Packing Boundary-Anchored Rectangles

Therese Biedl* Ahmad Biniaz† Anil Maheshwari† Saeed Mehrabi*

Abstract

In this paper, we study the *boundary-anchored rectangle packing* problem in which we are given a set P of points on the boundary of an axis-aligned square Q. The goal is to find a set of disjoint axis-aligned rectangles in Q such that each rectangle is anchored at some point in P, each point in P is used to anchor at most one rectangle, and the total area of the rectangles is maximized. We show how to solve this problem in linear-time in the number of points of P, provided that the points of P are given in sorted order along the boundary of Q. The solvability of the general version of this problem, in which the points of P can also lie in the interior of Q, in polynomial time, is still open.

1 Introduction

Let Q be an axis-aligned square in the plane and P be a set of points in Q. Call a rectangle r *anchored* at a point $p \in P$ if p is a corner of r. The *anchored rectangle packing* (ARP) problem is to find a set S of disjoint axis-aligned rectangles in Q such that each rectangle in S is anchored at some point in P, each point in P is a corner of at most one rectangle in S, and the total area of the rectangles in S is maximized; see Figure 1(a). It is not known whether or not this problem is NP-hard. The best known approximation algorithm for this problem, which achieves ratio $7/12 - \varepsilon$, is due to Balas et al. [1]. They also studied several variants of this problem.

In this paper, we study a simpler variant of the anchored rectangle packing problem in which all the points of P lie on the boundary of Q. We refer to this variant as the *boundary-anchored rectangle packing* (BARP) problem; see Figure 1(b). We present a simple algorithm that solves the BARP problem in linear time, provided that the points of P are given in sorted order along the boundary of Q. Despite the simplicity of our algorithm, its correctness proof is non-trivial. We present our algorithm in Section 3, and prove its correctness in Section 4.

*Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada. `biedl@uwaterloo.ca`, `smehrabi@uwaterloo.ca`. Research of TB supported by NSERC. Part of this work was done while SM was visiting Carleton University.

†School of Computer Science, Carleton University, Ottawa, Canada. `ahmad.biniaz@gmail.com`, `anil@scs.carleton.ca`

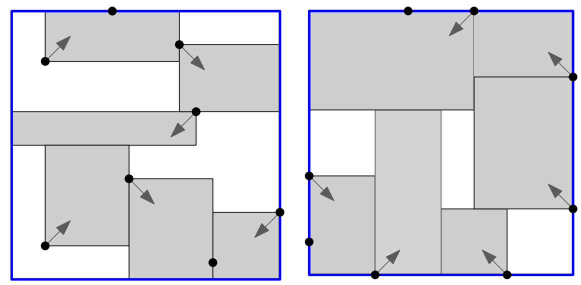

(a) (b)

Figure 1: Instances of (a) the ARP problem, and (b) the BARP problem.

Related results. The rectangle packing problem is related to strip packing and bin packing problems, which are well-known optimization problems in computational geometry. Rectangle packing problems have applications in map labelling [4, 7]. Balas et al. [1] studied several variants of the anchored rectangle packing problem, namely, the *lower-left anchored rectangle packing* problem in which points of P are required to be on the lower-left corners of the rectangles in R, the *anchored square packing* problem in which every anchored rectangles is required to be a square, and the *lower-left anchored square packing* problem which is a combination the two previous problems. For the lower-left rectangle packing problem, Freedman [6] conjectured that there is a solution that covers 50% of the area of Q. The best known lower bound of 9.1% of the area of Q is due to Dumitrescu and Tóth [3]. Balas et al. [1] presented approximation algorithms with ratios $(7/12 - \varepsilon)$ and $5/32$ for anchored rectangles and anchored square, respectively. They also presented a 1/3-approximation algorithm for the lower-left anchored square packing problem, and proved that this lower bound is tight. Balas and Tóth [2] studied the combinatorial structure of maximal anchored rectangle packings and showed that the number of such distinct packings with the maximum area can be exponential in the number n of points of P; they give an exponential upper bound of $2^n C_n$, where C_n denotes the nth Catalan number.

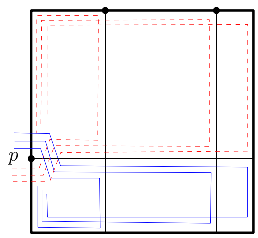

Figure 2: BARP can be solved via maximum-weight independent set in an outer-string graph.

2 An outline

We first briefly argue that BARP is sovable in polynomial time. It is easy to see [1] that in any rectangle packing the boundaries of rectangles must lie on the *grid* Λ obtained by extending rays inwards from all points until they hit the opposite boundary. For each point $p \in P$, there are $O(n^2)$ potential rectangles of Λ anchored at p and so we have $O(n^3)$ candidate rectangles, of which we must pick an independent set (among their intersection graph) such that the sum of the weights (defined to be the area of each rectangle) is maximized. If all points are on the boundary, then it is easy to represent each rectangle as a *string* (i.e., a Jordan curve) such that all strings have a point on the infinite face and two strings intersect if and only if not both rectangles should be taken, see Figure 2. It is known that maximum-weighted independent set is solvable in $O(N^3)$ time on an outer-string graph with a geometric representation of $O(N)$ [5]. As such, BARP is solvable in $O(n^9)$ time, but this is rather slow.

In this section, we give key insights that lead to faster algorithms. Define a *cell* to be a maximal rectangle not intersected by lines of grid Λ. Given an optimum solution S, define a *hole* of S to be a maximal connected region of Q that is not covered by S, see Figure 3(b). We show the following in Section 4:

Insight 1 *An optimal solution S either covers all of Q, or it has exactly one hole which is a single cell.*

It is quite easy to test whether all of Q can be covered (see Lemma 10). If this is not possible, then we want to minimize the hole. However, there are a quadratic number of cells, and more crucially, not all cells are feasible (i.e., can be holes in a packing). The second key result is therefore the following (by Theorem 2):

Lemma 1 *For any cell ψ, we can test in $O(1)$ time whether some packing covers $Q - \psi$.*

This immediately gives an $O(n^2 \log n)$ algorithm to find the best solution of type $Q - \psi$: sort the cells by increasing area, and test for each of them whether it is feasible until we succeed. However, it is not necessary to test each cell individually. We can characterize exactly when a cell ψ is feasible, based solely on where the supporting lines of ψ (which are either the boundary of Q or rays emanating from some points) have their endpoints. Hence we need not look at individual cells, only at the list of points on the four sides, to find the minimum area hole.

3 A Linear-Time Algorithm

Before stating this characterization, we need a few definitions. We write $P_B/P_L/P_T/P_R$ for the points of P on the bottom/left/top/right side. For a point p in the plane, we denote by $x(p)$ and $y(p)$ the x- and y-coordinates of p, respectively. The following theorem proved in Section 4 characterizes possible optimal solutions; Figure 7 on page 5 illustrates these configurations.

Theorem 2 *Any BARP instance has an optimal solution S with $i \leq 4$ rectangles. Moreover (up to rotating the instance by a multiple of $90°$ and/or reflecting horizontally) the anchor-points $p_1, \ldots, p_i$ used by S satisfy one of the following:*

1. *$i = 1$, and p_1 is the leftmost point of $P_L \cup P_B$.*
2. *$i = 2$, and one of the following holds:*
   (a) *p_1 is the bottommost point of P_L and p_2 is the leftmost point of $P_T \cup P_B$, or*
   (b) *p_1 and p_2 are the two points of $P_T \cup P_B$ with the closest x-coordinates.*
3. *$i = 3$, $p_1 \in P_B$ and $p_2 \in P_T \cup P_B$ have closest x-coordinates with $x(p_1) < x(p_2)$, and p_3 is the lowest point in P_L.*
4. *$i = 4$, $p_1 \in P_L$ and $p_3 \in P_R$ have closest y-coordinates with $y(p_1) > y(p_3)$, and $p_2 \in P_T$ and $p_4 \in P_B$ have the closest x-coordinates with $x(p_4) < x(p_2)$.*

Algorithm. Our algorithm proceeds as follows. For each of the four rotations, for each of the two reflections and for each rule 1, 2(a), 2(b), 3, and 4 in Theorem 2, compute the corresponding point set. Each of these up to 40 point sets defines a cell H, and a packing that covers $Q - H$ (see also Lemma 8). The algorithm returns the one that has the smallest hole H.

Having P_L, P_T, P_R, and P_B sorted along the boundary of Q, we can also compute sorted lists of $P_L \cup P_R$ and $P_T \cup P_B$ in linear time. The closest pair within each or between two of them can be computed in linear time. This implies our claimed running time.

The correctness will be proved in Section 4, and does not use that Q is a square, only that it is an axis-aligned rectangle. We hence have:

Theorem 3 *The boundary anchored rectangle packing problem for n points, given in sorted order on the boundary of a rectangle, can be solved in $O(n)$ time.*

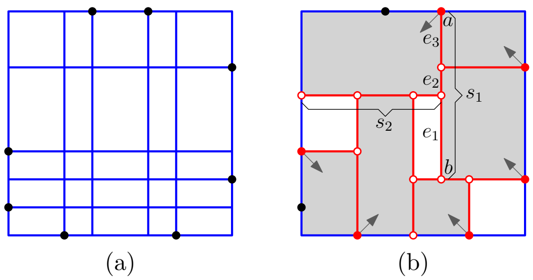

Figure 3: (a) The grid Λ. (b) White regions are holes. Graph $G(S)$ is in red (thick); filled vertices are points of P. The max-segment s_1 is introduced while s_2 is not.

4 Correctness of the Algorithm

We first eliminate some simple cases.

Observation 1 *Assume one of the following holds.*

(i) there exists a point $p_1 \in P$ on a corner of Q, or

(ii) there exist two points in $p_1, p_2 \in P_L \cup P_R$ that have the same y-coordinates, or

(iii) there exist two points in $p_1, p_2 \in P_T \cup P_B$ that have the same x-coordinates.

Then we can cover all of Q with anchored rectangles.

Proof. In case (i), one rectangle anchored at p_1 can cover all of Q. In case (ii) and (iii), two rectangles anchored at p_1, p_2 can cover all of Q. □

Since these conditions are easily tested, we assume for most of the remaining section that none of (i-iii) holds. (We will see that this implies that there must be a hole.)

We need some notation. Throughout this section, let S be a solution for the BARP problem. The term "rectangle" now means one of the rectangles used by S. Define $G(S)$ to be the graph whose vertices are the rectangle-corners that are not corners of Q, and whose edges are coincident with the rectangle-sides not on the boundary of Q; see Figure 3(b).

We define a *max-segment* of $G(S)$ to be a maximal chain s of collinear edges of $G(S)$. We say that s is *introduced* if at least one endpoint of s belongs to P and is used as anchor-point for some rectangle of S. Every edge e belongs to exactly one max-segment s_e; we say that e is *introduced* if s_e is. See Figure 3(b) We already know [1] that all boundaries of rectangles can be assumed to lie on the grid Λ, but we need to strengthen this a bit and prove the following:

Lemma 4 *There exists an optimal solution S such that all max-segments of S are introduced.*

Proof. Let S be an optimal solution that, among all optimal solutions, minimizes the number of max-segments. Assume for contradiction that there exists

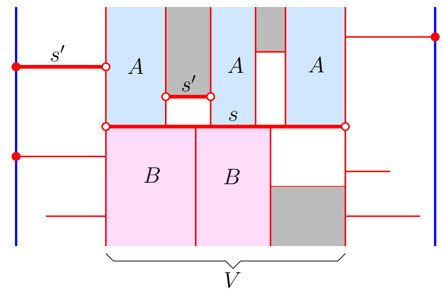

Figure 4: Illustration of the proof of Lemma 4.

a max-segment s that is not introduced. After rotation we may assume that s is horizontal. Let V be the vertical slab defined by the two vertical lines through the endpoints of s; see Figure 4.

Consider moving s upward in parallel, i.e., shortening the rectangles A with their bottom sides on s and lengthening the rectangles B with their top sides on s. Observe first that these rectangles indeed can be shortened/lengthened, because none of them can be anchord at a point on s: the only points of s that are possibly in P are its ends, but neither of them anchors a rectangle since s is not introduced. If this move of s increases the coverage, then S was not optimal, a contradiction. If this decreases the coverage, then moving downward in parallel would increase the coverage, a contradiction. So the covered area must remain the same during the move. Shift s up until it hits either the boundary of Q or intersects some other horizontal max-segment s' of $G(S)$. If s hits the boundary of Q, then s disappears and will be deleted from $G(S)$. If s intersects s' of $G(S)$ (which may be inside V or only share an endpoint with the translated s) then the two max-segments merge into one. Either way we decrease the number of max-segments, which contradicts the choice of S and proves the lemma. □

From now on, without further mentioning, we assume that S is an optimum solution where all max-segments are introduced. We also assume that, among all such optimal solutions, S minimizes the number of rectangles.

Lemma 5 *Every internal vertex of $G(S)$ has degree three or four.*

Proof. Every internal vertex b of $G(S)$ resides on the corner(s) of axis-aligned rectangle(s), and so has degree at least 2 and at most 4. Assume for contradiction that b has degree exactly 2, and let a and c be its neighbours. After possible rotation, we may assume that a lies to the left of b, and c lies above b, as depicted in Figure 5. Thus, b is the bottom-right corner of some rectangle r_1, and no other rectangle has b on its boundary. This

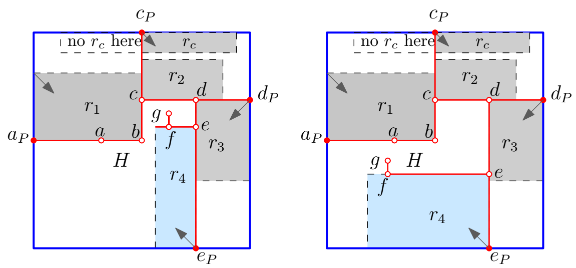

Figure 5: Illustration of the proof of Lemma 5.

implies that the region to the right of bc and below ab is a hole H. So rectangle r_1 is anchored either on the left or the top side of Q; after a possible diagonal flip we assume that it is anchored on the left.

Define a_P and c_P be the points of P that introduced ab and cb, respectively; we know that these must be on P_L respectively P_T since b has degree 2. By definition of "introduced" some rectangle r_c is anchored at c_P. We claim that r_c cannot have c_P as its top-right corner. Assume for contradiction that it did. Then we can expand r_c (if needed) to cover the entire rectangle spanned by a_P and c_P; this can only increase the coverage. In particular, the expanded r_c covers all of r_1. We know that $r_1 \neq r_c$ since r_1 was anchored on the left side of Q. This contradicts that S has the minimum number of rectangles, so r_c has c_P as its top-left corner.

If the right side $\mathrm{rs}(r_1)$ of r_1 is a sub-segment of bc, then we can stretch r_1 to the right to increase the coverage of S, contradicting optimality. So $\mathrm{rs}(r_1)$ must be a strict super-segment of bc, which in particular implies that c is interior and has no leftward edge. Since c is a vertex, it must have a rightward edge; let d be the vertex of H to the right of c. Let r_2 be the rectangle whose bottom-left corner is c; this exists since edge cd is the boundary of some rectangle(s), but the area below cd belongs to hole H. Rectangle r_2 cannot be anchored on the right, because otherwise we could expand r_c to cover all of r_2 and reduce the number of rectangles, a contradiction. So r_2 is anchored on the top, which implies that $r_2 = r_c$, else they would overlap.

If the bottom side $\mathrm{bs}(r_2)$ of r_2 is a sub-segment of cd, then we can stretch r_2 down to increase the coverage of S. So $\mathrm{bs}(r_2)$ is a strict super-segment of cd, which implies that d is interior. We iterate this process three times as follows. (i) Let e be the vertex of H that is below d, and let r_3 be the rectangle whose top-left corner is d. Argue as before that r_3 is anchored at the right endpoint d_P of the max-segment through cd, therefore the left side $\mathrm{ls}(r_3)$ is a strict super-segment of de and e is interior. (ii) Let f be the vertex of H that is to the left of e, and let r_4 be the rectangle whose top-right corner is e. Argue as before that r_4 is anchored at the bottom endpoint e_P of the max-segment through de, therefore the top side $\mathrm{ts}(r_4)$ is a strict super-segment of ef and f is interior. (iii) Finally, let g be the vertex of H that is above f (possibly $g = a$). Now observe that the max-segment through fg cannot reach the boundary of Q without intersecting r_4, r_1 or r_2. Therefore, fg is not introduced, a contradiction. □

We assumed that neither (ii) nor (iii) of Observation 1 holds, which means that any grid-line of grid Λ has exactly one end in P. So, we can direct the edges of the grid (and with it the edges of $G(S)$) from the end in P to the end not in P. See also Figure 7. Define a *guillotine cut* to be a max-segment of $G(S)$ for which both endpoints are on the boundary Q.

Lemma 6 *If there is no guillotine cut, then S has a hole H. Furthermore, H is a rectangle, H is not incident to the boundary of Q, and the boundary of H is a directed cycle of $G(S)$.*

Proof. We claim that no vertex w of $G(S)$ on the boundary of Q is a sink. For if the unique edge incident to w were directed $v \to w$, then by Lemma 4 and the way we directed the edges of $G(S)$, the point p that introduced vw would be on the opposite side and hence the max-segment pw would be a guillotine cut. Likewise no interior vertex w can be a sink, because $\deg(w) \geq 3$ by the previous lemma, which implies that two incident edge of w have the same orientation (horizontal or vertical). One of them then becomes outgoing at w since we direct edges along grid-lines. So $G(S)$ has no sink, which implies that it has a directed cycle C. The region enclosed by C has no point on the boundary, so no rectangle anchored on the boundary can cover parts of it without intersecting C. So the interior region of C is a hole H not incident to the boundary. We know that H is a rectangle since it has no vertex of degree 2 by the previous lemma, hence in particular no reflex vertex. □

This lemma serves as base-case for a stronger claim.

Lemma 7 *If S has holes, then it has a hole H that is a rectangle. Furthermore, every interior corner of H has an incoming edge that lies on H.*

Proof. If there is no guillotine cut, then Lemma 6 gives a rectangular hole that is interior and whose boundary is a directed cycle; this satisfies all claims. So, assume that there is a guillotine-cut aa', say it is horizontal. Since (ii) does not hold, not both a and a' can belong to P, say $a' \notin P$. Segment aa' divides Q into two rectangles Q_1 and Q_2 with Q_1 above Q_2; see Figure 6(a). There is a rectangle r_1 that is anchored at a; up to a vertical flip we may assume that r_1 is inside Q_1. Observe that r_1 must cover all of Q_1, else we could find a solution with

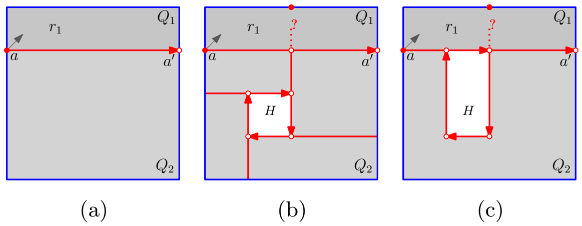

Figure 6: With a guillotine cut, a hole can be found in Q_2 recursively.

more coverage or fewer rectangles. Thus $S' := S \setminus \{r_1\}$ is an anchored-rectangle packing for Q_2 with anchor-points in $P \setminus \{a\}$. S' must be optimal for Q_2, else we could get a better packing for Q by adding r_1 to it. It cannot cover all of Q_2 since S had holes. So, induction applies to S', and it has a hole H.

Assume first that some vertical edge e of H is in the interior and directed downward, see Figure 6(b) and (c). Since e is introduced, the max-segment s_e containing it must then extend to the top of Q. This is impossible since s_e would intersect r_1. So all interior vertical edges of H are directed upwards.

This immediately shows that H cannot be in the interior of Q_2, because then its edges form a directed cycle and one of the vertical ones is directed downward. Likewise it is impossible that both vertical sides and the bottom side of H are interior to Q_2, since the tail-end of the bottom side has an incoming edge from H, which hence must be a downward vertical edge. Therefore H shares at least one side with the boundary of Q.

It remains to argue that any interior corner c of H has an incoming edge on H. If c was interior to Q_2 as well then this holds by induction. If c is interior to Q, but not to Q_2, then c lies on aa' but $c \neq a, a'$. Then the vertical edge of H incident to c is interior to Q_2, so it is directed upward as argued above and hence incoming to c as desired. □

Hence, hole H must satisfy this *hole-condition* on the edge-directions (at least for some optimal solution S); that is, every interior corner of H has an incoming edge that lies on H. It turns out that this condition is also sufficient.

Lemma 8 *Let H be a rectangle whose sides lie on $Q \cup \Lambda$. If every interior corner of H has an incoming edge that lies on H, then there exists a packing that covers $Q \setminus H$.*

Proof. Let $p_1, \ldots, p_i$ (for some $i \leq 4$) be the points of P that defined the grid-lines on which the sides of H reside. We distinguish cases (1-4) depending on how many sides of H are interior, where (2) splits further into (2a) and (2b) depending on whether the sides are adjacent or parallel. After possible rotation, the hole is situated as shown in Figure 7. Every interior corner of H has an incoming edge that is on H, which (up to reflection) forces the location of some of $p_1, \ldots, p_i$ as indicated in the figure. In all cases, one verifies that i rectangles anchored at $p_1, \ldots, p_i$ suffice to cover $Q \setminus H$. □

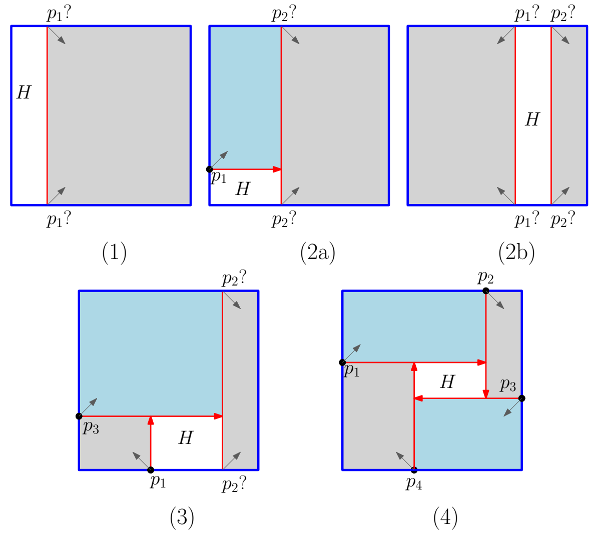

Figure 7: Any rectangle whose boundary is directed suitably can be realized as hole.

We are now ready to prove Insight 1. To this end, we first show the following:

Lemma 9 *If S has holes, then it has exactly one hole H, and H is a cell of Λ.*

Proof. Lemma 7 shows we may assume H to be a rectangle where all interior corners have incoming edges on H. By Lemma 8, we can cover $Q \setminus H$ with anchored rectangles, which by maximality of S means that H is unique.

If H is not a cell, then it is bisected by some grid-line ℓ into two pieces H_1 and H_2. If some $H' \in \{H_1, H_2\}$ satisfies the hole-condition (i.e., all interior corners have incoming edges on H'), then we can create a packing that covers $Q \setminus H' \supset Q \setminus H$, which contradicts minimality of S. In fact, by inspecting the possible configurations of H in cases 1, 2a, 2b, 3, and 4, as well as possible placements of the "undecided" anchor-points and the orientation/direction of ℓ (see Figure 8, which shows all but one case), we observe that H_1 satisfies this condition as we can cover $Q \setminus H_1$ in each of these cases. So, there is a contradiction in all cases, and H must be one cell. □

By Lemma 9, we have characterized solutions that have holes. It remains to characterize solutions that do

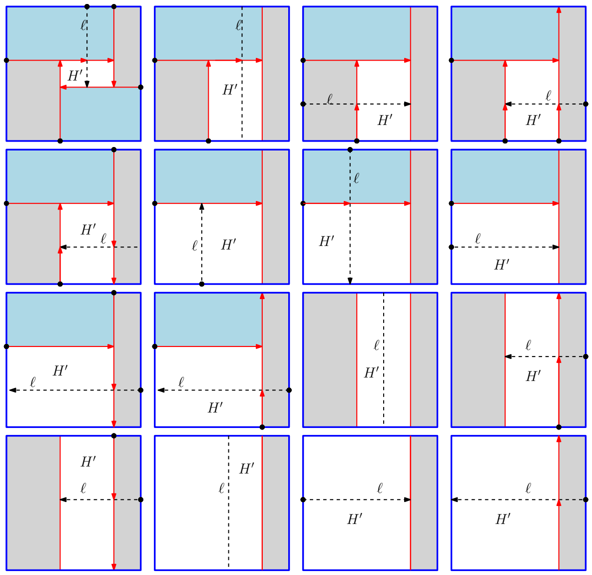

Figure 8: Any hole bisected by a grid-line ℓ gives rise to another hole H'.

not have holes; i.e., to show that the conditions (i-iii) of Observation 1 are necessary.

Lemma 10 *If Q can be covered with anchored rectangles, then one of (i-iii) holds.*

Proof. Let S be a packing that covers all of Q. If $G(S)$ has no edge, then all of Q must be covered by one rectangle, which hence must be anchored at a corner of Q and (i) holds. So assume that $G(S)$ has edges. By Lemma 6, since S has no hole there must be a guillotine-cut aa', say it is horizontal. If both a and a' are in P then (ii) holds and we are done, so assume $a \in P$ and $a' \notin P$.

Define Q_1, Q_2 and r_1 as in Lemma 7 and observe that $S' := S \setminus \{r_1\}$ covers all of Q_2 using anchor-points in $P' := P \setminus \{a\}$. Apply induction to S', P', Q_2. If (i) holds for them, then P' has a point on a corner of Q_2, which by $a, a' \notin P'$ is also a corner of Q and we are done. If (ii) holds for them, then two points in $P' \subset P$ have the same y-coordinate and we are done. Finally (iii) cannot hold for S', P', Q_2 because the top side of Q_2 has no point of P' on it since $a' \notin P$. □

We are finally ready to prove Theorem 2. Let S be the optimum solution with the minimum number of rectangles. If S covers all of Q, then by Lemma 10 one of (i-iii) holds. If (i) holds, then the corner in P will be chosen under rule (1). (In these and all other cases, "chosen" means "after a suitable rotation and/or reflection".) If (ii) or (iii) holds then the two points with the coinciding coordinate will be chosen under rule (2b).

If S has holes, then by Lemma 7 its unique hole H is a cell such that all interior corners of H have incoming edges on H. Let $p_1, \ldots, p_i$ be the points that introduce interior sides of H. We know that H has one of the types shown in Figure 7, and $p_1, \ldots, p_i$ hence will be considered under the corresponding rule. Moreover, all point sets that fit the type can be realized by Lemma 8. So H must be the one that minimizes the area, which corresponds to the points minimizing the x-distance resp. y-distance. So one of rules 1, 2a, 2b, 3 or 4 applies to the points $p_1, \ldots, p_i$ and Theorem 2 holds.

5 Conclusion

In this paper, we considered a variant of the anchored rectangle packing in which all points are on the boundary of the square Q. By exploiting the properties of an optimal solution, we gave an optimal linear-time exact algorithm for this problem. Observe that our algorithm covers nearly everything for large n (contrasting with the fraction of $7/12-\varepsilon$ achieved in the non-boundary case [1]). For there are (up to rotation) at least $n/2$ points in $R_B \cup P_T$, which define $n/2+1$ vertical slabs. Rule (1) or (2b) will consider the narrowest of them as hole, which has area at most $1/(n/2+1)$ if Q has area 1. So we cover a fraction of $1 - O(\frac{1}{n})$ of Q.

The most interesting open question is the status of arbitrary (non-boundary) anchored-rectangle packing. Is this polynomial-time solvable? As a first step, it would be interesting to characterize which polygonal curves on $Q \cup \Lambda$ could be boundaries of a hole in a solution.

References

[1] K. Balas, A. Dumitrescu, and C. D. Tóth. Anchored rectangle and square packings. In *SoCG 2016, Boston, MA, USA*, pages 13:1–13:16, 2016.

[2] K. Balas and C. D. Tóth. On the number of anchored rectangle packings for a planar point set. *Theor. Comput. Sci.*, 654:143–154, 2016.

[3] A. Dumitrescu and C. D. Tóth. Packing anchored rectangles. *Combinatorica*, 35(1):39–61, 2015.

[4] K. G. Kakoulis and I. G. Tollis. Labeling algorithms. In R. Tamassia, editor, *Handbook on Graph Drawing and Visualization.*, pages 489–515. Chapman and Hall/CRC, 2013.

[5] J. M. Keil, J. S. B. Mitchell, D. Pradhan, and M. Vatshelle. An algorithm for the maximum weight independent set problem on outerstring graphs. *Comput. Geom.*, 60:19–25, 2017.

[6] W. Tutte. Recent progress in combinatorics. In *proc. of the 3rd Waterloo Conference on Combinatorics*, 1968.

[7] M. J. van Kreveld, T. Strijk, and A. Wolff. Point labeling with sliding labels. *Comput. Geom.*, 13(1):21–47, 1999.