General-Dimensional Constrained Delaunay and Constrained Regular Triangulations, I: Combinatorial Properties

Jonathan Richard Shewchuk

July 9, 2007

Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, CA 94720

Abstract

Two-dimensional constrained Delaunay triangulations are geometric structures that are popular for interpolation and mesh generation because they respect the shapes of planar domains, they have "nicely shaped" triangles that optimize several criteria, and they are easy to construct and update. The present work generalizes constrained Delaunay triangulations (CDTs) to higher dimensions and describes constrained variants of regular triangulations, here christened *weighted CDTs* and *constrained regular triangulations*. CDTs and weighted CDTs are powerful and practical models of geometric domains, especially in two and three dimensions.

The main contributions are rigorous, theory-tested definitions of constrained Delaunay triangulations and piecewise linear complexes (geometric domains that incorporate nonconvex faces with "internal" boundaries), a characterization of the combinatorial properties of CDTs and weighted CDTs (including a generalization of the Delaunay Lemma), the proof of several optimality properties of CDTs when they are used for piecewise linear interpolation, and a simple and useful condition that guarantees that a domain has a CDT. These results provide foundations for reasoning about CDTs and proving the correctness of algorithms. Later articles in this series discuss algorithms for constructing and updating CDTs.

Supported in part by the National Science Foundation under Awards CMS-9318163, ACI-9875170, CMS-9980063, CCR-0204377, CCF-0430065, and EIA-9802069, 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 an Alfred P. Sloan Research Fellowship, and in part by gifts from the Okawa Foundation and Intel. The views in this document are those of the author. They are not endorsed by the sponsors or the U. S. Government.

Keywords: discrete geometry, constrained Delaunay triangulation, constrained regular triangulation, weighted Delaunay triangulation, piecewise linear complex

Contents

1 Introduction

Many geometric applications can benefit from triangulations that have properties similar to Delaunay triangulations, but are constrained to contain specified edges or faces. Delaunay triangulations have virtues when they are used to interpolate multivariate functions [13, 28, 37, 50], including a tendency to favor "round" simplices over "skinny" ones. However, some applications rely on the presence of faces that represent specified discontinuities, as illustrated in Figure 1, and the Delaunay triangulation might not respect these constraints. Triangulations also serve as meshes that represent objects for rendering or for the numerical solution of partial differential equations. For these purposes, Delaunay triangulations have many advantages, but the triangulations are required to assume the shapes of the objects being modeled, and perhaps to resolve interfaces where different materials meet or where boundary conditions are applied.

In two dimensions, there are two popular alternatives for creating a Delaunay-like triangulation that respects constraints. In either case, the input is a *planar straight line graph* (PSLG), such as the one illustrated in Figure 2 (left). A PSLG X is a set of vertices and segments (constraining edges) that satisfy two restrictions: both endpoints of every segment in X are members of X, and a segment in X may intersect other segments and vertices in X only at its endpoints. A triangulation is sought that contains the vertices in X and respects the segments in X.

The first alternative is to form a *conforming Delaunay triangulation* (Figure 2, second from right). 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 [20] show that a PSLG X can be triangulated with the addition of $O(m^2n)$ augmenting vertices, where

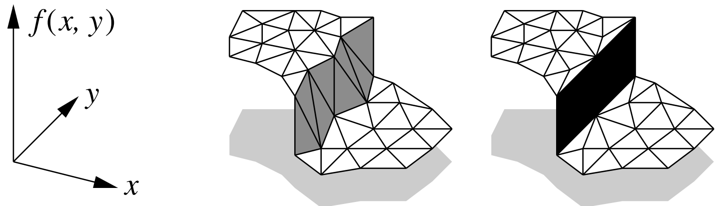

Figure 1: A triangulation that respects a discontinuity in a function (right) can be a better interpolating surface than one that does not (left).

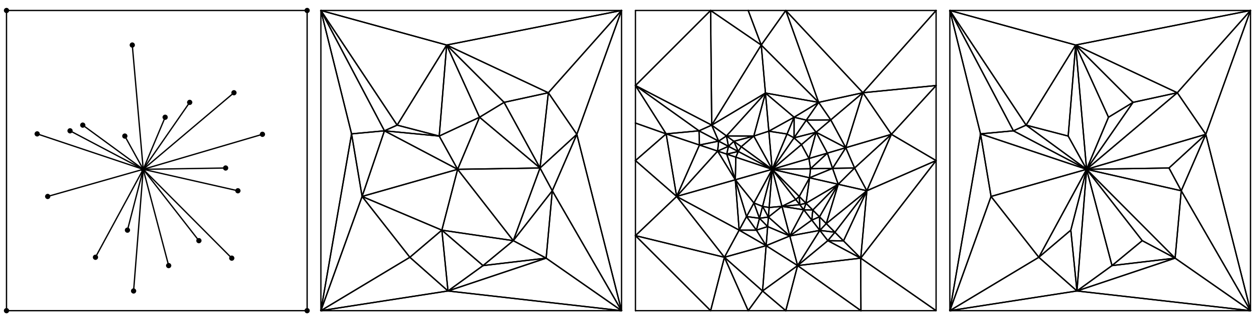

Figure 2: The Delaunay triangulation (second from left) of the vertices of a PSLG (far left) might not respect the segments of the PSLG. These segments can be incorporated by adding vertices to obtain a conforming Delaunay triangulation (second from right), or by forgoing Delaunay triangles in favor of constrained Delaunay triangles (far right).

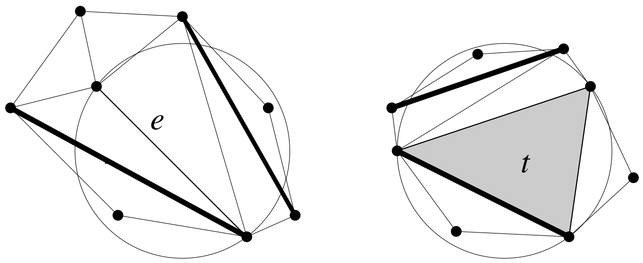

Figure 3: The edge e and the triangle t are both constrained Delaunay. Bold lines represent segments.

m is the number of segments in X, and n is the number of vertices in X. For many PSLGs, their algorithm uses far fewer augmenting vertices, but the numbers required in practice are often undesirably large. PSLGs are known that have no conforming Delaunay triangulation with fewer than $\Theta(mn)$ augmenting vertices. Closing the gap between the $O(m^2n)$ and $\Omega(mn)$ bounds remains an open problem.

The second alternative is to form a *constrained Delaunay triangulation* (CDT) [29, 9, 43], illustrated at far right in Figure 2. A CDT of X has no vertices not in X, and every segment in X is a single edge of the CDT. However, a CDT, despite its name, is not a Delaunay triangulation. In an ordinary Delaunay triangulation, every simplex (triangle, edge, or vertex) 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 in X. (Any number of vertices is permitted on the circle.) 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 in X that is *visible* from any point in the relative interior of the simplex—here visibility is occluded only by segments in X—and furthermore, the simplex does not "cross" any segment. (For a formal definition, see Section 1.1.)

Figure 3 demonstrates examples of a constrained Delaunay edge e and a constrained Delaunay triangle t. Segments in X appear as bold lines. Although there is no empty circle that encloses 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 sole circumcircle of t encloses 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 vertex other than those in X. The advantage of a conforming Delaunay triangulation is that its triangles are Delaunay, whereas those of a CDT are not. Nevertheless, CDTs retain 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 [29].

We live in a three-dimensional world, and those who model it have a natural interest in constructing constrained and conforming triangulations in three or more dimensions. Algorithms by Murphy, Mount, and Gable [33], Cohen-Steiner, Colin de Verdière, and Yvinec [12], Cheng and Poon [8], and Pav and Walkington [34] can construct a conforming Delaunay tetrahedralization of any three-dimensional polyhedron by inserting carefully chosen vertices on the boundary of the polyhedron. (Their algorithms work not only on polyhedra, but also on a more general input called a *piecewise linear complex*, defined below.) These algorithms might introduce a huge number of new vertices. No known algorithm for finding conforming Delaunay tetrahedralizations is guaranteed to introduce only a polynomial number of new vertices, and no algorithm of any complexity has been offered for four- or higher-dimensional conforming Delaunay triangulations.

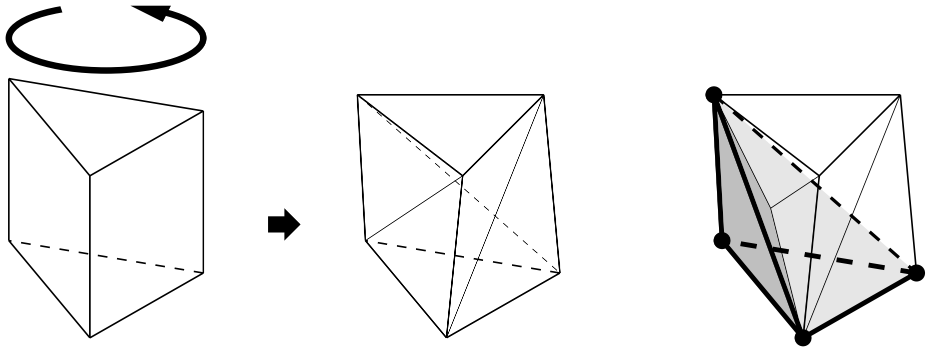

Figure 4: Schönhardt's untetrahedralizable polyhedron (center) is formed by rotating one end of a triangular prism (left), thereby creating three diagonal reflex edges. Every tetrahedron defined on the vertices of Schönhardt's polyhedron sticks out (right).

Prior to the present work (in its first incarnation [45]), CDTs had 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 [41] furnishes a three-dimensional example depicted in Figure 4 (center). The easiest way to envision this polyhedron is to begin with a triangular prism (Figure 4, left). Imagine twisting the prism so that the top triangular face rotates slightly like the lid of a jar, while the bottom triangular face is fixed in place. 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 (formerly) square face are separated by a reflex edge, and the line segment connecting them is outside 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 does not lie entirely within the polyhedron, as illustrated at right in Figure 4. Schönhardt's polyhedron cannot be tetrahedralized without an additional vertex. (One extra vertex at its center will do.)

Ruppert and Seidel [40] add to the difficulty by proving that it is NP-hard to determine whether a three-dimensional polyhedron is tetrahedralizable. Even among polyhedra that can be triangulated without additional vertices, there is not always a triangulation that is in any reasonable sense "constrained Delaunay."

What features of polytopes make them amenable to being triangulated with Delaunay-like simplices? This article offers a partial answer by proposing a conservative extension of the definition of CDT to higher dimensions, and by demonstrating that there is an easily tested and enforced, sufficient (but not necessary) condition that guarantees that a CDT exists. This article also shows that CDTs optimize several criteria for the accuracy of piecewise linear interpolation of certain classes of functions. These results extend to *weighted CDTs* (a constrained generalization of regular triangulations), wherein each vertex is assigned a numerical weight that influences the triangulation.

There is more than one way in which the notion of "constrained Delaunay" might generalize to three or more dimensions. The choices made here yield useful CDTs and efficient algorithms for their construction, though other generalizations of CDTs might be discovered in the future.

This article is the first in a three-part series. The second article discusses sweep and gift-wrapping algorithms for constructing the CDT of any piecewise linear complex that has one, except for a class of difficult, "nongeneric" inputs. It also demonstrates the NP-completeness of determining whether a nongeneric polyhedron has a CDT. The third article discusses algorithms for updating a CDT to reflect the insertion or deletion of a $(d-1)$-facade, a vertex, or several vertices, as well as an incremental algorithm for constructing CDTs that have a property called "ridge protection" (described in the next section).

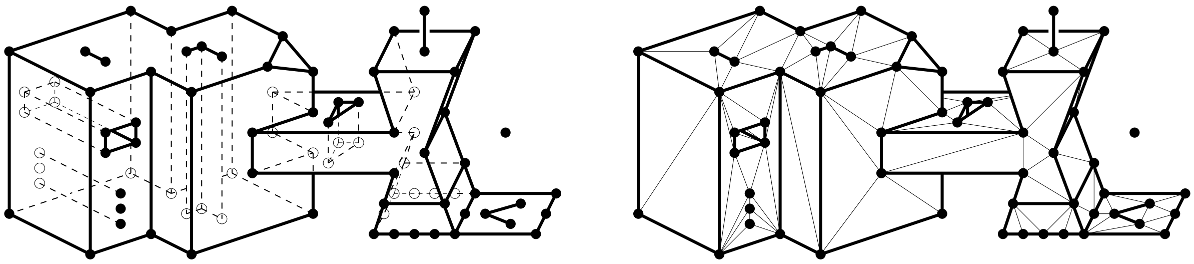

Figure 5: Each facade of a PLC (left) may have holes, slits, and interior vertices, which are used to constrain a triangulation or to support intersections with other facades. The illustration at right is the constrained Delaunay triangulation of the PLC at left. It is a PLC, too.

1.1 Summary of Results

The input is a *piecewise linear complex* (PLC), following Miller, Talmor, Teng, Walkington, and Wang [32].[1] A PLC is a finite set of *facades* in an ambient space E^d. A facade is a polytope (roughly speaking) of any dimension from zero to d, possibly with holes and lower-dimensional facades inside it. Figure 5 illustrates a three-dimensional PLC. As the figure shows, a facade may have any number of sides, may be nonconvex, and may have holes, slits, or vertices inside it. A *k-facade* is a k-dimensional facade. Zero-facades are vertices, and 1-facades are segments. Observe that a PSLG is a two-dimensional PLC without 2-facades.

PLCs have restrictions similar to those of PSLGs or any other type of complex. For each facade f in a PLC X, the boundary of f must be composed of lower-dimensional facades in X. Nonempty intersections of facades in X must be facades in X. For details, see Section 2.1, where the terms *facade* and *PLC* are defined with full mathematical rigor.

The purpose of most facades is to constrain a triangulation. A d-dimensional PLC typically includes d-facades, whose purpose is to specify what region the triangulation should fill. The union of all the facades in a PLC X is the *triangulation domain* $|X|$, the portion of space a user wishes to triangulate. The specification of a triangulation domain is sometimes crucial, because there are PLCs for which a CDT of the triangulation domain exists but a CDT of its convex hull does not. For example, it is easy to tetrahedralize the region sandwiched between Schönhardt's polyhedron and a suitable bounding box, even though the interior of the polyhedron is not tetrahedralizable.

The complement of the triangulation domain, $E^d \backslash |X|$, is called the *exterior domain* and includes any hollow cavities enclosed by the triangulation domain, as well as outer space. Because X is a complex, some of its $(d-1)$-facades separate the interior of the triangulation domain from the exterior domain. But not all $(d-1)$-facades play this role. Figure 5 includes several *dangling* lower-dimensional facades that are not part of any d-facade. Some facades are *internal facades*, which do not lie on the boundary of the exterior domain. These facades allow PLCs to represent multiple-component domains and domains with nonmanifold boundaries.

The goal of this work is to subdivide a domain into simplices. A *k-simplex* is a k-dimensional simplex—the convex hull of $k+1$ affinely independent points. A *triangulation* or *simplicial complex* T is a finite set of simplices that intersect each other "nicely": T contains every face of every simplex in T, and the intersection of any two simplices in T is either empty or a face of both simplices. A triangulation T *fills* a PLC X if $\bigcup_{t \in T} t = \bigcup_{f \in X} f$; that is, if the union of simplices in T is the triangulation domain $|X|$.

[1]Miller et al. call it a *piecewise linear system*, but their construction is so obviously a complex that a change in name seems obligatory.

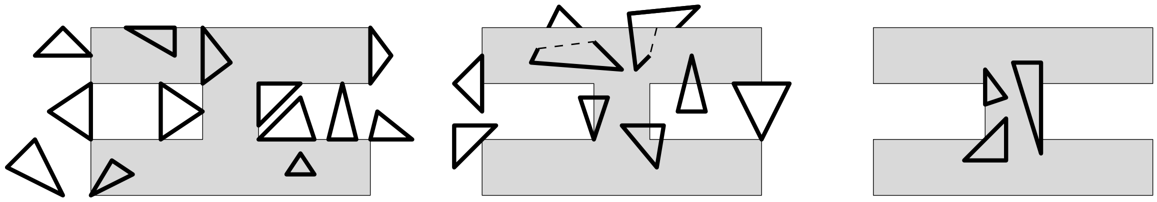

Figure 6: Left: examples of triangles that respect a shaded facade. Center: examples of triangles that do not respect the facade. Right: examples of triangles that respect the facade, but do not respect all its edges and vertices (which are facades themselves).

Of course, not all triangulations that fill X are equally good. Facades constrain what sort of simplex is acceptable. A simplex s *respects* a facade f if $s \cap f$ is a union of faces of s (possibly empty). As Figure 6 illustrates, the intersection of a nonconvex facade and a triangle that respects it might be the empty set, a vertex, an edge, the entire triangle, the union of two or all three edges, the union of two or all three vertices, or the union of an edge and opposite vertex of the triangle. Loosely speaking, if s respects f, then s cannot "cross" f or f's boundary.

A simplex s *respects* X if $s \subseteq |X|$ and s respects every facade in X, except perhaps some of the vertices. (Weighted CDTs may omit some of the vertices in X, unlike ordinary CDTs. But some designated vertices must be respected; see Section 2.3 for details.)

A triangulation T is a *triangulation of* X if T fills X, every simplex in T respects X, and every vertex in T is in X. This definition implies that every facade in X (except perhaps the vertices) is a union of simplices in T. (See Section 2.3 for a discussion of why the definition does not explicitly require every vertex in X to be in T. This requirement arises implicitly if every vertex is designated as one that must be respected.)

Sometimes it is desirable to permit a triangulation to have vertices not present in X—and sometimes it is necessary, as Schönhardt demonstrates. A triangulation T is a *conforming triangulation* or *Steiner triangulation* of X if T fills X and every simplex in T respects X. This article is devoted to pure triangulations in which extra vertices are not permitted, but Steiner triangulations are investigated elsewhere [47, 49].

Within a PLC X, the visibility between two points p and q is *occluded* if $pq \nsubseteq |X|$ or there is a facade between p and q whose affine hull contains neither p nor q. (Note, however, that some vertices do not obstruct visibility—namely those that the triangulation is not required to respect. See Section 2.4.) The points p and q are *visible* from each other if $pq \subseteq |X|$ and X contains no occluding facade.

Let s be a k-simplex (for any k) whose vertices are in X (though s is not necessarily a facade in X). Let S be a (full-dimensional) hypersphere 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 exists a circumsphere S of s that encloses no vertex in X. The simplex s is *strongly Delaunay* if there exists a circumsphere S of s such that no vertex in X lies inside *or on* S, except the vertices of s. Every 0-simplex (vertex) is trivially strongly Delaunay.

A simplex s is *constrained Delaunay* if

- the vertices of s are in X,
- s respects X, 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.

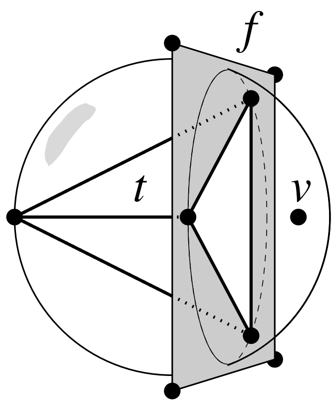

Figure 7: A constrained Delaunay tetrahedron t.

Figure 7 depicts a constrained Delaunay tetrahedron t in E^3. The intersection of t with the facade f is a face of t, so t respects X. The circumsphere of t encloses one vertex v, but v is not visible from any point in the interior of t, because f occludes its visibility.

A *constrained Delaunay triangulation* T of X is a triangulation of X in which every d-simplex is constrained Delaunay. If X has dangling facades, this characterization is insufficient, and we must resort to the (less readable) true definition: a CDT T of X is a triangulation of X in which each simplex is constrained Delaunay "within" the lowest-dimensional facade in X that includes it. For example, a three-dimensional CDT fills each 2-facade (dangling or not) with triangles that are constrained Delaunay "within" that 2-facade, and collectively comprise a two-dimensional CDT of the 2-facade. However, those triangles might not be constrained Delaunay within the three-dimensional PLC—they might have empty circumcircles, but not empty circumspheres. (That is why the 2-facade is there—to enforce the presence of triangles that might otherwise be absent.) The definition of "CDT" is therefore recursive in the dimension. See Section 2.4 for details.

The first main result of this article is a characterization of many basic properties of constrained Delaunay and weighted constrained Delaunay triangulations, analogous to the well-known properties of Delaunay triangulations. (Weighted CDTs are defined in Section 2.4.) For example, every face of a constrained Delaunay simplex is itself constrained Delaunay within some facade. A CDT of a facade includes CDTs of all the facade's faces (Section 3.1). If a PLC has no $d + 2$ vertices lying on a common hypersphere, then its constrained Delaunay simplices have disjoint relative interiors and form a simplicial complex, and it has at most one CDT (Section 3.3). The Delaunay Lemma, which guarantees that a triangulation of a vertex set is Delaunay if and only if its facets are locally Delaunay [14], generalizes to CDTs (Section 3.2). The Delaunay Lemma is a fundamental tool for verifying that a triangulation is a CDT, and for dynamically maintaining the CDT of a PLC whose vertices are moving or changing their weights.

The second main result is that CDTs are optimal by several criteria (described in Section 4) when they are used for piecewise linear interpolation. This fact is among the reasons why CDTs are so valuable.

The third main result is a condition that guarantees the existence of a CDT. The main impediment to the existence of CDTs is the difficulty of respecting facades of dimension $d - 2$ or less. Figure 8 offers an example of a three-dimensional PLC with no CDT. There is one segment that runs through the interior of the PLC. There is only one tetrahedralization of this PLC—composed of three tetrahedra encircling the central segment—and its tetrahedra are not constrained Delaunay, because each of them has a visible vertex inside its circumsphere. If the central segment were removed, the PLC would have a CDT made up of two tetrahedra.

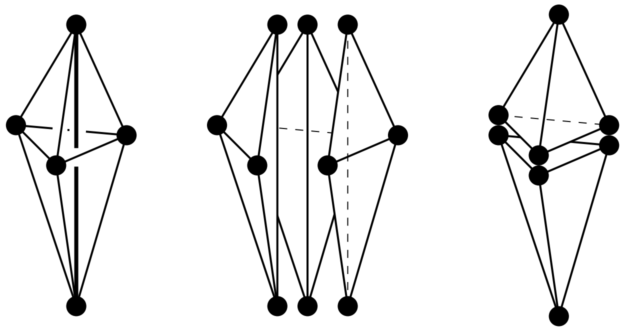

Figure 8: Left: a PLC with no CDT. Center: the sole tetrahedralization of this PLC. Its three tetrahedra are not constrained Delaunay. Right: the two Delaunay tetrahedra do not respect the central segment.

The condition that guarantees that a PLC has a CDT is easiest to describe, and easiest to enforce, in three dimensions. A three-dimensional PLC X is *ridge-protected* if every segment (1-facade) in X is strongly Delaunay. (See Section 2.4 for the general-dimensional definition.) Every ridge-protected PLC has a CDT. This result, called the *CDT Theorem*, makes three-dimensional CDTs useful in geometric modeling applications.

It is not sufficient for every segment to be Delaunay. If Schönhardt's polyhedron is embedded so that all six of its vertices lie on a common sphere, then all of its edges (and its triangular 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).

Here is a stronger and even more useful version of the CDT Theorem. In three dimensions a segment may serve as a boundary to several 2-facades, which can be sorted by their rotary order around the segment. A segment is *grazeable* if two consecutive 2-facades in the rotary order are separated by an interior angle of 180° or more, or if the segment is included in fewer than two 2-facades and is internal, not dangling. (An *interior angle* subtends the interior of the triangulation domain. Exterior angles of 180° or more are irrelevant to the CDT Theorem.) Only the grazeable segments need to be strongly Delaunay to guarantee a CDT. A three-dimensional PLC X is *weakly ridge-protected* if every grazeable segment in X is strongly Delaunay. Every weakly ridge-protected PLC has a CDT.

Segments that are not grazeable are common. For instance, in a complex of convex polyhedra, no segment is grazeable. The stronger result exempts the segments of the complex from the need to be strongly Delaunay.

Testing whether a PLC is ridge-protected, or weakly ridge-protected, is straightforward. See the comments following Definition 23.

This article's results extend to weighted CDTs, which are described in Section 2.4. Weighted CDTs are central in the design of flip algorithms for updating and constructing CDTs; see the third article in this series. Several researchers have shown that weighted Delaunay triangulations are useful for three-dimensional mesh generation, because some undesirable tetrahedra can be removed by adjusting the vertex weights [6, 7, 16]. Weighted CDTs share this virtue and are even more powerful, because of the ease with which they respect the shape of a domain.

The definition of "ridge-protected" generalizes to weighted PLCs, and every weakly ridge-protected, weighted PLC has a weighted CDT. Interestingly, even in two dimensions there are weighted PLCs that do not have weighted CDTs.

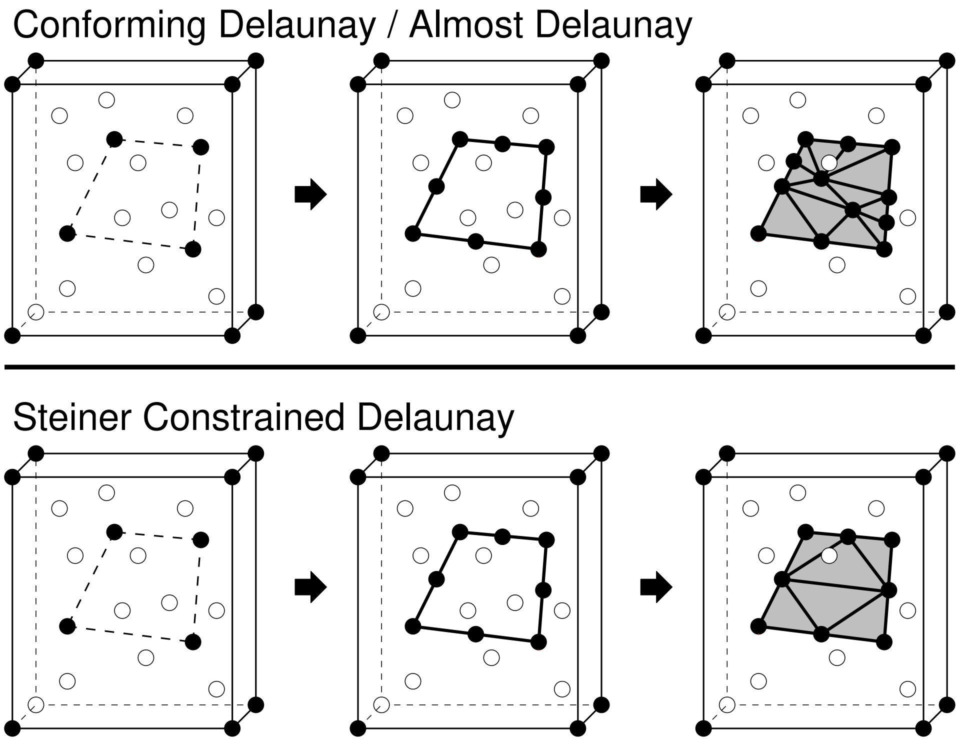

Figure 9: Two methods for recovering a 2-facade in the interior of a cubical triangulation domain. The initial Delaunay tetrahedralization does not respect the facade. (For clarity, the tetrahedra are not shown.) Both methods insert new vertices to recover missing segments. Next, the customary method is to insert more vertices to recover missing 2-facades (top), but no additional vertices are needed if constrained Delaunay tetrahedra are used (bottom).

1.2 Benefits of the CDT Theorem

Why is it useful to know that weakly ridge-protected PLCs have CDTs? Although a given PLC X might not be weakly ridge-protected, the insertion of additional vertices can transform it into a weakly (or fully) 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 a *conforming CDT* or *Steiner CDT* of X: "conforming" or "Steiner" because boundary conformity is obtained by inserting new vertices (Steiner points), and "CDT" because the simplices of the Steiner CDT are constrained Delaunay (rather than Delaunay).

Compare this idea to the most common methods of recovering missing facades in three-dimensional Delaunay-based mesh generation algorithms, which insert additional vertices into all the missing facades. Some of these algorithms produce conforming Delaunay meshes [8, 34, 39], and some recover the missing facades by bisecting and flipping tetrahedra, yielding a mesh that is not necessarily Delaunay nor constrained Delaunay, although you might say it is "almost" Delaunay [22, 27, 52, 53]. Figure 9 illustrates the advantage of a Steiner CDT. All the procedures use vertex insertions to recover missing grazeable segments, but the customary approaches require additional vertex insertions to recover missing 2-facades and non-grazeable segments. A Steiner CDT does not need these extra vertices.

Figure 10 (left) depicts an example of a PLC for which a Steiner CDT is much more effective than a conforming Delaunay tetrahedralization. In the interior of the box, many oddly shaped 2-facades adjoin a single shared segment. The triangulation domain is the entire box. Vertices inserted to recover one 2-facade—so that it is a union of triangular faces of the Delaunay tetrahedralization—are likely to knock out

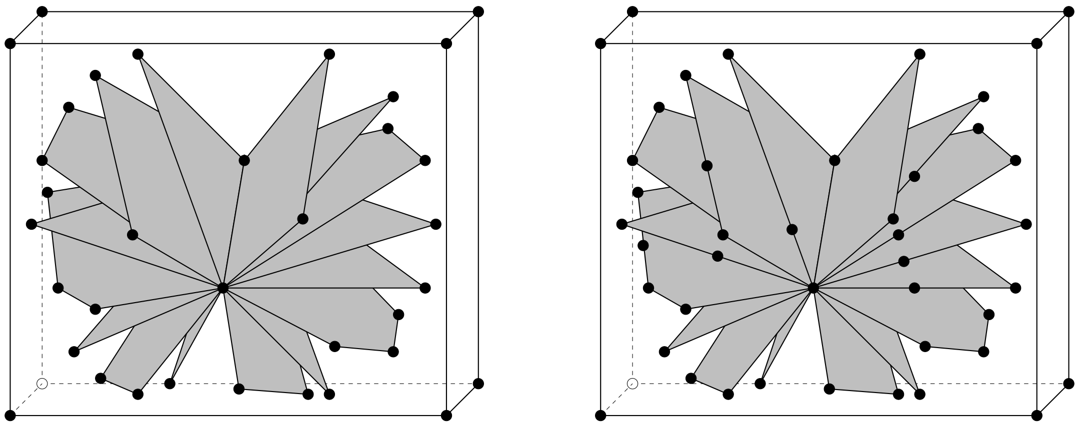

Figure 10: Left: it is difficult to mesh the interior of this box with Delaunay tetrahedra that conform to all the facades. Right: the box can be meshed with constrained Delaunay tetrahedra with the addition of just the vertices shown.

triangles from the adjacent 2-facades. The aforementioned algorithms of Murphy et al. and others [12, 8, 33, 34] can construct conforming Delaunay tetrahedralizations of this PLC, but they require many more vertices than are needed to form a Steiner CDT, most of them in the 2-facade interiors. The PLC augmented with a modest number of vertices (Figure 10, right) is weakly ridge-protected and has a CDT.

I conjecture that for the worst three-dimensional PLCs, conforming Delaunay triangulations need asymptotically more vertices than Steiner CDTs. It is an open question whether this is true, but based on the two-dimensional complexity results, it seems like a safe gamble.

An algorithm that decides how to choose new vertices so that there are provably good bounds on the edge lengths of the Steiner CDT (i.e. edges are not made unnecessarily short) is described elsewhere [47]. This algorithm does not guarantee a polynomial bound on the number of new vertices, but its guarantees on edge lengths are in some ways more useful, because the Steiner CDT is an excellent starting triangulation for several algorithms for three-dimensional mesh generation. One algorithm uses the constrained Delaunay property to guarantee its ability to tetrahedralize any PLC [46], and another uses it to establish provable bounds on the quality of the tetrahedra it produces and on the edge lengths of the final mesh [44]. The results in this article underpin those algorithms.

Why does this article take piecewise linear complexes as the input rather than, for simplicity, boundary triangulations? Consider finding a tetrahedralization of a cube. The edges of the cube are strongly Delaunay, so the CDT Theorem guarantees that the cube has a CDT. By contrast, consider a boundary triangulation of a cube. Any boundary triangulation bisects each square face of the cube with a diagonal edge. These diagonals are not strongly Delaunay, so the CDT Theorem does not apply. Moreover, a tetrahedralization respecting the boundary triangulation might not exist (depending on the choice of diagonals). Thus, the option to specify facades more general than simplices is an advantage both for the theorem and for CDT construction algorithms, which can choose a compatible set of diagonals.

If a PLC is ridge-protected, its CDT can be built by a simple incremental facade insertion algorithm described in the third article in this series. PLCs that are not ridge-protected (but have CDTs) currently require a more complicated sweep algorithm or a slower gift-wrapping algorithm, described in the second article in this series.

2 Complexes

This section defines the geometric constructions and ideas at the center of this work. The input structures—facades and piecewise linear complexes—are formalized in Section 2.1. The output structures, a generalization of CDTs called *weighted CDTs*, are described in Sections 2.2–2.4. Definitions are often a perfunctory part of a mathematics article, so it is worth noting that eight years of trial and error led to the definitions given here. "Constrained Delaunay" and the notion of visibility are defined differently here than in the earlier incarnation of this work [45], and the present definitions are more sound. These and other definitions in this article evolved with the proofs of the theorems here and in the sequel articles.

Throughout this article, the terms "simplex," "triangle," "tetrahedron," and "convex hull" refer to closed, convex sets of points; for instance, a "triangle" is not just three edges, but the points inside as well. The notation $\mathrm{conv}(S)$ represents the convex hull of the point set S.

Some simplices of specific dimensions have their own names. Of course, a vertex is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex. In a d-dimensional ambient space, a $(d-2)$-dimensional convex polytope or $(d-2)$-simplex is called a *ridge*, and a $(d-1)$-dimensional convex polytope or $(d-1)$-simplex is called a *facet*.

The notation pq denotes a line segment with endpoints p and q. The notation $p \cdot q$ denotes the Euclidean inner product, $|p| = \sqrt{p \cdot p}$ is the Euclidean norm, and $|pq| = |p - q|$ is the Euclidean length of pq. It might help the reader to know that this article strictly distinguishes between the verbs *contain* for set membership ($\ni$) and *include* for set inclusion ($\supseteq$).

2.1 Piecewise Linear Complexes

Consider points in an ambient space E^d. A *k-flat* (k-dimensional flat) is the affine hull of $k+1$ affinely independent points. (A flat is also known as an *affine subspace*—unlike a true subspace it is not required to contain the origin. For readers familiar with flats but not affine hulls, the *affine hull* of a point set is the lowest-dimensional flat that includes it.) A set of points $S \subseteq E^d$ is *k-dimensional* if the affine hull of S is a k-flat. (In other words, S contains $k+1$ affinely independent points, but does not contain $k+2$ affinely independent points.) A *hyperplane* is a $(d-1)$-flat. The set of points on one designated side of a hyperplane, excluding every point of the hyperplane itself, is an *open halfspace*. By contrast, a *closed halfspace* includes the hyperplane as well.

An *open convex k-polyhedron* is the nonempty intersection of a k-flat and a finite number of open halfspaces. It is *bounded* if it does not include a ray (equivalently, if its diameter is finite). A *closed convex k-polyhedron* is the closure of an open convex k-polyhedron. The *closure* of a polyhedron has its usual meaning from real analysis—the set of all the points and accumulation points of the polyhedron—and more intuitively is a point set containing all the points of the polyhedron, plus all the points on its boundary.

Definition 1 (facade)**.** *An* open k-facade *is the union of a finite number of bounded, open, convex k-polyhedra, all included in some common k-flat. A* closed k-facade *is the closure of an open k-facade.*

Observe that a facade is not required to be connected. A 0-facade (open or closed—there is no difference) is a vertex, and a 1-facade is either a segment or a sequence of collinear segments.

A closed facade is equivalent to Hadwiger's classic polyhedron [26], which is defined to be a union of closed convex polyhedra. It is the open facades that motivate the new name. In geometric modeling, open facades are more versatile than closed facades as abstractions of geometric domains and their boundaries,

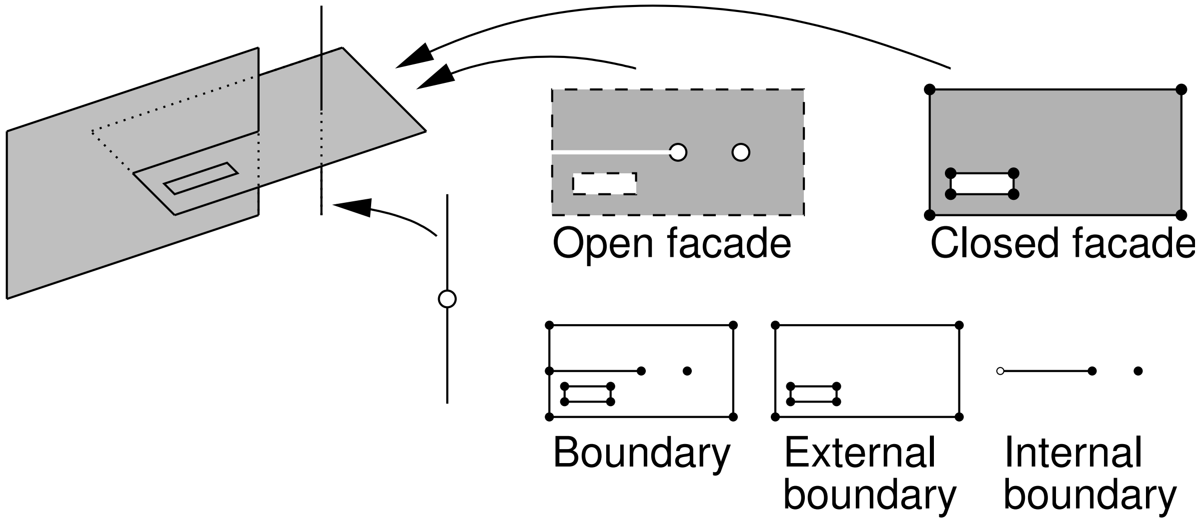

Figure 11: At left are two connected 2-facades and a 1-facade (composed of two segments). At center and right appear one of the 2-facades, represented as both an open and a closed facade. Dashed lines and open circles represent points that are not part of the open facade. The internal boundary includes a slit and an isolated vertex, both of which are needed to support contacts with other facades. The internal boundary cannot be inferred from the closed facade alone.

because an open facade can have internal boundaries. Internal boundaries serve at least two purposes: they support intersections between surfaces, as Figure 11 illustrates, and they constrain the permissible triangulations of the facade—for instance, to support the application of boundary conditions to a finite element mesh, or to model discontinuities in the lighting of a surface for computer graphics. Internal boundaries are necessary to model some domains with nonmanifold boundaries, like the domain in Figure 5.

Definition 2 (external and internal boundaries)**.** *The* external boundary *of a facade is the boundary of the closure of the facade. (Observe that the external boundary includes boundaries of holes.) The* internal boundary *of an open facade is the boundary of the open facade minus the external boundary. (Equivalently, it is the intersection of the boundary with the relative interior of the closure of the facade.) See Figure 11.*

Throughout this article, *relative interior* has its usual meaning from real analysis, but *boundary* is used as shorthand for *relative boundary*, and *open* for *relatively open*.

The faces of a facade are defined in a fundamentally different way than the faces of a convex polyhedron. The faces of a convex polyhedron are an intrinsic property of the polyhedron, whereas the faces of a facade are defined only in the context of a piecewise linear complex. Compare the following two definitions.

Definition 3 (face of a convex polyhedron)**.** *The* faces *of a closed, convex k-polyhedron P are P and every polyhedron found by taking the intersection of P with a hyperplane that does not intersect the relative interior of P. The* proper faces *of P are the faces of dimensionalities zero through $k-1$.*

This standard construction also defines the faces of a simplex. For example, the faces of a tetrahedron include its four vertices, its six edges, its four triangular faces, and the tetrahedron itself. By convention, the empty set is considered to be a (-1)-dimensional face of every polyhedron. This article makes no use of this convention, but in some circumstances it is convenient to assume that $\emptyset$ is a member of every nonempty piecewise linear complex and triangulation.

Piecewise linear complexes and the faces of a facade are defined in a way that gives a geometric model the power to constrain how the boundary of a facade can be triangulated.

Definition 4 (piecewise linear complex; face of a facade)**.** *An* open piecewise linear complex *(PLC) X is a set containing a finite number of open facades that satisfy the following two restrictions.*

- *For every facade $f \in X$, the boundary of f is a union of facades in X.*[2] *For example, X contains both endpoints of every segment in X, and every 2-facade's boundary is a union of segments and vertices in X.*
- *For any two facades $f, g \in X$, $f \cap g = \emptyset$.*

The faces *of a facade f are $\{g \in X : g \subseteq \text{closure}(f)\}$. They include f itself and its vertices. The* proper faces *of f are all its faces except f and $\emptyset$.*

For any open piecewise linear complex X, $\{\text{closure}(f) : f \in X\}$ is a closed piecewise linear complex.

It is possible to reverse the transformation and convert a closed piecewise linear complex into an open piecewise linear complex by subtracting from each facade every facade of lower dimension. Hence, for a closed facade in a closed PLC, define the *internal boundary* of the closed facade to be the internal boundary of the corresponding open facade. The internal boundary of a closed facade is not really part of the boundary of the closed facade, and it is defined only in the context of a PLC.

Definition 5 (triangulation domain)**.** *Let $|X|$ denote the union of facades $\bigcup_{f \in X} f$. $|X|$ is called the* triangulation domain*, or simply the* domain*. (It is also known as the* underlying space of *X.)*

A corollary of the definition of PLC is that $\bigcup_{f \in X} f$ is the same for an open PLC and the corresponding closed PLC. Another corollary is that a closed PLC Y satisfies the restrictions that Miller et al. [32] specified when they introduced the notion of a PLC.

- For every facade $f \in Y$, the boundary of f is a union of facades in Y.
- For any two facades $f, g \in Y$, $f \cap g$ is a union of facades in Y. (Usually $f \cap g$ is a single facade or the empty set, but imagine two nonconvex 2-facades that intersect each other at several isolated vertices and along several line segments. Each of these vertices and line segments must be in Y.)
- For any two facades $f, g \in Y$, if $f \cap g$ has the same dimensionality as f, then $f \subset g$, and f is of lower dimensionality than g.

Miller's third restriction is somewhat cryptic; its main effect is to prevent two facades of the same dimensionality from having overlapping relative interiors. The formulation of PLCs in terms of open facades is more elegant, because no similarly cryptic restriction is needed. However, closed PLCs offer a more elegant model for the incremental update of a PLC (discussed in the third article of this series). The insertion or deletion of a facade in a closed PLC can imply several modifications to the corresponding open PLC. For instance, when a vertex is added to an open PLC, if a facade contains the vertex, the facade must have that point removed.

This formal hair-splitting between open and closed facades is necessary because it is the open facades that determine the facade boundaries, but it is the closed facades that occlude visibility, and simplices must

[2]The boundary of a vertex is the empty set, which is a union of zero facades.

respect the closed facades. The rest of this article maintains an uneasy duality, wherein every use of the word "facade" refers to both the open and the closed versions of the facade. Fortunately, the bijective map between open and closed PLCs usually makes it unnecessary to specify which type of PLC is under discussion.

The reader should be aware that every reference to the "boundary of a facade" or the "faces of a facade" regards the boundary of the open facade, including the internal boundary. Similarly, the "relative interior of a facade" refers to the open facade. However, wherever this article states that a facade contains a point, a facade obstructs visibility, or a simplex respects a facade, the closed facade is implied.

It makes no difference to most of this article's results whether or not facades are connected. An open facade made up of n connected components can be replaced with n separate facades without changing any essential properties of the PLC. Some components of a facade may be grazeable while others are not, so breaking up a facade into its components may improve the prospects for having a weakly ridge-protected PLC. However, there is an important convention for weighted CDTs. If a vertex in a PLC is an endpoint of two collinear segments and is not needed to support their intersection with some other facade, it is usually better to think of the two segments as parts of a single 1-facade, because the vertex might be absent from the weighted CDT. (See Definition 12 in Section 2.3 for details.)

There appear to be few publications exploring the properties of geometric partitions that permit the existence of faces with internal boundaries. An interesting exception by Grunbaum and Shephard [25] shows how to reliably compute Euler characteristics for a class of objects more general than PLCs. One can convert an open PLC into a "relatively open convex dissection" by partitioning its open facades into open convex facades (polyhedra), whereupon its Euler characteristic is easy to calculate. This method is particularly interesting when applied to an open facade with a complicated internal boundary, or to a subset of an open PLC that allows faces of facades to be absent.

The notion of a PLC generalizes to complexes of curved manifolds. For example, every semialgebraic or subanalytic set of points can be partitioned into a *stratification*—a set of *strata* (which generalize open facades), each of which is a manifold. See Gomes [23] for an excellent introduction to the topic.

How might a PLC be represented as a data structure? Here are a few suggestions. A 0-facade (vertex) is represented by its d coordinates. For $j \geq 1$, a j-facade f is most easily represented by a list of its proper faces. To conserve space, f can be represented by a list of every proper face of f that is not a proper face of a proper face of f; the unlisted faces can be inferred by reading the listed faces' lists. This representation differs in several ways from the usual face lattice representation of polyhedra and polyhedral complexes. First, the faces in f's list are not necessarily all $(j-1)$-faces, because f's internal boundary may include lower-dimensional faces that are not included in any $(j-1)$-face. For example, a 2-facade may have an isolated vertex inside it. Second, this representation is technically not a lattice. For example, two 2-facades might intersect at two separate vertices that are included in no other facades, so a pair of facades do not necessarily have a unique meet and join, contrary to the definition of "lattice." See Ziegler [54] for a definition and discussion of face lattices.

Within the affine hull of a j-facade f, each $(j-1)$-face of f has two sides. In an implementation of the sweep algorithm or gift-wrapping algorithm for CDT construction, the list of f's $(j-1)$-faces should include annotations that indicate which side (or sides) of each $(j-1)$-face adjoins f. A $(j-1)$-face on f's internal boundary adjoins f on both sides. If an open $(j-1)$-face is composed of several connected components, it needs one annotation for each side of each connected component.

For the algorithms described in the sequel articles, it is unnecessary to specify the d-facades explicitly as part of the input. Instead, each side of each $(d-1)$-facade should bear an annotation that indicates whether it adjoins the exterior domain or the interior of the triangulation domain. A $(d-1)$-facade is part of the internal boundary of a PLC if both sides adjoin the triangulation domain, part of the external boundary if one side adjoins the exterior domain, and a dangling facade if both sides adjoin the exterior domain.

Definition 6 (dangling facade)**.** *Let X be a d-dimensional PLC. A facade in X is a* dangling facade *if it is not a face of any d-facade in X.*

To a programmer, the distinction between open and closed facades is almost irrelevant. Any reasonable PLC data structure simultaneously represents both.

2.2 Weighted Delaunay Triangulations

This section reviews known facts about weighted Delaunay triangulations [2] and introduces new terminology as a preliminary to introducing weighted CDTs in Section 2.4. Consider the Euclidean space E^{d+1}, and let $x_1, x_2, \ldots, x_{d+1}$ be the coordinate axes. E^d is the subspace of E^{d+1} orthogonal to the x_{d+1}-axis. In the space E^{d+1}, a d-flat is *vertical* if it includes a line parallel to the x_{d+1}-axis.

Definition 7 (polyhedral complex; triangulation)**.** *A* polyhedral complex *T is a set containing a finite number of closed, convex polyhedra that satisfy the following two restrictions.*

- *For every polyhedron $s \in T$, every face (in the sense of Definition 3) of s is in T.*
- *For any two polyhedra $s, t \in T$, if s and t are not disjoint, then $s \cap t$ is a face of both s and t.*

A triangulation *or* simplicial complex *is a polyhedral complex whose members are all simplices.*

Every polyhedral complex is a PLC. Observe that polyhedral complexes are less general than PLCs whose facades are all convex, because they use a different definition of "face." In a PLC, one side of a tetrahedron might be subdivided into several triangular faces, and a (closed) tetrahedron might have an edge passing through its interior. In a polyhedral or simplicial complex, both circumstances are forbidden: a side of a tetrahedron is represented by exactly one triangular face, and a tetrahedron's interior intersects no other simplex of equal or lesser dimension.

A d-dimensional triangulation or polyhedral complex is *regular* if it is the vertical projection of one "side" of some convex $(d+1)$-polyhedron.

Definition 8 (downward-facing; underside; regular)**.** *Let P be a convex $(d+1)$-polyhedron in E^{d+1}. A face f of P is* downward-facing *if no point in P is directly below any point in f (i.e. having the same x_1- through x_d-coordinates but a lesser x_{d+1}-coordinate). The* underside *of P is the set of all its downward-facing faces.*

A d-dimensional triangulation or polyhedral complex is regular *if it can be formed by vertically projecting the underside of some convex $(d+1)$-polyhedron P into E^d (by dropping the x_{d+1}-coordinate of each vertex).*

The best-known regular triangulation is the Delaunay triangulation. The regularity of most Delaunay triangulations is demonstrated by the well-known *parabolic lifting map* of Seidel [42, 18] (inspired by a spherical lifting map suggested by Brown [5]). Let V be a set of vertices in E^d for which a Delaunay triangulation is sought. The lifting map maps each vertex in V to a vertex on a paraboloid in a space one dimension higher, as Figure 12 illustrates. Specifically, each vertex $v = (v_{x_1}, v_{x_2}, \ldots, v_{x_d}) \in V$ maps to a point $v^+ = (v_{x_1}, v_{x_2}, \ldots, v_{x_d}, v_{x_1}^2 + v_{x_2}^2 + \cdots + v_{x_d}^2)$ in E^{d+1}.

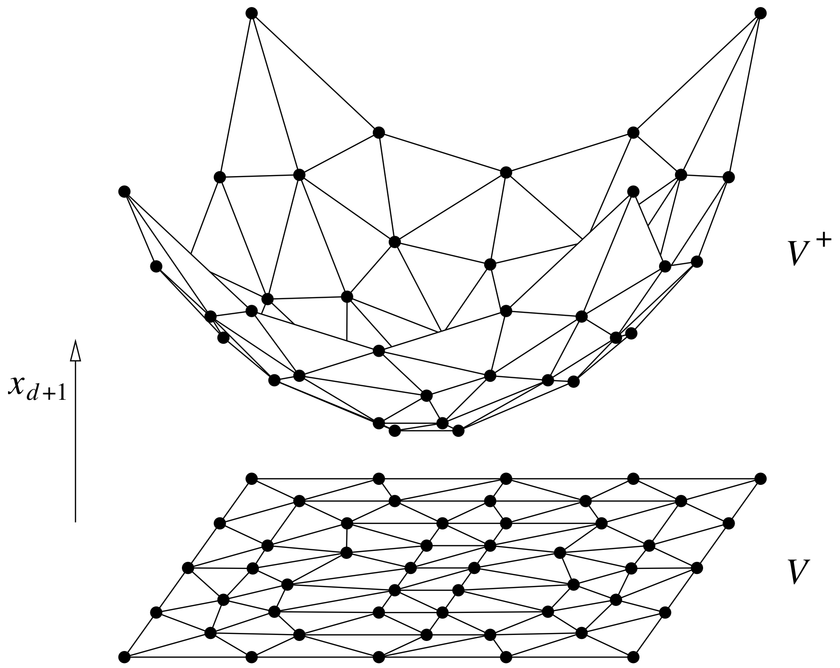

Figure 12: The parabolic lifting map. In this illustration a two-dimensional vertex set V is lifted to a paraboloid in E^3. The underside of the convex hull of the lifted vertices is a lifted Delaunay triangulation.

Definition 9 (companion)**.** *The pair of vertices v and v^+ are called* companions: *v^+ is the* lifted companion *of v, and v is the* projected companion *of v^+.*

If s is a k-simplex with vertices $v_0, v_1, \ldots, v_k$, then its lifted companion s^+ is the k-simplex embedded in E^{d+1} whose vertices are $v_0^+, v_1^+, \ldots, v_k^+$; and s is the projected companion of s^+. Note that s^+ is flat, and does not curve to hug the paraboloid.

Let $V^+ = \{v^+ : v \in V\}$. The Delaunay triangulation of V is regular because it has the same combinatorial structure as the underside of the convex hull of V^+, as the forthcoming Theorem 2 shows. Each downward-facing simplex of $\mathrm{conv}(V^+)$ projects to a Delaunay simplex of V. This connection is routinely used to transform any $(d+1)$-dimensional convex hull construction algorithm into a d-dimensional Delaunay triangulation construction algorithm.

Lemma 1. *Let S be a hypersphere in E^d. Let $S^+ = \{p^+ : p \in S\}$ be the ellipsoid found by lifting S to the paraboloid. Then the points of S^+ lie on a non-vertical d-flat h. (Recall that a d-flat is vertical if it is parallel to the x_{d+1}-axis.) Furthermore, a point p inside S lifts to a point p^+ below h, and a point p outside S lifts to a point p^+ above h. Therefore, testing whether a point p is inside, on, or outside S is equivalent to testing whether the lifted point p^+ is below, on, or above h.*

Proof. Let O and r be the center and radius of S, respectively. Let p be a point in E^d. The x_{d+1}-coordinate of p^+ is $|p|^2$. By expanding $|O - p|^2$, we have that $|p|^2 = 2O \cdot p - |O|^2 + |Op|^2$.

With O and r fixed and $x \in E^d$ varying, the equation $x_{d+1} = 2O \cdot x - |O|^2 + r^2$ defines a non-vertical d-flat h in E^{d+1}. For every point $p \in S$, $|Op| = r$, so $S^+ \subset h$. For every point $p \notin S$, if $|Op| < r$, then the lifted point p^+ lies below h, and if $|Op| > r$, then p^+ lies above h. ■

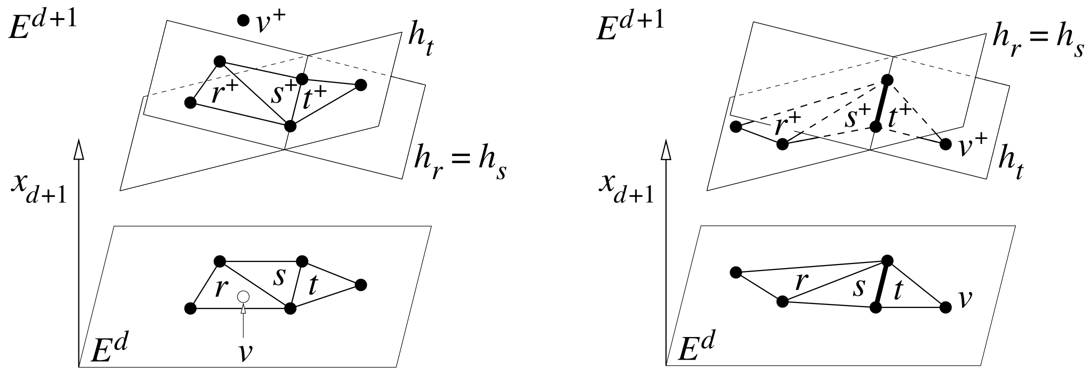

(a) (b)

Figure 13: (a) The triangles r, s, and t are all semiregular, but only t is regular. Triangles r and s have the same witness d-flat $h_r = h_s$, and t has a different witness h_t. The vertex v is submerged. (b) The bold edge is a constraining segment. The triangles r, s, and t are all constrained semiregular, but only t is constrained regular. No triangle is semiregular.

Theorem 2 (Seidel [42]). *Let s be a simplex whose vertices are in V, and let s^+ be its lifted companion. Then s is Delaunay if and only if s^+ is included in some face of the underside of* conv(V^+)*. The simplex s is strongly Delaunay if and only if s^+ is a face of the underside of* conv(V^+) *and no vertex in V^+ lies on s^+ except the vertices of s^+.*

Proof. If s is Delaunay, there is a circumsphere S of s such that no vertex of V lies inside S. Let h be the unique d-flat in E^{d+1} that includes S^+. By Lemma 1, no vertex in V^+ lies below h. The d-flat h includes s^+ because the vertices of s^+ are in S^+. Therefore, s^+ is included in a downward-facing face of the convex hull of V^+. If s is strongly Delaunay, no vertex in V^+ lies below h, and no vertex in V^+ lies on h except the vertices of s^+. Therefore, s^+ is a downward-facing face of the convex hull of V^+.

The converse implications follow by reversing the argument. ■

A *weighted Delaunay triangulation* is like a Delaunay triangulation, but each vertex $v \in V$ is assigned a real-valued *weight* w_v. A vertex v lifts to a companion $v^+ = (v_{x_1}, v_{x_2}, \ldots, v_{x_d}, v_{x_1}^2 + v_{x_2}^2 + \cdots + v_{x_d}^2 - w_v)$. The x_{d+1}-coordinate $|v|^2 - w_v$ is called the *height* of v. The weighted Delaunay triangulation of V is the projection to E^d of the underside of conv(V^+). It follows that a weighted Delaunay triangulation is regular.

Some faces of conv(V^+) might not be simplices, because some selection of $d + 2$ or more of the lifted vertices might lie on a common non-vertical d-flat. (Observe that vertices that lie on a common vertical d-flat do not cause trouble, because a vertical face cannot be downward-facing. This is good news, because a typical real-world vertex set V includes large groups of cohyperplanar vertices.) These non-simplicial faces can be filled with any compatible triangulation, so V has more than one weighted Delaunay triangulation. However, some faces can be triangulated with triangulations that are not regular, so not all weighted (or unweighted) Delaunay triangulations are regular! Section 6 describes a simple way to perturb the weights to simulate the circumstance that no $d + 2$ vertices in V^+ lie on a common non-vertical d-flat.

If its weight is sufficiently small, a lifted vertex v^+ might not be downward-facing—it might not lie on the underside of conv(V^+)—in which case the vertex v is absent from the weighted Delaunay triangulation of V, as illustrated in Figure 13(a). Then v is said to be *submerged*. If every vertex has a weight of zero, the weighted Delaunay triangulation is the Delaunay triangulation, and no vertex is submerged, because every point on the paraboloid is on the underside of the convex hull of the paraboloid.

Weights necessitate a generalization of the notion of a "Delaunay simplex."

Definition 10 (semiregular; witness; weighted Delaunay triangulation)**.** *A simplex s whose vertices are in V is* semiregular *if s^+ is included in a downward-facing face of* conv(V^+). *In other words, there exists a non-vertical d-flat $h_s \subset E^{d+1}$ such that h_s includes s^+, and no vertex in V^+ lies below h_s. The d-flat h_s is called a* witness[3] *to the semiregularity of s.*

A weighted Delaunay triangulation *of V is a simplicial complex that fills* conv(V) *wherein every simplex is semiregular.*

Figure 13(a) illustrates three semiregular triangles and their witnesses. All their edges and vertices are semiregular as well, but the submerged vertex v is not semiregular.

Definition 11 (regular)**.** *A simplex s is* regular *if s^+ is a downward-facing face of* conv(V^+)*, and no vertex in V^+ lies on s^+ except the vertices of s^+. In other words, there exists a non-vertical d-flat $h_s \subset E^{d+1}$ that is a* witness *to the regularity of s: h_s includes s^+, and every vertex in V^+ lies above h_s, except the vertices of s^+.*

A triangulation is regular *if there exists an assignment of weights to its vertices for which every simplex is regular.*

Of the three triangles in Figure 13(a), only t is regular. All the edges are regular except the edge shared by r and s. All the vertices are regular except v.

In a weighted Delaunay triangulation, a witness serves the same purpose that a circumsphere serves in an ordinary Delaunay triangulation. Theorem 2 shows that if all the weights are zero, "semiregular" is equivalent to "Delaunay" and "regular" is equivalent to "strongly Delaunay." If a simplex s is semiregular, it appears in at least one weighted Delaunay triangulation of V. If s is regular, it appears in *every* weighted Delaunay triangulation of V (see Theorem 19).

2.3 Triangulations of Piecewise Linear Complexes

For some geometric applications, the notion of a "constrained triangulation" of a PLC should permit some vertices to be left out, just as weighted Delaunay triangulations submerge vertices with insufficient weight. However, some vertices cannot be omitted, because they support other facades. The following definition identifies vertices that could conceivably be submerged.

Definition 12 (submersible)**.** *A vertex v in a closed PLC X is* submersible *if v is a proper face of some other facade (i.e. v is not isolated), and the removal of v from X (and possibly the merging of two collinear 1-facades) yields a valid closed PLC. Equivalently, either*

- *v lies on the internal boundary of a facade $f \in X$ such that f is a face of every facade (except v) that contains v, or*
- *v is an endpoint of two collinear 1-facades in X, and the condition above is satisfied by merging them into a single 1-facade f. In this case, X should be modified to reflect the merger. A row of collinear segments might comprise one 1-facade with many submersible vertices in it.*

The user of a PLC triangulation algorithm can arbitrarily designate vertices as being non-submersible, but a vertex can be designated as submersible only if Definition 12 permits it.

[3] A witness for a semiregular or regular simplex is also known as a *supporting hyperplane* of conv(V^+), but a witness for a *constrained* semiregular simplex is not necessarily a supporting hyperplane.

Definition 13 (fill; respect; triangulation of a piecewise linear complex)**.** *Let T be a set of simplices. T* fills *X if $|T| = |X|$, meaning that $\bigcup_{s \in T} s = \bigcup_{f \in X} f$.*

Let f be a closed facade. Let s be a simplex or convex polyhedron. Then s respects *f if $s \cap f$ is a union of faces of s.*

There is an equivalent definition that is less clear, but easier to use in proofs: s respects f if, for every face t of s whose relative interior intersects f, $t \subseteq f$.

If f is an open facade, s is said to respect f if s respects the closure of f.

A simplex (or convex polyhedron) s respects *a PLC X if $s \subseteq |X|$ and s respects every facade in X except perhaps the submersible vertices—after agglomerating the segments of X into 1-facades as described in Definition 12.*

A triangulation T respects a PLC X if every simplex in T respects X.

A triangulation T is a triangulation of a PLC *X if T fills and respects X, and T has no vertex not in X. A triangulation that fills and respects X, but may have vertices not present in X, is a* conforming triangulation *or* Steiner triangulation *of X.*

This definition allows a triangulation T of X to submerge vertices in X. However, submersibility is a nuisance when it is not needed. For some applications, such as unweighted PLCs and ordinary CDTs (in which vertices are never submerged), it does no harm to designate every vertex in X as non-submersible. Then Definition 13 implicitly requires that if T is a triangulation of X, then T and X have exactly the same vertices, because T must respect every vertex in X.

Why must adjoining collinear segments be agglomerated for Definition 13? If a vertex is submerged, then a triangulation lacking that vertex cannot respect a segment that terminates at that vertex, but it can respect a 1-facade that passes through the vertex.

2.4 Weighted Constrained Delaunay Triangulations

Before considering the formal definition of CDT, let us try to see intuitively what a CDT is, in terms of the parabolic lifting map. Suppose T is a CDT of a PLC X. Let $T^+ = \{s^+ : s \in T\}$ be the simplicial complex, embedded in E^{d+1}, defined by lifting T. As Figure 14 illustrates, the lifted triangulation T^+ graphs a continuous piecewise linear function but, in general, is not the underside of a convex polyhedron: each facet of the CDT that is not constrained Delaunay is mapped to a reflex ridge in the lifted surface. (A $(d-1)$-simplex is called a *facet* if it exists in the ambient space E^d, and a *ridge* if it exists in the ambient space E^{d+1}.)

However, from any point p in the interior of a d-facade, the portions of the CDT visible from p appear convex on the lifting map. Only facets included in $(d-1)$-facades can lift to reflex ridges; every other facet is constrained Delaunay.

The next several definitions build toward the definition of a CDT or, more generally, a weighted CDT, which is a triangulation of a weighted PLC.

Definition 14 (weighted piecewise linear complex)**.** *A* weighted PLC *is a PLC in which each vertex is assigned a real-valued weight.*

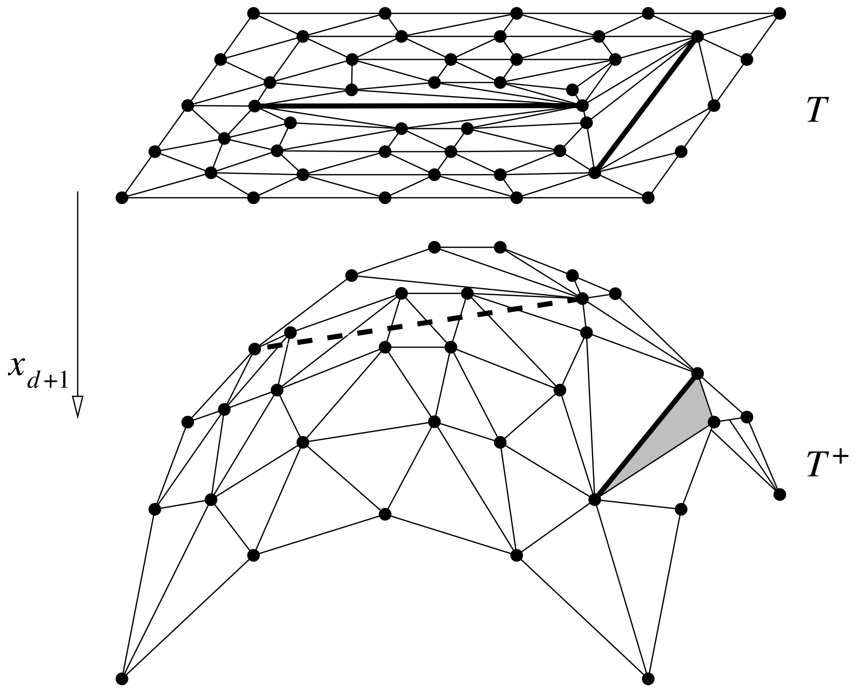

Figure 14: A lifted CDT. The paraboloid is inverted to more clearly show its topography. The bold edges are constraining edges that are not Delaunay. They map to reflex edges of the lifted surface.

Sections 3.1 and 3.3 study the relationship between the weighted CDT of a weighted PLC and the weighted CDTs of its facades. Consider computing a triangulation of a 2-dimensional PLC. Some algorithms need to "triangulate" the 1-facades of the PLC first—in other words, to decide which vertices on the 1-facades are submerged. The 1-facades may have both submersible and non-submersible vertices. A 1-facade in isolation does not reveal which of its vertices are submersible in the 2-dimensional PLC. Therefore, it is best to think of submersibility as a global property of a vertex which remains fixed across all contexts, and is determined by the highest-dimensional PLC that contains the vertex. These observations motivate the following two policies. First, the internal boundary of a 1-facade may contain both submersible and non-submersible vertices (whereas the external boundary is a set of non-submersible vertices). Second, non-submersible vertices occlude visibility and submersible vertices do not. This policy ensures that the weighted CDT of a 1-facade is consistent with the weighted CDT of any higher-dimensional facade that includes the 1-facade.

Visibility is occluded by constraining facades.

Definition 15 (constraining facade)**.** *A* constraining facade *in a d-dimensional PLC X is any facade in X that is not a submersible vertex or a d-facade.*

Definition 15 omits submersible vertices because they do not occlude visibility or constrain the triangulation. It omits d-facades because they do not occlude visibility, and because a simplex or polyhedron that respects all the lower-dimensional facades automatically respects the d-facades.

Definition 16 (occlusion; visibility)**.** *Within a PLC X, the visibility between two points p and q is* occluded *if $pq \nsubseteq |X|$; or if there is a (closed) constraining facade $f \in X$ such that the line segment pq intersects f, and neither p nor q lie on the affine hull of f. See Figure 15. The points p and q are* visible *from each other (equivalently, can* see *each other) if $pq \subseteq |X|$ and X places no constraining facade between them.*

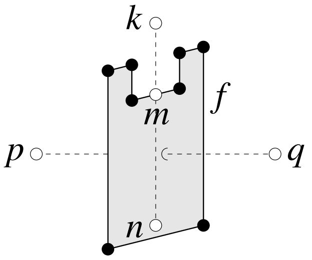

Figure 15: In this three-dimensional example the 2-facade f occludes the visibility between p and q. The point m can see both k and n, but the visibility between k and n is occluded—not by f, but by an edge of f.

If no vertex is submersible, a more elegant characterization is that p and q can see each other if there is an open facade $f \in X$ that includes the open line segment pq. Open facades thus act as conductors of visibility. In this interpretation, the d-facades play an essential role.[4]

There is a close relationship between visibility and the notion of respecting a PLC.

Theorem 3. *If a (closed) simplex or convex polyhedron s respects X, every point in s can see every other point in s.*

Proof. Suppose for the sake of contradiction that the visibility between two points $p, q \in s$ is occluded by some facade f. Then pq intersects f at a point m, but f contains neither p nor q. Let t be the face of s whose relative interior contains m; then $pq \subseteq t$. Because s respects f, and f intersects the relative interior of a face t of s, it follows that $t \subseteq f$, contradicting the fact that f contains neither p nor q. ■

Simplices in CDTs have the following property.

Definition 17 (constrained semiregular)**.** *A simplex s is* constrained semiregular *within X if*

- *the vertices of s are in X,*
- *s respects X, and*
- *there exists a d-flat $h_s \subset E^{d+1}$ that includes s^+, such that no vertex $v \in X$ that is visible from a point in the relative interior of s lifts to a point v^+ below h_s. The d-flat h_s is a witness to the constrained semiregularity of s.*

The third condition is a bit difficult to visualize, because one must simultaneously picture the vertices in the ambient space E^d where visibility is determined, and in the ambient space E^{d+1} where witness d-flats are defined, as Figure 13(b) illustrates. Think of it this way: if some lifted vertex v^+ lies below the d-flat that includes a lifted d-simplex s^+, then s is not semiregular, because s^+ is not on the underside of the convex hull of the lifted vertices. But if some facade occludes the view of v from inside s, s may be constrained semiregular anyway and appear in the weighted CDT. The triangle s in Figure 13(b) is an example: although v^+ lies below the witness h_s, v is not visible from the interior of s, so s is constrained semiregular. The shaded triangle in Figure 14 is an example in an unweighted CDT (but note that the paraboloid in the figure is inverted for clarity, so "below" is "above").

[4]An attractive alternative formulation of a weighted PLC extends this characterization to PLCs with submersible vertices. Express a weighted PLC as two separate sets: a PLC X with no submersible vertices, and a set V of submersible vertices. In this formulation, the open facades of X are both conductors and occluders of visibility, and there is a more elegant definition of "respect": a triangulation respects X if every *open* simplex of the triangulation is included in an *open* facade of X.

In Figure 13(b) all three triangles are constrained semiregular, and all the edges are constrained semiregular except the bold constraining segment.

Definition 18 (constrained regular)**.** *A simplex s is* constrained regular *within X if*

- *the vertices of s are in X,*
- *s respects X, and*
- *there exists a d-flat $h_s \subset E^{d+1}$ that includes s^+, such that every vertex $v \in X$ that is visible from a point in the relative interior of s, but is not a vertex of s, lifts to a point v^+ above h_s.*

Of the three triangles in Figure 13(b), only t is constrained regular. Neither the edge shared by r and s nor the constraining segment shared by s and t is constrained regular, but the other edges are.

The following implications hold. Statements in brackets are equivalent to the statements immediately above them in the unweighted case (i.e. when all the vertex weights are zero). *Locally semiregular* and *locally regular* are defined in Section 3.2 and apply to $(d-1)$-simplices only.

s is regular and respects X [s is strongly Delaunay and respects X]	$\longrightarrow$	s is semiregular and respects X [s is Delaunay and respects X]
$\downarrow$		$\downarrow$
s is constrained regular	$\longrightarrow$	s is constrained semiregular [s is constrained Delaunay]
$\downarrow$		$\downarrow$
s is locally regular	$\longrightarrow$	s is locally semiregular [s is locally Delaunay]

The statements in the right column become equivalent to the corresponding statements in the left column when the following condition holds. (Section 6 discusses a perturbation technique that enforces it.)

Definition 19 (genericity)**.** *A d-dimensional PLC X is* generic *if no $d+2$ vertices in X lift to a common non-vertical d-flat (in the ambient space E^{d+1}).*

If X is unweighted (or all the weights are equal), an equivalent statement is that no $d+2$ vertices in X lie on a common hypersphere (in the ambient space E^d).

Notions like constrained regularity are defined in the context of a specific PLC. The definition of "CDT" uses the notion that a simplex can be constrained semiregular within the context of some facade f of a PLC X, yet not be constrained semiregular within the context of X itself.

Definition 20 (facade PLC)**.** *Let f be a k-facade in a PLC X (for any value of k). The* facade PLC Y_f *is a k-dimensional PLC containing f and all the faces of f (taken from X).*

The vertices in a facade PLC often have coordinates from an ambient space E^d whose dimensionality is higher than that of the facade PLC itself (i.e. $d > k$). However, it is the latter dimensionality that defines constraining facades (facades of dimension $k-1$ or less that are not submersible vertices) and ridge protection (the protection of facades of dimension $k-2$ or less; see Definition 23) within Y_f. A simplex that is regular within Y_f might not be regular within X, and a segment that is grazeable within X might not be grazeable within Y_f. Hence, the word *within* is used wherever the context is not clear. Occasionally, this article will say that a simplex is "semiregular within the facade f" as shorthand for saying it is semiregular within the facade PLC Y_f. Likewise, a "triangulation of f" is a triangulation of Y_f.

At last, a definition of this article's central object of study.

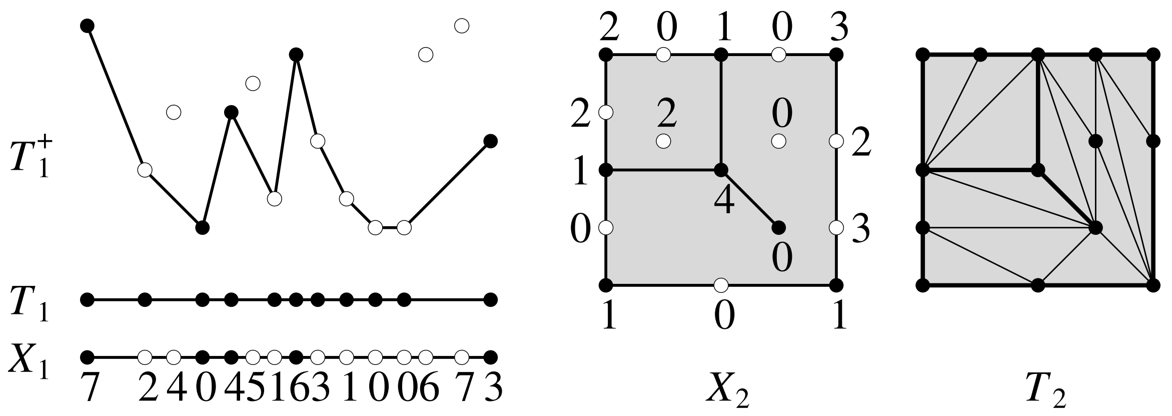

Figure 16: T_1 and T_2 are weighted CDTs of the one- and two-dimensional weighted PLCs X_1 and X_2. White vertices are submersible; black vertices are non-submersible. The number by each vertex is the height (x_{d+1}-coordinate) to which it is lifted.

Definition 21 (weighted constrained Delaunay triangulation)**.** *A* weighted constrained Delaunay triangulation *of a weighted PLC X is a simplicial complex that fills X wherein every simplex is constrained semiregular within the lowest-dimensional facade of X that includes it.*

A constrained Delaunay triangulation *of an unweighted PLC is a weighted CDT for which all the vertices in the PLC are implicitly assigned a weight of zero.*

Figure 16 gives two examples of weighted CDTs, in one and two dimensions. In both triangulations, some vertices are submerged, and some collinear segments of the PLC are agglomerated into single edges of the triangulation. Observe that the lifted one-dimensional triangulation T_1^+ is a sequence of convex hull undersides separated by non-submersible vertices. Note that $d = 1$ is the only dimensionality in which a PLC might have a new CDT if a vertex changes from submersible to non-submersible. For a higher-dimensional PLC with no dangling 1-facades, such a change might cause the PLC to have fewer CDTs (if a submerged vertex is proclaimed non-submersible), but it cannot cause the PLC to have a CDT it did not have before. (This claim is a consequence of the Delaunay Lemma in Section 3.2.)

In an unweighted CDT X (equivalently, if all the weights are equal), every vertex is regular and constrained regular, hence no vertex is submerged.

Definition 21 gives no reason to believe that the eligible simplices (those that are constrained semiregular within the lowest-dimensional facades that include them) can gel together to form a complex. Fortunately, if every facade can be filled with constrained regular simplices, Corollary 18 in Section 3.3 establishes that the facade CDTs match each other where they meet. Not every facade can be thus filled (recall Schönhardt's polyhedron). The next few definitions describe a class of PLCs that are guaranteed to have CDTs.

Definition 22 (grazeable; grazing triangle)**.** *A facade f is* grazeable *if there is an open* grazing triangle $L = \triangle pqr \subset |X|$ *such that*

- *p can see every point in the open triangle L,*
- *pq intersects the open version of f (i.e. f with its external and internal boundaries removed), and*
- *neither p nor q lie on the affine hull of f.*

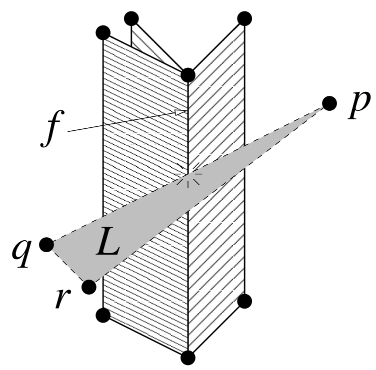

Figure 17: Example of a grazeable segment f. Here p and q cannot see each other, but p sees every point in the open triangle L, so there is a line of visibility that grazes f.

Every point in an open grazing triangle $\triangle pqr$ is visible from p, but q is not (its visibility is occluded by f); so loosely speaking, there is a line of visibility that grazes f. If f is a $(d-2)$-facade, Definition 22 is equivalent to the 180° angle condition described in Section 1.1, as Figure 17 shows. Definition 22 extends the idea to facades of dimension less than $d-2$. Note that the proper faces of a grazeable facade are not necessarily grazeable themselves.

Recall from Section 1.1 that a three-dimensional PLC X is *ridge-protected* if every segment in X is strongly Delaunay. The extension of this definition to weighted PLCs accounts for the possibility that vertices might be submerged: X is ridge-protected if every 1-facade is a union of regular edges, and every non-submersible vertex is regular. (Submersible vertices do not need to be regular, because it is okay to let them be submerged.) The extension of this definition to higher dimensions requires that all constraining facades of dimension $d-2$ or less be "regular," but the definition of "regular" applies only to simplices. It suffices if the facades can be broken up into regular simplices that respect X.

Definition 23 (ridge protection)**.** *A facade $f \in X$ is* protected *if there exists a triangulation of f whose simplices are regular within X and respect X.*

A simpler definition is that f is protected *if f is a union of simplices that are regular within X and respect X. (The equivalence of this definition with the first follows from the upcoming Theorem 4′ and Corollary 18.)*

A weighted PLC X is weakly ridge-protected *if every grazeable constraining facade in X of dimension $d-2$ or less is protected.*

X is ridge-protected *if* every *constraining facade in X of dimension $d-2$ or less is protected.*

How can you tell if a facade f is protected? A weighted Delaunay triangulation T (unconstrained) of the vertices in X contains every simplex that is regular within X (by Theorem 19 in Section 3.3). So the answer is to construct T and search it for a subset of faces that fill f. If T contains such faces, check whether they respect f's faces and are regular. If X is not generic, the trickiest part is distinguishing the regular simplices in T from the merely semiregular. Dafna Talmor (personal communication) points out that simplices that are semiregular but not regular dualize to degenerate faces of the power diagram [2] (the Voronoi diagram if all the weights are zero). This observation does not offer the most numerically effective way to test them, though, and this is not the place to describe a better way. But the simplest approach is to perturb the vertex weights as described in Section 6 before constructing T. Then all the simplices in T are regular, and there is no need to test. Theorem 31 in Section 6 shows that the CDT of the perturbed PLC is a CDT of the unperturbed PLC.

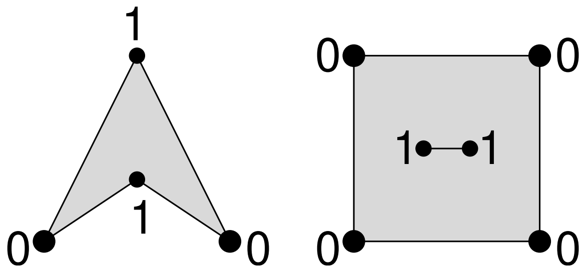

Figure 18: Two weighted PLCs that do not have weighted CDTs. Imagine that you are viewing the lifted vertices from directly underneath, and larger vertices are closer to you. The number by each vertex is the height (x_{d+1}-coordinate) to which it is lifted (i.e. its distance from you).

Ridge protection implies that T respects all the constraining k-facades in X for $k \leq d-2$, but might not respect the $(d-1)$-facades. Weak ridge protection implies that T respects the grazeable constraining facades of dimension $d-2$ or less (and their faces, whether grazeable or not), but perhaps not the other facades. One of the main results of this article is that every weakly ridge-protected weighted PLC has a weighted CDT, so the missing facades can be recovered without any need for additional vertices. See Section 5 for a proof.

Ridge protection requires non-submersible vertices to be regular. For $d = 2$, this is the sole requirement that defines ridge protection. In an unweighted PLC, every vertex is regular, which is why every unweighted two-dimensional PLC has a CDT. In the weighted PLC X_2 in Figure 16, the sole grazeable non-submersible vertex is regular, so X_2 is weakly ridge-protected and has a CDT. (The vertex at the center of X_2 is not regular, but it is not grazeable.) Figure 18 depicts two two-dimensional weighted PLCs that are not weakly ridge-protected, and do not have weighted CDTs. Both examples include a grazeable non-submersible vertex that is not regular.

A recently proposed way to model domains like these ones and Schönhardt's polyhedron is to generalize simplicial complexes to *pseudo-simplicial complexes* composed of nonconvex pseudo-simplices. Aichholzer, Aurenhammer, Krasser, and Brass [1] define constrained regular pseudotriangulations that generalize the two-dimensional constrained regular triangulations defined here, and are defined for every choice of vertex weights. Their lifted surface is not necessarily continuous, and is not guaranteed to interpolate all the vertex heights. Aurenhammer and Krasser [3] show that the approach generalizes to higher-dimensional nonconvex polyhedra, but pseudo-simplicial complexes representing polyhedra in three dimensions or more must sometimes introduce additional vertices.

Throughout the rest of this article, the terms "PLC" and "CDT" refer to both unweighted and weighted PLCs and CDTs, except where otherwise noted.

3 Foundations

This section proves several fundamental properties of CDTs and weighted CDTs. Among these are the fact that every face of a constrained semiregular simplex is constrained semiregular within some facade (Section 3.1), the fact that constrained regular simplices have disjoint relative interiors and form a complex, and the fact that a generic PLC has at most one CDT (Section 3.3). The Delaunay Lemma offers a powerful alternative characterization of what a CDT is (Section 3.2). Readers who seek the minimum background for understanding the CDT construction algorithms in the sequel articles may safely skip to Section 6.

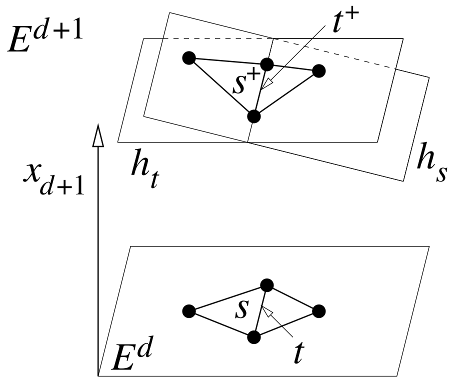

Figure 19: The simplex s is regular because every lifted vertex lies above some witness d-flat h_s for s^+, except the vertices of s^+. Let t be any face of s. Tilting h_s using t^+ as a hinge yields a witness d-flat h_t that shows that t is regular too.

3.1 Faces of Simplices Inherit Semiregularity and Constrained Semiregularity

CDTs (unweighted and weighted) have properties that allow proofs and algorithms to work in a top-down fashion: if a domain can be filled with a complex of constrained semiregular d-simplices, the lower-dimensional faces "work themselves out."

Let s be a simplex that is constrained semiregular within some PLC X. Let t be a face of s. If t is not included in a constraining facade in X, then t is also constrained semiregular. What if t *is* included in a constraining facade? Then t might not be constrained semiregular within X, but t is constrained semiregular within the lowest-dimensional facade that includes t (and is not a submersible vertex). It follows that the act of filling a d-facade with a complex of constrained semiregular d-simplices automatically fills all of its proper faces with lower-dimensional CDTs.

First consider unconstrained semiregularity.

Theorem 4. *Every face of a semiregular simplex is semiregular.*

Theorem 4′. *Every face of a regular simplex is regular.*

Proof. Let s be a semiregular simplex, and let t be a face of s as in Figure 19. Let h_s be a witness to the semiregularity of s. That is, h_s is a d-flat that includes s^+, and no vertex in X lifts to a point below h_s. Clearly, h_s is also a witness to the semiregularity of t, so Theorem 4 holds.

Suppose s is regular. Then every vertex in X lifts above h_s except the vertices of s. Let h_t be a d-flat found by tilting h_s by a tiny amount (as illustrated), so that h_t includes t^+ but lies below the vertices of s^+ not shared by t^+. If the tilt is small enough, the other vertices in X still lift to points above h_t. Hence, h_t is a witness to the regularity of t, and Theorem 4′ holds. ■

Theorem 5. *Let X be a PLC. Let s be a simplex, and let t be a face of s that is not included in a (closed) constraining facade in X. If s is constrained semiregular, then t is constrained semiregular.*

Theorem 5′. *Under the assumptions of Theorem 5, if s is constrained regular, then t is constrained regular.*

Proof. Because s is constrained semiregular, s respects X. As t is a face of s, t also respects X.

Observe that every vertex visible from the relative interior of t is visible from the relative interior of s. Specifically, suppose a vertex v is visible from a point p in the relative interior of t. Because p does not lie in a constraining facade in X, Lemma 6 below implies that some point p' in the relative interior of s sees v.

The rest of the proof is identical to the proof of Theorems 4 and 4′, except that only vertices visible from the relative interior of t are considered, and every occurrence of "semiregular" or "regular" is thus replaced with "constrained semiregular" or "constrained regular." ■

The following lemma (which is used frequently in this article) tells us a way to perturb a point p without occluding its visibility from another point q.

Definition 24 (ϵ-neighbor)**.** *A point p' is an ϵ-neighbor[5] of a point p, with respect to a point q and a closed PLC X, if $p' \in |X|$, $|pp'| \leq \epsilon$, and every (closed) constraining facade in X that contains p contains either p' or q.*

Lemma 6. *Let p and q be two points that can see each other within a PLC X. There is a positive constant ϵ such that every ϵ-neighbor of p can see q.*

Proof. Any facade whose affine hull contains q cannot occlude the visibility between p' and q. Every facade that contains p contains either p' or q, and thus cannot occlude the visibility between p' and q.

What about the other facades? The line segment pq does not intersect any of them. There is a finite gap between pq and any facade that does not intersect pq, and p must move some non-infinitesimal distance to close the gap. A sufficiently small choice of ϵ ensures that every ϵ-neighbor of p is visible from q. ■

Observe that if p lies in a constraining facade f, but p' and q do not, then p' is not an ϵ-neighbor of p, and f might occlude the visibility between p' and q.

Next, consider the circumstance where a face of a simplex is included in a constraining facade. The case of a semiregular simplex is considered first (that's *un*constrained semiregular, albeit in the context of a PLC), followed by the case of a constrained semiregular simplex.

Theorem 7. *Let s be a simplex, and let t be a face of s. Suppose a constraining k-facade $f \in X$ includes t. Let Y_f be the k-dimensional facade PLC for f (recall Definition 20).*

If s is semiregular within X, then t is semiregular within Y_f.

Theorem 7′. *Under the assumptions of Theorem 7, if s is regular within X, then t is regular within Y_f.*

Proof. For intuition's sake, consider first the special case where s is a Delaunay tetrahedron, illustrated in Figure 20. No vertex lies inside the circumsphere of s. If a triangular face t of s lies within a 2-facade f, then t is Delaunay within the two-dimensional PLC Y_f. Why? Because the circumcircle of t is a cross-section of the circumsphere of s, and therefore it encloses no vertex. If s is strongly Delaunay, t is strongly Delaunay.

Figure 21 extends this reasoning to weighted CDTs. Let s be a semiregular simplex. There is a witness d-flat h_s that includes s^+ such that no lifted vertex lies below h_s. Because the face t of s is included in a k-facade f, h_s yields a witness to the fact that t is semiregular within Y_f as follows.

Let F be the affine hull of f. Think of F as the affine space in which Y_f is defined. Let $F^+ = \{\langle p, \alpha \rangle \in E^{d+1} : p \in F, \alpha \in \mathbb{R}\}$. F^+ is a vertical $(k+1)$-flat in E^{d+1}, as Figure 21 shows. Think of F^+ as the affine space in which witnesses for Y_f are defined. Then $h_t = h_s \cap F^+$ is a witness k-flat within F^+ that includes t^+.

[5]It would be more apt to call this an (ϵ, q, X)-neighbor of p, but it would clutter the writing.

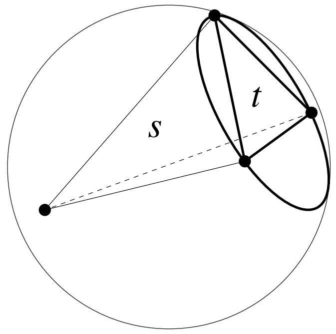

Figure 20: An unweighted example where $d = 3$. If a tetrahedron s is Delaunay, each of its faces has an empty circumcircle, because each face's circumcircle is a cross-section of the tetrahedron's circumsphere.

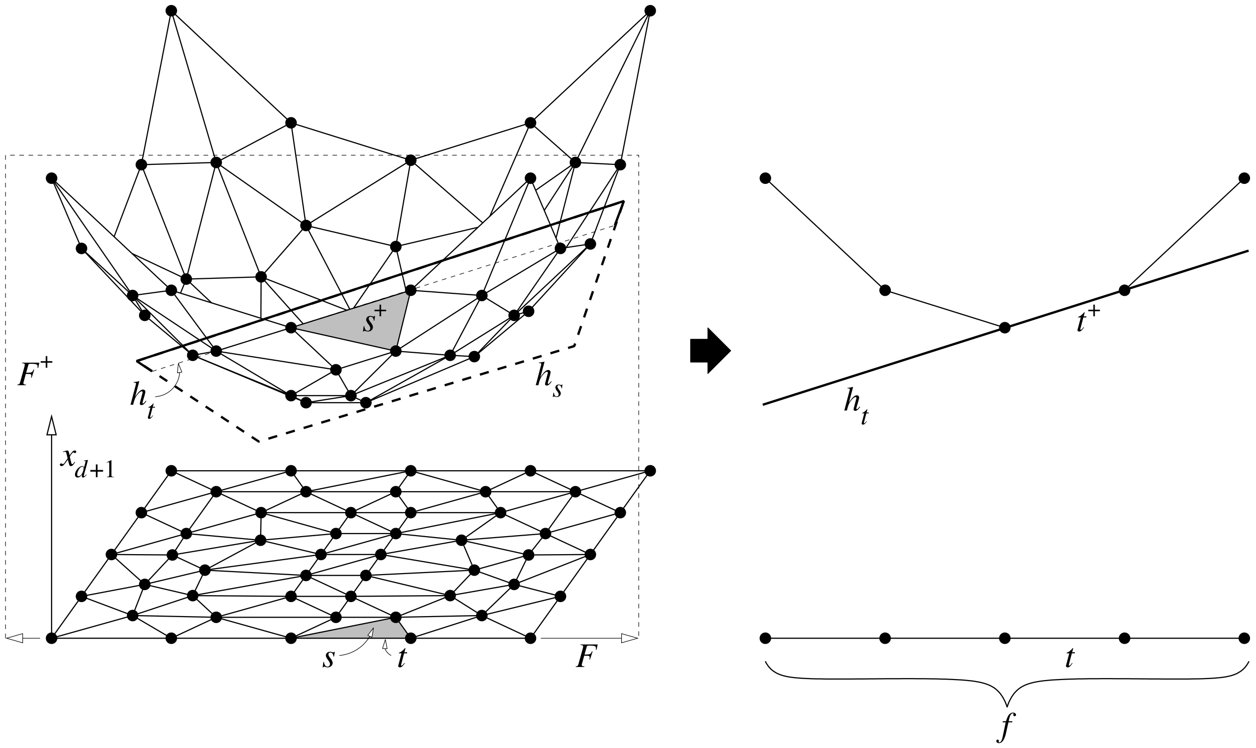

Figure 21: A weighted example where $d = 2$. If a simplex s is semiregular (no lifted vertex lies below h_s), any face t of s that lies in a facade f is semiregular within f (no lifted vertex lies below h_t).

Because no vertex in X lifts to a point below h_s, no vertex in Y_f lifts to a point below h_t, so t is semiregular within Y_f and Theorem 7 holds.

If s is regular, the lifted companion of every vertex in X lies above h_s, except the vertices of s^+ (which lie on h_s). Thus the lifted companion of every vertex in Y_f lies above h_t, except the vertices of s^+. If every vertex of s in Y_f is also a vertex of t, then h_t is a witness to the regularity of t within Y_f. Otherwise, h_t contains at least one vertex of s^+ that is not a vertex of t^+, but that is no obstacle. By tilting slightly as described in the proof of Theorem 4′, h_t becomes a witness to the regularity of t within Y_f. Thus Theorem 7′ holds. ■

The next theorem generalizes Theorem 5, and is the constrained analog of Theorem 7.

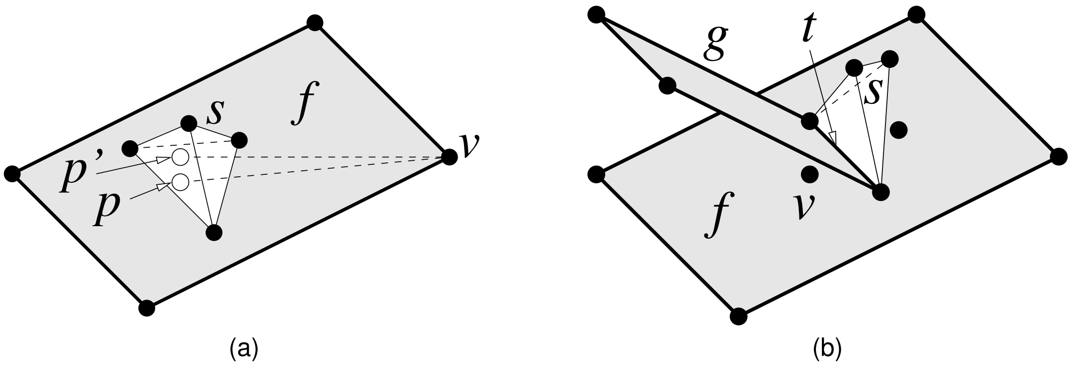

Figure 22: (a) Example in which f is a 2-facade inside a three-dimensional PLC. The tetrahedron s (which is not a facade) intersects f at a triangular face of s. The point p lies in the relative interior of the triangular face, and p' lies in the interior of s. Both p and p' can see v. (b) Here s intersects f at an edge t. Although v is visible from every point on t, v is not visible from inside s.

Theorem 8. *Let s be a simplex, and let t be a face of s. Let f be the lowest-dimensional facade in X that includes t and is not a submersible vertex.*

If s is constrained semiregular within X, then t is constrained semiregular within Y_f.

Theorem 8′. *Under the assumptions of Theorem 8, if s is constrained regular within X, then t is constrained regular within Y_f.*

Proof. Because s is constrained semiregular, s respects X. As t is a face of s, t also respects X. Moreover, $t \subseteq f = |Y_f|$, so t respects Y_f.

Let v be any vertex in Y_f that is visible from some point p in the relative interior of t. The following reasoning shows that v is also visible from some point in the relative interior of s. See Figure 22(a). As p can see v, Lemma 6 guarantees that there is an $\epsilon > 0$ such that every ϵ-neighbor of p can also see v. Because t respects f's faces and f is the lowest-dimensional non-submersible facade that includes t, the relative interior of t does not intersect any proper face of f, with the possible exception of submersible vertices. Therefore, every constraining facade that contains p has f for a face and contains v as well. It follows that every point in $|X|$ within a distance of ϵ from p is an ϵ-neighbor of p (with respect to v). Because p is on the boundary of s, v is visible from some point in the relative interior of s.

The rest of the proof is identical to the proof of Theorems 7 and 7′, except that only vertices that are in Y_f and visible from the relative interior of t are considered, and every occurrence of "semiregular" or "regular" is thus replaced with "constrained semiregular" or "constrained regular." ■

Figure 22(b) demonstrates why Theorems 8 and 8′ do not apply if f is not the lowest-dimensional facade (other than a submersible vertex) that includes t. In this example, an edge t of a tetrahedron s is a constraining segment on the internal boundary of a 2-facade f. Although s is constrained regular within X, t is not constrained regular within Y_f, because the vertex v is visible from every point on t. The 2-facade g occludes the visibility of v from every point in the interior of s, allowing s to be constrained regular.

The next two theorems simplify proving that a triangulation is a CDT, by putting the burden on the highest-dimensional simplices.

Theorem 9. *Let X be a d-dimensional PLC with no dangling facades (i.e. each facade in X is included in a d-facade in X). Let T be a simplicial complex that fills X. Suppose every d-simplex in T is constrained semiregular.*

Then T is a CDT of X. Furthermore, for every facade f in X except submersible vertices, $\{t \in T : t \subseteq f\}$ is a CDT of Y_f.

Proof. By the definition of "constrained semiregular," every d-simplex in T respects X; and as T is a simplicial complex with no dangling simplices, every simplex in T respects X.

Let t be any simplex in T. Let f be the lowest-dimensional facade in X that includes t. If f is a vertex, then $t = f$ and t is trivially constrained semiregular within Y_f. Otherwise, let s be a d-simplex in T having t for a face. (Some such d-simplex must exist, because T is a simplicial complex filling a PLC with no dangling facades.) By assumption, s is constrained semiregular, so t is constrained semiregular within Y_f by Theorem 8.

Therefore, every simplex in T is constrained semiregular within the lowest-dimensional facade that includes it. By definition, T is a CDT of X. Because T fills and respects X, for every non-submersible facade $f \in X$, the subcomplex $\{t \in T : t \subseteq f\}$ fills and respects Y_f, and thus is a CDT of Y_f. ■

The next theorem generalizes Theorem 9 to cover PLCs of mixed dimensionality.

Theorem 10. *Let X be a PLC (possibly with dangling facades). Let T be a simplicial complex that fills X. Suppose that for every $k \geq 1$, for every k-facade $f \in X$ that is not a face of a higher-dimensional facade, every k-simplex of T included in f is constrained semiregular within Y_f.*

Then T is a CDT of X. Furthermore, for every facade f in X except submersible vertices, $\{t \in T : t \subseteq f\}$ is a CDT of Y_f.

Proof. Identical to the proof of Theorem 9, except that s is a k-simplex in T having t for a face, where k is the dimensionality of the highest-dimensional facade that includes t. (By assumption, s is constrained semiregular within the k-facade that includes s.) ■

This section concludes with two corollaries of Theorem 7′.

Corollary 11. *If X is ridge-protected, every facade in X is ridge-protected. (That is, if $f \in X$, then its facade PLC Y_f is ridge-protected.)*

Proof. Ridge protection holds trivially for a PLC of dimension less than two, so let f be any facade in X of dimension $k \geq 2$. Let Y_f be f's facade PLC. Let d be the dimensionality of X. Because X is ridge-protected and $Y_f \subseteq X$, every constraining facade in Y_f of dimension $d - 2$ or less has a triangulation whose simplices respect X and are regular within X. By Theorem 7′, these simplices are also regular within Y_f. Therefore, every constraining facade in Y_f of dimension $k - 2$ or less has a triangulation whose simplices respect Y_f and are regular within Y_f. ■

Corollary 12. *If X is weakly ridge-protected, every facade in X is weakly ridge-protected.*

Proof. Let f be a facade in X, and let Y_f be f's facade PLC. It is apparent from Definition 22 that if a face of f is grazeable within Y_f, then the face is grazeable within X too. The rest of the proof is identical to the proof of Corollary 11, except that only every *grazeable* constraining facade in Y_f of dimension $d - 2$ or less has a triangulation whose simplices respect X (and therefore Y_f) and are regular within X (and therefore within Y_f). ■

3.2 The Delaunay Lemma

A well-known and important property of Delaunay triangulations is that "local optimality" is equivalent to "global optimality," in the following sense. A facet shared by two d-simplices s and t is said to be *locally Delaunay* if the apex of s (not shared by t) is not inside the circumsphere of t (equivalently, the apex of t is not inside the circumsphere of s). If a triangulation is Delaunay, every facet of the triangulation is locally Delaunay. Conversely, if every facet of a triangulation of a point set is locally Delaunay, then the triangulation is Delaunay (i.e. every simplex is Delaunay). Boris Delaunay [14] himself was the first to make this observation.

This section shows that this equivalence generalizes to weighted CDTs, with the change that facets included in constraining facades need not be locally Delaunay (or locally semiregular). This result is valuable because it provides an inexpensive way to test whether a triangulation is a weighted CDT: check that it fills and respects the PLC, check every non-constraining facet for local semiregularity, and check each submerged vertex to ensure it really should be submerged. (A *non-constraining facet* is a facet that is not included in a constraining facade.) The Delaunay Lemma offers an alternative answer to the question, "What does it mean for a PLC X to have no CDT?" It means that no triangulation of X fulfills these requirements.

Definition 25 (locally regular; locally semiregular)**.** *Let T be a triangulation, and let s and t be two d-simplices in T that share a facet f. The facet f is* locally regular *within T if the lifted d-simplices s^+ and t^+ adjoin each other at a dihedral angle, measured from above, of less than* 180°. *In other words, the apex of t^+ lies above the witness d-flat of s, and vice versa, as illustrated in Figure 13(a).*

The facet f is locally semiregular *within T if the upper dihedral angle where s^+ meets t^+ is less than or equal to* 180°. *In other words, either f is locally regular, or s and t have the same witness d-flat.*

If a facet f is constrained regular, then f is locally regular, because the apices of s^+ and t^+ lie above some witness d-flat of f. If f is constrained semiregular, f is locally semiregular.

Theorem 13 (Delaunay Lemma)**.** *Let X be a (weighted) PLC with no dangling facades. A triangulation T is a (weighted) CDT of X if and only if T has the following four properties.*

A. T fills X.

B. T respects X.

C. Every facet in T is either locally semiregular or included in a constraining $(d-1)$*-facade of X.*

D. If a vertex v in X is missing from T (submerged), then v is in a d-simplex s of T such that v^+ lies on or above s^+.

If X is unweighted, Property D reads, "No vertex is submerged."

The proof of the Delaunay Lemma relies on a lemma that is worth stating separately because it is reused in Sections 3.3 and 5.1. The lemma uses the following definitions.

Definition 26 (overlaps)**.** *Let z be an arbitrary point in E^d. Let s and t be two simplices (each of any dimensionality). Say that s* overlaps *t* from the viewpoint *z if some point of s not shared by t lies between z and t, as Figure 23 illustrates. In other words, there exists a point $p_s \in s \setminus t$ and a point $p_t \in t$ such that $p_s \in zp_t$.*

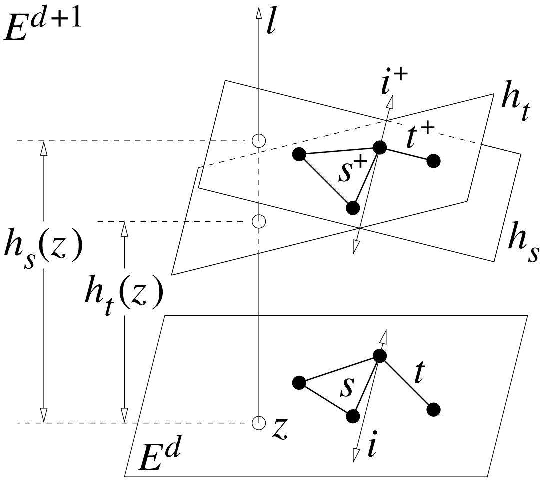

Figure 23: If s overlaps t from the viewpoint z, then $h_s(z) > h_t(z)$.

Definition 27 (witness function)**.** *Let $h \subset E^{d+1}$ be a non-vertical d-flat. The* witness function $h(p)$ *is the linear function that maps each point $p \in E^d$ to the x_{d+1}-coordinate such that $\langle p, h(p)\rangle \in h$. In other words, if $\ell \subset E^{d+1}$ is the vertical line (parallel to the x_{d+1}-axis) that contains $\langle p, 0\rangle$, then $h(p)$ is the x_{d+1}-coordinate of $h \cap \ell$, as Figure 23 illustrates.*

Lemma 14. *Let s and t be two simplices, each of any dimensionality. Suppose there is a non-vertical d-flat h_s that includes s^+ such that every vertex of t^+ lies on or above h_s. Suppose there is a non-vertical d-flat h_t that includes t^+ such that every vertex of s^+ lies* strictly *above h_t, except the vertices shared by t^+. Then the following statements hold.*

- *If s and t are not disjoint, then $s \cap t$ is a face of both s and t.*
- *Let z be an arbitrary point in E^d. If s overlaps t from the viewpoint z, then $h_s(z) > h_t(z)$.*

Proof. If s is a face of t, both results follow immediately. (In this case, s does not overlap t from any viewpoint.) Otherwise, s^+ has a vertex that t^+ lacks. This vertex lies on h_s and above h_t, so $h_s \neq h_t$. The d-flats h_s and h_t must intersect, because some vertex of s^+ lies above h_t and some vertex of t^+ lies on or above h_s. Let i^+ be the $(d-1)$-flat $h_s \cap h_t$. Let $i = \{p \in E^d : h_s(p) = h_t(p)\}$ be the vertical projection of i^+ into E^d, as illustrated in Figure 23.

The hyperplane i cuts E^d into two halfspaces. Every vertex of s lies in the closed halfspace $\{p \in E^d : h_s(p) \geq h_t(p)\}$. Therefore, so does every point in s. Likewise, every point in t lies in the closed halfspace $\{p \in E^d : h_s(p) \leq h_t(p)\}$. Any vertex v of s that lies on i has a lifted companion v^+ that lies on h_t, so by assumption, v must be a vertex of t. Therefore, $s \cap t$ is the convex hull of the vertices of s that lie on i, which is a face of both s and t. Furthermore, any point of s not shared by t cannot lie on i.

If s overlaps t from the viewpoint z, then some point $p_s \in s \backslash t$ lies between z and t. The point p_s lies in the open halfspace $\{p \in E^d : h_s(p) > h_t(p)\}$, so z must lie there too. ■

Lemma 14 is similar to theorems of Edelsbrunner [15] and Edelsbrunner and Shah [19], which they use to prove the *acyclicity* of every regular triangulation T: for any fixed viewpoint z, the overlap relation among regular simplices is a partial order. The function $h_s(z)$ imposes a total order on the simplices in T

such that no simplex overlaps another simplex that appears later in the order. This acyclicity property does not extend to CDTs, but it does apply to the regular simplices that comprise the lower-dimensional facades in a ridge-protected PLC (see Section 5.1).

Proof of the Delaunay Lemma. The "only if" implication is straightforward. If T is a CDT of X, Properties A and B follow by the definition of CDT. Property D follows because every d-simplex in a CDT is constrained semiregular. Property C follows because each facet in a CDT is constrained semiregular—unless it is included in a constraining $(d-1)$-facade of X—and every constrained semiregular simplex is locally semiregular.

Not surprisingly, the "if" implication takes more work to prove. Suppose T is a triangulation with all four properties. Let s be any d-simplex in T. The following argument establishes that s is constrained semiregular.

Let v be any vertex in X that is visible from some point p in the interior of s. It is helpful if the line segment vp does not intersect any simplex in T of dimension less than $d-1$, except at the vertex v. If this is not true, then by Lemma 6 there is a neighborhood of p from which every point can see v. Choose from this neighborhood a point p' such that p' is in the interior of s and vp' does not intersect any simplex in T of dimension less than $d-1$, except at v.

T is a simplicial complex that fills X by Property A, so the line segment vp' intersects the interiors of a contiguous sequence of d-simplices $s_1, s_2, \ldots, s_k = s$, with $v \in s_1$. Let f_i denote the facet shared by s_i and s_{i+1}. Because vp' does not intersect any lower-dimensional faces of T (except at v), it passes through the relative interiors of the facets $f_1, f_2, \ldots, f_{k-1}$. Because v is visible from p', none of these facets is included in a constraining facade, so by Property C all of them are locally semiregular.

Because f_1 is locally semiregular, either $h_{s_1} = h_{s_2}$ or f_1 is locally regular. In the latter case, $h_{s_1}(v) > h_{s_2}(v)$ by Lemma 14; in either case, $h_{s_1}(v) \geq h_{s_2}(v)$. The same reasoning holds for $f_2, \ldots, f_{k-1}$, so $h_{s_1}(v) \geq h_{s_2}(v) \geq \cdots \geq h_{s_k}(v) = h_s(v)$. If v is a vertex of s_1, then the height (x_{d+1}-coordinate) of v^+ is $h_{s_1}(v)$; otherwise, v is submerged, and by Property D the height of v^+ is at least $h_{s_1}(v)$.[6] In either case, $v_{x_{d+1}} \geq h_{s_1}(v) \geq h_s(v)$, so v^+ cannot lie below the witness d-flat h_s. Because this is true of every vertex v that is visible from the interior of s, and because s respects X by Property B, s is constrained semiregular.

By assumption, X has no dangling facades, so by Theorem 9, T is a CDT of X. ∎

If X has dangling facades, T may be cut into subcomplexes of different dimensionalities so that each subcomplex has no dangling simplices. Then the Delaunay Lemma can be applied to each piece separately, thereby showing the constrained semiregularity of the whole. In a k-dimensional portion of the triangulation, only the local semiregularity of the $(k-1)$-faces needs to be checked.

To make good on the title of this article, the following definition offers the constrained analog of a regular triangulation. A constrained regular triangulation is a projection of a polyhedron whose ridges are locally convex everywhere except where the constraining facades permit them to be reflex.

Definition 28 (constrained regular triangulation)**.** *A triangulation T is* constrained regular *relative to an unweighted PLC X if T fills and respects X, and there exists an assignment of weights to the vertices in X such that every non-constraining facet in T is locally regular.*[7]

[6] The vertex v might lie in several d-simplices of T (on a shared boundary), and Property D explicitly applies to only one of them. But the lifted surface T^+ is continuous where simplices of T meet, so Property D holds for all the simplices in T that contain v.

[7] Obviously, there is always an assignment of weights to the vertices of X missing from T that satisfies Property D of the Delaunay Lemma. Just make their weights really small.

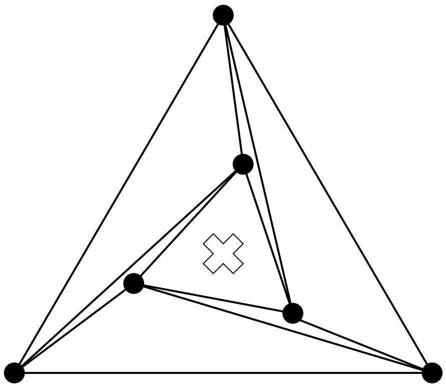

Figure 24: A triangulation that is not regular. From the viewpoint at the center, the three outer triangles form a mutually overlapping cycle.

Every generic CDT (recall Definition 19) is a constrained regular triangulation. This fact is a consequence of the Delaunay Lemma and the fact that in a generic CDT, constrained regularity and constrained semiregularity are the same. However, not *every* CDT is a constrained regular triangulation. For example, let T be the triangulation illustrated in Figure 24, which is not regular. If all the vertex heights are zero, T is a valid weighted Delaunay triangulation and (relative to a compatible PLC) a valid weighted CDT. But only the simplices on the boundary of T are regular; the rest are only semiregular. No assignment of weights can make every edge of T regular. But if T is a triangulation of a PLC X, and X includes one of the long internal edges as a constraining segment, then T is constrained regular with respect to X.

The viewpoint at the center of the triangulation T in Figure 24 demonstrates that T does *not* have the acyclic property established by Edelsbrunner and Shah [19] for regular triangulations. However, constrained regular triangulations have a limited acyclicity property. Say that *s visibly overlaps t* from the viewpoint z if there exists a point p_t in t's relative interior that is visible from z, and a point $p_s \in s \backslash t$ such that $p_s \in zp_t$. For any viewpoint z, the visible overlap relation among simplices is a partial order. This fact follows from Lemma 14 by the same inductive step used to prove the Delaunay Lemma, with the inequalities replaced by the strict inequalities $h_{s_i}(z) > h_{s_{i+1}}(z)$.

Linear programming can determine whether a triangulation T that fills and respects X is constrained regular relative to X. The variables of the linear program are the vertex weights and a variable δ. For each non-constraining facet f in T, write a linear constraint enforcing the local regularity of f. Specifically, f is a facet of two d-simplices s and t; the linear constraint requires that the apex of s^+ (not shared by t^+) be a distance of at least δ above t's witness d-flat. The objective is to maximize δ subject to the facet constraints. If this linear program has a feasible point with $\delta > 0$, T is constrained regular relative to X.

3.3 The Omnipresent Complex of Constrained Regular Simplices

A property of every PLC X is that its constrained regular simplices (of all dimensionalities, within all the facades in X) have disjoint relative interiors and form a simplicial complex, even if X has no CDT. Another property is that if X does have a CDT—perhaps several CDTs—then every constrained regular simplex appears in every CDT of X. This property implies that if X is generic, it has at most one CDT.

These properties do not hold for semiregular simplices. If some selection of $d + 2$ or more vertices of a PLC lift to a common non-vertical d-flat, the PLC might have more than one CDT, and its semiregular simplices might have intersecting interiors.

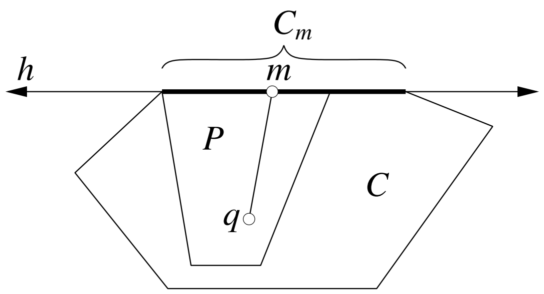

Figure 25: If $P \subseteq C$ but $P \nsubseteq C_m$, then m is on the boundary of P.

Because the CDT of a generic PLC contains every constrained regular simplex and no other simplex, CDT construction algorithms can work in a bottom-up fashion, from low dimensionalities to high: if an algorithm obtains the CDT of each constraining facade in a generic PLC X (perhaps by calling itself recursively), it can construct the constrained regular d-simplices with confidence that they will match the facade triangulations. For a nongeneric PLC, however, the CDTs of different constraining facades might be incompatible with each other, causing a CDT construction algorithm to fail to find a CDT of the whole PLC even when one exists. Section 6 offers a perturbation method that enforces genericity, so that CDT construction algorithms may avoid this fate.

The proofs rely on the following lemma, which is also used heavily in Section 5.

Lemma 15. *Let P and C be closed, convex polyhedra (not necessarily of the same dimensionality) with $P \subseteq C$. Let m be a point in the relative interior of P. Let C_m be the face of C whose relative interior contains m. Then $P \subseteq C_m$.*

Proof. If $C_m = C$ the results follows immediately. Otherwise, by Definition 3, there is a hyperplane h such that $C_m = C \cap h$ and h does not intersect the relative interior of C, which implies that $C \backslash C_m$ lies entirely on one side of h. Clearly, $m \in h$. Suppose for the sake of contradicting the lemma that $P \nsubseteq C_m$, as illustrated in Figure 25. Then P contains a point q in $C \backslash C_m$. Thus q is on the same side of h as C, and no point in P is on the other side of h. Because P is convex, P includes the line segment qm, but any extension of the line segment qm past m lies outside P. Therefore, m is on the boundary of P. This contradicts the assumption that m is in the relative interior of P, so $P \subseteq C_m$. ■

The following theorem, which generalizes half of Lemma 14, shows that a constrained semiregular simplex and a constrained regular simplex can intersect only at a shared face.

Theorem 16. *Let s and t be simplices. Suppose that s is constrained semiregular within f and t is constrained regular within f', where f and f' are facades in a PLC X (possibly with $f = f'$), and neither f nor f' is a submersible vertex. If s and t are not disjoint, then $s \cap t$ is a face of both s and t.*

Proof. Suppose s and t are not disjoint. Let p be a point in the relative interior of $s \cap t$, as illustrated in Figure 26. Let g be the lowest-dimensional facade in X that contains p and is not a submersible vertex. (Either p is in the relative interior of g, or p coincides with an isolated submersible vertex of g's internal boundary. Note that g might be of any dimension from zero to d.) Because $p \in f$ and $p \in f'$, g is a face (not necessarily a proper face) of both f and f'.

Each of s and t has one face whose relative interior contains p. Call these faces $\hat{s}$ and $\hat{t}$, respectively. Because s and t respect X, so do $\hat{s}$ and $\hat{t}$. It follows that every facade that contains p (and is not a submersible vertex) includes $\hat{s}$ and $\hat{t}$. Three such facades are f, f', and g.

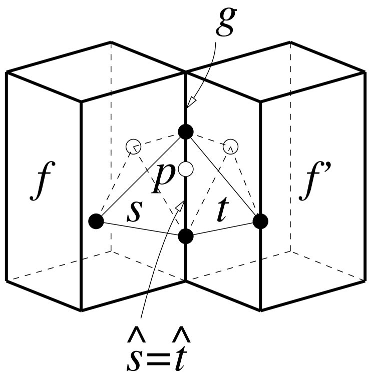

Figure 26: A constrained semiregular simplex s and a constrained regular simplex t can intersect only at a shared face.

By Theorem 8, $\hat{s}$ is constrained semiregular within g. By Theorem 8′, $\hat{t}$ is constrained regular within g. By Lemma 14 (applied within the facade PLC Y_g), $\hat{s} \cap \hat{t}$ is a face of both $\hat{s}$ and $\hat{t}$. But p is in the relative interiors of both $\hat{s}$ and $\hat{t}$, and $p \in \hat{s} \cap \hat{t}$, so $\hat{s} = \hat{s} \cap \hat{t} = \hat{t}$.

Because $\hat{s} = \hat{t}$ is a face of both s and t, $\hat{t} \subseteq s \cap t$. By Lemma 15 (substituting p for m, t for C, $\hat{t}$ for C_m, and $s \cap t$ for P), $s \cap t \subseteq \hat{t}$. Therefore, $\hat{s} = \hat{t} = s \cap t$, verifying that $s \cap t$ is a face of both s and t. ■

Corollary 17. *The constrained regular simplices of a PLC have disjoint relative interiors.* ■

Corollary 18. *Let X be a PLC. Let T be the set that contains every simplex that is constrained regular within X or within a constraining facade in X. T is a simplicial complex.*

Proof. By Theorem 8′, every face of every simplex in T is constrained regular within some facade that is not a submersible vertex. Therefore, T contains every face of every simplex in T. By Theorem 16, the intersection of any two simplices in T is either empty or a shared face of the two simplices. Hence T is a simplicial complex. ■

A consequence of Corollary 18 is that if a PLC does not have a CDT, one or more of its facades has a gap that is not covered by constrained regular simplices. The next theorem shows that if a PLC has several CDTs, they share the same constrained regular simplices, and differ only by the simplices that are constrained semiregular but not constrained regular.

Theorem 19. *Every CDT of a PLC X contains every simplex that is constrained regular within X or within a constraining facade in X.*

Proof. Let t be any simplex that is constrained regular within some facade f in X, where f is not a submersible vertex. (If a simplex is constrained regular within X, it is constrained regular within some d-facade in X.) Let p be a point in the relative interior of t.

Let T be a CDT of X. Because T fills X, T contains a simplex s that contains p and is not a submersible vertex. By the definition of CDT, s is constrained semiregular within the lowest-dimensional facade that includes it.

By Theorem 16, $s \cap t$ is a face of both s and t. But $s \cap t$ contains p, which is in the relative interior of t, so $s \cap t = t$. Therefore, t is a face of s, and $t \in T$. This conclusion holds for every CDT T of X and every constrained regular simplex t. ■

Corollary 20. *A generic PLC has at most one CDT.*

Proof. By Theorem 19, every CDT of a PLC X contains every simplex that is constrained regular within a facade in X, except perhaps within a submersible vertex. By the definition of CDT, no CDT of X contains a simplex that is not constrained semiregular within a facade in X. If X is generic, constrained regularity and constrained semiregularity are equivalent. Therefore, two CDTs of X can differ from each other only in the choice of submersible vertices. But a CDT fills X, so the choice of submersible vertices is uniquely determined by the higher-dimensional simplices. Therefore, X has at most one CDT. ■

This corollary and Corollary 18 together imply that if a PLC is generic and has a CDT, a CDT construction algorithm can triangulate each facade of the PLC, starting with the 1-facades and working up to the d-facades, and rest assured that the facade triangulations of different dimensions all match.

4 Interpolation Criteria Optimized by CDTs

Among all triangulations of a fixed two-dimensional vertex set, the Delaunay triangulation is optimal by a variety of criteria—maximizing the smallest angle in the triangulation [28], minimizing the largest circumcircle among the triangles [4], and minimizing a property called the *roughness* of the triangulation [37, 35]. A two-dimensional CDT shares these same optimality properties, if it is compared with every other *constrained* triangulation of the same PSLG [29, 4].

Delaunay triangulations in higher dimensions also have optimality properties that generalize to CDTs and offer some of the reasons why higher-dimensional CDTs are such worthy objects of study. Rippa [38] investigates the use of two-dimensional triangulations for piecewise linear interpolation of a bivariate function of the form $Ax^2 + By^2 + Cx + Dy + E$, and concludes that if $A = B$, the Delaunay triangulation minimizes the interpolation error measured in the L_q-norm for every $q \geq 1$ (compared with all other triangulations of the same vertices). Melissaratos [31] generalizes Rippa's result to higher dimensions. D'Azevedo and Simpson [13] show that a two-dimensional Delaunay triangulation minimizes the radius of the largest min-containment circle of its simplices, and Rajan [36] generalizes this result to Delaunay triangulations and min-containment spheres of any dimensionality. The *min-containment sphere* of a simplex is the smallest hypersphere that encloses the simplex. If the center of the circumsphere of a simplex lies in the simplex, then the min-containment sphere is the circumsphere. Otherwise, the min-containment sphere is the min-containment sphere of some face of the simplex.

Rajan's result and a theorem of Waldron [51] together imply a second optimality result related to multivariate piecewise linear interpolation. Suppose you must choose a triangulation to interpolate an unknown function (not necessarily convex), and you wish to minimize the largest pointwise error in the domain. After you choose the triangulation, an adversary will choose the worst possible smooth function for your triangulation to interpolate, subject to a fixed upper bound on the absolute curvature (i.e. second directional derivative) of the function anywhere in the domain. The Delaunay triangulation is your optimal choice.

This section shows that Melissaratos' and Rajan's results generalize to CDTs (when CDTs exist). Melissaratos' result also generalizes to any monotonic norm and, with help from weighted CDTs, to any convex function. Rajan's result is particular to unweighted CDTs—the paraboloid is the right choice of heights to minimize the largest min-containment sphere. The proofs given here are similar to Fortune's presentation for unconstrained Delaunay triangulations [21], and are substantially simpler than Melissaratos' and Rajan's.

Consider multivariate piecewise linear interpolation on a weighted CDT. Let X be a PLC, and let $f(p)$ be a convex scalar function defined over the triangulation domain $|X|$. Assign each vertex $v \in X$ the weight

$|v|^2 - f(v)$, so that the x_{d+1}-coordinate of v^+ is $f(v)$. Let T be a weighted CDT of X, if one exists. The triangulation T and the vertex heights $f(v)$ define a piecewise linear surface $T^+ = \{s^+ : s \in T\}$. By analogy to witness functions (Definition 27), think of T^+ as a continuous piecewise linear function $T^+(p)$, which maps each point $p \in |X|$ to a real value. Because f is convex, every vertex in X is semiregular, so T^+ interpolates the lifted companion of every vertex in X, even if some vertices in X are missing from T.

Let $e(p) = T^+(p) - f(p)$ be the error in the interpolated function T^+ as an approximation of the true function f. At each vertex v in X, $e(v) = 0$. Because f is convex, the error satisfies $e(p) \geq 0$ for all $p \in |X|$.

Consider the unconstrained case first. T is the weighted Delaunay triangulation of the vertices in X, so T^+ is the underside of the convex hull of the lifted vertices. The intuition (formalized in Theorem 21 below) is that for any point $p \in |X|$, there is no way to triangulate the lifted vertices that yields a lesser value of $T^+(p)$ than the underside of the convex hull. Melissaratos' result follows immediately: T minimizes $\|e\|_{L_q}$ for every Lebesgue norm L_q.

The constrained case is only a little more complicated.

Theorem 21. *Let $f(p)$ be a function defined over the domain $|X|$ of a PLC X. Assign each vertex $v \in X$ the height $f(v)$—i.e. the weight $|v|^2 - f(v)$. If X has a weighted CDT, then at every point $p \in |X|$, every weighted CDT T of X minimizes $T^+(p)$ among all triangulations of X.*

Proof. Let T be a weighted CDT of X. Suppose for the sake of contradiction that there is a triangulation S of X and a point p such that $S^+(p) < T^+(p)$. Let s be the simplex in S whose relative interior contains p. Let t be a simplex in T that contains p and is not a submersible vertex. Let f be the lowest-dimensional facade in X that includes t. Because t is not a submersible vertex, f is not one either, so s respects f. Because p is in both f and the relative interior of s, $s \subseteq f$. Because s respects X and $p \in s$, the vertices of s are visible from p by Theorem 3.

Define the point $p_s = \langle p, S^+(p)\rangle \in E^{d+1}$. Thus $p_s \in s^+ \in S^+$, and p is the projected companion of p_s. Because $S^+(p) < T^+(p)$, p_s lies below t^+. For every witness d-flat h_t that includes t^+, at least one vertex of s^+ lies below h_t, because s^+ is a simplex that contains p_s. Therefore, t is not constrained semiregular within f. But by assumption, T is a weighted CDT of X, so t is constrained semiregular within f. By contradiction, there is not a triangulation S and a point p such that $S^+(p) < T^+(p)$. ■

Corollary 22. *Let $f(p)$ be a convex function defined over the domain $|X|$ of a PLC X. Assign each vertex $v \in X$ the height $f(v)$. If X has a weighted CDT, then at every point $p \in |X|$, every weighted CDT T of X minimizes the interpolation error $|T^+(p) - f(p)|$ among all triangulations of X.* ■

Because the weighted CDT minimizes the error $e(p)$ at every point, the weighted CDT minimizes e in every norm that is monotonic in e, including the Lebesgue norms. With the right choice of weights, this result holds for any convex function. Rippa also investigates the special case of interpolating $f(p) = Ax^2 + By^2 + Cx + Dy + E$ where $A \neq B$. For a function of this form, an anisotropic triangulation (with long, thin triangles) is optimal. Rippa suggests handling such functions by affinely mapping the vertices in E^d to a “stretched” space over which $f(p)$ is isotropic, finding the Delaunay triangulation of the mapped vertices, and mapping the triangulation back to the original space. Corollary 22 suggests an alternative: use weights to achieve the same effect as Rippa's mapping. This approach obtains exactly the same results when $f(p)$ is parabolic, but it is more flexible as it can adapt to other convex functions as well.

Corollary 22 plays a part in showing that Rajan's result generalizes to CDTs.

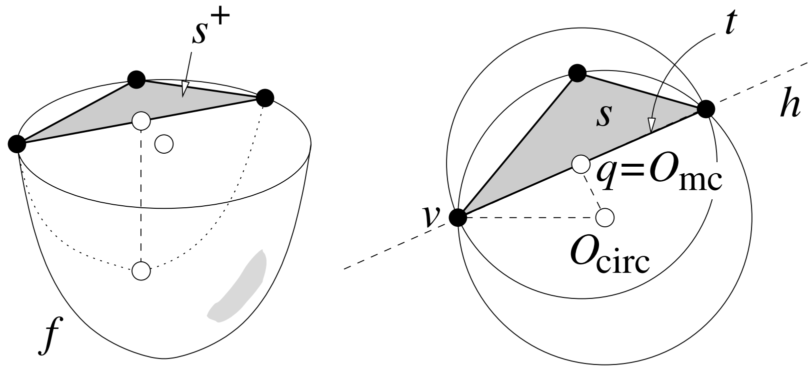

Figure 27: Left: within s, the error $e(p)$ is maximized at the point nearest the circumcenter of s. Right: top view of s, its circumcircle, and its min-containment circle.

Theorem 23. *If X has an unweighted CDT, then every unweighted CDT of X minimizes the largest min-containment sphere, compared with all other triangulations of X.*

Proof. Recall that $e(p) = T^+(p) - f(p)$. As X is unweighted, $f(p) = |p|^2$.

Over any single d-simplex s, there is an explicit expression for $e(p)$. Recall from the proof of Lemma 1 that the witness d-flat h_s that includes s^+ has the witness function $h_s(p) = 2O_{\mathrm{circ}} \cdot p - |O_{\mathrm{circ}}|^2 + r_{\mathrm{circ}}^2$, where O_{circ} and r_{circ} are the circumcenter and circumradius of s, and $p \in E^d$ varies freely. (The *circumcenter* and *circumradius* of s are the center and radius of s's circumsphere.) Hence, for all $p \in s$,

$$\begin{aligned} e(p) &= h_s(p) - f(p) \\ &= 2O_{\mathrm{circ}} \cdot p - |O_{\mathrm{circ}}|^2 + r_{\mathrm{circ}}^2 - |p|^2 \\ &= r_{\mathrm{circ}}^2 - |O_{\mathrm{circ}}p|^2. \end{aligned}$$

Figure 27 (left) illustrates the functions $h_s(p)$ and $f(p)$ over a triangle s. The error $e(p)$ is the vertical distance between the two functions. At which point p in s is $e(p)$ largest? At the point nearest the circumcenter, because $|O_{\mathrm{circ}}p|^2$ is smallest there. (The error is maximized at the circumcenter if the circumcenter is in s; Figure 27 gives an example where it is not.) Let O_{mc} and r_{mc} be the center and radius of the min-containment sphere of s, respectively. Lemma 24 below shows that the point in s nearest O_{circ} is O_{mc}, and $r_{\mathrm{mc}}^2 = e(O_{\mathrm{mc}})$.

It follows that the square of the min-containment radius of s is $\max_{p \in s} e(p)$, and thus the largest min-containment sphere of the entire triangulation has a squared radius of $\max_{p \in |T|} e(p)$. By Corollary 22, the unweighted CDT T minimizes this quantity among all triangulations of X. ■

Lemma 24. *Let O_{circ} and r_{circ} be the circumcenter and circumradius of a d-simplex s. Let O_{mc} and r_{mc} be the center and radius of the min-containment sphere of s. For $p \in s$, define the function $e(p) = r_{\mathrm{circ}}^2 - |O_{\mathrm{circ}}p|^2$. Let q be the point in s nearest O_{circ}. Then $O_{\mathrm{mc}} = q$ and $r_{\mathrm{mc}}^2 = e(q)$.*

Proof. Let t be the face of s whose relative interior contains q. The face t is not a vertex, because the vertices of s are s's furthest points from O_{circ}. Because q is the point in t nearest O_{circ}, and because q is in the relative interior of t, the line segment $O_{\mathrm{circ}}q$ is orthogonal to t. (This is true even if $t = s$, in which case $O_{\mathrm{circ}} - q = \mathbf{0}$.) This fact, plus the fact that O_{circ} is equidistant from all the vertices of t, implies that q is equidistant from all the vertices of t (as Figure 27 demonstrates). Let r be the distance between q and any vertex of t. Because $q \in t$, there is no containing sphere of t (or s) with radius less than r, because there is no direction q can move without increasing its distance from one of the vertices of t. Therefore, q and r are the center and radius of the min-containment sphere of t.

By the following reasoning, s has the same min-containment sphere as t. If $q = O_{\mathrm{circ}}$, this conclusion is immediate. Otherwise, let h be the hyperplane through q orthogonal to $O_{\mathrm{circ}}q$. Observe that h includes t. No point in s is on the same side of h as O_{circ}: if there were such a point w, there would be a point in s (between w and q) closer to O_{circ} than q, contradicting the fact that q is closest. Observe that h cuts the circumsphere into two pieces, and that the smaller piece encloses s and is enclosed by the min-containment sphere of t. Therefore, q and r are the center and radius of the min-containment sphere of s.

Let v be any vertex of t. Pythagoras' Law on $\triangle O_{\mathrm{circ}}qv$ (see Figure 27) yields $r_{\mathrm{circ}}^2 = r^2 + |O_{\mathrm{circ}}q|^2$, and therefore $r^2 = e(q)$. ■

For an algebraic proof of Lemma 24 (based on quadratic program duality), see Lemma 3 of Rajan [36].

The optimality of the CDT for controlling the largest min-containment radius dovetails nicely with an error bound for piecewise linear interpolation derived by Waldron [51]. Let C_c be the space of scalar functions defined over $|X|$ that have C^1 continuity and whose absolute curvature nowhere exceeds c. In other words, for every $f \in C_c$, every point $p \in |X|$, and every unit direction vector $\mathbf{d}$, the magnitude of the second directional derivative $f''_{\mathbf{d}}(p)$ is at most c. This is a common starting point for analyses of piecewise linear interpolation error. In contrast with Corollary 22, C_c is not restricted to convex functions.

Let f be a function in C_c. Let $s \subseteq |X|$ be a simplex (of any dimensionality) with min-containment radius r_{mc}. Let h_s be a linear function that interpolates f at the vertices of s. Waldron shows that for all $p \in s$, the absolute error $|e(p)| = |h_s(p) - f(p)|$ is at most $cr_{\mathrm{mc}}^2/2$. Furthermore, this bound is sharp: for every simplex s with min-containment radius r_{mc}, there is a function $f \in C_c$ and a point $p \in s$ such that $|e(p)| = cr_{\mathrm{mc}}^2/2$. (That function is $f(p) = c|p|^2/2$, as illustrated in Figure 27.)

Theorem 25. *Every unweighted CDT T of X (if any exist) minimizes*

$$\max_{f \in C_c} \max_{p \in |X|} |T^+(p) - f(p)|,$$

the worst-case pointwise interpolation error, among all triangulations of X.

Proof. For any triangulation T, $\max_{f \in C_c} \max_{p \in |X|} |T^+(p) - f(p)| = cr_{\max}^2/2$, where $r_{\max}$ is the largest min-containment radius among all simplices in T. The result follows immediately from Theorem 23. ■

One of the reasons why CDTs are important is because, in the senses of Corollary 22 and Theorem 25, the CDT is an optimal piecewise linear interpolating surface. Of course, $e(p)$ is not the only criterion for the merit of a triangulation used for interpolation. Many applications need the interpolant to approximate the gradient—that is, not only must $T^+(p)$ approximate $f(p)$, but $\nabla T^+(p)$ must approximate $\nabla f(p)$ well too. For the goal of approximating $\nabla f(p)$ in three or more dimensions, the weighted CDT is sometimes far from optimal even for simple functions like the paraboloid $f(p) = |p|^2$. Still, the CDT is a good starting point for mesh improvement algorithms [6, 7, 10, 11, 16, 30, 46, 48] that create a triangulation that is appropriate for approximating both $f(p)$ and $\nabla f(p)$.

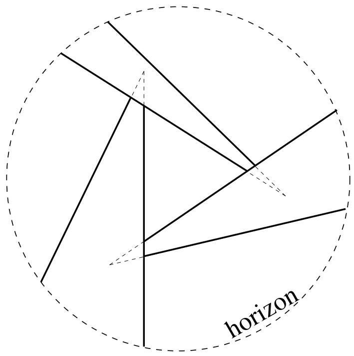

Figure 28: Spherical projection of the halfspace above your vantage point.

5 Proof of the CDT Theorem for Generic PLCs

Theorem 26. *Let X be a generic, weakly ridge-protected, d-dimensional PLC (weighted or not). X has a CDT (a weighted CDT if X is weighted).*

This section is devoted to the proof of Theorem 26, the generic version of the *CDT Theorem*. A lot of ink must be spilt for it, and readers who are not feeling athletic are invited to skip to Section 6, where the genericity requirement is removed from Theorem 26.

Half the work is already done: Corollary 18 states that the constrained regular simplices form a simplicial complex, and Theorem 10 states that if this complex fills X, it is a CDT of X. The most difficult part of the proof is to show that if X is generic and weakly ridge-protected, the complex fills X. The forthcoming Theorem 30 shows that every point in a weakly ridge-protected PLC lies in some constrained semiregular simplex. Unfortunately, several long proofs are needed to build up to that result.

5.1 Visibility Lemmata

One potential difficulty for the CDT Theorem is illustrated in Figure 28. Imagine that you are standing at a point p in the interior of a three-dimensional domain, scanning the halfspace "above" p for a visible vertex. Looking up into the sky, you see the three illustrated 2-facades, each of which occludes the apical vertex of another. The remaining vertices of these facades are below the horizon (in the halfspace below you). No vertex in the halfspace is visible from your vantage point, so there is no constrained semiregular simplex that contains p.

To prove the existence of a CDT, one must show that this possibility is precluded if X is weakly ridge-protected. Fortunately, Lemma 14 does exactly that. By the definition of "weakly ridge-protected," every grazeable constraining facade in X of dimension $d-2$ or less is a union of regular simplices. In Figure 28, observe that the inner edges of the three facades form a cycle of overlapping edges. These edges are grazeable. However, Lemma 14 implies that the overlap relation among regular simplices (from a fixed viewpoint) constitutes a partial order. The regular edges bounding the 2-facades cannot form a cycle. This fact is the key to proving two key lemmata for weakly ridge-protected PLCs.

For each regular simplex s, let h_s be a witness to the regularity of s. Every lifted vertex lies above h_s, except the vertices of s^+. Recall from Definition 27 the witness function $h_s(p)$, a linear function that maps each point $p \in E^d$ to the x_{d+1}-coordinate such that $\langle p, h_s(p)\rangle \in h_s$. If s is not d-dimensional, it has infinitely many witness d-flats; choose one arbitrarily so that $h_s(p)$ is consistently defined.

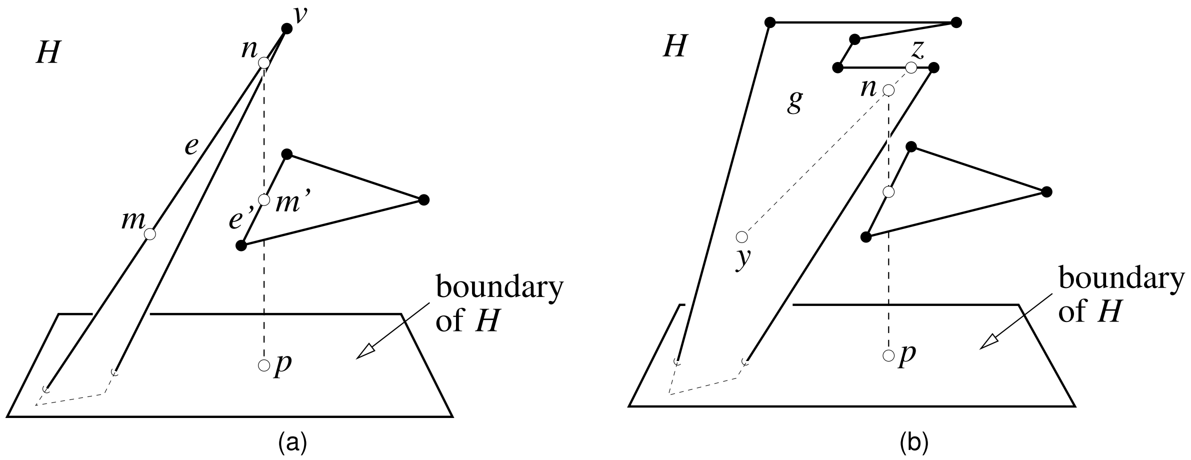

Figure 29: The supposition that no vertex in H is visible from p leads to a contradiction.

Lemma 27. *Let X be a weakly ridge-protected, d-dimensional PLC. Let p be a point in the interior of $|X|$. Let H be an open d-dimensional halfspace whose closure contains p. At least one vertex of X is in H and visible from p.*

Proof. Suppose for the sake of contradiction that no vertex of X is in H and visible from p. Let A be the set containing every simplex e that has the following properties:

- e respects X and is regular within X, and
- there is a point m in e's relative interior such that $m \in H$ and m is visible from p.

A is empty—suppose for the sake of contradiction that it is not. Because no vertex of X is in H and visible from p, A contains no vertex. Let e be the simplex in A that maximizes $h_e(p)$. Let m be a point in e's relative interior that is in H and visible from p. Because e is a simplex that intersects H, at least one vertex v of e is in H, as Figure 29(a) shows. (The other vertices of e might lie below the horizon, outside H.)

By assumption, v is not visible from p, although m is. 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 on a "walk" from m to v.[8] The line segment pn must intersect some occluding facade of X at some point m'. If several facades occlude the view of n from p, consider only the facade that intersects pn closest to p, so that m' is visible from p.

Let f be the face of that facade whose relative interior contains m'. (In Figure 29, f is the edge e'.) Because n is the first occluded point on the walk from m to v, f must have dimension $d-2$ or less (i.e. m' cannot lie in the relative interior of a $(d-1)$-facade). Because no vertex is in H and visible from p, f is not a vertex. The grazing triangle $\triangle pnm$ demonstrates that f is grazeable. As X is weakly ridge-protected, f has a triangulation whose simplices respect X and are regular within X. Let e' be the simplex in that triangulation

[8] How do we know that there *is* a first occluded point on the walk from m to v, rather than a last visible point? On the walk, there is at least one transition from points p can see to points p cannot see. Let n be the point where the first such transition occurs. Is n visible from p? There are two ways that a transition might occur. One possibility is an interposing facade that occludes the visibility of n, as in Figure 29(a). The second possibility is that n lies on a facade f and is visible from p, but the points following n on the walk are occluded by f. To exclude this possibility, observe that e is convex, m is in e's relative interior, $v \in e$, $n \in mv$, and $n \neq v$. Therefore, n must lie in e's relative interior. Because e respects X, every facade that contains n includes e, and therefore f cannot occlude the visibility of any point in e from anywhere. These details are dreary, but the proof depends on them.

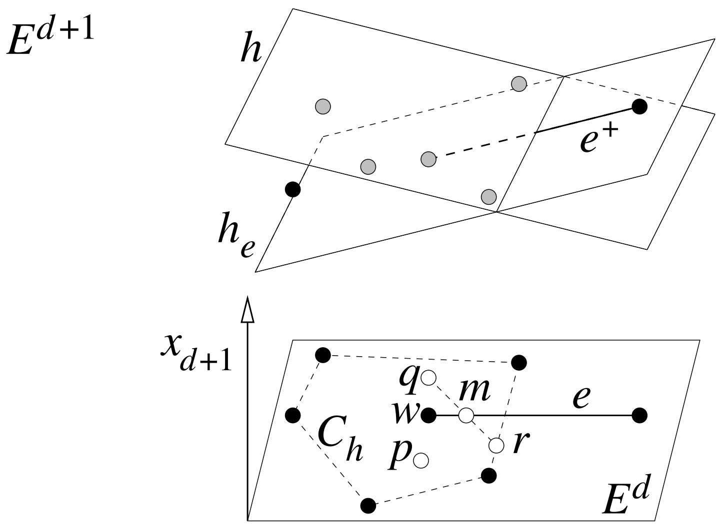

Figure 30: Because m is between q and r, $h(m) > h_e(m)$.

whose relative interior contains m'. Because n lies in H and p lies in its closure, m' lies in H, so $e' \in A$ (by the definition of A). Because $n \in e$, e' overlaps e from the viewpoint p, and therefore $h_{e'}(p) > h_e(p)$ by Lemma 14.

But this contradicts the assumption that e maximizes $h_e(p)$ among all members of A. It follows that A is empty.

Because p is in the interior of $|X|$, at least one facade in X intersects H. Let g be the lowest-dimensional facade in X whose relative interior contains a point y that is in H and visible from p. By assumption, g is not a vertex.

Because g intersects H, at least one vertex of g is in H. Imagine shooting a ray from y toward that vertex. Let z be the first point on the boundary of g struck by the ray, as illustrated in Figure 29(b). As g might not be convex, z might not be the vertex, but z is in H. Because g is the lowest-dimensional facade whose relative interior contains a point in H visible from p, and z lies in the relative interior of a proper face of g, z is not visible from p. Let n be the first occluded point encountered on a "walk" from y to z. By a repetition of the reasoning above, some simplex in A is interposed between p and n. But A is empty, so this is a contradiction.

It follows that some vertex of X is in H and visible from p. ■

A second lemma reveals a more subtle (and barely comprehensible) property of visibility in PLCs.

Lemma 28. *Let X be a weakly ridge-protected, d-dimensional PLC. Let $h \subset E^{d+1}$ be a non-vertical d-flat. Let $V_h = \{v \in X : v$ is a vertex and v^+ is on or below $h\}$, and let $C_h = \mathrm{conv}(V_h)$. (See Figure 30. Note that V_h and C_h are sets of points in E^d, not E^{d+1}.)*

Let p be a point in C_h. Suppose that no vertex in X visible from p lifts to a point below h.

Let $f \in X$ be a grazeable constraining facade of dimension $d - 2$ or less. Suppose that some point $m_f \in f \cap C_h$ is visible from p.

Then f includes the face of C_h whose relative interior contains m_f. (This face may be C_h itself.)

Proof. Because X is weakly ridge-protected, f has a triangulation whose simplices respect X and are regular within X. Let t be the simplex in this triangulation whose relative interior contains m_f.

Let A be the set containing every simplex e that has the following properties:

- e respects X and is regular within X, and
- there is a point m in e's relative interior such that
  - m is visible from p,
  - $m \in C_h$, and
  - e does not include the face of C_h whose relative interior contains m.

If A is empty, then $t \notin A$, so t includes the face of C_h whose relative interior contains m_f, and the lemma holds. Suppose for the sake of contradiction that A contains at least one simplex.

Let e be the simplex in A that maximizes $h_e(p)$. As $e \in A$, there is a point m in the relative interior of e such that $m \in C_h$ and m is visible from p. Because e is regular, there is a witness d-flat $h_e \subset E^{d+1}$ that includes e^+, as illustrated in Figure 30. Each vertex of e lifts to a point on h_e. Every other vertex in X lifts to a point above h_e.

For each vertex $v \in V_h$, v^+ lies on or below h, and on or above h_e, so $h(v) \geq h_e(v)$. If v is in V_h but not in e, then v^+ lies strictly above h_e, so $h(v) > h_e(v)$.

Because C_h is the convex hull of V_h, and h and h_e are linear functions, it follows that for each point $q \in C_h$, $h(q) \geq h_e(q)$, and if q is not in e, then $h(q) > h_e(q)$.

Let C_m be the face of C_h whose relative interior contains m. By assumption, e does not include C_m, so some point $q \in C_m$ is not in e. Because m is in the relative interior of C_m, there is a point $r \in C_m$ such that m is between q and r. (See Figure 30.) Thus $h(q) > h_e(q)$ and $h(r) \geq h_e(r)$, so by the linearity of h and h_e, $h(m) > h_e(m)$.

Because e is a simplex that contains m, there must be at least one vertex w of e for which $h(w) > h_e(w)$. Because w^+ lies on h_e, w^+ lies below h, so $w \in V_h$ (by the definition of V_h). By assumption, no vertex visible from p lifts to a point below h, so w is not visible from p. But recall that $m \in e$ is visible from p. Can m be visible from p if w is not?

Let n be the point nearest m on the line segment mw that cannot see p, as illustrated in Figure 31. The line segment pn must intersect some facade in X at some point m'. If there are several facades occluding the view of n from p, consider only the facade that intersects pn closest to p, so that m' is visible from p.

Let g be the face of that facade whose relative interior contains m'. (In Figure 31, g is the edge e'.) Because n is the first occluded point on the walk from m to w, g must have dimension $d - 2$ or less (i.e. m' cannot lie in the relative interior of a $(d-1)$-facade). The grazing triangle $\triangle pnm$ demonstrates that g is grazeable. As X is weakly ridge-protected, g has a triangulation whose simplices respect X and are regular within X. Let e' be the simplex in that triangulation whose relative interior contains m'. Observe that $n \in C_h$ because n lies between m and w, which are both in C_h. Moreover, $m' \in C_h$ because m' lies between n and p. Let $C_{m'}$ be the face of C_h whose relative interior contains m'. By Lemma 15 (substituting C_h for C, $C_{m'}$ for C_m, and pn for P), $pn \subseteq C_{m'}$. Because g occludes the visibility between p and n, e' contains neither p nor n. It follows that $C_{m'} \nsubseteq e'$.

By the definition of A, $e' \in A$. Because e' overlaps e from the viewpoint p, $h_{e'}(p) > h_e(p)$ by Lemma 14. But this contradicts the assumption that e maximizes $h_e(p)$ among all members of A. It follows that A is empty, and the lemma holds. ■

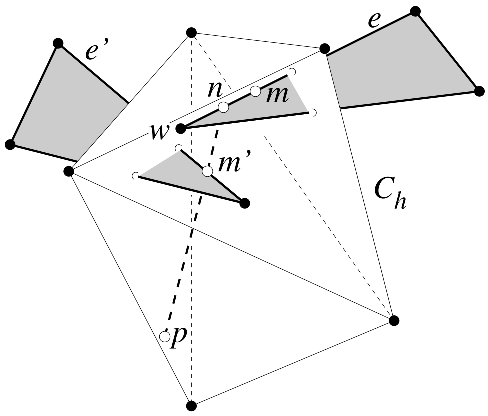

Figure 31: Because m is visible from p and w is not, some simplex e' must overlap e.

5.2 Ridge-Protected PLCs Are Filled

This section completes the proof of Theorem 26. Most of the effort is spent proving that if a PLC is weakly ridge-protected, every point in the triangulation domain lies in some constrained semiregular simplex. The proof is made easier by considering a subset of the triangulation domain first—a set of points from which visibility is particularly well behaved.

For a d-dimensional PLC X, let N be the set containing every point in the interior of $|X|$ that is not cohyperplanar with any d affinely independent vertices in X. No point in N lies on any constraining facade, nor on any k-simplex whose vertices are in X for $k < d$, nor on their affine hulls. The closure of N is the union of the d-facades in X.

Lemma 29. *Let X be a weakly ridge-protected, d-dimensional PLC. Define N as above, and let p be a point in N. Some constrained semiregular d-simplex contains p.*

Proof. Let W be the set of all vertices in X visible from p (the black vertices in Figure 32, far left). The following reasoning establishes that p is in $\text{conv}(W)$. Suppose for the sake of contradiction that it is not. Then there is an open halfspace H such that p lies on the boundary of H and $W \cap H = \emptyset$.[9] But by Lemma 27, some vertex of X is in H and visible from p. This vertex is in $W \cap H$, a contradiction.

Let s be the d-simplex that contains p in a weighted Delaunay triangulation of W (Figure 32, second from left). Because $p \in \text{conv}(W)$, some such simplex must exist. The rest of this proof shows that s is constrained semiregular within X, so the lemma holds.

Let h_s be the unique witness to the semiregularity of s within W. No vertex in W lifts to a point below h_s, so no vertex in X visible from p lifts to a point below h_s. Let $V_h = \{v \in X : v \text{ is a vertex and } v^+ \text{ lies on or below } h_s\}$. Let C_h be the convex hull of V_h (Figure 32, center). Observe that the vertices of s are in V_h, so $s \subseteq C_h$ and $p \in C_h$.

Let Z be the closure of the set of all points that p can see in the triangulation domain $|X|$ (Figure 32, second from right). Because $p \in N$, p lies in the interior of $|X|$, which implies that Z is d-dimensional with p in its interior. The vertices of s are in Z.

Because Z is the closure of points visible from p, the shadows cast by constraining facades of dimension $d - 2$ or less have no effect on Z. Z is a star-shaped polyhedron (not generally convex) with two types

[9]This claim is intuitive, but its formal proof is tricky. It is the well-known Farkas Lemma; see Ziegler [54] for a proof.

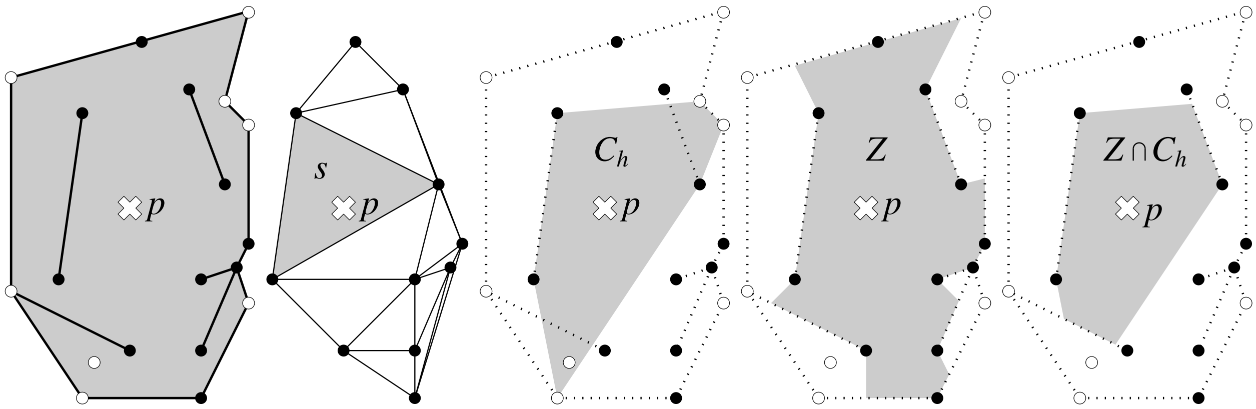

Figure 32: From left to right: a PLC X wherein the vertices visible from p (the set W) are colored black, and the vertices not visible from p are colored white. The weighted Delaunay triangulation of W, and the simplex s therein that contains p. C_h is the convex hull of the vertices whose lifted companions lie on or below the witness d-flat for s. Z is the closure of the set of points visible from p. $Z \cap C_h$ is convex and respects X.

of facets: portions of $(d-1)$-facades, and *shadow facets* that are cohyperplanar with p because they are boundaries of shadows cast by occluding $(d-1)$-facades.

The rest of this proof is a sequence of claims and their justifications.

Claim: no constraining $(d-1)$*-facade in X intersects the interior of Z.* Because $p \in N$, p is not cohyperplanar with any $(d-1)$-facade, so every $(d-1)$-facade casts a shadow (occludes visibility from p) and no $(d-1)$-facade intersects the interior of Z.

Claim: $Z \cap C_h$ *is a star-shaped d-polyhedron.* Because $p \in N \cap C_h$, p is in the interior of C_h. Because Z and C_h are both closed star-shaped d-polyhedra with p in their kernels and in their interiors, so is $Z \cap C_h$.

Claim: no constraining facade in X intersects the interior of $Z \cap C_h$. Suppose for the sake of contradiction that a constraining facade f intersects the interior of $Z \cap C_h$. Let m be a point in the intersection of f's relative interior and the interior of $Z \cap C_h$. Assume without loss of generality that m is visible from p—if it is not, then m's visibility is occluded by some other constraining facade that intersects the interior of $Z \cap C_h$ closer to p (because $Z \cap C_h$ is star-shaped with p in its kernel), so f and m can be replaced by the occluding facade and the closer intersection point.

Because no constraining $(d-1)$-facade intersects the interior of Z, f must have dimension $d-2$ or less. To show that f is grazeable, choose an open grazing triangle L that does not intersect any constraining facade, such that one boundary edge of L contains m. Does such a triangle always exist? If L has m on its boundary and is sufficiently small, the only constraining facades that can intersect L are those that contain m. These facades intersect the interior of Z, so they have dimension $d-2$ or less. Almost every plane (2-flat) through m intersects these facades only at the point m. (Here, "almost every" is used in the analytic sense: for any $(d-2)$-facade g that contains m, the set of planes through m that intersect $g\backslash\{m\}$ has measure zero in the space of planes through m.) Therefore, almost every sufficiently small open triangle with m on its boundary intersects no constraining facade, so f has a grazing triangle.

By Lemma 28, $f \supseteq C_h$. This contradicts the fact that C_h is d-dimensional and f is at most $(d-2)$-dimensional, so no constraining facade intersects the interior of $Z \cap C_h$.

Claim: $Z \cap C_h$ *is convex.* (See Figure 32, far right.) Suppose for the sake of contradiction that $Z \cap C_h$ is not convex. Then there exist two points q and r in the interior of $Z \cap C_h$ such that $qr \nsubseteq Z \cap C_h$. Because

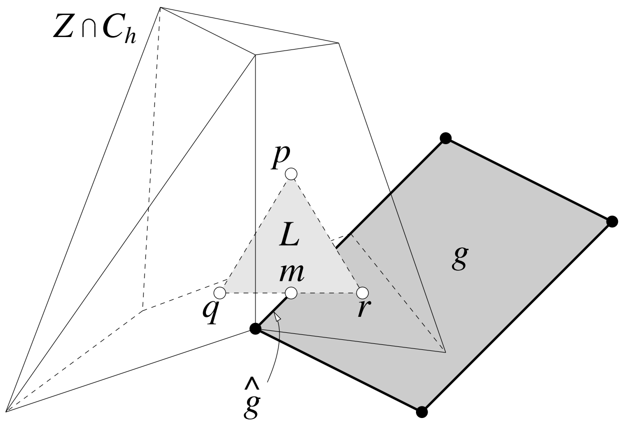

Figure 33: If $Z \cap C_h$ is not convex, its boundary incorporates a grazeable facade $\hat{g}$.

$Z \cap C_h$ is star-shaped with p in its kernel, $Z \cap C_h$ includes both pq and pr, so the three points p, q, and r cannot be collinear. Continuously move q and r toward p until $\triangle pqr \subset Z \cap C_h$, but qr still intersects the boundary of $Z \cap C_h$, as illustrated in Figure 33. Let m be the point nearest q on qr that lies on the boundary of $Z \cap C_h$. (That point is neither q nor r, which are in the interior.) Loosely speaking, $Z \cap C_h$ is locally reflex at m. Because C_h is convex with q and r in its interior, m also lies in the interior of C_h, so m must lie on the boundary of Z.

Because the open triangle $L = \triangle pqr$ is included in the interior of $Z \cap C_h$, which intersects no constraining facade, L is a grazing triangle for m, and m is visible from p. Because m lies on the boundary of Z, but the open line segment pm does not intersect Z's boundary, m lies on at least one facet of Z that is not a shadow facet. Therefore, m lies on some $(d-1)$-facade g, as illustrated. Because g intersects neither the open triangle L nor the open line segment qm, m must lie on the boundary of g.

Let $\hat{g}$ be the face of g whose relative interior contains m. Because m is on g's boundary, $\hat{g}$ has dimension $d - 2$ or less. L demonstrates that $\hat{g}$ is grazeable. By Lemma 28, $\hat{g} \supseteq C_h$. This contradicts the fact that C_h is d-dimensional and $\hat{g}$ is at most $(d - 2)$-dimensional, so $Z \cap C_h$ is convex.

Claim: $Z \cap C_h$ has no shadow facets. This claim follows because shadow facets are cohyperplanar with p, but $Z \cap C_h$ is a convex d-polyhedron with p in its interior.

Claim: $Z \cap C_h$ respects X. Suppose for the sake of contradiction that some facade $f \in X$ (that is not a submersible vertex) intersects the relative interior of a face P of $Z \cap C_h$, but f does not include P. Let y be a point in the intersection of f with the relative interior of P, as illustrated in Figure 34.

Because f is closed and does not include P, there is a point z in the relative interior of P that is not in f. Let m be the point nearest z in $f \cap yz$, as illustrated. Because y and z are in the relative interior of P, so is m. Let y' be a point in P such that m is between y' and z. (The choice $y' = y$ will do if $y \neq m$; but if $y = m$, choose y' just past m on the ray $\vec{zm}$.) Let $\hat{f}$ be the face of f (possibly f itself) whose relative interior contains m. This choice guarantees that y' and z do not lie on the affine hull of $\hat{f}$, and $\hat{f}$ cannot have dimension d.

The facade $\hat{f}$ cannot have dimension $d - 1$, either. If it did, then it would intersect the interior of $Z \cap C_h$, because m is in the relative interiors of both $\hat{f}$ and $y'z$, $y'z$ is on the boundary of $Z \cap C_h$, and $y'z$ does not lie on the same hyperplane as $\hat{f}$. But no $(d - 1)$-facade intersects the interior of Z. Therefore, $\hat{f}$ has dimension $d - 2$ or less. To show that $\hat{f}$ is grazeable, choose an open grazing triangle L such that L is included in the interior of $Z \cap C_h$, and $y'z$ is an edge of (the closure of) L. No constraining facade intersects the interior of $Z \cap C_h$, so L is indeed a grazing triangle.

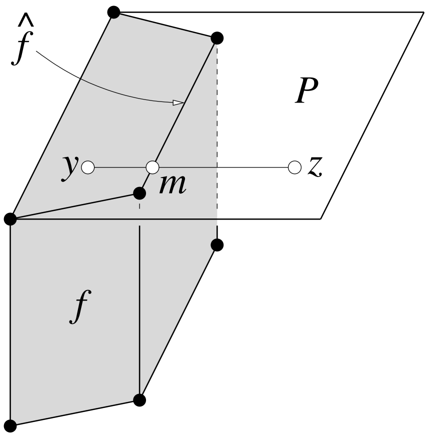

Figure 34: The circumstance depicted here, where a facade f intersects the interior of a face P of $Z \cap C_h$ but does not include P in its entirety, cannot happen in a weakly ridge-protected PLC.

Let C_m be the face of C_h whose relative interior contains m. By Lemma 28, $\hat{f} \supseteq C_m$. Recall that m lies in the relative interior of P, which is a face of $Z \cap C_h$, which implies that $P \subseteq C_h$. By Lemma 15, $P \subseteq C_m$. Thus $z \in P \subseteq C_m \subseteq \hat{f} \subseteq f$, contradicting the fact that z is not in f. The claim that $Z \cap C_h$ respects X follows.

Claim: $s \subseteq Z \cap C_h$. This claim follows because both Z and C_h contain all the vertices of s, and $Z \cap C_h$ is convex.

Claim: s respects X. Let t be any face of s, and let m be any point in the relative interior of t. Suppose some facade $f \in X$ (that is not a submersible vertex) contains m. As $m \in t \subseteq s \subseteq Z \cap C_h$, let C_m be the face of $Z \cap C_h$ whose relative interior contains m. Because t and $Z \cap C_h$ are convex with $t \subseteq Z \cap C_h$, it follows from Lemma 15 (substituting t for P and $Z \cap C_h$ for C) that $t \subseteq C_m$.

Recall that $Z \cap C_h$ respects f: if f intersects the relative interior of a face of $Z \cap C_h$, then f includes the whole face. Because f intersects the relative interior of C_m (at m), f includes C_m, which implies that $t \subseteq f$. This relationship holds for any face t of s, any point m, and any facade $f \in X$ that satisfy the assumptions, so s respects X.

Claim: Every point in $Z \cap C_h$ can see every other point in $Z \cap C_h$, but no point in the interior of $Z \cap C_h$ can see any vertex of V_h not in $Z \cap C_h$. The first half of this claim follows from Theorem 3 because $Z \cap C_h$ respects X. For the second half of the claim, let q be a point in the interior of $Z \cap C_h$, and let v be a vertex in V_h that is not in $Z \cap C_h$. Some facet F of $Z \cap C_h$ lies between q and v. Because v is in C_h (which is convex) and q is in its interior, F is not on the boundary of C_h. Thus F must lie on the boundary of Z. Because $Z \cap C_h$ has no shadow facets, F must be included in some $(d-1)$-facade in X, which occludes the visibility of v from q. Therefore, no point in the interior of $Z \cap C_h$ can see any vertex of V_h not in $Z \cap C_h$.

Claim: s is constrained semiregular. Because $p \in Z \cap C_h$, p sees every vertex in $Z \cap C_h$. By construction, no vertex visible from p has a lifted companion below the witness h_s; therefore, no vertex in $Z \cap C_h$ has one. By the definition of V_h, every vertex in X whose lifted companion is below h_s is in V_h. By the previous claim, no point in the interior of $s \subseteq Z \cap C_h$ can see any vertex of V_h not in $Z \cap C_h$. Therefore, no point in the interior of s can see any vertex whose lifted companion is below h_s. Moreover, s respects X, so s is constrained semiregular. ■

Theorem 30. *Let X be a weakly ridge-protected, d-dimensional PLC. Let p be a point in a d-facade in X. Some constrained semiregular d-simplex contains p.*

Proof. If $p \in N$, the result follows from Lemma 29. What about points not in N? Every point in N lies in some closed constrained semiregular d-simplex, and the closure of N is the union of all the d-facades in X. It follows that every point in every d-facade in X lies in some constrained semiregular d-simplex. ■

Theorem 30 provides the machinery to prove Theorem 26: if X is a generic, weakly ridge-protected, d-dimensional PLC, then X has a CDT.

Proof of Theorem 26. Let T be the set that contains every simplex that is constrained semiregular within X or within a constraining facade in X. Because X is generic, every constrained semiregular simplex is constrained regular, and Corollary 18 guarantees that T is a simplicial complex.

Let p be any point in the triangulation domain $|X|$. Let f be the highest-dimensional facade in X that contains p, and let k be the dimensionality of f. If $k = d$, Theorem 30 states that there exists a constrained semiregular d-simplex that contains p. By the definition of T, this d-simplex is in T.

If $k < d$, f is a dangling facade. Let Y_f be the k-dimensional facade PLC for f. By Corollary 12, Y_f is weakly ridge-protected. Therefore Theorem 30 applies, with Y_f substituted for X and k substituted for d. In this case the theorem states that some k-simplex exists that contains p and is constrained semiregular within Y_f. This k-simplex is in T.

Because such a simplex exists for every point $p \in |X|$, T fills X. By Theorem 10, T is a CDT of X. ■

Theorem 26 requires X to be generic only to ensure that Corollary 18 applies. If X is nongeneric, T may contain constrained semiregular simplices whose interiors overlap. Theorem 30, however, holds even for nongeneric X.

6 Nongeneric PLCs, Weight Perturbations, and the CDT Theorem

It is well known that the Delaunay triangulation is not unique when $d + 2$ or more vertices lie on a common empty hypersphere. Every affinely independent subset of these cospherical vertices yields a Delaunay simplex. Some of the Delaunay simplices have mutually overlapping relative interiors, so some Delaunay simplices must be omitted to form a proper triangulation. Different choices yield different Delaunay triangulations. Likewise, a weighted Delaunay triangulation is not unique when the underside of the convex hull of the lifted vertices has a facet that is not a simplex.

The story is a bit more complicated for CDTs and weighted CDTs. The triangulation domain of a PLC might have a polyhedral gap (not necessarily convex) that is not covered by constrained regular simplices. Sometimes this happens simply because the PLC has no CDT, but sometimes the gap can be triangulated with constrained semiregular simplices. A gap might have several such triangulations, yielding multiple CDTs of one PLC. If a gap is shaped like Schönhardt's polyhedron, it cannot be triangulated at all.

A generic PLC has at most one weighted CDT (by Corollary 20), consisting of every constrained regular simplex (by Theorem 19), so it is pleasingly unambiguous. A nongeneric PLC raises the question of whether there exists a set of constrained semiregular simplices that fill the gaps and complete the triangulation. Because there may be several choices of constrained semiregular simplex to cover any point in a gap, determining whether a CDT exists is like solving a jigsaw puzzle with extra, useless pieces included.

Surprisingly, the problem of determining whether a three-dimensional nongeneric PLC has a CDT is NP-complete [24], even for an unweighted PLC. By contrast, it is always possible to determine whether a generic PLC has a CDT in polynomial time—by attempting to construct it. (See the second article in this series for further discussion.)

This section removes the genericity requirement from the CDT Theorem by perturbing the vertex weights so that no $d+2$ vertices lift to a common non-vertical d-flat. The vertex coordinates are not perturbed. If the perturbed PLC has a CDT, the latter is also a CDT of the original, unperturbed PLC. The method works even for unweighted PLCs, by temporarily assigning each vertex a tiny weight. This idea first appears in the work of Edelsbrunner and Mücke [17, Section 5.4].

The weight perturbation method serves a practical function as well as a theoretical one. The third article in this series describes an easy way to implement the perturbations to ensure the correctness of algorithms for constructing and updating CDTs. There is a catch, though. Will a PLC that has a CDT still have a CDT after it is perturbed? Not necessarily. Perturbations cannot circumvent the NP-hardness result.

The perturbations are *symbolic*—the magnitudes of the perturbations are not explicitly specified. Following Edelsbrunner and Mücke, the ith vertex weight could be perturbed by ϵ^{2^i} for a sufficiently small ϵ, but the proofs are simpler if the perturbations are implicitly chosen by the following procedure instead.

Let X be a d-dimensional PLC. Let V be the set of vertices in X. Consider all the $(d+1)$-simplices, including degenerate ones, that can be defined by taking subsets of $d+2$ lifted vertices from V^+. Call these the *orientation simplices*. Assume that the vertices of each orientation simplex are listed in some canonical order. The signed volume of an orientation simplex $\langle v_0^+, v_1^+, \ldots, v_{d+1}^+\rangle$ is $1/(d+1)!$ times the determinant of the matrix with column vectors $v_1^+ - v_0^+, v_2^+ - v_0^+, \ldots, v_{d+1}^+ - v_0^+$. Each signed volume varies linearly with the vertex weights. Every question about whether a lifted vertex lies above a witness d-flat for a d-simplex is a question about the sign of the volume of an orientation simplex. A volume of zero indicates cohyperplanarity.

Perturb the weights of the vertices in V one at a time, in some arbitrary order, each by a tiny negative or positive amount (different for each vertex). To perturb the weight of a vertex v, choose the magnitude of the perturbation to be sufficiently small that no orientation simplex's signed volume changes from positive to nonpositive, or from negative to nonnegative. Some signed volumes may change from zero to nonzero—that is the goal of the perturbations. Once a signed volume becomes nonzero, subsequent perturbations are not permitted to change its sign. The idea is to move vertices off of witnesses, but never to move a vertex from above a witness to below, nor vice versa. For each vertex in turn, it is always possible to choose a nonzero perturbation small enough to satisfy these restrictions. Perturb every vertex once.

Theorem 31. *Let X be a PLC. Let X' be a weighted PLC defined by perturbing every vertex weight in X as described above. (If X is unweighted, assign each vertex a weight of zero before perturbing it.) The following statements hold.*

A. If a simplex s is regular within X, it is regular within X'.
B. If a simplex s is constrained regular within a facade in X, it is constrained regular within the same facade in X'.
C. If a simplex s is regular within X', it is semiregular within X.
D. If a simplex s is constrained regular within a facade in X', it is constrained semiregular within the same facade in X.
E. X' is generic and has at most one CDT.
F. If X is ridge-protected, so is X'.
G. If X is weakly ridge-protected, so is X'.
H. If X' has a CDT, the CDT of X' is a CDT of X.
I. If X is generic, X and X' have the same CDT (or lack thereof).

Proof. A lifted vertex lies above, on, or below the witness for a d-simplex according to whether the signed volume of some orientation simplex is positive, zero, or negative. Similarly, the regularity of any lower-dimensional simplex depends on the volumes of certain orientation simplices all having the right sign. Because a perturbation never changes the volume of any orientation simplex from positive to nonpositive or from negative to nonnegative, Statements A, B, C, and D hold by induction on the sequence of perturbations.

Perturbing the height of a vertex v moves v^+ off of any non-vertical d-flat that it lay on before the perturbation. No perturbation, of v or any other vertex, can move v^+ onto a witness d-flat that v^+ did not lie on before the perturbation, because that would imply that the volume of some orientation simplex changes from nonzero to zero. Therefore, v^+ does not lie on any witness immediately after it is perturbed, except the witnesses that by definition pass through v^+; and subsequent perturbations preserve this claim. By induction on the sequence of vertex perturbations, the claim holds for every vertex in X', and X' is generic. By Corollary 20, X' has at most one CDT.

Statements F and G follow from A. Statement H follows from D and the genericity of X'.

If X is generic, then constrained regularity and constrained semiregularity are equivalent. Thus, Statement B implies that any CDT of X is a CDT of X', just as Statement H says that any CDT of X' is a CDT of X. Either X and X' both have the same CDT, or both have no CDT. ■

A CDT of X' is a CDT of X, but if X is nongeneric, different perturbations of X (i.e. perturbing the vertices in a different order, or using different mixtures of positive and negative perturbations) may yield different CDTs of X, or no CDT at all. Nevertheless, any choice of perturbation faithful to the procedure described above suffices to excise the genericity requirement from Theorem 26.

Theorem 32 (CDT Theorem)**.** *Let X be a weakly ridge-protected, d-dimensional PLC (weighted or not). X has a CDT (a weighted CDT if X is weighted).*

Proof. Let X' be the perturbed weighted PLC defined in Theorem 31. By the theorem, X' is generic and weakly ridge-protected, so by Theorem 26, X' has a CDT T. By Theorem 31, T is a CDT of X. ■

7 Conclusions

In their article on two-dimensional conforming Delaunay triangulations, Edelsbrunner and Tan [20] 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.

Three-dimensional CDTs shift the emphasis back to the former of these two problems. An algorithm that could create a Steiner CDT by inserting only a polynomial number of additional vertices would be an exciting development.

Some applications of finite element methods use meshes that have open slits, which are infinitesimally thin fissures across which information does not flow. The ideas in this article seem to extend in a straightforward way to topological PLCs wherein open slits are modeled by topological holes in the domain. Unfortunately, it is difficult to describe these PLCs in simple geometric terms, because of the need to distinguish topologically distinct points that have the same coordinates. For example, an internal $(d-1)$-facade can be converted into an open slit by making a topologically distinct copy of the facade that coincides with the original. Both the original and the copy adjoin the exterior domain (the infinitesimally thin hole), but they adjoin each other only along their external boundaries. The internal vertices in the original facade are topologically distinct from the internal vertices in the copy (and may or may not coincide), thereby supporting the interpolation of discontinuous functions as illustrated in Figure 1 (right). The open question is how to formulate these topological PLCs rigorously, and how to extend the results in this article to them.

Several other questions deserve investigation. Is there a simply stated and tested condition that is both sufficient and necessary for a generic PLC to have a CDT? The NP-hardness result suggests that there is no such condition for nongeneric PLCs. Is there a less conservative definition of "constrained Delaunay" (perhaps giving more power to constraining facades of dimension less than $d-1$) that admits useful, well-defined triangulations over a larger class of PLCs? Is there a better approach to assuring the existence of a CDT than to make a PLC weakly ridge-protected? Finally, when do curved manifold complexes (e.g. the stratifications mentioned in Section 2.1) have CDTs?

Acknowledgments

I thank Dafna Talmor and Herbert Edelsbrunner for helpful discussions. In particular, Dafna pointed out the duality between degenerate faces of the Voronoi diagram and simplices that are Delaunay but not strongly Delaunay.

References

[1] Oswin Aichholzer, Franz Aurenhammer, Hannes Krasser, and Peter Brass. *Pseudotriangulations from Surfaces and a Novel Type of Edge Flip*. SIAM Journal on Computing **32**(6):1621–1653, 2003.

[2] Franz Aurenhammer. *Power Diagrams: Properties, Algorithms, and Applications*. SIAM Journal on Computing **16**(1):78–96, February 1987.

[3] Franz Aurenhammer and Hannes Krasser. *Pseudo-Tetrahedral Complexes*. Proceedings of the Twenty-First European Workshop on Computational Geometry (Eindhoven, The Netherlands), pages 85–88, March 2005.

[4] Marshall Bern and David Eppstein. *Mesh Generation and Optimal Triangulation*. Computing in Euclidean Geometry (Ding-Zhu Du and Frank Hwang, editors), Lecture Notes Series on Computing, volume 1, pages 23–90. World Scientific, Singapore, 1992.

[5] Kevin Q. Brown. *Voronoi Diagrams from Convex Hulls*. Information Processing Letters **9**:223–228, 1979.

[6] Siu-Wing Cheng and Tamal Krishna Dey. *Quality Meshing with Weighted Delaunay Refinement*. Proceedings of the Thirteenth Annual Symposium on Discrete Algorithms (San Francisco, California), pages 137–146. Association for Computing Machinery, January 2002.

[7] Siu-Wing Cheng, Tamal Krishna Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. *Sliver Exudation*. Journal of the ACM **47**(5):883–904, September 2000.

[8] Siu-Wing Cheng and Sheung-Hung Poon. *Graded Conforming Delaunay Tetrahedralization with Bounded Radius-Edge Ratio*. Proceedings of the Fourteenth Annual Symposium on Discrete Algorithms (Baltimore, Maryland), pages 295–304. Society for Industrial and Applied Mathematics, January 2003.

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

[10] ———. *Guaranteed-Quality Triangular Meshes*. Technical Report TR-89-983, Department of Computer Science, Cornell University, 1989.

[11] ———. *Guaranteed-Quality Delaunay Meshing in 3D*. Proceedings of the Thirteenth Annual Symposium on Computational Geometry (Nice, France), pages 391–393. Association for Computing Machinery, June 1997.

[12] David Cohen-Steiner, Éric Colin de Verdière, and Mariette Yvinec. *Conforming Delaunay Triangulations in 3D*. Proceedings of the Eighteenth Annual Symposium on Computational Geometry (Barcelona, Spain), pages 199–208. Association for Computing Machinery, June 2002.

[13] Ed F. D'Azevedo and R. Bruce Simpson. *On Optimal Interpolation Triangle Incidences*. SIAM Journal on Scientific and Statistical Computing **10**:1063–1075, 1989.

[14] Boris Nikolaevich Delaunay. *Sur la Sphère Vide*. Izvestia Akademia Nauk SSSR, VII Seria, Otdelenie Matematicheskii i Estestvennyka Nauk **7**:793–800, 1934.

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

[16] Herbert Edelsbrunner and Damrong Guoy. *An Experimental Study of Sliver Exudation*. Tenth International Meshing Roundtable (Newport Beach, California), pages 307–316. Sandia National Laboratories, October 2001.

[17] Herbert Edelsbrunner and Ernst Peter Mücke. *Simulation of Simplicity: A Technique to Cope with Degenerate Cases in Geometric Algorithms*. ACM Transactions on Graphics **9**(1):66–104, 1990.

[18] Herbert Edelsbrunner and Raimund Seidel. *Voronoi Diagrams and Arrangements*. Discrete & Computational Geometry **1**:25–44, 1986.

[19] Herbert Edelsbrunner and Nimish R. Shah. *Incremental Topological Flipping Works for Regular Triangulations*. Algorithmica **15**(3):223–241, March 1996.

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

[21] 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.

[22] Paul-Louis George and Houman Borouchaki. *Delaunay Triangulation and Meshing: Application to Finite Elements*. Hermès, Paris, 1998.

[23] Abel J. P. Gomes. *A Concise B-Rep Data Structure for Stratified Subanalytic Objects*. Symposium on Geometry Processing (Aachen, Germany) (Leif Kobbelt, Peter Schröder, and Hugues Hoppe, editors), pages 83–93. Eurographics Association, June 2003.

[24] Nicolas Grislain and Jonathan Richard Shewchuk. *The Strange Complexity of Constrained Delaunay Triangulation*. Proceedings of the Fifteenth Canadian Conference on Computational Geometry (Halifax, Nova Scotia, Canada), pages 89–93, August 2003.

[25] Branko Grünbaum and Geoffrey C. Shephard. *A New Look at Euler's Theorem for Polyhedra*. The American Mathematical Monthly **101**(2):109–128, February 1994.

[26] Hugo Hadwiger. *Vorlesungen über Inhalt, Oberfläche und Isoperimetrie*. Springer-Verlag, Berlin, 1957.

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

[28] Charles L. Lawson. *Software for C^1 Surface Interpolation*. Mathematical Software III (John R. Rice, editor), pages 161–194. Academic Press, New York, 1977.

[29] Der-Tsai Lee and Arthur K. Lin. *Generalized Delaunay Triangulations for Planar Graphs*. Discrete & Computational Geometry **1**:201–217, 1986.

[30] Xiang-Yang Li and Shang-Hua Teng. *Generating Well-Shaped Delaunay Meshes in 3D*. Proceedings of the Twelfth Annual Symposium on Discrete Algorithms (Washington, D.C.), pages 28–37. Association for Computing Machinery, January 2001.

[31] Elefterios A. Melissaratos. *L_p Optimal d Dimensional Triangulations for Piecewise Linear Interpolation: A New Result on Data Dependent Triangulations*. Technical Report RUU-CS-93-13, Department of Computer Science, Utrecht University, Utrecht, The Netherlands, April 1993.

[32] 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.

[33] Michael Murphy, David M. Mount, and Carl W. Gable. *A Point-Placement Strategy for Conforming Delaunay Tetrahedralization*. Proceedings of the Eleventh Annual Symposium on Discrete Algorithms, pages 67–74. Association for Computing Machinery, January 2000.

[34] Steven E. Pav and Noel J. Walkington. *Robust Three Dimensional Delaunay Refinement*. Thirteenth International Meshing Roundtable (Williamsburg, Virginia), pages 145–156. Sandia National Laboratories, September 2004.

[35] P. L. Powar. *Minimal Roughness Property of the Delaunay Triangulation: A Shorter Approach*. Computer Aided Geometric Design **9**(6):491–494, December 1992.

[36] V. T. Rajan. *Optimality of the Delaunay Triangulation in $\mathbb{R}^d$*. Proceedings of the Seventh Annual Symposium on Computational Geometry (North Conway, New Hampshire), pages 357–363, June 1991.

[37] Shmuel Rippa. *Minimal Roughness Property of the Delaunay Triangulation*. Computer Aided Geometric Design **7**(6):489–497, November 1990.

[38] ———. *Long and Thin Triangles Can Be Good for Linear Interpolation*. SIAM Journal on Numerical Analysis **29**(1):257–270, February 1992.

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

[40] Jim Ruppert and Raimund Seidel. *On the Difficulty of Triangulating Three-Dimensional Nonconvex Polyhedra*. Discrete & Computational Geometry **7**(3):227–253, 1992.

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

[42] Raimund Seidel. *Voronoi Diagrams in Higher Dimensions*. Diplomarbeit, Institut für Informationsverarbeitung, Technische Universität Graz, 1982.

[43] ———. *Constrained Delaunay Triangulations and Voronoi Diagrams with Obstacles*. 1978–1988 Ten Years IIG (H. S. Poingratz and W. Schinnerl, editors), pages 178–191. Institute for Information Processing, Graz University of Technology, 1988.

[44] 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.

[45] ———. *A Condition Guaranteeing the Existence of Higher-Dimensional Constrained Delaunay Triangulations*. Proceedings of the Fourteenth Annual Symposium on Computational Geometry (Minneapolis, Minnesota), pages 76–85. Association for Computing Machinery, June 1998.

[46] ———. *Mesh Generation for Domains with Small Angles*. Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong), pages 1–10. Association for Computing Machinery, June 2000.

[47] ———. *Constrained Delaunay Tetrahedralizations and Provably Good Boundary Recovery*. Eleventh International Meshing Roundtable (Ithaca, New York), pages 193–204. Sandia National Laboratories, September 2002.

[48] ———. *Delaunay Refinement Algorithms for Triangular Mesh Generation*. Computational Geometry: Theory and Applications **22**(1–3):21–74, May 2002.

[49] Hang Si and Klaus Gärtner. *Meshing Piecewise Linear Complexes by Constrained Delaunay Tetrahedralizations*. Proceedings of the Fourteenth International Meshing Roundtable (San Diego, California) (Byron W. Hanks, editor), pages 147–163, September 2005.

[50] Robin Sibson. *A Brief Description of Natural Neighbor Interpolation*. Interpreting Multivariate Data (V. Barnett, editor), pages 22–36. John Wiley, 1981.

[51] Shayne Waldron. *The Error in Linear Interpolation at the Vertices of a Simplex*. SIAM Journal on Numerical Analysis **35**(3):1191–1200, 1998.

[52] 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.

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

[54] Günter M. Ziegler. *Lectures on Polytopes*, revised first edition. Springer-Verlag, New York, 1995.