On the Union Complexity of Families of Axis-Parallel Rectangles with a Low Packing Number

Chaya Keller* Shakhar Smorodinsky†

Department of Mathematics
Ben-Gurion University of the Negev
Be'er-Sheva, Israel

{kellerc,shakhar}@math.bgu.ac.il

Submitted: Feb 2, 2017; Accepted: Oct 24, 2018; Published: Nov 16, 2018

Abstract

Let $\mathcal{R}$ be a family of n axis-parallel rectangles with packing number $p-1$, meaning that among any p of the rectangles, there are two with a non-empty intersection. We show that the union complexity of $\mathcal{R}$ is at most $O(n+p^2)$, and that the $(k-1)$-level complexity of $\mathcal{R}$ is at most $O(n+kp^2)$. Both upper bounds are tight.

Mathematics Subject Classifications: 52C45, 52C15

1 Introduction

For a finite family $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ of geometric objects in general position in the plane, the *union complexity* of $\mathcal{U}(\mathcal{C}) = \cup_{i=1}^{n} C_i$ (or, in short, the union complexity of $\mathcal{C}$) is the number of *vertices* on the boundary $\partial(\mathcal{U}(\mathcal{C}))$, where a vertex is an intersection point of the boundaries of two objects $C_i, C_j \in \mathcal{C}$.[1] More generally, for any $k \geqslant 0$, the *k-level complexity* of $\mathcal{C}$ is the number of vertices that are contained in the interior of exactly k elements of $\mathcal{C}$.

*Research partially supported by Grant 635/16 from the Israel Science Foundation, by the Shulamit Aloni Post-Doctoral Fellowship of the Israeli Ministry of Science and Technology, and by the Kreitman Foundation Post-Doctoral Fellowship.

†Research partially supported by Grant 635/16 from the Israel Science Foundation.

[1] Formally, the definition of the union complexity is slightly more complex: it is the total number of faces of all dimensions of the arrangement of the boundaries of the objects, which lie on the boundary of the union (see [1]). We use our simpler definition as in our context, both definitions are clearly equivalent up to a constant factor.

Bounding the union complexity of families of geometric objects is useful for analyzing the running time of various algorithms, and has applications to linear programming, robotics, molecular modeling, and many other fields. In particular, Clarkson and Varadarajan [5] showed that if the union complexity of a family $\mathcal{R}$ of r ranges with VC dimension δ is sufficiently close to $O(r)$, then $\mathcal{R}$ has an ϵ-net of size smaller than $O(\frac{\delta}{\epsilon}\log\frac{\delta}{\epsilon})$; Smorodinsky [11] showed that bounds on the union complexity and on the level-1 complexity of families of geometric objects in the plane can be used in computing the proper chromatic number and the conflict-free chromatic number of the corresponding hypergraph. Union complexity of various families has been studied extensively; k-level complexity was also the subject of extensive study and is sometimes considered harder to evaluate. For example, finding the k-level complexity of a family of n half-planes is the well-known 'k-set problem' posed by Lovász (1971) and Erdős et al. (1973), which is still wide open (see, e.g., [6, 12]). For more on union complexity and k-level complexity, see the survey [1] and [10, Chapter 11].

For several families of geometric objects, it was shown that the union complexity is asymptotically lower than the trivial $O(n^2)$ bound. In particular, Kedem et al. [9] showed that the union complexity of any family of n pseudo-discs in the plane is at most $6n-12$, and Alt et al. [2] and Efrat et al. [7] proved a similar bound for any family of fat wedges. An almost linear bound for families of γ-fat triangles was obtained by Ezra et al. [8].

For a general family of axis-parallel rectangles in the plane, the union complexity can be quadratic – e.g., if the family is an $\frac{n}{2}$-by-$\frac{n}{2}$ grid of long and thin rectangles. However, one may note that such a family contains as many as $n/2$ pairwise disjoint sets. Hence, it is natural to ask whether any family of axis-parallel rectangles with a quadratic union complexity must contain a linear-sized sub-family whose elements are pairwise disjoint.

In this note we answer this question in the affirmative. We show that the union complexity of any family $\mathcal{R}$ of axis-parallel rectangles is sub-quadratic if the *packing number* of the family is sub-linear. Recall that the packing number of $\mathcal{R}$, denoted $\nu(\mathcal{R})$, is $p-1$ if p is the smallest number such that among any p elements of $\mathcal{R}$, two have a non-empty intersection. Our main result is the following:

Theorem 1. *Let $\mathcal{R}$ be a family of axis-parallel rectangles in general position with packing number $\nu(\mathcal{R})$. Then for any $k \geqslant 1$, the $(k-1)$-level complexity of $\mathcal{R}$ is $O(n + k\nu(\mathcal{R})^2)$. In particular, the union complexity of $\mathcal{R}$ is $O(n + \nu(\mathcal{R})^2)$.*

Both results are tight, as we show by an explicit example.[2]

2 Proof of Theorem 1

The proof of Theorem 1 consists of several steps, and for convenience we divide them into separate subsections. We start with a few definitions and notations.

[2]We note that our upper bound on the union complexity is not hereditary, in the sense that there may exist a sub-family of $\mathcal{R}$ (of size $\Theta(p)$, where $\nu(\mathcal{R}) = p-1$ is the packing number of $\mathcal{R}$) whose union complexity is quadratic in its number of elements. Another non-hereditary bound on the union complexity, for specific families of discs in the plane, was obtained by Aronov et al. [3].

2.1 Definitions and Notations

Throughout this note, $\mathcal{R}$ denotes a family of axis-parallel rectangles in the plane, and we assume that $\mathcal{R}$ is in *general position*, meaning that no two rectangles have more than 4 common points (i.e., no two rectangles share a segment of the boundary; this implies that no three boundaries intersect at the same point). Put $\nu(\mathcal{R}) = p - 1$, so any p rectangles in $\mathcal{R}$ contain two with a non-empty intersection.

For any $x \in \mathbb{R}^2$, the *depth* of x, denoted $\text{depth}(x)$, is the number of rectangles in $\mathcal{R}$ that contain x as an interior point. For $k \geqslant 0$, let Y_k be the set of *vertices* (i.e., intersections of pairs of boundaries) of depth k. Of course, $|Y_0|$ is the union complexity of $\mathcal{R}$ and $|Y_k|$ is the k-level complexity of $\mathcal{R}$.

2.2 Partition of the rectangles into floors

Let $R_1 \in \mathcal{R}$ be the rectangle whose upper boundary is the lowest (i.e., has the smallest y coordinate) among the rectangles in $\mathcal{R}$. If there are several such rectangles, we choose one of them arbitrarily. Denote by ℓ_1 the horizontal line that contains the upper boundary of R_1.

Define inductively a sequence $\{\ell_i\}_{2 \leqslant i \leqslant p'}$, for some $p' \leqslant p$, as follows. Let R_i be the rectangle whose upper boundary is the lowest between all elements of $\mathcal{R}$ whose lower boundary is above ℓ_{i-1}. (Again, if there are several such rectangles, we pick one of them arbitrarily.) Denote by ℓ_i the horizontal line that contains the upper boundary of R_i. If there are no rectangles in $\mathcal{R}$ whose lower boundary is above ℓ_{i-1}, take ℓ_i to be an arbitrary horizontal line above ℓ_{i-1}, set $p' = i$, and stop the process. Note that by the construction, the rectangles $\{R_i\}_{1 \leqslant i \leqslant p'}$ are pairwise disjoint. As $\nu(\mathcal{R}) = p - 1$, this implies that the process in indeed finite and that $p' \leqslant p$.

We now define the partition of $\mathcal{R}$ into floors: we say that $R \in \mathcal{R}$ belongs to floor i, $1 \leqslant i \leqslant p' - 1$, if the upper boundary of R is above or contained in ℓ_i and *lower* than ℓ_{i+1}. We denote the set of all rectangles in floor i $(1 \leqslant i \leqslant p' - 1)$ by $\mathcal{F}_i$. It is clear from the construction that $\{\mathcal{F}_i\}_{1 \leqslant i \leqslant p'-1}$ is a partition of $\mathcal{R}$ into $p' - 1 \leqslant p - 1$ pairwise disjoint families. In addition, we need the following observation:

Observation 2. *For any $1 \leqslant i \leqslant p' - 1$, if $R \in \mathcal{F}_i$ then $R \cap \ell_i \neq \emptyset$. Furthermore, i is the largest index such that R intersects ℓ_i.*

Proof. Let $R \in \mathcal{F}_i$. If the lower boundary of R is above ℓ_i then by the definition of ℓ_{i+1}, the upper boundary of R cannot lie strictly below ℓ_{i+1}, a contradiction. Hence, the lower boundary of R is either below ℓ_i or on ℓ_i. As the upper boundary of R is either on ℓ_i or above ℓ_i and also lower than ℓ_{i+1}, the assertion follows. □

Observation 2 implies that $\mathcal{R}$ is *pierced* by the set of lines $\mathcal{L} = \{\ell_1, \ldots, \ell_{p'-1}\}$, meaning that each $R \in \mathcal{R}$ has a non-empty intersection with (at least) one of the lines. A similar argument shows that there exists a set $\mathcal{H} = \{h_1, h_2, \ldots, h_{p''-1}\}$ (for some $p'' \leqslant p$) of vertical lines (arranged in increasing order of the x coordinate) that pierces $\mathcal{R}$. We may assume w.l.o.g. that $p'' = p$. The set $\mathcal{H}$ will be used, along with $\mathcal{L}$, in the sequel.

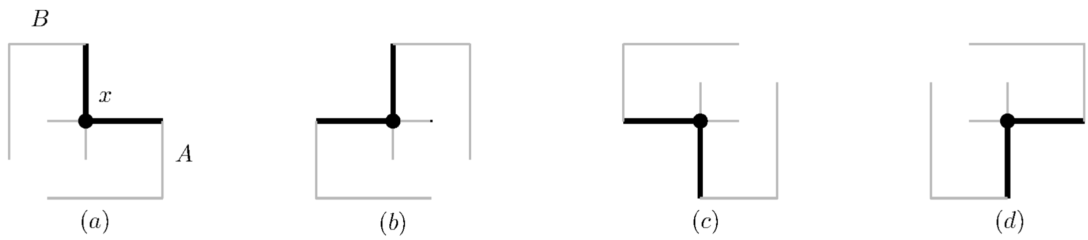

Figure 1: Types of intersection of pairs of axis-parallel rectangles in general position.

2.3 Classification of the intersection points of type L

Intersection points of boundaries of two axis-parallel rectangles can be partitioned into four types, depicted in Figure 1. The type described in Figure 1(a) (in which the intersection point is the rightmost-upmost point of the intersection of the rectangles) will be called *type L* intersection. We denote by X_k the set of all points of type L in Y_k. In what follows, we obtain an upper bound on $|X_k|$. By symmetry, this will imply an upper bound on the k-level complexity of $\mathcal{R}$. As a preparation, we classify the intersection points of type L.

For any intersection point x of type L, we denote by A_x the rectangle to whose *upper* boundary x belongs, and by B_x the rectangle to whose *right* boundary x belongs.

Definition 3. Let $x \in X_k$. Denote by h_x the rightmost amongst the vertical lines in the set $\{h \in \mathcal{H} : h \cap B_x \neq \emptyset\}$. We say that x is *(A_x, h_x)-contributed.*

For $A \in \mathcal{R}$, we say that x is A-contributed if there exists $h \in \mathcal{H}$ such that x is (A, h)-contributed. Conversely, for $h \in \mathcal{H}$, we say that x is h-contributed if there exists $A \in \mathcal{R}$ such that x is (A, h)-contributed (see Figure 2(a)).

Observation 4. *1. For any given A, h, k, there exists at most a single point x with* $\text{depth}(x) = k$ *that is (A, h)-contributed.*

2. It may be that x is (A, h)-contributed, while $A \cap h = \emptyset$ (see Figure 2(b)).

Definition 5. An (A, h)-contributed point x is called an *inner contribution* of A if there exist points y, z and lines $h \neq h', h'' \in \mathcal{H}$, such that:

- y is (A, h')-contributed and z is (A, h'')-contributed, and
- x lies strictly between y and z. (Note that all of x, y, z belong to the upper boundary of A. This induces a natural ordering between them.)

If there are no such points, x is called an *extremal contribution* of A (see Figure 2(c)).

The following observation is crucial in the sequel.

Observation 6. *Let $x \in X_k$ be an (A, h_i)-contributed intersection point. If x is an inner contribution of A, then A intersects both h_i and h_{i+1}.*

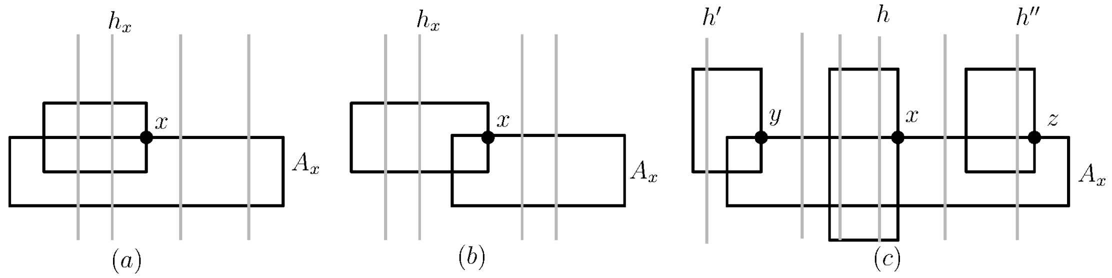

Figure 2: An auxiliary figure for Section 2.3. In (a) and (b), the point x is (A_x, h_x)-contributed. In (c), $x \in X_0$ is an inner contribution of A.

Proof. Denote the vertical lines that contain the left and right boundaries of A by l_A and r_A, respectively. Note that if for some m there exists an (A, h_m)-contributed point $\bar{x}$, then the line h_{m+1} must lie to the right of l_A (as otherwise, $B_{\bar{x}}$ must intersect h_{m+1}, contradicting the assumption that $\bar{x}$ is contributed by h_m). On the other hand, h_m must lie to the left of r_A, since it intersects $B_{\bar{x}}$ and the right boundary of $B_{\bar{x}}$ is to the left of r_A (as the intersection point $\bar{x}$ is of type L, see Figure 1(a)).

In our case, as x is an inner contribution of A, there exist some $s_1, s_2 \geqslant 1$ and points y, z such that y is (A, h_{i-s_1})-contributed and z is (A, h_{i+s_2})-contributed. By the previous paragraph, the former implies that h_{i-s_1+1} lies to the right of l_A while h_{i+s_2} lies to the left of r_A. As $s_1, s_2 \geqslant 1$, this implies that both h_i and h_{i+1} lie to the right of l_A and to the left of r_A, and thus, both intersect A, as asserted. □

2.4 Upper bound on 'inner contributions' to the k-level complexity

In this subsection we obtain an upper bound on the number of elements of X_k that are inner contributions, by considering pairs of the form (Floor $\mathcal{F}_i$, vertical line h_j) separately, and for each such pair, upper bounding the number of (A, h_j)-contributed points for $A \in \mathcal{F}_i$ that are inner contributions.

Proposition 7. *For $k \geqslant 0$, $1 \leqslant i \leqslant p' - 1$, and $1 \leqslant j \leqslant p - 1$, let*

$$S_k^{i,j} = \{x \in X_k : \exists A \in \mathcal{F}_i, x \text{ is } (A, h_j)\text{-contributed and } x \text{ is an inner contribution of } A\}.$$

(Informally, $S_k^{i,j}$ is the set of all contributions to the k-level complexity, that are contributed by h_j on the i-th floor in an 'inner' way). Then for all i, j,

$$|S_k^{i,j}| \leqslant k + 1. \tag{1}$$

Proof. Fix $1 \leqslant j \leqslant p - 1$. Define, for any $1 \leqslant i \leqslant p' - 1$,

$$\mathcal{A}_i = \{A \in \mathcal{F}_i : \exists (A, h_j)\text{-contributed } x \in X_k \text{ that is an inner contribution of } A\}.$$

(Informally, $\mathcal{A}_i$ is the set of all rectangles on the i-th floor, whose upper edge contains an inner contribution to the k-level complexity, contributed by h_j.) Denote $|\mathcal{A}_i| = m_i$, and

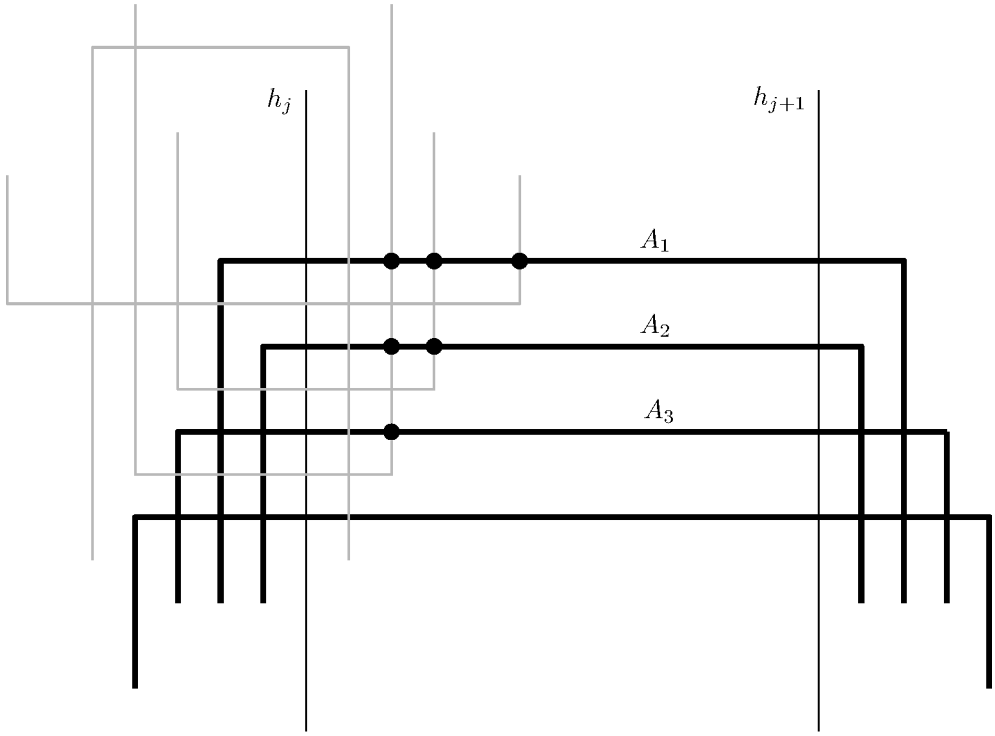

Figure 3: An illustration to the proof of Proposition 7.

let the elements of $\mathcal{A}_i = \{A_1, A_2, \ldots, A_{m_i}\}$ be ordered in descending order of the height of the upper boundary, as demonstrated in Figure 3. (So, A_1 is the rectangle whose upper boundary is the highest, A_2's upper boundary is the second highest, etc.. Note that equality cannot occur here as by Observation 6, any $A \in \mathcal{A}_i$ intersects both h_j and h_{j+1}, and so, if two of these rectangles had upper boundaries of the same height, they would share part of the boundary, contradicting the assumption that the elements of $\mathcal{R}$ are in general position.)

For each $1 \leqslant l \leqslant m_i$, denote

$$Q_l = \{x \in X_k : x \text{ is } (A_l, h_j)\text{-inner contributed}\}.$$

By Observation 4, for each l we have $|Q_l| \leqslant 1$. It is clear that

$$|S_k^{i,j}| = |\{x \in X_k : \exists A \in \mathcal{A}_i \text{ such that } x \text{ is } (A, h_j)\text{-inner contributed}\}| = \sum_{l=1}^{m_i} |Q_l|. \quad (2)$$

Hence, the assertion will follow once we show that $Q_l = \emptyset$ for all $l > k + 1$. To see this, we prove that for each l, each $x \in Q_l$, and each $1 \leqslant r \leqslant l - 1$, x is an interior point of A_r (and thus, each $x \in Q_l$ is of depth $\geqslant l - 1$). We use several simple observations.

1. Any (A_l, h_j)-contributed x lies between the lines h_j (inclusive) and h_{j+1} (non-inclusive). Indeed, as x lies on the right boundary of B_x and h_j intersects B_x,

x must lie either on h_j or to the right of h_j. On the other hand, if x lies on h_{j+1} or on the right of h_{j+1}, then B_x must intersect h_{j+1}, a contradiction.

2. Any such x lies above or on the line ℓ_i, since it belongs to the upper boundary of $A_l \in \mathcal{F}_i$.

3. Each of the rectangles $A_1, \ldots, A_{m_i}$ intersects ℓ_i by Observation 2, and intersects both h_j and h_{j+1} by Observation 6.

By the simple observations, for each $1 \leqslant r \leqslant l-1$, the rectangle A_r intersects ℓ_i, h_j and h_{j+1}, and its upper boundary lies above x (since x lies on the upper boundary of A_l). As x lies between the lines h_j and h_{j+1} and above ℓ_i, it follows that x is an interior point of A_r. This completes the proof. □

2.5 Finalizing the proof of Theorem 1

Now we are ready to prove Theorem 1. Actually, we prove the following exact version of the theorem:

Theorem 8. *Let $\mathcal{R}$ be a family of n axis-parallel rectangles in general position with $\nu(\mathcal{R}) = p-1$. For any $k \geqslant 0$, the k-level complexity of $\mathcal{R}$ is at most $8n + 4(p-1)(p-3)(k+1)$. In particular, the union complexity of $\mathcal{R}$ is at most $8n + 4(p-1)(p-3)$.*

Proof. By symmetry considerations, the k-level complexity of $\mathcal{R}$ is at most $4|X_k|$, so it is sufficient to prove

$$|X_k| \leqslant 2n + (p-1)(p-3)(k+1). \tag{3}$$

We prove (3) by upper bounding the inner contributions and the extremal contributions separately.

Inner contributions. By Proposition 7, for each i, j, the number of inner contributions that correspond to $\mathcal{F}_i$ and h_j is at most $k+1$. For $j \in \{1, p-1\}$, any h_j-contributed x is an extremal contribution. Hence, the number of inner contributions that correspond to $\mathcal{F}_i$ is at most $(p-3)(k+1)$, and so, the total number of inner contributions is at most $(p-1)(p-3)(k+1)$.

Extremal contributions. Let $A \in \mathcal{R}$. By the definition of inner and extremal contributions, all A-contributed points that are extremal contributions belong to one of two vertical lines. By Observation 4, for any single pair (A, h), X_k contains at most one (A, h)-contributed point. Therefore, there are at most two A-contributed points that are extremal contributions. It follows that the total number of extremal contributions is at most $2n$. This completes the proof. □

Remark 9. Theorem 8 implies that the $(\leqslant k)$-level complexity of $\mathcal{R}$ is at most $8(k+1)n + 2(p-1)(p-3)(k+1)(k+2)$. We note that a similar bound on the $(\leqslant k)$-level complexity can be achieved by first obtaining an upper bound on the union complexity of $\mathcal{R}$ and then applying the classical technique of Clarkson and Shor [4] (which bounds the $(\leqslant k)$-level

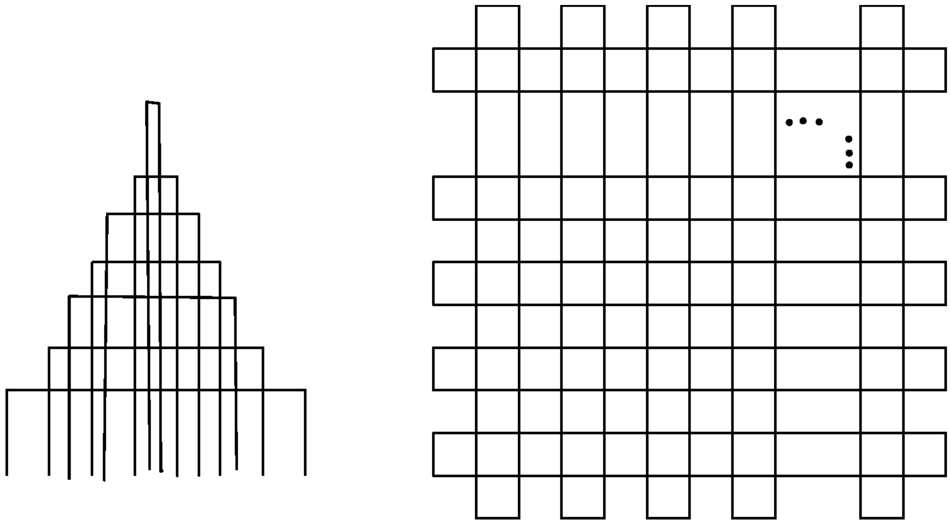

Figure 4: A family of axis-parallel rectangles that demonstrates the tightness of Theorem 1.

complexity of a family $\mathcal{R}'$ using a bound on the union complexity of a random sub-family in which each element of $\mathcal{R}'$ is selected with probability $1/k$). It seems, however, that the Clarkson-Shor technique does not provide an effective bound on the k-level complexity, and thus is not sufficient for proving Theorem 8.

3 Tightness of Theorem 1

In this section we present a family $\mathcal{R}$ of n axis-parallel rectangles with $\nu(\mathcal{R}) = p - 1$ whose $(k-1)$-level complexity is $\Theta(n + kp^2)$, thus showing that Theorem 1 is tight (up to a constant factor). We note that in our construction, we assume that $k + 1 \leqslant \frac{n}{4(p-2)}$.

The family $\mathcal{R}$, presented in Figure 4, is a disjoint union of two subfamilies of $n/2$ rectangles each.

The subfamily drawn in the left of the figure consists of a sequence of pairwise-intersecting rectangles in which each rectangle is taller and thinner than its successor. This subfamily contributes $O(n)$ points to the k-level complexity of $\mathcal{R}$.

The subfamily drawn in the right of the figure is based on an $(p-2)$-by-$(p-2)$ grid of long thin rectangles. We replace each rectangle in the basic grid with $\frac{n}{4(p-2)}$ nested copies to obtain a family of $n/2$ rectangles (for simplicity, we assume $4(p-2)|n$; note that only the basic grid is depicted in the figure). This subfamily contributes $\Theta(kp^2)$ points to the k-level complexity of $\mathcal{R}$. (Here we use the assumption $\frac{n}{4(p-2)} \geqslant k + 1$.)

Hence, the $(k-1)$-level complexity of $\mathcal{R}$ is $\Theta(n + kp^2)$, as asserted.

References

[1] P. K. Agarwal, J. Pach and M. Sharir, State of the union, of geometric objects: A review, in Proc. Joint Summer Research Conf. on Discrete and Computational Geometry: 20 Years Later, *Contemp. Math.* **452** (2008), AMS, pp. 9–48.

[2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, *Algorithmica* **8** (1992), pp. 391–406.

[3] B. Aronov, M. Dulieu, R. Pinchasi, and M. Sharir, On the union complexity of diametral disks, *Electr. J. Comb.* **20(2)** (2013), #P53.

[4] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Disc. Comput. Geom.* **4** (1989), pp. 387–421.

[5] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set covers, *Disc. Comput. Geom.* **37** (2007), pp. 43–58.

[6] T. L. Dey, Improved bounds for planar k-sets and related problems, *Disc. Comput. Geom.* **19(3)** (1998), pp. 373–382.

[7] A. Efrat, G. Rote, and M. Sharir. On the union of fat wedges and separating a collection of segments by a line, *Comput. Geom. Theory Appl.* **3** (1993), pp. 277–288.

[8] E. Ezra, B. Aronov, and M. Sharir, Improved bound for the union of fat triangles, proceedings of SODA'2011, pp. 1778–1785.

[9] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles in the plane, *Disc. Comput. Geom.* **1** (1986), pp. 59–71.

[10] J. Matoušek, Lectures on Discrete Geometry, Springer-Verlag New York, 2002.

[11] S. Smorodinsky, On the chromatic number of some geometric hypergraphs, *SIAM J. Disc. Math.*, **21(3)** (2007), pp. 676–687.

[12] G. Tóth, Point sets with many k-sets, *Disc. Comput. Geom.* **26(2)** (2001), pp. 187–194.