Dihedral Bounds for Mesh Generation in High Dimensions

Marshall Bern* Paul Chew† David Eppstein‡ Jim Ruppert§

Abstract

We show that any set of n points in $\mathbb{R}^d$ has a Steiner Delaunay triangulation with $O(n^{\lceil d/2 \rceil})$ simplices, none of which has an obtuse dihedral angle. This result improves a naive bound of $O(n^d)$. No bound depending only on n is possible if we require the maximum dihedral angle to measure at most $90° - \epsilon$ or the minimum dihedral to measure at least ϵ.

1 Introduction

A *mesh* is a partition of a geometric domain into small simple cells (such as boxes or simplices) called *elements*. Automatically generating high quality meshes is a crucial preprocessing step in most numerical methods for physical simulation.

A typical domain for automatic mesh generation consists of a complicated polyhedral region of $\mathbb{R}^d$. In the current state of the art, arbitrary polygonal regions in $\mathbb{R}^2$ can be handled fairly easily, but higher-dimensional mesh generation remains quite difficult [5]. As a simpler version of the higher-dimensional problem, here we consider the problem of computing a simplicial mesh for a point set in $\mathbb{R}^d$. The given point set appears as a subset of the mesh's vertex set; additional vertices, called *Steiner points*, may be added in order to form better quality elements. The mesh covers a convex solid, namely the convex hull of all input and Steiner points. The *complexity* of a mesh is measured by its number of simplices. For efficiency we naturally wish to minimize the complexity of a mesh.

Mathematical folklore has long held that the best meshes are those with elements of low aspect ratio, or "round" elements. An exception occurs in the area of fluid dynamics; long thin elements have proved very efficient in discretizing laminar flows, provided that these elements are aligned with the simulated flow. Babuška and Aziz [1] justified this practice by proving that bounding angles away from 180° (but not away from 0°) in a two-dimensional triangular mesh suffices to guarantee convergence of the finite element method. A bound of 90° on the maximum angle in a triangular mesh also has special importance in the literature. This tradition is justified by both numerical and geometric observations [2], especially in the case of problems with physical characteristics (for example, thermal conductivity) that vary greatly over the domain [17].

The most convenient and appropriate generalization of these angle considerations to higher-dimensional simplicial meshes involves dihedral angles, that is, angles between adjacent facets ($(d-1)$-dimensional faces). Vavasis [17] in fact provides theoretical justification for a dihedral bound of 90° in arbitrary dimension. In this paper we look at all three important cases: dihedral angles bounded away from 0°, bounded away from 180°, and bounded at 90°.

1.1 Previous Work. Most of the relevant previous work has concentrated on triangulations in two dimensions. Assuming that we allow point holes, polygonal input generalizes the case of point sets. For polygonal input, there are algorithms based on grids and quadtrees [2, 6] and others based on Delaunay triangulation [9, 16] that bound angles away from 0°. Two of these algorithms [6, 16] have complexity guarantees as well: the complexity is within a constant factor of the minimum possible for triangulations with bounded minimum angle. There are also algorithms [2, 13] that simultaneously bound the minimum angle away from 0° and the maximum at 90°. For point set input only, there is an algorithm [6] that bounds all angles between 36° and 80°.

The complexity of any of these triangulations, however, depends not only on n—the number of input vertices—but also on the geometry of the input. This dependence is indeed necessary to bound the mini-

*Xerox PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304.

†Department of Computer Science, Cornell University, Ithaca, NY 14853. Work supported by AFOSR Grant AFOSR-91-0328, by ONR Grant N00014-89-J-1946, and by ARPA under ONR contract N00014-88-K-0591.

‡Department of Information and Computer Science, University of California, Irvine, CA 92717. Work done in part while visiting Xerox PARC, and supported in part by NSF grant CCR-9258355.

§IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120. Work performed while the author was at the NASA Ames Research Center as an employee of Computer Sciences Corporation, under NASA contract NAS 2-12961.

mum angle away from 0° (by the complexity guarantee mentioned above). This dependence, however, is not necessary to bound only the maximum angle; and—taking the viewpoint traditional to theoretical computer science—it is fundamental to find triangulations of complexity polynomial in n whenever possible. For point set input, $O(n)$-complexity triangulations are possible with all angles acute (some by an amount dependent upon the geometry) [6]. Since any bound below $90° - \epsilon$ would lead to a bound of $\epsilon/2$ on the minimum angle, no further improvement is possible. For polygonal input, Bern and Eppstein [4] gave an $O(n^2)$-complexity triangulation that bounds the maximum angle at 90°. This was improved very recently to linear complexity by Bern, Mitchell, and Ruppert [7]. For a description of related work see the survey by Bern and Eppstein [5].

In higher dimensions, there arise new types of angles, between faces of different dimensions. Mitchell and Vavasis [14] generalized the quadtree method [6] to three dimensional polyhedra, in order to bound solid angles at vertices away from zero, which implies bounds on all other types of angles. As in two dimensions, this bound requires nonpolynomial complexity. Do polynomial bounds exist for less stringent angle requirements? In this paper, we answer this question for point set inputs by characterizing exactly which no-bad-angle problems can be polynomially solved. Our lower bounds also hold for the more difficult case of polyhedral inputs.

1.2 New Results. The first listing below formulates a family of no-bad-angle problems for simplicial meshes in d dimensions, thereby generalizing the two-dimensional minimum and maximum angle problems. The first three listings justify our focus on dihedrals.

1. Each k-face in a d-simplex defines an angle, and all angles that involve the entire simplex are so defined. To measure an angle, we place a small sphere around the k-face and determine what fraction of the sphere lies interior to the simplex. The sphere is $(d-k-1)$-dimensional and lies in a $(d-k)$-flat perpendicular to the k-face; its center is the projection of the face onto the flat. If $k = d-2$, we call the angle a *dihedral angle* and we can write its measure in degrees in the usual way. If $k = 0$, the angle is a *solid angle*.

   We can now define a family of problems: for each k, we can demand either no small angles, that is, all angles must be bounded away from zero, or no large angles—all angles must be bounded away from flat (half the sphere). We show that this family forms two sequences of difficulty. No small angle at a 0-face (no small solid angle) implies no small angle at a 1-face, which implies no small angle at a 2-face, and so forth up to no small angle at a $(d-2)$-face. No-large-angle problems reverse the order: no large dihedral implies no large angle at a $(d-3)$-face, which implies no large angle at a $(d-4)$-face, and so forth down to no large solid angle. There are two bridges between the sequences. No small solid angle implies no large dihedral, and hence good angles of all types. In the opposite direction, a large solid angle implies a small dihedral, and hence bad angles of all types.

2. If all dihedrals are required to be larger than some fixed $\epsilon > 0$, then for some point sets a simplicial mesh must have $\Omega(n \log A)$ complexity, where A is a parameter that depends on geometry. Thus no strongly polynomial bound (that is, polynomial in n alone) is possible for the no-small-dihedral problem or any other no-small-angle problem. On the other hand, a method based on quadtrees achieves $O(n \log A)$ complexity for the hardest problem, no small solid angles.

3. Quadtree triangulation gives $O(n)$-complexity meshes with all dihedrals smaller than $180° - \epsilon$, where ϵ is a constant depending only on the dimension. Thus the no-large-dihedral problem and hence all no-large-angle problems admit polynomial solutions.

4. Any point set in d dimensions can be triangulated with $O(n^{\lceil d/2 \rceil})$ simplices with no obtuse dihedrals. Our triangulation uses only self-centered simplices (those containing their circumcenters) and thus is a Delaunay triangulation of its vertices [15]. Moreover each face of each simplex is self-centered, so that the mesh allows barycentric subdivision, the natural generalization of perpendicular planar dual embedding [4].

5. Bounding the dihedrals of a simplex to $90° - \epsilon$ forces it to have bounded aspect ratio. A simplicial mesh with simplices of bounded aspect ratio must have $\Omega(n \log A)$ complexity. Thus result 4 gives the best possible (fixed) angle bound.

We regard statement 4 above as our main result. This result betters a naive bound of $O(n^d)$ simplices, resulting from passing axis-aligned planes through each input point and then triangulating the resulting boxes. Our $O(n^{\lceil d/2 \rceil})$-complexity algorithm uses a product construction to produce an *unstructured mesh*, meaning a mesh in which vertices have nonisomorphic local neighborhoods. It is interesting to note that for this problem an unstructured mesh offers significant—by a square

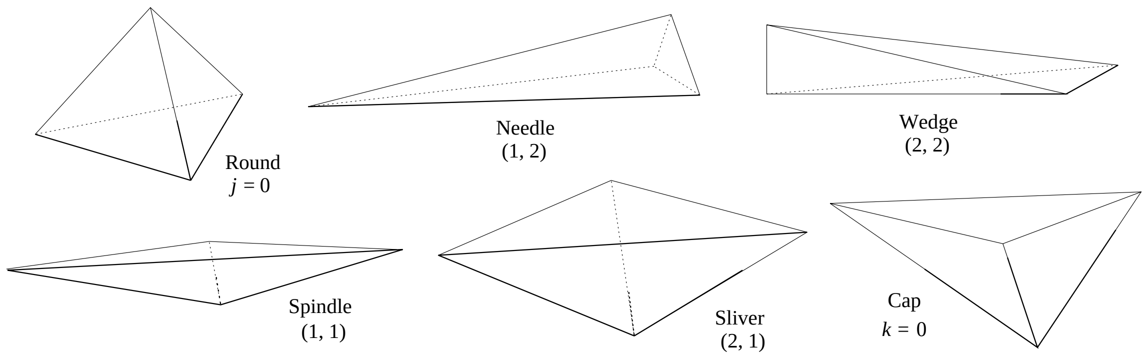

Figure 1: Tetrahedra classified by bad angles.

root—savings in complexity, more than compensating for the usual constant-factor space and time advantages of structured meshes [5]. Finally, we remark that Haiman [12] uses a similar sort of product construction for the problem of triangulating the d-cube with few simplices.

2 Why Dihedrals?

Assume that we have defined angles within a d-simplex S as above, and the measure of a flat angle (one subtending a hemisphere) is $\frac{1}{2}$. Let $NS(k, \epsilon)$ denote the property that no angle at a k-face within S measures less than ϵ. Similarly let $NL(k, \epsilon)$ denote the property that no angle at a k-face within S measures more than $\frac{1}{2} - \epsilon$. The following theorem states that the NS and NL properties are each linearly ordered by implication. We then connect the two sequences.

Theorem 2.1. *There exists $\epsilon' > 0$ depending only on d and ϵ, such that $NS(k, \epsilon)$ implies $NS(k+1, \epsilon')$ and $NL(k+1, \epsilon)$ implies $NL(k, \epsilon')$ for each k between 0 and $d-3$.*

Proof. To prove that $NS(k, \epsilon)$ implies $NS(k+1, \epsilon')$, place a small $(d-k-1)$-sphere around any k-face of a $(k+1)$-simplex within S. Slicing this sphere perpendicular to the $(k+1)$-simplex gives a $(d-k-2)$-sphere that measures the angle at the $(k+1)$-simplex. The fraction of the $(d-k-1)$-sphere covered by S can be no greater than the covered fraction of the $(d-k-2)$-sphere—imagine integrating over many such slices—so $NS(k+1, \epsilon')$ holds with $\epsilon' = \epsilon$.

To prove that $NL(k+1, \epsilon)$ implies $NL(k, \epsilon')$, consider the contrapositive. If an angle at a k-face f covers almost all of a $(d-k-1)$-hemisphere, then its slices nearly cover $(d-k-2)$-hemispheres, so the angles at all the $(k+1)$-faces incident to f must be large. □

Theorem 2.2. *There exists $\epsilon' > 0$ depending only on d and ϵ, such that $NS(0, \epsilon)$ implies $NL(d-2, \epsilon')$ and $NS(d-2, \epsilon)$ implies $NL(0, \epsilon')$.*

Proof. To prove the first part, assume for contradiction that simplex S has a large dihedral angle at a $(d-2)$-face f. Then the edge opposite f subtends almost a semicircle, so the solid angles at the endpoints of this edge are arbitrarily small in at least one direction.

To prove the second part, assume that S has a large solid angle at vertex v. Then the perpendicular h from v to the opposite facet f is very short and lands well interior to f, so all the dihedrals subtended by h (that is, all dihedrals at facets of f) must be small. □

Theorems 2.1 and 2.2 can be used to classify simplices. A (j, k)-bad-angle simplex, $0 \le j, k \le d-1$, allows small angles at faces of dimension up to $j-1$ and large angles at faces of dimension down to k. The case $j = 0$ allows no small angles, and hence no bad angles at all, so k is irrelevant. On the other extreme, $k = 0$ allows a large solid angle and renders j irrelevant. Except for these extreme cases all other combinations of j and k are possible, so there are $2 + (d-1)^2$ different types of simplices.

Figure 1 illustrates this classification for the case $d = 3$. A "needle" allows small solid angles, but not small or large dihedrals; a "wedge" allows small but not large dihedrals, and so forth. Similar but less systematic classifications have appeared in the mesh generation literature [3, 11].

3 No Small Dihedrals

We now consider mesh generation with no small dihedrals, the weakest condition in the NS sequence. Let the *width* of a simplex be the minimum distance between parallel supporting hyperplanes; equivalently it is the minimum distance between nonadjacent faces. We first claim that a lower bound on dihedrals imposes a constraint on the shapes of our simplices.

Lemma 3.1. *Let simplex S have width w and no dihedral angle smaller than some fixed $\epsilon > 0$. Then every facet of S has width $\Theta(w)$.*

Proof. We observe that a narrow-width facet implies that S has the same or even narrower width. Thus every facet of S has width $\Omega(w)$.

Let π_1 and π_2 be the supporting hyperplanes of S with minimum separation w. Any facet forming an angle greater than $\epsilon/2$ with π_1 must have width $O(w)$. Let F be the set of facets of large width, $\Omega(w)$ for a sufficiently large constant. Each member of F is nearly parallel to π_1 and each pair of facets in a simplex meet at a dihedral; hence two members of F must form a dihedral angle either less than ϵ or greater than $180° - \epsilon$. By assumption the former case cannot occur, so pairs of facets from F meet at very flat dihedral angles. Thus in a perpendicular projection of (interiors of) facets of F onto π_1, points of π_1 are covered at most once.

On the other hand, all points of π_1, except for a set of measure zero, are covered an even number of times by projections of facets in S. Thus the small-width facets (those in $S \setminus F$) cover the large-width facets in the projection. But the projection of a facet has width no greater than the width of the facet, so each projection of a large-width facet is covered by projections of width $O(w)$. Since there are only $d + 1$ facets, all facets must have width $O(dw) = O(w)$. □

We define the *aspect ratio* of a simplex to be the ratio between the largest and smallest separations of parallel supporting hyperplanes. Alternate definitions include the ratio between circumradius and inradius, but a bound on one such ratio implies a bound on the other, so the exact definition is immaterial for results such as the one below.

Theorem 3.1. *Triangulation with no dihedral angle smaller than ϵ can require $\Omega(n \log A)$ simplices, where A is the maximum aspect ratio of a simplex in the Delaunay triangulation of the input point set.*

Proof. We first show the result for $n = d + 2$. It can be extended to any n by making many copies of the construction.

Consider the vertices of a unit-volume regular simplex, together with one extra point at small distance δ from a simplex vertex. The Delaunay triangulation of this point set has a simplex of aspect ratio $A = \Theta(1/\delta)$. We wish to show that any Steiner triangulation of this point set that has no small dihedrals must have $\Omega(\log A)$ simplices.

Suppose we have a Steiner triangulation of the point set that has no small dihedral angles. Consider the line segment between the extra point and the nearby vertex. Either it is entirely within some simplex of the triangulation, or it crosses some simplex face disjoint from the extra point. In either case we have found a simplex in our triangulation of width $O(1/A)$. On the other hand, the volume of the intersection of a simplex of width w with the initial, unit-volume, regular simplex is $O(w)$, so if there are a total of s simplices in the triangulation, some simplex must have width $\Omega(1/s)$.

Since our triangulation covers a connected region of $\mathbb{R}^d$, we can find a path of simplices, joined facet to facet, connecting any pair of simplices. By Lemma 3.1, the width of each facet—and hence each simplex—in such a path can grow by only a constant factor at each step. Thus a path from the width-$O(1/A)$ simplex to the width-$\Omega(1/s)$ simplex must pass through $\Omega(\log \frac{A}{s})$ distinct simplices. Hence s is $\Omega(\log A) - O(\log s)$, which implies that s is $\Omega(\log A)$. □

A matching $O(n \log A)$ upper bound can be achieved by using quadtrees to find a triangulation with bounded aspect ratio [6].

4 No Large Dihedrals

In this section, we use methods of Bern, Eppstein, and Gilbert [6] to solve the hardest problem in the *NL* sequence—bounding the largest dihedral away from 180°. We construct a higher-dimensional analog of a quadtree (a 2^d-ary tree) for the point set. To do this, we begin with a d-dimensional root cube containing all the points, and then recursively subdivide each cube until none contains more than one point. We *balance* the quadtree by further subdividing cubes until no cube is orthogonally adjacent to one of side length less than half its own.

If a sequence of cube subdivisions continues for some constant number of steps without separating any points, we can "shortcut" the subdivision process by clustering points as in [6]. Briefly, we collect a subset of points that have small diameter relative to their distance to the rest of the input points. We recursively triangulate this subset and then treat the resulting triangulated quadtree as a unit back in our original quadtree.

Once the quadtree has been sufficiently subdivided, we further subdivide all boxes a constant number of times, and then merge boxes locally so that each point or recursively triangulated cluster is situated near the middle of an unsubdivided cube. We then triangulate each cube of the quadtree in a way consistent with its neighbors. We omit the proof of the following theorem, because the complexity bound is similar to the previous analysis [6] and the dihedral bound follows from a rather uninteresting case analysis.

Theorem 4.1. *The algorithm sketched above gives a Steiner triangulation of linear complexity, in which the maximum dihedral measures at most $180° - \epsilon$, where ϵ is a constant depending only on dimension.*

We do not know the best possible bounds for Theorem 4.1. For $\mathbb{R}^3$, it should be fairly straightforward to achieve maximum dihedral arbitrarily close to $120°$.

5 No Obtuse Dihedrals

For some applications, dihedrals bounded away from $180°$ may suffice. But for others, and in particular in the "wild coefficient" model of Vavasis [17], the dihedrals should be no larger than $90°$.

5.1 Nonobtuse Simplices. We start by describing the building blocks of our nonobtuse triangulations: nonobtuse simplices. Coxeter [10] describes a general class of nonobtuse simplices, formed as follows. Any tree T with $d+1$ vertices and d edges can be embedded in $\mathbb{R}^d$ so that the edge lengths are preserved and any two edges are at right angles to each other: treat distinct edges as vectors of the appropriate lengths parallel to distinct coordinate axes, place one vertex v_0 at the origin, and place each remaining vertex v_i by adding the vectors on the path in the tree from v_i to v_0. The convex hull of these $d+1$ vertices forms a simplex S_T in $\mathbb{R}^d$.

Definition 1. *An orthogonal simplex is a simplex S_T for some tree T.*

Lemma 5.1. (Coxeter [10]) *Dihedral angles in S_T opposite edges of T are acute; all remaining dihedral angles are right.*

Coxeter proves a number of other metric properties of orthogonal simplices. Since orthogonal simplices have a maximum number of right-angled dihedrals, they can be considered an appropriate generalization of right triangles in the plane. One important property of right triangles, however, does not always hold for orthogonal simplices.

Definition 2. *A simplex is self-centered if it contains its circumcenter.*

A triangle is self-centered if and only if it has no obtuse angle. Thus an alternate generalization of nonobtuse triangles would be self-centered simplices. This generalization alone, however, is insufficient for mesh generation, because self-centered simplices—for example the "sliver" in Figure 1—can have dihedrals arbitrarily close to $180°$. An example of an orthogonal simplex that is not self-centered is the convex hull of the three standard unit vectors in $\mathbb{R}^3$.

Self-centered triangulations (those containing only self-centered triangles) boast two useful properties. A self-centered triangulation is necessarily a Delaunay triangulation, and the planar dual of a self-centered triangulation can be embedded such that dual edges cross at right angles. Rajan [15] proved that the first property holds in arbitrary dimension. For higher-dimensional generalizations of the second property, however, we require something stronger.

Definition 3. *A simplex is fully self-centered if each face is self-centered.*

The following generalization of dual embedding holds: any fully self-centered simplex can be triangulated into $(d+1)!$ fully self-centered simplices (actually path simplices, as defined below) by barycentric subdivision. Barycentric triangulations of neighboring simplices fit together to form a refinement of the original triangulation.

We now define a special type of orthogonal simplex that fulfills all our needs. These simplices have no obtuse dihedrals, a maximum number of right dihedrals, and are fully self-centered.

Definition 4. *A path simplex is a simplex S_T for a tree T that is a path.*

Definition 1 above defines an orthogonal simplex by the position of its vertices. For a path simplex, there is also a nice dual description in terms of facets.

Lemma 5.2. *Any path simplex is isometric to a simplex of the form*

$$\{(x_1, x_2, \ldots, x_d) \;\mid\; 0 \le c_1x_1 \le c_2x_2 \le \ldots \le c_dx_d \le 1\}$$

for some set of positive constants $c_1, c_2 \ldots c_d$, and every simplex of this form is a path simplex.

Proof. Let S_T be a path simplex for path T. Rename the coordinate axes so that edge $d+1-i$ in T is parallel to the ith coordinate axis; let c_i be the reciprocal of the length of this edge; and define simplex S' by the inequalities above. It is not hard to see that the vertices of T are the extremal points of S'. Both S_T and S' are the convex hulls of their extremal points, so S_T is isometric to S'.

In the other direction, let S' be defined by the inequalities above and let the extremal points v_i of S' be defined as $v_i = (0, \ldots, 0, 1/c_i, 1/c_{i+1}, \ldots, 1/c_d)$. Then in the path T with vertices $v_0, v_1, \ldots, v_d$, distinct edges are parallel to distinct coordinate axes, and hence $S' = S_T$. $\square$

The next two lemmas establish path simplices as the closest generalization of right triangles.

LEMMA 5.3. *Every path simplex is fully self-centered.*

Proof. Let S_T be a path simplex, characterized as in Lemma 5.2. Note that the vertices of S_T are a subset of the vertices of the box

$$\{ (x_1, x_2, \dots x_d) \mid 0 \le c_i x_i \le 1 \}.$$

The circumcenters of the simplex and the box are the same, namely the point $(\frac{1}{2c_1}, \frac{1}{2c_2}, \dots \frac{1}{2c_d})$ halfway between the first and last vertices, v_0 and v_d, of the path T. The circumcenter lies in S_T, hence S_T itself is self-centered. If a face of S_T includes both v_0 and v_d, its circumcenter will be the same as that of all of S_T, since the circumsphere of S_T is the diameter sphere of $v_0 v_d$. If a face f of S_T does not include vertex v_d, then f is also a face of the path simplex formed by removing v_d from T. Similarly, if a face does not include vertex v_0, then it is also a face of a smaller path simplex. By induction on dimension, these smaller path simplices are fully self-centered, so the given face is self-centered. Thus all faces of S_T are self-centered. □

LEMMA 5.4. *Any orthogonal simplex that is self-centered is a path simplex.*

Proof. For $d \le 2$, all orthogonal simplices are path simplices. For $d = 3$, the only orthogonal simplex that is not a path simplex is a corner of a box; it is not hard to see that such a simplex cannot be self-centered.

Now suppose $d > 3$ and the tree T associated with the orthogonal simplex is not a path. Then T contains a vertex v of degree at least three; let T' be the subtree containing v and three of its neighbors. This subtree generates an orthogonal simplex $S_{T'}$ in $\mathbb{R}^3$ that is not a path simplex, hence not self-centered. We now show that this implies that S_T cannot be self-centered. The circumsphere of S_T intersects the three-dimensional subspace spanned by $S_{T'}$ in a sphere that touches the vertices of $S_{T'}$, in other words, in the circumsphere of $S_{T'}$. So in the perpendicular projection onto the three-dimensional subspace, the circumcenter of S_T projects onto the circumcenter of $S_{T'}$. But S_T itself projects perpendicularly onto $S_{T'}$, so S_T must be disjoint from its circumcenter. □

5.2 Products of Simplices. For any point set in $\mathbb{R}^d$, we can compute a Steiner triangulation with $O(n^d)$ path simplices by forming a grid of n^d boxes and partitioning each box into $d!$ path simplices. Our aim in the remainder of Section 5 is to show that this naive bound can be greatly improved. We reduce the number of simplices to $O(n^{\lceil d/2 \rceil})$, roughly the square root of the naive bound.

Let S_1 be a set in $\mathbb{R}^{d_1}$, and S_2 be a set in $\mathbb{R}^{d_2}$. We define their *product* $S_1 \times S_2$ to be the set

$$\{ (x_1, x_2, \dots, x_{d_1}, y_1, y_2, \dots, y_{d_2}) \mid \\ (x_1, x_2, \dots x_{d_1}) \in S_1 \text{ and } (y_1, y_2, \dots y_{d_2}) \in S_2 \}.$$

Products can also be defined in terms of inequalities. If S_1 is the convex region satisfying a set I_1 of inequalities, and if S_2 is similarly the region satisfying a set I_2 of inequalities, such that no variable appears in both I_1 and I_2, then $S_1 \times S_2$ is the region satisfying $I_1 \cup I_2$.

We now show that products of path simplices are well-behaved, in that they can themselves be decomposed into unions of path simplices. In this subsection and the next, P_1 and P_2 denote path simplices with d_1 and d_2 dimensions respectively. We call the triangulation given by the following lemma the *product triangulation* of $P_1 \times P_2$.

LEMMA 5.5. *Let P_1 and P_2 be path simplices with d_1 and d_2 dimensions respectively. Then $P_1 \times P_2$ can be triangulated with $\binom{d_1+d_2}{d_1}$ path simplices.*

Proof. $P_1 \times P_2$ is a region defined by a set of inequalities of the form $c_i x_i \le c_j x_j$, with all variables also satisfying $0 \le c_i x_i \le 1$. These inequalities give a partial ordering on the variables that can be extended to $\binom{d_1+d_2}{d_1}$ total orderings. By Lemma 5.2, these total orderings yield distinct path simplices that together triangulate $P_1 \times P_2$. □

As a special case of Lemma 5.5, a box in $\mathbb{R}^d$ (which is a product of d one-dimensional simplices) can be triangulated by $d!$ path simplices. Incidentally, every triangulation of the product of two simplices has the same number of simplices [8, 12].

5.3 Oriented Path Simplices. In our final triangulation algorithm we will need to glue together several product triangulations. In order to ensure consistency, we orient simplex edges.

DEFINITION 5. *An orientation of a path simplex is an assignment of directions to edges, in such a way that the resulting edge digraph is acyclic and the path defining the simplex is a directed path.*

Since the edges of a simplex form a complete graph, a path simplex is oriented if and only if its edges are acyclically oriented with the source and sink equal to the two vertices v_0 and v_d on the long diagonal. We say that a triangulation of path simplices is oriented when the edges are given directions so that each simplex is oriented. The next two lemmas show that product triangulation can be extended to oriented path simplices. Finally, Lemma 5.8 will be used by our algorithm to produce an initial oriented path-simplex triangulation.

LEMMA 5.6. *Every simplex of the product triangulation of $P_1 \times P_2$ contains vertices $v_0 \times w_0$ and $v_{d_1} \times w_{d_2}$, where $v_0 v_{d_1}$ is the long diagonal of P_1 and $w_0 w_{d_2}$ is the long diagonal of P_2.*

Proof. Each such simplex is defined by some inequalities $c_i x_i \le c_j x_j$, sufficient to totally order the variables, together with the inequalities $0 \le c_i x_i \le 1$. These inequalities are all satisfied if every $c_i x_i$ is zero, giving point $v_0 \times w_0$, or if every $c_i x_i$ is one, giving point $v_{d_1} \times w_{d_1}$. □

LEMMA 5.7. *If P_1 and P_2 are each oriented with source at the origin, then the simplices in the product triangulation of $P_1 \times P_2$ can all be oriented consistently with each other.*

Proof. By Lemma 5.6, the long diagonal of each such simplex runs from $v_0 \times w_0$ to $v_{d_1} \times w_{d_2}$. Each simplex has a unique acyclic orientation such that the endpoints of the long diagonal are source and sink; moreover, since every undirected graph has some acyclic orientation with a given source and sink, these unique orientations give an acyclic orientation of the graph formed by all edges in the triangulation. □

LEMMA 5.8. *Let Δ be a barycentric subdivision of a fully self-centered triangulation. Then Δ can be oriented.*

Proof. In barycentric subdivision, one vertex is added in the center of each face of the original triangulation. The new vertex inside face f is then connected to the centers of all the lower-dimensional faces bounding f. Directing each edge from the higher-dimensional face to the lower-dimensional face gives an acyclic graph, in which orthogonal paths—those paths that visit centers of faces of each dimension—are consistently oriented. □

5.4 Triangulation of Point Sets. We are now ready to describe our triangulation algorithm.

THEOREM 5.1. *Any point set in $\mathbb{R}^d$ can be triangulated with $O(n^{\lceil d/2 \rceil})$ path simplices.*

Proof. We prove by induction that there is such an oriented triangulation. For $d = 1$ the result is trivial. For $d = 2$ either a quadtree-based algorithm [6] or the recent nonobtuse triangulation algorithm for polygons [7] can be used to produce a linear complexity triangulation with no obtuse triangles. Lemma 5.8 shows that the barycentric subdivision of a nonobtuse triangulation is an oriented triangulation of two-dimensional path simplices (right triangles).

For higher values of d, we project the points perpendicularly to two coordinate axes to produce a point set in $\mathbb{R}^{d-2}$, which by induction has an oriented triangulation Δ with $O(n^{\lceil d/2 \rceil - 1})$ path simplices. If we project the original point set perpendicularly to the remaining axes we get a point set in $\mathbb{R}^2$, which has an oriented triangulation Δ' with linear complexity.

The product $\Delta \times \Delta'$ is a polyhedral subdivision containing $O(n^{\lceil d/2 \rceil})$ products of path simplices and the original input point set as a subset of the vertices. By Lemmas 5.5 and 5.7 we can further triangulate each cell in $\Delta \times \Delta'$ using path simplices, and orient each simplex consistently with the orientations in $\Delta \times \Delta'$. Each face shared by a pair of adjacent cells in $\Delta \times \Delta'$ is subdivided identically by each cell, due to the consistent orientations, so the triangulations of cells can be put together to form a triangulation of the entire product. Since the orientations in Δ and Δ' are consistent, we also have a consistent orientation in $\Delta \times \Delta'$, thus proving the induction hypothesis and the theorem. □

We note that consistent orientations are crucial: as shown in Figure 2, five isosceles right triangles arranged around a common center point form a two-dimensional triangulation that cannot be consistently oriented, and the product of this triangulation by an interval produces a collection of right prisms that cannot be triangulated by path simplices without additional Steiner points.

6 Only Acute Dihedrals

One might ask whether Theorem 5.1 can be strengthened to bound the dihedrals below $90° - \epsilon$ for some $\epsilon > 0$. We now show that if we retain a strongly polynomial complexity bound, such improvements are not possible.

LEMMA 6.1. *There is some constant ϵ' (depending on ϵ and d) such that every simplex in $\mathbb{R}^d$ with no dihedral angle larger than $90° - \epsilon$ has aspect ratio at most $1/\epsilon'$.*

Proof. Let S be a simplex with all dihedrals smaller than $90° - \epsilon$. Then the dihedrals of any facet f of S must be smaller than $90° - \epsilon$. By induction on dimension, facet f has bounded aspect ratio.

Now let us consider a specific f, the facet of S most distant from its opposite vertex v. Vertex v cannot be too distant from f (more than a constant times the diameter of f), because it lies in the simplex bounded by hyperplanes forming dihedral angles of $90° - \epsilon$ with f. And vertex v cannot be too close to f (less than a

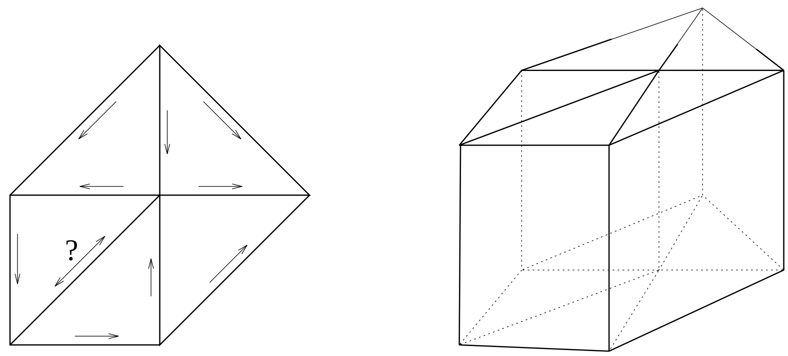

Figure 2: (a) A path-simplex triangulation that cannot be oriented. (b) Its product with an interval.

constant times the diameter of f), or else v would not be most distant from its opposite facet. □

Bern et al. [6] gave a lower bound for bounded aspect ratio triangulation. Combining this lower bound with Lemma 6.1 gives the following result.

Theorem 6.1. *There are point sets such that every triangulation with maximum dihedral at most $90° - \epsilon$ has complexity $\Omega(n \log A)$, where A is the maximum aspect ratio of a simplex in the Delaunay triangulation.* □

Quadtree-based methods can be used to achieve a matching $O(n \log A)$ upper bound on the complexity of a bounded aspect ratio triangulation, although it is unknown whether these triangulations can also achieve dihedrals below $90° - \epsilon$. The analogous problem in two dimensions has a solution: Bern et al. [6] showed how to compute two-dimensional triangulations of complexity $O(n \log A)$ with maximum angle at most 80°.

7 Conclusions and Open Problems

We have extended work on two-dimensional bounded-angle triangulation to point sets in arbitrary dimension. A nontrivial part of this work was simply finding the appropriate generalization of nonobtuse triangulation; path simplices now seem inevitable. An interesting and difficult open problem is to extend bounded-angle triangulation to polyhedra.

There are also many open questions specifically concerning the techniques used in this paper. What are the best dihedral bounds that can be achieved with the linear-complexity quadtree technique? And with the $O(n \log A)$-complexity quadtree technique?

Can we improve our $O(n^{\lceil d/2 \rceil})$ bounds on the complexity of nonobtuse triangulation? If we could achieve an improvement in some dimension using path (or other fully self-centered) simplices, the improvement would carry over to higher dimensions by the product construction of Theorem 5.1.

References

[1] I. Babuška and A. Aziz. On the angle condition in the finite element method, *SIAM J. Numer. Analysis* 13 (1976) 214–227.

[2] B.S. Baker, E. Grosse, and C.S. Rafferty. Nonobtuse triangulation of polygons. *Disc. and Comp. Geom.* 3 (1988) 147–168.

[3] T.J. Baker. Element quality in tetrahedral meshes. *7th Int. Conf. on Finite Element Models in Flow Problems* (1989) Huntsville, Alabama.

[4] M. Bern and D. Eppstein. Polynomial-size nonobtuse triangulation of polygons. *Int. J. Computational Geometry & Applications* 2 (1992) 241–255.

[5] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry*, D.-Z. Du and F.K. Hwang, eds. World Scientific (1992) 23–90.

[6] M. Bern, D. Eppstein, and J. R. Gilbert. Provably good mesh generation. *J. Comp. System Science* 48 (1994) 384–409.

[7] M. Bern, S. Mitchell, and J. Ruppert. Linear-size nonobtuse triangulation of polygons. In *Proc. 10th Symp. Computational Geometry* (1994) 221–230.

[8] L.J. Billera, R. Cushman, and J.A. Sanders. The Stanley decomposition of the harmonic oscillator. *Nederl. Akad. Wetensch. Proc. Ser. A* 91 (1988) 375–393.

[9] L.P. Chew. Guaranteed-Quality Mesh Generation for Curved Surfaces. In *Proc. 9th ACM Symp. Computational Geometry* (1993) 274–280.

[10] H.S.M. Coxeter. Orthogonal trees. In *Proc. 7th ACM Symp. Computational Geometry* (1991) 89–97.

[11] T. Dey, C. Bajaj, and K. Sugihara. On good triangulations in three dimensions. *Int. J. Comp. Geometry and Applications* 2 (1992) 75–95.

[12] M. Haiman. A simple and relatively efficient triangulation of the n-cube. *Disc. and Comp. Geometry* 6 (1991) 287–289.

[13] E. Melissaratos and D. Souvaine. Coping with inconsistencies: a new approach to produce quality triangulations of polygonal domains with holes. In *Proc. 8th ACM Symp. Comp. Geometry* (1992) 202–211.

[14] S.A. Mitchell and S.A. Vavasis. Quality mesh generation in three dimensions. In *Proc. 8th ACM Symp. Comp. Geometry* (1992) 212–221.

[15] V.T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. In *Proc. 7th ACM Symp. Computational Geometry* (1991) 357–363.

[16] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proc. 4th ACM-SIAM Symp. Discrete Algorithms* (1993) 83–92.

[17] S. Vavasis. Stable finite elements for problems with wild coefficients. Tech. Rep. 93-1364, Dept. of Computer Science, Cornell Univ., 1993.