Minimum perimeter-sum partitions in the plane

Citation for published version (APA):

Abrahamsen, M., de Berg, M., Buchin, K., Mehr, M., & Mehrabi, A. D. (2020). Minimum perimeter-sum partitions in the plane. *Discrete and Computational Geometry*, 63(2), 483-505. <https://doi.org/10.1007/s00454-019-00059-0>

Document license:
TAVERNE

DOI:
[10.1007/s00454-019-00059-0](https://doi.org/10.1007/s00454-019-00059-0)

Document status and date:
Published: 01/03/2020

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Minimum Perimeter-Sum Partitions in the Plane

Mikkel Abrahamsen¹ · Mark de Berg² · Kevin Buchin² · Mehran Mehr² · Ali D. Mehrabi²

Received: 31 January 2018 / Revised: 21 December 2018 / Accepted: 13 January 2019 /
Published online: 6 February 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Let P be a set of n points in the plane. We consider the problem of partitioning P into two subsets P_1 and P_2 such that the sum of the perimeters of $\text{CH}(P_1)$ and $\text{CH}(P_2)$ is minimized, where $\text{CH}(P_i)$ denotes the convex hull of P_i . The problem was first studied by Mitchell and Wynters in 1991 who gave an $O(n^2)$ time algorithm. Despite considerable progress on related problems, no subquadratic time algorithm for this problem was found so far. We present an exact algorithm solving the problem in $O(n \log^2 n)$ time and a $(1 + \varepsilon)$ -approximation algorithm running in $O(n + 1/\varepsilon^2 \cdot \log^2(1/\varepsilon))$ time.

Keywords Computational geometry · Clustering · Minimum-perimeter partition · Convex hull

Mathematics Subject Classification 68W05

Editor in Charge: Kenneth Clarkson

A preliminary version of this paper appeared at the 33rd International Symposium on Computational Geometry (SoCG 2017).

Mikkel Abrahamsen
miab@di.ku.dk

Mark de Berg
M.T.d.Berg@tue.nl

Kevin Buchin
k.a.buchin@tue.nl

Mehran Mehr
mehran.mehr@gmail.com

Ali D. Mehrabi
admehrabi@gmail.com

¹ Department of Computer Science, University of Copenhagen, 2100 Copenhagen, Denmark

² Department of Computer Science, TU Eindhoven, 5600 MB Eindhoven, The Netherlands

1 Introduction

The clustering problem is to partition a given data set into clusters (that is, subsets) according to some measure of optimality. We are interested in clustering problems where the data set is a set P of points in Euclidean space. Most of these clustering problems fall into one of two categories: problems where the maximum cost of a cluster is given and the goal is to find a clustering consisting of a minimum number of clusters, and problems where the number of clusters is given and the goal is to find a clustering of minimum total cost. In this paper we consider a basic problem of the latter type, where we wish to find a bipartition (P_1, P_2) of a planar point set P . Bipartition problems are not only interesting in their own right, but also because bipartition algorithms can form the basis of hierarchical clustering methods.

There are many possible variants of the bipartition problem on planar point sets, which differ in how the cost of a clustering is defined. A variant that received a lot of attention is the 2-center problem [8, 11, 12, 15, 21], where the cost of a partition (P_1, P_2) of the given point set P is defined as the maximum of the radii of the smallest enclosing disks of P_1 and P_2 . Other cost functions that have been studied include the maximum diameter of the two point sets [5] and the sum of the diameters [14]; see also the survey by Agarwal and Sharir [3] for some more variants.

A natural class of cost functions considers the size of the convex hulls $\text{CH}(P_1)$ and $\text{CH}(P_2)$ of the two subsets, where the size of $\text{CH}(P_i)$ can either be defined as the area of $\text{CH}(P_i)$ or as the perimeter $\text{per}(P_i)$ of $\text{CH}(P_i)$. (The perimeter of $\text{CH}(P_i)$ is the length of the boundary $\partial \text{CH}(P_i)$.) This class of cost functions was already studied in 1991 by Mitchell and Wynters [17]. They studied four problem variants: minimize the sum of the perimeters, the maximum of the perimeters, the sum of the areas, or the maximum of the areas. In three of the four variants the convex hulls $\text{CH}(P_1)$ and $\text{CH}(P_2)$ in an optimal solution may intersect [17, full version]—only in the *minimum perimeter-sum problem* the optimal bipartition is guaranteed to be a so-called *line partition*, that is, a solution with disjoint convex hulls. For each of the four variants they gave an $O(n^3)$ algorithm that uses $O(n)$ storage and that computes an optimal line partition; for all except the minimum area-maximum problem they also gave an $O(n^2)$ algorithm that uses $O(n^2)$ storage. Note that (only) for the minimum perimeter-sum problem the computed solution is an optimal bipartition. Around the same time, the minimum-perimeter sum problem was studied for partitions into k subsets for $k > 2$; for this variant Capote et al. [7] presented an algorithm with running time $O(n^{6k})$. Arkin et al. [4] studied the same problem and gave a similar algorithm. Very recently, Abrahamsen et al. [1] gave an algorithm for that problem running in time $O(n^{28})$, even when k is part of the input. Unless $P = NP$, this result refutes a conjecture by Arkin et al. [4] that the problem is NP-complete.

Mitchell and Wynters mentioned the improvement of the space requirement of the quadratic-time algorithms for the bipartition problems as an open problem, and they stated the existence of a subquadratic algorithm for any of the four variants as the most prominent open problem.

Rokne et al. [19] made progress on the first question, by presenting an $O(n^2 \log n)$ algorithm that uses only $O(n)$ space for the line-partition version of each of the four problems. Devillers and Katz [10] gave algorithms for the min-max variant of the

problem, both for area and perimeter, which run in $O((n+k)\log^2 n)$ time. Here k is a parameter that is only known to be in $O(n^2)$, although Devillers and Katz suspected that k is subquadratic. They also gave linear-time algorithms for these problems when the point set P is in convex position and given in cyclic order. Segal [20] proved an $\Omega(n \log n)$ lower bound for the min-max problems. Very recently, and apparently unaware of some of the earlier work on these problems, Bae et al. [6] presented an $O(n^2 \log n)$ time algorithm for the minimum-perimeter-sum problem and an $O(n^4 \log n)$ time algorithm for the minimum-area-sum problem (considering all partitions, not only line partitions). Despite these efforts, the main question is still open: is it possible to obtain a subquadratic algorithm for any of the four bipartition problems based on convex-hull size?

1.1 Our Contribution

We answer the question above affirmatively by presenting a subquadratic algorithm for the minimum perimeter-sum bipartition problem in the plane.

As mentioned, an optimal solution (P_1, P_2) to the minimum perimeter-sum bipartition problem must be a line partition. A straightforward algorithm would generate all $\Theta(n^2)$ line partitions and compute the value $\text{per}(P_1) + \text{per}(P_2)$ for each of them. If the latter is done from scratch for each partition, the resulting algorithm runs in $O(n^3 \log n)$ time. The algorithms by Mitchell and Wynters [17] and Rokne et al. [19] improve on this by using the fact that the different line bipartitions can be generated in an ordered way, so that subsequent line partitions differ in at most one point. Thus the convex hulls do not have to be recomputed from scratch, but they can be obtained by updating the convex hulls of the previous bipartition. To obtain a subquadratic algorithm a fundamentally new approach is necessary: we need a strategy that generates a subquadratic number of candidate partitions, instead of considering all line partitions. We achieve this as follows.

We start by proving that an optimal bipartition (P_1, P_2) has the following property: either there is a set of $O(1)$ canonical orientations such that P_1 can be separated from P_2 by a line with a canonical orientation, or the distance between $\text{CH}(P_1)$ and $\text{CH}(P_2)$ is $\Omega(\min(\text{per}(P_1), \text{per}(P_2)))$. There are only $O(n)$ bipartitions of the former type, and finding the best among them is relatively easy. The bipartitions of the second type are much more challenging. We show how to employ a compressed quadtree to generate a collection of $O(n)$ canonical 5-gons—intersections of axis-parallel rectangles and canonical halfplanes—such that the smaller of $\text{CH}(P_1)$ and $\text{CH}(P_2)$ (in a bipartition of the second type) is contained in one of the 5-gons.

Even though the number of such bipartitions is linear, we cannot afford to compute their perimeters from scratch. We therefore use the data structure of Oh and Ahn [18] to quickly compute $\text{per}(P \cap Q)$, where Q is a query canonical 5-gon. Given a set $\mathcal{O}$ of k orientations, Oh and Ahn described how to create a data structure using $O(nk^3 \log^2 n)$ time and space to answer queries of the following type in time $O(k \log^2 n)$: Given a convex polygon Q where each edge has an orientation in $\mathcal{O}$, what is $\text{per}(P \cap Q)$? In our case, each query polygon Q is the intersection of an axis-parallel square and a canonical halfplane bounded by a line with one of $C = O(1)$ different orientations.

We therefore make C different instances of the data structure, where each instance has as orientations $\mathcal{O}$ the two axis-parallel directions and one of the C different orientations of the canonical halfplanes (i.e., $k = 3$).¹

To sum up, our main result is an exact algorithm for the minimum perimeter-sum bipartition problem that runs in $O(n \log^2 n)$ time. As our model of computation we use the real RAM (with the capability of taking square roots) so that we can compute the exact perimeter of a convex polygon—this is necessary to compare the costs of two competing clusterings. We furthermore make the (standard) assumption that the model of computation allows us to compute a compressed quadtree of n points in $O(n \log n)$ time; see footnote 3 in Sect. 2.2.2.

Besides our exact algorithm, we present a linear-time $(1 + \varepsilon)$ -approximation algorithm. Its running time is $O(n + T(1/\varepsilon^2)) = O(n + 1/\varepsilon^2 \cdot \log^2(1/\varepsilon))$, where $T(1/\varepsilon^2)$ is the running time of an exact algorithm on an instance of size $1/\varepsilon^2$.

2 The Exact Algorithm

In this section we present an exact algorithm for the minimum-perimeter-sum partition problem. We first prove a separation property that an optimal solution must satisfy, and then we show how to use this property to develop a fast algorithm.

Let P be the set of n points in the plane for which we want to solve the minimum-perimeter-sum partition problem. An optimal partition (P_1, P_2) of P has the following two basic properties: P_1 and P_2 are non-empty, and the convex hulls $\text{CH}(P_1)$ and $\text{CH}(P_2)$ are disjoint [17, full version]. In the remainder, whenever we talk about a partition of P , we refer to a partition with these two properties.

2.1 Geometric Properties of an Optimal Partition

Consider a partition (P_1, P_2) of P . Define $\mathcal{P}_1 := \text{CH}(P_1)$ and $\mathcal{P}_2 := \text{CH}(P_2)$ to be the convex hulls of P_1 and P_2 , respectively, and let ℓ_1 and ℓ_2 be the two inner common tangents of $\mathcal{P}_1$ and $\mathcal{P}_2$. The lines ℓ_1 and ℓ_2 define four wedges: one containing P_1 , one containing P_2 , and two empty wedges. We call the opening angle of the empty wedges the *separation angle* of P_1 and P_2 . Furthermore, we call the distance between $\mathcal{P}_1$ and $\mathcal{P}_2$ the *separation distance* of P_1 and P_2 .

Theorem 2.1 *Let P be a set of n points in the plane, and let (P_1, P_2) be a partition of P that minimizes $\text{per}(P_1) + \text{per}(P_2)$. Then the separation angle of P_1 and P_2 is at least $\pi/6$ or the separation distance is at least $c_{\text{sep}} \cdot \min(\text{per}(P_1), \text{per}(P_2))$, where $c_{\text{sep}} := 1/250$.*

The remainder of this section is devoted to proving Theorem 2.1. To this end let (P_1, P_2) be a partition of P that minimizes $\text{per}(P_1) + \text{per}(P_2)$. Let ℓ_3 and ℓ_4 be the

¹ In a preliminary version of this paper [2], we described a less efficient data structure answering these queries in time $O(\log^4 n)$, resulting in the total running time $O(n \log^4 n)$. After that Oh and Ahn [18] developed a more efficient data structure that, as they already observed, can be used to speed up our algorithm.

Fig. 1 The setup in the proof of Theorem 2.1

outer common tangents of $\mathcal{P}_1$ and $\mathcal{P}_2$. We define α to be the angle between ℓ_3 and ℓ_4 . More precisely, if ℓ_3 and ℓ_4 are parallel we define $\alpha := 0$, otherwise we define α as the opening angle of the wedge defined by ℓ_3 and ℓ_4 containing $\mathcal{P}_1$ and $\mathcal{P}_2$. We denote the separation angle of $\mathcal{P}_1$ and $\mathcal{P}_2$ by β ; see Fig. 1.

The idea of the proof is as follows. Suppose that the separation distance and the separation angle β are both relatively small. Then the region A in between $\mathcal{P}_1$ and $\mathcal{P}_2$ and bounded from the bottom by ℓ_3 and from the top by ℓ_4 is relatively narrow. But then the left and right parts of ∂A (which are contained in $\partial \mathcal{P}_1$ and $\partial \mathcal{P}_2$) would be longer than the bottom and top parts of ∂A (which are contained in ℓ_3 and ℓ_4), thus contradicting the assumption that $(\mathcal{P}_1, \mathcal{P}_2)$ is an optimal partition. To make this idea precise, we first prove that if the separation angle β is small, then the angle α between ℓ_3 and ℓ_4 must be large. Second, we show that there is a value $f(\alpha)$ such that the distance between $\mathcal{P}_1$ and $\mathcal{P}_2$ is at least $f(\alpha) \cdot \min(\text{per}(\mathcal{P}_1), \text{per}(\mathcal{P}_2))$. Finally we argue that this implies that if the separation angle is smaller than $\pi/6$, then (to avoid the contradiction mentioned above) the separation distance must be relatively large. Next we present our proof in detail.

Let c_{ij} be the intersection point between ℓ_i and ℓ_j , where $i < j$. If ℓ_3 and ℓ_4 are parallel, we choose c_{34} as a point at infinity on ℓ_3 . Assume without loss of generality that neither ℓ_1 nor ℓ_2 separate $\mathcal{P}_1$ from c_{34} , and that ℓ_3 is the outer common tangent such that $\mathcal{P}_1$ and $\mathcal{P}_2$ are to the left of ℓ_3 when traversing ℓ_3 from c_{34} to an intersection point in $\ell_3 \cap \mathcal{P}_1$. Assume furthermore that c_{13} is closer to c_{34} than c_{23} .

For two lines, rays, or segments r_1, r_2 , let $\angle(r_1, r_2)$ be the angle we need to rotate r_1 in a counterclockwise direction until r_1 and r_2 are parallel. For three points a, b, c , let $\angle(a, b, c) := \angle(ba, bc)$. For $i = 1, 2$ and $j = 1, 2, 3, 4$, let s_{ij} be a point in $\mathcal{P}_i \cap \ell_j$. Let $\partial \mathcal{P}_i$ denote the boundary of $\mathcal{P}_i$ and $\text{per}(\mathcal{P}_i)$ the perimeter of $\mathcal{P}_i$. Furthermore, let $\partial \mathcal{P}_i(x, y)$ denote the portion of $\partial \mathcal{P}_i$ from $x \in \partial \mathcal{P}_i$ counterclockwise to $y \in \partial \mathcal{P}_i$, and $\text{length}(\partial \mathcal{P}_i(x, y))$ denote the length of $\partial \mathcal{P}_i(x, y)$.

Lemma 2.2 Let p_0 and q be points and $\mathbf{v}$ be a unit vector. Let $p(t) := p_0 + t \cdot \mathbf{v}$ and $d(t) := |p(t)q|$ and assume that $p(t) \neq q$ for all $t \in \mathbb{R}$. Then $d'(t) = \cos(\angle(q, p(t), p(t) + \mathbf{v}))$ if the points $q, p(t), p(t) + \mathbf{v}$ make a left-turn and $d'(t) = -\cos(\angle(q, p(t), p(t) + \mathbf{v}))$ otherwise.²

Proof We prove the lemma for an arbitrary value $t = t_0$. By reparameterizing p , we may assume that $t_0 = 0$. Furthermore, by changing the coordinate system, we can without loss of generality assume that $p_0 = (0, 0)$ and $q = (x, 0)$ for some value $x > 0$.

Let $\phi := \angle((x, 0), (0, 0), \mathbf{v})$. Assume that $\mathbf{v}$ has positive y -coordinate—the case that $\mathbf{v}$ has negative y -coordinate can be handled analogously. We have proved the lemma if we manage to show that $d'(0) = -\cos \phi$. Note that since $\mathbf{v}$ has positive y -coordinate, we have $p(t) = (t \cos \phi, t \sin \phi)$ for every $t \in \mathbb{R}$. Hence

$$d(t) = \sqrt{(t \cos \phi - x)^2 + t^2 \sin^2 \phi}$$

and

$$d'(t) = \frac{t - x \cos \phi}{\sqrt{t^2 - 2tx \cos \phi + x^2}}.$$

Evaluating at $t = 0$, we get

$$d'(0) = -\frac{x \cos \phi}{|x|} = -\cos \phi,$$

where the last equality follows since $x > 0$. □

Lemma 2.3 We have $\alpha + 3\beta \geq \pi$.

Proof Since $\text{per}(\mathcal{P}_1) + \text{per}(\mathcal{P}_2)$ is minimum, we know that

$$\text{length}(\partial \mathcal{P}_1(s_{13}, s_{14})) + \text{length}(\partial \mathcal{P}_2(s_{24}, s_{23})) \leq \Psi,$$

where $\Psi := |s_{13}s_{23}| + |s_{14}s_{24}|$. Furthermore, we know that $s_{11}, s_{12} \in \partial \mathcal{P}_1(s_{13}, s_{14})$ and $s_{21}, s_{22} \in \partial \mathcal{P}_1(s_{24}, s_{23})$. We thus have

$$\text{length}(\partial \mathcal{P}_1(s_{13}, s_{14})) + \text{length}(\partial \mathcal{P}_2(s_{24}, s_{23})) \geq \Phi,$$

where $\Phi := |s_{13}s_{11}| + |s_{11}s_{12}| + |s_{12}s_{14}| + |s_{24}s_{21}| + |s_{21}s_{22}| + |s_{22}s_{23}|$. Hence, we must have

$$\Phi \leq \Psi. \tag{1}$$

Now assume that $\alpha + 3\beta < \pi$. We will show that this assumption, together with inequality (1), leads to a contradiction, thus proving the lemma. To this end we will

² Note that $\angle(q, p(t), p(t) + \mathbf{v}) = \angle(q, p(t), p(t) - \mathbf{v})$ by the definition of $\angle(\cdot, \cdot, \cdot)$ which is the reason that there are two cases in the lemma.

argue that if (1) holds, then there exist points s'_{ij} for $i = 1, 2$ and $j = 1, 2, 3, 4$, where s'_{ij} is a point on ℓ_j , with the following properties:

- (i) $\Phi' \leq \Psi'$, where Φ' and Ψ' are defined as Φ and Ψ when each point s_{ij} is replaced by s'_{ij} ,
- (ii) s'_{21} or s'_{22} coincides with c_{12} , and
- (iii) s'_{11} or s'_{12} coincides with c_{12} .

To finish the proof it then suffices to observe that properties (i)–(iii) together contradict the triangle inequality.

Note that the point s'_{ij} is not required to be contained in P_i . In particular, the points s'_{13} and s'_{14} will in some cases be on the other side of c_{34} than the points s_{13} and s_{14} . In that case there is no pair of convex polygons with outer common tangents defined by (s'_{13}, s'_{23}) and (s'_{14}, s'_{24}) . The contradiction applies to distances between a configuration of points that need not be realizable as the supporting points of the common tangents of two convex polygons.

To prove the existence of the points s'_{ij} with the claimed properties, we initially define $s'_{ij} := s_{ij}$, so that property (i) is satisfied. Then we will move the points s'_{ij} (where each s'_{ij} moves on ℓ_j) so that property (i) is preserved throughout the movements and properties (ii) and (iii) are satisfied at the end of the movements.

We first show how to create a situation where (ii) holds, and (i) still holds as well. Let $\gamma_{ij} := \angle(\ell_i, \ell_j)$. We consider two cases.

- *Case (A): $\gamma_{32} < \pi - \beta$.*

We observe that moving s'_{23} along ℓ_3 away from s'_{13} increases Ψ' more than it increases Φ' , so property (i) is preserved by such a movement. Note that $\angle(xs'_{23}, \ell_2) \geq \gamma_{32}$ for any $x \in s'_{22}c_{12}$. However, by moving s'_{23} sufficiently far away we can make $\angle(xs'_{23}, \ell_2)$ arbitrarily close to γ_{32} . We therefore move s'_{23} so far away that $\angle(xs'_{23}, \ell_2) < \pi - \beta$ for any point $x \in s'_{22}c_{12}$. We now consider what happens as we let a point x move at unit speed from s'_{22} towards c_{12} . To be more precise, let $T := |s'_{22}c_{12}|$, let $\mathbf{v}$ be the unit vector with direction from c_{23} to c_{12} , and for any $t \in [0, T]$ define $x(t) := s'_{22} + t \cdot \mathbf{v}$. Note that $x(0) = s'_{22}$ and $x(T) = c_{12}$.

Let $a(t) := |x(t)s'_{23}|$ and $b(t) := |x(t)s'_{21}|$. Lemma 2.2 gives that

$$a'(t) = -\cos(\angle(x(t)s'_{23}, \ell_2)) \quad \text{and} \quad b'(t) = \cos(\angle(\ell_2, x(t)s'_{21})).$$

Since $\angle(x(t)s'_{23}, \ell_2) < \pi - \beta$ for any value $t \in [0, T]$, we get $a'(t) < -\cos(\pi - \beta)$. Furthermore, we have $\angle(\ell_2, x(t)s'_{21}) \geq \pi - \beta$ and hence $b'(t) \leq \cos(\pi - \beta)$. Therefore, $a'(t) + b'(t) < 0$ for any t and we conclude that $a(T) + b(T) \leq a(0) + b(0)$. This is the same as $|s'_{21}c_{12}| + |c_{12}s'_{23}| \leq |s'_{21}s'_{22}| + |s'_{22}s'_{23}|$, so we now move s'_{22} to c_{12} and are ensured that (i) still holds.

- *Case (B): $\gamma_{32} \geq \pi - \beta$.*

Using our assumption $\alpha + 3\beta < \pi$ we get $\gamma_{32} > \alpha + 2\beta$. Note that $\gamma_{14} = \pi - \gamma_{32} + \alpha + \beta$. Hence, $\gamma_{14} < \pi - \beta$. By first moving s'_{24} away from s'_{14} and then s'_{21} towards c_{12} , we can argue, similarly to Case (A), that we can reach a situation where (i) still holds and s'_{21} coincides with c_{12} .

Fig. 2 Illustration for Lemma 2.4. Φ is the total length of the four segments t_1m , t_2m , b_1m , b_2m , and Ψ is the total length of the two fat segments

We conclude that in both cases we can ensure (ii) without violating (i).

Since $\gamma_{13} \leq \gamma_{14}$ and $\gamma_{42} \leq \gamma_{32}$, we likewise have $\gamma_{13} < \pi - \beta$ or $\gamma_{42} < \pi - \beta$. Hence, by first moving s'_{13} or s'_{14} and since then s'_{11} or s'_{12} , we can in a similar way reach a situation where s'_{11} or s'_{12} coincides with c_{12} without violating (i), thus ensuring (iii) and finishing the proof. $\square$

The following technical lemma is illustrated in Fig. 2. The lemma will be used in the proof of the subsequent Lemma 2.5. The overall idea in the two lemmata is that we consider pushing $\mathcal{P}_2$ towards $\mathcal{P}_1$ until they touch. In the configuration where they touch, m in Lemma 2.4 corresponds to a common point, r_1 , r_2 correspond to the outer common tangents, and b_1 , t_1 , resp. b_2 , t_2 , correspond to the points where $\mathcal{P}_1$, resp. $\mathcal{P}_2$, supports r_1 , r_2 . The lemma then gives a lower bound on how much cheaper it would be to unite $\mathcal{P}_1$ and $\mathcal{P}_2$. This in turn implies a lower bound on how far we pushed $\mathcal{P}_2$ (using that $(\mathcal{P}_1, \mathcal{P}_2)$ was assumed to be an optimal bipartition), which is a lower bound on the original distance between $\mathcal{P}_1$ and $\mathcal{P}_2$, as stated in Theorem 2.1.

Lemma 2.4 *Let x be a point and r_1 and r_2 be two rays starting at x such that $\angle(r_1, r_2) = \delta$, and assume that $\delta \leq \pi$. Let $b_1, b_2 \in r_1$ and $t_1, t_2 \in r_2$ be such that $b_1 \in xb_2$ and $t_1 \in xt_2$, and let m be a point in the wedge bounded by r_1 and r_2 . Then*

$$\Phi - \Psi \geq \frac{(1 - \cos(\delta/2)) \cdot \sin(\delta/2)}{1 + \sin(\delta/2)} \cdot (|b_1m| + |t_1m|),$$

where $\Phi := |b_1m| + |t_1m| + |b_2m| + |t_2m|$ and $\Psi := |b_1b_2| + |t_1t_2|$.

Proof First note that

$$|b_1m| + |b_2m| \geq |b_1b_2| \quad (2)$$

and

$$|t_1m| + |t_2m| \geq |t_1t_2|. \quad (3)$$

Let r_3 be the angular bisector of r_1 and r_2 . Assume without loss of generality that m lies in the wedge defined by r_1 and r_3 . Then $\angle(m, t_1, t_2) \geq \delta/2$.

We now consider two cases.

- *Case (A):* $|t_1 m| \geq \frac{\sin(\delta/2)}{1+\sin(\delta/2)} \cdot (|b_1 m| + |t_1 m|)$.
Our first step is to prove that

$$|t_1 m| + |t_2 m| - |t_1 t_2| \geq (1 - \cos(\delta/2)) \cdot |t_1 m|. \quad (4)$$

Let p be the orthogonal projection of m on r_2 . Note that $|t_2 m| \geq |t_2 p|$. Consider first the case that p is on the same side of t_1 as x . In this case $|t_2 p| \geq |t_1 t_2|$ and therefore

$$|t_1 m| + |t_2 m| - |t_1 t_2| \geq |t_1 m| \geq (1 - \cos(\delta/2)) \cdot |t_1 m|,$$

which proves (4).

Assume now that p is on the same side of t_1 as t_2 . In this case, we have $\angle(m, t_1, t_2) \leq \pi/2$ and thus $|t_1 p| = \cos(\angle(m, t_1, t_2)) \cdot |t_1 m| \leq \cos(\delta/2) \cdot |t_1 m|$. Hence we have

$$|t_1 m| + |t_2 m| - |t_1 t_2| \geq |t_1 m| + |t_2 p| - (|t_1 p| + |t_2 p|) \geq (1 - \cos(\delta/2)) \cdot |t_1 m|,$$

and we have proved (4).

We now have

$$\begin{aligned} \Phi - \Psi &= |b_1 m| + |t_1 m| + |b_2 m| + |t_2 m| - |b_1 b_2| - |t_1 t_2| \\ &\geq |b_1 m| + |b_2 m| - |b_1 b_2| + (1 - \cos(\delta/2)) \cdot |t_1 m| \quad \text{by (4)} \\ &\geq (1 - \cos(\delta/2)) \cdot \frac{\sin(\delta/2)}{1+\sin(\delta/2)} \cdot (|b_1 m| + |t_1 m|) \quad \text{by (2)} \end{aligned}$$

where the last step uses that we are in Case (A). Thus the lemma holds in Case (A).

- *Case (B):* $|t_1 m| < \frac{\sin(\delta/2)}{1+\sin(\delta/2)} \cdot (|b_1 m| + |t_1 m|)$.
The condition for this case can be rewritten as

$$|b_1 m| > \frac{1}{1 + \sin \delta/2} \cdot (|b_1 m| + |t_1 m|). \quad (5)$$

To prove the lemma in this case we first argue that $\angle(b_2, b_1, m) > \pi/2$. To this end, assume for a contradiction that $\angle(b_2, b_1, m) \leq \pi/2$. It is easy to verify that for a given length of $t_1 m$ (and assuming $\angle(b_2, b_1, m) \leq \pi/2$), the fraction $|b_1 m|/(|b_1 m| + |t_1 m|)$ is maximized when segment $t_1 m$ is perpendicular to r_2 , and $m \in r_3$, and $b_1 = x$. But then

$$\frac{|b_1 m|}{|b_1 m| + |t_1 m|} \leq \frac{1}{1 + \sin \delta/2},$$

which would contradict (5). Thus we indeed have $\angle(b_2, b_1, m) > \pi/2$. Hence, $|b_2m| \geq |b_1b_2|$, and so $|b_1m| + |b_2m| - |b_1b_2| \geq |b_1m|$. We can now derive

$$\begin{aligned}\Phi - \Psi &= |b_1m| + |t_1m| + |b_2m| + |t_2m| - |b_1b_2| - |t_1t_2| \\ &\geq |b_1m| + |t_1m| + |t_2m| - |t_1t_2| \quad \text{by the above} \\ &\geq \frac{1}{1+\sin \delta/2} \cdot (|b_1m| + |t_1m|) \quad \text{by (3) and (5)} \\ &\geq (\sin(\delta/2) \cdot (1 - \cos(\delta/2))) \cdot \frac{1}{1+\sin \delta/2} \cdot (|b_1m| + |t_1m|).\end{aligned}$$

Thus the lemma also holds in Case (B). $\square$

Let $\text{dist}(\mathcal{P}_1, \mathcal{P}_2) := \min_{(p,q) \in \mathcal{P}_1 \times \mathcal{P}_2} |pq|$ denote the separation distance between $\mathcal{P}_1$ and $\mathcal{P}_2$. Recall that α denotes the angle between the two common outer tangents of $\mathcal{P}_1$ and $\mathcal{P}_2$; see Fig. 1. We are now ready to give a lower bound on the separation distance increasing in the angle α between the outer common tangents ℓ_3 and ℓ_4 . The lemma will be used when there is a positive lower bound on α , which in turn implies a lower bound on $\text{dist}(\mathcal{P}_1, \mathcal{P}_2)$.

Lemma 2.5 *We have*

$$\text{dist}(\mathcal{P}_1, \mathcal{P}_2) \geq f(\alpha) \cdot \text{per}(\mathcal{P}_1), \quad (6)$$

where $f: [0, \pi] \rightarrow \mathbb{R}$ is the increasing function

$$f(\varphi) := \frac{\sin(\varphi/4)}{1 + \sin(\varphi/4)} \cdot \frac{\sin(\varphi/2)}{1 + \sin(\varphi/2)} \cdot \frac{1 - \cos(\varphi/4)}{2}.$$

Proof The statement is trivial if $\alpha = 0$ so assume $\alpha > 0$. Let $p \in \mathcal{P}_1$ and $q \in \mathcal{P}_2$ be points so that $|pq| = \text{dist}(\mathcal{P}_1, \mathcal{P}_2)$ and assume without loss of generality that pq is a horizontal segment with p being its left endpoint. Let ℓ_1^{vert} and ℓ_2^{vert} be vertical lines containing p and q , respectively. Note that $\mathcal{P}_1$ is in the closed halfplane to the left of ℓ_1^{vert} and $\mathcal{P}_2$ is in the closed halfplane to the right of ℓ_2^{vert} . Recall that s_{ij} denotes a point on $\partial \mathcal{P}_i \cap \ell_j$. $\square$

Claim: There exist two convex polygons $\mathcal{P}'_1$ and $\mathcal{P}'_2$ satisfying the following conditions:

1. $\mathcal{P}'_1$ and $\mathcal{P}'_2$ have the same outer common tangents as $\mathcal{P}_1$ and $\mathcal{P}_2$, namely ℓ_3 and ℓ_4 .
2. $\mathcal{P}'_1$ is to the left of ℓ_1^{vert} and $p \in \partial \mathcal{P}'_1$; and $\mathcal{P}'_2$ is to right of ℓ_2^{vert} and $q \in \partial \mathcal{P}'_2$.
3. $\text{per}(\mathcal{P}'_1) = \text{per}(\mathcal{P}_1)$.
4. $\text{per}(\mathcal{P}'_1) + \text{per}(\mathcal{P}'_2) \leq \text{per}(\text{CH}(\mathcal{P}'_1 \cup \mathcal{P}'_2))$.
5. There are points $s'_{ij} \in \mathcal{P}'_i \cap \ell_j$ for all $i \in \{1, 2\}$ and $j \in \{3, 4\}$ such that $\partial \mathcal{P}'_1(s'_{13}, p)$, $\partial \mathcal{P}'_1(p, s'_{14})$, $\partial \mathcal{P}'_2(s'_{24}, q)$, and $\partial \mathcal{P}'_2(q, s'_{23})$ each consist of a single line segment.
6. Let $s'_{2j}(\lambda) := s'_{2j} - (\lambda, 0)$ and let $\ell'_j(\lambda)$ be the line through s'_{1j} and $s'_{2j}(\lambda)$ for $j \in \{3, 4\}$. Then $\angle(\ell'_3(|pq|), \ell'_4(|pq|)) \geq \alpha/2$. (Looking at Fig. 3, one might believe that this inequality even holds for α instead of $\alpha/2$. The reason for using $\alpha/2$ will be explained later.)

Proof of the claim Let $\mathcal{P}'_1 := \mathcal{P}_1$ and $\mathcal{P}'_2 := \mathcal{P}_2$, and let s'_{ij} be a point in $\mathcal{P}'_i \cap \ell_j$ for all $i \in \{1, 2\}$ and $j \in \{3, 4\}$. We show how to modify $\mathcal{P}'_1$ and $\mathcal{P}'_2$ until they have all the required conditions. Of course, they already satisfy conditions 1–4. We first show how to obtain condition 5, namely that $\partial \mathcal{P}'_1(s'_{13}, p)$ and $\partial \mathcal{P}'_1(p, s'_{14})$ —and similarly $\partial \mathcal{P}'_2(s'_{24}, q)$ and $\partial \mathcal{P}'_2(q, s'_{23})$ —each consist of a single line segment, as depicted in Fig. 3. To this end, let v_{ij} be the intersection point $\ell_i^{\text{vert}} \cap \ell_j$ for $i \in \{1, 2\}$ and $j \in \{3, 4\}$. Let $s' \in s'_{14}v_{14}$ be the point such that $\text{length}(\partial \mathcal{P}'_1(p, s')) = |ps'| + |s's'_{14}|$. Such a point exists since

$$|ps'_{14}| \leq \text{length}(\partial \mathcal{P}'_1(p, s'_{14})) \leq |pv_{14}| + |v_{14}s'_{14}|.$$

The only condition that $(\mathcal{P}'_1, \mathcal{P}'_2)$ might not satisfy is condition 6. Let $s'_{2j}(\lambda) := s'_{2j} - (\lambda, 0)$ and let $\ell_j(\lambda)$ be the line through $s'_{2j}(\lambda)$ and s'_{1j} for $j \in \{3, 4\}$. Clearly, if the slopes of ℓ_3 and ℓ_4 have different signs (as in Fig. 3), the angle $\angle(\ell_3(\lambda), \ell_4(\lambda))$ is increasing for $\lambda \in [0, |pq|]$, and condition 6 is satisfied. However, if the slopes of ℓ_3 and ℓ_4 have the same sign, the angle might decrease.

Consider the case where both slopes are positive—the other case is analogous. Changing $\mathcal{P}'_2$ by replacing $\partial \mathcal{P}'_2(s'_{23}, s'_{24})$ by the line segment $s'_{23}s'_{24}$ makes the sum $\text{per}(\mathcal{P}'_1) + \text{per}(\mathcal{P}'_2)$ and $\text{per}(\text{CH}(\mathcal{P}'_1 \cup \mathcal{P}'_2))$ decrease equally much and hence condition 4 is preserved. This clearly has no influence on the other conditions. We thus assume that $\mathcal{P}'_2$ is the triangle $qs'_{23}s'_{24}$. Consider what happens if we move s'_{23} along the line ℓ_3 away from c_{34} with unit speed. Then $|s'_{13}s'_{23}|$ grows with speed exactly 1 whereas $|qs'_{23}|$ grows with speed at most 1. We therefore preserve condition 4, and the other conditions are likewise not affected.

We now move s'_{23} sufficiently far away so that $\angle(\ell_3, \ell_3(|pq|)) \leq \alpha/4$. Similarly, we move s'_{24} sufficiently far away from c_{34} along ℓ_4 to ensure that $\angle(\ell_4, \ell_4(|pq|)) \leq \alpha/4$. It then follows that $\angle(\ell_3(|pq|), \ell_4(|pq|)) \geq \angle(\ell_3, \ell_4) - \alpha/2 = \alpha/2$, and condition 6 is satisfied. $\square$

Note that condition 2 in the claim implies that $\text{dist}(\mathcal{P}'_1, \mathcal{P}'_2) = \text{dist}(\mathcal{P}_1, \mathcal{P}_2) = |pq|$, and hence inequality (6) follows from condition 3 if we manage to prove $\text{dist}(\mathcal{P}'_1, \mathcal{P}'_2) \geq f(\alpha) \cdot \text{per}(\mathcal{P}'_1)$. Therefore, with a slight abuse of notation, we assume from now on that $\mathcal{P}_1$ and $\mathcal{P}_2$ satisfy the conditions in the claim, where the points s_{ij} play the role as s'_{ij} in conditions 5 and 6.

We now consider a copy of $\mathcal{P}_2$ that is translated horizontally to the left over a distance λ ; see Fig. 3. Let $s_{24}(\lambda)$, $s_{23}(\lambda)$, and $q(\lambda)$ be the translated copies of s_{24} , s_{23} , and q , respectively, and let $\ell_j(\lambda)$ be the line through s_{1j} and $s_{2j}(\lambda)$ for $j \in \{3, 4\}$. Furthermore, define

$$\Phi(\lambda) := |s_{13}p| + |s_{14}p| + |s_{23}(\lambda)q(\lambda)| + |s_{24}(\lambda)q(\lambda)|$$

and

$$\Psi(\lambda) := |s_{13}s_{23}(\lambda)| + |s_{14}s_{24}(\lambda)|.$$

Note that $\Phi(\lambda) = \Phi$ is constant. By conditions 4 and 5, we know that

$$\Phi \leq \Psi(0). \quad (7)$$

Note that $q(|pq|) = p$. We now apply Lemma 2.4 to get

$$\Phi - \Psi(|pq|) \geq \sin(\delta/2) \cdot \frac{1 - \cos(\delta/2)}{1 + \sin(\delta/2)} \cdot (|s_{13}p| + |s_{14}p|), \quad (8)$$

where $\delta := \angle(\ell_3(|pq|), \ell_4(|pq|))$. By condition 6, we know that $\delta \geq \alpha/2$. The function $\varphi \mapsto \sin(\varphi/2) \cdot \frac{1 - \cos(\varphi/2)}{1 + \sin(\varphi/2)}$ is increasing for $\varphi \in [0, \pi]$ and hence inequality (8) also holds when δ is replaced by $\alpha/2$.

When λ increases from 0 to $|pq|$ with unit speed, the value $\Psi(\lambda)$ decreases with speed at most 2, i.e., $\Psi(\lambda) \geq \Psi(0) - 2\lambda$. Using this and inequalities (7) and (8), we get

$$2|pq| \geq \Psi(0) - \Psi(|pq|) \geq \Phi - \Phi + \sin(\alpha/4) \cdot \frac{1 - \cos(\alpha/4)}{1 + \sin(\alpha/4)} \cdot (|s_{13}p| + |s_{14}p|),$$

and we conclude that

$$|pq| \geq \frac{1}{2} \cdot \sin(\alpha/4) \cdot \frac{1 - \cos(\alpha/4)}{1 + \sin(\alpha/4)} \cdot (|s_{13}p| + |s_{14}p|). \quad (9)$$

By the triangle inequality, $|s_{13}p| + |s_{14}p| \geq |s_{13}s_{14}|$. Furthermore, for a given length of $s_{13}s_{14}$, the fraction $|s_{13}s_{14}|/(|s_{14}c_{34}| + |c_{34}s_{13}|)$ is minimized when $s_{13}s_{14}$ is

perpendicular to the angular bisector of ℓ_3 and ℓ_4 . (Recall that c_{34} is the intersection point of the outer common tangents ℓ_3 and ℓ_4 ; see Fig. 3.) Hence

$$|s_{13}s_{14}| \geq \sin(\alpha/2) \cdot (|s_{14}c_{34}| + |c_{34}s_{13}|). \quad (10)$$

We now conclude

$$\begin{aligned} |s_{13}p| + |s_{14}p| &= \frac{\sin(\alpha/2)}{1+\sin(\alpha/2)} \cdot \left(\frac{|s_{13}p|+|s_{14}p|}{\sin(\alpha/2)} + |s_{13}p| + |s_{14}p| \right) \\ &\geq \frac{\sin(\alpha/2)}{1+\sin(\alpha/2)} \cdot \left(\frac{|s_{13}s_{14}|}{\sin(\alpha/2)} + |s_{13}p| + |s_{14}p| \right) \quad \text{triangle inequality} \\ &\geq \frac{\sin(\alpha/2)}{1+\sin(\alpha/2)} \cdot (|s_{14}c_{34}| + |c_{34}s_{13}| + |s_{13}p| + |s_{14}p|) \quad \text{by (10)} \\ &\geq \frac{\sin(\alpha/2)}{1+\sin(\alpha/2)} \cdot \text{per}(\mathcal{P}_1), \end{aligned}$$

where the last inequality follows because $\mathcal{P}_1$ is fully contained in the quadrilateral $s_{14}, c_{34}, x_{13}, p$. The statement (6) in the lemma now follows from (9). $\square$

We are now ready to prove Theorem 2.1.

Proof of Theorem 2.1 If the separation angle of P_1 and P_2 is at least $\pi/6$, we are done. Otherwise, Lemma 2.3 gives that $\alpha > \pi/2$, and Lemma 2.5 gives that $\text{dist}(\mathcal{P}_1, \mathcal{P}_2) \geq f(\pi/2) \cdot \text{per}(\mathcal{P}_1) \geq (1/250) \cdot \min(\text{per}(\mathcal{P}_1), \text{per}(\mathcal{P}_2))$. $\square$

2.2 The Algorithm

Theorem 2.1 suggests to distinguish two cases when computing an optimal partition: the case when the separation angle is large (namely at least $\pi/6$) and the case when the separation distance is large (namely at least $c_{\text{sep}} \cdot \min(\text{per}(P_1), \text{per}(P_2))$). As we will see, the first case can be handled in $O(n \log n)$ time and the second case in $O(n \log^2 n)$ time, leading to the following theorem.

Theorem 2.6 *Let P be a set of n points in the plane. Then we can compute a partition (P_1, P_2) of P that minimizes $\text{per}(P_1) + \text{per}(P_2)$ in $O(n \log^2 n)$ time using $O(n \log^2 n)$ space.*

2.2.1 The Best Partition with Large Separation Angle

Define the *orientation* of a line ℓ , denoted by $\phi(\ell)$, to be the counterclockwise angle that ℓ makes with the positive y -axis. If the separation angle of P_1 and P_2 is at least $\pi/6$, then there must be a line ℓ separating P_1 from P_2 that does not contain any point from P and such that $\phi(\ell) = j \cdot \pi/7$ for some $j \in \{0, 1, \dots, 6\}$. For each of these seven orientations we can compute the best partition in $O(n \log n)$ time, as explained next.

Without loss of generality, consider separating lines ℓ with $\phi(\ell) = 0$, that is, vertical separating lines. Let X be the set of all x -coordinates of the points in P . For any x -value $x \in X$ define $P_1(x) := \{p \in P \mid p_x \leq x\}$, where p_x denotes the x -coordinate of a point p , and define $P_2(x) := P \setminus P_1(x)$. Our task is to find the best partition of the form $(P_1(x), P_2(x))$ over all $x \in X$. To this end we first compute the

values $\text{per}(P_1(x))$ for all $x \in X$ in $O(n \log n)$ time in total, as follows. We compute the lengths of the upper hulls of the point sets $P_1(x)$, for all $x \in X$, using Graham's scan [9], and we compute the lengths of the lower hulls in a second scan. (Graham's scan goes over the points from left to right and maintains the upper (or lower) hull of the encountered points; it is trivial to extend the algorithm so that it also maintains the length of the hull.) By combining the lengths of the upper and lower hulls, we get the values $\text{per}(P_1(x))$.

Computing the values $\text{per}(P_2(x))$ can be done similarly, after which we can easily find the best partition of the form $(P_1(x), P_2(x))$ in $O(n)$ time. Thus the best partition with large separation angle can be found in $O(n \log n)$ time.

2.2.2 The Best Partition with Large Separation Distance

Next we show how to compute the best partition with large separation distance. We assume without loss of generality that $\text{per}(P_2) \leq \text{per}(P_1)$. It will be convenient to treat the case where P_2 is a singleton separately.

Lemma 2.7 *The point $p \in P$ minimizing $\text{per}(P \setminus \{p\})$ can be computed using $O(n \log n)$ time.*

Proof The point p we are looking for must be a vertex of $\text{CH}(P)$. First we compute $\text{CH}(P)$ in $O(n \log n)$ time [9]. Let $v_0, v_1, \dots, v_{m-1}$ denote the vertices of $\text{CH}(P)$ in counterclockwise order. Let Δ_i be the triangle with vertices $v_{i-1}v_iv_{i+1}$ (with indices taken modulo m) and let P_i denote the set of points lying inside Δ_i , excluding v_i but including v_{i-1} and v_{i+1} . Note that any point $p \in P$ is present in at most two sets P_i . Hence, $\sum_{i=0}^{m-1} |P_i| = O(n)$. It is not hard to compute the sets P_i in $O(n \log n)$ time in total. After doing so, we compute all convex hulls $\text{CH}(P_i)$ in $O(n \log n)$ time in total. Since

$$\text{per}(P \setminus \{v_i\}) = \text{per}(P) - |v_{i-1}v_i| - |v_iv_{i+1}| + \text{per}(P_i) - |v_{i-1}v_{i+1}|,$$

we can now find the point p minimizing $\text{per}(P \setminus \{p\})$ in $O(n)$ time. $\square$

It remains to compute the best partition (P_1, P_2) with $\text{per}(P_2) \leq \text{per}(P_1)$ whose separation distance is at least $c_{\text{sep}} \cdot \text{per}(P_2)$ and where P_2 is not a singleton. Let (P_1^*, P_2^*) denote this partition. Define the *size* of a square³ σ to be its edge length. A square σ is a *good square* if (i) $P_2^* \subset \sigma$, and (ii) $\text{size}(\sigma) \leq c^* \cdot \text{per}(P_2^*)$, where $c^* := 18$. Our algorithm globally works as follows.

1. Compute a set S of $O(n)$ squares such that S contains a good square.
2. For each square $\sigma \in S$, construct a set H_σ of $O(1)$ halfplanes such that the following holds: if $\sigma \in S$ is a good square then there is a halfplane $h \in H_\sigma$ such that $P_2^* = P(\sigma \cap h)$, where $P(\sigma \cap h) := P \cap (\sigma \cap h)$.

³ Whenever we speak of squares, we always mean axis-parallel squares.

3. For each pair (σ, h) with $\sigma \in S$ and $h \in H_\sigma$, compute $\text{per}(P \setminus P(\sigma \cap h)) + \text{per}(P(\sigma \cap h))$, and report the partition $(P \setminus P(\sigma \cap h), P(\sigma \cap h))$ that gives the smallest sum.

Step 1: Finding a good square. To find a set S that contains a good square, we first construct a set S_{base} of so-called *base squares*. The set S will then be obtained by expanding the base squares appropriately.

We define a base square σ to be *good* if (i) σ contains at least one point from P_2^* , and (ii) $c_1 \cdot \text{diam}(P_2^*) \leq \text{size}(\sigma) \leq c_2 \cdot \text{diam}(P_2^*)$, where $c_1 := 1/4$ and $c_2 := 4$ and $\text{diam}(P_2^*)$ denotes the diameter of P_2^* . Note that $2 \cdot \text{diam}(P_2^*) \leq \text{per}(P_2^*) \leq 4 \cdot \text{diam}(P_2^*)$. For a square σ , define $\bar{\sigma}$ to be the square with the same center as σ and whose size is $(1 + 2/c_1) \cdot \text{size}(\sigma)$.

Lemma 2.8 *If σ is a good base square then $\bar{\sigma}$ is a good square.*

Proof The distance from any point in σ to the boundary of $\bar{\sigma}$ is at least

$$\frac{\text{size}(\bar{\sigma}) - \text{size}(\sigma)}{2} \geq \text{diam}(P_2^*).$$

Since σ contains a point from P_2^* , it follows that $P_2^* \subset \bar{\sigma}$. Since $\text{size}(\sigma) \leq c_2 \cdot \text{diam}(P_2^*)$, we have

$$\text{size}(\bar{\sigma}) \leq (2/c_1 + 1) \cdot c_2 \cdot \text{diam}(P_2^*) = 36 \cdot \text{diam}(P_2^*) \leq c^* \cdot \text{per}(P_2^*). \quad \square$$

To obtain S it thus suffices to construct a set S_{base} that contains a good base square. To this end we first build a compressed quadtree for P . For completeness we briefly review the definition of compressed quadtrees; see also Fig. 4 (left).

Assume without loss of generality that P lies in the interior of the unit square $U := [0, 1]^2$. Define a *canonical square* to be any square that can be obtained by subdividing U recursively into quadrants. A *compressed quadtree* [13] for P is a hierarchical subdivision of U , defined as follows. In a generic step of the recursive process we are given a canonical square σ and the set $P(\sigma) := P \cap \sigma$ of points inside σ . (Initially $\sigma = U$ and $P(\sigma) = P$.)

- If $|P(\sigma)| \leq 1$ then the recursive process stops and σ is a square in the final subdivision.
- Otherwise there are two cases. Consider the four quadrants of σ . The first case is that at least two of these quadrants contain points from $P(\sigma)$. (We consider the quadrants to be closed on the left and bottom side, and open on the right and top side, so a point is contained in a unique quadrant.) In this case we partition σ into its four quadrants—we call this a *quadtree split*—and recurse on each quadrant. The second case is that all points from $P(\sigma)$ lie inside the same quadrant. In this case we compute the smallest canonical square, σ' , that contains $P(\sigma)$ and we partition σ into two regions: the square σ' and the so-called *donut region* $\sigma \setminus \sigma'$. We call this a *shrinking step*. After a shrinking step we only recurse on the square σ' , not on the donut region.

A compressed quadtree for a set of n points can be computed in $O(n \log n)$ time in the appropriate model of computation⁴ [13]. The idea is now as follows. Let $p, p' \in P_2^*$ be a pair of points defining $\text{diam}(P_2^*)$. The compressed quadtree hopefully allows us to zoom in until we have a square in the compressed quadtree that contains p or p' and whose size is roughly equal to $|pp'|$. Such a square will be then a good base square. Unfortunately this does not always work since p and p' can be separated too early. We therefore have to proceed more carefully: we need to add five types of base squares to S_{base} , as explained next and illustrated in Fig. 4 (right).

- (B1) Any square σ that is generated during the recursive construction—note that this not only refers to squares in the final subdivision—is put into S_{base} .
- (B2) For each point $p \in P$ we add a square σ_p to S_{base} , as follows. Let σ be the square of the final subdivision that contains p . Then σ_p is a smallest square that contains p and that shares a corner with σ .
- (B3) For each square σ that results from a shrinking step we add an extra square σ' to S_{base} , where σ' is the smallest square that contains σ and that shares a corner with the parent square of σ .
- (B4) For any two regions in the final subdivision that touch each other—we also consider two regions to touch if they only share a vertex—we add at most one square to S_{base} , as follows. If one of the regions is an empty square, we do not add anything for this pair. Otherwise we have three cases.
  - (B4.1) If both regions are non-empty squares containing single points p and p' , respectively, then we add a smallest enclosing square for the pair of points p, p' to S_{base} .
  - (B4.2) If both regions are donut regions, say $\sigma_1 \setminus \sigma'_1$ and $\sigma_2 \setminus \sigma'_2$, then we add a smallest enclosing square for the pair σ'_1, σ'_2 to S_{base} .
  - (B4.3) If one region is a non-empty square containing a single point p and the other is a donut region $\sigma \setminus \sigma'$, then we add a smallest enclosing square for the pair p, σ' to S_{base} .

Lemma 2.9 *The set S_{base} has size $O(n)$ and contains a good base square. Furthermore, S_{base} can be computed in $O(n \log n)$ time.*

Proof A compressed quadtree has size $O(n)$ so we have $O(n)$ base squares of type (B1) and (B3). Obviously there are $O(n)$ base squares of type (B2). Finally, the number of pairs of final regions that touch is $O(n)$ —this follows because we have a planar rectilinear subdivision of total complexity $O(n)$ —and so the number of base squares of type (B4) is $O(n)$ as well. The fact that we can compute S_{base} in $O(n \log n)$ time follows directly from the fact that we can compute the compressed quadtree in $O(n \log n)$ time [13].

It remains to prove that S_{base} contains a good base square. We call a square σ *too small* when $\text{size}(\sigma) < c_1 \cdot \text{diam}(P_2^*)$ and *too large* when $\text{size}(\sigma) > c_2 \cdot \text{diam}(P_2^*)$; otherwise we say that σ has the *correct size*. Let $p, p' \in P_2^*$ be two points with $|pp'| = \text{diam}(P_2^*)$, and consider a smallest square $\sigma_{p,p'}$, in the compressed quadtree

⁴ In particular we need to be able to compute the smallest canonical square containing two given points in $O(1)$ time. See the book by Har-Peled [13] for a discussion.

Fig. 4 A compressed quadtree and some of the base squares generated from it. In the right figure, only the points are shown that are relevant for the shown base squares

that contains both p and p' . Note that $\sigma_{p,p'}$ cannot be too small, since $c_1 = 1/4 < 1/\sqrt{2}$. If $\sigma_{p,p'}$ has the correct size, then we are done since it is a good base square of type (B1). So now suppose $\sigma_{p,p'}$ is too large.

Let $\sigma_0, \sigma_1, \dots, \sigma_k$ be the sequence of squares in the recursive subdivision of $\sigma_{p,p'}$ that contain p ; thus $\sigma_0 = \sigma_{p,p'}$ and σ_k is a square in the final subdivision. Define $\sigma'_0, \sigma'_1, \dots, \sigma'_{k'}$ similarly, but now for p' instead of p . Suppose that none of these squares has the correct size—otherwise we have a good base square of type (B1). There are three cases.

- *Case (i): σ_k and $\sigma'_{k'}$ are too large.*

We claim that σ_k touches $\sigma'_{k'}$. To see this, assume without loss of generality that $\text{size}(\sigma_k) \leq \text{size}(\sigma'_{k'})$. If σ_k does not touch $\sigma'_{k'}$, then $|pp'| \geq \text{size}(\sigma_k)$, which contradicts the assumption that σ_k is too large.

Hence, σ_k indeed touches $\sigma'_{k'}$. But then we have a base square of type (B4.1) for the pair p, p' and since $|pp'| = \text{diam}(P_2^*)$ this is a good base square.

- *Case (ii): σ_k and $\sigma'_{k'}$ are too small.*

In this case there are indices $0 < j \leq k$ and $0 < j' \leq k'$ such that σ_{j-1} and $\sigma'_{j'-1}$ are too large and σ_j and $\sigma'_{j'}$ are too small. Note that this implies that both σ_j and $\sigma'_{j'}$ result from a shrinking step, because $c_1 < c_2/2$ and so the quadrants of a too-large square cannot be too small. We claim that σ_{j-1} touches $\sigma'_{j'-1}$. Indeed, similarly to Case (i), if σ_{j-1} and $\sigma'_{j'-1}$ do not touch then $|pp'| > \min(\text{size}(\sigma_{j-1}), \text{size}(\sigma'_{j'-1}))$, contradicting the assumption that both σ_{j-1} and $\sigma'_{j'-1}$ are too large. We now have two subcases.

- The first subcase is that the donut region $\sigma_{j-1} \setminus \sigma_j$ touches the donut region $\sigma'_{j'-1} \setminus \sigma'_{j'}$. Thus a smallest enclosing square for σ_j and $\sigma'_{j'}$ has been put into S_{base} as a base square of type (B4.2). Let σ^* denote this square. Since the segment pp' is contained in σ^* we have

$$c_1 \cdot \text{diam}(P_2^*) < \text{diam}(P_2^*)/\sqrt{2} = |pp'|/\sqrt{2} \leq \text{size}(\sigma^*).$$

Furthermore, since σ_j and $\sigma'_{j'}$ are too small we have

$$\begin{aligned} \text{size}(\sigma^*) &\leq \text{size}(\sigma_j) + \text{size}(\sigma'_{j'}) + |pp'| \leq 3 \cdot \text{diam}(P_2^*) \\ &< c_2 \cdot \text{diam}(P_2^*), \end{aligned} \quad (11)$$

and so σ^* is a good base square.

- The second subcase is that $\sigma_{j-1} \setminus \sigma_j$ does not touch $\sigma'_{j'-1} \setminus \sigma_{j'}$. This can only happen if σ_{j-1} and $\sigma'_{j'-1}$ just share a single corner, v . Observe that σ_j must lie in the quadrant of σ_{j-1} that has v as a corner, otherwise $|pp'| \geq \text{size}(\sigma_{j-1})/2$ and σ_{j-1} would not be too large. Similarly, $\sigma'_{j'}$ must lie in the quadrant of $\sigma'_{j'-1}$ that has v as a corner. Thus the base squares of type (B3) for σ_j and $\sigma'_{j'}$ both have v as a corner. Take the largest of these two base squares, say σ_j . For this square σ^* we have

$$c_1 \cdot \text{diam}(P_2^*) < \text{diam}(P_2^*)/2\sqrt{2} = |pp'|/2\sqrt{2} \leq \text{size}(\sigma^*),$$

since $|pp'|$ is contained in a square of twice the size of σ^* . Furthermore, since σ_j is too small and $|pv| < |pp'|$ we have

$$\text{size}(\sigma^*) \leq \text{size}(\sigma_j) + |pv| \leq (c_1 + 1) \cdot \text{diam}(P_2^*) < c_2 \cdot \text{diam}(P_2^*). \quad (12)$$

Hence, σ^* is a good base square.

- *Case (iii): neither (i) nor (ii) applies.*

In this case σ_k is too small and $\sigma'_{k'}$ is too large (or vice versa). Thus there must be an index $0 < j \leq k$ such that σ_{j-1} is too large and σ_j is too small. We can now follow a similar reasoning as in Case (ii): First we argue that σ_j must have resulted from a shrinking step and that σ_{j-1} touches $\sigma'_{k'}$. Then we distinguish two subcases, namely where the donut region $\sigma_j \setminus \sigma_{j-1}$ touches $\sigma'_{k'}$ and where it does not touch $\sigma'_{k'}$. The arguments for the two subcases are similar to the subcases in Case (ii), with the following modifications. In the first subcase we use base squares of type (B4.3) and in (11) the term $\text{size}(\sigma'_{j'})$ disappears; in the second subcase we use a type (B3) base square for σ_j and a type (B2) base square for p' , and when the base square for p' is larger than the base square for σ_j then (12) becomes $\text{size}(\sigma^*) \leq 2|p'v| < c_2 \cdot \text{diam}(P_2^*)$. $\square$

Step 2: Generating halfplanes. Consider a good square $\sigma \in S$. Let Q_σ be a set of $4 \cdot c^*/c_{\text{sep}} + 1 = 18001$ points placed equidistantly around the boundary of σ . Note that the distance between two neighbouring points in Q_σ is less than $c_{\text{sep}}/c^* \cdot \text{size}(\sigma)$. For each pair q_1, q_2 of points in Q_σ , add to H_σ the two halfplanes defined by the line through q_1 and q_2 .

Lemma 2.10 *For any good square $\sigma \in S$, there is a halfplane $h \in H_\sigma$ such that $P_2^* = P(\sigma \cap h)$.*

First a remark: We do not claim that the line ℓ bounding the halfplane h separates P_1^* and P_2^* globally, but only in σ —indeed, ℓ might intersect $\text{CH}(P_1^*)$.

Proof In the case where $\sigma \cap P_1^* = \emptyset$, two points in Q_σ from the same edge of σ define a halfplane h such that $P_2^* = P(\sigma \cap h)$, so assume that σ contains one or more points from P_1^* .

We know that the separation distance between P_1^* and P_2^* is at least $c_{\text{sep}} \cdot \text{per}(P_2^*)$. Moreover, $\text{size}(\sigma) \leq c^* \cdot \text{per}(P_2^*)$. Hence, there is an empty open strip O with a width of at least $c_{\text{sep}}/c^* \cdot \text{size}(\sigma)$ separating P_2^* from P_1^* . Since σ contains a point from P_1^* , we know that $\sigma \setminus O$ consists of two pieces and that the part of the boundary of σ inside O consists of two disjoint portions B_1 and B_2 each of length at least $c_{\text{sep}}/c^* \cdot \text{size}(\sigma)$. Hence the sets $B_1 \cap Q_\sigma$ and $B_2 \cap Q_\sigma$ contain points q_1 and q_2 , respectively, that define a halfplane h as desired. $\square$

Step 3: Evaluating candidate solutions. In this step we need to compute for each pair (σ, h) with $\sigma \in S$ and $h \in H_\sigma$, the value $\text{per}(P \setminus P(\sigma \cap h)) + \text{per}(P(\sigma \cap h))$. Given a set $\mathcal{O}$ of k orientations, Oh and Ahn [18] described how to create a data structure using $O(nk^3 \log^2 n)$ time and space to answer queries of the following type in time $O(k \log^2 n)$: Given a convex polygon Q where each edge has an orientation in $\mathcal{O}$, what is $\text{per}(P \cap Q)$? In our case, we need to compute the perimeter of the points in canonical 5-gons and their complements, i.e., $\text{per}(P(\sigma \cap h))$ and $\text{per}(P \setminus P(\sigma \cap h))$ for a given pair (σ, h) . Recall that the bounding lines of the halfplanes h we must process have $O(1)$ different orientations. For each such orientation o , we make an instance of the data structure of Oh and Ahn which has as orientations $\mathcal{O}$ the two axis-parallel directions and o . We can then clearly compute $\text{per}(P(\sigma \cap h))$ in time $O(\log^2 n)$. Note that the complement $P \setminus P(\sigma \cap h)$ is the disjoint union of the points in four axis-parallel rectangles and the complementary canonical 5-gon $\sigma \setminus (\sigma \cap h)$. For each of these four rectangles and the 5-gon, we can compute the convex hull of the points inside it in $O(\log^2 n)$ time, using the data structure of Oh and Ahn [18]. This gives us five convex hulls, represented as balanced trees. We can then compute, for each pair of convex hulls, the outer common tangents in $O(\log n)$ time [18, Lemma 3], from which we can compute the overall convex hull and its perimeter. The total time to compute $\text{per}(P \setminus (P(\sigma \cap h)))$ is thus likewise $O(\log^2 n)$.

We thus obtain the following result, which finishes the proof of Theorem 2.6.

Lemma 2.11 *Step 3 can be performed in $O(n \log^2 n)$ time and space.*

3 The Approximation Algorithm

Theorem 3.1 *Let P be a set of n points in the plane and let (P_1^*, P_2^*) be a partition of P minimizing $\text{per}(P_1^*) + \text{per}(P_2^*)$. Suppose we have an exact algorithm for the minimum perimeter-sum problem running in $T(k)$ time for instances with k points. Then for any given $\varepsilon > 0$ we can compute a partition (P_1, P_2) of P such that $\text{per}(P_1) + \text{per}(P_2) \leq (1 + \varepsilon) \cdot (\text{per}(P_1^*) + \text{per}(P_2^*))$ in $O(n + T(1/\varepsilon^2))$ time.*

Proof Consider the axis-parallel bounding box B of P . Let w be the width of B and let h be its height. Assume without loss of generality that $w \geq h$. Our algorithm works in two steps.

- *Step 1: Check if $\text{per}(P_1^*) + \text{per}(P_2^*) \leq w/16$. If so, compute the exact solution.*
 We partition B vertically into four strips with width $w/4$, denoted B_1, B_2, B_3 , and B_4 from left to right. If B_2 or B_3 contains a point from P , we have $\text{per}(P_1^*) + \text{per}(P_2^*) \geq w/2 > w/16$ and we go to Step 2. If B_2 and B_3 are both empty, we consider two cases.
  - *Case (i): $h \leq w/8$.*
 In this case we simply return the partition $(P \cap B_1, P \cap B_4)$. To see that this is optimal, we first note that any subset $P' \subset P$ that contains a point from B_1 as well as a point from B_4 has $\text{per}(P') \geq 2 \cdot (3w/4) = 3w/2$. On the other hand, $\text{per}(P \cap B_1) + \text{per}(P \cap B_4) \leq 2 \cdot (w/2 + 2h) \leq 3w/2$.
  - *Case (ii): $h > w/8$.*
 We partition B horizontally into four rows with height $h/4$, numbered R_1, R_2, R_3 , and R_4 from bottom to top. If R_2 or R_3 contains a point from P , we have $\text{per}(P_1^*) + \text{per}(P_2^*) \geq h/2 > w/16$, and we go to Step 2. If R_2 and R_3 are both empty, we overlay the vertical and the horizontal partitioning of B to get a 4×4 grid of cells $C_{ij} := B_i \cap R_j$ for $i, j \in \{1, \dots, 4\}$. We know that only the corner cells $C_{11}, C_{14}, C_{41}, C_{44}$ contain points from P . If three or four corner cells are non-empty, $\text{per}(P_1^*) + \text{per}(P_2^*) \geq 6h/4 > w/16$, and we go to Step 2. Hence, we may without loss of generality assume that any point of P is in C_{11} or C_{44} . We now return the partition $(P \cap C_{11}, P \cap C_{44})$, which is easily seen to be optimal.
- *Step 2: Handle the case where $\text{per}(P_1^*) + \text{per}(P_2^*) > w/16$.*
 The idea is to compute a subset $\hat{P} \subset P$ of size $O(1/\varepsilon^2)$ such that an exact solution to the minimum perimeter-sum problem on $\hat{P}$ can be used to obtain a $(1 + \varepsilon)$ -approximation for the problem on P .
 We subdivide B into $O(1/\varepsilon^2)$ rectangular cells of width and height at most $c := \varepsilon w / (64\pi\sqrt{2})$. For each cell C where $P \cap C$ is non-empty we pick an arbitrary point in $P \cap C$, and we let $\hat{P}$ be the set of selected points. For a point $p \in \hat{P}$, let $C(p)$ be the cell containing p . Intuitively, each point $p \in \hat{P}$ represents all the points $P \cap C(p)$. Let $(\hat{P}_1, \hat{P}_2)$ be a partition of $\hat{P}$ that minimizes $\text{per}(\hat{P}_1) + \text{per}(\hat{P}_2)$. We assume we have an algorithm that can compute such an optimal partition in $T(|\hat{P}|)$ time. For $i = 1, 2$, define

$$P_i := \bigcup_{p \in \hat{P}_i} P \cap C(p).$$

Our approximation algorithm returns the partition (P_1, P_2) . (Note that the convex hulls of P_1 and P_2 are not necessarily disjoint.) It remains to prove the approximation ratio.

First, note that $\text{per}(\hat{P}_1) + \text{per}(\hat{P}_2) \leq \text{per}(P_1^*) + \text{per}(P_2^*)$ since $\hat{P} \subseteq P$. For $i = 1, 2$, let $\tilde{P}_i$ consist of all points in the plane (not only points in P) within a distance of at

Fig. 5 The crossed points are the points of $\hat{P}$. The left gray region is $\tilde{P}_1$ and the right gray region is $\tilde{P}_2$. The left dashed polygon is the convex hull of P_1 and the right dashed polygon is the convex hull of P_2

most $c\sqrt{2}$ from $\text{CH}(\hat{P}_i)$. In other words, $\tilde{P}_i$ is the Minkowski sum of $\text{CH}(\hat{P}_i)$ with a disk D of radius $c\sqrt{2}$ centered at the origin; see Fig. 5. Note that if $p \in \hat{P}_i$, then $q \in \tilde{P}_i$ for any $q \in P \cap C(p)$, since any two points in $C(p)$ are at most $c\sqrt{2}$ apart from each other. Therefore $P_i \subset \tilde{P}_i$ and hence $\text{per}(P_i) \leq \text{per}(\tilde{P}_i)$. Note also that $\text{per}(\tilde{P}_i) = \text{per}(\hat{P}_i) + 2c\pi\sqrt{2}$. These observations yield

$$\begin{aligned} \text{per}(P_1) + \text{per}(P_2) &\leq \text{per}(\tilde{P}_1) + \text{per}(\tilde{P}_2) \\ &= \text{per}(\hat{P}_1) + \text{per}(\hat{P}_2) + 4c\pi\sqrt{2} \\ &\leq \text{per}(P_1^*) + \text{per}(P_2^*) + 4c\pi\sqrt{2} \\ &= \text{per}(P_1^*) + \text{per}(P_2^*) + 4\pi\sqrt{2} \cdot (\varepsilon w / (64\pi\sqrt{2})) \\ &\leq \text{per}(P_1^*) + \text{per}(P_2^*) + \varepsilon w / 16 \\ &\leq (1 + \varepsilon) \cdot (\text{per}(P_1^*) + \text{per}(P_2^*)). \end{aligned}$$

As all the steps can be done in linear time, the time complexity of the algorithm is $O(n + T(n_\varepsilon))$ for some $n_\varepsilon = O(1/\varepsilon^2)$. $\square$

4 Concluding Remarks

We note that in the exact algorithm, for each of the $O(n)$ base squares $\sigma \in S$, the number of values $\text{per}(P(\sigma \cap h))$ that we query is approximately $2 \cdot \binom{4}{2} \cdot 4500^2 \approx 2.4 \cdot 10^8$. Although it is surely possible to modify the algorithm to get smaller constants (to which we made no attempt), we expect that the algorithm will remain impractical.

Consider the degenerate case where all the input points P are on the x -axis. Then the minimum-perimeter sum problem reduces to the well-known maximum-gap problem, where the goal is to find the largest difference between two consecutive numbers in sorted order. Lee and Wu [16] gave a lower bound of $\Omega(n \log n)$ for that problem in the algebraic computation tree model, which therefore also holds for the minimum-perimeter sum problem in that model.

The question by Mitchell and Wynters [17] about the existence of sub-quadratic algorithms for the minimum-perimeter maximum, minimum-area sum, and minimum-

area maximum problems remain interesting open problems. To our knowledge, the only published algorithm for any of these problems is the $O(n^4 \log n)$ -time algorithm by Bae et al. [6] for the minimum-area sum problem, since the algorithms by Mitchell and Wynters consider line partitions only.

Acknowledgements This research was initiated when the first author visited the Department of Computer Science at TU Eindhoven during the winter 2015–2016. He wishes to express his gratitude to the other authors and the department for their hospitality.

Funding M. Abrahamsen is supported by the Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research under the Sapere Aude research career programme. M. de Berg, K. Buchin, M. Mehr, and A. D. Mehrabi are supported by the Netherlands' Organisation for Scientific Research (NWO) under Project No. 024.002.003, 612.001.207, 022.005025, and 612.001.118 respectively.

References

1. Abrahamsen, M., Adamaszek, A., Bringmann, K., Cohen-Addad, V., Mehr, M., Rotenberg, E., Roytman, A., Thorup, M.: Fast fencing. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC'18), pp. 564–573. ACM, New York (2018)
2. Abrahamsen, M., de Berg, M., Buchin, K., Mehr, M., Mehrabi, A.D.: Minimum perimeter-sum partitions in the plane. In: Aronov, B., Katz, M.J. (eds.) Proceedings of the 33rd International Symposium on Computational Geometry (SoCG'17), Article No. 4. LIPIcs. Leibniz International Proceedings in Informatics, vol. 77. Schloss Dagstuhl, Leibniz-Zentrum für Informatik, Wadern (2017)
3. Agarwal, P.K., Sharir, M.: Efficient algorithms for geometric optimization. *ACM Comput. Surv.* **30**(4), 412–458 (1998)
4. Arkin, E.M., Khuller, S., Mitchell, J.S.B.: Geometric knapsack problems. *Algorithmica* **10**(5), 399–427 (1993)
5. Asano, T., Bhattacharya, B., Keil, M., Yao, F.: Clustering algorithms based on minimum and maximum spanning trees. In: Proceedings of the 4th Annual Symposium on Computational Geometry (SoCG'88), pp. 252–257. ACM, New York (1988)
6. Bae, S.W., Cho, H.-G., Evans, W., Saeedi, N., Shin, C.-S.: Covering points with convex sets of minimum size. *Theor. Comput. Sci.* **718**, 14–23 (2018)
7. Capriles, V., Rote, G., Woeginger, G.: Geometric clusterings. *J. Algorithms* **12**(2), 341–356 (1991)
8. Chan, T.M.: More planar two-center algorithms. *Comput. Geom.* **13**(2), 189–198 (1999)
9. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.: *Computational Geometry: Algorithms and Applications*, 3rd edn. Springer, Berlin (2008)
10. Devillers, O., Katz, M.J.: Optimal line bipartitions of point sets. *Int. J. Comput. Geom. Appl.* **9**(1), 39–51 (1999)
11. Drezner, Z.: The planar two-center and two-median problems. *Transp. Sci.* **18**(4), 351–361 (1984)
12. Eppstein, D.: Faster construction of planar two-centers. In: Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'17), pp. 131–138. ACM, New York (1997)
13. Har-Peled, S.: *Geometric Approximation Algorithms*. Mathematical Surveys and Monographs, vol. 173. American Mathematical Society, Providence (2011)
14. Hershberger, J.: Minimizing the sum of diameters efficiently. *Comput. Geom.* **2**(2), 111–118 (1992)
15. Jaromczyk, J.W., Kowaluk, M.: An efficient algorithm for the Euclidean two-center problem. In: Proceedings of the 10th ACM Symposium on Computational Geometry (SoCG'94), pp. 303–311. ACM, New York (1994)
16. Lee, D.T., Wu, Y.F.: Geometric complexity of some location problems. *Algorithmica* **1**(2), 193–211 (1986)
17. Mitchell, J.S.B., Wynters, E.L.: Finding optimal bipartitions of points and polygons. In: Dehne, F., et al. (eds.) *Algorithms and Data Structures (WADS'91)*. Lecture Notes in Computer Science, vol. 519, pp. 202–213. Springer, Berlin (1991). <http://www.ams.sunysb.edu/~jsbm/>
18. Oh, E., Ahn, H.-K.: Polygon queries for convex hulls of points. In: Wang, L., Zhu, D. (eds.) *Computing and Combinatorics (COCOON'18)*. Lecture Notes in Computer Science, vol. 10976, pp. 143–155. Springer, Cham (2018)

19. Rokne, J., Wang, S., Wu, X.: Optimal bipartitions of point sets. In: Proceedings of the 4th Canadian Conference on Computational Geometry (CCCG'92), pp. 11–16. Memorial University of Newfoundland, St. John's (1992)
20. Segal, M.: Lower bounds for covering problems. *J. Math. Model. Algorithms* **1**(1), 17–29 (2002)
21. Sharir, M.: A near-linear algorithm for the planar 2-center problem. *Discrete Comput. Geom.* **18**(2), 125–134 (1997)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.