Exact and approximation algorithms for computing optimal α-fat decompositions

Mirela Damian-Iordache*

Abstract

The minimum α-fat decomposition problem is the problem of decomposing a simple polygon into the fewest subpolygons, each with aspect ratio at most α. This paper considers the version of this problem that disallows Steiner points. The main result is a polynomial time algorithm to solve this problem. For any input polygon, the algorithm produces an optimal α-fat decomposition if there is one, and reports fail otherwise. A faster approximation method is then used to produce, for any $\varepsilon > 0$, an optimal $(\alpha + \varepsilon)$-fat decomposition.

1 Introduction

The *diameter* of a simple polygon is the diameter of a smallest enclosing circle. The *width* of a simple polygon is the diameter of a largest enclosed circle. The *aspect ratio* of a simple polygon is the ratio of its diameter to its width. Polygons with a small aspect ratio have been called *fat*. In particular, for any positive real number α, a simple polygon is said to be *α-fat* if its aspect ratio does not exceed α. Other related definitions of fatness are also possible [10].

In this paper we consider the *α-fat decomposition* problem: given a polygon P and positive real α, decompose P into the fewest α-fat subpolygons. For any real α, not all polygons have an α-fat decomposition if Steiner points are disallowed. We say that a polygon P is *α-fat decomposable* whenever P has an α-fat decomposition. In this paper we focus on the version of the α-fat decomposition problem that disallows Steiner points. In such a version of the problem, we need to choose a minimum subset of the diagonals of the polygon that induces an α-fat decomposition.

The α-fat decomposition problem is a classic example of the type of problems considered in the polygon decomposition literature. For example, Keil et al. [6] show that a decomposition of a simple polygon with n vertices into a minimum number of convex polygons, can be computed in $O(n + r^2 \min\{r^2, n\})$ time. Here r denotes the number of *reflex* vertices, that is, vertices whose interior angle is greater than π. Keil [4] also develops polynomial-time algorithm for decomposing a simple polygon into fewest spiral polygons, star-shaped polygons and monotone polygons. These algorithms do not use any Steiner points. Allowing Steiner points, Chazelle and Dobkin [2] give an intricate algorithm to decompose a simple polygon into the minimum number of convex components. For a summary of results on polygon decomposition, refer to the survey papers by Keil [5] and Bern [1].

The α-fat decomposition problem finds motivation in collision detection. Suri et al. [9] show that the bounding box heuristic in collision detection performs well if the colliding objects are fat and have approximately the same size. Several other papers report reductions in combinatorial complexity for a variety of algorithms that process fat objects [7, 9, 10, 11]. An algorithm that solves the α-fat decomposition problem can be used as a crucial preprocessing step in algorithms for problems in which fatness is useful.

In this paper we present two results. The main result is a polynomial time algorithm to solve the α-fat decomposition problem. A similar result of slightly greater complexity has been reported in [3], but no algorithmic details have been presented. As far as we know, this is the first polynomial time algorithm for the problem. To reduce the time complexity of the exact algorithm we devise an algorithm that produces, for any $\varepsilon > 0$, an $(\alpha + \varepsilon)$-fat decomposition that has as few polygons as an optimal α-fat decomposition of P.

2 Computing optimal α-fat decompositions

Let P be a simple polygon in the plane with n vertices, r of which are reflex. Let the vertices of P be labeled $1, 2, \cdots n, n+1$ in counterclockwise order, where 1 and $n+1$ are identical vertices. Before we show how to compute an optimal α-fat decomposition of P, we introduce some terminology. A *diagonal* of P is a line segment that connects two vertices of P and lies entirely inside P. For ease of presentation, we refer to edges of P as diagonals also. The visibility graph of P is a graph whose nodes are vertices of P and whose edges are diagonals of P. We denote by m be the number of edges in the visibility graph of P. For any diagonal

*Department of Computer Science, Villanova University, Villanova, PA 19085, U.S.A. e-mail: mirela.damian@villanova.edu

(i, j) of P, with $i < j$, we denote by P_{ij} the subpolygon of P defined by vertices $i, i+1, \cdots j$ (as an example, the shaded polygon in Figure 2(a) is P_{ij}).

To solve the minimum α-fat decomposition problem we use dynamic programming. We consider diagonals (i, j), with $i < j$, in increasing order of $j - i$, and compute an optimal α-fat decomposition O_{ij} of each P_{ij}. P_{1n} is the entire polygon and therefore the α-fat decomposition O_{1n} of P_{1n} is a solution to the minimum α-fat decomposition problem.

2.1 Properties of optimal α-fat decompositions

Focus on an optimal α-fat decomposition O_{ij} of P_{ij} (for the purpose of the discussion in this section, assume that one exists). Let Q be the polygon in O_{ij} adjacent to (i, j) and let C and I be the smallest circumscribed circle and the largest inscribed circle respectively, as illustrated in Figure 1. Then clearly the boundary of Q

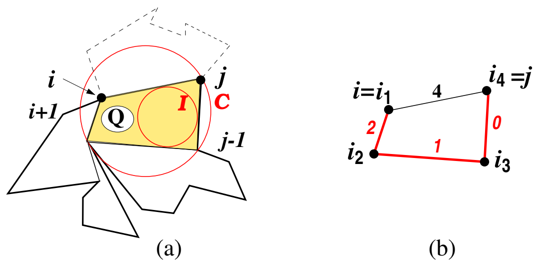

Figure 1: Definitions (a) Properties of O_{ij} (b) The weighted graph $G_{ij}^{(C,I)}$

lies inside C and outside I. Moreover, the ratio between the diameter of C and the diameter of I does not exceed α, since Q is an α-fat polygon. Let T be a homogeneous set containing all reflex vertices of P and all edges in the visibility graph of P. Note that C is tangent to at least two vertices of P and is well defined by at most three such non reflex vertices. Similarly, I is well defined by at most three elements in T.

Let $G_{ij}^{(C,I)}$ be the subgraph of the visibility graph of P induced by diagonals (k, l) of P_{ij} such that (k, l) lies inside C and P_{kl} lies entirely outside I. For example, Figure 1(b) shows the graph $G_{ij}^{(C,I)}$ corresponding to P_{ij}, C and I shown in Figure 1(a). Figure 2(b) shows the graph $G_{ij}^{(C,I)}$ corresponding to P_{ij}, C and I shown in Figure 2(a). To make another important property of O_{ij} apparent, we associate to each edge (k, l) in $G_{ij}^{(C,I)}$ a weight w_{kl} as follows. If P_{kl} is α-fat decomposable, then w_{kl} is the size of an optimal α-fat decomposition O_{kl} of P_{kl}, that is, the number of polygons in O_{kl}. Otherwise, $w_{kl} = \infty$. By convention, $w_{kl} = 0$ if $l = k + 1$. An example of a weighted graph $G_{ij}^{(C,I)}$ is given in Figure 1(b). Also Figure 2(b) shows the graph $G_{ij}^{(C,I)}$ prior to the computation of the weight of (i, j). The reason for which the edge (i, j) in Figure 2(b) has an infinite weight will become clear in the description of our algorithm in section 2.2.

A *path* in a graph is a sequence of vertices, with successive vertices connected by edges. Define the *weight* of a path as the sum of the weights of its constituent edges. It is important to note that the path in $G_{ij}^{(C,I)}$ induced by the vertices of Q is a path of minimum weight from i to j. Figure 2(b) shows with thick lines a path of minimum weight from i to j (weight $2 + 1 + 0 = 3$).

We will use these properties to determine the subpolygon Q in a minimum α-fat decomposition of P_{ij} that contains (i, j). The dynamic programming approach ensures that the optimal α-fat decomposition of the polygonal pieces of P_{ij} obtained by removing Q are available at this point. Then the union of these optimal α-fat decompositions and Q will give us an optimal α-fat decomposition of P_{ij}.

2.2 Computing O_{ij}

We use the properties stated in section 2.1 to compute an optimal α-fat decomposition O_{ij} of P_{ij} as follows. For $j = i + 1$, O_{ij} is the empty set. For $j > i + 1$, first check if P_{ij} is α-fat. If so, O_{ij} is simply P_{ij}. Otherwise consider all circles C passing through three non-reflex vertices of P or whose diameter connects two non-reflex vertices of P, in arbitrary order. There are $O((n - r)^3)$ such circles. For each such circle C of diameter D, consider all circles I of diameter $\frac{D}{\alpha}$ tangent to two elements in T. There are $O(m^2)$ such circles. Note

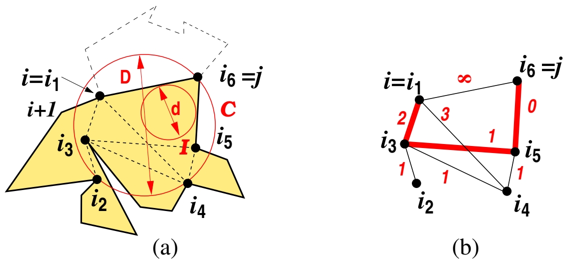

Figure 2: Computing O_{ij} (a) A pair (C, I): C tangent to i_2, i_4, i_6; I tangent to (i_1, i_6), (i_5, i_6); $d = D/\alpha$ (b) $\{i_1, i_3, i_5, i_6\}$ is a path of minimum weight from i to j in $G_{ij}^{(C,I)}$

that since the diameter of I is fixed, two elements in T suffice to completely define I, as illustrated in Figure 2(a). We process all pairs of circles (C, I) for which the diagonal (i, j) lies entirely inside C and outside I. For each such pair (C, I), let $G_{ij}^{(C,I)}$ be a weighted graph as defined in section 2.1 (refer to Figure 2(b)). At this stage of the dynamic programming algorithm we have optimal decompositions O_{kl} for all polygons P_{kl} that are smaller than P_{ij}. This implies that all the weights in $G_{ij}^{(C,I)}$ are well defined, except for w_{ij}, which we set to ∞.

For each pair of circles (C, I), compute a path $p_{ij}^{(C,I)}$ of minimum weight from i to j in $G_{ij}^{(C,I)}$. Since (i, j) is an edge in the graph, such a path always exists, however, its weight may be ∞. Pick the path p_{ij} of smallest weight taken over all pairs of circles (C, I), $p_{ij} = \min_{(C,I)} p_{ij}^{(C,I)}$. It is possible that the weight of p_{ij} is ∞. In such a case, O_{ij} is set to $\perp$. Thus, $\perp$ represents the absence of an α-fat decomposition. Otherwise, note that the vertices of p_{ij} induce an α-fat polygon Q. In this case we set $O_{ij} = \cup_{(k,l)\subset E(p_{ij})} O_{kl} \cup Q$, where $E(p_{ij})$ denotes the set of edges along the path p_{ij}.

The only piece missing in the computation of O_{ij} is an algorithm to compute $p_{ij}^{(C,I)}$ for $j > i + 1$. This can be done in linear time as follows. Pick a point ℓ adjacent to both i and j, with $i < \ell < j$, that minimizes the weight of the path $p_{i\ell}^{(C,I)} \cup p_{\ell j}^{(C,I)}$ (note that both $p_{i\ell}^{(C,I)}$ and $p_{\ell j}^{(C,I)}$ have been computed in previous stages of the dynamic programming algorithm). If we can find ℓ, then clearly $p_{ij}^{(C,I)} = p_{i\ell}^{(C,I)} \cup p_{\ell j}^{(C,I)}$. Otherwise the only path from i to j is $\{i, j\}$ and has infinite weight.

Theorem 1 *O_{ij} is a minimum size α-fat decomposition of P_{ij}.*

Proof: The proof is by induction on $(j - i)$. The base case is when $j - i = 1$. In this case P_{ij} is a degenerate polygon with only one edge and therefore O_{ij} is the empty set.

The inductive hypothesis is that for some natural number s and for all i, j, such that $1 \le j - i \le s$, O_{ij} is a minimum size α-fat decomposition of P_{ij}. To prove the inductive step we consider the α-fat decomposition O_{ij} for some i, j such that $j - i = s + 1$.

If P_{ij} has no α-fat decomposition, then clearly $O_{ij} = \perp$. So suppose that P_{ij} has an α-fat decomposition. If P_{ij} is an α-fat polygon, then the decomposition is the polygon itself. Otherwise let Q^* be the polygon in a *minimum* α-fat decomposition O_{ij}^* of P_{ij} adjacent to (i, j).

Since Q^* is α-fat, there is a pair of circles (C, I) considered by our algorithm such that the boundary of Q^* lies inside C and outside I. This implies that $G_{ij}^{(C,I)}$ contains all vertices of Q^*. Moreover, the vertices of Q^* induce a path p_{ij}^* in $G_{ij}^{(C,I)}$. The path p_{ij}^* is a candidate for the minimum weight path p_{ij} computed by our algorithm and therefore the size of $O_{ij} = \cup_{(k,l)\subset E(p_{ij})} O_{kl} \cup Q$ is no greater than the size of $O_{ij}^* = \cup_{(k,l)\subset E(p_{ij}^*)} O_{kl} \cup Q^*$, which is a minimum size α-fat decomposition. It follows that O_{ij} is a minimum size α-fat decomposition. □

Theorem 2 *O_{1n} is a minimum α-fat decomposition of P that can be computed in time $O(m^3(n - r)^3 n)$.*

Proof: The algorithm computes α-fat decompositions for $O(m)$ polygons P_{ij}. For each polygon P_{ij}, there are $O(n - r)^3$ circles C and $O(m^2)$ circles I to consider. For a given pair of circles (C, I), computing $p_{ij}^{(C,I)}$ takes $O(n)$ time. Thus the running time of the algorithm is $O(m) \times O(n - r)^3 \times O(m^2) \times O(n) = O(m^3(n - r)^3 n)$. This however requires precomputing the visibility graph of P, which can be done in $O(m + n)$ time and precomputing the set of polygons P_{ij} which are α-fat. To determine whether P_{ij} is α-fat, we need to compute the diameter and the width of P_{ij}. The diameter can be easily computed in $O(n^2)$ time. To determine the width of P_{ij}, compute in $O(n^2)$ time the medial axis transform of P_{ij}, which is the locus of centers of all maximal inscribed circles [8]. Thus the set of all α-fat polygons P_{ij} can be precomputed in $O(mn^2)$ time. □

3 Computing optimal $(\alpha + \varepsilon)$-fat decompositions

The algorithm described in the previous section proves that the α-fat decomposition problem is solvable in polynomial time, however, it is not yet practical. The bottleneck of the algorithm is the large number of pairs of circles (C, I) that we need to consider. Our goal in this section is to derive an algorithm that computes $(\alpha + \varepsilon)$-fat decompositions considering a reduced number of pairs (C, I), where ε is a user specified tolerance value. We also impose the additional restriction that the subpolygons in the decomposition have diameter no smaller than a given lower bound L. This prevents the algorithm from generating polygons that are too small.

We start by placing an infinite grid of $\beta \times \beta$ cells on the plane. As we will shortly see, β is a positive real that is chosen based on the aspect ratio α, the user tolerance ε and the lower bound L. Next we consider diagonals (i, j) and compute an optimal $(\alpha + \varepsilon)$-fat decomposition D_{ij} of each P_{ij}. The algorithm that computes D_{ij} is very similar to the algorithm that computes O_{ij} (see Section 2.2), and we outline only the changes here. In computing D_{ij}, we consider only those circles C of diameter D no less than L. For each such circle C, we consider all circles I of diameter $(\frac{D}{\alpha} - 2\sqrt{2}\beta)$, centered at one of the grid points. The rest of the algorithm is the same.

Note that for each pair of circles (C, I), any polygon whose boundary lies outside C and inside I has aspect ratio $AR = \frac{D}{\frac{D}{\alpha} - 2\sqrt{2}\beta}$. If we pick $\beta = \frac{L\varepsilon}{2\sqrt{2}\alpha(\alpha+\varepsilon)}$, then the aspect ratio becomes $AR = \frac{D\alpha(\alpha+\varepsilon)}{D\alpha+(D-L)\varepsilon} \le \alpha+\varepsilon$ and therefore the subpolygon in D_{ij} adjacent to (i, j) has aspect ratio no greater than $\alpha+\varepsilon$. Based on this idea, it is easy to show that D_{ij} is an $(\alpha+\varepsilon)$-fat decomposition.

Next we show that D_{ij} is optimal in the sense that it is no larger than an optimal α-fat decomposition O_{ij} of P_{ij} in which all polygons have diameter bounded below by L. Focus on the subpolygon Q in O_{ij} that is adjacent to (i,j). To show that D_{ij} is optimal, it suffices to show that our algorithm considers a pair of circles (C_1, I_1) such that the boundary of Q lies entirely inside C_1 and outside I_1. This implies that Q is a candidate for the polygon in D_{ij} adjacent to (i,j) and therefore D_{ij} is no larger than O_{ij}.

Let C be the smallest circle enclosing Q and I the largest circle enclosed in Q. Our algorithm processes the circle $C_1 = C$, since C is either tangent to three non-reflex vertices of P or its diameter connects two non-reflex vertices of P. Also the diameter of C is no less than L. Denote by D and d the diameter of C and I respectively. Then clearly $\frac{D}{d} \le \alpha$, and simple calculus shows that $d \ge \frac{D}{\alpha} > 2\sqrt{2}\beta$. This implies that there is a $\beta \times \beta$ grid square S entirely contained inside C that contains the center of I. Any corner of S is at distance at most $\sqrt{2}\beta$ from the center of I. This implies that the circle I_1 of diameter $d_1 = \frac{D}{\alpha} - 2\sqrt{2}\beta \le d - 2\sqrt{2}\beta$ centered at a grid point adjacent to S is wholly contained inside I. This shows that Q lies entirely outside I_1. Since the algorithm processes all circles centered at a grid point of diameter d_1, it also processes the pair (C_1, I_1).

Theorem 3 *Let P be a polygon of diameter D and let* OPT *be the size of an optimal α-fat decomposition of P in which all polygons have diameter no less than L. For any real $\varepsilon > 0$, there is algorithm that computes an $(\alpha+\varepsilon)$-fat decomposition of P of size no greater than* $|\mathrm{OPT}|$ *in time $O(mn(n-r)^3(\frac{2\sqrt{2}D\alpha(\alpha\varepsilon)}{L\varepsilon})^2)$.*

Proof: The dynamic programming algorithm discussed in this section considers $O(m)$ polygons P_{ij}, $O((n-r)^3)$ circles C and $O((\frac{D^*}{\beta})^2)$ circles I, where D^* is the diameter of P. Also the algorithm takes $O(n)$ time to compute a shortest path in a weighted graph and therefore the running time of the algorithm is $O(mn(n-r)^3(\frac{D^*}{\beta})^2)$, where $\beta = \frac{L\varepsilon}{2\sqrt{2}\alpha(\alpha+\varepsilon)}$ as defined above. □

4 Conclusions

We have presented the first polynomial time algorithms for the α-fat decomposition problem. To obtain faster algorithms, we are investigating more sophisticated techniques to prune the dynamic programming search space.

References

[1] M. Bern. Triangulations. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 413–428. CRC Press, Boca Raton, 1997.

[2] B. Chazelle and D. P. Dobkin. Optimal convex decompositions. *Computational Geometry, G. Toussaint*, 1985.

[3] M. Damian and S. Pemmaraju. Computing optimal α-fat and α-small decompositions. *Proc. of the 12th Ann. ACM-SIAM Symp. on Discr. Algh.*, pages 338–340, 2001.

[4] J. M. Keil. Decomposing a polygon into simpler components. *SIAM Journal on Computing*, 14:799–817, 1985.

[5] J. M. Keil. Polygon decomposition. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*. Elsevier Science B.V. North-Holland, 1999.

[6] J. M. Keil and J. Snoeyink. On the time bound for convex decomposition of simple polygons. *Proc. of the 10th Canad. Conf. Comp. Geom.*, pages 338–340, 1998.

[7] Mark H. Overmars and A. Frank van der Stappen. Range searching and point location among fat objects. *Journal of Algorithms*, 21(3):629–656, 1996.

[8] F. Preparata. The medial axis of a simple polygon. *Proc. of the 6th Symp. Math. Foundations of Comp. Sci.*, pages 443–450, 1977.

[9] S. Suri, P. M. Hubbard, and J. F. Hughes. Collision detection in aspect and scale bounded polyhedra. *Proc. of the 9th Ann. Symp. on Discr. Alg.*, pages 127–136, 1998.

[10] A. Frank van der Stappen and Mark H. Overmars. Motion planning amidst fat obstacles. *Proc. of the 10th ACM Symposium on Computational Geometry*, pages 31–40, 1994.

[11] Y. Zhou and S. Suri. Collision detection using bounding boxes: Convexity helps. *Proc. of the 8th Annual European Symposium on Algorithms*, 2000.