

Duals of Orphan-Free Anisotropic Voronoi Diagrams are Embedded Meshes

Guillermo D. Canas
Massachusetts Institute of Technology, and
Italian Institute of Technology
Cambridge, MA
guilledc@mit.edu

Steven J. Gortler
Harvard University
Cambridge, MA
sjg@seas.harvard.edu

ABSTRACT

Given an anisotropic Voronoi diagram, we address the fundamental question of when its dual is embedded. We show that, by requiring only that the primal be orphan-free (have connected Voronoi regions), its dual is always guaranteed to be an embedded triangulation. Further, the primal diagram and its dual have properties that parallel those of ordinary Voronoi diagrams: the primal's vertices, edges, and faces are connected, and the dual triangulation has a simple, closed boundary. Additionally, if the underlying metric has bounded anisotropy (ratio of eigenvalues), the dual is guaranteed to triangulate the convex hull of the sites. These results apply to the duals of anisotropic Voronoi diagrams of *any* set of sites, so long as their Voronoi diagram is orphan-free. By combining this general result with existing conditions for obtaining orphan-free anisotropic Voronoi diagrams, a simple and natural condition for a set of sites to form an embedded anisotropic Delaunay triangulation follows.

Categories and Subject Descriptors

I.3.5 [Computational Geometry and Object Modeling]: Geometric algorithms, languages, and systems.

General Terms

Algorithms, Theory.

Keywords

Anisotropic Voronoi diagrams, Delaunay triangulations, Orphan-free.

1. INTRODUCTION

Voronoi diagrams and their dual Delaunay triangulations are fundamental constructions with numerous associated guarantees, and extensive application in practice [6, 8]. At their heart is the use of a distance between points, which in the

original version is taken to be Euclidean. This suggests that, by considering generalizations of the Euclidean distance, it may be possible to obtain variants which can be well-suited to a wider range of applications.

Attempts in this direction have been met with some success. Power diagrams [9] generalize Euclidean distance by associating a bias-term to each site. The duals of these diagrams are guaranteed to be embedded triangulations, in any number of dimensions. Although this is a strict generalization of Euclidean distance, it is a somewhat limited one. The effect of the bias term is to locally enlarge or shrink the region associated to each site, loosely-speaking “equally in every direction”. It allows some freedom in choosing local scale, with no preference for specific directions.

Another way to generalize Voronoi diagrams is to endow Euclidean space with a continuously varying Riemannian metric, and use its associated (geodesic) distance in the definition of the Voronoi diagram. This generalization carries a significant amount of freedom, allowing for local scale and “directionality” to be freely specified at each point.

Despite being of potentially great interest, this latter approach has faced several obstacles. The dual of such an anisotropic Voronoi diagram is, in general, *not* an embedded triangulation, even if the primal is orphan-free (the Voronoi regions are connected), and may produce element inversions and edge crossings. An additional obstacle is that the geodesic distance is extremely expensive to compute in practice, rendering traditional methods for constructing Voronoi diagrams impractical.

Two notable attempts to resolve these problems have been proposed. Independently, Labelle and Shewchuk [13], and Du and Wang [7] propose two efficient approximations of geodesic distance from a site to any point in the domain. Although their associated Voronoi diagrams are, in general, no longer orphan-free, Labelle and Shewchuk show that a set of sites exists with an orphan-free diagram, whose dual is embedded, in two dimensions. They accomplish this by proposing an iterative site-insertion algorithm that, for any given metric, constructs one such set of sites. Some recent work making use of the above definitions are [2, 1].

In this paper, we show that if a set of sites produces an orphan-free anisotropic Voronoi diagram, using the definition of [7], then its dual is always an embedded triangulation (or a polygonal mesh with convex faces in general), in two dimensions (Thm. 6.2). This effectively states that, regardless of the sites' positions, if the primal is well-behaved, then the dual is as well. Further, in a way that parallels the ordinary Delaunay case, the dual has no degenerate el-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'12, June 17–20, 2012, Chapel Hill, North Carolina, USA.
Copyright 2012 ACM 978-1-4503-1299-8/12/06 ...\$10.00.

ements (Lem. 4.9), its elements (vertices, edges, faces) are unique (Cor. 6.3), and, under mild assumptions on the metric, is guaranteed to triangulate the convex hull of the sites (Thm. 4.7). Note that, while [13] prove a property of the output of an algorithm, the results in this paper are fundamental properties of the diagrams themselves (independent of how the orphan-freedom of the primal was obtained).

In practice, we may combine our results with those of [4] to conclude that duals of anisotropic Voronoi diagrams of appropriate ϵ -nets are always embedded triangulations (Cor. 6.4). This may be particularly useful in (asymptotically-optimal) function approximation applications, where we are often interested in constructing the anisotropic Delaunay triangulations of appropriate ϵ -nets [5, 12], and was the initial motivation for the current work. Finally, we note that, as discussed in Sec. 7, in practice, algorithms for constructing anisotropic diagrams of the type discussed here may have a more numerical flavor (i.e. front-propagation, fast marching methods), in contrast with the more combinatorial nature of other Voronoi diagrams.

2. SETUP

Given a finite set $V \subset \mathbb{R}^2$ of sites on the plane, a Voronoi diagram decomposes Euclidean space into regions, each distinguished by the site its points are closest to. The notion of closeness is defined as follows: consider a continuous metric (in coordinates: $Q : \mathbb{R}^2 \rightarrow \mathbb{R}^{2 \times 2}$, symmetric, positive definite) over two-dimensional Euclidean space and, following [7], define the (asymmetric) “distance” between a site ($v \in V$) and point ($p \in \mathbb{R}^2$) as

$$D(v, p) = [(p - v)^t Q_p (p - v)]^{1/2}$$

The Voronoi region of site v is the set of points no further from v than from any other site:

$$R(v) = \{p \in \mathbb{R}^2 : D(v, p) \leq D(w, p), \forall w \in V\}$$

Note that, because V is a set, and not a multiset, there are no coincident sites.

This definition of anisotropic Voronoi diagram results in diagrams that take the usual shape, but whose regions may have curved boundaries, and tend to be elongated along certain directions, depending on the metric Q .

Following the notation of [13], we say that a diagram is *orphan-free* if its regions $R(v)$ are connected, for all $v \in V$ (equivalently: each *cell*, or connected component of a Voronoi region, contains its generating site).

An important distinction in this construction is that the “interfaces” in-between Voronoi regions do not, in general, have null measure (and points equidistant to three or more sites are not always isolated), unlike geodesic-distance diagrams [14] (see [5] Lemma 5.2), and those of [13].

Note that, given one such degenerate diagram, any interior point p of a Voronoi edge $R(v) \cap R(w)$ of non-zero measure has an open neighborhood $N_p \subset R(v) \cap R(w)$. By adding to Q an arbitrarily small perturbation supported on N_p , an orphan is created. That is, with respect to the metric, any degenerate diagram is arbitrarily close to a diagram that has orphans, and therefore the non-degeneracy requirement is only slightly more restrictive than the orphan-freedom one.

We begin with three simple lemmas, proved in Appendix A, which provide an introduction to some of the techniques used in the sequel. The first is a simple extension of Lemma

2.1 of [7], and the second follows directly from the continuity of Q and the fact that every site is strictly closer to itself than to all other sites.

LEMMA 2.1. *Given two sites $v, w \in V$, $v \neq w$, the only point equidistant to v, w in their supporting line is the midpoint $(v + w)/2$.*

LEMMA 2.2. *Every site is an interior point of its corresponding Voronoi region.*

LEMMA 2.3. *Every Voronoi region of an orphan-free anisotropic Voronoi diagram in \mathbb{R}^2 is simply connected.*

In the sequel, unless otherwise stated, proofs are omitted in the interest of conciseness, but they can be found in a longer version of this paper in [3].

Primal diagram. The Voronoi diagram is a collection of Voronoi regions $R(v)$, one for every site in V , as well as a structure induced by the sets of points closest and equidistant to two or more sites. Let \tilde{P} be the primal Voronoi diagram embedded on the sphere \mathbb{S}^2 by stereographically projecting the plane onto the punctured sphere and completing it with a “point at infinity” p_∞ . \tilde{P} includes not just regions, but also embedded edges \tilde{E}_{vw} , which are *connected* sets of points closest and equidistant to two sites v, w , as well as vertices (points closest and equidistant to three or more sites). Since we are assuming that sets of equidistant points to two sites have null measure, and that points equidistant to three or more sites are isolated, we can consider the abstract mesh $P = (V_p, E_p, F_p)$ with structure derived from \tilde{P} , where V_p, E_p, F_p are the vertices, edges, and faces of P .

A simple induction argument reveals that every edge in E_p connects two vertices in V_p . To see this, start with an orphan-free diagram with two sites. The only edge of \tilde{E} is the unbounded set of points equidistant to the sites, passing through their midpoint (Lem. 2.1), and with endpoints at p_∞ . The insertion of an additional site can only split existing edges of \tilde{E} at points where two edges cross (a vertex by definition), and thus the original edge is split into pieces whose endpoints are now either p_∞ or an edge crossing (a vertex).

Clearly, \tilde{P} is a planar embedded diagram since two edges cannot meet except at vertices: any point p where edge $D(u, p) = D(v, p)$ crosses edge $D(v, p) = D(w, p)$ is equidistant to three sites, u, v, w , and thus a vertex. Therefore P is planar.

Dual of an orphan-free Voronoi diagram. The dual of the primal mesh structure P is another abstract mesh $G = (V, E, F)$, with one vertex per site, and edges connecting adjacent regions of \tilde{P} . Its face structure is derived from that of P by duality and, in particular, (V, E) is planar, since (V_p, E_p) is planar. Using the embedded primal \tilde{P} , we can define an embedding \tilde{G} of G on the plane, with vertices coinciding with sites (and curved edges). Although \tilde{P} was defined on the sphere, we embed \tilde{G} on the plane instead, and therefore we exclude (from both G and \tilde{G}) the unbounded face dual to the point at infinity p_∞ .

Since (V, E) is planar, G derives the following properties from the primal:

- (i) Every edge in E is incident to two dual faces, except for *boundary edges* (whose corresponding primal edge in \tilde{P}

connects unbounded regions), which are only incident to one dual face.

- (ii) To every dual face in F , connecting vertices v_1, \dots, v_m , corresponds a primal vertex $c \in R(v_1) \cap \dots \cap R(v_m)$ of \tilde{P} , which is equidistant to the corresponding sites v_1, \dots, v_m .

The property (ii) requires that the Voronoi diagram be orphan-free, as noted in the remark after Thm. 2.7.

Notice that either \tilde{G} or \tilde{P} could be multigraphs (e.g. \tilde{P} would be a multigraph if an edge \tilde{E}_p of \tilde{P} has multiple connected components). However, the following lemma ensures that both are simple graphs.

LEMMA 2.4. *\tilde{G} and \tilde{P} are simple (have no multi-edges or self-loops).*

By its definition, an edge of \tilde{P} is a connected set of points equidistant and closest to two sites v, w . Since \tilde{P} has no multi-edges, this implies the following:

COROLLARY 2.5. *The set of points closest and equidistant to two sites of an orphan-free diagram is connected.*

Property (ii) above can be restated in a more useful way that parallels an equivalent property of duals of ordinary Voronoi diagrams.

DEFINITION 2.6 (**Empty circum-ellipse property**). *A face of G incident to vertices v_1, \dots, v_m satisfies the empty circum-ellipse (ECE) property if there is some ellipse that circumscribes the sites corresponding to v_1, \dots, v_m , and contains no site in its interior.*

Note that *any* empty circum-ellipse serves as witness to the ECE property. We show that the ellipse centered at a Voronoi vertex c , with axis given by Q_c , is always an empty circum-ellipse of its dual polygon.

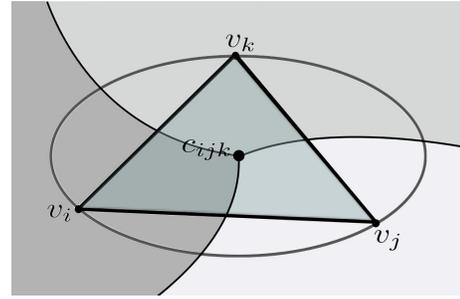
THEOREM 2.7. *Every face of G satisfies the empty circum-ellipse property.*

PROOF. The proof is constructive. To every face in G connecting vertices with corresponding sites v_1, \dots, v_m , by property (ii), corresponds a Voronoi vertex $c \in R(v_1) \cap \dots \cap R(v_m)$ of \tilde{P} that is equidistant to v_1, \dots, v_m . If $\mu = D(v_1, c) = \dots = D(v_m, c)$, then the open ellipse

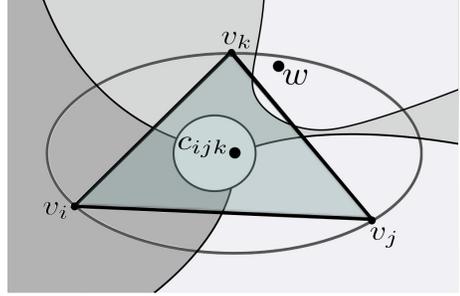
$$\theta = \{p \in \mathbb{R}^2 : D(p, c) < \mu\} = \{p \in \mathbb{R}^2 : (p-c)^t Q_c (p-c) < \mu^2\}$$

circumscribes the v_1, \dots, v_m and does not contain any site (otherwise, c would be closer to that site than to v_1, \dots, v_m , and therefore not in $R(v_1) \cap \dots \cap R(v_m)$, a contradiction). \square

Although this may not at first seem apparent, the above property uses the fact that the primal diagram is orphan free. The critical fact in the proof of Thm. 2.7 is not that there is a Voronoi vertex c equidistant to v_1, \dots, v_m , but rather that it is $c \in R(v_1) \cap \dots \cap R(v_m)$. Consider Figs. 1(a) and 1(b), in which sites v_i, v_j, v_k form a face of G . In 1(a), the diagram is orphan-free, and the fact that $c_{ijk} \in R(v_i) \cap R(v_j) \cap R(v_k)$ implies Thm. 2.7. Figure 1(b) shows part of a Voronoi diagram that has an orphan cell (corresponding to site w) “covering” c_{ijk} . If the dual is constructed such that the face v_i, v_j, v_k is not in the dual then a “hole” in the triangulation may occur. If the face v_i, v_j, v_k is in the dual, then,



(a)



(b)

Figure 1: In an orphan-free diagram (a), every face satisfies the empty-circum ellipse (ECE) property. If the diagram is not orphan-free, then some faces may not satisfy the ECE property.

despite the fact that c_{ijk} is equidistant to v_i, v_j, v_k , this c_{ijk} is no longer in $R(v_i) \cap R(v_j) \cap R(v_k)$, but in the Voronoi region of w . Since c_{ijk} is closer to w than to v_i, v_j, v_k , any open ellipse centered at c_{ijk} , and having v_i, v_j, v_k on its boundary, contains w . Therefore, in Fig. 1(b), the ECE property does not hold for the dual face of c_{ijk} . Ensuring that Thm. 2.7 holds is one of the main reasons for requiring orphan-freedom.

3. SUMMARY OF RESULTS AND OUTLINE

The two main results in this paper assume that we are given an orphan-free anisotropic Voronoi diagram (in two dimensions) and prove that: 1) *the dual is an embedded triangulation* (Thm. 6.2), and 2) *if the metric has bounded anisotropy (ratio of eigenvalues), the dual triangulates the convex hull of the sites* (Thm. 6.1).

We prove a number of additional results. Possibly the most important shows that the elements of an orphan-free diagram (vertices, edges, faces), are connected sets, and therefore, in some sense, unique (Cor. 6.3). This is important, since this is a natural well-behavedness condition on the primal, which, together with the above results, indicates that orphan-freedom is sufficient to guarantee well-behavedness of both primal and dual.

A smaller result of interest is: the dual faces satisfy an *empty circum-ellipse* property (Thm. 2.7), which parallels the empty-circumcircle property of ordinary Voronoi diagrams, and could have further practical implications from the ones described here.

The order of the proofs is, however, different from the one just stated. In effect, we begin by proving the most

restrictive case: that the dual of an orphan-free diagram with metric of bounded anisotropy is embedded (Sec. 5), and that its boundary is the boundary of the convex hull of the sites (Sec. 4). We then relax the bounded anisotropy condition (Sec. 6), and show that, in the general case, we lose the convex hull property, but the dual remains embedded. Finally, once these results are established, we prove that vertices of the primal are unique (uniqueness of edges was proven in Sec. 2, and uniqueness of primal faces is the same as orphan-freeness).

We assume in the remainder of the paper that there are more than two sites, not all of which are colinear, and relegate the (considerably simpler) colinear case to Appendix B.

4. DUAL OF ORPHAN-FREE DIAGRAM (PART I: BOUNDARY)

In this section, we assume that the metric has bounded anisotropy, and conclude that the boundary of the dual of an orphan-free diagram is the same as the boundary of the convex hull of the sites (and in particular is simple and closed). If $R_p \Lambda_p R_p^t$ is the eigendecomposition of Q_p at $p \in \mathbb{R}^2$, with $\Lambda_p = \text{diag}[\lambda_2(p), \lambda_1(p)]$, $\lambda_2(p) \geq \lambda_1(p) > 0$, and R_p orthonormal, then we assume that there is some bound γ on the anisotropy of Q , such that: $1 \leq \lambda_2(p)/\lambda_1(p) < \gamma^2$ for $p \in \mathbb{R}^2$. Note that this condition may commonly hold in practice, if the metric is sampled on a compact domain (and possibly extended to the plane by reusing sampled values only).

We begin by defining \bar{G} to be the straight-edge drawing of G with vertices at the sites. For the moment, we assume that every Voronoi vertex is equidistant to no more than three sites, and therefore that all faces in \bar{G} are triangles (we extend the results to the general case in Appendix G). We associate to \bar{G} a mapping (from the vertices, edges, and faces of \bar{G} to \mathbb{R}^2) for which, because all faces of \bar{G} are triangles, it is well defined whether a point in \mathbb{R}^2 belongs to any given face of \bar{G} , a fact that will be used in the proofs of Sec. 5. This mapping will be shown in Sec. 5 to be an embedding. In the sequel, it is assumed that \bar{G} encompasses both the mesh structure, and the mapping.

The boundary vertices of G are those whose corresponding primal regions in \tilde{P} are unbounded, while boundary edges of G connect boundary vertices. Note that this is a *topological* property of G , rather than a geometric one (boundary elements of G may, in principle, not lie in the boundary of the convex hull of the sites). For convenience, we call $B \subseteq E$ the set of boundary edges of G .

The convex hull $\mathcal{CH}(V)$ of V is the minimal (w.r.t. set containment) convex set that contains V . For convenience, we name $W = \{w_i \in V : i = 1, \dots, m\}$ the sites that are part of the boundary of the convex hull $\mathcal{CH}(V)$, and order them in clock-wise order around $\mathcal{CH}(V)$. The boundary \mathcal{B} of the convex hull is a simple circular chain $\mathcal{B} = \{(w_i, w_{i \oplus 1}) : i = 1, \dots, m\}$. We prove that it is $B = \mathcal{B}$ (loosely speaking: the topological boundary of G , and the geometric boundary of its straight-edge embedding \bar{G} , are the same), which implies that \bar{G} covers the convex hull of the sites, and its boundary edges form a simple, closed polygonal chain.

LEMMA 4.1. *To every boundary edge (v, w) of G corresponds a segment in the boundary of $\mathcal{CH}(V)$. [$B \subseteq \mathcal{B}$]*

We now turn to the converse claim: that to every segment $(w_i, w_j) \in \mathcal{B}$ corresponds a boundary edge $(w_i, w_j) \in B$ in G . Since B is the set of boundary edges of G , whose primal edges in \tilde{P} are unbounded, the claim is equivalent to proving that, to every segment (w_i, w_j) in the boundary of the convex hull corresponds an edge of the dual G , whose primal edge \tilde{E}_{w_i, w_j} is unbounded.

The proof proceeds as follows. First, assume w.l.o.g. that the origin is in the interior of $\mathcal{CH}(V)$. Let $C(\sigma) = \{x \in \mathbb{R}^2 : \|x\| = \sigma\}$ be a sufficiently large origin-centered circle. We define two functions:

$$\begin{aligned} \pi : C(\sigma) &\rightarrow \partial\mathcal{CH}(V), & \pi(p) &= \underset{r \in \mathcal{CH}(V)}{\text{argmin}} D(r, p) \\ \nu : \partial\mathcal{CH}(V) &\rightarrow C(\sigma), & \nu(r) &= \sigma \cdot r / \|r\| \end{aligned}$$

π simply projects every point in C to its closest in $\mathcal{CH}(V)$ (with respect to the distance D), and ν projects a point back to C .

LEMMA 4.2. *$\pi(p) = \underset{x \in \mathcal{CH}(V)}{\text{argmin}} D(x, p)$ is a continuous function in \mathbb{R}^2 .*

PROOF. We first prove by contradiction that $\pi(p)$ is unique. Let M_p be the unique [4], symmetric positive definite square-root matrix such that $M_p^t M_p^t = Q_p$. Consider distinct $q_1, q_2 \in \mathcal{CH}(V)$ closest to p . Since $q_1 \neq q_2$, by the convexity of the Euclidean norm, and the positive-definiteness of M_p , it is

$$\begin{aligned} D((q_1 + q_2)/2, p) &= \|M_p((q_1 + q_2)/2 - p)\| \\ &< \|M_p(q_1 - p)\| = \|M_p(q_2 - p)\| \end{aligned}$$

Thus p is closer to $(q_1 + q_2)/2$ than to q_1, q_2 , a contradiction. Therefore the closest point $\pi(p)$ to p is unique.

If π were not continuous at p , then there is $\epsilon > 0$ such that for all $\delta > 0$ there is a point $\bar{p} = \bar{p}(\delta)$ such that $\|p - \bar{p}\| < \delta$ and $\|\pi(p) - \pi(\bar{p})\| \geq \epsilon$. Consider the sequence $\{\kappa_i = \pi(\bar{p}(1/i)) : i \in \mathbb{N}\}$ of points in $\mathcal{CH}(V)$. Because $\mathcal{CH}(V)$ is compact, $\{\kappa_i\}$ has a subsequence that converges to some $\kappa \in \mathcal{CH}(V)$. By continuity of D (which follows from the continuity of Q), it is

$$\begin{aligned} D(\pi(p), p) &= \min_{x \in \mathcal{CH}(V)} D(x, p) = \lim_{i \rightarrow \infty} \min_{x \in \mathcal{CH}(V)} D(x, \bar{p}_i) \\ &= \lim_{i \rightarrow \infty} D(\kappa_i, \bar{p}_i) = D(\kappa, p) \end{aligned}$$

and therefore the closest point in $\mathcal{CH}(V)$ to p is not unique, a contradiction. \square

Clearly, ν is continuous in $\partial\mathcal{CH}(V)$. Note that, because $\mathcal{CH}(V)$ contains the origin, then, as shown in Fig. 2, ν projects every point $\pi(p) \in (w_i, w_j)$ on a segment of $\partial\mathcal{CH}(V)$, *outwards* from the convex hull (and on the empty side of (w_i, w_j)); that is, so that $\nu(\pi(p)) \in H_{ij}^+ \cap C$ ($\nu(\pi(p))$ is in the empty half-space of (w_i, w_j)).

Conveniently, Lem. 4.3 shows that, if $\pi(p) = (w_i + w_j)/2$, then $p \in \tilde{E}_{w_i, w_j}$, and so the claim now reduces to showing that for each segment (w_i, w_j) of $\partial\mathcal{CH}(V)$, and for every sufficiently large circle C , there is $p \in C$ with $\pi(p) = (w_i + w_j)/2$. Since this implies that \tilde{E}_{w_i, w_j} is unbounded, it means that the corresponding edge (w_i, w_j) is in B .

The proof is by contradiction. Lem. 4.5 uses Brouwer's fixed point theorem to show that, for every segment (w_i, w_j) of \mathcal{B} , if there were no $p \in C$ with $\pi(p) = (w_i + w_j)/2$, then the function $\nu \circ \pi : C \rightarrow C$ must have a point $q \in C$ such that

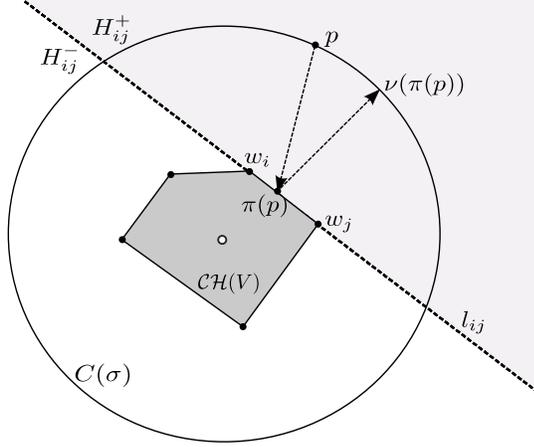


Figure 2: The construction for the proof of Thm. 4.7.

$\nu(\pi(q)) = -q$, or equivalently such that q is “behind” the segment $(w_i, w_j) \in \partial\mathcal{CH}(V)$ to which it is closest ($q \in H_{ij}^-$). On the other hand, Lem. 4.4 shows that, for all sufficiently large circles C , no point $q \in C$ can be closest to some segment $(w_i, w_j) \in \partial\mathcal{CH}(V)$ that it is *behind* of, creating a contradiction.

As mentioned above, we are interested in identifying points in $C(\sigma)$ that are also in some primal edge \tilde{E}_{w_i, w_j} . The following lemma will be used to show that, for sufficiently large σ , such a point p can be characterized by satisfying $\pi(p) = (w_i + w_j)/2$.

LEMMA 4.3. *If $\pi(p) = (w_i + w_j)/2$, with $p \in \mathbb{R}^2$, $w_i, w_j \in W$, then $D(w_i, p) = D(w_j, p)$.*

The following technical lemma is the key in constructing a contradiction by showing that, for sufficiently large circles $C(\sigma)$, no point $q \in C(\sigma)$ can be closest to some segment $(w_i, w_j) \in \partial\mathcal{CH}(V)$ that it is behind of ($q \notin H_{ij}^-$); where, as before, the open half space H_{ij}^- is chosen to be the only of the two half spaces on either side of the supporting line l_{ij} of (w_i, w_j) such that $H_{ij}^- \cap V \neq \emptyset$.

LEMMA 4.4. *There is ρ such that, for any segment $(w_i, w_j) \in \mathcal{B}$, with supporting line l_{ij} , every $p \in H_{ij}^-$ with $\|p\| > \rho$ whose closest point in l_{ij} is $m_p \in \overline{w_i w_j}$ is closer to a site in $V \setminus \{w_i, w_j\}$ than to l_{ij} .*

LEMMA 4.5. *Every continuous function $F : \mathbb{S}^n \rightarrow \mathbb{S}^n$ that is not onto has a fixed point.*

LEMMA 4.6. *To every segment in the boundary of $\mathcal{CH}(V)$ corresponds a boundary edge of G . [$\mathcal{B} \supseteq \mathcal{B}$]*

Therefore,

THEOREM 4.7. *If the metric Q has bounded ratio of eigenvalues, then the boundary edges of \bar{G} are the same as the boundary edges of the convex hull of V . [$\mathcal{B} = \mathcal{B}$]*

COROLLARY 4.8. *If the metric Q has bounded ratio of eigenvalues, then the boundary edges of \bar{G} form a simple, closed, convex polygonal chain.*

The ECE property is sufficient to show that no face in \bar{G} is degenerate, a fact that will be used in the following section.

LEMMA 4.9. *\bar{G} has no degenerate (null area) elements.*

In Thm. 5.9, we show that, even if Voronoi vertices are incident to more than three Voronoi regions (equidistant to more than three sites), every face in \bar{G} is a (strictly) convex polygon, and therefore can be trivially triangulated (say in a fan arrangement). The resulting triangulation has no degenerate elements, and all its triangles satisfy the ECE condition (using the same witness ellipsoid as the convex polygon from which they are triangulated).

5. DUAL OF ORPHAN-FREE DIAGRAM (PART II: INTERIOR)

In this section, we assume that the (topological) boundary of \bar{G} is simple and closed, and prove that \bar{G} must be embedded. The main argument in the proof uses Theorems 2.7 and 4.7, Cor. 4.8, as well as the theory of discrete one-forms on graphs, to show that there are no “edge foldovers” in \bar{G} (edges whose two incident faces are on the same side of its supporting line), and use this to conclude that \bar{G} is embedded (Thm. 5.9). As in Sec. 4, we assume that not all sites are colinear (the simpler colinear case was addressed in Appendix B). We distinguish between the sites $W \subseteq V$ that lie on the boundary of the convex hull, and the remaining, or *interior sites* ($V \setminus W$).

The following definition, from [11], assumes that, for each edge (v_i, v_j) in G , we distinguish the two opposing half-edges (v_i, v_j) and (v_j, v_i) .

DEFINITION 5.1 (GORTLER ET AL. [11]). *A non-vanishing (discrete) one-form ξ is an assignment of a real value $\xi_{ij} \neq 0$ to each half edge (v_i, v_j) in G , such that $\xi_{ji} = -\xi_{ij}$.*

Since \bar{G} has the same structure as G , we can construct a non-vanishing one-form over \bar{G} as follows. Given some unit direction vector $n \in \mathbb{S}^1$ (in coordinates $n = [n_1, n_2]^t$), we assign a real value $z(v) = n^t v$ to each vertex v in \bar{G} , and define $\xi_{ij} = z(v_i) - z(v_j)$, which clearly satisfies $\xi_{ji} = -\xi_{ij}$. The one-form, denoted by ξ^n , is non-vanishing if, for all edges $(v_i, v_j) \in E$, it is $\xi_{ij} = n^t(v_i - v_j) \neq 0$. That is, if the direction n is not orthogonal to any edge. The set of edges has cardinality $|E| \leq |V|(|V| - 1)/2$, and in particular it is finite. Therefore *almost all* directions $n \in \mathbb{S}^1$ generate a non-vanishing one-form ξ^n .

Since $G = (V, E, F)$ is a planar graph with a well-defined face structure, there is, for each face $f \in F$, a cyclically ordered set ∂f of half-edges round the face. Likewise, for each vertex $v \in V$, the set δv of cyclically ordered (oriented) half-edges emanating from each vertex is well-defined.

DEFINITION 5.2 (GORTLER ET AL. [11]). *Given non-vanishing one-form ξ^n corresponding to $n \in \mathbb{S}^1$, the index of vertex v with respect to ξ^n is $\text{ind}_{\xi^n}(v) = 1 - \text{sc}_{\xi^n}(v)/2$, where $\text{sc}_{\xi^n}(v)$ is the number of sign changes of ξ^n as one visits the half-edges of δv in order.*

The index of face f is $\text{ind}_{\xi^n}(f) = 1 - \text{sc}_{\xi^n}(f)/2$ where $\text{sc}_{\xi^n}(f)$ is the number of sign changes of ξ^n as one visits the half-edges of ∂f in order.

Note that, by definition, it is always $\mathbf{ind}_{\xi^n}(v) \leq 1$. A discrete analog of the Poincaré-Hopf index theorem relates the two indices above:

THEOREM 5.3 (GORTLER ET AL. [11]). *For any non-vanishing one-form ξ^n , it is*

$$\sum_{v \in V} \mathbf{ind}_{\xi^n}(v) + \sum_{f \in F} \mathbf{ind}_{\xi^n}(f) = 2$$

Note that this follows from Theorem 3.5 of [11] because the unbounded, outside face, which is not in G , is assumed in this section to be closed and simple, and therefore would have null index. Note that the machinery from [11] to deal with degenerate cases isn't needed here because vertices, by definition, cannot coincide (V is not a multiset).

The one-forms defined above satisfy the following property.

LEMMA 5.4. *Given a non-vanishing ξ^n , the sum of indices of interior vertices ($V \setminus W$) of \bar{G} is non-negative.*

The next two lemmas relate the presence of edge foldovers and the ECE property of Definition 2.6 to the indices of vertices in \bar{G} .

LEMMA 5.5. *If \bar{G} has an edge foldover, then there is a non-vanishing one-form ξ^n such that $\mathbf{ind}_{\xi^n}(v) < 0$ for some interior vertex $v \in V \setminus W$.*

LEMMA 5.6. *Given $n \in \mathbb{S}^1$ and non-vanishing one-form ξ^n , if \bar{G} has an interior vertex $v \in V \setminus W$ with index $\mathbf{ind}_{\xi^n}(v) = 1$, then there is a face f of G that does not satisfy the empty circum-ellipse property.*

The above provides the necessary tools to prove the following key lemma.

LEMMA 5.7. *\bar{G} has no edge foldovers.*

Finally, the absence of edge foldovers, together with a simple and closed boundary, is sufficient to show that \bar{G} is embedded.

LEMMA 5.8. *If its (topological) boundary is simple and closed, then the straight-line dual \bar{G} of an orphan-free diagram, with vertices incident to at most three sites, is an embedded triangulation.*

As shown in Appendix C, even if there are non-generic Voronoi vertices that are incident to more than three sites, the dual is composed of faces each of which is convex and satisfies the ECE condition (Def. 2.6). Each convex face can be triangulated (e.g. in a fan arrangement) in such a way that individual triangles satisfy the ECE condition with the same witness ellipse as the face from which they are derived. This leads to the following:

THEOREM 5.9. *If its (topological) boundary is simple and closed, then the straight-line dual \bar{G} of an orphan-free diagram is an embedded polygonal mesh with convex faces.*

6. FINAL RESULTS

We can combine Corollary 4.8 and Theorem 5.9 into

THEOREM 6.1. *If the metric Q has bounded ratio of eigenvalues, then the dual of an orphan-free Voronoi diagram with respect to Q is an embedded polygonal mesh with convex faces, and covers the convex hull of the sites.*

This result can be generalized, dropping the bounded anisotropy condition on Q , but at the cost of losing the convex hull property, as shown in Appendix D:

THEOREM 6.2. *The dual of an orphan-free Voronoi diagram is an embedded polygonal mesh with convex faces.*

From the above results, Corollary 2.5, and the definition of the dual, it follows (see Appendix D) that

COROLLARY 6.3. *An orphan-free anisotropic Voronoi diagram is composed of unique (connected) vertices, edges, and faces.*

Finally, we note that Thm. 5.9 can be combined with existing conditions for orphan-freedom [4], resulting in a simple and natural condition for a set of sites to induce an embedded polygonal mesh as the dual of their anisotropic Voronoi diagram:

COROLLARY 6.4. *If V is an asymmetric ϵ -net w.r.t. D , Q a continuous metric with metric variation σ , and $\epsilon\sigma \leq 0.09868$, then the dual of the anisotropic Voronoi diagram of V is an embedded polygonal mesh with convex faces.*

where an asymmetric ϵ -net is simply a weaker form of ϵ -net defined on non-symmetric functions D , which can be computed with the iterative algorithm of [10], and the metric variation σ is a Lipschitz-type condition on Q [4]. The above condition is known to be conservative, and there may be simpler conditions to achieve orphan-freedom. As a practical observation, Du and Wang [7] report orphans to be a rare occurrence in their experiments.

7. PROOF-OF-CONCEPT IMPLEMENTATION

Though not aiming for an efficient implementation, we implemented a simple proof-of-concept that constructs anisotropic Voronoi diagrams like the ones considered in this paper, and their duals (Fig. 3). A closed-form metric, which has bounded ratio of eigenvalues, is discretized on a fine regular grid, and linearly interpolated inside grid elements, resulting in a continuous metric. The sites are generated randomly (Figs. 3(a) and 3(b)), or using a combination of random, and equispaced points forming an asymmetric ϵ -net (remaining figures).

The primal diagram was obtained using front propagation from the sites outwards, until fronts meet at Voronoi edges. The runtime is proportional to the grid size, since every grid-vertex is visited exactly six times (equal to their valence). The implementation does not guarantee the correctness of the diagram unless it is orphan-free, and serves to verify the claims of the paper since well-behavedness of the dual is predicated on that of the primal.

The two main claims of the paper are clearly illustrated in these examples. In all examples, the dual covers the convex hull of the vertices (Thm. 4.7), is a single cover, embedded with straight edges without edge crossings (Thm. 5.9), and has no degenerate faces (Lem. 4.9).

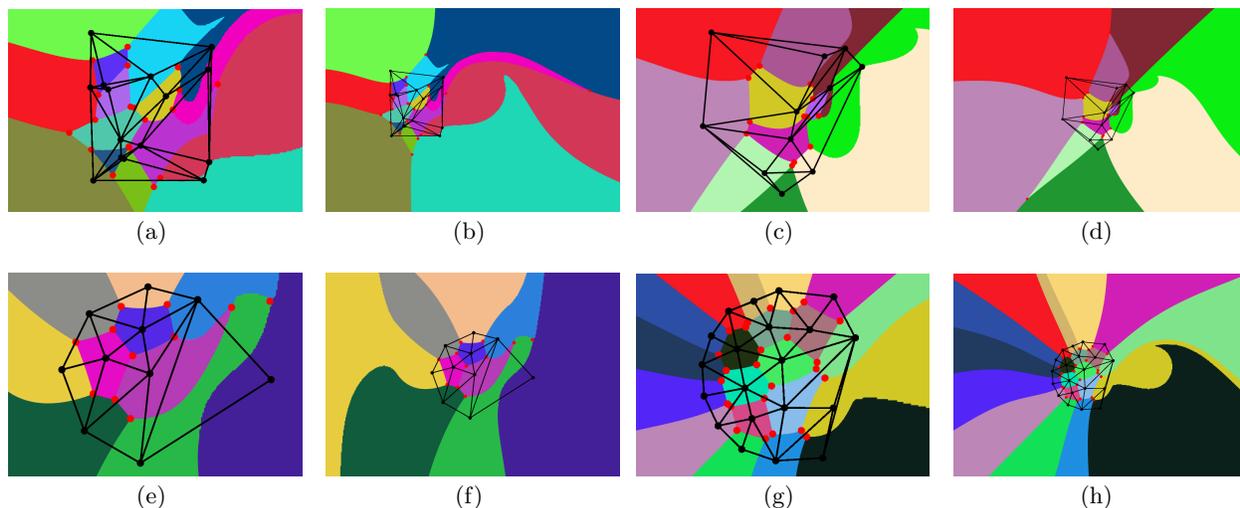


Figure 3: Anisotropic Voronoi diagrams, and their duals generated by our proof-of-concept implementation. Voronoi vertices are marked in red, while dual vertices (sites) and edges are drawn in black.

By focusing on the primal diagrams (second and fourth column), further claims in the paper become apparent, namely that Voronoi regions are simply connected (Lem. 2.3), Voronoi vertices are unique (Cor. 6.3), Voronoi edges are connected (Cor. 2.5), unbounded regions correspond to boundary dual vertices, and unbounded edges of the Voronoi diagram correspond to boundary dual edges.

8. CONCLUSION AND OPEN QUESTIONS

We studied the properties of duals of orphan-free anisotropic Voronoi diagrams, for the purposes of constructing triangulations on the plane. The main result (Theorems 5.9 and 4.7, Cor. 6.4) is that the dual, with straight edges and having the sites as vertices, is embedded and covers the convex hull of the sites, mirroring similar results for ordinary Voronoi diagrams and their duals.

A few, somewhat less important properties are proven, including the fact that every primal region is simply connected, that elements of the primal are unique (Cor. 6.3, but perhaps most significantly that every face in the dual satisfies an *empty circum-ellipse* property that has a direct parallel in the *empty circum-circle* property of ordinary diagrams, and is the basis for proving that it is embedded with straight edges.

Perhaps the most important outstanding question may be whether these ideas extend to higher dimensions. The results in Secs. 4 and 5, except for Lem. 2.3, can be trivially extended to n dimensions. Sec. 4 has been written only for the two-dimensional case, but a similar construction, and the same argument would work in higher dimensions (Lem. 4.5 being a hint of this). It is the argument in Sec. 5 that becomes problematic. While the ECE property is shown to be sufficient to prevent foldovers in the triangulation, it is not sufficient in higher dimensions. In particular, fixing the boundary to be simple and convex, there are simple arrangements of tetrahedra in \mathbb{R}^3 that contain face foldovers but do not break the ECE property. We plan to study these question next.

9. REFERENCES

- [1] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Curved Voronoi diagrams. In *Effective Computational Geometry for Curves and Surfaces*, Mathematics + Visualization, pages 67–116. Springer, 2007.
- [2] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, SCG '08, pages 270–277, New York, NY, USA, 2008. ACM.
- [3] G. D. Cañas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are triangulations. *arXiv preprint 1102.3673v1, cs.CG*, 2011.
- [4] G. D. Cañas and S. J. Gortler. Orphan-free anisotropic Voronoi diagrams. *Discrete & Computational Geometry*, 46(3):526–541, 2011.
- [5] K. L. Clarkson. Building triangulations using epsilon-nets. In *STOC 2006: Proceedings of the Thirty-eighth Annual SIGACT Symposium*, 2006.
- [6] B. N. Delaunay. Sur la sphère vide. *Bulletin of Academy of Sciences of the USSR*, (6):793–800, 1934.
- [7] Q. Du and D. Wang. Anisotropic centroidal Voronoi tessellations and their applications. *SIAM Journal of Scientific Computing*, 26(3):737–761, 2005.
- [8] H. Edelsbrunner. Triangulations and meshes in computational geometry. *Acta Numerica*, pages 133–213, 2000.
- [9] H. E. F. Aurenhammer. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition 17 (1984)*, 251–257.
- [10] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.
- [11] S. J. Gortler, C. Gotsman, and D. Thurston. Discrete one-parameter on meshes and applications to 3d mesh parameterization. *Computer Aided Geometric Design*, 23:83–112, 2006.
- [12] P. M. Gruber. Optimum quantization and its

applications. *Advances in Mathematics*, 186(2):456 – 497, 2004.

- [13] F. Labelle and J. R. Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SCG '03: Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, pages 191–200, New York, NY, USA, 2003. ACM.
- [14] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for riemannian manifolds. In *SCG '00: Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, pages 341–349, New York, NY, USA, 2000. ACM.

APPENDIX

A. PRELIMINARY PROOFS

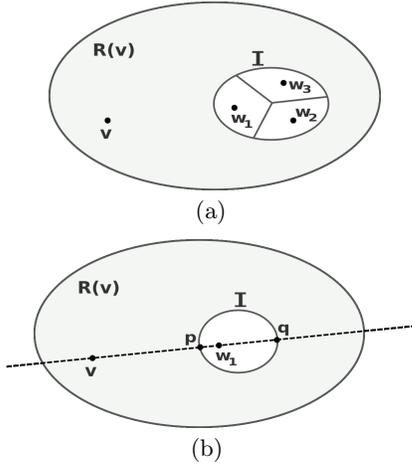


Figure 4: Diagram for the proof of Lem. 2.3.

Lemma 2.1 *Given two sites $v, w \in V$, $v \neq w$, the only point equidistant to v, w in their supporting line is the midpoint $(v + w)/2$.*

PROOF. If p is in the supporting line of v, w , then $p = (1 - \lambda)v + \lambda w$, $\lambda \in \mathbb{R}$. p is equidistant to v, w (i.e. $D(v, p) = D(w, p)$) if and only if $\lambda^2(v - w)^t Q_p(v - w) = (1 - \lambda)^2(v - w)^t Q_p(v - w)$.

Since Q is positive definite and $v \neq w$, it is $(v - w)^t Q_p(v - w) > 0$, and so the above equality holds iff $\lambda = 1/2$ (p is the midpoint). \square

Lemma 2.3 *Every Voronoi region of an orphan-free anisotropic Voronoi diagram in \mathbb{R}^2 is simply connected.*

PROOF. A multiply connected region $R(v) \subset \mathbb{R}^2$ is path connected, and is such that there is a map $f : \mathbb{S}^1 \rightarrow R(v)$ that cannot be continuously contracted to a point. We can assume f injective (simple), since a non-injective f can always be broken up into injective pieces, at least one of which must be such that it cannot be continuously contracted to a point (or else f would be). By the Jordan curve theorem, f encloses a bounded set B . Since f cannot be continuously contracted to a point, there must be $I \subseteq B$ with $I \subset \mathbb{R}^2 \setminus R(v)$. Since I is a subset of B , and B is bounded, I is bounded.

I is part of the Voronoi diagram, so it must be composed of one or more Voronoi cells. Because the diagram is orphan-free, each cell in the Voronoi diagram contains its generating

site: $I = \cup_i R(w_i)$, $i = 1, \dots, m$ (Fig 4(a)), for some number $m > 0$.

We can now consider what the diagram would look like if we remove, one by one, all of the sites w_i in I , until only one (w_1) is left. The new set of sites is $V' = V \setminus \{w_i : i = 2, \dots, m\}$. From the definition of Voronoi diagram it is clear that the region corresponding to site w_1 in the new diagram is $R'(w_1) \subseteq I$ (since no point is *strictly* closer to w_1 than to $\{w_i : i = 1, \dots, m\}$), and $R'(w_1) \neq \phi$ (since, by Lem. 2.2, it is $w_1 \in \text{int } R'(w_1)$).

In particular, since $R'(w_1) \neq \phi$, and $R'(w_1)$ is connected, $R'(v)$ is *still* multiply-connected, with $R'(w_1)$ playing the role of I in the new diagram (Fig. 4(b)). We show that this cannot be the case, resulting in a contradiction.

Since I is bounded, then $R'(w_1) \subseteq I$ must be bounded as well. And since w_1 is an interior point of $R'(w_1)$, and v is an interior point of $R'(v)$ (and therefore an interior point of the complement of $R'(w_1)$), then the line passing through v and w_1 must intersect $R'(v)$ at least at two distinct points p and q .

Since the boundary between $R'(v)$ and $R'(w_1)$ is composed of equidistant points to v, w_1 , then both p, q must be equidistant to v, w_1 . By Lem. 2.1, only one point on the line connecting v, w_1 can be equidistant to them, a contradiction.

\square

B. COLINEAR SITES

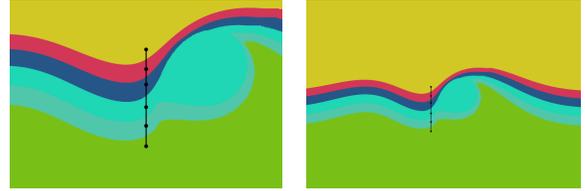


Figure 5: If all sites are colinear, the dual is always a polygonal chain.

If all the sites are colinear then the structure of the Voronoi diagram \tilde{P} is greatly simplified, and always has the form shown in Fig. 5. In particular, \tilde{P} has no vertices aside from p_∞ since vertices are equidistant to three or more sites, and no point $p \in \mathbb{R}^2$ can be equidistant to three colinear sites (since points equidistant to p form an ellipse, and a line intersects an ellipse at most twice).

Consider the set of sites $V = \{w_l : l = 1, \dots, m\}$ ordered linearly along their supporting line. We show that the graph dual to P , having V as vertices, has edges $\{(w_l, w_{l+1}) : l = 1, \dots, m - 1\}$.

For every pair of sites w_l, w_{l+1} , because they are consecutive, and by Lem. 2.1, the midpoint $m_{l,l+1} = (w_l + w_{l+1})/2$ is closest and equidistant to w_l, w_{l+1} , and therefore the edge (w_l, w_{l+1}) is in the dual.

We finally show that every dual edge (w_i, w_j) is of the form (w_l, w_{l+1}) . Assume otherwise, and therefore that, since w_i, w_j are not consecutive, there is some w_k between them. Because there is a dual edge (w_i, w_j) , the corresponding primal edge \tilde{E}_{w_i, w_j} in \tilde{P} is not empty, and therefore there is some point $p_{ij} \in \tilde{E}_{w_i, w_j}$ that is closest and equidistant to w_i, w_j . By the convexity of $D(\cdot, p_{ij})$, and the fact that

$w_k \in \overline{w_i w_j}$, it is $D(w_k, p_{ij}) < D(w_i, p_{ij}) = D(w_j, p_{ij})$, a contradiction. Therefore (w_i, w_j) are consecutive sites.

Since we have characterized the set of edges and vertices, and there are no (interior) faces in the dual, this completely determines the dual when all sites are colinear.

Perhaps interestingly, this means that, in the colinear case, the structure of the dual does not depend on the metric.

C. GENERAL CASE

So far we have assumed that Voronoi vertices are incident (equidistant) to no more than three sites. In general, however, Voronoi vertices may be incident to three or more sites, and the straight-edge dual \bar{G} will be a polygonization, instead of a triangulation (simplicial complex). We show here that, even in this, more general case, the polygonization \bar{G} is embedded, and can be easily triangulated into an embedded simplicial complex. The argument is quite simple and is given in summary.

Theorem 5.9 *If its (topological) boundary is simple and closed, then the straight-line dual \bar{G} of an orphan-free diagram is an embedded polygonal mesh with convex faces.*

PROOF. First assume that only Voronoi vertex c is incident to more than three sites $\{v_i : i = 1 \dots m\}$, $m > 3$, ordered in, say, clockwise order around c (which is possible since they are equidistant to c and therefore lie at the boundary of the ellipse $\theta_c(v_1)$). Note that, since the $\{v_i : i = 1 \dots m\}$ are equidistant to c , then any polygon (or triangle) connecting any of the $\{v_i\}$ satisfies the empty circum-ellipse property (with empty circum-ellipse $\theta_c(v_1)$). The Voronoi regions $\{R_i : i = 1, \dots, m\}$, corresponding to those sites, are the only ones incident to c . We can order these regions in clockwise order around c : $\{R_{i_k} : k \in K\}$ (for some index sequence K with $k \in K$ satisfying $k \in \{1, \dots, m\}$), and therefore the polygon Π dual to c will connect vertices $\{v_{i_k} : k \in K\}$, in that order. Now choose any v_j in Π and triangulate Π in a fan arrangement centered at v_j . As pointed out above, both Π and every triangle in the triangulation of Π satisfies the empty circum-ellipse property. Therefore the new, triangulated \bar{G} (which replaces Π by its triangulation), satisfies all the conditions needed for Theorems 5.9 and 4.7 (as well as Lem. 4.9) to hold, and in particular, it is embedded.

The claim now easily follows: any closed polygon with vertices $\{v_{i_k} : k \in K\}$, such that all of its fan-arranged triangulations centered at every one of its vertices is embedded, must be convex (given a polygon with non-convex vertex w , triangulate it in a fan centered at a vertex incident to w : this triangulation is not embedded). Lastly, since vertices $\{v_i : i = 1 \dots m\}$ are equidistant from c , the only polygon connecting them that is also convex is the one where vertices are arranged in clockwise (counter-clockwise) order around c .

The argument has been made for a single Voronoi vertex incident to more than three sites, but the same argument applies to diagrams in which there are more such vertices (the key being that the triangulations of each dual polygon are independent of one another). \square

D. BOUNDED ANISOTROPY

We show here that we may drop the bounded anisotropy condition from Thm. 6.1, at the expense of losing the prop-

erty of covering the convex hull of the sites; thus concluding that the dual of an orphan-free anisotropic diagram is an embedded polygonal mesh with convex faces. The proof proceeds by reducing this case to that of Thm. 6.1. Because of its simplicity, it is given here in summary.

Assume given a continuous metric Q defined over \mathbb{R}^2 , and a set V of sites forming an orphan-free anisotropic Voronoi diagram V_Q . For now, assume that the set of Voronoi vertices is finite, and therefore bounded. Since there is $\rho > 0$ such that all Voronoi vertices are inside the origin-centered ball $B(0; \rho)$ of radius ρ , we can construct a new metric $Q'(p) = Q(p) \cdot (1 - \lambda(p)) + I \cdot \lambda(p)$, where

$$\lambda(p) = \begin{cases} 0, & \text{if } \|p\| < \rho \\ \|p\| - \rho, & \text{if } \rho \leq \|p\| \leq 2\rho \\ 1, & \text{if } \|p\| > 2\rho \end{cases}$$

Clearly, Q' is continuous, since both Q and λ are. Since it is $Q' = I$ outside of $B(0; 2\rho)$, and $B(0; 2\rho)$ is compact, then Q' has bounded anisotropy (ratio of eigenvalues). We may then apply Thm. 6.1 to conclude that the dual $\bar{G}_{Q'}$ of $V_{Q'}$ is an embedded polygonal mesh with convex faces. Since $B(0; \rho)$ includes all Voronoi vertices of V_Q , and it is $Q' = Q$ inside $B(0; \rho)$, then all Voronoi vertices in V_Q are also in $V_{Q'}$. Therefore, by duality, all faces of \bar{G}_Q are also in $\bar{G}_{Q'}$. Finally, since $\bar{G}_{Q'}$ is embedded with convex faces, and removing faces from a polygonal mesh preserves both properties, then \bar{G}_Q is also embedded with convex faces.

The only assumption we have made is that the set of Voronoi vertices of V_Q is finite, which can be justified as follows.

First, we show that Voronoi vertices cannot be isolated, and are always incident to some Voronoi edge. Consider a Voronoi vertex c_{ijk} closest and equidistant to sites v_i, v_j, v_k (the reasoning for vertices equidistant to more sites is analogous). Since Q , and therefore D , is continuous, there is a (possibly small) closed ball $B(c_{ijk}, \epsilon)$ centered at c_{ijk} where all points are *strictly* closer to v_i, v_j, v_k than to any other site, and therefore in $B(c_{ijk}, \epsilon)$, the Voronoi diagram of V is the same as that of $\{v_i, v_j, v_k\}$. We consider the Voronoi diagrams $\{v_i, v_j\}_Q$ and $\{v_j, v_k\}_Q$ of $\{v_i, v_j\}$ and $\{v_j, v_k\}$ (restricted to $B(c_{ijk}, \epsilon)$), respectively. By Corollary 2.5, the only Voronoi edge of $\{v_i, v_j\}_Q$ is connected, and likewise for $\{v_j, v_k\}_Q$. Since c_{ijk} is in their intersection, then it must be incident to both, and therefore it is not isolated.

By Corollary 2.5, the edges (connected sets of points equidistant to two given sites) of V_Q are unique (thus finite since V is finite), and by the induction argument of Sec. 2, each edge connects two Voronoi vertices. If there are n_e Voronoi edges then, since Voronoi vertices are always incident to some Voronoi edge, and each edge connects two vertices, there cannot be more than $2n_e$ Voronoi vertices. In particular, the number of Voronoi vertices is finite. This concludes the proof.

Finally, note that, from the way we have defined the dual \bar{G}_Q , and the fact that it is embedded, we can conclude that Voronoi vertices are unique (otherwise, multiple Voronoi vertices would result in coincident dual polygons, which would contradict the fact that \bar{G}_Q is embedded). This, together with Corollary 2.5, and the orphan-freedom assumption means that, as stated in Corollary 6.3, orphan-freedom is sufficient to ensure that all the elements (vertices, edges, faces) of the primal diagram are unique.