

A Lower Bound for Shallow Partitions

Wolfgang Mulzer*

Daniel Werner†

Abstract

We give a lower bound of $\Omega(\log(n/k)/\log\log(n/k))$ for the crossing number of shallow partitions in the plane.

1 Introduction

Range searching is a fundamental problem in computational geometry that has long driven innovation in the field [3]: given a set of n points in d dimensions, find a data structure such that all points inside a given query range can be found efficiently. Depending on the precise nature of the query range and on the dimension, many different versions of the problem can be studied. Consequently, a wide variety of techniques have been developed to address them. Among these tools we can find such classics as range trees and kd -trees [5, Chapter 5], ε -nets and cuttings [7], spanning trees with small crossing number [13], geometric partitions [9], and many more. For several problems, almost matching lower bounds are known (in certain models of computation) [7].

Geometric partitions provide the most effective means for solving the *simplex range searching* problem, where the query range is given by a d -dimensional simplex [6, 9]. They provide a way to subdivide a point set into parts of roughly equal size, such that (i) each part is contained in a simplex; and (ii) any given hyperplane intersects only few of these simplices. This makes it possible to construct a tree-like data structure in which each node corresponds to a simplex in an appropriate geometric partition. With a careful implementation, one can achieve query time $O(n^{1-1/d} + z)$ with linear space [6] (here z is the output size, i.e., the number of reported points).

If the query simplex degenerates to a half-space, we can do better [10]. For this, we need a more specialized version of geometric partitions, called *shallow partitions*. Again, these partitions provide a way for subdividing a d -dimensional point set into parts of roughly equal size, such that each part is contained in a simplex and such that a hyperplane intersects only few of these simplices. This time, however, we restrict

ourselves to *shallow* hyperplanes. Such hyperplanes have only few points to one side. Thus, we only have the guarantee that any shallow hyperplane will intersect few simplices of the partition (see below for details). This makes it possible to decrease the number of simplices that are intersected and to achieve better bounds for halfspace range searching. Namely, one can obtain for $d \geq 4$ a linear-space data structure that answers a query in time $O(n^{1-1/\lfloor d/2 \rfloor} + z)$, where z is the output size [6] (for $d = 2, 3$, one can achieve query time $O(\log n + z)$ and linear space [2]).

Shallow partitions (as well as their cousins—shallow cuttings) have proved invaluable tools in computational geometry and have found numerous further applications. Nonetheless, there still remain some open questions. As mentioned above, we would like every shallow hyperplane to intersect as few simplices of the shallow partition as possible. But what exactly is possible? For dimension $d \geq 4$, the original bound by Matoušek [9] is known to be asymptotically tight. For lower dimensions, however, Matoušek asked whether his result could be improved. It took almost 20 years until Afshani and Chan [2] provided the first lower bound in three dimensions, almost matching the upper bound. For the plane, however, so far no non-trivial lower bounds appear in the literature.

Here, we will give a construction that provides such a lower bound for shallow partitions in two dimensions. Our result almost matches the upper bound and also gives an alternative proof for the bound of Afshani and Chan [2]. A similar construction has been discovered independently by Afshani [1].

2 Shallow partitions

We begin by providing the details of Matoušek’s shallow partition theorem in two dimensions. Let $P \subseteq \mathbb{R}^2$ be a planar n -point set in general position. Let $k \in \{1, \dots, n\}$ be a parameter. A k -partition \mathcal{P} for P consists of two parts: (i) a sequence $P_1, P_2, \dots, P_{\lfloor n/k \rfloor}$ of pairwise disjoint subsets of P such that $\bigcup_i P_i = P$ and $|P_i| = k$ for $i = 1, \dots, \lfloor n/k \rfloor$; and (ii) a sequence $\Delta_1, \Delta_2, \dots, \Delta_{\lfloor n/k \rfloor}$ of triangles such that $P_i \subseteq \Delta_i$ for all i .

Now let ℓ be a line that does not contain any point in P , and let ℓ^+ denote the open halfplane above ℓ . We say that ℓ is k -shallow if $|\ell^+ \cap P| \leq k$. Given a k -partition \mathcal{P} of P , the *crossing number* of \mathcal{P} is the maximum number of triangles in \mathcal{P} that are intersected by

*Institut für Informatik, Freie Universität Berlin, 14195 Berlin, Germany mulzer@inf.fu-berlin.de

†Institut für Informatik, Freie Universität Berlin, 14195 Berlin, Germany werner@inf.fu-berlin.de. Supported by Deutsche Forschungsgemeinschaft within the Research Training Group (Graduiertenkolleg) “Methods for Discrete Structures”.

any k -shallow line. For any given k , the goal is to find a k -partition of P whose crossing number is as small as possible. Matoušek [10, Theorem 3.1] proved the following theorem.

Theorem 1 *Let P be a planar n -point set in general position and let $k \in \{1, \dots, n\}$. Then there exists a k -partition of P with crossing number $O(\log(n/k))$.*

Matoušek’s original proof uses cuttings and a variant of the iterative reweighting technique (also known as the multiplicative weights update method [4]), and it readily generalizes to higher dimensions. More recently, Har-Peled and Sharir [8, Lemma 3.3] give an approach for proving Theorem 1 with elementary means, but it is not clear whether their technique can be applied to higher dimensions. As mentioned in the introduction, Matoušek [10] asked whether the crossing number in Theorem 1 can be improved to $O(1)$. He conjectured that the answer is no. Afshani and Chan [2] proved that for any k there are arbitrarily large point sets in \mathbb{R}^3 such that the crossing number of any k -partition for them is $\Omega(\frac{\log(n/k)}{\log \log(n/k)})$. However, their construction does not apply for two dimensions. Hence, we will describe here a different—and arguably simpler—construction that yields the same lower bound for the plane. Independently, Afshani [1] used very similar ideas to obtain the same lower bound.

3 The Lower Bound

Let $a(n, k)$ be the minimum crossing number that a k -partition can achieve for any planar n -point set in general position. For the lower bound, we shall consider the dual setting. We use the standard duality transform along the unit paraboloid that maps the point $p : (p_x, p_y)$ to the line $p^* : y = 2p_x x - p_y$ and vice versa [11].

A point set P dualizes to a set P^* of planar lines. We now define the k -level of P^* , $\text{lev}_k(P^*)$ [12]. It is the closure of the set of all points that lie on a line of P^* and that have exactly k lines of P^* beneath them. We observe that $\text{lev}_k(P^*)$ is an x -monotone polygonal curve whose edges and vertices come from the arrangement of P^* . Let \mathcal{C} be the upper convex hull of $\text{lev}_k(P^*)$. For each vertex v of \mathcal{C} , we let $P_v^* \subseteq P^*$ denote the set of lines beneath it. We call P_v^* the *conflict set* of v . We have $|P_v^*| = k$,¹ hence v is dual to a k -shallow line v^* in the primal plane.

Now we can interpret shallow partitions in the dual plane:

¹Note that $\text{lev}_k(P^*)$ may also contain vertices with only $k-1$ lines of P^* beneath them, but these vertices cannot appear on \mathcal{C} , since they correspond to a concave bend in $\text{lev}_k(P^*)$.

Proposition 2 *Let \mathcal{C} be an x -monotone downward convex chain, and let L be a set of n lines such that for each vertex v of \mathcal{C} the conflict set L_v has cardinality k . Then there exists a coloring of L such that (i) each color class has size at most k ; and (ii) each conflict set L_v contains at most $a(n, k) + 1$ different colors.*

Proof. Consider the primal plane, where $L = P^*$ corresponds to a point set P . By assumption, there exists a k -partition \mathcal{P} of P with crossing number $a(n, k)$. Each vertex v of \mathcal{C} corresponds to a k -shallow line v^* , and at most one triangle of \mathcal{P} can be wholly contained in v^{*+} . Thus, the claim follows from the properties of the duality transform. \square

We are now ready to describe the construction. Let $m = 2^\beta$ be a power of 2 and let \mathcal{C} be an x -monotone convex chain with m vertices. We denote these vertices by v_1, \dots, v_m , from left to right. Now, for $j = 0, \dots, \beta$, let L_j be a set of $m/2^j$ lines such that the first line in L_j lies exactly below the vertices v_1 to v_{2^j} , the second line lies below v_{2^j+1} to $v_{2 \cdot 2^j}$, the third line lies below $v_{2 \cdot 2^j+1}$ to $v_{3 \cdot 2^j}$, etc. We set $L' := \bigcup_{j=0}^{\beta} L_j$. See Fig. 1.




Figure 1: Sets of lines L_j .

Assume for now that k is a multiple of $\beta + 1$, and let L consist of $k/(\beta + 1)$ copies of L' . We perturb the lines in L such that they are all distinct while their

relationship with the vertices of \mathcal{C} remains unchanged. It follows that L has exactly $n := (2m - 1)k/(\beta + 1)$ lines, with exactly k lines in each conflict set L_{v_i} (recall that by definition $\beta = \log m$).

By Proposition 2, there is a coloring of L such that each color class has size at most k and such that each conflict set contains at most $a(n, k) + 1$ colors. The structure of L lets us interpret this coloring as follows: let T be a complete binary tree with $2m - 1$ nodes and height β . We label the leaves of T with the vertices v_1, \dots, v_m , from left to right. Thus, every node w of T corresponds to an interval of consecutive vertices of \mathcal{C} , namely the leaves of the subtree rooted in w . By assigning to w the lines that lie exactly below the vertices in this interval, we obtain a partition of L into sets of size $k/(\beta + 1)$. This leads to an interpretation of shallow partitions as multi-colorings of trees.




Figure 2: A tree with height $\beta = \log n$. The leaves correspond to the vertices v_i , and the level of height i corresponds to the lines in L_i . A line ℓ is stored in the node whose subtree corresponds to the vertices that have ℓ below them.

Proposition 3 *Let T be a complete binary tree with height $\beta = \log m$ and $2m - 1$ nodes, and let k be a multiple of $\log m + 1$. Then there exists a multi-coloring of the nodes of T with the following properties: (i) every node is associated with a multiset of $k/(\beta + 1)$ colors; (ii) each color class has at most k elements; (iii) along each root-leaf path there are at most $a(n, k) + 1$ distinct colors, where $n = (2m - 1)k/(\beta + 1)$.*

Proof. Properties (i) and (ii) follow immediately from Proposition 2 and the construction. For Property (iii), observe that the lines encountered along a root-leaf path are exactly the lines below the vertex of \mathcal{C} corresponding to the leaf. \square

We can now prove the desired lower bound.

Lemma 4 *Let T be a complete binary tree with height $\beta = \log m$ and $2m - 1$ nodes, and let k be a multiple of $\log m + 1$. Consider a multi-coloring of T such that (i) every node is associated with a multiset of $k/(\beta + 1)$ colors; and (ii) each color class has at most $2k$ elements. Then there exists a root leaf-path with $\Omega(\log m / \log \log m)$ distinct colors.*

Proof. We subdivide the nodes of T into *slices*. The first slice consists of the first $\lceil \log(3\beta) \rceil$ levels of T , the

second slice consists of the following $\lceil \log(6\beta) \rceil$ levels, the third slice has the next $\lceil \log(9\beta) \rceil$ levels, and so on. In general, the i th slice consists of $\lceil \log(3i\beta) \rceil$ consecutive levels of T .

We claim that there exists a root-leaf path that has at least one distinct color for each slice that it crosses, except for the last one. To see this, we first consider a complete subtree T' of T that has its root in the first level of a slice i and its leaves in the last level of the same slice. As a complete binary tree with $\lceil \log(3i\beta) \rceil$ levels, T' has at least $3i\beta - 1 \geq 2i\beta + 2i$ nodes. Therefore, our multi-coloring needs to assign at least $2(i\beta + i)k/(\beta + 1)$ colors in T' . Since each color class has size at most $2k$, this requires at least i distinct colors.

We now construct the required root-leaf path slice by slice. Throughout, we maintain the invariant that after i slices have been considered, the path contains at least i distinct colors. This is certainly true at the root. Now suppose that we have constructed a partial path Q_{i-1} that ends at a node z in the last level of the $(i - 1)$ th slice. If Q_{i-1} contains at least i distinct colors, we arbitrarily extend it to a path Q_i that ends at the bottom of the i th slice. Otherwise, we pick an arbitrary child z' of z . As noted above, the complete subtree that is rooted at z' and restricted to the i th slice contains at least i distinct colors. Thus, we can extend Q_{i-1} through z' to a path Q_i that goes to the bottom of the i th slice and that meets at least i distinct colors. The claim follows.

It remains to calculate a lower bound for the number of slices b . By construction, we must have

$$\sum_{i=1}^b \lceil \log(3i\beta) \rceil \geq \beta + 1.$$

Now,

$$\begin{aligned} \sum_{i=1}^b \lceil \log(3i\beta) \rceil &\leq \sum_{i=1}^b \log(4i\beta) \\ &\leq b(2 + \log b + \log \beta) \\ &\leq 3b \log \beta, \end{aligned}$$

since clearly $b \leq \beta$. Hence,

$$b \geq \frac{\beta + 1}{3 \log \beta} = \Omega\left(\frac{\log m}{\log \log m}\right),$$

as desired. \square

We now indicate how to drop the assumption that k is a multiple of $\beta + 1$. Indeed, suppose that this is not the case, but $k \geq \beta + 1$. We first perform the above construction with $k' := \lfloor k/(\beta + 1) \rfloor (\beta + 1)$ instead of k . Note that since $k \geq \beta + 1$, we have $k \geq k'$. Then we add $k - k'$ suitably perturbed copies of L_β (the set containing a line in conflict with all vertices of \mathcal{C}).

Let L be the resulting set of lines. By Proposition 2, there exists a coloring of L such that each color class has at most $k \leq 2k'$ elements and such that each conflict set has at most $a(|L|, k) + 1$ distinct colors. The tree T corresponding to L has the same structure as before, but now each non-leaf node except the leaf is associated with $k'/(\beta + 1)$ colors, while the leaves have $k - k'$ additional colors. This suffices for the argument of Lemma 4 to go through.

Theorem 5 *There is a constant $c > 0$ such that the following holds. For every n and $k \in \{\log n, \dots, n/4\}$, there exists a planar n -point set P such that the crossing number for any k -partition of P is at least $c \log(n/k) / \log \log(n/k)$. Thus,*

$$a(n, k) = \Omega\left(\frac{\log(n/k)}{\log \log(n/k)}\right).$$

Proof. Let $\beta \in \mathbb{N}$ be maximum with $(2^{\beta+1} - 1)/(\beta + 1) \leq n/2k$. Set $m := 2^\beta$ and $k' := \lfloor k/(\beta + 1) \rfloor(\beta + 1)$.

From Propositions 2 and 3 and Lemma 4, it follows that by taking the dual we obtain a set P' of $n' := (2m - 1)k'/(\beta + 1) + k - k'$ points such that any k -partition of P' has crossing number at least $c' \log m / \log \log m$, for some constant $c' > 0$.

First note that $\beta < \log n$ and $k \geq \beta + 1$. Hence, $k' \leq k \leq 2k'$ and $k - k' \leq \log n$. Thus, we can conclude that

$$n' = \frac{2m - 1}{\log m + 1} k' + k - k' \leq \frac{n}{2} + \log n \leq n.$$

and

$$n' = \frac{2m - 1}{\log m + 1} k' + k - k' \geq \frac{n}{4k} \cdot \frac{k}{2} = \frac{n}{8}.$$

Thus, by adding at most $7n/8$ points that are contained in no k -shallow halfplane, we can obtain from P' a point set P with n points and crossing number at least $c \log m / \log \log m$. Finally, observe that

$$m \geq \frac{n'}{k'} - k \geq \frac{n}{9k},$$

so P also has crossing number at least $c \cdot \frac{\log(n/k)}{\log \log(n/k)}$, for some $c > 0$. The result follows. \square

Note that our construction also implies a similar lower bound in \mathbb{R}^3 by embedding the plane into three-dimensional space and perturbing the points slightly. This provides an alternative proof of the result by Afshani and Chan [2].

4 Conclusion and Open Problems

We have given a simple construction that give a lower bound of $\Omega\left(\frac{\log(n/k)}{\log \log(n/k)}\right)$ for the crossing number of

any shallow partition of a planar point set. Matoušek's result gives an upper bound of $O(\log(n/k))$. Thus, there still remains a factor of $\log \log(n/k)$ to be settled. Can we show that Matoušek's analysis is tight? Or, perhaps more interestingly, can we construct shallow partitions with crossing number $O\left(\frac{\log(n/k)}{\log \log(n/k)}\right)$?

Acknowledgments

We would like to thank Nabil Mustafa for drawing our attention to shallow partitions and for enlightening conversations. We would also like to thank Timothy Chan for answering our questions about shallow partitions.

References

- [1] P. Afshani. Unpublished manuscript.
- [2] P. Afshani and T. M. Chan. Optimal halfspace range reporting in three dimensions. In *Proc. 20th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 180–186, 2009.
- [3] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In *Advances in discrete and computational geometry*, volume 223 of *Contemp. Math.*, pages 1–56. 1999.
- [4] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and its applications. *To appear*.
- [5] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational geometry: algorithms and applications*. Springer-Verlag, Berlin, third edition, 2008.
- [6] T. M. Chan. Optimal partition trees. In *Proc. 26th Annu. ACM Sympos. Comput. Geom. (SoCG)*, pages 1–10. 2010.
- [7] B. Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, Cambridge, 2000.
- [8] S. Har-Peled and M. Sharir. Relative (p, ϵ) -approximations in geometry. *Discrete Comput. Geom.*, 45(3):462–496, 2011.
- [9] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8(3):315–334, 1992.
- [10] J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2(3):169–186, 1992.
- [11] K. Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice-Hall, Englewood Cliffs, 1994.
- [12] M. Sharir and P. K. Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, New York, NY, USA, 1995.
- [13] E. Welzl. On spanning trees with low crossing numbers. In *Data structures and efficient algorithms*, volume 594 of *Lecture Notes in Comput. Sci.*, pages 233–249. Springer, Berlin, 1992.