our construction are pyramids over 4-gons, triangular prisms, and cubes. The following lemma summarizes the necessary facts concerning triangulations of these polytopes.

Lemma 2.2.

- (i) *Let P be a pyramid over an n -gon G . Every triangulation of G induces a unique triangulation of P without interior edges.*
- (ii) *The triangular prism has a (up to symmetry) unique triangulation. In this triangulation the three newly added edges on the boundary are connected.*
- (iii) *The cube has two minimal triangulations (both of which have no interior edges and hence exactly five tetrahedra). In these triangulations the six newly added edges on the boundary form the edge-skeleton of a tetrahedron.*

Proof. (i) If a triangulation T_G of the n -gon G is given, the only way to extend a triangle of T_G to a tetrahedron of a triangulation of P is to form the convex hull with the apex of the pyramid. This uniquely determines the triangulation of P .

(ii) Any segment that joins two vertices of the triangular prism P lies on the boundary of P . Thus the triangulation has to arise by first truncating a vertex and then by truncating a 3-valent vertex of the remaining bipyramid over a triangle. It is easy to check that in this case the newly added edges always form a chain.

(iii) A complete classification of the triangulations of a cube can be found, for instance, in [1]. The unique minimal triangulation can be obtained by truncating every second vertex of the cube. After these truncations a central tetrahedron remains. The two possible triangulations arise from the two possibilities of choosing the central tetrahedron. \square

Figure 1 shows a possible and an impossible pattern of boundary edges for the triangulations of the triangular prism. Figure 2 shows the two minimal triangulations of a cube. All three types of polytopes mentioned in the above lemma play crucial roles in our construction.

The two ways in which the cubes can be triangulated are used to mimic the *switches*. Chains of cubes are used as kinds of wires to *transport information*. The fact that not all newly added edges on the boundary of a triangular prism can have the same “orientation” is used to mimic the logical connection for the NOT-ALL-EQUAL-3SAT. The fact that pyramids form no obstruction at all for triangulations is used to prevent the flow of information whenever necessary. Such pyramids serve as a kind of *isolation* between different wires.

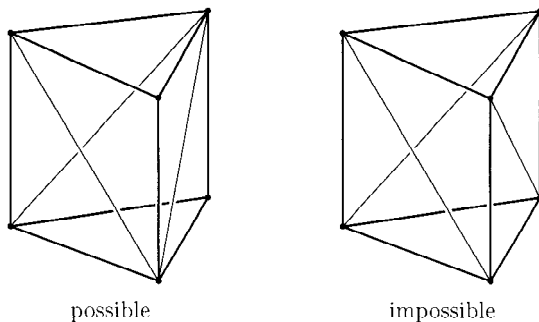


Fig. 1. Possible and impossible edge pattern in the boundary of the prism.

2.3. Transporting Information

Consider a three-dimensional polytopal complex that is exclusively formed by cubes. We are interested in the minimal triangulations of such a complex. For each cube in the complex there are by Lemma 2.1 exactly two ways to form a minimal triangulation. However, the choices for the individual cubes are not independent from each other. Whenever two cubes are adjacent along a 4-gon F , they have to induce the same diagonal on F . This implies that adjacent cubes have to use “opposite” triangulations (similar to the situation in Fig. 2 if we glue the two cubes together). Hence, a chain C_1, C_2, \dots, C_n of adjacent cubes in which C_i is adjacent to C_{i+1} and no other adjacencies occur has only two ways of being minimally triangulated. Figure 3 shows such a chain. In a sense such chains are devices for transporting information: the last cube of the chain “knows” how the first cube was triangulated.

It is even possible to form loops of n cubes such that there is no triangulation of the corresponding polytopal complex that achieves the lower bound of $5n$ tetrahedra. For this consider the complex K shown on the left of Fig. 4. It consists of six cubes (the central triangular prism is not part of the complex). There is no triangulation of K that uses the minimal number of five tetrahedra for every cube of K . To see this consider the diagram drawn on the right of Fig. 4. There the situation of the complex K is shown by looking only at the “top.” In a triangulation that uses only five tetrahedra for each cube




Fig. 2. There are two minimal triangulations of a cube.




Fig. 3. Minimal triangulation of a chain of cubes.

the diagonals of adjacent quadrangles have to form a “zigzag” pattern. It is easy to check that, no matter how we start, there is no consistent way to close the cycle.

2.4. *Isolating the Wires*

Consider a cube with one additional point in the interior (see Fig. 5). We consider the six pyramids that can be formed by this point and a face of the cube. The polytopal complex formed by these six pyramids has a boundary similar to the cube. However, if we are looking for a minimal triangulation of this polytopal complex it induces no obstructions on the boundary. This is an immediate consequence of Lemma 2.1(i).

Consider a polytopal complex K that contains some cubes. If we replace any cube C in K by the complex shown in Fig. 5, we “destroy” the obstructions that come from minimal triangulations of this cube.

2.5. *The Main Construction*

We now describe the *frame of reference* in which the whole construction is embedded. Consider a prism P over a regular triangle. We first slice this prism P into small polytopes




Fig. 4. No triangulation with $6 \cdot 5 = 30$ tetrahedra.




Fig. 5. Polytopal complex formed by six square pyramids.

by cutting it with hyperplanes that are parallel to its facets. For this let the height of the base triangle be equal to 1. For given $n > 1$ and a quadrangular face F of P we introduce n cutting hyperplanes parallel to F at distances $i/(3n + 2)$; $i = 1, \dots, n$. Let h be the height of the prism. For $m > 1$ we introduce $m - 1$ cutting hyperplanes parallel to the base triangle at a distance of $(i \cdot h)/m$; $i = 1, \dots, m - 1$. Figure 6 shows the situation for $n = 2$ and $m = 4$.

The cutting hyperplanes chop the original prism P into small polytopal pieces that form a polytopal complex $K_{n,m}$. All cells of $K_{n,m}$ are either combinatorial cubes or triangular prisms (all prisms meet the central axis of the figure). There are $3m(n^2 + n)$ cubical cells and m prisms. So every triangulation of $K_{n,m}$ requires at least $5 \cdot (3m(n^2 + n)) + 3 \cdot m$ tetrahedra. However, for $n > 1$ there is clearly no triangulation of $K_{n,m}$ that uses only that many tetrahedra, since $K_{n,m}$ contains many substructures of the type shown in Fig. 4 and hence it impossible to use five tetrahedra for each of the cubes.

Now, we are ready for the main construction. We start with an instance of a NOT-ALL-EQUAL-3SAT problem S on n variables and m clauses. The polyhedral complex $K_{4n,2m+1}$ serves as a frame of reference for our construction. In this complex we “destroy” most of the cubes by substituting them with the complex shown in Fig. 5 (containing six quadrangular pyramids). Cubes that are not destroyed by our construction are called “active.” The resulting polytopal complex will consist of a cubes, b prisms, and c quadrangular pyramids. It will have the property that one can triangulate it with $5a + 3b + 2b$




Fig. 6. A polytopal complex consisting of cubes and triangular prisms.




Fig. 7. Labels of the cubes.

tetrahedra if and only if the original NOT-ALL-EQUAL-3SAT problem S had a solution. The only information that we have to give is which cubes have to be *destroyed*.

In order to be able to refer to a particular 3-cell of our polyhedral complex we assign labels to them by the following rule. We orient our cutting hyperplanes in a way such that the planes parallel to the base triangle all point upward. The remaining cells are oriented such that they all point away from the center. To each cell we assign a quadruple (i, j, k, l) of integers that indicates how many planes of the different parallel classes point toward the cell. The first three numbers i, j, k refer to the three quadrangular faces of the original prism. The last number l indicates the layer of the cell. Figure 7 shows the numbers i, j, k for the case $n = 2$ (in any layer). If the cubes (i_1, j_1, k_1, l_1) and (i_2, j_2, k_2, l_2) share four supporting hyperplanes, we denote by $[(i_1, j_1, k_1, l_1), (i_2, j_2, k_2, l_2)]$ the set of all cubes in the convex hull of (i_1, j_1, k_1, l_1) and (i_2, j_2, k_2, l_2) . We call such a set an *interval of cubes*. The number of cubes in an interval is the *length* of the interval.

Each clause of S will be associated with a specific layer. The triangular prism in this layer will serve as the logical gate. The variables will correspond to belts of cubes around the center of the construction. In addition layer 1 will be used to produce copies and negations of each variable. All cubes of $K_{4n,2m+1}$ will be destroyed except those given in a set

$$A = \left(\bigcup_{i=1}^n A_i \right) \cap B \cap \left(\bigcup_{i=1}^m C_i \right)$$

described below. The full information of our construction is carried by the description




Fig. 8. The first belt in the base layer.

of this set. We first consider the lowest layer. For $i = 1, \dots, n$ we set

$$\begin{aligned}
 A_i = & [(0, 4i - 2, 4i - 2, 1), (4i - 2, 4i - 2, 0, 1)] \\
 & \cup [(4i - 2, 4i - 2, 0, 1), (4i - 2, 0, 4i, 1)] \\
 & \cup [(4i - 2, 0, 4i, 1), (0, 4i, 4i, 1)] \\
 & \cup [(0, 4i, 4i, 1), (4i, 4i, 0, 1)] \\
 & \cup [(4i, 0, 4i, 1), (0, 4i, 4i, 1)].
 \end{aligned}$$

In the lowest layer only those cubes that appear in one of the A_i remain active in our construction. The cubes of each A_i form a spiral chain that cycles $1\frac{2}{3}$ times around the center. Figure 8 shows the situation in the top of the lowest layer for A_1 . Cubes that are destroyed are marked by a dot. Now assume that we have a triangulation of the lowest layer that uses only five tetrahedra for each active cube. Under this assumption for each i there are only two possibilities to triangulate the cubes in A_i (the choice of one cube in A_i determines the triangulations of all other cubes in A_i). In Fig. 8 the diagonals drawn in the top faces of the active cubes indicate one possibility for the minimal triangulation. Switching *all* these diagonals describes the other possibility. Now consider the cubes with indices

$$\begin{aligned}
 (0, 4i - 2, 4i - 2, 1), & (4i - 2, 4i - 2, 0, 1), (4i - 2, 0, 4i - 2, 1), \\
 (0, 4i, 4i, 1), & (4i, 4i, 0, 1), (4i, 0, 4i, 1).
 \end{aligned}$$

We associate each of these cubes with a literal in our 3-SAT problem. The diagonals of the top faces of these cubes are darkened in Fig. 8. Either these diagonals point to

the nearby vertex of the central triangle (state T) or they are parallel to the opposite side of the central triangle (state F). We associate state T with the corresponding literal being *true*, and state F with the corresponding literal being *false*. By our construction the cubes

$$pos_i = \{(0, 4i - 2, 4i - 2, 1), (4i, 4i, 0, 1), (4i - 2, 0, 4i - 2, 1)\}$$

are always in the same state, and the cubes

$$neg_i = \{(0, 4i, 4i, 1), (4i - 2, 4i - 2, 0, 1), (4i, 0, 4i, 1)\}$$

are always in the opposite state of the cubes in pos_i . We associate each of the cubes in pos_i with the literal x_i of our 3-SAT problem and the cubes in neg_i with the literal $\neg x_i$.

This finishes the construction in the bottom layer. Now we provide each layer with “copies” of each variable. For this we make sure that the following cubes in the set B are not destroyed.

$$B = \bigcup_{i=1}^{2n} [(2i, 2i, 0, 1), (2i, 2i, 0, 2m + 1)] \\ \cup [(2i, 0, 2i, 1), (2i, 0, 2i, 2m + 1)] \\ \cup [(0, 2i, 2i, 1), (0, 2i, 2i, 2m + 1)].$$

Each of the intervals in B forms a pile of cubes over one of the cubes in the lowest layer, which represent the literals. Thus every layer in our construction is supplied with three copies of each literal. In particular, among all cubes of the form $(i, i, 0, l)$ for fixed l and $i \in 1, \dots, 4n$ we find copies of each literal. A similar statement holds for the cubes of the form $(i, 0, i, l)$ and $(0, i, i, l)$. One has to be a bit careful, since the role of the negated and the original literals changes with the parity of the layer. In layers with even l the cubes from B will be the only active cubes in our construction. The destroyed cubes in the even layers serve as an isolation between the layers with odd index.

We finally come to the encoding of a particular clause of our 3-SAT. We have altogether m clauses. The construction for a specific clause G_j will be realized in the layer $2j + 1$ and is as follows. Let $G_j = (a, b, c)$ be the clause, where a, b , and c are literals taken from the set $\{x_1, \dots, x_n, \neg x_1, \dots, \neg x_n\}$. We set

$$C_i^1 = \begin{cases} [(0, 4i - 2, 4i - 2, 2j + 1), (0, 4i - 2, 0, 2j + 1)] \\ \cup [(0, 4i - 2, 0, 2j + 1), (0, 1, 0, 2j + 1)] & \text{if } a = x_i, \\ [(0, 4i, 4i, 2j + 1), (0, 4i, 0, 2j + 1)] \\ \cup [(0, 4i, 0, 2j + 1), (0, 1, 0, 2j + 1)] & \text{if } a = \neg x_i; \end{cases}$$

$$C_i^2 = \begin{cases} [(4i - 2, 0, 4i - 2, 2j + 1), (0, 0, 4i - 2, 2j + 1)] \\ \cup [(0, 0, 4i - 2, 2j + 1), (0, 0, 1, 2j + 1)] & \text{if } b = x_i, \\ [(4i, 0, 4i, 2j + 1), (0, 0, 4i, 2j + 1)] \\ \cup [(0, 0, 4i, 2j + 1), (0, 0, 1, 2j + 1)] & \text{if } b = \neg x_i; \end{cases}$$

$$C_i^3 = \begin{cases} [(4i, 4i, 0, 2j + 1), (4i, 0, 0, 2j + 1)] \\ \cup [(4i, 0, 0, 2j + 1), (1, 0, 0, 2j + 1)] & \text{if } c = x_i, \\ [(4i - 2, 4i - 2, 0, 2j + 1), (4i - 2, 0, 0, 2j + 1)] \\ \cup [(4i - 2, 0, 0, 2j + 1), (1, 0, 0, 2j + 1)] & \text{if } c = \neg x_i. \end{cases}$$




Fig. 9. Connecting the literals to a clause.

The cubes in $C_i = C_i^1 \cup C_i^2 \cup C_i^3$ are the only active cubes in layer $2j + 1$. They connect representatives of the literals in the clause G_i to the cubes $(0, 1, 0, 2j + 1)$, $(0, 0, 1, 2j + 1)$, and $(1, 0, 0, 2j + 1)$, which are next to the central triangular prism of this layer. Figure 9 shows the situation for a (slightly artificial) clause $(\neg x_1, \neg x_2, \neg x_2)$. Again this picture represents the top faces of this layer. The active cubes in the construction are marked by a consistent choice for the diagonals in a triangulation that uses five tetrahedra for each cube. The figure shows the situation $x_1 = true, x_2 = false$.

Our construction is such that we can now read off the logical state of the literals a, b and c from the triangulation induced in the cubes $(0, 1, 0, 2j + 1)$, $(0, 0, 1, 2j + 1)$, and $(1, 0, 0, 2j + 1)$ (these are the cubes adjacent to the quadrangles of the central prism). We orient the diagonal on their top face toward the face where it meets the central triangle. For each of these cubes we say that its state is *true* if the diagonal is oriented clockwise with respect to the central axis of our construction. Otherwise we say it is in state *false*. Under our general assumption that we use five tetrahedra for the triangulation of each cube, the state of the cubes $(0, 1, 0, 2j + 1)$, $(0, 0, 1, 2j + 1)$, and $(1, 0, 0, 2j + 1)$ is true if and only if the corresponding literal is *true*. The crucial lemma that makes the construction work can be stated as follows.

Lemma 2.3. *A triangulation of $X = \{(0, 1, 0, 2j + 1), (0, 0, 1, 2j + 1), \text{ and } (1, 0, 0, 2j + 1)\}$ in which five tetrahedra are used for each cube is compatible with a triangulation of the central prism $(0, 0, 0, 2j + 1)$ if and only if at least one cube in X is in state true and at least one cube in X is in state false.*

Proof. If all diagonals of the cubes are oriented clockwise (counterclockwise), then the diagonals of the quadrangular faces of $(0, 0, 0, 2j + 1)$ correspond to the “impossible” situation in Fig. 1. By Lemma 2.1 this cannot be extended to a triangulation of $(0, 0, 0, 2j + 1)$. \square

We obtain immediately

Lemma 2.4. *Let S be an instance of a NOT-ALL-EQUAL-3SAT Problem. Let A_i , B , and C_i be defined as described by the above construction, and let $K(S)$ be a polytopal complex that arises from $K_{4n, 2m+1}$ by destroying all cubes except those in*

$$A = \left(\bigcup_{i=1}^n A_i \right) \cup B \cup \left(\bigcup_{i=1}^m C_i \right).$$

Then K can be triangulated with exactly five tetrahedra for every active cube if and only if S had an admissible assignment.

Proof. Assume that there is an assignment to the variables x_1, \dots, x_n such that in each clause of S there is at least one *true* and one *false* literal. In this case we choose the triangulation that corresponds to this choice of truth values. In the layers in which the clauses are encoded the cubes adjacent to the triangles are by construction in the same state as the corresponding literals. By Lemma 2.3, the corresponding prism can be triangulated consistently. Conversely, if there is a triangulation of $K(S)$ that uses exactly five tetrahedra in each active cube, then by Lemma 2.3 it cannot be possible that all cubes around a prism are in the same state. Hence from the triangulation one can read off an admissible assignment of variables for S \square

It is easy to check that one can derive coordinates for the vertices of $K(S)$ together with a combinatorial description of the face lattice in polynomial time from the data of S . Thus we obtain:

Theorem 2.5. *The problem of determining whether a polytopal 3-ball (given by coordinates of the vertices and a description of the face lattice) can be triangulated using k tetrahedra is NP-complete.*

Proof. Our above translation of NOT-ALL-EQUAL-3SAT to a triangulation problem for $K(S)$ proves that $K(S)$ can be triangulated with $3p + 5a + 12d$ tetrahedra if and only if S had an admissible assignment. Here t , a , and d are the number of prisms, alive cubes, and destroyed cubes in $K(S)$, respectively. This proves the NP-hardness part. To see that the problem is NP-complete we just observe that it can be checked in polynomial time whether a given cell decomposition of $K(S)$ into tetrahedra is indeed a polytopal complex. \square

3. Boundaries of 4-Polytopes

Now we prove that our construction can also be applied to the cell complexes that appear as boundaries of 4-polytopes. We first show how one can lift the original complex $K_{4n,2m+1}$ (without any destroyed cubes) to a substructure of the boundary of a suitable 4-polytope. This can be easily done since $K_{4n,2m+1}$ arises as a substructure of an arrangement of hyperplanes $H_{4n,2m+1}$ (the cutting hyperplanes together with the supporting hyperplanes of the original prism). The lifting process can be performed by assigning to each hyperplane a canonical piecewise linear convex function. The sum of all these functions describes a convex function whose graph has the desired properties.

Lemma 3.1. *There is a 4-polytope $P_{4n,2m+1}$ that contains a cell complex (combinatorially) isomorphic to $K_{4n,2m+1}$ in its boundary.*

Proof. Assume the hyperplane arrangement $H_{4n,2m+1}$ consists of k hyperplanes h_1, \dots, h_k . Each of these hyperplanes is given as the zero set of an affine functional $f_i(x, y, z) = 0$. Without loss of generality we may assume that $f_i(0, 0, 0) < 0$ for $i = 1, \dots, k$. Let

$$\varphi(x) := \begin{cases} x & \text{if } x > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\varphi(f_i(x, y, z))$ is a convex piecewise linear function. It has exactly two linear components. These components are divided by h_i . The sum

$$f(x, y, z) = \sum_{i=1}^k \varphi(f_i(x, y, z))$$

is still a piecewise linear convex function. Over each cell of $K_{4n,2m+1}$ this function is linear. The set

$$\{(x, y, z, w) \mid (x, y, z) \in K_{4n,2m+1}; w > f(x, y, z)\}$$

is an (unbounded) convex four-dimensional polyhedron. Intersecting it with half-space

$$\left\{ (x, y, z, h) \in \mathbb{R}^4 \mid h < \max_{(x,y,z) \in K_{4n,2m+1}} (f(x, y, z)) \right\}$$

produces a convex 4-polytope that contains a structure combinatorially isomorphic to $K_{4n,2m+1}$ in its boundary. □

Figure 10 shows a corresponding picture for one dimension less. It represents a lifting of a hyperplane arrangement in the plane to a 3-polytope. Actually the arrangement shown in this figure corresponds to a two-dimensional subconfiguration of our construction.

What remains is to provide a method that models the process of “destroying cubes” in our main construction from Section 2. We do this by stellating the facets of the polytope $P_{4n,2m+1}$.




Fig. 10. Lifting hyperplane arrangements to polytopes.

Assume $P_{4n,2m+1}$ has altogether l facets $\mathcal{F} = \{F_1, \dots, F_l\}$. For each facet $F_i \in \mathcal{F}$ we choose a point $p_i \in \text{relint}(F_i) \subset \mathbb{R}^4$ in its relative interior (for instance its center of gravity). Let n_i be an outer normal vector of F_i . For sufficiently small $\varepsilon > 0$ the points $q_i = \varepsilon \cdot n_i + p_i$ with $i = 1, \dots, l$ stellate the polytope, i.e., the polytope $Q_{4n,2m+1} = \text{conv}(P_{4n,2m+1} \cup \{q_1, \dots, q_l\})$ has as vertex set all vertices of $P_{4n,2m+1}$ together with all points q_1, \dots, q_l . All faces of $Q_{4n,2m+1}$ are pyramids over n -gons (the apex of such a pyramid is a point q_i ; the base face corresponds to a 2-face of F_i). It is easy to check that one can calculate concrete rational coordinates for the points q_i in polynomial time from the data of $P_{4n,2m+1}$.

Since all facets of $Q_{4n,2m+1}$ are pyramids Lemma 2.2(i) tells us that every triangulation of the 2-faces of $Q_{4n,2m+1}$ is compatible with a suitable triangulation of the boundary of $Q_{4n,2m+1}$. Now let S be a concrete instance of a NOT-ALL-EQUAL-3SAT problem and let $K(S)$ be the corresponding 3-ball of our construction. Let \mathcal{F}' be the set of facets of $P_{4n,2m+1}$ that correspond to active cubes of $K(S)$ or that correspond to triangular prisms of $K(S)$. Let $\mathcal{I}(S) = \{i \mid 0 < i \leq l; f_i \notin \mathcal{F}'\}$ be the index set of all other faces. We consider the polytope

$$Q(S) = \text{conv} \left(P_{4n,2m+1} \cup \bigcup_{i \in \mathcal{I}(S)} \{q_i\} \right).$$

Lemma 3.2. *Let F_1, F_2, \dots, F_r be the facets of $Q(S)$ and assume that for $i = 1, \dots, r$ the facet F_i has v_i vertices. The boundary of the polytope $Q(S)$ can be triangulated with exactly $\sum_{i=1}^r v_i - 3r$ tetrahedra if and only if S has an admissible assignment.*

Proof. By construction the boundary of the polytope $Q(S)$ contains a polytopal subcomplex that is combinatorially equivalent to $K(S)$. This part of the boundary of $Q(S)$ has a triangulation that uses only five tetrahedra for each cube if and only if $K(S)$ had such a triangulation. The remaining facets of $Q(S)$ are pyramids over n -gons and introduce no additional obstructions. Hence Lemma 2.4 implies the desired result. \square

As an immediate consequence we obtain the polytopal counterpart of Theorem 2.5.

Theorem 2.5. *The problem of determining whether the boundary of a 4-polytope (given by coordinates of the vertices of the polytope) can be triangulated using k tetrahedra is NP-complete.*

Proof. The proof is analogous to the proof of Theorem 2.5. □

References

1. J. De Loera: Computing minimal and maximal triangulations of convex polytopes, Manuscript, 1999, 14pp.
2. M.J. Garey and D.S. Johnson: *Computers and Intractability*, Freeman, San Francisco, CA, 1979.
3. N.E. Mnëv, The universality theorems on the classification problem of configuration varieties and convex polytopes varieties, in: *Topology and Geometry—Rohlin Seminar* (O. Ya Viro, ed.), Lecture Notes in Mathematics, vol. 1346, Springer-Verlag, Berlin, 1988, pp. 527–544.
4. J. Richter-Gebert, *Realization Spaces of Polytopes*, Lecture Notes in Mathematics, vol. 1643, Springer-Verlag, Berlin, 1996.
5. J. Richter-Gebert, Testing orientability for matroids is NP-complete, *Advances in Applied Mathematics*, **23** (1999), 78–90.
6. P. Shor, Stretchability of pseudolines is NP-hard, in: *Applied Geometry and Discrete Mathematics—The Victor Klee Festschrift* (P. Gritzmann, B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4, American Mathematical Society, Providence, RI, 1991, pp. 531–554.

Received February 17, 1999, and in revised form October 20, 1999. Online publication May 16, 2000.

Note added in proof. Very recently substantial progress has been made on the question that originally motivated this paper. A. Below, J. de Loera, and L. Richter-Gebert proved the following result.

Theorem. *It is NP-complete to decide whether a 3-polytope can be triangulated with less than k simplices.*

In principle the results of the present paper could be derived as a corollary of this theorem by building a pyramid over the constructed 3-polytope. Nevertheless there are substantial differences in the proving techniques that make the present paper still valuable. The above theorem is proved essentially based on visibility arguments in the interior of the polytope. By this a structurally very complicated polytope is obtained and the proof requires many considerations that take care of the actual coordinate sizes. Compared with this the proof presented in this paper is structurally very simple and relies much more on combinatorial than on metrical considerations. In fact, inspection into the proofs of the main theorems of the present paper shows that the results are still valid even if all used 3-faces are just cubes, pyramids, or triangular prisms. A similar result would not be achievable by the visibility methods that are crucial for the proof of the above theorem.