Separating Balls with a Hyperplane

Michael Hoffmann* Vincent Kusters* Tillmann Miltzow

Abstract

Let $\mathcal{D}$ be a set of n pairwise disjoint unit balls in $\mathbb{R}^d$ and P the set of their center points. A hyperplane $\mathcal{H}$ is an *m-separator* for $\mathcal{D}$ if each closed halfspace bounded by $\mathcal{H}$ contains at least m points from P. This generalizes the notion of halving hyperplanes, which correspond to $n/2$-separators. The analogous notion for point sets has been well studied. Separators have various applications, for instance, in divide-and-conquer schemes. In such a scheme any ball that is intersected by the separating hyperplane may still interact with both sides of the partition. Therefore it is desirable that the separating hyperplane intersects a small number of balls only.

We present two deterministic algorithms to bisect or approximately bisect a given set of disjoint unit balls by a hyperplane: Firstly, we present a simple linear-time algorithm to construct an αn-separator for balls in $\mathbb{R}^d$, for any $0 < \alpha < 1/2$, that intersects at most $cn^{(d-1)/d}$ balls, for some constant c that depends on d and α. The number of intersected balls is best possible up to the constant c. Secondly, we present a near-linear time algorithm to construct an $(n/2 - o(n))$-separator in $\mathbb{R}^d$ that intersects $o(n)$ balls.

1 Introduction

Let $\mathcal{D}$ be a set of n pairwise disjoint unit balls in $\mathbb{R}^d$ and P the set of their center points. A hyperplane $\mathcal{H}$ is an *m-separator* for $\mathcal{D}$ if each closed halfspace bounded by $\mathcal{H}$ contains at least m points from P. This generalizes the notion of halving hyperplanes, which correspond to $n/2$-separators. The analogous notion of separating hyperplanes for point sets has been well studied (see, e.g, [12] for a survey). Separators have various applications, for instance in divide-and-conquer schemes (we discuss some explicit examples below). In such a scheme any ball that is intersected by the separating hyperplane may still interact with both sides of the partition. Therefore it is desirable that the separating hyperplane intersects a small number of balls only.

Alon, Katchalski and Pulleyblank [1] prove that for any set $\mathcal{D}$ in $\mathbb{R}^2$, there exists a direction such that every line with this direction intersects $O(\sqrt{n \log n})$ disks. In particular, this guarantees the existence of a halving line that intersects at most $O(\sqrt{n \log n})$ disks. Löffler and Mulzer [11] observed that this proof gives a randomized linear-time algorithm.

We develop a generic algorithm in $\mathbb{R}^d$ that can be instantiated with different parameters to obtain Theorem 1 and Theorem 2. Note that Theorem 2 improves the separation of the center points (compared to Theorem 1) at the cost of increasing the running time slightly.

Theorem 1 *Given a set $\mathcal{D}$ of n pairwise disjoint unit balls in $\mathbb{R}^d$ and $\alpha \in (0, 1/2)$, one can construct in $O((1-2\alpha)n)$ time a hyperplane $\mathcal{H}$ that intersects $O((n/(1-2\alpha))^{(d-1)/d})$ balls from $\mathcal{D}$ and such that each closed halfspace bounded by $\mathcal{H}$ contains at least αn centers of balls from $\mathcal{D}$. The constants hidden by the asymptotic notation depend on d only.*

Theorem 2 *Given a set $\mathcal{D}$ of n pairwise disjoint unit balls in $\mathbb{R}^d$ and a function $f(n) \in \omega(1) \cap O(\log n)$, one can construct in $O(nf(n))$ time a hyperplane $\mathcal{H}$ such that each closed halfspace bounded by $\mathcal{H}$ contains at least $\frac{n}{2}(1 - 1/f(n)) = \frac{n}{2}(1 - o(1))$ balls from $\mathcal{D}$.*

Related work. Bereg, Dumitrescu and Pach [4, 5] strengthen the initial result of Alon, Katchalski and Pulleyblank slightly by proving that there exists a direction such that any line with this direction has at most $O(\sqrt{n \log n})$ disks *within constant distance*[1]. They use this lemma to prove that one can always move a set of n unit disks from a start to a target configuration in $3n/2 + O(\sqrt{n \log n})$ moves. Their algorithm runs in $O(n^{3/2}(\log n)^{-1/2})$ time, which Theorem 1 improves to $O(n \log n)$.

Held and Mitchell [8] introduced a paradigm for modeling data imprecision where the location of a point in the plane is not known exactly. For each point, however, we are given a unit disk that is guaranteed to contain the point. The authors show that after preprocessing the disks in $O(n \log n)$ time, they can construct a triangulation of the actual point set in linear time. Löffler and Mulzer [11] follow the same model to construct the onion layer of an imprecise point set. They observed that the proof by Alon et al. immediately gives a randomized expected linear-time algorithm in the following fashion. Pick an angle $\beta \in [0, \pi]$ uniformly at random and compute a halving line for the disks with slope β. This

*Partially supported by the ESF EUROCORES programme EuroGIGA, CRP GraDR and the Swiss National Science Foundation, SNF Project 20GG21-134306.

[1]This result also appeared in [14, Lemma 9.3.2].

halving line intersects at most $O(\sqrt{n \log n})$ unit disks with probability at least 1/2. Löffler and Mulzer use this algorithm to compute a (α, β)-*space decomposition tree*: a data structure similar to a binary space partition in which every line is an αk-separator that intersects at most k^β disks. They show that such a $(1/2+\varepsilon, 1/2+\varepsilon)$ space decomposition tree can be computed in $O(n \log n)$ expected time, for every $\varepsilon > 0$. Theorem 1 can be used to improve this to $O(n \log n)$ deterministic time. Next, they describe a deterministic $O(n \log n)$ algorithm to compute a line ℓ such that there are at least $n/2 - cn^{5/6}$ disks completely to each side of ℓ. The algorithm uses an r-partition of the plane [13] to find good candidate lines. Theorem 1 can be used to improve the running time of this algorithm to $O(n)$.

Finally, the question has a continuous counterpart that has been solved recently [7].

2 Algorithm

In this section, we develop a generic algorithm to compute a separator for a given set of pairwise disjoint unit balls in $\mathbb{R}^d$. Using this generic algorithm, we will give two algorithms to compute an approximately halving hyperplane that intersects a sublinear number of balls.

Besides the set $\mathcal{D}$ of n balls in $\mathbb{R}^d$, the generic algorithm has two more parameters. First, a number $b \in \{1, \ldots, n\}$ that quantifies the quality of the approximation: we will show that the hyperplane constructed by the algorithm forms an $(n-b)/2$-separator for $\mathcal{D}$. The main step of the algorithm consists in finding a direction d such that we are guaranteed to find a desired separator that is orthogonal to d. A second parameter $k \in \mathbb{N}$ of the algorithm specifies the number of different directions to generate and test during this step. As a rule of thumb, generating more directions results in a better solution, but the runtime of the algorithm increases proportionally. The algorithm works for certain combinations of these parameters only, as detailed in the following theorem.

Theorem 3 *Given a set $\mathcal{D}$ of n pairwise disjoint unit balls in $\mathbb{R}^d$ and parameters $b \in \{1, \ldots, n\}$ and $k \in \mathbb{N}$ that satisfy the conditions*

$$dn \leq kb \quad \text{and} \quad (1)$$

$$t := \left(\frac{2^{d-1}\, V_d}{d^{(d-2)/2}}\right)^{1/d} \frac{n^{1/d}}{k^{2-1/d}} - 2 \; > \; 2, \qquad (2)$$

(where V_d is the volume of the d-dimensional unit ball), one can construct in $O(kn)$ time a hyperplane $\mathcal{H}$ that intersects at most $2b/(t-2)$ balls from $\mathcal{D}$ and such that each closed halfspace bounded by $\mathcal{H}$ contains at least $(n-b)/2$ centers of balls from $\mathcal{D}$.

Maybe more interesting than Theorem 3 in its full generality are the special cases stated as Theorem 1 and Theorem 2 above. Theorem 1 describes the case that k is constant. It can be obtained by choosing $b = \lfloor(1-2\alpha)n\rfloor$ and $k = \lceil(1-2\alpha)d\rceil$. Theorem 2 describes the case that k is a very slowly growing function $f(n)$. It can be obtained by choosing $b = n/f(n)$ and $k = df(n)$.

Overview of the algorithm. Our algorithm consists of two steps. In the first step, we find a direction d in which the balls from $\mathcal{D}$ are "spread out nicely". More precisely, for an arbitrary (oriented) line ℓ consider the set P of points that results from orthogonally projecting all centers of balls from $\mathcal{D}$ onto ℓ. Denote by $p_1, \ldots, p_n$ the order of points from P sorted along ℓ. We want to find an $(n-b)/2$-separator orthogonal to ℓ. This means that the separating hyperplane $\mathcal{H}$ must intersect ℓ somewhere in between $p_{(n-b)/2}$ and $p_{(n+b)/2}$.

However, we also need to guarantee that not too many points from P are within distance one of $\mathcal{H}$, which may or may not be possible depending on the choice of ℓ. Therefore we try several possible directions/lines and select the first one among them that works. In order to evaluate the quality of a line, we use as a simple criterion the *spread*, defined to be the distance between $p_{(n-b)/2}$ and $p_{(n+b)/2}$. Given a line ℓ with sufficient spread, we can find a suitable $(n-b)/2$-separator orthogonal to ℓ in the second step of our algorithm, as the following lemma demonstrates. Note the safety cushion of width one to the remaining balls of $\mathcal{D}$.

Lemma 4 *Given a set P of b (one-dimensional) points in an interval $[\ell, r]$ of length $w = r - \ell > 2$, we can find in $O(b)$ time a point $p \in (\ell+1, r-1)$ such that at most $2b/(w-2)$ points from P are within distance one of p.*

Proof. We select $\lceil(w-2)/2\rceil$ pairwise disjoint closed sub-intervals of length two in (ℓ, r). By the pigeonhole principle at least one interval contains at most $b/\lceil(w-2)/2\rceil \leq 2b/(w-2)$ points from P. Select p to be the midpoint of such an interval.

Algorithmically, we can find such an interval using a kind of binary search on the intervals: We maintain a set of points and a range of intervals. At each step consider the median interval I and test for every point whether it lies in I, to the left of I, or to the right of I. Then either I contains at most $2b/(w-2)$ points from P and we are done, or we recurse on the side that contains fewer points, after discarding all points and intervals on the other side. The process stops as soon as the current range of intervals contains at most $2b/(w-2)$ points from P, at which point any of the remaining intervals can be chosen. Given that we maintain the ratio between the number of points and the number of intervals, the process terminates with an interval of the desired type. As the number of points decreases by a constant factor in each iteration, the overall number of comparisons can be bounded by a geometric series and the resulting runtime is linear. ☐

How to find a good direction. Our algorithm tries k different directions and stops as soon as it finds a direction with spread at least t (see Theorem 3). For a given direction the spread can be computed in $O(n)$ time using linear time rank selection [6]. In the remainder of this section, we will discuss how to select an appropriate set of directions such that one direction is guaranteed to have spread at least t.

Given hyperplanes $\mathcal{H}_1, \dots, \mathcal{H}_d$ and $w_1, \dots, w_d > 0$, let the *slab* S_i be the set of points with distance $w_d/2$ to $\mathcal{H}_i$. Below we give a formula based on a volume argument to bound the number of balls in $S_1 \cap \dots \cap S_d$. This formula in turn motivates our choice of directions, which we will explain afterwards.

Lemma 5 *Let $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d \in S^{d-1} \subset \mathbb{R}^d$ be linearly independent directions and $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_d$ be hyperplanes with corresponding normal directions. Given $w_1, \dots, w_d > 0$, the number of pairwise disjoint unit balls completely contained in $S_1 \cap \dots \cap S_d$ is upper bounded by*

$$\frac{w_1 \dots w_d}{|\det(\vec{v}_1, \dots, \vec{v}_d)|V_d},$$

where V_d denotes the volume of the d-dimensional unit ball.

Proof. Since each ball has volume V_d and they are pairwise disjoint, it is sufficient to bound the volume of $S := S_1 \cap \dots \cap S_d$. The volume of S depends linearly on $w_1, \dots, w_d$, so we scale them all to one. We can map the linearly independent vectors $(\vec{v}_1, \dots, \vec{v}_d)$ to the standard basis $(e_1, \dots, e_d)$ by multiplying with the matrix $(\vec{v}_1, \dots, \vec{v}_d)^{-1}$. The volume changes by this transformation by a factor of $1/\det(\vec{v}_1, \dots, \vec{v}_d)$. This transformation yields the unit cube, which has volume 1. Reversing the transformations gives the stated bound. ☐

The bound in Lemma 5 depends on the determinant formed by the d direction vectors, which corresponds to the volume of the $(d-1)$-simplex spanned by them. In order to obtain a good upper bound, we must guarantee that this volume does not become too small. Ensuring this reduces to the *Heilbronn Problem*: Given $k \in \mathbb{N}$ and a compact region $P \subset \mathbb{R}^d$ of unit volume, how to select k points from P as to maximize the area of the smallest d-simplex formed by these points? Heilbronn posed this question for $d = 2$ and the natural generalization to higher dimension was studied by Barequet [3] and Lefmann [9]. We use the following simple explicit construction that goes back to Erdős in $\mathbb{R}^2$ and was generalized to higher dimension by Barequet.

Lemma 6 ([3, 15]) *Given a prime k, let $P = \{p_0, \dots, p_{k-1}\} \subset [0,1]^d$ with*

$$p_i = \frac{1}{k}\left(i, i^2 \bmod k, \dots, i^d \bmod k\right).$$

Then the smallest d-simplex spanned by $d+1$ points from P has volume at least $1/(d!k^d)$.

Assuming k to be prime is not a restriction: If k is not prime, then by Bertrand's postulate there is a prime $k' \le 2k$. We can compute k' efficiently, for instance, in $O(k/\log\log k)$ time using Atkin's sieve [2]. In order to obtain the desired direction vectors we proceed as follows: Use Lemma 6 to generate k points $p_0, \dots, p_{k-1}$ in $[0,1]^{d-1}$. Then lift the points to $S^{d-1} \subset \mathbb{R}^d$ using the map

$$f : (x_1, \dots, x_{d-1}) \mapsto \frac{(x_1 - \frac{1}{2}, \dots, x_{d-1} - \frac{1}{2}, \frac{1}{2})}{||(x_1 - \frac{1}{2}, \dots, x_{d-1} - \frac{1}{2}, \frac{1}{2})||}$$

and denote the resulting set of directions by $D = \{\vec{v}_0, \dots, \vec{v}_{k-1}\}$ with $\vec{v}_i = f(p_i)$.

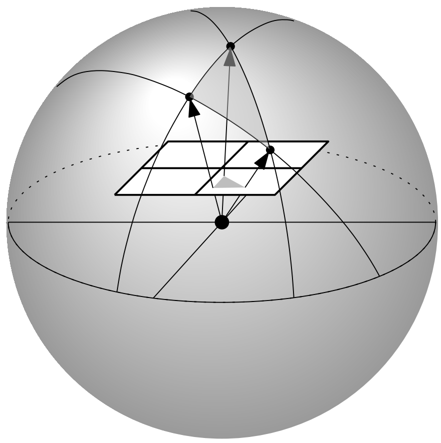

Lemma 7 *For any d vectors $\vec{v}_{i_1}, \dots, \vec{v}_{i_d}$ from D we have $|\det(\vec{v}_{i_1}, \dots, \vec{v}_{i_d})| \ge 2^{d-1}/((d-1)!d^{d/2}k^{d-1})$.*

We only give the proof idea of Lemma 7 here. Lemma 5 states that we need to pick directions such that no d of them are close to collinear. Lemma 6 gives us points that do not form small simplices. When those points are lifted to dimension $d+1$, the size of a spanned simplex corresponds directly to how close the vectors are to collinearity. We are now ready to prove Theorem 3.

Proof. The algorithm goes as follows. Compute directions $\vec{v}_1, \dots, \vec{v}_k$ as in Lemma 7. For each $i \in \{1, \dots, k\}$ consider the sequence of center points of the balls in $\mathcal{D}$, sorted according to direction $\vec{v}_i$. Let S_i be the slab orthogonal to $\vec{v}_i$ bounded by the points of rank $(n-b)/2$ and $(n+b)/2$ on $\vec{v}_i$: then S_i contains the middle b projected center points on $\vec{v}_i$. Denote by w_i the width of S_i in direction $\vec{v}_i$ (which is the spread of $\vec{v}_i$) and by $|S_i|$ the number of center points in S_i. We claim that $w_i \geq t$ for some $i \in \{1, \ldots, k\}$. Intuitively, if all w_i were small, then the slabs would overlap very little (for the directions we chose), which is impossible since each of them contains b center points.

We can bound

$$kb = \sum_{i=1}^{k} |S_i| \leq (d-1)n + \sum_{i_1 < \ldots < i_d} |S_{i_1} \cap \ldots \cap S_{i_d}|,$$

noting that a point that is contained in at most $d-1$ sets S_i is counted $d-1$ times on the right hand side, whereas a point that is contained in $a \geq d$ sets is counted $d - 1 + \binom{a}{d} \geq a$ times.

For the purpose of contradiction assume $w_i < t$ for all $i \in \{1, \ldots, k\}$. Together with Lemma 5 (noting that a center point is in S_i if and only if its ball is completely contained in S_i widened by one on both sides) and Lemma 7 we get $kb = \sum_{i=1}^{k} |S_i|$

$$\begin{aligned}
&\overset{5}{\leq} && (d-1)n + \sum_{i_1 < \ldots < i_d} \frac{(w_{i_1} + 2) \ldots (w_{i_d} + 2)}{|\det(\vec{v}_{i_1}, \ldots, \vec{v}_{i_d})| V_d} \\
&\overset{7}{<} && (d-1)n + \sum_{i_1 < \ldots < i_d} \frac{(t+2)^d}{V_d} \frac{(d-1)! d^{d/2} k^{d-1}}{2^{d-1}} \\
&= && (d-1)n + \binom{k}{d} \frac{(t+2)^d (d-1)! d^{d/2} k^{d-1}}{2^{d-1} V_d} \\
&\leq && (d-1)n + \frac{d^{(d-2)/2}}{2^{d-1} V_d} (t+2)^d k^{2d-1}.
\end{aligned}$$

In combination with Condition (1) we get

$$dn \leq kb < (d-1)n + \frac{d^{(d-2)/2}}{2^{d-1} V_d} (t+2)^d k^{2d-1}$$

and so

$$(t+2)^d > \frac{V_d \; 2^{d-1}}{d^{(d-2)/2}} \frac{n}{k^{2d-1}},$$

in contradiction to the definition of t in Condition (2). Therefore, our assumption $w_i < t$ for all $i \in \{1, \ldots, k\}$, was wrong and there is some $w_j \geq t$.

Using Lemma 4 on the center points from S_j projected to a line in direction $\vec{v}_j$ we obtain a hyperplane $\mathcal{H}$ orthogonal to $\vec{v}_j$ that intersects at most $2b/(w_j - 2) \leq 2b/(t-2)$ balls from $\mathcal{D}$. By Lemma 4 the hyperplane $\mathcal{H}$ has distance greater than one to any ball in $\mathcal{D}$ whose center is not in S_j, and so $\mathcal{H}$ is the desired separator.

Regarding the runtime bound, as stated above we can compute the spread of any direction in $O(n)$ time, which yields $O(kn)$ time for k directions. The second step of finding $\mathcal{H}$ can be done in $O(b) = O(n)$ time by Lemma 4. Therefore the overall runtime is $O(kn)$. $\square$

Acknowledgments. We want to thank Marek Elias, Jiří Matoušek, Edgardo Roldán-Pensado and Zuzana Safernová for interesting discussions and referring us to related work.

References

[1] Noga Alon, Meir Katchalski, and William R. Pulleyblank. Cutting disjoint disks by straight lines. *Discrete & Computational Geometry*, 4:239–243, 1989. 1

[2] A. Oliver L. Atkin and Daniel J. Bernstein. Prime sieves using binary quadratic forms. *Math. Comput.*, 73(246):1023–1030, 2004. 2

[3] Gill Barequet. A lower bound for Heilbronn's triangle problem in d dimensions. *SIAM Journal on Discrete Mathematics*, 14(2):230–236, 2001. 2, 6

[4] Sergey Bereg, Adrian Dumitrescu, and János Pach. Sliding disks in the plane. In *Discrete and Computational Geometry*, pages 37–47. Springer, 2005. 1

[5] Sergey Bereg, Adrian Dumitrescu, and János Pach. Sliding disks in the plane. *International Journal of Computational Geometry & Applications*, 18(05):373–387, 2008. 1

[6] Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, and Robert E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973. 2

[7] Luca Esposito, Vincenzo Ferone, Bernd Kawohl, Carlo Nitsch, and Cristina Trombetti. The longest shortest fence and sharp Poincaré–sobolev inequalities. *Archive for Rational Mechanics and Analysis*, 206(3):821–851, 2012. 1

[8] Martin Held and Joseph SB Mitchell. Triangulating input-constrained planar point sets. *Information Processing Letters*, 109(1):54–56, 2008. 1

[9] Hanno Lefmann. On Heilbronn's problem in higher dimension. *Combinatorica*, 23(4):669–680, 2003. 2

[10] Chi-Yuan Lo, Jiří Matoušek, and William L. Steiger. Algorithms for ham-sandwich cuts. *Discrete & Computational Geometry*, 11:433–452, 1994.

[11] Maarten Löffler and Wolfgang Mulzer. Unions of onions. *CoRR*, abs/1302.5328, 2013. 1, 1

[12] Horst Martini and Anita Schöbel. Median hyperplanes in normed spaces – a survey. *Discrete Applied Mathematics*, 89(1):181–195, 1998. 1

[13] Jiří Matoušek. Efficient partition trees. *Discrete & Computational Geometry*, 8(1):315–334, 1992. 1

[14] János Pach and Micha Sharir. *Combinatorial geometry and its algorithmic applications: The Alcalá lectures*, volume 152 of *Mathematical Surveys and Monographs*. Amer. Math. Soc., 2009. 1

[15] Klaus F. Roth. On a problem of Heilbronn. *J. London Math. Soc.*, 26(3):198–204, 1951. 6

[16] Helge Tverberg. A seperation property of plane convex sets. *Mathematica Scandinavica*, 45:255–260, 1979.