Title:
Weighted Skeletal Structures for Computing Variable-Radius Offsets

Authors:
Martin Held, held@cs.sbg.ac.at, University of Salzburg, Salzburg, Austria
Stefan de Lorenzo, slorenzo@cs.sbg.ac.at, University of Salzburg, Salzburg, Austria

Keywords:
Voronoi diagram, multiplicative weights, wavefront propagation, variable-radius skeleton, variable-radius offset

DOI: 10.14733/cadconfP.2020.46-50

Introduction:

Consider a set S of n points in the plane. The Voronoi diagram, $\mathcal{VD}(S)$, of S partitions the plane into interior-disjoint, convex areas — so-called Voronoi regions — such that the Voronoi region $\mathcal{VR}(s, S)$ of $s \in S$ contains all points of the plane that are closer to s under the Euclidean distance metric than to any other point of S; see Figure 1a. It is well-known that $\mathcal{VD}(S)$ has $O(n)$ nodes and (straight-line) edges.

In the prairie-fire analogy, $\mathcal{VD}(S)$ is given by those points of the plane where fire waves meet if (1) fires are ignited simultaneously at all points of S and if (2) all fires spread uniformly at the same speed. If we assign multiplicative positive weights to the points of S and let them influence the expansion speeds of the fires then we get the multiplicatively weighted Voronoi diagram of S; see Figure 1b. More formally, the weighted distance $d_w(p, s)$ between a point $p \in \mathbb{R}^2$ and a weighted point $s \in S$ is defined as $d_w(p, s) := \frac{d(p,s)}{w(s)}$, where $d(p, s)$ denotes the standard Euclidean distance between p and s, and $w(s)$ is the positive weight of s. Hence, the larger the weight of a point the quicker its fire spreads.

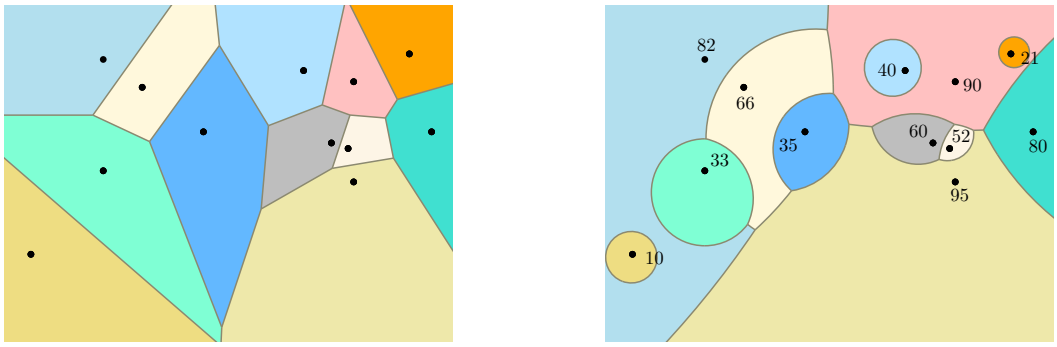

(a) (b)

Fig. 1: (a) Voronoi diagram and (b) multiplicatively weighted Voronoi diagram of a set of points. The multiplicative weights are written next to the points. Each Voronoi region is indicated by a colored area.

Held et al. [3] generalize this concept to generalized weighted Voronoi diagrams (GWVDs) of a set S

of points and straight-line segments: They assign weights to the end-points a and b of an input segment s and then obtain the weight of a point q on s by a linear interpolation of the weights of a and b. Then $d_w(p, s) := \min_{q \in s} d_w(p, q)$, and the corresponding GWVD $\mathcal{VD}_w(S)$ is defined accordingly. See Figure 2a for a sample GWVD.

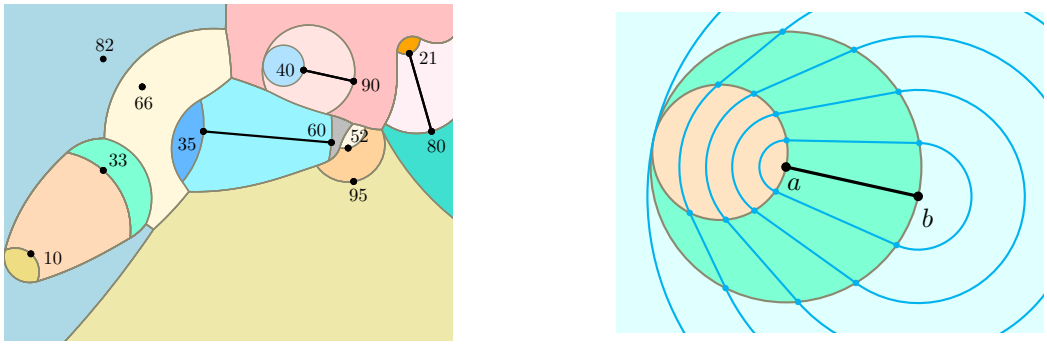

(a) (b)

Fig. 2: (a) Generalized weighted Voronoi diagram (GWVD) of a set of weighted points and straight-line segments (highlighted in black). (b) Wavefronts of one weighted straight-line segment $\overline{ab}$.

As discussed in [3], a fire wavefront of a weighted straight-line segment $\overline{ab}$ is formed by two circular arcs, which are induced by the end-points a and b, and two straight-line segments; see Figure 2b. A similar result holds if identical weights are assigned to the end-points of a circular arc. Wavefronts of weighted segments and arcs are known as variable-radius offsets; see Figure 3. Among other fields, variable-radius offsets find applications in brush-stroke modeling and the generation of ornamental seams.

Edelsbrunner and Seidel [2] establish the connection between Voronoi diagrams in $\mathbb{R}^d$ and lower envelopes in $\mathbb{R}^{d+1}$. Setter et al. [4] use this connection to add a powerful tool to CGAL [1] for computing Voronoi diagrams via lower or upper envelopes of suitably formed surfaces in 3D. Held et al. [3] argue that these surface are sub-surfaces of right conoids in the case of weighted points and segments. Their implementation is based on this tool provided by CGAL.

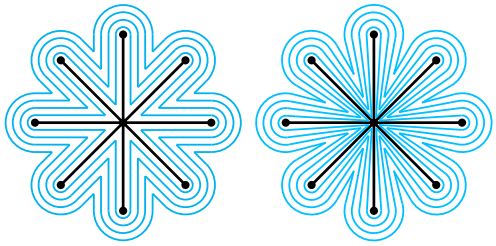

(a) (b)

Fig. 3: (a) Constant-radius offsets of unweighted segments and (b) variable-radius offsets of weighted segments. The weight of the center point is half of the weight of all other end-points.

The obvious practical problem of a GWVD is that a Voronoi region may consist of several connected components even if the edges of S form a polygon P; see Figure 4(a). If such a diagram is used to generate ornamental seams then we would get seam curves at locations inside of P where one would not expect to see them, see Figure 4(b). In the sequel we sketch how a skeletal structure used for variable-radius offsetting can be obtained such that all Voronoi-like regions remain connected, without resorting to intricate computations in higher dimensions.

Variable-Radius Skeleton:

Assume that the weighted points and straight-line segments of S form the vertices and edges of a polygon P. Recall that a (positive) weight $w(v)$ is assigned to every vertex v of P. A vertex is *dominant* if its

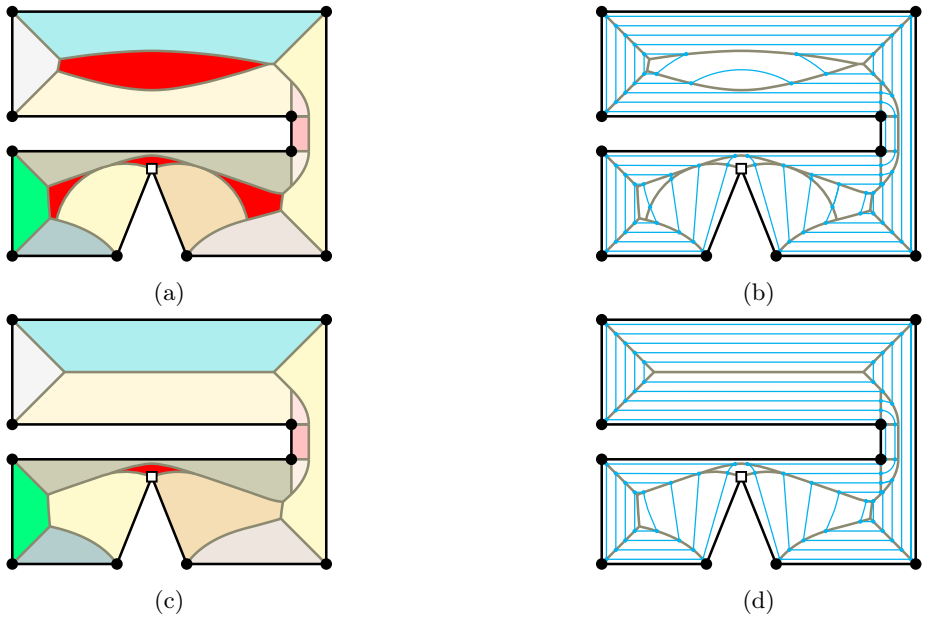

Fig. 4: (a) GWVD and (b) the corresponding offsets inside a polygon P where the node indicated by a square has three times the weight of all other nodes. Note that the Voronoi region of the "square" node consists of four connected components shown in red. (c) VRS and (d) the corresponding offsets.

weight is greater than the weight of one of its two neighbors.

We call a vertex v *reflex* (*convex*) if the interior angle at v is greater (smaller, resp.) than 180°. (For the sake of descriptional simplicity we assume that no interior angle equals 180°.) Let A denote the area bounded by P. We now explain how we obtain a Voronoi-like structure inside of A, the variable-radius skeleton (VRS) $\mathcal{S}_w(P)$, such that all regions induced by this structure are connected; see Figure 4.

To compute $\mathcal{S}_w(P)$ we apply the well-known idea of wavefront propagation: Starting at $t := 0$, the expansion of the wavefront $\mathcal{W}(P,t)$ is simulated by continuously increasing the time (or weighted distance) t. We know that $\mathcal{W}(P,t)$ is formed by (possibly several) closed curvilinear chains that consist of straight-line segments and circular arcs as wavefront edges. Every straight-line segment of the wavefront corresponds to an edge of P while every circular arc corresponds to a vertex of P. For $t = 0$ we have $\mathcal{W}(P,t) = P$. For a sufficiently small time t_0, a counter-clockwise traversal of $\mathcal{W}(P,t_0)$ will match a counter-clockwise traversal of P such that every straight-line segment of $\mathcal{W}(P,t_0)$ corresponds to exactly one edge of P and every arc of $\mathcal{W}(P,t_0)$ corresponds to exactly one reflex vertex of P, and vice versa. This *initial wavefront* is shown in Figure 5(a) for a sample polygon.

It is obvious that the vertices of $\mathcal{W}(P,t)$ change their locations as t is increased. During the wavefront propagation process, we keep track of the traces of these so-called *moving intersections*. However, a change of the locations of the wavefront is not the only change that will happen. Rather, new wavefront edges may appear, old wavefront edges may disappear, and the wavefront may split into two or more connected *wavefront components*.

These combinatorial changes are witnessed by the following three types of events; see Figure 5. A *split event* occurs at time t' if an arc of $\mathcal{W}(P,t')$ collides with another arc or a segment of $\mathcal{W}(P,t')$, there splitting $\mathcal{W}(P,t)$ into two wavefront components for $t > t'$. An *edge event* occurs at time t' if the two end-points of an edge e of $\mathcal{W}(P,t')$ coincide in $\mathcal{W}(P,t')$, thus causing e to disappear and to be absent

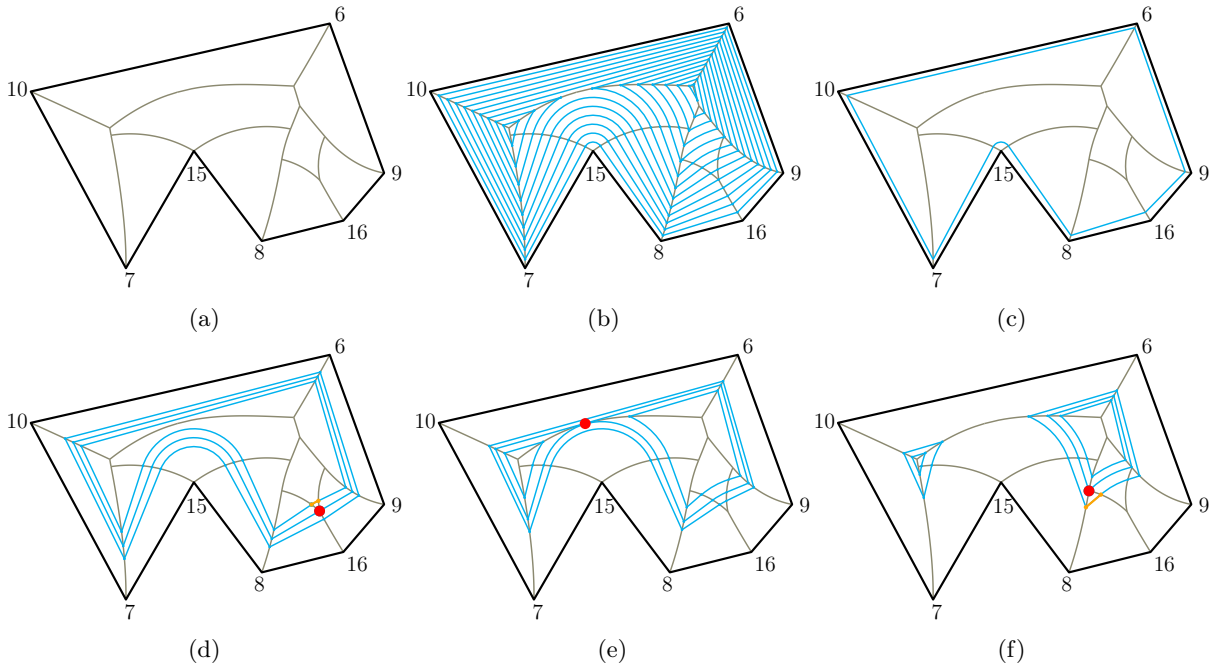

Fig. 5: (a) Sample variable-radius skeleton (VRS) of a polygon and (b) a family of offsets/wavefronts. The numbers next to the vertices of the polygon indicate their weights. (c) Initial wavefront; (d) break-through event; (e) split event; (f) edge event. The location of each event is marked by a red dot, and the wavefront edge that appears of disappears is drawn in orange.

from $\mathcal{W}(P,t)$ for $t > t'$. A *break-through event* occurs at time t' if, for a dominant convex vertex v_i of P, the wavefront segments induced by the polygon edges $\overline{v_{i-1}v_i}$ and $\overline{v_iv_{i+1}}$ are collinear at time t'. In this case $\mathcal{W}(P,t)$ will contain a circular arc centered at v_i for $t > t'$. That is, a break-through event causes a circular arc to appear on the wavefront.

The locations of edge events and break-through events coincide with nodes of $\mathcal{S}_w(P)$. At an edge event one of the faces induced by $\mathcal{S}_w(P)$ is closed, while a new face is started at a break-through event. It is easy to see that we get at most one face per edge or vertex of P. Since $\mathcal{S}_w(P)$ is planar, this implies that the numbers of edges and nodes of $\mathcal{S}_w(P)$ are linear in the number of edges/vertices of P.

All events are stored in a priority queue $\mathcal{Q}$ ordered by the times of their occurrences. Every wavefront arc holds a pointer to each of its current neighbors in the corresponding wavefront component. Initially, the times of all potential split events as well as all break-through events are computed and pushed to $\mathcal{Q}$. Additionally, for every wavefront edge that shrinks to zero length an edge event is inserted into $\mathcal{Q}$. After the initialization phase is completed the individual events are successively popped from $\mathcal{Q}$. Since a break-through event results in the insertion of a circular wavefront arc into the appropriate wavefront component, checks for split events between this arc and other edges of the wavefront component are carried out. The wavefront propagation terminates once all wavefront components have disappeared.

Discussion:

With minor modifications the wavefront propagation to compute a VRS can be extended to any set of weighted points, weighted straight-line segments and constantly-weighted circular arcs as input sites: We

only need to demand that no pair of input sites intersects in a point other than a common end-point. A similar approach allows to compute a GWVD.

We have been working on a C++ implementation of the wavefront-based construction of both the GWVD as well as the VRS. Our implementation operates entirely in two dimensions and is based on conventional IEEE 754 floating-point arithmetic. The skeletons and offsets shown in Figure 6 and the other figures were generated by our implementation. Since our implementation has not yet been engineered for speed or robustness, we refrain from presenting any performance statistics here.

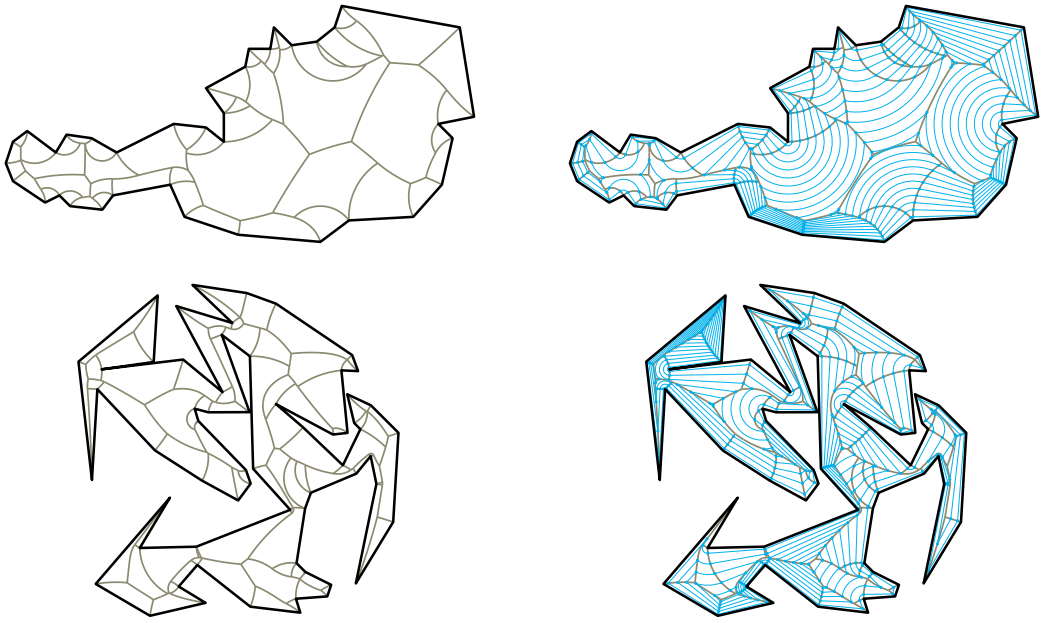

Fig. 6: Variable-radius skeletons and families of variable-radius offsets.

Acknowledgements:
Work supported by Austrian Science Fund (FWF): Grant P31013-N31.

References:

[1] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org/.

[2] Edelsbrunner, H.; Seidel, R.: Voronoi Diagrams and Arrangements. Discrete & Comp. Geom., 1(1), 25-44, 1986. http://doi.org/10.1007/BF02187681.

[3] Held, M.; Huber, S.; Palfrader, P.: Generalized Offsetting of Planar Structures using Skeletons. Comput. Aided Design & Appl., 13(5), 712-721, 2016. http://doi.org/10.1080/16864360.2016.1150718.

[4] Setter, O.; Sharir, M.; Halperin, D.: Constructing Two-Dimensional Voronoi Diagrams via Divide-and-Conquer of Envelopes in Space. In M.L. Gavrilova; C.J.K. Tan; F. Anton, eds., Transactions on Computational Science IX, vol. 6290 of LNCS, 1-27. Springer, 2010. http://doi.org/10.1007/978-3-642-16007-3_1.