Optimal Convex Partitions of Point Sets with Few Inner Points

Andreas Spillner*

Abstract

We present a fixed-parameter algorithm for the Minimum Convex Partition and the Minimum Weight Convex Partition problem. On a set P of n points the algorithm runs in $O(2^k k^3 n^3 + n \log n)$ time. The parameter k is the number of points in P lying in the interior of the convex hull of P.

1 Introduction

Let P denote a set of n points in the plane. By $CH(P)$ we denote the convex hull of P. To ease argumentation we will assume that no three points in P are collinear and no two points in P have the same x-coordinate.

A *convex partition* of P is a set E of straight line segments with endpoints in P, called edges, such that edges do not cross each other and partition $CH(P)$ into a set $\mathcal{R}(E)$ of empty convex regions. A region is *empty* if it does not contain a point of P in its interior. An example of a convex partition is given in Figure 1. Note that the edges of $CH(P)$ are contained in every convex partition of P.

The *Minimum Convex Partition* problem (MCP) is to compute a convex partition E of P such that the number of regions in $\mathcal{R}(E)$ is minimum. Lingas has shown that the related problem of partitioning a polygon with n vertices by diagonals into a minimum number of convex pieces is NP-hard for polygons with holes [10]. For polygons without holes Keil and Snoeyink give an $O(n^3)$ time algorithm [8]. Fevens et al. have shown that MCP can be solved in $O(n^{3h+3})$ time if the points in P lie on h nested convex hulls [3].

The *Minimum Weight Convex Partition* problem (MWCP) is to compute a convex partition E of P such that the total length of the edges in E is minimum. Again we have the related problem of partitioning a polygon with n vertices into convex pieces such that the total length of the diagonals used for the partition is minimum. This related problem is NP-hard for polygons with holes, as shown by Keil [7]. But it can be solved in $O(n^4)$ time for polygons without holes as Keil and Snoeyink note in [8]. There are also polynomial time constant-factor approximation algorithms for MWCP by Plaisted and Hong [11] and Levcopoulos and Krznaric [9].

*Institute of Computer Science, Friedrich-Schiller-University Jena, Germany, spillner@minet.uni-jena.de

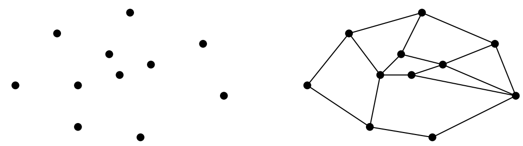

Figure 1: A point set P and a convex partition of P.

However, in general MCP and MWCP seem to be hard problems and we are not aware of any algorithms solving these problems for every point set P in polynomial time. One possible way to deal with this situation is to consider *fixed-parameter algorithms*. Such an algorithm solves the problem under consideration in $O(f(k)p(n))$ time. The function f is only allowed to depend on the so called parameter k. This should be a quantity associated with the input such that k can be small even for large input size n. The function p must be a polynomial of constant degree in the size of the input. A problem that admits a fixed-parameter algorithm for parameter k is called *fixed-parameter tractable* with respect to parameter k. For more details about the theory of parameterized complexity we refer to the book of Downey and Fellows [2]. Note that the algorithm of Fevens et al. in [3] is not a fixed-parameter algorithm with respect to the number h of nested convex hulls, since the degree of the polynomial in the bound on the running time is not constant but depends on h.

Grantson and Levcopoulos present a fixed-parameter algorithm for MCP running in $O(k^{6k-5} 2^{16k} n)$ time [5]. For MWCP Grantson gives a fixed-parameter algorithm running in $O(k^{4k-8} 2^{13k} n^3)$ time [4]. The parameter k in both algorithms is the number of *inner points* in P, which are the points in P lying in the interior of $CH(P)$. Recently the Traveling Salesman Problem (TSP) and the Minimum Weight Triangulation problem (MWT) have been shown to be fixed-parameter tractable with respect to this parameter too. For TSP Deineko et al. give an $O(2^k k^2 n)$ time algorithm [1] and for MWT Hoffmann and Okamoto present an $O(6^k n^5 \log n)$ time algorithm [6].

In this paper we show that the approach of Hoffmann and Okamoto for MWT can be adapted to yield a new fixed-parameter algorithm for MCP and MWCP running in $O(2^k k^3 n^3 + n \log n)$ time. Keeping in mind that fixed-parameter algorithms are proposed as a possible way to cope with hard problems, the aim is at practical and efficient algorithms. So our work could be seen as

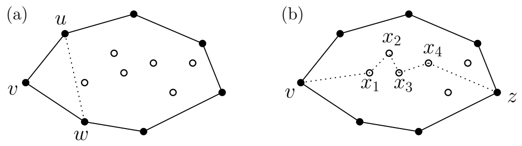

Figure 2: Either we can cut off an empty triangle or we have an x-monotone path starting at v.

just one further step to reach this goal for MCP and MWCP.

2 Outline of the algorithm

Since our algorithm does not exploit any geometric properties that depend on the objective function, it will work for both problems under consideration. So let E_{opt} be an optimal convex partition of P, that is a solution for MCP or MWCP. Let v be the point in P with minimum x-coordinate. Then v is a vertex of $CH(P)$. We will call the vertices of $CH(P)$ *outer points* for short. The *neighbors* of an outer point are the two other outer points to which it is connected by an edge of $CH(P)$. Let u and w denote the neighbors of v. There are two possible cases.

Case 1: The two edges between v and its neighbors u and w are the only edges incident to v in E_{opt}. Then the triangle with vertices u, v and w must be empty and it suffices to find an optimal convex partition of $P \setminus \{v\}$. This situation is indicated in Figure 2(a).

Case 2: There are more than two edges incident to v in E_{opt}. Then v is connected in E_{opt} to an outer point z by an *x-monotone path* Π, that is a sequence of points $v = x_0, x_1, x_2, \ldots, x_l, x_{l+1} = z$ from P such that x_i has smaller x-coordinate than x_{i+1} and there is an edge in E_{opt} between x_i and x_{i+1}. The points $x_1, \ldots, x_l$ are inner points. Such a path Π exists since the regions in $\mathcal{R}(E_{opt})$ are convex polygons and in a convex polygon all interior angles are smaller than 180^o. The path Π divides P into two independent subproblems as indicated in Figure 2(b). Note that the concept of x-monotone paths forms the basis of the fixed-parameter algorithm for MWT by Hoffmann and Okamoto [6].

One can imagine our algorithm building possible convex partitions edge by edge. Thus in an intermediate stage we have a set E of edges that have already been selected to be in the convex partition and we check possible ways to add a further edge. The subproblems that arise are the connected components of the set of points that belong to $CH(P)$ but to none of the edges in E. Such a connected component can be described by the edges of E that form its boundary. As mentioned in Section 1 we can suppose that E contains all the edges of $CH(P)$. So in the description of subproblems we focus on those edges in E which are not edges of $CH(P)$.

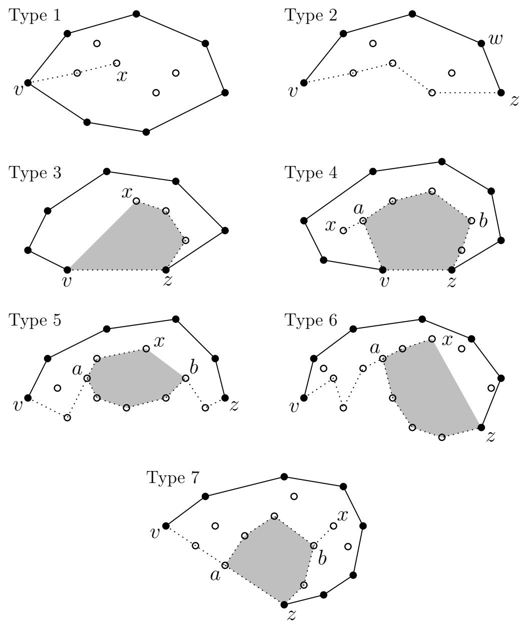

Figure 3: Examples for the types of subproblems.

In Figure 3 we give an example for each type of subproblem we will encounter. Inner points are indicated by empty circles. Shading indicates that a region is empty. We use solid lines for edges of $CH(P)$ and dotted lines for those edges which are not edges of $CH(P)$. Note that these dotted edges can be artificial in the sense that they are not contained in the optimal convex partition of P to be computed. But this is not a problem as long as we only want to know which edges must still be added to partition the subproblem optimally. It is only when combining solutions of subproblems that we must take care not to include any artificial edges. But this is not hard to manage.

We will say that a point q is between points p and r, if q has larger x-coordinate than p and smaller x-coordinate than r. In the same situation we will also say that p is to the left of q and r is to the right of q.

From the case analysis above we derive the first two types of subproblems. The other types arise when we process subproblems. An important thing is that we process subproblems in a way that only new subproblems are created that have one of the seven listed types too, at least after reflection on the x-/y-axis or rotation by 180^o. Let $\mathcal{R}_{opt}$ denote the set of convex regions in an optimal solution of the subproblem under consideration.

Type 1: An x-monotone path Π is growing through $CH(P)$ starting at an outer point v. The path Π ends at an inner point x and it suffices to consider all possible ways to extend Π by one edge beyond x. If we extend Π to another inner point we obtain again a subproblem of Type 1. If we reach an outer point we have created two subproblems of Type 2.

Type 2: Such a subproblem is defined by an x-monotone path Π between two outer points v and z. In processing the subproblem we will concentrate on point z. Similarly we could discuss the situation around point v but it suffices to do it only for one of the endpoints of Π. Note that Π can degenerate to a single edge e with endpoints v and z. Furthermore in such a situation edge e can stem from cutting off an empty triangle and then e would be an artificial edge as mentioned above.

Let us first discuss how we treat a subproblem of Type 2 where indeed Π consists only of a single edge e. Let w denote the neighbor of z present in the subproblem. There is a unique region R in $\mathcal{R}_{opt}$ such that e is an edge of R. If w is a vertex of R then the triangle with vertices v, z and w is empty and we consider the subproblem of Type 2 obtained by cutting off this triangle. If w is not a vertex of R there is an inner point x such that the straight line segment between z and x is an edge of R. So we consider all the subproblems where z is connected by an edge to an inner point x and the triangle with vertices v, z and x is empty. This leads us to subproblems of Type 3.

Now we can assume that Π consists of at least two edges. If all these egdes belong to the boundary of a single region in $\mathcal{R}_{opt}$ then Π and the straight line segment with endpoints v and z form the boundary of an empty convex polygon. We obtain a subproblem of Type 2 by substituting a single (artificial) edge between v and z for Π. Otherwise there are essentially two possible cases. The first case is that there are two inner points a and b on Π such that a is to the left of b, the part of Π between a and b belongs to the boundary of a single region R in $\mathcal{R}_{opt}$, a is the leftmost point of R and b is the rightmost point of R. This leads us to subproblems of Type 5. The second case is that there is an inner point a on Π such that the part of Π between a and z belongs to the boundary of a single region R in $\mathcal{R}_{opt}$ and a is the leftmost point of R. This way we get subproblems of Type 6.

Type 3: We grow the boundary of a region R in $\mathcal{R}_{opt}$ edge by edge. We extend this boundary at the point x by adding a new edge to a point x'. We consider each suitable point x', in particular the triangle with vertices v, x and x' must be empty. If x' is an inner point we obtain again a subproblem of Type 3. So let's assume x' is an outer point. Depending on the shape of the part Ψ of the boundary of R from z to x' we are led to different types of subproblems. If all inner points on Ψ are between z and x' we obtain a subproblem of Type 2. If there is an inner point on Ψ to the right of z but no inner point on Ψ to the left of x' or there is an inner point on Ψ to the left of x' but no inner point on Ψ to the right of z then we obtain a subproblem of Type 7. If there is an inner point on Ψ to the right of z and an inner point on Ψ to the left of x' then we obtain a subproblem of Type 4. Finally if $x' \neq v$ we obtain a second subproblem of Type 2 by adding an (artificial) edge between the outer points x' and v.

Type 4: We have grown a part Ψ of the boundary of a region in $\mathcal{R}_{opt}$ and have reached an outer point v. The leftmost point a and the rightmost point b on Ψ are inner points. From a we grow an x-monotone path Π which we extend edge by edge to the left. We consider each point x' to the left of x and connect the current endpoint x of Π to x' by an edge. If x' is an inner point we obtain again a subproblem of Type 4. If x' is an outer point we obtain two subproblems, one of Type 2 and another one of Type 7.

Type 5: We grow the boundary of a region R in $\mathcal{R}_{opt}$ edge by edge. Point a is the leftmost point of R and point b is the rightmost point of R. We extend the boundary of R at x by adding an edge to a point x'. We consider each suitable point x' that is between x and b, in particular the triangle with vertices b, x and x' must be empty. If x' is an inner point we obtain a subproblem of Type 5. If x' is an outer point we obtain two subproblems: a subproblem of Type 2 and a subproblem of Type 6. The subproblem of Type 6 is created by adding an (artificial) edge between x' and b. Finally we consider the subproblem of Type 2 obtained by adding an edge connecting x and b.

Type 6: Here again we grow the boundary of a region R in $\mathcal{R}_{opt}$ edge by edge. We must extend this boundary at the point x by adding a new edge to a point x'. We consider each suitable point x', in particular the triangle with vertices z, x and x' must be empty. If x' is an inner point we obtain again a subproblem of Type 6. So suppose x' is an outer point. If x' is the rightmost point among the points on the part of the boundary of R from a to x' then we obtain a subproblem of Type 2. Otherwise we obtain a subproblem of Type 7. Furthermore if $x' \neq z$ we create a second subproblem of Type 2 by adding an (artificial) edge connecting the outer points z and x'.

Type 7: We have grown a part Ψ of the boundary of a region in $\mathcal{R}_{opt}$ and have reached an outer point z. The rightmost point b of Ψ is an inner point. From b we grow an x-monotone path Π. We consider each point x' to the right of the current endpoint x of Π and add an edge connecting x and x'. If x' is an inner point we obtain a subproblem of Type 7. Otherwise we obtain two subproblems of Type 2.

Now the collection of types of subproblems is complete and we have outlined how to process each type of subproblem. Note that when processing a subproblem we indeed create only smaller subproblems that have one of the seven listed types too.

3 Analysis

We want to employ the dynamic programming technique to avoid re-computation of optimal convex partitions for subproblems. Then we could assume that we can look up optimal convex partitions for the subproblems we have already computed in a table in constant time. Of course, such an approach only pays off if the number of distinct subproblems can be bounded appropriately. So first we analyze the number of subproblems for each of the seven types listed in Section 2.

A key observation of Hoffmann and Okamoto in [6] is that an x-monotone path is uniquely defined by the points of P on the path. Now our types of subproblems are not only defined by x-monotone paths. However, we can obtain bounds on their number in a similar way. We will discuss this only for subproblems of Type 5. The same bound can be shown to hold for the other types as well by similar arguments.

So, how many subproblems of Type 5 are there? There are $O(n^2)$ possible choices for the outer points v and z. When we have any subset S of the inner points in P we can number them $x_1, x_2, \ldots, x_l$ according to increasing x-coordinates. There are $O(l^2)$ possible ways to choose $x_i = a$ and $x_j = b$ such that $i < j$. A choice for a and b can only lead to a subproblem of Type 5 if the inner points $a = x_i, x_{i+1}, \ldots, x_{j-1}, x_j = b$ are the vertices of an empty convex polygon. But then this empty convex polygon and therefore the boundary of it is uniquely defined. By the observation of Hoffmann and Okamoto the x-monotone path from v to a is uniquely defined by $v, x_1, \ldots, x_i$ and the x-monotone path from b to z is uniquely defined by $x_j, \ldots, x_l, z$. Since $x_1, \ldots, x_l$ are inner points we have $l \leq k$. To summarize: there are only $O(n^2 2^k k^2)$ subproblems of Type 5.

Finally we have to give a bound on the time needed to compute an optimal convex partition for a subproblem. Since we essentially grow x-monotone paths or boundaries of regions in the optimal convex partition edge by edge and we use dynamic programming, $O(n)$ time suffices in most cases. Only when we have to deal with a subproblem of Type 2 we might consider $O(k^2)$ subproblems of Type 5. Thus we get a time bound of $O(n + k^2)$. However this is only valid under the assumption that we use $O(1)$ time to create a subproblem. To achieve this we preprocess P in $O(kn^3)$ time using $O(n^3)$ space such that for each triangle with vertices in P we can look up in $O(1)$ time whether it is empty.

Now we can state our result. The $n \log n$-term in the bound on the running time is for checking whether P contains any inner points at all.

Theorem 1 *Given a set P of n points in the plane, we can compute solutions to MCP and MWCP for P in $O(2^k k^3 n^3 + n \log n)$ time using $O(2^k k^2 n^2 + n^3)$ space. The parameter k is the number of inner points in P.*

4 Conclusion

It is not hard to adapt our fixed-parameter algorithm such that it will work in situations where in addition to the point set P a set of edges E_{in} is given as part of the input and an optimal convex partition of P among those containing E_{in} must be computed. With a suitable preprocessing of E_{in} we can achieve the same time bound as in Theorem 1.

It seems unlikely that we can improve the 2^k-term in the bound on the running time significantly only relying on a concept based on x-monotone paths.

From the viewpoint of parameterized complexity the question whether the problems under consideration admit a reduction to a problem kernel [2] suggests itself. Since we can always find a convex partition E of P such that $\mathcal{R}(E)$ contains only $O(k)$ regions, MCP seems to be a good candidate for such a reduction.

Last but not least we mention the important issue of implementation and experimentation. Do our purely theoretical results really lead to practical and efficient algorithms?

References

[1] V. G. Deineko, M. Hoffmann, Y. Okamoto, and G. J. Woeginger. The traveling salesman problem with few inner points. *Operations Research Letters*. to appear.

[2] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer, 1999.

[3] T. Fevens, H. Meijer, and D. Rappaport. Minimum convex partition of a constrained point set. *Discrete Appl. Math.*, 109:95–107, 2001.

[4] M. Grantson. Fixed-parameter algorithms and other results for optimal partitions. Licentiate Thesis, Lund University, Sweden, 2004.

[5] M. Grantson and C. Levcopoulos. A fixed-parameter algorithm for the minimum number convex partition problem. In *JCDCG*, LNCS. Springer, 2004. to appear.

[6] M. Hoffmann and Y. Okamoto. The minimum weight triangulation problem with few inner points. In *IWPEC*, volume 3162 of *LNCS*, pages 200–212. Springer, 2004.

[7] J. M. Keil. *Decomposing a Polygon into Simpler Components*. PhD thesis, University of Toronto, 1983.

[8] J. M. Keil and J. Snoeyink. On the time bound for convex decomposition of simple polygons. *Internat. J. Comput. Geom. Appl.*, 12:181–192, 2002.

[9] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *Journal of Algorithms*, 27:303–338, 1998.

[10] A. Lingas. The power of non-rectilinear holes. In *Internat. Colloq. Automata Lang. Program.*, volume 140 of *LNCS*, pages 369–383. Springer, 1982.

[11] D. A. Plaisted and J. Hong. A heuristic triangulation algorithm. *Journal of Algorithms*, 8:405–437, 1987.