A Condition Guaranteeing the Existence of Higher-Dimensional Constrained Delaunay Triangulations

Jonathan Richard Shewchuk
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
jrs@cs.cmu.edu

Abstract

Let X be a complex of vertices and piecewise linear constraining facets embedded in E^d. Say that a simplex is *strongly Delaunay* if its vertices are in X and there exists a sphere that passes through its vertices but passes through and encloses no other vertex. Then X has a d-dimensional constrained Delaunay triangulation if each k-dimensional constraining facet in X with $k \leq d - 2$ is a union of strongly Delaunay k-simplices.

This theorem is especially useful in E^3 for forming tetrahedralizations that respect specified planar facets. If the bounding segments of these facets are subdivided so that the subsegments are strongly Delaunay, then a constrained tetrahedralization exists. Hence, fewer vertices are needed than in the most common practice in the literature, wherein additional vertices are inserted in the relative interiors of facets to form a conforming (but *un*constrained) Delaunay tetrahedralization.

1 Introduction

Many applications can benefit from triangulations that have properties similar to Delaunay triangulations, but are constrained to contain specified edges or faces. For instance, Delaunay triangulations have desirable properties when used for function interpolation, but a triangulation might be required to conform to specified facets so that discontinuities can be represented. Delaunay triangulations can also serve as meshes that represent objects for rendering or the numerical solution of partial differential equations. In these cases, each triangulation is required to conform to the shape of the object being modeled.

In two dimensions, there are two popular alternatives for creating a Delaunay-like triangulation that conforms to constraints. In either case, the input is a *planar straight line graph* (PSLG) X, which is a set of vertices and segments (constraining edges) as illustrated in Figure 1 (upper left). A PSLG is required to contain both endpoints of every segment it contains, and a segment may intersect vertices and other segments only at its endpoints. A triangulation is sought that includes the vertices of X and respects the segments of X.

The first alternative is to form a *conforming Delaunay triangulation* (Figure 1, lower left). The vertices of X are augmented by additional vertices (sometimes called *Steiner points*) carefully chosen so that the Delaunay triangulation of the augmented vertex set conforms to all the segments—in other words, so that each segment is represented by a contiguous linear sequence of edges of the triangulation. Edelsbrunner and Tan [4] show that any PSLG can be triangulated with the addition of no more than $\mathcal{O}(m^2n)$ augmenting vertices, where m is the number of segments in X, and n is the number of vertices. Not every PSLG requires this many augmenting vertices, but the numbers required in practice may nonetheless seem undesirably large for some applications.

The second alternative is to form a *constrained Delaunay triangulation* (CDT) [1] (Figure 1, lower right). A CDT of X has no vertices not in X, and every segment of X is a single edge of

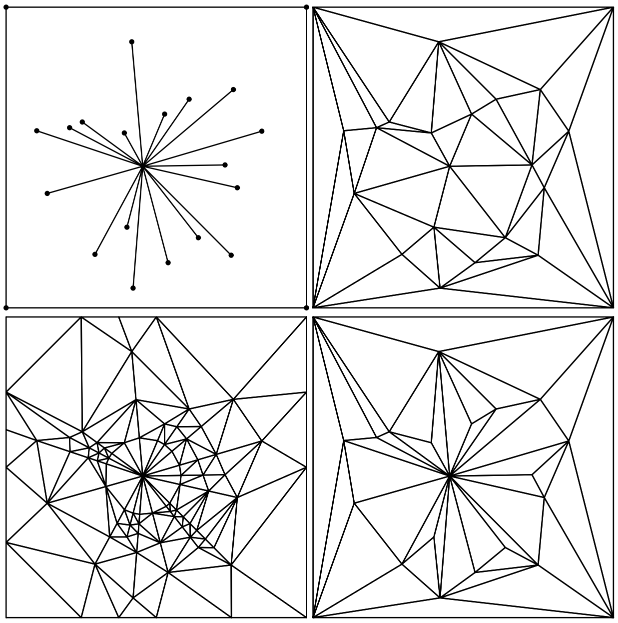

Figure 1: The Delaunay triangulation (upper right) of the vertices of a PSLG (upper left) might not respect the segments of the PSLG. These segments can be incorporated by adding vertices to obtain a conforming Delaunay triangulation (lower left), or by forgoing Delaunay triangles in favor of constrained Delaunay triangles (lower right).

Supported in part by the Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, USAF under agreement number F30602-96-1-0287, in part by the National Science Foundation under Grant CMS-9318163, and in part by a grant from the Intel Corporation.

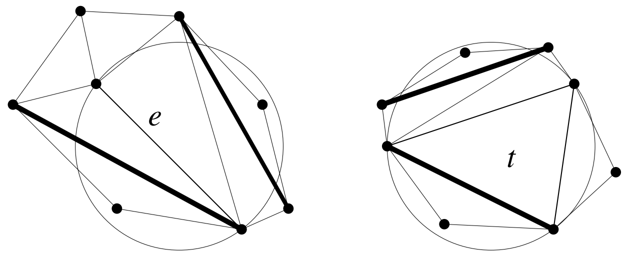

Figure 2: The edge e and triangle t are each constrained Delaunay. Bold lines represent segments.

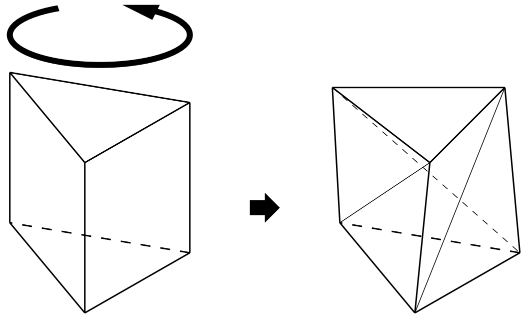

Figure 3: Schönhardt's untetrahedralizable polyhedron (right) is formed by rotating one end of a triangular prism (left), thereby creating three diagonal reflex edges.

the CDT. However, a CDT is not a Delaunay triangulation. In an ordinary Delaunay triangulation, every simplex (of any dimensionality) is *Delaunay*. A simplex is Delaunay if its vertices are in X and there exists a *circumcircle* of the simplex—a circle that passes through all its vertices—that encloses no vertex of X (although any number of vertices is permitted on the circle itself). In a CDT, this requirement is waived, and instead every simplex must either be a segment specified in X or be *constrained Delaunay*. A simplex is constrained Delaunay if it has a circumcircle that encloses no vertex of X that is *visible* from any point in the relative interior of the simplex; and furthermore, the relative interior of the simplex does not intersect any segment. Visibility is occluded only by segments of X.

Figure 2 demonstrates examples of a constrained Delaunay edge e and a constrained Delaunay triangle t. Input segments appear as bold lines. Although there is no empty circle that contains e, the depicted circumcircle of e encloses no vertex that is visible from the relative interior of e. There are two vertices inside the circle, but both are hidden behind segments. Hence, e is constrained Delaunay. Similarly, the circumcircle of t contains two vertices, but both are hidden from the interior of t by segments, so t is constrained Delaunay.

The advantage of a CDT over a conforming Delaunay triangulation is that it has no vertices other than those in X. However, a conforming Delaunay triangulation's triangles are Delaunay, but those of a CDT are not. Nevertheless, a CDT retains many of the desirable properties of Delaunay triangulations. For instance, a two-dimensional CDT maximizes the minimum angle in the triangulation, compared with all other constrained triangulations of X [7].

Unfortunately, CDTs have not been generalized to dimensions higher than two. One reason is that in three or more dimensions, there are polytopes that cannot be triangulated at all without additional vertices. Schönhardt [12] furnishes a three-dimensional example depicted in Figure 3 (right). The easiest way to envision this polyhedron is to begin with a triangular prism. Imagine grasping the prism so that one of its two triangular faces cannot move, while the opposite triangular face is rotated slightly about its center without moving out of its plane. As a result, each of the three square faces is broken along a diagonal *reflex edge* (an edge at which the polyhedron is locally nonconvex) into two triangular faces. After this transformation, the upper left corner and lower right corner of each (former) square face are separated by a reflex edge and are no longer visible to each other within the polyhedron. Any four vertices of the polyhedron include two separated by a reflex edge; thus, any tetrahedron whose vertices are vertices of the polyhedron will not lie entirely within the polyhedron. Therefore, Schönhardt's polyhedron cannot be tetrahedralized without additional vertices.

Ruppert and Seidel [11] add to the difficulty by proving that it is NP-hard to determine whether a polyhedron is tetrahedralizable. Even among polytopes that can be triangulated without additional vertices, there might not always be a triangulation that is in any reasonable sense "constrained Delaunay." Work is needed to determine what features of polytopes make them amenable to being triangulated with Delaunay-like simplices. This paper takes a first step by demonstrating that there is an easily tested condition that guarantees that a CDT exists, for a conservative definition of CDT. Furthermore, there are useful consequences to this guarantee.

The input for which a triangulation is sought is called a *piecewise linear complex* (PLC), following Miller, Talmor, Teng, Walkington, and Wang [8]. Before considering PLCs in their full generality, consider the special case in which a PLC X is a set of vertices and *constraining simplices* in E^d. Each constraining simplex is a k-dimensional simplex (henceforth, k-simplex) having $k+1$ affinely independent vertices, where $1 \leq k \leq d-1$. Each constraining simplex must appear as a face of the final triangulation. (Observe that a PSLG is a two-dimensional PLC.)

PLCs have restrictions similar to those of PSLGs. If X contains a simplex s, then X must contain every lower-dimensional face of s, including its vertices. Any two simplices of a PLC, if one is not a face of the other, may intersect only at a shared lower-dimensional face or vertex.

Several definitions are needed prior to a formal statement of the main result. Say that the visibility between two points p and q in E^d is *occluded* if there is a $(d-1)$-simplex s of X such that p and q lie on opposite sides of the hyperplane that contains s, and the line segment pq intersects s (either in the boundary or in the relative interior of s). If either p or q lies in the hyperplane containing s, then s does not occlude the visibility between them. Simplices in X of dimension less than $d-1$ do not occlude visibility. The points p and q can *see* each other if there is no occluding $(d-1)$-simplex of X.

Let s be a k-simplex (for any k) whose vertices are in X (but s is not necessarily a constraining simplex of X). Let S be a (full-dimensional) sphere in E^d; S is a *circumsphere* of s if S passes through all the vertices of s. If $k = d$, then s has a unique circumsphere; otherwise, s has infinitely many circumspheres. The simplex s is *Delaunay* if there is a circumsphere S of s such that no vertex of X lies inside S. The simplex s is *strongly Delaunay* if there is a circumsphere S of s such that no vertex of X lies inside *or on* S, except the vertices of s. Every 0-simplex is trivially strongly Delaunay.

The simplex s is *constrained Delaunay* if no constraining simplex of X intersects the interior of s unless it contains s in its entirety, and there is a circumsphere S of s such that no vertex of X inside S is visible from any point in the relative interior of s.

A PLC X is said to be *ridge-protected* if each constraining simplex in X of dimension $d-2$ or less is strongly Delaunay.

The main result of this paper is that if X is ridge-protected, and if no $d+2$ vertices of X lie on a common sphere, then the

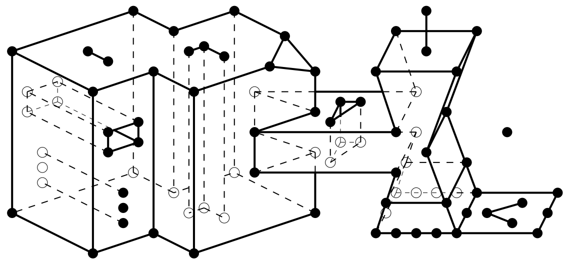

Figure 4: Any facet of a PLC may contain holes, slits, and vertices, which may be used to enforce the presence of specific faces or to support intersections with other facets.

constrained Delaunay d-simplices defined on the vertices of X collectively form a triangulation of X. This triangulation is called the *constrained Delaunay triangulation* of X. The condition that X be ridge-protected holds trivially in two dimensions (hence the success of two-dimensional CDTs), but not in three or more.

Note that it is not sufficient for each lower-dimensional constraining simplex to be Delaunay; if Schönhardt's polyhedron is specified so that all six of its vertices lie on a common sphere, then all of its edges (and its faces as well) are Delaunay, but it still does not have a tetrahedralization. It is not possible to place the vertices of Schönhardt's polyhedron so that all three of its reflex edges are strongly Delaunay (though any two may be).

If some $d + 2$ vertices of X are cospherical, the existence of a CDT can be established through perturbation arguments (not presented here), but the CDT is not necessarily unique. There may be constrained Delaunay d-simplices whose interiors are not disjoint. Some of these simplices must be omitted to yield a proper triangulation.

It is appropriate, now, to consider a more general definition of PLC. To illustrate the limitation of the definition given heretofore, consider finding a tetrahedralization of a three-dimensional cube. If each square face is represented by two triangular constraining simplices, then the definition given above requires that the diagonal edge in each face be a constraining edge of X. However, the result of this paper may be obtained even if these edges are not strongly Delaunay. Hence, a PLC may contain constraining *facets*, which are more general than constraining simplices. Each facet is a polytope of any dimension from one to $d - 1$, possibly with holes and lower-dimensional facets inside it. A k-simplex s is said to be a *constraining simplex in* X if s appears in the k-dimensional CDT of some k-facet in X. If X is ridge-protected, then the CDT of a facet is just its Delaunay triangulation, because the boundary simplices of each facet are strongly Delaunay.

Figure 4 illustrates a three-dimensional PLC. As the figure illustrates, a facet may have any number of sides, may be nonconvex, and may have holes, slits, or vertices inside it. Two facets, if one is not a boundary facet of the other, may intersect only at shared lower-dimensional facets and vertices. See Miller et al. [8] for a complete list of restrictions.

The advantage of employing whole facets, rather than individual constraining simplices, is that only the lower-dimensional boundary facets of a $(d - 1)$-facet need be included in the PLC (including interior boundaries, at holes and slits). The lower-dimensional faces that are introduced in the relative interior of a $(d - 1)$-facet when it is triangulated do not need to be strongly Delaunay for this paper's result to hold.

This advantage does not extend to lower-dimensional facets, because a ridge-protected k-facet, for $k \leq d-2$, must be composed of strongly Delaunay k-simplices, and it is easy to show that any lower-dimensional face of a strongly Delaunay simplex is strongly Delaunay. Hence, in a ridge-protected PLC, all simplicial faces in the Delaunay triangulation of a facet are required to be strongly Delaunay, except for the faces in the triangulation of a $(d-1)$-facet that are not in the facet's boundary.

Testing whether a PLC is ridge-protected is straightforward. Form the Delaunay triangulation of the vertices of the PLC. If a constraining simplex s is missing from the triangulation, then s is not strongly Delaunay. Otherwise, s is Delaunay; testing whether s is strongly Delaunay is a local operation equivalent to determining whether the dual face of s in the corresponding Voronoi diagram is nondegenerate.

Why is it useful to know that ridge-protected PLCs have CDTs? Although a given PLC X may not be ridge-protected, it can be made ridge-protected by splitting simplices that are not strongly Delaunay into smaller simplices, with the insertion of additional vertices. The result is a new ridge-protected PLC Y, which has a CDT. The CDT of Y is not a CDT of X, because it has vertices that X lacks, but it is what I call a *conforming constrained Delaunay triangulation* (CCDT) of X: conforming because of the additional vertices, and constrained because its simplices are constrained Delaunay (rather than Delaunay).

One advantage of a CCDT over a conforming Delaunay triangulation is that the number of additional vertices needed is generally smaller. Once a PLC X has been augmented to form a ridge-protected PLC Y, the Delaunay triangulation of the vertices of Y contains all the constraining simplices of Y of dimension $d - 2$ or smaller (because they are strongly Delaunay), but does not necessarily respect the $(d - 1)$-simplices of Y. The result of this paper implies that the CDT of Y may be formed without adding any more vertices to Y. This idea stands in apposition to the most common method of recovering unrepresented facets in three-dimensional Delaunay-based mesh generation algorithms [14, 6, 15, 10], wherein additional vertices are inserted within the facet (for instance, where an edge of the tetrahedralization intersects the missing facet).

Because of the large disparity between the number of vertices required for two-dimensional constrained and conforming Delaunay triangulations, I suspect that asymptotically more vertices are needed (in the worst case) to produce a Delaunay tetrahedralization that conforms to a set of segments and facets in E^3 than are needed to produce a Delaunay tetrahedralization that merely conforms to the segments. Because the latter can be converted into a CCDT, the former may not be needed.

The utility of the CCDT is further buttressed by a practical algorithm for three-dimensional mesh generation that I describe elsewhere [13]. One variant of the Delaunay refinement algorithm described therein uses the constrained Delaunay property of the tetrahedra it produces to prove guaranteed bounds on the quality of the final mesh. The results of this paper are the necessary underpinnings of that algorithm.

2 Proof of Existence of the CDT

Throughout the proof, the terms "simplex" and "convex hull" refer to closed, convex sets of points; hence, they include all the points on their boundaries and in their interiors.

Let X be a ridge-protected PLC in E^d, and suppose that some subset of $d + 1$ vertices in X is affinely independent. The set of constrained Delaunay d-simplices defined on X forms a triangulation of the vertices of X if X contains no $d + 2$ vertices that lie on a common sphere. This fact may be proven in two steps: first, by showing that every point in the convex hull of X is contained in some constrained Delaunay d-simplex of X; second, by showing that any two constrained Delaunay d-simplices have disjoint interiors, and that the d-simplices are aligned on their $(d - 1)$-faces. Hence, the union of the constrained Delaunay d-simplices is the convex hull of X, and the constrained Delaunay d-simplices (and

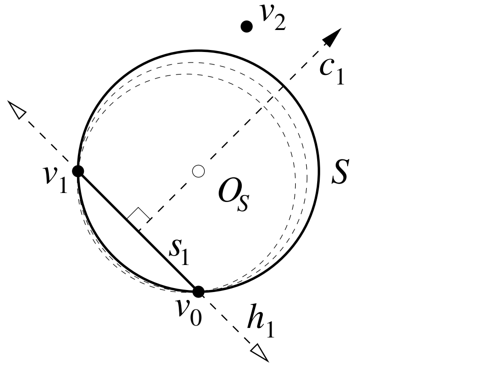

Figure 5: An empty sphere S that circumscribes s_1, expanding in search of another vertex v_2.

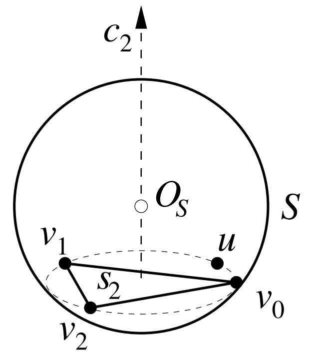

Figure 6: The point u lies on S, but is ignored because it is not affinely independent from $\{v_0, v_1, v_2\}$.

their faces) form a simplicial complex. Only the second step requires the assumption that no $d + 2$ vertices are cospherical.

Both steps are based on a straightforward procedure for "growing" a simplex, vertex by vertex. The procedure is presumed to be in possession of a constrained Delaunay k-simplex s_k, whose vertices are $v_0, \ldots, v_k$, and produces a constrained Delaunay $(k+1)$-simplex s_{k+1} that possesses a face s_k and an additional vertex v_{k+1}.

Because s_k is constrained Delaunay, it has a circumsphere S that encloses no vertex visible from the interior of s_k. The growth procedure expands S like a bubble, so that its center O_S moves in a direction orthogonal to the k-dimensional hyperplane h_k that contains s_k, as illustrated in Figure 5. Because O_S is moving in a direction orthogonal to h_k, O_S remains equidistant from the vertices of s_k, and hence it is always possible to choose a radius for S that ensures that S continues to pass through all the vertices of s_k. The expansion ends when S contacts an additional vertex v_{k+1} that is visible from some point in the interior of s_k. Once it is found, the forthcoming Theorem 5 guarantees that v_{k+1} is visible from every point in s_k. The simplex s_{k+1} defined by v_{k+1} and the vertices of s_k is constrained Delaunay, a fact attested to by the circumsphere S, which can be shown (again by Theorem 5) to enclose no vertex visible from any point in the interior of s_{k+1}.

The motion of the sphere center O_S is governed by a *direction vector* c_k, which is constrained to be orthogonal to h_k, but may otherwise be specified freely. In the limit as O_S moves infinitely far away, S will approach the $(d-1)$-dimensional hyperplane that contains s_k and h_k and is orthogonal to c_k. The region enclosed by S will approach an open half-space bounded by this hyperplane. A point is said to be *above* h_k if it lies in this open half-space. Any vertex outside S that S comes in contact with while expanding must lie above h_k; the portion of S below h_k is shrinking toward the inside of S.

A special case occurs if the sphere S already contacts a vertex u not in s_k before S begins expanding. If u is affinely independent of s_k, and is visible from the interior of s_k, then it is immediately accepted as v_{k+1}. Otherwise, u is ignored. (The affinely dependent case can only occur for $k \geq 2$; for instance, when a triangle grows to become a tetrahedron, as illustrated in Figure 6, there may be ignored vertices on the triangle's planar circumcircle.)

If the procedure is to succeed, c_k must satisfy the following *visibility hypothesis*. The open half-space above h_k must contain a vertex of X that is visible from the interior of s_k. In the case of unconstrained Delaunay triangulations, this simply means that the half-space contains a vertex of X. Theorem 4, to follow, clarifies the case of constrained Delaunay triangulations by showing that if X is ridge-protected and there is a vertex above h_k, then there is a vertex above h_k that is visible from an arbitrarily chosen point in s_k. Hence, for a ridge-protected PLC, the visibility hypothesis holds in the constrained case whenever it holds in the unconstrained case.

This growth procedure is the basis for the following proof.

Theorem 1 *Let p be any point in the convex hull of X. If X is ridge-protected, then some constrained Delaunay d-simplex of X contains p.*

Proof: The proof is based on a two-stage constructive procedure similar to one used by Fortune [5] to prove the existence of unconstrained Delaunay triangulations. The first stage finds an arbitrary constrained Delaunay d-simplex, and involves $d + 1$ steps, numbered zero through d. For step zero, define a sphere S whose center is an arbitrary vertex v_0 in X, and whose radius is zero. The vertex v_0 seeds the starting simplex $s_0 = \{v_0\}$.

During the remaining steps, illustrated in Figure 7, S expands according to the growth procedure described above. At step $k + 1$, the direction vector c_k is chosen so that some vertex of X lies above h_k, and thus the visibility hypothesis is established. Such a choice of c_k is always possible because X contains $d + 1$ affinely independent vertices, and thus some vertex is not in h_k. Theorem 4 guarantees that some vertex is visible above h_k from the interior of s_k. If there are several vertices visible, let v_{k+1} be the first such vertex contacted by the expanding sphere, so that no vertex inside the sphere is visible from the interior of s_k. Theorem 5 shows that the new simplex s_{k+1} is constrained Delaunay.

After step d, s_d is a constrained Delaunay d-simplex. If s_d contains p, the procedure is finished; otherwise, the second stage begins.

Consider the directed line segment qp, where q is an arbitrary point in the interior of s_d chosen so that qp does not intersect any simplices (defined on the vertices of X) of dimension $d - 2$ or smaller, except possibly at the point p. (Such a choice of q is always possible; the set of points in s_d that *don't* satisfy this condition has measure zero.) The second stage "walks" along the segment qp, traversing a sequence of constrained Delaunay d-simplices that intersect qp (illustrated in Figure 8), until it finds one that contains p.

Wherever qp exits a d-simplex, the next d-simplex may be constructed as follows. Let s be the face through which qp exits the current d-simplex. The $(d-1)$-simplex s is either a constraining simplex in X, or is constrained Delaunay. s is the base from which another constrained Delaunay d-simplex is grown, with the direction vector chosen to be orthogonal to s and directed out of the previous d-simplex. Is the visibility hypothesis satisfied? Clearly, the point p is above s. Some vertex of X must lie above s, because if none did, the convex hull of X would not intersect the half-space

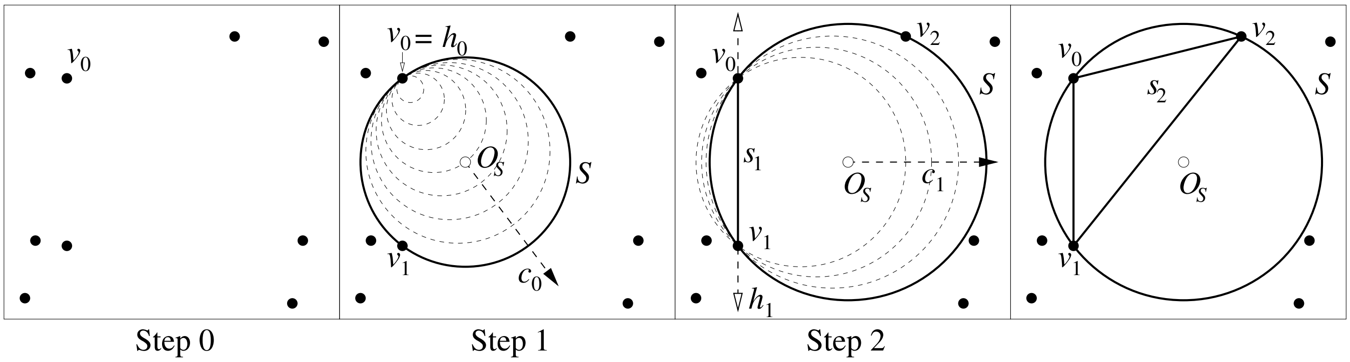

Figure 7: Growing a Delaunay simplex. If $d > 2$, c_2 will point directly out of the page.

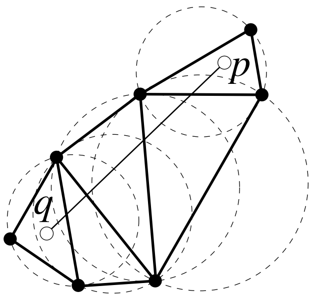

Figure 8: Walking to the d-simplex that contains p.

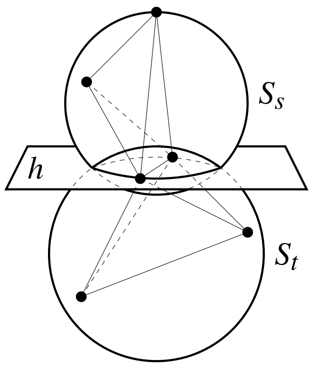

Figure 9: The constrained Delaunay d-simplices s and t intersect at a lower-dimensional shared face (in this illustration, an edge).

above s and thus would not contain p. Hence, by Theorem 4, some vertex above s is visible from the interior of s.

If s is constrained Delaunay, a new constrained Delaunay d-simplex may be formed as in the first stage. If s is a constraining $(d-1)$-simplex in X, but is not constrained Delaunay, then the growth procedure must be modified slightly to account for the fact that s does not have a circumsphere that encloses no vertex visible from s. In this case, the "expanding" sphere S begins with its center infinitely far below s (so that the "inside" of S is the open half-space below s), and the portion of S above s is initially empty but expands as usual. A different proof (Theorem 8) is needed to show that the new d-simplex is constrained Delaunay.

Because the new d-simplex is above s, it is distinct from the previous d-simplex. Because each successive d-simplex intersects a subsegment of qp having nonzero length, each successive $(d-1)$-face intersects qp closer to p, and thus no simplex is visited twice. Since only a finite number of simplices can be defined (over a finite set of vertices), the procedure must terminate; and since the procedure will not terminate unless the current d-simplex contains p, there exists a Delaunay d-simplex that contains p. ■

This procedure is recognizable as the basis for a well-known algorithm, called *gift-wrapping*, *graph traversal*, or *incremental search*, for constructing Delaunay triangulations [2]. Gift-wrapping begins by finding a single Delaunay d-simplex, which is used as a seed upon which the remaining Delaunay d-simplices crystallize one by one. Each $(d-1)$-face of a Delaunay d-simplex is used as a base from which to search for the vertex that serves as the apex of an adjacent d-simplex. Theorem 1 shows that gift-wrapping can be used to produce CDTs as well, at least for ridge-protected PLCs.

As every point in the convex hull of X is contained in a constrained Delaunay d-simplex, it remains only to show that the set of constrained Delaunay d-simplices do not occupy common volume or fail to intersect neatly. It is only here that the assumption that no $d+2$ vertices are cospherical is needed.

Theorem 2 *Suppose that no $d+2$ vertices of X lie on a common sphere. Then the constrained Delaunay d-simplices of X (and their faces) collectively form a simplicial complex.*

Proof: First, I show that constrained Delaunay d-simplices have disjoint interiors. Suppose for the sake of contradiction that some point p lies in the interior of two distinct constrained Delaunay d-simplices s and t. Because s and t are constrained Delaunay, every vertex of s and t is visible from p.

Let S_s and S_t be the circumspheres of s and t. S_s and S_t cannot be identical, because s and t each have at least one vertex not shared by the other; if S_s and S_t were the same, at least $d+2$ vertices would lie on S_s. S_s cannot enclose S_t, nor the converse, because S_s and S_t enclose no vertices visible from p. Hence, either S_s and S_t are entirely disjoint (and thus so are s and t), or their intersection is a $(d-1)$-dimensional circle or point and is contained in a $(d-1)$-dimensional hyperplane h, as Figure 9 illustrates. (If S_s and S_t intersect at a single point, h is chosen to be tangent to both spheres.) Without loss of generality, suppose h is oriented horizontally, with the center of S_s directly above the center of S_t. Because p lies in the interiors of s and t, either some vertex of s lies below h, or some vertex of t lies above h. In the former case, there is a vertex of s inside S_t that is visible from a point (p) inside t, so t is not constrained Delaunay. In the latter case, there is a vertex of t inside S_s that is visible from a point inside s, so s is not constrained Delaunay. Either case implies a contradiction, so s and t have disjoint interiors and can intersect only at their boundaries.

Recall from the proof of Theorem 1 that if s is a constrained Delaunay d-simplex, then from any $(d-1)$-face of s not in the

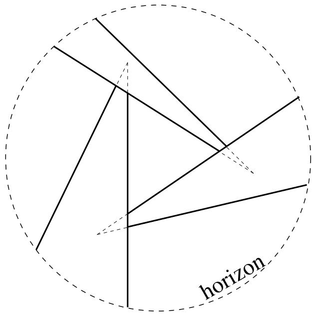

Figure 10: Spherical projection of the half-space above your vantage point.

boundary of the convex hull of X, one can find an adjoining constrained Delaunay d-simplex that shares that $(d-1)$-face. No other constrained Delaunay d-simplex may occupy the same volume. It follows that constrained Delaunay d-simplices meet neatly on their $(d-1)$-faces; and since the union of all constrained Delaunay d-simplices is the convex hull of X, the constrained Delaunay d-simplices collectively form a simplicial complex. ∎

It is now straightforward to show that every constraining simplex in X appears as a face of the CDT. Simply note that from any constraining simplex in X, one may grow a constrained Delaunay d-simplex whose faces include the constraining simplex.

The remainder of this section is devoted to completing the proof of Theorem 1. The first step is to prove the visibility hypothesis. One potential difficulty is illustrated (for the three-dimensional case) in Figure 10. Imagine that you are standing on a $(d-1)$-simplex, scanning the half-space above the simplex for a vertex that can serve as the apex of a d-simplex. Looking up into the sky, you see the three illustrated $(d-1)$-facets, each of which occludes the apical vertex of another; the remaining vertices of these facets are hidden below the horizon (in the half-space below you). Hence, no vertex in the half-space is visible from your vantage point.

To prove the existence of a constrained Delaunay triangulation, one must show that this possibility is precluded if X is ridge-protected. In Figure 10, observe that the inner edges of the three facets form a cycle of overlapping simplices. The proof operates by attacking the possibility that the strongly Delaunay faces that bound the occluding $(d-1)$-facets can form such a cycle.

Let p be an arbitrary vantage point in E^d. Let s and t be any two strongly Delaunay simplices. Say that s *overlaps* t from the viewpoint p if some point of s not shared by t lies directly between p and t. In other words, there exists a point p_s of s and a point p_t of t such that $p_s \notin t$ and p_s lies between p and p_t.

Lemma 3 *From any fixed vantage point p, X contains no cycle of consecutively overlapping strongly Delaunay constraining simplices.*

Proof: For each strongly Delaunay simplex s of X, let S_s be a sphere that passes through the vertices of s, but neither passes through nor encloses any other vertex. Hence, S_s attests to the fact that s is strongly Delaunay. There may be many such spheres, but one should be chosen arbitrarily and used for all pairwise comparisons of simplices herein.

Let O_s and r_s be the center and radius of S_s, respectively. Consider the function $\Psi_p(s) = |pO_s|^2 - r_s^2$ defined over the set of strongly Delaunay simplices relative to a vantage point p. The proof stands on the fact that if s and t are strongly Delaunay simplices and s overlaps t from the viewpoint p, then $\Psi_p(s) < \Psi_p(t)$. Hence, the overlap relation among strongly Delaunay simplices

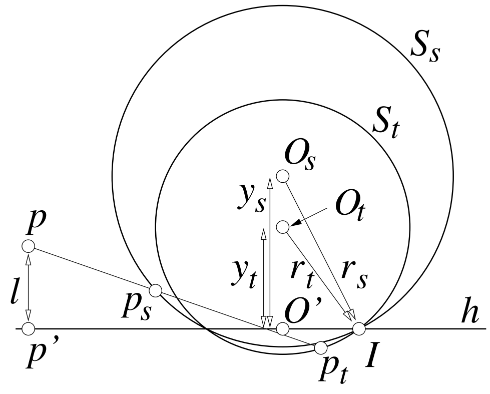

Figure 11: A two-dimensional cross-section of E^d that passes through p, O_s, and O_t. The hyperplane h is orthogonal to O_sO_t, and thus extends orthogonally out of the page. Note that p_s and p_t do not necessarily lie in this cross-section; they are depicted here as a reminder of why p must lie above h.

defines a partial order, and no cycle of consecutively overlapping strongly Delaunay simplices is possible.

It remains only to show that $\Psi_p(s) < \Psi_p(t)$ if s overlaps t from the viewpoint p. Because s contains a point that t lacks (namely p_s), s must also possess a vertex that t lacks; hence, S_s (which passes through this vertex) and S_t (which does not) are distinct. If S_s and S_t intersect in a $(d-1)$-dimensional circle, let h be the $(d-1)$-dimensional hyperplane that passes through the circle of intersection (recall Figure 9). Otherwise, let h be the $(d-1)$-dimensional hyperplane tangent to S_s at the point of S_s nearest S_t.

Assume without loss of generality that h is oriented horizontally, with the center of S_s directly above the center of S_t. Observe that S_s encloses any portion of S_t above h, and S_t encloses any portion of S_s below h. Because s and t are strongly Delaunay, every vertex of s, and hence every point in s, lies in or above h; and every point in t lies in or below h. Every vertex of s or t that lies in h must belong to both s and t; hence, every point of s or t that lies in h belongs to both. Recall that there exists a point p_s of s and a point p_t of t such that $p_s \notin t$ and p_s lies between p and p_t. The point p_s must lie strictly above h, and p_t lies in or below h, so p lies strictly above h.

Consider Figure 11, which depicts the two-dimensional cross-section through E^d that passes through p and the centers O_s and O_t of the spheres S_s and S_t. Let y_s be the signed height of O_s above h; y_s is negative if O_s is below h. Similarly, let y_t be the signed height of O_t above h, and let l be the signed height of p above h. Let r_s and r_t be the radii of S_s and S_t, respectively. Let p' be the orthogonal projection of p onto h (in other words, pp' is orthogonal to h), and let O' be the orthogonal projection of O_s onto h, which coincides with the orthogonal projection of O_t onto h. Let I be any point of intersection of S_s, S_t, and h; if the spheres do not intersect, let $I = O'$. Observe that if S_s and S_t intersect, then $r_t^2 = y_t^2 + |O'I|^2$; otherwise, $r_t^2 < y_t^2$ and $|O'I| = 0$, as Figure 12 illustrates.

From these relationships and the Pythagorean Theorem, we have

$$\begin{aligned}\Psi_p(t) - \Psi_p(s) &= |pO_t|^2 - |pO_s|^2 - r_t^2 + r_s^2 \\ &\geq [(y_t - l)^2 + |p'O'|^2] \\ &\quad -[(y_s - l)^2 + |p'O'|^2] \\ &\quad -[y_t^2 + |O'I|^2] + [y_s^2 + |O'I|^2] \\ &= 2l(y_s - y_t).\end{aligned}$$

Because l and $y_s - y_t$ are both positive, $\Psi_p(t) > \Psi_p(s)$. ∎

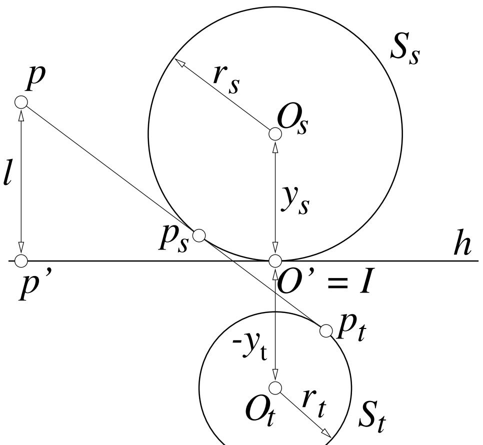

Figure 12: The case where S_s and S_t do not intersect.

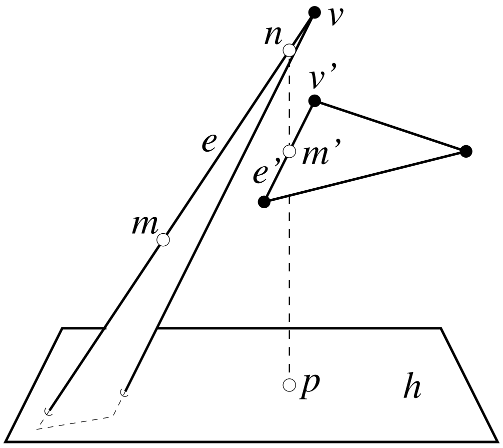

Figure 13: Some vertex above h is visible from p.

Edelsbrunner [3] presents an acyclicity theorem that is nearly identical[1] to Lemma 3. The proof given here is much simpler.

The impossibility of a cycle of overlapping strongly Delaunay simplices is the key to proving the visibility hypothesis for ridge-protected PLCs.

Theorem 4 (Visibility Hypothesis) *Let h be a $(d-1)$-dimensional hyperplane, and let p be a point in h. Suppose there is at least one vertex of X in the open half-space above h. If X is ridge-protected, then at least one vertex of X in the open half-space above h is visible from p.*

Proof: There is at least one vertex above h. Either it is visible from p and the result follows, or there must be a $(d-1)$-facet of X occluding its visibility. Hence, at least a portion of the boundary of at least one $(d-1)$-facet lies above h. It follows that some simplex e (of dimension $d-2$ or less) in the boundary of a $(d-1)$-facet is at least partly visible from p in the half-space above h. (The only possible alternative is a single facet blotting out the whole sky, which would only be possible if facets could be curved). Let m be a point in e that is above h and visible from p, as illustrated in Figure 13. Assume without loss of generality that m is in the relative interior of e by choosing e to be of as low dimension as possible; for instance, if m lies in an edge of a tetrahedron in E^5, choose e to be the edge and not the tetrahedron. If e is a vertex the result follows, so assume e is of dimension at least one.

The proof proceeds by "walking" from m toward a vertex of e, replacing e with any simplex that occludes its visibility from p, and continuing the walk on the new simplex. This process is repeated until a vertex visible from p is found.

Begin by observing that because e is a simplex that contains m, e must have at least one vertex v above h. (The other vertices of e might lie on or below h.) If v is visible from p, the result follows. Otherwise, m is visible from p but v is not.

Let n be the point nearest m on the line segment mv that is not visible from p. (In other words, n is the first occluded point encountered when walking from m to v.) The segment pn must intersect some $(d-1)$-facet of X at some point m'. (If there are several $(d-1)$-facets occluding the view from p to n, consider only the facet that intersects pn closest to p, so that m' is visible from p). Since n is the first occluded point on mv, m' must lie in the boundary of the occluding facet. Generally, m' will lie in the interior of a $(d-2)$-simplex bounding the facet, but in some cases m' may lie in a lower-dimensional boundary; let e' be the lowest-dimensional boundary simplex that contains m'. If e' is a vertex, the proof is complete, so assume e' is not a vertex.

Because e and e' lie in boundaries of facets in X, e and e' are strongly Delaunay. Clearly, e' overlaps e from the viewpoint p.

Let v' be a vertex of e' that is above h. If v' is visible from p, the proof is complete. Otherwise, find another facet-bounding simplex e'' of dimension $d-2$ or less that overlaps e' from the viewpoint p, and search for one of its vertices. Repeat the process until a vertex visible from p is found. Because X contains only a finite number of constraining simplices, and because Lemma 3 rules out the possibility of a cycle of consecutively overlapping strongly Delaunay simplices, the search must end. ■

The visibility hypothesis tells us that if the open half-space above a k-simplex s_k contains a vertex, then at least one candidate vertex u above s_k is visible from the interior of s_k. Can we find a constrained Delaunay simplex s_{k+1} by taking the convex hull of s_k and u? There are two catches. First, the candidate vertex u, although visible from at least one point inside s_k, might not be visible from *all* points inside s_k. Second, some vertex visible from inside s_{k+1} might prevent s_{k+1} from being constrained Delaunay. The following theorem shows that these problems do not arise in the case that matters: when a sphere circumscribing s_k and u encloses no vertex visible from the interior of s_k.

Theorem 5 *Let X be a ridge-protected PLC. Let s_k be a constrained Delaunay k-simplex, for some $k < d$. Let u be a vertex of X that is affinely independent of s_k and is visible from some point p in the relative interior of s_k. Suppose that there is a sphere S that passes through u and all the vertices of s_k, and encloses no vertex of X that is visible from p. Let s_{k+1} be the convex hull of s_k and u. Then no constraining facet of X intersects the relative interior of s_{k+1} unless it contains s_{k+1}. Furthermore, no vertex inside S is visible from any point in the relative interior of s_{k+1}; hence, s_{k+1} is constrained Delaunay.*

The proof requires the following two lemmata.

Lemma 6 *Let S be a sphere, and let H_S be the convex hull of all the vertices of X that lie on or inside S. Let t be a strongly Delaunay simplex. Suppose that some point of t lies between two points of H_S not in t. Then at least one vertex of t lies inside S.*

Proof: Suppose, for the sake of contradiction, that all vertices of t lie on or outside S. Because t is strongly Delaunay, there is some

[1] Edelsbrunner [3, Section 3] investigates relationships among Delaunay, rather than strongly Delaunay, simplices, and arrives at roughly the same conclusion as the present Lemma 3. I believe that his conclusion is incorrect in the case of simplices that share a common circumsphere.

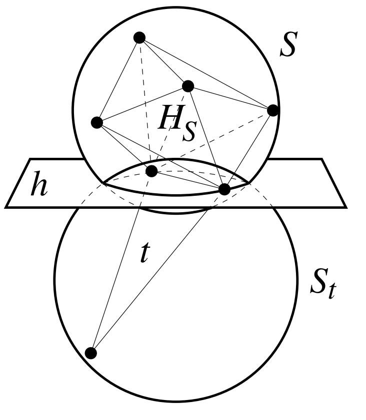

Figure 14: No point of t can lie between two points of H_S not in t.

sphere S_t that circumscribes t, but has no other vertices on or inside it. Because there are points of H_S not in t, a portion of S must lie outside S_t. By assumption, some point of t falls between two points of H_S; this point lies inside S and hence cannot be a vertex of t, so t has at least two vertices. Because these vertices are on or outside S, and S_t cannot be identical to S, a portion of S_t must lie outside S. It follows that the spheres S and S_t intersect in a $(d-1)$-dimensional circle.

Let h be the $(d-1)$-dimensional hyperplane that contains the intersection of S and S_t, as illustrated in Figure 14. Without loss of generality, suppose h is oriented horizontally with the center of S directly above the center of S_t. The vertices of t lie on S_t but not inside S, so all points of t must lie on or below h. Because t is strongly Delaunay, every vertex of H_S lies on or above h, and only those vertices of H_S that are shared with t may lie on h. Therefore, any point of H_S not in t lies strictly above h. Hence, no point of t can lie between two points of H_S not in t. The result follows by contradiction. ■

Lemma 7 *Let S be a sphere, and let H_S be the convex hull of all the vertices of X that lie on or inside S. Let q and z be two points in H_S with the property that qz intersects a strongly Delaunay simplex e that contains neither q nor z. Let m be the intersection point, and let p be any point in H_S from which m is visible. If X is ridge-protected, then there is a vertex of X inside S that is visible from p.*

Proof: See Figure 15. By assumption, q and z do not lie in e. Because the point m of e lies between two points that are in H_S but not in e, Lemma 6 implies that at least one vertex v of e lies inside S. If v is visible from p, the result follows. Otherwise, observe that m and v are both in H_S and are both in e. The point m is visible from p, but v is not.

Let n be the point nearest m on the line segment mv that cannot see p. The line segment pn must intersect the boundary of some $(d-1)$-facet of X at some point m' (if there are several $(d-1)$-facets occluding the view from p to n, consider only the facet that intersects pn closest to p). Let e' be the lowest-dimensional boundary simplex that contains m'. Because e' lies in the boundary of a facet, e' is strongly Delaunay. Clearly, e' overlaps e from the viewpoint p. Observe that p and n both lie in H_S, but they cannot lie in e'; if either of them did, the facet containing e' would not obstruct the visibility between them. Again, Lemma 6 implies that at least one vertex v' of e' lies inside S. If v' is visible from p, the result follows. Otherwise, find another simplex e'' that overlaps e' from the viewpoint p and continue.

The act of iterating in this manner yields a sequence of strongly Delaunay simplices, each overlapping the previous one. These iterations must terminate because the number of constraining simplices in X is finite and, by Lemma 3, a cycle of consecutively overlapping strongly Delaunay simplices is not possible. Hence, the procedure will eventually yield a vertex visible from p. ■

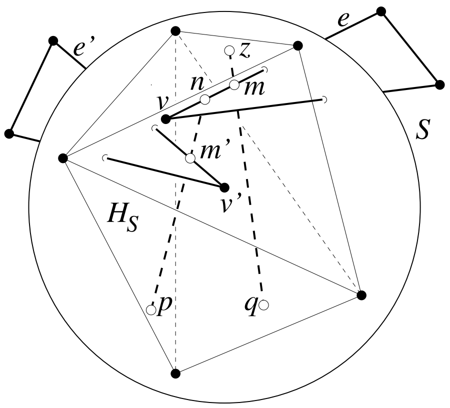

Figure 15: Because qz intersects a strongly Delaunay simplex e at m, and m is visible from p, some vertex inside S is visible from p.

Proof of Theorem 5: Consider the first claim. Suppose, for the sake of contradiction, that some j-facet f (for any $j < d$) intersects the relative interior of s_{k+1} but does not contain s_{k+1}.

Let s be a j-simplex of the Delaunay triangulation of f that intersects the relative interior of s_{k+1}. Because s_k is constrained Delaunay, it follows that if s contains any point in the interior of s_k, then s contains s_k entirely. By assumption, f does not contain s_{k+1}, so s lacks at least one vertex of s_{k+1}. Hence, s either does not contain u or does not contain any point in the relative interior of s_k.

There exists a simplex t, which is s or a face of s, that intersects the relative interior of s_{k+1} and contains neither u nor any point in the relative interior of s_k. How do we know this? If s contains u, then s does not contain any point in the relative interior of s_k; let t be the $(j-1)$-face of s that does not have u for a vertex. Because s intersects the relative interior of s_{k+1}, so must t. On the other hand, if s contains s_k, then s does not contain u; let t be any $(j-1)$-face of s that lacks one vertex of s_k and intersects the relative interior of s_{k+1}. (There must be at least one such face between u and the relative interior of $s \cap s_{k+1}$.) Otherwise, let $t = s$. In each of these cases, t contains neither u nor any point in the relative interior of s_k.

Let r be a point in the relative interior of s_k such that ru intersects t. Note that t contains neither r nor u. Let q be the point nearest p on the line segment pr such that qu intersects a constraining simplex or facet of dimension $d-2$ or less that contains neither q nor u, as Figure 16 illustrates. To see that such a choice of q exists, imagine sliding q along pr from p (which can see u) toward r, stopping when the line segment qu first intersects an appropriate simplex. If t is a constraining simplex of dimension $d-2$ or less, then q will stop sliding no later than when it reaches r. Alternatively, if t is not such a simplex, then f must be a $(d-1)$-facet (as illustrated), with t lying within it. Because p can see u, the line segment pu does not intersect f. Hence, the 2-simplex $\blacktriangle pru$ must intersect a boundary of f, so q will stop sliding no later than when qu meets this boundary.

Let m be the intersection point of qu and the constraining simplex, as illustrated. The point m is visible from p, because the relative interior of $\blacktriangle pqu$ intersects no obstructing facet (otherwise, q would have stopped earlier). Lemma 7 holds that there is a vertex of X inside S that is visible from p. This contradicts the assump-

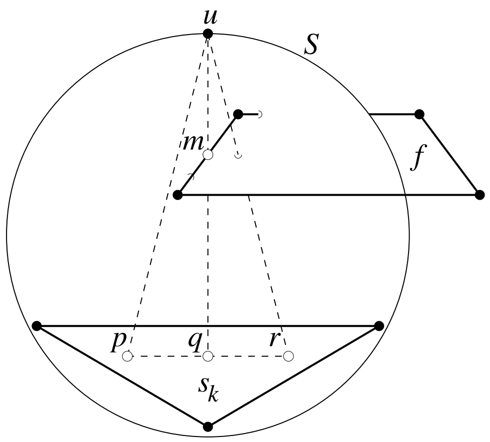

Figure 16: The circumstance depicted here, wherein p can see u but r cannot, could not occur if the vertices of f were on or outside S and the edges of f were strongly Delaunay.

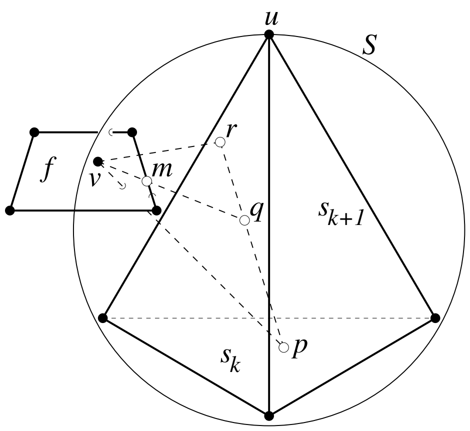

Figure 17: The circumstance depicted here, wherein p cannot see v but r can, could not occur if the vertices of f were on or outside S and the edges of f were strongly Delaunay.

tion that no vertex inside S is visible from p; hence, no constraining facet of X intersects the relative interior of s_{k+1} unless it contains s_{k+1}.

Next, consider the claim that the $(k+1)$-simplex s_{k+1} formed by s_k and u is constrained Delaunay. Suppose, for the sake of contradiction, that there is some point r in the interior of s_{k+1} that can see a vertex v of X inside S, as Figure 17 illustrates. By assumption, v is not visible from p.

Let f be the $(d-1)$-facet that obstructs the visibility of v from p; if there are several such facets, choose f to be the facet that intersects pv closest to p.

Let q be the point nearest p on the line segment pr such that qv intersects the boundary of a $(d-1)$-facet that contains neither q nor v, and the intersection point m is visible from p. To see that such a choice of q exists, imagine sliding q along pr from p (whose view of v is obstructed by f) toward r (from which v is visible). As q moves, either qv will strike the boundary of a $(d-1)$-facet lying between q and f, or the intersection of qv and f will remain visible from p until qv intersects a boundary of f (which will occur before q reaches r).

Lemma 7 holds that there is a vertex of X inside S that is visible from p. By contradiction, s_{k+1} is constrained Delaunay. ■

Theorem 5 shows that a constrained Delaunay simplex can be "grown" into a higher-dimensional constrained Delaunay simplex. However, a separate result is needed to show that a constraining $(d-1)$-simplex of X—which might not be constrained Delaunay—can also be "grown" into a constrained Delaunay d-simplex.

Theorem 8 *Let s be a constraining $(d-1)$-simplex of X. Let u be a vertex of X that is above s and visible from some point p in the interior of s. Suppose that there is a sphere S that passes through u and all d vertices of s, and encloses no vertex of X that is above s and visible from a point in the interior of s. Let s_d be the convex hull of s and u. Then no constraining facet of X intersects the interior of s_d, and s_d is constrained Delaunay.*

Proof: The proof is similar to that of Theorem 5, with a few modifications to account for the fact that there may be vertices of X inside S below s. Consider the first claim. Suppose, for the sake of contradiction, that some facet intersects the interior of s_d.

Let p' be a point on the line segment pu that is extremely close to p. Specifically, p' is such a small distance above s that p' cannot see any vertex of X below s (because they are all occluded by s). Furthermore, for any vertex w of X above s, the line wp' passes through the interior of s. Such a choice of p' is always possible because p is in the relative interior of s. Additionally, p' is close enough to p that the convex hull of s and p' intersects no facet above s. Such a choice of p' is always possible because no facet of X can intersect the interior of s (except the facet that contains s).

Because s is a Delaunay triangle of a ridge-protected facet, no vertex of X coplanar with s is inside S. Hence, if there is a vertex inside S that is visible from p', it lies above s. Observe that p' lies in H_S and can see u.

By repeating the argument from the proof of Theorem 5, with p replaced by p', one may deduce that there is a vertex w of X inside S that is visible from p'; w must be above s. Because the line wp' passes through the relative interior of s, w is visible from some point in the relative interior of s, a contradiction. Hence, no facet intersects the interior of s_d.

Next, consider the claim that the d-simplex s_d formed from s and u is constrained Delaunay. Suppose, for the sake of contradiction, that there is some point r in the interior of s_d from which a vertex v of X inside S is visible.

If v is below s, let $p'' = p'$. If v is above s, let p'' be a point on the line segment pv that is extremely close to p, satisfying all the same conditions listed above for p' as well as the condition that p'' is in the interior of s_d. In either case, p'' lies in the interiors of s_d and H_S, and can see r but not v.

By repeating the argument from the proof of Theorem 5, with p replaced by p'', one may deduce that there is a vertex w of X inside S that is visible from p''. Because the line wp'' passes through the relative interior of s, w is visible from some point in the interior of s, a contradiction. Hence, s_d is constrained Delaunay. ■

3 Algorithms for Constructing CDTs

Starting with one constrained Delaunay d-simplex, a naïve implementation of the foregoing gift-wrapping algorithm might construct each additional d-simplex by considering all of the vertices as candidates and testing the visibility of each one against every facet. The running time is thus $\mathcal{O}(d^4 n_s n_v n_f)$, where n_s is the number of d-simplices in the output, n_v is the number of input vertices, and n_f is the number of $(d-1)$-simplices in the triangulations of the input $(d-1)$-facets. (The factor of d^4 is the cost of testing whether a $(d-1)$-simplex obstructs the visibility between two points.) This leaves much room for improvement, which I hope will be filled by future research. For instance, the sophisticated search algorithm and analysis techniques applied by Dwyer [2] to unconstrained gift-wrapping might be generalizable to the constrained case as well.

Here, I offer a few practical suggestions without asymptotic guarantees.

For most applications, the fastest way to form the CDT of a ridge-protected PLC (albeit not in the worst case) is to use the best available algorithm to find an unconstrained Delaunay triangulation of the input vertices, then recover the $(d-1)$-facets one by one. Each $(d-1)$-facet f may be recovered by deleting the d-simplices whose interiors it intersects, then retriangulating the polytopes now left empty on each side of f. It is easy to show that all of the simplices not thus deleted are still constrained Delaunay. Since a CDT of the new configuration exists, each empty polytope can be triangulated with constrained Delaunay simplices. If these polytopes are typically small, the performance of the algorithm used to triangulate them is not critical, and gift-wrapping will suffice.

Vertices may be incrementally inserted into and deleted from a CDT just like an ordinary Delaunay triangulation, so long as the underlying PLC (which changes incrementally with the triangulation) remains ridge-protected. When a vertex is inserted, the simplices that are no longer constrained Delaunay are deleted. When a vertex is deleted, the simplices that contain it are deleted. In either case, the resulting polytopal hole is retriangulated to complete the new CDT. As with facet recovery, the existence of a CDT of the entire underlying PLC ensures that a CDT of the hole can be produced. Hence, the best approach to triangulating a PLC might be to start with a Delaunay triangulation of the vertices of the $(d-1)$-facets, then recover the $(d-1)$-facets themselves, and then finally insert the remaining vertices incrementally.

The ability to incrementally insert and delete vertices is also useful for mesh generation, especially in circumstances where the constrained Delaunay property can be used to establish provable properties of the meshing algorithm [13].

Unfortunately, subsets of $d+2$ or more cospherical vertices can cause real difficulties for gift-wrapping. A gift-wrapping algorithm may make decisions that are mutually inconsistent, and find itself unable to complete the triangulation. For an example affecting *unconstrained* Delaunay tetrahedralizations in E^3, imagine a large vertex set that includes six cospherical vertices whose surroundings have been inadvertently tetrahedralized so as to form a hollow space shaped like Schönhardt's polyhedron.

This problem can be solved by using symbolic perturbation to simulate general position, thereby ensuring that all decisions made by gift-wrapping are mutually consistent. Additionally, cospherical or nearly-cospherical vertices create the need for exact arithmetic when performing the *insphere* tests associated with Delaunay triangulation.

4 Conclusions

In their paper on two-dimensional conforming Delaunay triangulations, Edelsbrunner and Tan [4] write:

> A seemingly difficult open problem is the generalization of our polynomial bound to three dimensions. The somewhat easier version of the generalized problem considers a graph whose vertices are embedded as points in $\mathbb{R}^3$, and edges are represented by straight line segments connecting embedded vertices. More relevant, however, is the problem for the crossing-free embedding of a complex consisting of vertices, edges, *and* triangles.

The present result shifts the emphasis back to the former of these two problems. An algorithm that could create a CCDT by inserting only a modest number of additional vertices on input segments might have great practical importance.

Several other questions also deserve investigation. Are there higher-dimensional constrained Delaunay triangulation algorithms that have the same running time as optimal algorithms for unconstrained Delaunay triangulations? Do higher-dimensional CDTs have optimality properties such as minimizing the largest min-containment sphere, as higher-dimensional Delaunay triangulations do [9]? Is there a less conservative definition of "constrained Delaunay" (perhaps allowing visibility to be affected by constraining simplices of dimension less than $d-1$) that defines triangulations over a larger class of PLCs?

Acknowledgments

I thank Dafna Talmor for her many helpful comments, and in particular for suggesting an intuitive way to recognize whether a simplex is strongly Delaunay.

References

[1] L. Paul Chew. *Constrained Delaunay Triangulations*. Algorithmica **4**(1):97–108, 1989.

[2] Rex A. Dwyer. *Higher-Dimensional Voronoi Diagrams in Linear Expected Time*. Discrete & Computational Geometry **6**(4):343–367, 1991.

[3] Herbert Edelsbrunner. *An Acyclicity Theorem for Cell Complexes in d Dimension*. Combinatorica **10**(3):251–260, 1990.

[4] Herbert Edelsbrunner and Tiow Seng Tan. *An Upper Bound for Conforming Delaunay Triangulations*. Discrete & Computational Geometry **10**(2):197–213, 1993.

[5] Steven Fortune. *Voronoi Diagrams and Delaunay Triangulations*. Computing in Euclidean Geometry (Ding-Zhu Du and Frank Hwang, editors), Lecture Notes Series on Computing, volume 1, pages 193–233. World Scientific, Singapore, 1992.

[6] Carol Hazlewood. *Approximating Constrained Tetrahedrizations*. Computer Aided Geometric Design **10**:67–87, 1993.

[7] D. T. Lee and A. K. Lin. *Generalized Delaunay Triangulations for Planar Graphs*. Discrete & Computational Geometry **1**:201–217, 1986.

[8] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, Noel Walkington, and Han Wang. *Control Volume Meshes using Sphere Packing: Generation, Refinement and Coarsening*. Fifth International Meshing Roundtable (Pittsburgh, Pennsylvania), pages 47–61, October 1996.

[9] V. T. Rajan. *Optimality of the Delaunay Triangulation in $\mathbb{R}^d$*. Proceedings of the Seventh Annual Symposium on Computational Geometry, pages 357–363. ACM, 1991.

[10] James Martin Ruppert. *Results on Triangulation and High Quality Mesh Generation*. Ph.D. thesis, University of California at Berkeley, Berkeley, California, 1992.

[11] Jim Ruppert and Raimund Seidel. *On the Difficulty of Triangulating Three-Dimensional Nonconvex Polyhedra*. Discrete & Computational Geometry **7**(3):227–254, 1992.

[12] E. Schönhardt. *Über die Zerlegung von Dreieckspolyedern in Tetraeder*. Mathematische Annalen **98**:309–312, 1928.

[13] Jonathan Richard Shewchuk. *Delaunay Refinement Mesh Generation*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1997. Available as Technical Report CMU-CS-97-137.

[14] Nigel P. Weatherill and O. Hassan. *Efficient Three-Dimensional Grid Generation using the Delaunay Triangulation*. Proceedings of the First European Computational Fluid Dynamics Conference (Brussels, Belgium) (Ch. Hirsch, J. Périaux, and W. Kordulla, editors), pages 961–968, September 1992.

[15] Nigel P. Weatherill, O. Hassan, D. L. Marcum, and M. J. Marchant. *Grid Generation by the Delaunay Triangulation*. Von Karman Institute for Fluid Dynamics 1993–1994 Lecture Series, 1994.