


Forschungsschwerpunkt S92

Industrial Geometry

<http://www.ig.jku.at>



FSP Report No. 47

Weighted Skeletons and Fixed-Share Decomposition

F. Aurenhammer

June 2006

FWF

Der Wissenschaftsfonds.



Weighted Skeletons and Fixed-Share Decomposition*

FRANZ AURENHAMMER

Institute for Theoretical Computer Science

University of Technology, Graz, Austria

auren@igi.tugraz.at

Abstract

We introduce the concept of weighted skeleton of a polygon and present various decomposition and optimality results for this skeletal structure when the underlying polygon is convex.

1 Introduction

Polygon decomposition is a major issue in computational geometry. Its relevance stems from breaking complex shapes (modeled by polygons) into sub-polygons that are easier to manipulate, and from subdividing areas of interest into parts that satisfy certain containment requirements and/or optimality properties. We refer to [13] for a nice survey on this topic. In particular, a rich literature exists on decomposition into convex polygons. Convex decompositions are most natural in some sense. They have many applications and can be computed efficiently; see e.g. [7, 14, 16].

In this paper, we focus on the problem of decomposing a *convex* polygon such that predefined constraints are met. More specifically, the goal is to partition a given convex n -gon P into n convex parts, each part being based on a single side of P and containing a specified 'share' of P . The share may relate, for example, to the spanned area, to the number of contained points from a given point set, or to the total edge length covered from a given set of curves. Possible applications of such fixed-share decompositions include priority-based or fair facility allocation, which may concern real estate or access to power lines, aquifers, or oil wells.

We introduce the concept of weighted skeleton for a convex polygon P and exploit it to prove existence of various types of fixed-share decompositions of P . Moreover, the decompositions induced by weighted skeletons will be shown to have several optimality properties. For instance, the average normal distance to the respective sides of P is minimized, among all decompositions of P into polygons with fixed areas. Our approach is inspired by results on power diagrams in [5]. However, these substantially differ from our setting because of the boundary conditions imposed by P . Algorithmic methods in [5] can be generalized to compute fixed-share decompositions but are inherently slow, at least from the theoretical point of view. For this reason, we outline a simple and efficient divide-and-conquer algorithm that produces such decompositions, but without further optimality properties.

Weighted skeletons are generalizations of the straight skeleton of a polygon introduced in [2]. They can be defined for arbitrary simple polygons, by individually tuning the translation speed of polygon sides

*Research supported by the FWF Joint Research Project 'Industrial Geometry', S9205-N12.




Figure 1: Weighted skeleton for a convex 6-gon

in the defining shrinking process. For a restricted class of non-convex polygons, such skeletons have been used in [3] for different purposes. Unfortunately (for reasons to be explained later), weighted skeletons do not lead to corresponding decomposition results when convexity of the input polygon is dropped. However, generalizations to convex polytopes in \mathbb{R}^d are possible; our methods for proving the existence of fixed-share decompositions do not depend on the underlying dimension.

2 Weighted skeletons

We start with recalling the definition of a straight skeleton. Weighted versions of this structure are then introduced, and some of their basic properties discussed.

Let P be a polygon with n sides in the plane. The straight skeleton [2] of P is a skeletal structure in the interior of P which is defined procedurally as follows. Move inwards the boundary of P in a self-parallel fashion and at constant speed, until it contracts to area zero, possibly after having split at places of prior self-contact. During this shrinking process, the breakpoints of the boundary of P trace out the edges of the skeleton, which are pieces of angle bisectors for P . The straight skeleton exists for arbitrary simple polygons P . If P is convex then exactly the medial axis [15] of P is obtained: For each point x on an edge of the skeleton, the minimum distance from x to the boundary of P is realized by two or more points.

For the purposes of the present paper, a weighted version of straight skeleton for a *convex* polygon is required. Let $W = (w_1, \dots, w_n)$ be an n -tuple that assigns an individual weight $w_i \geq 0$ to each side e_i of the given convex polygon P . Weight w_i expresses the speed at which (the current portion of) e_i is translated in the shrinking process. The structure resulting from this process is termed the *weighted skeleton* of P , denoted by $SK_W(P)$. The portion of the plane swept out by a given side e_i is called the region, $reg(e_i)$, of e_i . We have $reg(e_i) = e_i$ if and only if $w_i = 0$. (We say that $reg(e_i)$ is degenerate in this case.) Clearly, $SK_W(P)$ is the classical straight skeleton of P if and only if all weights are positive and equal. Figure 1 illustrates a weighted skeleton for a convex polygon with six sides. Numbers at sides denote side weights.

The translate of a side e_i of P cannot re-appear after having shrunk to length zero. This implies that $reg(e_i)$ is a connected set. If $reg(e_i)$ is non-degenerate then it has e_i as a boundary edge. We conclude that the edges of $SK_W(P)$ form a tree with exactly n leaf vertices (the vertices of P) and at most $n - 2$ non-leaf vertices.




Figure 2: Discontinuous region change

When all weights are positive then $SK_W(P)$ can be interpreted as a *weighted medial axis* of P . Define the weighted distance of a point $x \in P$ to a side e_i as $d(x, e_i) = \nu(x, e_i)/w_i$, where $\nu(x, e_i)$ denotes the normal distance of x from the line supporting e_i . Then, for any $i \neq j$, the two regions $reg(e_i)$ and $reg(e_j)$ are separated by the straight line defined by all points x of equal weighted distance from both e_i and e_j . This shows that each region $reg(e_i)$ is the intersection of $n - 1$ halfplanes, $n - 1$ coming from other sides, and one coming from the line supporting e_i . We conclude that $reg(e_i)$ is either degenerate or a convex polygon with at most n edges. By convention, regions are considered as closed sets.

The shrinking process we used above to define $SK_W(P)$ generalizes directly for arbitrary simple polygons P . A tree-like structure that decomposes P into (possibly non-convex) polygonal regions is obtained. However, when altering the weights for the individual sides of P , we face the following peculiarity: The shape of regions does not change continuously with the weights. Figure 2 gives an example. The entire shaded area displays $reg(e_1)$ for a given tuple $(w_1, w_2, w_3, w_4, w_5)$ of weights. A slight increase of w_5 clips off from $reg(e_1)$ the portion shaded in light grey, and enlarges the regions of e_5 , e_4 , and e_3 accordingly. This effect cannot be observed when P is a convex polygon – a property that will be crucial for the proof of Theorem 1 in the subsequent section.

3 Decomposition results

Weighted skeletons are capable of generating tailor-made decompositions of convex polygons. In particular, as will be shown in this section, fixed-share decompositions of various kinds can be obtained.

Let ρ be any continuous density function on a given convex polygon P , and let μ be the measure defined by ρ on P . We have the following general theorem.

Theorem 1 *Let P be a convex polygon in the plane, having sides e_1, \dots, e_n . For any non-negative real numbers A_1, \dots, A_n whose sum is 1, there exists an n -tuple W of weights such that $\mu(reg(e_i)) = A_i$ holds for each region of $SK_W(P)$.*

Proof. For any positive number $\lambda \in \mathbb{R}$, we have $SK_{\lambda W}(P) = SK_W(P)$. This directly follows from the procedural definition of a weighted skeleton. Restricting the domain of weights $D \subset \mathbb{R}^n$ by requiring $\sum_{i=1}^n w_i = 1$ thus means no loss of generality. Note that, by non-negativity of weights, D is just the convex hull spanned by the n unit vectors of \mathbb{R}^n . In particular, D is a bounded and closed set. Observe next that,

for each fixed index i , the region $reg(e_i)$ continuously expands when w_i is increased. This property is evident from the distance-based definition of $SK_W(P)$. Thus, by the assumed continuity of ρ , the measure $\mu(reg(e_i))$ is a continuous function of w_i . Consider the function $\Psi : D \rightarrow \mathbb{R}$, given by

$$\Psi(W) = \sum_{i=1}^n |\Delta_i|, \text{ where } \Delta_i = \mu(reg(e_i)) - A_i.$$

Ψ is a continuous and nonnegative function on a compact domain, so Ψ achieves a global minimum. Let W^* be a corresponding n -tuple. We are left with proving $\Psi(W^*) = 0$. Assume, to the contrary, that $\Psi(W^*) > 0$. Then, as $\sum A_i = \sum \mu(reg(e_i)) = 1$, there exist indices j and k such that $\Delta_j < 0$ and $\Delta_k > 0$. Consider all the regions of the former type (i.e., whose measures are too small), and let Q be the unique connected component of their union which contains $reg(e_j)$. Expand Q by multiplying the weight for each region in Q by a fixed factor of $c > 1$, where c is chosen such that $\mu(Q)$ increases by an arbitrarily small amount. This is possible because $\mu(Q)$ continuously grows with c . Thereby, some regions adjacent to Q (which have been too large) decrease in measure, whereas regions in Q only can increase, because separators between them do not change. For the resulting n -tuple W' , we thus have $\Psi(W') < \Psi(W^*)$. But this contradicts our assumption that W^* is a global minimum. \square

If the density function ρ is non-vanishing all over P , then the weighted skeleton that achieves the required measures in Theorem 1 is unique: Any change in ratio between weights results in a change of measure for at least two regions. Taking ρ as a suitable constant gives uniform distribution on P , and we obtain:

Corollary 1 *For any choice of n areas that sum up to the area of P , there exists a unique weighted skeleton for P whose regions realize this choice.*

Theorem 1 can be shown to hold for certain non-continuous density functions as well. Let S be an arbitrary subset of P , and consider the indicator function ψ of S in P . That is, $\psi(x) = c > 0$ if $x \in S$, and $\psi(x) = 0$, otherwise. The constant c is chosen such that ψ is indeed a density function on P . We may express ψ as the limit of a series of continuous density functions $(\rho_j)_{j \geq 1}$ as follows. Let $K_j(x)$ be a vertical and downwards directed cone, with aperture angle $\frac{1}{j}$ and apex at (fixed) height h above the point $x \in P$. The upper envelope of the cones $K_j(x)$, for all $x \in S$, is the graph of a continuous function, κ_j , on the plane. For a suitable choice of h , the pointwise maximum of κ_j with the zero function gives a continuous density function, ρ_j , on P , and we have

$$\lim_{j \rightarrow \infty} \rho_j = \psi.$$

We conclude that Theorem 1 is valid for the measure μ defined by ψ .

Various implications are obtained when the dimension of the point set S that underlies ψ is varied. For example, S may be the union of finitely many two-dimensional objects, like triangles or disks. A prescribed assignment of their areas to the sides of P exists — a result with possible applications for priority-based (or fair) facility allocation. This also applies to one-dimensional objects, for instance, to a finite set of length-measurable curves that might model power lines or aquifers to be accessed. The measure $\mu(Q)$ of a subset $Q \subseteq P$ then becomes the total curve length in Q . More specifically, we have:

Corollary 2 *Let μ be defined by a finite set of curves (of total length 1) in a convex n -gon P . For any non-negative real numbers ℓ_1, \dots, ℓ_n that sum up to 1, there exists a weighted skeleton $SK_W(P)$ with $\mu(reg(e_i)) \geq \ell_i$, for $i = 1, \dots, n$.*

Observe that inequality arises above because portions of curves might have to be shared by two (closed) regions of $SK_W(P)$. Finally, in the case where S is a finite point set, the quantity $\mu(Q) \cdot |S|$ counts the points of S in a given subset Q of P .




Figure 3: Fixed-area decompositions are not unique

Corollary 3 *Let a convex polygon P enclose a set S of m points. For every choice of non-negative integers m_1, \dots, m_n whose sum is m , there exists a weighted skeleton $\mathcal{SK}_W(P)$ such that $|\text{reg}(e_i) \cap S| \geq m_i$, for $i = 1, \dots, n$.*

4 Optimality properties

Fixed-share decompositions of a convex polygon are, in general, not unique, even when decomposition into convex pieces is required. Interestingly, those obtained from weighted skeletons show several optimality properties. The present section addresses this issue.

Let us call a decomposition of a convex n -gon P *proper* if it consists of (exactly n) simple polygons, each having a single side in common with P . Clearly, every weighted skeleton for P gives rise to a proper decomposition of P , provided there are no degenerate regions. Corollary 3 implies a result proved earlier, and by different means, in [10]: Given P and a finite set S of points inside P , there always exists a proper decomposition of P into *convex* polygons, each containing a prescribed number of points from S . Points of S on polygon boundaries may have to be assigned appropriately. Even when boundary ambiguities do not arise, such fixed-cardinality decompositions need not be unique with respect to the assignment of points in S to polygons. In fact, the decompositions constructed in [10] induce assignments which are, in general, different from those we obtain via weighted skeletons.

As has to be expected, prescribing the *areas* of the polygons in a proper and convex decomposition also does not lead to a unique solution. Figure 3 exemplifies this fact. It shows, in full lines, some proper decomposition Π of P . Dashed lines delineate the (unique) weighted skeleton of P that realizes the same polygon areas. The fact that Π is not a weighted skeleton of P can also be seen from Lemma 1 given at the end of this section: The three straight lines through edge b , side e_1 , and side e_3 , respectively, do not concur in a point.

We are going to prove that weighted skeletons exhibit certain optimality properties. Let $\nu(x, e_i)$ denote the normal distance of a point x to the line supporting a given side e_i of P . Then the average normal distance is minimized in the following respect.

Theorem 2 Let $\Pi : P \rightarrow \{e_1, \dots, e_n\}$ be a proper (not necessarily convex) decomposition of the convex polygon P , having fixed polygon areas A_1, \dots, A_n . The weighted skeleton of P that assumes these areas (uniquely) minimizes the expression

$$\int_{x \in P} \nu(x, \Pi(x)) \, dx \quad (1)$$

over all such decompositions Π .

Proof. Let $\mathcal{SK}_W(P)$ be the weighted skeleton whose regions achieve the required areas A_1, \dots, A_n . $\mathcal{SK}_W(P)$ uniquely exists by Corollary 1. Fix a decomposition Π as above, but different from $\mathcal{SK}_W(P)$. For each index i , the polygon $\Pi^{-1}(e_i)$ as well as the region $\text{reg}(e_i)$ of $\mathcal{SK}_W(P)$ have e_i as one side, and both polygons have the same area, A_i . Thus there exists a cycle of consecutively overlapping polygons $\text{reg}(e_i), \Pi^{-1}(e_i), \text{reg}(e_j), \Pi^{-1}(e_j), \text{reg}(e_k), \Pi^{-1}(e_k), \dots, \text{reg}(e_i)$. Take some re-assignment of points from each polygon to its successor in the cycle, such that (1) polygon areas stay unchanged, (2) the area of $\text{reg}(e_t) \cap \Pi^{-1}(e_t)$ increases for all indices t involved, and (3) the resulting decomposition Λ is still proper. This can always be achieved when the intersection of two consecutive polygons is split appropriately with straight edges. Let now $W = (w_1, \dots, w_n)$. By definition, $\mathcal{SK}_W(P)$ assigns each point $x \in P$ to the side e_i of P that minimizes the weighted distance $\nu(x, e_i)/w_i$. Thus, when comparing the decompositions Λ and Π , we have

$$\int_{x \in \Lambda^{-1}(e_t)} \frac{\nu(x, e_t)}{w_t} \, dx < \int_{x \in \Pi^{-1}(e_t)} \frac{\nu(x, e_t)}{w_t} \, dx \quad (2)$$

for the indices t above. Clearly, for all remaining indices, equality holds in (2). Dividing by w_t , and summing up over all sides e_t , shows that Π does not minimize the integral in (1). The theorem follows. \square

Define the *altitude* of a polygon $Q = \Pi^{-1}(e_i)$ as the maximum normal distance $\nu(x, e_i)$ that occurs for a point $x \in Q$. Let the *altitude* of Π be the largest altitude of a polygon in Π . We have:

Corollary 4 Among all proper decompositions of P with fixed polygon areas, the weighted skeleton of P that assumes these areas is minimal in altitude.

Proof. Using the same construction as in the proof of Theorem 2, the cyclic re-assignment of points gives

$$\max_{x \in \Lambda^{-1}(e_t)} \frac{\nu(x, e_t)}{w_t} \leq \max_{x \in \Pi^{-1}(e_t)} \frac{\nu(x, e_t)}{w_t}.$$

Dividing by w_t , and taking the maximum over all sides e_t , implies that the altitude of Π , which is given by $\max_{x \in P} \nu(x, \Pi^{-1}(x))$, is minimized if $\Pi = \mathcal{SK}_W(P)$. \square

There also exist discrete variants of Theorem 2 and Corollary 4. They assert that the sum (or the maximum) of the normal distances in a given finite point set $S \subset P$ to the respective supporting lines of P is minimized over all assignments $S \rightarrow \{e_1, \dots, e_n\}$ of fixed subset cardinalities. The proofs are similar and are left to the interested reader.

Theorem 2 has an obvious geometric interpretation. The normal distance $\nu(x, e_i)$ is a linear function, g_i , on P . The corresponding plane $z = g_i(x)$ in \mathbb{R}^3 has slope 1 and contains the side e_i . For any given proper decomposition Π of P , the integral in (1) gives the sum of the volumes above $\Pi^{-1}(e_i)$ and below g_i , for $i = 1, \dots, n$. In this sense, $\mathcal{SK}_W(P)$ gives minimal volume for fixed region areas.

Related is a three-dimensional interpretation of $\mathcal{SK}_W(P)$, namely, as the lower envelope of the linear functions $f_i = \nu(x, e_i)/w_i$ on P (which now correspond to planes with slopes $\frac{1}{w_i}$ through the sides e_i). In

particular, if all weights are equal to 1 then we have $f_i = g_i$, and the volume expressed in (1) is minimum possible for *all* proper decomposition of P , without restrictions on areas. Clearly, $SK_W(P)$ is just the medial axis of P in this case.

The question arises when a given decomposition of a polygon P actually is the weighted skeleton of P , for a suitable set of weights. Based on their embedding in space, a complete characterization can be given for weighted skeletons of convex polygons. Let $L(s)$ denote the straight line containing a given line segment s .

Lemma 1 *Let Π be a (proper) decomposition of P into convex polygons Q_1, \dots, Q_n , such that Q_i shares side e_i with P . Then Π is the weighted skeleton of P for some set of weights if and only if, for each edge $b = Q_i \cap Q_j$ of Π , the condition $L(e_i) \cap L(e_j) \in L(b)$ holds.*

Proof. Let $f_i = \nu(x, e_i)/w_i$, as before. The condition $L(e_i) \cap L(e_j) \in L(b)$ implies that b is contained in the vertical projection of the intersection line of two planes $z = f_i$ and $z = f_j$, for suitable weights w_i and w_j . Thus the condition necessarily holds for any weighted skeleton, by the lower envelope picture described above. On the other hand, given Π with this condition holding for each edge b , we prove the existence of n planes $z = f_i$ whose lower envelope defines Π . This implies $\Pi = SK_W(P)$ for the resulting tuple $W = (w_1, \dots, w_n)$ of weights.

As the edges of Π define a tree inside P , at least one polygon of Π , say Q_k , is a triangle. Let side e_k be adjacent to sides e_i and e_j . Removing e_k and prolonging sides e_i and e_j gives an $(n - 1)$ -gon P' . (For $n \geq 4$, side e_k can always be chosen such that P' is bounded.) Consider the edge $b = Q_i \cap Q_j$ of Π . We have $L(e_i) \cap L(e_j) \in L(b)$, so prolonging b yields a proper decomposition Π' of P' . Assume inductively that the planes in question do exist for Π' and P' . This is obviously true in the base case $n = 3$. As Q_k has a single vertex v inside P , namely the one it shares with Q_i and Q_j , the desired plane $z = f_k$ for side e_k uniquely exists; it passes through e_k and the vertical projection of v onto the plane $z = f_i$ (or, equivalently, onto $z = f_j$). \square

Using their spatial interpretation, we get another property of weighted skeletons, which might be interesting from the physicist's (or the architect's) viewpoint.

Lemma 2 *Let SK be any weighted skeleton of a convex polygon P , with non-vanishing region areas. Then each internal edge e of SK can be associated with a positive tension $\tau(e)$ such that all internal vertices of SK are in equilibrium state.*

Proof. The lower envelope of f_1, \dots, f_n describes (together with P) the boundary of a convex polyhedron in \mathbb{R}^3 . Thus SK , being the vertical projection of this envelope to P , is a so-called Schlegel diagram; see [11]. It is well known that edge tensions as above do exist for such diagrams [9, 4]. \square

Finally, let us mention a negative result. See Figure 3. In the decomposition drawn in full lines, call it Π again, each internal vertex is incident to three angles of $\frac{2\pi}{3}$. Hence, Π is the Steiner minimal tree [12] of the vertices of the underlying polygon P , the shortest possible connection for the vertices of P . This shows that weighted skeletons, in general, do not achieve minimum total edge length over all proper decompositions with fixed polygon areas. Observe that a minimum-length proper decomposition of P (*without* area restrictions) does not always exist, namely, if the Steiner minimal tree of P 's vertices runs via sides of P . Note finally that tensions as in Lemma 2 *do* exist for the decomposition Π in Figure 3; for instance, take $\tau(e_i) = 1$ for $i = 1, \dots, 4$. Thus the existence of such tensions is not sufficient (but necessary) for a proper decomposition to be a weighted skeleton.




Figure 4: Equal-area decomposition of a disk

5 Algorithmic aspects

Let us turn to the problem of computing weighted skeletons for convex polygons under given requirements. As the easiest variant, the weighted skeleton $\mathcal{SK}_W(P)$ for a convex n -gon P and a given weight tuple W is to be constructed. From Section 4, we know an equivalent formulation of this problem: Construct the intersection C of n halfspaces in \mathbb{R}^3 , with the property that C intersects the plane containing P just in P . This is exactly the setting where the deterministic $O(n)$ -time algorithm in [1] applies. Clearly, the halfspace description of C can be derived from P and W in $O(n)$ time. A linear-time algorithm for computing $\mathcal{SK}_W(P)$ follows. From a practical point of view, the $O(n)$ -time randomized incremental algorithm in [8] may be preferable. Designed for computing the medial axis of a convex polygon, this algorithm (and its analysis) directly extends to weighted skeletons.

Computing weighted skeletons that realize fixed-share decompositions for a convex n -gon P is much harder. The main problem consists of finding a suitable weight tuple W that makes the regions of $\mathcal{SK}_W(P)$ contain the prescribed shares. (The existence of W is guaranteed by the results in Section 3.) If share is defined as the number of points from a given m -point set, then the incremental algorithm in [5] can be adapted to compute W . This algorithm inserts the given points one at a time, and adjusts the weights of P 's sides such that shares are not exceeded for any region. It runs in (roughly) $O(n^2m)$ time and optimal $O(n + m)$ space.

If share is based on a continuous measure μ on P , then W can be approximated using a gradient-descent method; see, e.g., [6]. Let us assume that the underlying density function ρ has a constant description within any given triangle $\Delta \subset P$. (For example, let μ measure area.) Recall the function $\Psi(W)$ defined in Section 3. If ρ is continuous then Ψ can be shown to be convex and smooth. The problem in question now amounts to finding the weight tuple W^* where Ψ attains its unique minimum $\Psi(W^*) = 0$. To solve this optimization problem, we utilize the gradient $\nabla\Psi$ of Ψ . $\nabla\Psi$ is given by (g_1, \dots, g_n) , where g_i is the partial derivative of Ψ in the variable w_i . That is, $g_i = (f_i)'$, where $f_i = |\mu(\text{reg}(e_i)) - A_i|$ is considered as a function of w_i . For fixed W , the functions f_i (and thus, g_i) can be derived from $\mathcal{SK}_W(P)$ in time proportional to the number of sides of $\text{reg}(e_i)$. So we get the gradient $\nabla\Psi(W)$ of Ψ at W in $O(n)$ time. Application of the iteration scheme

$$W_{k+1} = W_k + t_k \cdot \nabla\Psi(W_k)$$

for appropriate step sizes t_k guarantees convergence of W_k to the optimal solution W^* at a superlinear rate [6]. The iteration is stopped if $\Psi(W_k) < \varepsilon$, for a predefined accuracy $\varepsilon > 0$. As $\mu(\Delta)$ can be computed in $O(1)$ time for any triangle $\Delta \subset P$, we can calculate $\Psi(W_k)$ in $O(n)$ time, which gives linear time per iteration step.

As a nice property, for a fixed-*area* decomposition of a given triangle with shares A_1, A_2, A_3 , corresponding side weights can be calculated directly, by putting $w_i = \frac{A_i}{|e_i|}$ where $|e_i|$ denotes side length. Unfortunately, this property is lost for $n \geq 4$. Still, to get a good starting value for the iteration process above, it seems plausible to take the tuple $W_0 = (\frac{A_1}{|e_1|}, \dots, \frac{A_n}{|e_n|})$ if the measure μ is area.

An alternative to the gradient-descent method is to approximate the given measure μ in a discrete way: Randomly choose points drawn from the density function ρ , and compute $\mathcal{SK}_W(P)$ such that the fractions of points its regions contain are proportional to the prescribed shares. This approach also works in cases where ρ is discontinuous. The insertion algorithm mentioned before runs (nearly) linear in the number m of points, so satisfactory approximations seem achievable in reasonable time, especially when n , the size of P , is small.

If we do not insist on computing fixed-share decompositions via weighted skeletons, then more efficiency is possible; optimality properties related to weighted skeletons are then, of course, lost. We propose the divide-and-conquer algorithm below, which can be seen as an extension to general measures of the inductive method used in [10].

The input consists of a convex polygon P and a measure μ on P , along with non-negative shares A_1, \dots, A_n , to be attained by n convex polygons attached to the sides of P . Call a polygon side *active* if its assigned share is (strictly) positive.

Base: If P is a triangle, or if P has only two active sides, then construct the decomposition for P and its (at most three positive) shares for sides directly. Otherwise:

Divide: Split P with with a diagonal d that halves the set of its active sides. For the obtained sub-polygons P_1 and P_2 , calculate the sum of shares for their sides (except d). For one sub-polygon, say P_1 , this sum is at most $\mu(P_1)$. Let $A \geq 0$ be the difference between $\mu(P_1)$ and this sum.

Recur: Construct the decomposition, Π_1 , for P_1 recursively, using A as the share assigned to the side d of P_1 . Let P_d be the convex polygon (with measure A) constructed for d .

Recur: Construct the decomposition, Π_2 , for the convex polygon $P_2 \cup P_d$ recursively, taking share zero for all sides that came from P_d .

Conquer: Concatenate Π_1 and Π_2 . This yields the required fixed-share decomposition of P .

The runtime of this algorithm is determined by the base step and the divide step; the other steps can be handled in time linear in the size of the objects involved. Let us first analyse the total number of sides of all the polygons considered. This number is clearly $O(n)$ for the base steps, because the final decomposition, Π , of P , which has size $O(n)$, is a concatenation of objects constructed in base cases. As for the divide steps, any side considered is either a diagonal or a side of P , or an (inactive) side showing up in Π . The used diagonals pairwise do not cross, and P has n sides, so the number of sides *constructed* is $O(n)$ as well. Once constructed, a fixed side may be considered $O(\log n)$ times, as it can be part of the boundary of $O(\log n)$ polygons, at most one for each level of recursion. We conclude that the total size of all polygons considered in the divide steps is $O(n \log n)$.

For decomposition with respect to non-discrete measures μ , suppose for the moment that $\mu(\Delta)$ for a given triangle Δ can be computed in constant time. We show that, under this assumption, $O(s)$ time per polygon Q with s sides suffices, for both the base step and the divide step. This is obvious for the divide step, because Q can be triangulated in $O(s)$ time such that a given diagonal is included. For the base case

where Q has only two active sides, e and e' say, we use a triangulation of Q with diagonals incident to the vertex shared by e and e' . In $O(s)$ time, we single out the triangle Δ that splits Q into two sub-polygons of measures not exceeding the shares for e and e' , respectively. For Δ , an optimization problem of constant size remains, namely, where Δ has two active sides. Similar is the other possible base case, where Q already is a triangle, now with three active sides. Both optimization problems can be solved in $O(1)$ time. We conclude a total running time of $O(n \log n)$, as this is the sum of sizes of all the polygons considered.

To cover more substantial measures, let T denote the maximal time complexity for computing $\mu(\Delta)$, over all triangles $\Delta \subset P$. Along the lines above, the runtime of the algorithm becomes $O(T \cdot n \log n)$. This setting now fits scenarios where shares are to be taken from the union of $O(T)$ constant-size objects inside P , for example, triangles, disks, or curves of constant degree. Figure 4 shows how the algorithm decomposes a given disk inside P into pieces of equal area. The decomposition is drawn in full lines. The first divide step proceeds along the diagonal shown in dashed.

Assume now that the measure concerns the number of points from a given m -point set. Dividing mainly involves shuffling points from one side of a diagonal to the other. Thereby, each point is considered at most once at each level of recursion, which gives a total work of $O(m \log n)$. For analyzing the base cases, let the current polygon, Q , contain k points. If Q has only two active sides e and e' , then consider these points in angular order around the vertex $v = e \cap e'$, and use an $O(k)$ -time median algorithm to compute the angle of a line through v that splits Q into two sub-polygons containing the required shares of points. If Q is a triangle with three active sides, then split Q into three triangles based on these sides and containing the required shares, using the $O(k)$ -time splitting algorithm in [17]. No point is considered twice during all the base cases, hence they take $O(m)$ time in total. In conclusion, a runtime of $O(m \log n)$ results for the m -point set case.

6 Higher dimensions

The concept of weighted skeleton (or weighted medial axis) is not limited to two dimensions. Given a convex polytope P in \mathbb{R}^d and an assignment of positive weights to its facets, a unique convex cell complex inside P can be defined, either by the respective shrinking process, or based on weighted distances to the hyperplanes that support P . In fact, the decomposition and optimality results presented in Section 3 and Section 4, respectively, directly generalize to higher dimensions. Moreover, a weighted skeleton for a convex polytope in \mathbb{R}^d can be computed by intersecting n halfspaces in \mathbb{R}^{d+1} , one for each polytope facet. Equivalently, the convex hull of n points in \mathbb{R}^{d+1} has to be constructed. In the probably most interesting case $d = 3$, this leads to an $O(n^2)$ -time algorithm, which is optimal in the worst case as such a skeleton may consist of $\Theta(n^2)$ components.

The problem of analyzing and constructing weighted or unweighted skeletons for (well-behaved) non-convex polytopes in \mathbb{R}^d is left as a topic for further research.

References

- [1] A. Aggarwal, L.J. Guibas, J. Saxe, P.W. Shor. *A linear time algorithm for computing the Voronoi diagram of a convex polygon*. Discrete & Computational Geometry 4 (1989), 591-604.
- [2] O. Aichholzer, D. Alberts, F. Aurenhammer, B. Gärtner, *A novel type of skeleton for polygons*. Journal of Universal Computer Science 1 (1995), 752-761.

- [3] O. Aichholzer, F. Aurenhammer, B. Palop. *Quickest paths, straight skeletons, and the city Voronoi diagram*. Discrete & Computational Geometry 31 (2004), 17-35.
- [4] F. Aurenhammer. *Recognising polytopical cell complexes and constructing projection polyhedra*. J. Symbolic Computation 3 (1987), 249-255.
- [5] F. Aurenhammer, F. Hoffmann, B. Aronov. *Minkowski-type theorems and least-squares clustering*. Algorithmica 20 (1998), 61-76.
- [6] D.P. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*. Academic Press, New York, 1982.
- [7] B. Chazelle, D.P. Dobkin. *Decomposing a polygon into its convex parts*. Proc. 11th Ann. IEEE Symp. Found. Comput. Sci., 1979, 38-48.
- [8] P. Chew. *Building Voronoi diagrams for convex polygons in linear expected time*. Technical Report PCS-TR90-147, Dept. Math. Comput. Sci., Dartmouth College, Hanover, NH, U.S.A., 1986.
- [9] H. Crapo, W. Whiteley. *Plane self stresses and projected polyhedra*. Structural Topology 20 (1993), 55-78.
- [10] A. García, J. Tejel. *Dividendo una nube de puntos en regiones convexas*. Actas VI Encuentros de Geometría Computational, Barcelona, Spain, 1995, 169-174.
- [11] B. Grünbaum. *Convex Polytopes*. Wiley Interscience, London, 1967.
- [12] F.K. Hwang, D.S. Richards, P. Winter. *The Steiner tree problem*. Annals of Discrete Mathematics 53, Elsevier Science Publishers, Netherlands, 1992.
- [13] M. Keil. *Polygon decomposition*. In: Handbook of Computational Geometry, J.-R. Sack, J. Urrutia (eds), Elsevier, 2000, 491-518.
- [14] M. Keil, J. Snoeyink. *On the time bound for convex decomposition of simple polygons*. Proc. 10th Canadian Conf. Computational Geometry, 1998, 54-55.
- [15] D. Kirkpatrick. *Efficient computation of continuous skeletons*. Proc. 20th Ann. IEEE Found. Comput. Sci., 1979, 18-27.
- [16] J.-M. Lien, N.M. Amato. *Approximate convex decomposition of polygons*. Computational Geometry: Theory and Applications 35 (2006), 100-123.
- [17] K. Numata, T. Tokuyama. *Splitting a configuration in a simplex*. Algorithmica 9 (1993), 649-668.