

On the Difficulty of Triangulating Three-Dimensional Nonconvex Polyhedra*

Jim Ruppert and Raimund Seidel

Computer Science Division, University of California at Berkeley,
Berkeley, CA 94720, USA
ruppert@cs.berkeley.edu
seidel@lyra.berkeley.edu

Abstract. A number of different polyhedral *decomposition* problems have previously been studied, most notably the problem of triangulating a simple polygon. We are concerned with the *polyhedron triangulation* problem: decomposing a three-dimensional polyhedron into a set of nonoverlapping tetrahedra whose vertices must be vertices of the polyhedron. It has previously been shown that some polyhedra cannot be triangulated in this fashion. We show that the problem of deciding whether a given polyhedron can be triangulated is NP-complete, and hence likely to be computationally intractable. The problem remains NP-complete when restricted to the case of star-shaped polyhedra. Various versions of the question of how many Steiner points are needed to triangulate a polyhedron also turn out to be NP-hard.

1. Introduction

Polyhedron decomposition concerns the problem of dividing a d -dimensional polyhedron into simpler polyhedra. It has applications in robotics, computer-aided design, computer graphics, and other fields. Two main types of decomposition have been considered: *coverings* and *partitions*. In a covering the simpler polyhedra may overlap arbitrarily, whereas in a partition overlaps are allowed only in common faces of dimension $< d$ (for $d = 3$ this means the only overlaps are common vertices, edges, and polygonal faces). Usually, only vertices from the original polyhedron may be used as vertices of the subpolyhedra, although sometimes we may allow additional “Steiner” points to be used. We may require a

* This work was supported by National Science Foundation Grant CCR-8809040.




Fig. 1. The number of tetrahedra in a triangulation of a polyhedron is not unique.

decomposition into polyhedra that are *convex*, or *monotone*, or otherwise “simpler” than the original polyhedron. Here, we investigate the problem of partitioning a three-dimensional polyhedron into simplices, i.e., tetrahedra. We call this the *polyhedron triangulation* problem. Partitioning a two-dimensional polyhedron into simplices is known as *polygon triangulation*, and has been well studied. Many efficient algorithms are known for producing triangulations, with the recent linear-time algorithm of Chazelle [4] being asymptotically optimal.

In this paper we show that the three-dimensional triangulation problem is significantly more difficult than the two-dimensional triangulation problem. One difference between the problems lies in the size of the resulting partitions: every triangulation of an n -sided polygon produces $n - 2$ triangles, but the number of tetrahedra in a triangulation of a given polyhedron is not unique. For example, a bipyramid with a triangular base may be partitioned into either two or three tetrahedra (see Fig. 1).

Even more significant is the difference that any (non-self-intersecting) polygon may be triangulated, whereas there exist simple three-dimensional polyhedra which cannot be triangulated. The following example is due to Schönhardt [12] and is referred to in Chapter 10 of [9]. The six-vertex polyhedron P in Fig. 2 is




Fig. 2. A polyhedron that cannot be triangulated.

constructed as follows. Starting with a triangular prism, “twist” the top face by a small amount. The three side faces cannot remain planar, so allow them to “bend in” along the appropriate diagonals to become two triangular faces.

The polyhedron P cannot be triangulated because no tetrahedron formed by four of its vertices is wholly contained in P . If a tetrahedron was contained in P , then each pair of vertices of the tetrahedron would be able to “see” each other within P . In the case of P , the pairs of vertices which can see each other are exactly those which form edges of P . Since the edge-graph of P contains no complete subgraph on four vertices, P contains no tetrahedron.

O’Rourke has posed the problem of characterizing polyhedra that cannot be triangulated [9]. We show that it is unlikely that there exists such a characterization that is computationally useful. We investigate the related question of deciding whether a given polyhedron can be triangulated. Our main result states that this problem is NP-complete. Thus, we cannot expect (unless $P = NP$, see [8]) to find an algorithm that can decide whether a given polyhedron can be triangulated and that has a running time bounded by a polynomial in the number of vertices of the polyhedron. NP-completeness is shown by transformation from the Satisfiability problem [8, p. 259]. We show how for any Boolean formula in conjunctive normal form we can construct a three-dimensional polyhedron that can be triangulated iff the Boolean formula is satisfiable. The main tool in this construction is a gadget that we call a *niche*, which, when attached to a polyhedron, restricts the possible triangulations. In particular, niches can force certain tetrahedra to appear, and they can prevent certain pairs of tetrahedra from appearing simultaneously.

NP-completeness holds even for a fairly restricted class of polyhedra. The construction produces polyhedra without holes or dangling faces or edges (i.e., *simple polyhedra*). These polyhedra will also have triangular faces. The NP-completeness does not depend on coplanarities of faces or other degeneracies. We extend the proof to show that the problem is still NP-complete even if we restrict the input to star-shaped polyhedra. As a corollary of this, we are able to show NP-hardness for several problems concerning polyhedron triangulation when Steiner points are allowed.

Most other work in polyhedron decomposition has focused on two-dimensional problems, a number of which have been shown to be NP-hard [10], [6]. In the three-dimensional case Chazelle has investigated partitions into convex pieces [3], and Bajaj and Dey have recently given an algorithm with an improved running time [2]. Chazelle and Palios have shown that any three-dimensional n -vertex polyhedron with r reflex edges can be triangulated (if Steiner points are allowed) using $O(n + r^2)$ tetrahedra in time $O(nr + r^2 \log r)$ [5]. Chazelle also showed that $\Omega(n^2)$ tetrahedra are necessary in the worst case [3]. Triangulation of *convex* three-dimensional polyhedra is fairly easy [7]. Von Hohenbalken gives an algorithm that partitions a d -dimensional convex polyhedron into simplices in time that is linear in the number of simplices produced [13]. The problem of triangulating a set of *points* in three dimensions is investigated in [7] and [1].

A note about terminology: what we refer to here as the three-dimensional polyhedron *triangulation* problem has had a variety of names in the literature. The

three-dimensional case has been referred to as *tetrahedrization*, *tetrahedralization*, and *tetrahedralization*. *Triangulation* usually refers to the two-dimensional case or the d -dimensional case. Here we follow [5], and use *triangulation* for three dimensions. It will be noted when we are referring to a two-dimensional triangulation. We will, however, always use *triangle* to refer to a two-dimensional simplex and *tetrahedron* for a three-dimensional simplex. The general d -dimensional case of polyhedron decomposition has also been called *simplication* and *simplicial decomposition* because a simplicial complex is produced.

2. Niches and the Illuminant Lemma

In this section we describe a gadget called a *niche*, that can be used to force any triangulation of a polyhedron to satisfy certain conditions. Our construction is based on the following simple fact: In every triangulation of a polyhedron P every triangular face of the boundary of P must appear in exactly one tetrahedron. This tetrahedron will include a fourth vertex that is not in the face. This fourth vertex must be able to “see” all of the triangular face. We would like to take advantage of this by specifying certain triangular faces that can only form a tetrahedron with a certain vertex or one of a small set of vertices.

The “twisted prism” introduced in Section 1 can be adapted to serve this purpose. Note that though no vertex of the prism can see all of the “base” (the bottom triangular face), other points may be able to see all of the base, as for instance the center point of the top face (see Fig. 3). If the top face of the prism is removed, then there is a triangular cone of points that can all see the inside of the base of the prism. In Fig. 4 the cone is determined by the planes containing the shaded faces $q_1q_2p_2$, $q_2q_3p_3$, and $q_3q_1p_1$. This uncapped prism, consisting of seven triangles on six vertices, which by itself cannot be triangulated, but which might be triangulated using another vertex, we refer to as a *niche*. Niches will be used as parts of the boundaries of larger polyhedra. When a niche is part of a larger polyhedron, we refer to a point in the polyhedron that can see the entire




Fig. 3. Entire prism visible from the center x_0 of the top face.




Fig. 4. Visibility cone for a niche.




Fig. 5. A niche attached to a cube.

base of the niche as a *viewpoint* of the niche. We refer to the set of all viewpoints as the *illuminant* of the niche. Figure 5 shows a niche, a viewpoint of the niche, and the illuminant of the niche. For simplicity, we define the illuminant so that it does not include any of the interior of the niche. Note that the illuminant of a niche consists of a single connected portion of the interior of the polyhedron, incident to the open end of the niche. The illuminant may not be the intersection of a triangular cone with the polyhedron, as some points may be obstructed by other faces of the polyhedron, preventing them from seeing the entire base of the niche.

For our purposes it is important that vertices can be placed so that they can form tetrahedra with the bases of certain niches, and that we can construct niches so that their bases can form tetrahedra only with certain vertices. The following lemma shows that this is relatively easy to achieve. It shows how to construct niches with prespecified illuminants.

Lemma 1 (The Illuminant Lemma). *Let F be a face of a polyhedron P and let C be a triangular cone that intersects the relative interior of F in a triangle T . Let C_T be the set of points in $P \cap C$ that can see all of T . Then it is possible to attach a niche N to F such that $C_T \subset I \subset C$, where I is the illuminant of N . (See Fig. 6.)*

Proof. We show how to construct a niche N with the desired properties. Let $\tilde{p}_1, \tilde{p}_2, \tilde{p}_3$ be the vertices of the triangle T . Let q_1, q_2, q_3 be the intersections of the edges of the cone C with a plane parallel to the face F and slightly “outside” F (see Fig. 7). One at a time, move $\tilde{p}_1, \tilde{p}_2, \tilde{p}_3$ “out” from T on F to produce $p_1, p_2,$




Fig. 6. The setting for the Illuminant Lemma.




Fig. 7. Placing q_1, q_2, q_3 .




Fig. 8. Placing p_1, p_2, p_3 .




Fig. 9. The constructed niche.

p_3 in the following way: move \tilde{p}_1 a small positive distance in the direction $\overrightarrow{\tilde{p}_3\tilde{p}_1}$, staying within F . Also move \tilde{p}_2 in the direction $\overrightarrow{\tilde{p}_1\tilde{p}_2}$ and move \tilde{p}_3 in the direction $\overrightarrow{\tilde{p}_2\tilde{p}_3}$ (see Fig. 8). The amount of movement must be small enough so that the niche N specified below is well formed, i.e., any two of its triangular faces only intersect in a complete common edge, or a common vertex.

Let the niche N consist of the seven faces $q_3q_2q_1, p_1q_1p_2, q_1q_2p_2, p_2q_2p_3, q_2q_3p_3, p_3q_3p_1, q_3q_1p_1$. Attach N to F by removing the triangle $p_1p_2p_3$ from F , and triangulate (here we mean a two-dimensional polygon triangulation) the rest of F in any fashion (see Fig. 9).

Now note that the three faces $q_1q_2p_2, q_2q_3p_3, q_3q_1p_1$ of the niche N are contained in the boundary of the cone C , so every point in C is on the “inside” of these faces. The other three side faces of the niche, $p_1q_1p_2, p_2q_2p_3, p_3q_3p_1$, were moved “outward” from the cone C . By assumption every point in C_T can see all of T , and thus by construction every such point can see the entire base triangle $q_1q_2q_3$ of the niche N and is therefore in I , the illuminant of N . Thus $C_T \subset I$. (Usually $C_T = I$, but it is possible that a point in C may be obstructed from seeing all of T , and still see $q_1q_2q_3$.)

On the other hand, any point not in C will be outside the plane of one of the faces $q_1q_2p_2, q_2q_3p_3, q_3q_1p_1$, and will thus not be in the illuminant of N . Thus $I \subset C$. □

The illuminant is interesting because any vertex that can form a tetrahedron with the bottom triangular face of the niche must be contained in it. Often, we will want to triangulate the entire niche (i.e., form a tetrahedron with each of its triangular faces) from some vertex. For any given set of vertices in C_T , the niche can be constructed so that these vertices can see all of the side faces as well as the bottom face, and hence any of them could triangulate the entire niche. When we place vertices in illuminants, below, we will make sure to place them in the region C_T .

Niches and groups of niches will be used as “gadgets” that force any triangulation to have certain properties. For example, if exactly two vertices x, y of the polyhedron are in the illuminant of a niche that has base $q_1q_2q_3$, then exactly one of the tetrahedra $T_1 = xq_1q_2q_3$ or $T_2 = yq_1q_2q_3$ must appear in any triangulation (see Fig. 10). Since tetrahedra may not overlap in a triangulation (except in




Fig. 10. Tetrahedra T_1 and T_2 overlap; at most one of them can appear in any triangulation.

common faces), each of T_1 and T_2 may obstruct other tetrahedra from appearing with them in a triangulation. Other more complicated gadgets may be built up in a similar fashion.

3. NP-Completeness of the Decision Problem

Theorem 1. *It is NP-complete to decide whether a given three-dimensional polyhedron can be triangulated without using additional Steiner points.*

Proof. We use a polynomial-time transformation from instances of the Satisfiability problem. We assume that the Satisfiability instance is a Boolean expression in conjunctive normal form with n variables and m clauses. For example:

$$(X_1 + \overline{X_2} + \overline{X_3} + X_4)(X_2 + X_3)(\overline{X_3} + \overline{X_4} + X_5)$$

is an expression with five variables X_1, X_2, X_3, X_4, X_5 and three clauses. The first clause contains four literals, the negative literals $\overline{X_2}, \overline{X_3}$, and the positive literals X_1, X_4 .

Given an expression E that is an instance of Satisfiability, we show how to construct a polyhedron P such that

$$E \text{ is satisfiable} \iff P \text{ can be triangulated.}$$

We at first give a rough outline of the construction of P from which the implication

$$P \text{ can be triangulated} \Rightarrow E \text{ is satisfiable}$$

can be proved easily. Subsequently we refine this construction significantly so that the reverse implication

$$E \text{ is satisfiable} \Rightarrow P \text{ can be triangulated}$$

can also be proved. After stating some construction constraints that this refined construction must satisfy, we show how these constraints ensure the two-way




Fig. 11. Clause niches attached to polyhedron P .

implication. Lastly, we show that P can indeed be constructed to satisfy the constraints.

Outline of Construction. For the time being imagine the general shape of the polyhedron P we construct as that of a rectangular box, with tiny niches arranged on two sides of the box. There will be *clause niches* that correspond to clauses, and *variable niches* corresponding to variables. We also use two special kinds of vertices: *truth-setting vertices* and *literal vertices*.

The clause niches will be attached to the bottom of the box, and will be constructed (using the Illuminant Lemma) such that their illuminants form skinny vertical regions that do not intersect within the polyhedron, as shown in Fig. 11. There will be one literal vertex for each occurrence of each literal, with each literal vertex being placed on the top of the box in the illuminant of the corresponding clause. Each variable's literals will be arranged in two rows, one for the positive literals and one for the negative literals. Figure 12 shows a Satisfiability expression and the resulting clause niches and literal vertex placements.

The idea is that a clause niche may be triangulated only from its corresponding literal vertices, as they are the only vertices in its illuminant. The literal vertex that triangulated the niche corresponds to a literal that satisfies the clause in the expression E . We need a way to enforce a "truth assignment," to prevent a




Fig. 12. Regions of polyhedron corresponding to variables, niche for clause $(\bar{X}_1 + \bar{X}_2 + X_3)$, and placement of three corresponding literal vertices.




Fig. 13. Truth-setting vertices are viewpoints of variable niches.

variable’s positive and negative literal vertices from simultaneously being used to triangulate clause niches. Because of the way we have placed the literal vertices in rows, we can use a gadget like that of Fig. 10 in which one of two specified tetrahedra must be present in any triangulation, each of which “blocks” one row of literal vertices from seeing their corresponding clause niches. For each variable in the expression, we add to the polyhedron a variable niche along the back face, and two truth-setting vertices on the front face, as shown in Fig. 13. We use the Illuminant Lemma to ensure that only the two truth-setting vertices can be used to triangulate the variable niche. The use of the TRUE truth-setting vertex will represent the variable being set true, the FALSE vertex will represent false.

Given a triangulation of the polyhedron P , we can show that the expression E is satisfiable. We can interpret a truth assignment from the triangulation of the variable niches, since the base triangle of each variable niche forms a tetrahedron with either the TRUE truth-setting vertex or the FALSE truth-setting vertex. If the TRUE truth-setting vertex was used, the variable is set true, and the negative literals will be blocked from seeing their clause niches. If the FALSE truth-setting vertex was used, the variable is set false, and the positive literals will be blocked from seeing their clause niches. We know that this truth assignment is a satisfying truth assignment because of the following: The base triangle of each clause niche must have formed a tetrahedron with one of the literal vertices. This vertex must correspond to a literal set true, since vertices corresponding to false literals are blocked from triangulating their clause niches. The literal vertex that triangulates the niche represents a literal that satisfies the corresponding clause. Thus we have shown that the implication

$$\text{polyhedron } P \text{ can be triangulated} \Rightarrow \text{expression } E \text{ is satisfiable}$$

holds for this simple polyhedron construction.

The Refined Construction. We run into difficulties if we use the above construction and try to show that a satisfying truth assignment for E yields a triangulation of the polyhedron P . We would like to use the satisfying assignment to guide us in triangulating the clause niches and variable niches, but this yields a partial triangulation in which many tetrahedra stretch across the interior of the poly-

hedron. The remaining untriangulated portion is shaped very irregularly, and it is unclear whether the triangulation can be completed. Instead of trying to do this, we refine our construction to produce a polyhedron that is more complicated, but in which we can describe how to complete the triangulation.

In the refined construction, each variable will correspond to a portion of the polyhedron. One at a time, the truth setting of each variable is used to determine how to triangulate that variable's portion of the polyhedron. When all variables have been considered, the triangulation will be complete. Things are set up so that there is minimal interaction between the triangulations of the different portions, so the overall triangulation proceeds smoothly. By working incrementally, we also simplify the explanation, as we need only describe the triangulation of a single variable's portion of the polyhedron.

We assume the given Satisfiability expression has m clauses and n variables. To simplify the construction, we restrict the Satisfiability instances to those in which each variable appears exactly three times, once as a negative literal, and twice as a positive literal. Each appearance must be in a different clause. This restricted version of Satisfiability is easily shown to be NP-complete, by extending an argument on p. 259 of [8] (a variable with only one positive and one negative literal can be eliminated by resolution, i.e., by combining the two clauses that contain these two literals).

The starting point for our construction will be a sort of "distorted wedge." Figure 14 shows an example for the case in which $n = m = 2$. We sometimes refer to the (x, y, z) coordinates of points. Figure 14 shows the orientation of the coordinate axes and gives coordinates for several points. Other coordinates will




Fig. 14. Starting point for the polyhedron construction.

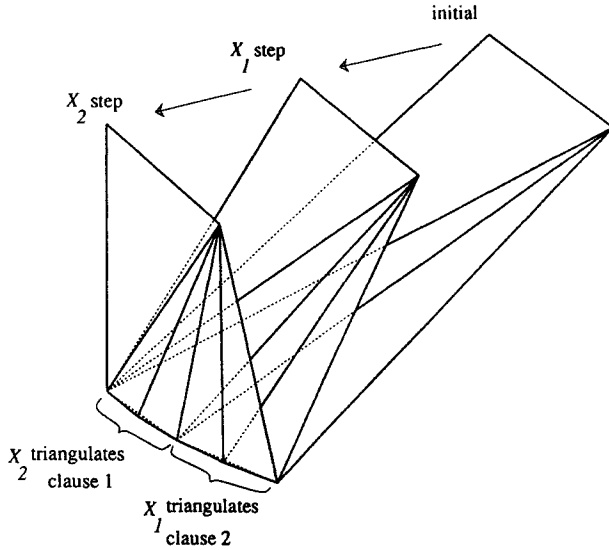


Fig. 15. The interface sweeping across the polyhedron. After each step, everything to the right of the interface has been triangulated.

be specified later. In the following we add niches and extra vertices and faces to this wedge to finish the construction.

The base of the wedge consists of $2m + 1$ vertices $c_1, c_2, \dots, c_{2m+1}$, lying on a parabola to be specified later. These vertices bound m triangular faces $c_1c_2c_3, c_3c_4c_5, \dots, c_{2m-1}c_{2m}c_{2m+1}$. We will later attach a clause niche to each of these triangles. The top of the polyhedron is a row of n squares, one per variable. To each of these squares, we will attach a “roof” containing the variable niche and the variable’s three literal vertices.

Figure 16 shows a single variable’s roof. For reference, we also include the chain of triangles from the base of the polyhedron. The front gable of the roof is the face $z_F z_A z_T$, and the back gable is $y_1 y_2 y_3$. Attached to the back gable will be the variable’s niche, which will have only two vertex viewpoints, z_F and z_T . On the rooftop will be the literal vertices x_1, x_2 , and x_3 , with x_1 and x_2 representing the two positive literals and x_3 representing the single negative literal (without loss of generality we assume that x_1 is contained in a lower-indexed clause than x_2). The idea is that if the variable is set true in the satisfying assignment, then z_T can be used to triangulate the variable niche, and then x_1 and x_2 can be used to triangulate their respective clause niches if needed. Using z_T to triangulate the variable niche will “block” x_3 (the negative literal, which is false) from triangulating its clause niche. The case when the variable is set false is handled symmetrically.

We will attach the superscript i to a roof vertex if we need to specify that it belongs to the roof of the i th variable. For example, notice that z_T^i and z_F^{i+1} are two names for the same vertex. Vertices y_3^i and y_1^{i+1} are also identified. Often, when referring to a generic variable, we drop the i superscripts.

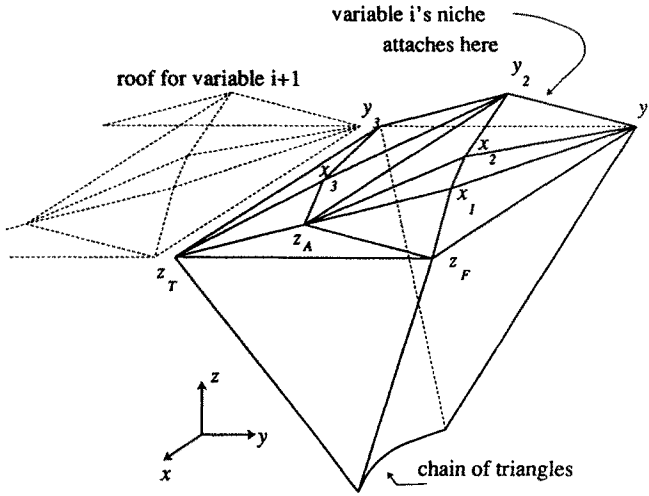


Fig. 16. A variable's roof (seen from the "right").

A number of geometric details must be dealt with before we can describe the actual triangulation corresponding to a satisfying truth assignment. We first give a brief idea of how the triangulation process would proceed. The "interface" between the triangulated and untriangulated portions of the polyhedron will "sweep" across the polyhedron, triangulating one variable's portion at a time (see Fig. 15).

After the i th step of the triangulation (n steps altogether, one per variable), the interface will have the following properties:

- It will have the form of a topological disk.
- It will consist of triangles.
- Each triangle contains the vertex z_T^i , which will be called the *hub* of the interface. One triangle is $z_T^i y_3^i c_1$, and the other triangles each include z_T^i and two vertices from $c_1, c_2, \dots, c_{2m+1}$, as specified below.
- The triangulation will include the niches of clauses satisfied so far. Since the clause niches are attached to triangles in the chain, we have:
 - If the j th clause is not satisfied by one of the first i variables, then the triangle $z_T^i c_{2j-1} c_{2j+1}$ will be part of the interface.
 - If the j th clause is satisfied by one of the first i variables, then the triangles $z_T^i c_{2j-1} c_{2j}$ and $z_T^i c_{2j} c_{2j+1}$ will be part of the interface.

Thus, the number of triangles in the interface will be $m + 1 + j$, where m is the total number of clauses and j is the number of clauses currently satisfied. In particular, after the last step, the interface will be bounded by the vertices $z_T^n, y_3^n, c_1, c_2, c_3, \dots, c_{2m+1}$, which is part of the boundary of the polyhedron (i.e., the entire polyhedron will have been triangulated). Writing $z_T^0 = z_F^1$ and $y_3^0 = y_1^1$, the initial interface is bounded by $z_F^1, y_1^1, c_1, c_3, c_5, \dots, c_{2m+1}$, which is the opposite portion of the polyhedron's boundary.

The simple structure of the interface makes the triangulation easier. What happens is that some vertex (z_F^i , x_1^i , x_2^i , x_3^i , or z_T^i) will be able to see all of the triangles in the interface. Using this vertex as an apex, and the triangles in the interface as bases, we form new tetrahedra to add to the triangulation, and advance the interface to the new vertex.

There are a number of things yet to be specified about the construction. We must lay out the chain of triangles and the clause niches attached to them. We must also specify the placement of the roof vertices and the variable niches attached to the roofs. First we list the geometric constraints that must be satisfied.

The Construction Constraints. For each variable, the following constraints must be satisfied in the construction. These constraints ensure that the truth-setting and literal vertices can see the appropriate variable and clause niches, or are blocked from seeing them, as necessary. The fifth constraint guarantees that the successive variables' portions of the triangulation will be separated by a nicely behaved interface. Since each constraint applies to all variables, we omit the i superscripts.

- **Variable-niche-filling:** Vertices z_F , z_T must be viewpoints of the variable's niche and they must be the only such vertices (i.e., one of them must be used to triangulate the niche).
- **Clause-niche-filling:** The vertices x_1 , x_2 , x_3 corresponding to literals must be viewpoints of their respective clause niches, but of no other niches.
- **Clause visibility:** The use of the vertex z_T to triangulate the variable niche should not prevent the positive literals' vertices from seeing the clause niches. Thus we require that the tetrahedron $z_T y_1 y_2 y_3$ must not intersect the tetrahedra $x_1 c_{2k-1} c_{2k} c_{2k+1}$ or $x_2 c_{2k-1} c_{2k} c_{2k+1}$ for $1 \leq k \leq m$. Similarly, the tetrahedron $z_F y_1 y_2 y_3$ must not intersect the tetrahedra $x_3 c_{2k-1} c_{2k} c_{2k+1}$ for $1 \leq k \leq m$.
- **Clause blocking:** To guarantee that x_1 cannot triangulate its clause niche when the variable is set false, it must be "blocked," i.e., it must not see all of the base face of the clause niche. Here we can assume only that z_F forms a tetrahedron with the base face $q_1 q_2 q_3$ of the variable niche (not the entire variable niche). We require that the tetrahedron $z_F q_1 q_2 q_3$ intersect every line segment from x_1 or x_2 to any point in the base triangles $c_{2k-1} c_{2k} c_{2k+1}$ for $1 \leq k \leq m$. Similarly, the tetrahedron $z_T q_1 q_2 q_3$ must intersect with every line segment from x_3 to any point in the base triangles. (This constraint is slightly stronger than necessary, in that every base triangle is invisible rather than just the bases of the clause niches within the base triangles.)
- **Interface visibility:** For the triangulation to "sweep" across the "slice" of the polyhedron corresponding to a single variable, we require that:
 - (a) x_1 be to the "left" (i.e., negative y -direction) of the planes $c_{2k-1} c_{2k} z_F$, $c_{2k} c_{2k+1} z_F$, and $c_{2k-1} c_{2k+1} z_F$ for $1 \leq k \leq m$.
 - (b) x_2 be to the left of the planes $c_{2k-1} c_{2k} x_1$, $c_{2k} c_{2k+1} x_1$, and $c_{2k-1} c_{2k+1} x_1$ for $1 \leq k \leq m$.
 - (c) x_3 be to the left of the planes $c_{2k-1} c_{2k} z_F$, $c_{2k} c_{2k+1} z_F$, and $c_{2k-1} c_{2k+1} z_F$ for $1 \leq k \leq m$.

- (d) z_T be to the left of the planes $c_{2k-1}c_{2k}x_2$, $c_{2k}c_{2k+1}x_2$, $c_{2k-1}c_{2k+1}x_2$, $c_{2k-1}c_{2k}x_3$, $c_{2k}c_{2k+1}x_3$, and $c_{2k-1}c_{2k+1}x_3$ for $1 \leq k \leq m$.
- **Roof convexity:** The roof must be convex, i.e., the convex hull of the points $z_F, z_A, z_T, x_1, x_2, x_3, y_1, y_2, y_3$ must include the faces shown in Fig. 16.

We postpone the remaining details of the construction of the polyhedron P . Instead, we first show that if P satisfies these constraints, then a satisfying assignment for the expression E can be converted into a triangulation of P .

The Triangulation. Assuming that the six construction constraints can be met in the refined construction of the polyhedron P , we now show that given the original Boolean formula E , we can construct a polyhedron P such that

$$E \text{ is satisfiable} \iff P \text{ can be triangulated.}$$

One direction is fairly easy: given a triangulation of P , we can interpret a satisfying assignment for E just as we did in the rough construction of the “box-shaped” polyhedron above, by seeing which of the truth-setting vertices (z_F^i or z_T^i) was used to triangulate the niche of the i th variable.

Showing that a satisfying assignment for E yields a triangulation of P requires listing quite a few tetrahedra, and relies heavily upon the six constraints developed above. The triangulation proceeds one variable at a time. We need only describe a single step, relying on the structure of the *interface* described above as an invariant that ensures that we can proceed.

We now describe the triangulation step for the i th variable X_i . There are two cases, depending on whether X_i is set TRUE or FALSE in the satisfying assignment. We handle the TRUE case first. Since we are dealing with the single variable X_i , we drop the i subscripts and superscripts.

Case 1: Variable X set TRUE. Referring to Fig. 17 we see that the “hub” of the interface starts at z_F (the invariant). Intuitively, the hub will “advance” to x_1 , then to x_2 , and finally to z_T , establishing the invariant for the next step (see Figs. 18–21).

Because the variable-niche-filling constraint was satisfied in the construction, z_T is a viewpoint of the variable X ’s niche. This means that z_T can see the base




Fig. 17. Shaded triangles to form tetrahedra with z_T .




Fig. 18. Shaded triangles to form tetrahedra with x_1 .




Fig. 19. Shaded triangles to form tetrahedra with x_2 .




Fig. 20. Shaded triangles to form tetrahedra with z_T .

face of the niche. Since z_T can see all of the triangle $y_1y_2y_3$ (the triangle T in the Illuminant Lemma), z_T will also be able to see the other faces of the niche. Thus we can form tetrahedra with the seven faces of the niche from z_T , producing the tetrahedra $z_Tq_3q_2q_1$, $z_Tp_1q_1p_2$, $z_Tq_1q_2p_2$, $z_Tp_2q_2p_3$, $z_Tq_2q_3p_3$, $z_Tp_3q_3p_1$, $z_Tq_3q_1p_1$. Here the p_j and q_j are as specified in Fig. 9. In the application of the Illuminant Lemma that created the niche, the back face $y_1y_2y_3$ of the variable's roof got a triangular "hole" $p_1p_2p_3$, and was retriangulated (two-dimensional triangulation), producing six new faces that must also form tetrahedra with z_T . From z_T we also form tetrahedra with the faces $z_Ax_1z_F$, $z_Ax_2x_1$, $z_Ay_2x_2$, $z_Ax_3y_2$, $x_3y_3y_2$, $x_2y_2y_1$ of the variable's roof (see Fig. 17). The vertex z_T can see all of these triangles because of the roof convexity constraint. Figure 18 shows the situation after removing these tetrahedra.

The next step is to form tetrahedra between x_1 and the interface triangles. This is possible by part (a) of the interface visibility constraint. We also use the tetrahedron $x_1z_Tc_{2m+1}z_F$. The shaded portion of Fig. 18 shows the triangles that must form tetrahedra with x_1 . Since x_1 is a literal vertex corresponding to a literal set TRUE in the satisfying assignment, we may use x_1 to triangulate its corresponding clause niche. We do this only if the clause niche has not been previously triangulated. The clause-niche-filling constraint ensures that x_1 can see all of its




Fig. 21. Final interface after step for variable X .

clause niche, and all of the base triangle containing the clause niche. Thus we use x_1 to form tetrahedra with the seven triangular faces of the clause niche and the six faces of the original base triangle (which was retriangulated when the niche was attached, causing a hole).

Now the hub of the interface moves from x_1 to x_2 . The shaded portion of Fig. 19 shows the triangles that must form tetrahedra with x_2 . Part (b) of the interface visibility constraint guarantees that these triangles can all be seen by x_2 . Then x_2 is used to triangulate its clause niche if necessary.

The situation is now as shown in Fig. 20. This figure shows the triangles that must form tetrahedra with z_T . Note that this includes the rear triangle $y_1y_3c_1$. These triangles are visible by part (d) of the interface visibility constraint. After these tetrahedra are added to the triangulation, the hub of the interface is now at the vertex z_T , as shown in Fig. 21. This re-establishes the invariant for the next variable's step of the triangulation.

Case 2: Variable X set FALSE. This case is analogous to Case 1. The variable X is set FALSE, so we use z_F to triangulate the variable niche and the back face of the roof as well as the back face of the entire slice, and also the faces of the roof shown in Fig. 22. Figure 23 shows the result. Next, x_3 is used to triangulate the faces shown in Fig. 23, as well as its clause niche (if not already triangulated), and the remainder of the base triangle containing the clause niche. This advances the hub of the interface from z_F to x_3 , leaving the situation shown in Fig. 24. Then z_T can be used to form tetrahedra with the triangles shown shaded in Fig. 24, to complete the triangulation of X 's portion of the polyhedron. Just as in Case 1, the hub of the interface has reached z_T , shown in Fig. 21.

Thus we have shown how the triangulation can “sweep” across the polyhedron P , and thus that

$$E \text{ is satisfiable} \iff P \text{ can be triangulated}$$

as long as P satisfies the construction constraints.

Satisfying the Construction Constraints. We now fill in the details of the construction of the polyhedron P , and show how the required constraints are satisfied. As




Fig. 22. Shaded triangles to form tetrahedra with z_F .




Fig. 23. Shaded triangles to form tetrahedra with x_3 .




Fig. 24. Shaded triangles to form tetrahedra with z_T .




Fig. 25. The cone C defined by the two similar triangles T_1 and T_2 .

usual, we omit the i superscripts when referring to the vertices of the i th variable's roof.

To satisfy the variable-niche-filling constraint, we use the Illuminant Lemma to construct the variable niche on the back face $y_1y_2y_3$ of the roof. The shaded region in Fig. 25 shows part of the triangular cone C . We want z_T and z_F to be the only vertices in the interior of C . We specify C by giving two similar triangles T_1 and T_2 , the intersections of C with the parallel planes $x = 0$ and $x = 1$, respectively. The corresponding edges of the two triangles are also parallel. Suppose the roof is of height C_1 , to be specified later. That is, the z -coordinate of z_A and y_2 is $1 + C_1$.

Figure 26 shows how we choose the triangle T_1 within the back face $y_1y_2y_3$ of the roof. Consider the triangle T formed by the midpoints of the edges y_1y_2 , y_2y_3 , and y_3y_1 . Let T_1 be this triangle, shrunk by a factor of C_2 (to be specified later) around its centerpoint.

Figure 27 shows how we choose the triangle T_2 within the $x = 1$ plane. T_2 will be similar to T_1 chosen above, with corresponding edges being parallel. The top edge t_1t_2 of T_2 is slightly above the $z = 1$ plane. This height is chosen so that the plane containing t_1t_2 and the corresponding edge of T_1 intersects the segment y_2z_F exactly three-quarters of the way from t_2 to z_F (see Fig. 27). This ensures that the




Fig. 26. In the $x = 0$ plane: the back gable of the roof (top), and the choice of T_1 within the back gable.




Fig. 27. Specification of T_2 . The cone C intersects segment y_2z_F exactly three-quarters of the way from y_2 to z_F .




Fig. 28. The shape of the clause niche illuminants. The clause niches themselves are very small, along the bottom.

literal vertices, placed later, will not be within the illuminant of the variable niche. We place t_1 at distance one-quarter to the left of z_T and t_2 at distance one-quarter to the right of z_F . This determines the position of t_3 , the third vertex of T_2 . Together, the triangles T_1 and T_2 determine a cone C . The only vertices in the interior of C are z_T and z_F .

Now the Illuminant Lemma can be applied. This constructs a niche on the back face of the roof, with the two vertex viewpoints z_T and z_F .

To satisfy the clause-niche-filling constraint, we again use the Illuminant Lemma, this time to place a clause niche within each of the triangles $c_{2k-1}c_{2k}c_{2k+1}$ along the base of the polyhedron. These niches will have very skinny illuminants (see Fig. 28) that do not intersect within the polyhedron. These will be regions in which we will later be able to place the literal vertices when we construct the roofs. For use in later constraints, we add the further requirement that the portions of these niches and illuminants that lie within the polyhedron must be between the planes $x = \frac{1}{2}$ and $x = 51/100$. To do this we must now specify the parabola containing the base vertices $c_1, c_2, \dots, c_{2m+1}$.

The parabola will be in the xy -plane and will be of the form $y = \alpha x^2 + \beta x + \gamma$. It must pass through the point $(\frac{1}{2}, 0)$ with slope 0. Thus we need $0 = \alpha/4 + \beta/2 + \gamma$ and $\alpha + \beta = 0$. The slope of the parabola should be positive and increasing for $x > \frac{1}{2}$. It will also be important later that the slope be small, say less than $\frac{1}{20}$ at $x = 51/100$. This requires $51\alpha/50 + \beta < \frac{1}{20}$. It can be checked that the values $\alpha = \frac{1}{4}$, $\beta = -1/4$, $\gamma = \frac{1}{16}$ satisfy these requirements. The points $c_1, c_2, \dots, c_{2m+1}$ are spaced evenly between $x = \frac{1}{2}$ and $x = 51/100$ on this parabola. Point c_k is placed at (a_k, b_k) , where

$$a_k = \frac{1}{2} + \frac{k-1}{200m} \quad \text{and} \quad b_k = \frac{1}{4}a_k^2 + \frac{-1}{4}a_k + \frac{1}{16}.$$

To construct the clause niches we must specify a cone C for each clause niche, and then apply the Illuminant Lemma. We do not give the exact specifications of these cones, just a rough idea. The triangle T needed for the Illuminant Lemma is as shown in Fig. 29. It is required that the illuminant of the k th niche contains




Fig. 29. Clause niches are attached to base triangles.

the line segment between the points $(a_{2k}, 0, 1)$ and $(a_{2k}, n, 1)$, where a_{2k} is as specified in the above paragraph. We can choose a second triangle containing this segment, which together with T determines the cone C needed to apply the Illuminant Lemma. These cones should be chosen so that they do not intersect within the polyhedron.

Constructing the i th Variable's Roof. Next we show how to construct the roof for the variable X_i so that the remaining constraints are satisfied. We pick a height for the roof (and thus place the vertices z_A and y_2), and then specify the locations of x_1 , x_2 , and x_3 . This requires choosing the two values C_1 and C_3 mentioned above. The shape of the roofs depend upon the variable index i , so we choose different values C_1^i and C_3^i for each roof. Since we will be referring to the i th variable throughout, we omit the i superscripts from the vertices.

If the roof is too tall, then some of x_1 , x_2 , and x_3 might not be able to see "under" the edge $z_T y_3$ to the clause niches. To avoid this problem we specify a point $p = (0, \frac{1}{2}, \frac{3}{4})$ and require that all vertices of the i th roof lie beneath the plane containing z_T , y_3 , and p (see Fig. 30). This is sufficient since all of the base triangles lie beneath this plane. Choosing $C_1^i = 1/200(n - i + 1)$ keeps z_A and y_2 well below this plane, and builds in some leeway we will utilize later.

Next we describe the placement of the roof vertices x_1 , x_2 , x_3 . Vertices x_1 and x_2 are placed slightly above the segment $y_2 z_P$, in the vertical plane containing the




Fig. 30. Front view of polyhedron, showing plane (dotted) bounding variable X_i 's roof (not to scale).




Fig. 31. Placement of x_1, x_2, x_3 (for i th variable, i superscripts omitted).

segment (see Fig. 31). Suppose the literal vertex x_1 corresponds to a literal contained in the k th clause. Then x_1 is placed at x -coordinate a_{2k} (the a_{2k} 's were specified above, and are between $\frac{1}{2}$ and $51/100$ for all $1 \leq k \leq m$), at a distance C_3^i above the segment $y_2 z_F$. The vertex x_2 is placed the same distance above this segment at the appropriate x -coordinate, and x_3 is placed similarly above the segment $y_2 z_T$.

We now show that choosing $C_3^i = 1/1000(n - i + 1)$ will place the literal vertices such that all the remaining constraints are satisfied.

Since the literal vertices x_1, x_2, x_3 will have z -coordinate $\leq 1 + C_1^i/2 + C_3^i$, they will be below the segment $z_A y_2$, which has z -coordinate $1 + C_1^i$ ($C_3^i < C_1^i/2$). Thus, x_1, x_2, x_3 will not be able to see "over" the top of the roof $z_A y_2$, and the faces of the convex hull of the roof vertices will be those shown in Fig. 16, satisfying the roof convexity constraint.

The literal vertices x_1, x_2, x_3 have been placed above the illuminant of their variable's niche, so they will not be viewpoints of the niche, and hence will not violate the variable-niche-filling constraint.

The clause-niche-filling constraint was satisfied by the choice of x -coordinate for the literal vertices, which placed them within their respective clause niches (but within no other clause niches, since the clause niches do not intersect).

Showing that the clause visibility constraint is satisfied is a little trickier. The x_3 case is easy: since it is to the "left" of the tetrahedron $z_F y_1 y_2 y_3$, the tetrahedron cannot prevent it from seeing the base triangles, which are even farther to the left. For x_1, x_2 we must show that they can see "beneath" the tetrahedron $z_T y_1 y_2 y_3$ to the base triangles. The critical thing is that they be able to see beneath the edge $z_T y_1$, and the worst-case placement of x_1 or x_2 is at coordinates $(\frac{1}{2}, n - i + \frac{3}{4}, 1 + C_1^i/2 + C_3^i)$ (the point q in Fig. 31). The most difficult point on the base triangles to see is c_1 . Thus we need to show that q is beneath the plane $c_1 z_T y_1$.

This orientation test can be done (see p. 43 of [11]) by computing the sign of the determinant

$$\Delta(c_1, z_T, y_1, q) = \begin{vmatrix} 1 & \frac{1}{2} & 0 & 0 \\ 1 & 1 & n-i & 1 \\ 1 & 0 & n-i+1 & 1 \\ 1 & \frac{1}{2} & n-i+\frac{3}{4} & 1 + \frac{C_1^i}{2} + C_3^i \end{vmatrix},$$

which evaluates to $-(986(n-i+1) + 7)/4000(n-i+1)$, which is negative since $i \leq n$. Since the point z_F clearly lies beneath the plane $c_1 z_T y_1$, and the determinant

$$\Delta(c_1, z_T, y_1, z_F) = -1$$

is also negative, the point q must also be beneath the plane $c_1 z_T y_1$.

To show that the clause blocking constraint is satisfied, we show how to choose the “shrink factor” C_2 , used in constructing the variable niches, so that x_1, x_2, x_3 are blocked in the appropriate way. We show this only for the x_1, x_2 cases, as the x_3 case is similar. We continue to omit the i superscript from vertices of the i th variable’s roof. Let R be the smallest rectangle in the $z = 0$ plane with sides parallel to the x and y axes that contains all of the base triangles. Let s_1 be the segment on which x_1 and x_2 were placed (see Fig. 31). We show that choosing $C_2 = 99/100$ will force every line of sight from s_1 to R to pass through the interior of the tetrahedron $z_F q_1 q_2 q_3$ (q_1, q_2, q_3 is the base triangle of variable i ’s niche). This means that if the variable is set FALSE, then x_1 and x_2 will not be able to see any of the clause niches at all, and will not be able to triangulate them.

We choose a point w in the $x = 0$ plane such that the segment $z_F w$ is contained in the tetrahedron $z_F q_1 q_2 q_3$ (see Fig. 32). Then we show that every line of sight from the segment s_1 to the rectangle R must pass beneath the segment $z_F w$. Since these lines of sight also pass above the segment $z_F q_1$, they must “pierce” the triangle $z_F q_1 w$, and hence must pass through the interior of the tetrahedron $z_F q_1 q_2 q_3$, which is what we need to show to establish the clause blocking constraint.




Fig. 32. Every line of sight from the segment s_1 to R must pass “beneath” the segment t .




Fig. 33. The slope of v_1v_2 must be greater than the slope of v_1v_3 .

If we choose the “shrink factor” C_2 to be $99/100$, then the upper left corner of the triangle T_1 shown in Fig. 26 will move in slightly toward the center of T . If the plane containing the points q_1, q_2, q_3 used for the Illuminant Lemma (see Fig. 7) is specified to be at distance $1/100$ from the back face $y_1y_2y_3$ of the roof, then the point

$$w = \left(0, n - i + \frac{1}{4} + \frac{1}{100}, 1 + \frac{C_1^i}{3} + \frac{96}{100} \frac{C_1^i}{6} \right)$$

is within the tetrahedron $z_Fq_1q_2q_3$. (The point w is four times closer to the center of T than is the upper left corner of the triangle T_1 of Fig. 26.) Let t be the subsegment of z_Fw between $x = \frac{1}{2}$ and $x = 51/100$ (i.e., the same x -range as s_1 and R , see Fig. 32). We ignore the x -dimension, by projecting to the $x = 0$ plane (see Fig. 33).

The most difficult line of sight from s_1 to R (passing beneath t) is from the highest, leftmost point of s_1 , under the lowest, rightmost point of t , to the leftmost edge of R (“leftmost” means minimum y -coordinate, “highest” means maximum z -coordinate). Letting v_1 be a point on the leftmost edge of R , letting v_2 be the lowest, rightmost point of t , and letting v_3 be the leftmost, highest point of s_1 , we have the situation shown in Fig. 33. Since the point v_2 is $49/100$ of the distance from z_F to w , and the point v_3 is (a distance C_3^i above the point that is) half of the way from z_F to y_2 , the (y, z) -coordinates of the points are

$$\begin{aligned} v_1 &= (0, 0), \\ v_2 &= \left(n - i + 1 - \frac{49}{100} \times \frac{74}{100}, 1 + \frac{49}{100} \left(\frac{C_1^i}{3} + \frac{96C_1^i}{600} \right) \right), \\ v_3 &= \left(n - i + 1 - \frac{1}{4}, 1 + \frac{C_1^i}{2} + C_3^i \right). \end{aligned}$$

Showing that the line of sight from v_3 to v_1 passes beneath v_2 is a two-dimensional orientation test, so we compute the determinant

$$\Delta(v_1, v_2, v_3) = \begin{vmatrix} 1 & 0 & 0 \\ 1 & n - i + 1 - \frac{49}{100} \times \frac{74}{100} & 1 + \frac{49}{100} \left(\frac{C_1^i}{3} + \frac{96C_1^i}{600} \right) \\ 1 & n - i + 1 - \frac{1}{4} & 1 + \frac{C_1^i}{2} + C_3^i \end{vmatrix},$$




Fig. 34. In the $z = 1$ plane: shaded region shows where interface planes will intersect.

which evaluates to $-(824963(n - i + 1) + 7840)/7500000(n - i + 1)$, which is negative since $i \leq n$. Thus, the lines of sight from s_1 to R pass beneath t , and hence through the tetrahedron $z_F q_1 q_2 q_3$. Hence, when the variable X_i is false, the tetrahedron $z_F q_1 q_2 q_3$ blocks x_1 and x_2 from triangulating the base faces of their clause niches. It can similarly be shown that the tetrahedron $z_T q_1 q_2 q_3$ blocks x_3 , so the clause blocking constraint is satisfied.

We only show that part (a) of the interface visibility constraint is satisfied. Part (b) is similar, and parts (c) and (d) are easier. For part (a), we need to show that x_1 is on the “left” side (i.e., lower y -coordinate) of all the triangles in the interface. It is sufficient to show that x_1 is to the left of the planes $z_F c_1 c_2$, $z_F c_{2m} c_{2m+1}$, and $z_F c_1 c_{2m+1}$. Consider the intersection of these planes with the $z = 1$ plane (see Fig. 34). Since the maximum slope of the base parabola was chosen to be $\frac{1}{20}$, and the minimum slope of the parabola is 0, the three planes will intersect the $z = 1$ plane in lines with slopes between 0 and $\frac{1}{20}$ (relative to the xz -plane). The vertical projection of x_1 to the $z = 1$ plane will be to the left of these lines, since x_1 lies on the segment $z_F y_2$, which projects to a segment with slope $\frac{1}{2}$. Since x_1 is above the $z = 1$ plane, and the three planes $z_F c_1 c_2$, $z_F c_{2m} c_{2m+1}$, and $z_F c_1 c_{2m+1}$ increase in y -coordinate as they increase in z -coordinate, x_1 will be to the left of all three.

Concluding the Proof. In order to ascertain NP-completeness of the triangulation decision problem, we need to show that the problem is actually in NP. First, a triangulation can be “guessed” nondeterministically since we can enumerate all possible collections of tetrahedra. We also need a polynomial-time algorithm to verify that a given collection of tetrahedra is indeed a triangulation of a given polyhedron P . This can be done as follows. Compare all pairs of tetrahedra to make sure that they only intersect in a common triangular face, an edge, a vertex, or the empty set. Next, compute the sum of their volumes and check that it equals the volume of P (which can be computed easily without knowing any triangulation). Lastly, check that each tetrahedron T lies within P , by computing the centroid of T , and performing a point-in-polyhedron test. Also, the intersection of T and the boundary of P should be a collection of faces, edges, and vertices of T .

We must also verify that the polyhedron construction we have given can be performed in polynomial time. In particular, the coordinates of all the points

produced are rational and can clearly be expressed using a polynomial number of bits, because every coordinate results from a constant number of arithmetic operations, starting with small integers. Each application of the Illuminant Lemma can certainly be performed in polynomial time. \square

A Final Remark on the Proof. The polyhedron we have constructed has only triangular faces, but some of them are coplanar, and many of the vertices are coplanar. These coplanarities are only incidental to our construction, and the difficulty of triangulating the polyhedron does not depend on the coplanarities. It would be fairly easy to modify the construction so that the vertices in the resulting polyhedron are in a nondegenerate position.

4. Restriction to Star-Shaped Polyhedra

If the polyhedra are restricted to being star-shaped, the problem is still NP-complete. (A polyhedron is *star-shaped* if there exists a point inside the polyhedron that can “see” all of the polyhedron.)

Theorem 2. *It is NP-complete to decide whether a given three-dimensional star-shaped polyhedron can be triangulated without using additional Steiner points.*

Proof (Sketch). We modify the construction of the previous section so that the polyhedron produced is star-shaped and contains a point p that can see all the polyhedron’s faces. The point p is the same as in Fig. 30, $p = (0, \frac{1}{2}, \frac{3}{4})$. To ensure that the point p can see all of the niches’ faces, we must construct the niches so that the intersection of their illuminants includes p . This is shown in Fig. 35, in which all of the illuminants are “tilted” slightly so that they intersect in a




Fig. 35. To produce a star-shaped polyhedron: illuminants of clause niches “tilt” slightly so as to intersect (triple dashed lines).




Fig. 36. Illuminants of variable niches “stretch” so as to include the point p .

needle-shaped region indicated by the dashed lines. This can be achieved by rearranging the literal vertices, and by slight modifications to the clause niche construction. The construction constraints can be maintained since the polyhedron changes only slightly, and the modified niches will not contain any unwanted vertices.

Next we must modify the variable niches so that their illuminants also contain the point p . Figure 36 gives a rough idea of how the illuminants are “stretched,” and Fig. 37 gives more detail. We describe the modification of the niche by modifying the cone C used in the construction of the niche. The cone C is determined by two similar triangles T_1 and T_2 , in the $x = 0$ and $x = 1$ planes, respectively. Referring now to Fig. 37, these triangles are modified as follows. The top edge of T_2 moves down, into the $z = 1$ plane. Next, the bottom vertex of T_2 is moved down and to the left until the cone includes the point p . In doing this, the corresponding sides of T_1 and T_2 must remain parallel, so fix the top edge of T_1 , and allow the bottom vertex of T_1 to decrease its y -coordinate, remaining within the back gable of the roof. If we also constrain the top right vertex of T_2 , then the top left vertex of T_2 must decrease its y -coordinate to allow the




Fig. 37. More details of variable niche stretching (the roof gables are shown dotted).

corresponding edges to remain parallel. Finally, slightly rotate both T_1 and T_2 counterclockwise, staying parallel to the yz -plane, so that T_2 contains the two appropriate truth-setting vertices, but no others. The rotation should be sufficiently small so that the point p remains within the cone C and T_1 remains within the back gable of the roof.

Next we argue that this modification can be done without violating any of the construction constraints, and check that p remains within the illuminants of the variable niches. Recall that being in the cone C does not ensure that p is in the niche's illuminant. The point p must also be able to see all of the triangle T_1 . This is guaranteed because the entire roof was placed beneath the plane $z_T y_3 p$ (see Fig. 30, which contains the same point p).

In general, the construction constraints are maintained because the shape of the polyhedron P changes only slightly. The variable-niche-filling constraint is met because the triangle T_2 was rotated to include only the desired truth-setting vertices. None of the vertices of the base triangles along the bottom ($z = 0$) of the polyhedron will fall in the variable niche illuminants, because the lowest point in these illuminants will be above $z = \frac{1}{4}$. The clause visibility constraint is unaffected, since, for each variable i , either z_F^i or z_T^i will still triangulate the entire back gable of the roof, exactly as before. The clause blocking constraint depended only upon the positions of z_F^i , z_T^i , and the upper edge of the triangle T_1 . The vertices z_F^i and z_T^i do not move, and the movement of the edge of T_1 can be made as small as necessary. The other constraints can easily be seen to hold.

We have modified the niches so that the point p can see all of the niches' faces. It remains to show that p can see all of the remaining faces of the polyhedron. The only worrisome faces are the faces on the variables' roofs, but p can see these faces because it is beneath all of their planes as constructed (see Fig. 30). \square

Since a star-shaped polyhedron contains a point p that sees every face, it can be triangulated by allowing this additional point p to be used as a vertex of the tetrahedra in the triangulation. (Simply triangulate all faces (polygon triangulation), and use each of the resulting triangles as the base of a tetrahedron with the point p as the apex.) Such additional points are called *Steiner points*.

The previous theorem has consequences for some triangulation problems that allow Steiner points to be used. If we take an n -vertex star-shaped polyhedron P constructed in the proof of Theorem 2 and attach k niches whose illuminants do not intersect the illuminant of any other niche and do not contain any vertex of P , we obtain a polyhedron P' with $n + 6k$ vertices, that can be triangulated using k Steiner points iff P could be triangulated without Steiner points. This yields the following theorems.

Theorem 3. *For any fixed integer $k > 0$ it is NP-hard to determine whether a given polyhedron can be triangulated with at most k Steiner points.*

Theorem 4. *There exists a real constant $C > \frac{1}{4}$ so that, for all positive $c < C$, it is NP-hard to decide whether an n -vertex polyhedron can be triangulated with at most cn Steiner points.*

In his book [9, p. 255] O'Rourke describes n -vertex three-dimensional polyhedra that require $\Omega(n^{3/2})$ "guards." These polyhedra also require $\Omega(n^{3/2})$ Steiner points in order to be triangulated. By attaching such a polyhedron to the polyhedron we have constructed above, the following stronger theorem can be proved.

Theorem 5. *There exists a real constant $C > 0$ so that, for all positive $c < C$, it is NP-hard to decide whether an n -vertex polyhedron can be triangulated with at most $cn^{3/2}$ Steiner points.*

Acknowledgments

Several of the figures in this paper were produced using the UNIGRAPH geometric modeling system developed at U.C. Berkeley. The authors would like to thank Antony Ng, who suggested that the three-dimensional triangulation decision problem might be NP-complete.

References

1. D. Avis and H. ElGindy. Triangulating simplicial point sets in space. *Discrete Comput. Geom.*, 2:99–111, 1987.
2. C. L. Bajaj and T. K. Dey. Convex Decompositions of Simple Polyhedra. Technical Report CSD-TR-833, Dept. Comput. Sci., Purdue University, 1989.
3. B. Chazelle. Computational Geometry and Convexity. Ph.D. thesis, Yale University, 1980. Also available as Technical Report CMU-CS-80-150, Dept. Comput. Sci., Carnegie Mellon University, July 1980.
4. B. Chazelle. Triangulating a simple polygon in linear time. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, IEEE, 1990, pages 220–230.
5. B. Chazelle and L. Palios. Triangulating a nonconvex polytope. *Discrete Comput. Geom.*, 5:505–526, 1990.
6. J. Culberson and R. Reckhow. Covering polygons is hard. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, IEEE, 1988, pages 601–611.
7. H. Edelsbrunner, F. P. Preparata, and D. B. West. Tetrahedrizing point sets in three dimensions. Technical Report UIUCDCS-R-86-1310, Dept. Comput. Sci., University of Illinois, 1986.
8. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.
9. J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, New York, 1987.
10. J. O'Rourke and K. Supowit. Some NP-hard polygon decomposition problems. *IEEE Trans. Inform. Theory*, 29:181–190, 1983.
11. F. P. Preparata and M. I. Shamos. *Computational Geometry—An Introduction*. Springer-Verlag, New York, 1985.
12. E. Schönhardt. Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ann.*, 98:309–312, 1928.
13. B. Von Hohenbalken. Finding simplicial subdivisions of polytopes. *Math. Programming*, 21:233–234, 1981.

Received March 19, 1990, and in revised form January 25, 1991.