Packing anchored rectangles*

Adrian Dumitrescu† Csaba D. Tóth‡

Abstract

Let S be a set of n points in the unit square $[0,1]^2$, one of which is the origin. We construct n pairwise interior-disjoint axis-aligned empty rectangles such that the lower left corner of each rectangle is a point in S, and the rectangles jointly cover at least a positive constant area (about 0.09). This is a first step towards the solution of a longstanding conjecture that the rectangles in such a packing can jointly cover an area of at least 1/2.

1 Introduction

We consider a rectangle packing problem proposed by Allen Freedman [19, p. 345] in the 1960s; see also [10, p. 113]. More recently, the problem was brought again to attention (including ours) by Peter Winkler [1, 20, 21, 22]. It is a one-round game between Alice and Bob. First, Alice chooses a finite point set S in the unit square $U = [0,1]^2$ in the plane, including the origin, that is, $(0,0) \in S$ (Fig. 1(a)). Then Bob chooses an axis-parallel rectangle $r(s) \subseteq U$ for each point $s \in S$ such that s is the lower left corner of $r(s)$, and the interior of $r(s)$ is disjoint from all other rectangles (Fig. 1(b)). The rectangle $r(s)$ is said to be *anchored* at s, but $r(s)$ contains no point from S in its interior. It is conjectured that for any finite set $S \subset U$, $(0,0) \in S$, Bob can choose such rectangles that jointly cover at least half of U. However, it has not even been known whether Bob can always cover at least a positive constant area. It is clear that Bob cannot always cover $\frac{1}{2} + \varepsilon$ area for any fixed $\varepsilon > 0$. If Alice chooses S to be a set of n equally spaced points along the diagonal $[(0,0),(1,1)]$, as in Fig. 1(c), then the total area of Bob's rectangles is at most $\frac{1}{2} + \frac{1}{2n}$. There has been no progress on this problem for more than 40 years, even though it appeared several times in the literature.

Outline. In this paper, we present two simple strategies for Bob that cover at least 0.09121 area. These are the GreedyPacking and the TilePacking algorithms described below. Both algorithms process the points in the same specific order, namely the decreasing order of the sum of the two coordinates, with ties broken arbitrarily (hence $(0,0)$ is the last point processed).

The GreedyPacking algorithm chooses a rectangle of maximum area for each point in S sequentially, in the above order.

The TilePacking algorithm partitions U into staircase-shaped tiles, and then chooses a rectangle of maximum area within each tile independently. We next describe how the tiling is obtained. Each tile is a staircase-shaped polygon, with a vertical left side, a horizontal bottom side, and a descending staircase connecting them. The lower left corner of each tile is a point in S. We say that the tile is *anchored* at that point. The algorithm maintains the invariant that the set of unprocessed points are in the interior of a staircase shaped polygon (super-tile), and in addition

*A preliminary version of this paper appeared in the *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms*, (SODA 2012), Kyoto, Japan, January 2012.

†Department of Computer Science, University of Wisconsin–Milwaukee, USA. Email: dumitres@uwm.edu. Supported in part by the NSF grant DMS-1001667.

‡Department of Mathematics, University of Calgary, Canada and Department of Computer Science, Tufts University, Medford, MA, USA. Email: cdtoth@ucalgary.ca. Supported in part by the NSERC grant RGPIN 35586 and the NSF grant CCF-0830734.

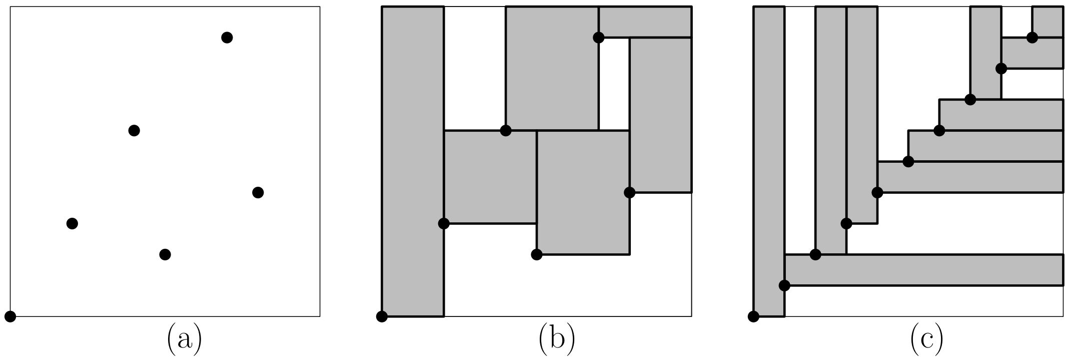

Figure 1: (a) A set S of 6 points in a unit square $[0,1]^2$, including the origin $(0,0)$. (b) A rectangle packing where the lower left corner of each rectangle is a point in S. (c) Ten equally spaced points along the diagonal $[(0,0),(1,1)]$, and a corresponding rectangle packing that covers roughly 1/2 area.

the *anchor* and possibly other points are on its left and lower sides. Processing a point amounts to shooting a horizontal ray to the right and a vertical ray upwards which together isolate a new tile anchored at that point, and the new staircase shaped polygon containing the remaining points is updated. Since $(0,0) \in S$, TilePacking does indeed compute a tiling of the unit square.

It will be shown shortly (Lemma 1) that the GreedyPacking algorithm covers at least as much area as TilePacking. Hence it suffices to analyze the performance of the latter. The bulk of the work is in the analysis of this simple TilePacking algorithm, which involves geometric considerations and a charging scheme.

Related work. Very little is known about anchored rectangle packing. Recently, Christ et al. [9] proved that *if* Alice can force Bob's share to be less than $\frac{1}{r}$, then $n \geq 2^{2^{\Omega(r)}}$. Our result indicates that this condition does not materialize for large r, since Bob can always cover at least a constant fraction of the area for any $n \in \mathbb{N}$.

Previous results on rectangle packing typically consider optimization problems, and are only loosely related to our work. While our focus here is not on the optimization version of the anchored rectangle packing problem, in which the total area of the anchored rectangles is to be maximized for a given set S of anchors, our algorithms do provide a constant-factor approximation.

In the classical *strip packing* problem, n given axis-aligned rectangles should be placed (without rotation or overlaps) in a rectangular container of width 1 and minimum height. This problem is APX-hard (by a reduction from bin packing). After a series of previous results (*e.g.*, [17, 18]), Harren et al. [12] recently found a $(5/3+\varepsilon)$-approximation. Jensen and Solis-Oba [14] devised an AFPTAS which packs the rectangles into a box of height at most $(1+\varepsilon)\mathrm{OPT}+1$ for every $\varepsilon > 0$. Bansal et al. [5] gave a 1.69-approximation algorithm for the 3-dimensional version.

Further related problems are the *2-dimensional knapsack* and *bin packing* problems. Given a set of axis-aligned rectangles and a box B, the *geometric 2D knapsack* problem asks for a subset of the rectangles of maximum total area that fit into B. In contrast, the *2D bin packing* asks for the minimum number of bins congruent to B that can accommodate all rectangles. Jansen and Prädel [16] designed a PTAS for the geometric 2D knapsack problem, although it does not admit a FPTAS. The weighted version does not admit an AFPTAS and an approximation algorithm by Jansen and Zhang [15] guarantees a ratio of $2+\varepsilon$ for every $\varepsilon > 0$. For the 2D bin packing, Jansen et al. [13] gave a 2-approximation, and Bansal et al. [3] designed a randomized algorithm with an *asymptotic* approximation ratio of about 1.525, improving the previous ratio 1.691 by Caprara [6, 7]. However, 2D bin packing does not admit an AFPTAS [4, 8]. Finally, we mention that Bansal et al. [4] gave a PTAS for the *rectangle placement* problem, which goes back to Erdős and Graham [11]. Here a given set of axis-aligned rectangles should be arranged (without rotation or overlaps) so that the area of their bounding box is minimized.

2 Constructing a rectangle packing

In this section we describe the two strategies for Bob and then compare their performance.

Ordering the points in S. Let S be a set of n distinct points in the unit square $[0,1]^2$ such that $(0,0) \in S$. Denote by $x(s)$ and $y(s)$, respectively, the x- and y-coordinates of each point $s \in S$. Order the points in S as $s_1, s_2, \ldots, s_n$ such that

$$x(s_j) + y(s_j) \leq x(s_i) + y(s_i)$$

for $1 \leq i < j \leq n$ (ties are broken arbitrarily). Equivalently, this order is given by a left-moving sweep-line with slope -1. See Fig. 2. Clearly, we have $s_n = (0,0)$. In GreedyPacking, Bob chooses rectangles of maximum area for $s_1, \ldots, s_n$ in this order.

> GreedyPacking. For $i = 1, \ldots, n$, choose an axis-aligned rectangle $r_i \subseteq [0,1]^2$ of maximum area such that the lower left corner of r_i is s_i, and r_i is interior-disjoint from any r_j, $j < i$.

Recall the partial order, called *dominance order*, among points in the plane. For two points, $p = (x_p, y_p)$ and $q = (x_q, y_q)$, we say that $p \preceq q$ (in words, q *dominates* p) if

$$x_p \leq x_q \qquad \text{and} \qquad y_p \leq y_q.$$

With this definition, an axis-aligned rectangle with lower left corner c_1 and upper right corner c_2 can be written as $\{p \in \mathbb{R}^2 : c_1 \preceq p \preceq c_2\}$. In particular, any point in $r(s)$ dominates s.

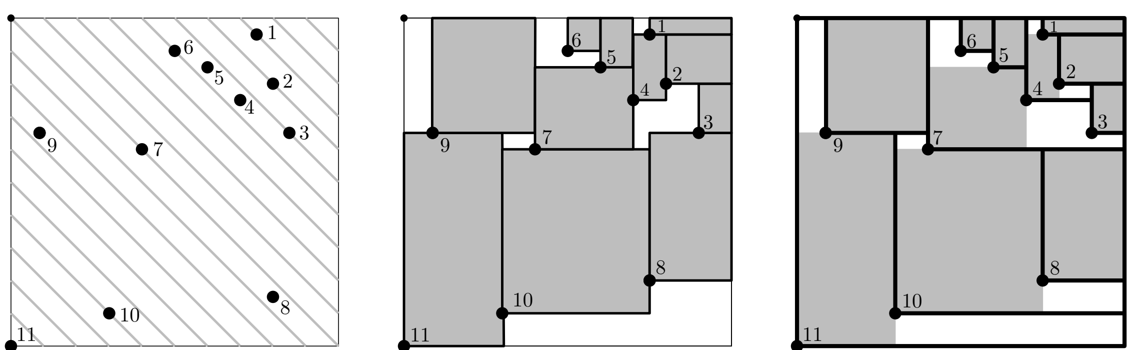

(a) (b) (c)

Figure 2: (a) Eleven points in $[0,1]^2$, sorted by decreasing order of the sum of coordinates. (b) Rectangles chosen greedily in this order. (c) The tiling of $[0,1]^2$ induced by the dominance order; and the maximum-area rectangles chosen from each tile. Note that rectangle 8 is smaller than the corresponding greedy rectangle.

We now define interior-disjoint *tiles* for the set $S = \{s_1, \ldots, s_n\}$ that jointly cover the unit square $U = [0,1]^2$. For $i = 1, 2, \ldots, n$, let tile t_i be the set of points in U that dominate s_i, but have not been covered by any previous tile t_j, $j < i$. Formally, let

$$t_i = \{p \in [0,1]^2 : s_i \preceq p \text{ and } s_j \not\preceq p \text{ for all } j < i\}.$$

The tiles are disjoint by definition, and they cover U since the origin is in S. Each tile is a *staircase polygon* with axis-aligned sides, bounded by one horizontal side from below, one vertical side from the left, and a monotone decreasing curve from the top and from the right. Observe that the axis-aligned rectangle spanned by the lower left corner s_i and any point $p \in t_i$ is contained in the tile t_i. That means that every maximum-area axis-aligned rectangle in the tile is incident to the lower left corner s_i. We can now describe our second strategy for Bob.

> TilePacking. Compute the tiling $U = \bigcup_{i=1}^{n} t_i$. For $i = 1, \ldots, n$, independently, choose an axis-aligned rectangle $r_i \subseteq t_i$ of maximum area.

By the above observation, the lower left corner of r_i is the lower left corner of the tile t_i, which is $s_i \in S$.

We now show that GreedyPacking always covers a greater or equal area than TilePacking.

Lemma 1 *For each point $s_i \in S$, GreedyPacking chooses a rectangle of greater or equal area than TilePacking.*

Proof. The rectangles chosen by the greedy tiling for s_j, $j < i$, are all disjoint from the tile t_i, because every point in $r(s_j)$ dominates s_j. Hence, GreedyPacking could choose any maximum-area axis-aligned rectangle from tile t_i, but it may choose a larger rectangle (such as $r(s_8)$ in Fig. 2). □

Remark. It is worth noting that GreedyPacking cannot give a better worst-case constant than TilePacking. If S is in "general position" (that is, no two points lie on the same sweep-line), we construct a set S', where $|S'| \leq 3n - 2$ as follows. Choose a sufficiently small $\varepsilon > 0$. Then for each point $s_i = (x_i, y_i)$ in the interior of U, we add two nearby points $(x_i - \varepsilon, y_i)$ and $(x_i, y_i - \varepsilon)$ (below the sweep-line incident to s_i, one to the left of s_i and one below s_i). Observe that on input S', GreedyPacking and TilePacking give the same set of rectangles. Moreover, the total area covered by TilePacking with S' and the total area covered by TilePacking with S will differ from each other by an arbitrarily small amount, i.e., by at most $2n\varepsilon$.

3 Analysis of TilePacking

Formally, define ρ as the worst-case performance of TilePacking. In this section, we show that TilePacking chooses a set of rectangles of total area $\rho \geq 0.09121$. In our analysis, we will use two variables, $\beta \geq 5$ and $0 < \lambda < 1$.

Let $r_i \subseteq t_i$ be the axis-aligned rectangle of maximum area in tile t_i, whose lower left corner is s_i, and let $R = \{r_i : i = 1, \ldots, n\}$ be this set of rectangles. If $\text{area}(r_i) \geq 0.09121 \cdot \text{area}(t_i)$ for every i, then our proof is complete. However, the ratio $\text{area}(r_i)/\text{area}(t_i)$ may be arbitrarily small because the tiles can be arbitrary staircase polygons.

For $\beta \geq 5$, we say that tile t_i is a *β-tile* if $\text{area}(r_i) < \frac{1}{\beta}\,\text{area}(t_i)$. We will give an upper bound $F(\beta, \lambda)$ on the total area of β-tiles for every $\beta \geq 5$ and $0 < \lambda < 1$. This immediately implies that, for every $\beta \geq 5$, the complement of all β-tiles cover at least $1 - F(\beta, \lambda)$ area, and the total area of all rectangles in R is

$$\sum_{i=1}^{n} \text{area}(r_i) \geq \rho \geq \frac{1 - F(\beta, \lambda)}{\beta}. \tag{1}$$

In Section 3.1 we study the properties of individual β-tiles, and in Section 3.2 we prove an upper bound on the total area of β-tiles (for every $\beta \geq 5$), which already gives a preliminary bound

$\rho \geq 0.07229$, using (1). By integrating over β, we improve this bound to $\rho \geq 0.09121$ in Section 3.3. Finally, Section 3.4 shows that TilePacking runs in $O(n \log n)$ time.

3.1 Properties of a β-tile

Let us introduce some notation for describing a single tile t_i (Fig. 3(a)). It is bounded from below by a horizontal side, denoted a_i, and from the left by a vertical side, denoted b_i. The *width* (resp., *height*) of t_i is the length of a_i (resp., b_i), denoted by $|a_i|$ (resp., $|b_i|$). We show next that a β-tile t_i has a much smaller area than its bounding box (recall that β-tiles are defined for $\beta \geq 5$).

Lemma 2 *Let $\beta \geq 1$ and t be a staircase polygon of height h and width w. If the area of every axis-aligned rectangle contained in t is less than* area$(t)/\beta$, *then*

$$\text{area}(t) < \frac{\beta}{e^{\beta-1}} \cdot hw, \tag{2}$$

and this bound is the best possible.

Proof. Assume, by translating t if necessary, that the lower left corner of t is the origin. Then the bounding box of t is $[0, w] \times [0, h]$. Let area$(t_i) = \beta u^2$, for some $u > 0$. Then all vertices of t lie strictly below the hyperbola arc $f(x) = u^2/x$, for $0 < x$. See Fig. 3(b). The area of the part of $[0, w] \times [0, h]$ below the curve $f(x) = u^2/x$ is

$$u^2 + \int_{u^2/h}^{w} \frac{u^2}{x}\, \mathrm{d}x = u^2 + \left[u^2 \ln x\right]_{u^2/h}^{w} = u^2(1 + \ln w - \ln(u^2/h)) = u^2(1 + \ln(hw) - \ln u^2).$$

We have shown that area$(t) = \beta u^2 < u^2(1+\ln(hw) - \ln u^2)$, or $\beta < 1 + \ln(hw) - \ln u^2$. Rearranging this inequality yields $e^{\beta-1}u^2 < hw$, which implies (2), as required.

If we approximate the shaded area in Fig. 3(b) with a staircase polygon t of height h and width w that lies strictly below the hyperbola, then the area of every axis-aligned rectangle contained in t is less than u^2, and area(t) can be arbitrarily close to $\frac{\beta}{e^{\beta-1}} \cdot hw$. □

Sectors and tips. Fix $\beta \geq 5$ and consider a β-tile t_i. Decompose t_i into rectangular *vertical sectors* by vertical lines passing though its vertices; see Fig. 3(c). Each sector is part of some maximum-area axis-aligned rectangle in t_i. Hence the area of each sector is less than area$(t_i)/\beta$, where near equality is possible for the leftmost sector. Similarly, we can decompose t_i into rectangular *horizontal sectors* by horizontal lines passing through its vertices, and the area of each sector is less than area$(t_i)/\beta$.

Decompose each β-tile t_i into three parts, called *right tip*, *upper tip*, and *main body*, as follows. The right tip of t_i is cut off from t_i by the right-most vertical line that passes through a vertex of t_i such that the area of the right part is at least area$(t_i)/\beta$. Similarly, the upper tip of t_i is cut off from t_i by the upper-most horizontal line that passes through a vertex of t_i such that the area of the part above is at least area$(t_i)/\beta$. The remaining part, denoted by t_i', is the main body of t_i. All three parts are staircase polygons. Both the right and the upper tips are unions of some sectors of t_i. Since the area of each sector is less than area$(t_i)/\beta$, the area of each tip is at least $\frac{1}{\beta}$area(t_i) but less than $\frac{2}{\beta}$area(t_i). In particular, since $\beta \geq 5$, the right tip of t_i is disjoint from the upper tip of t_i.

Let a_i' and b_i', respectively, be the lower and left side of t_i'. Note that the topmost horizontal side and the rightmost vertical side of t_i' each contains some point from S, because each contains a reflex vertex of the original tile t_i.

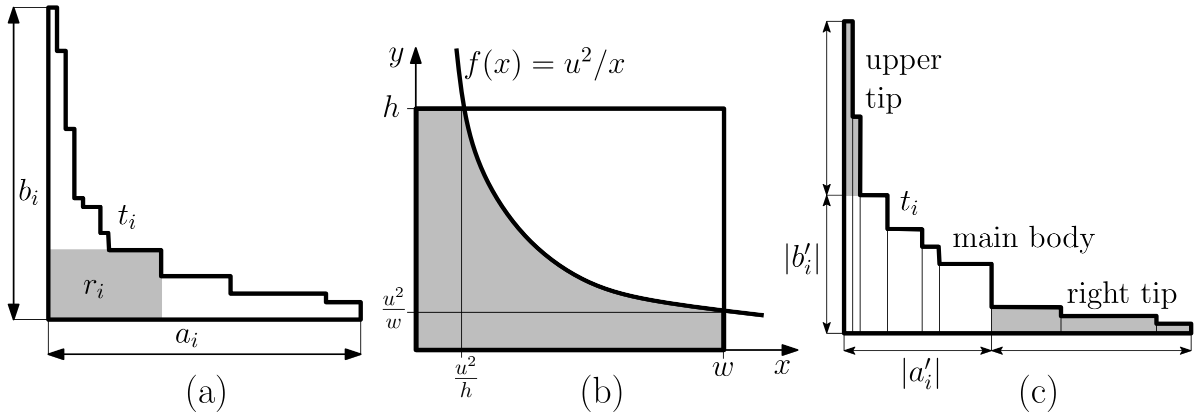

Figure 3: (a) A β-tile t_i of width $|a_i|$ and height $|b_i|$. (b) The portion of the rectangle $[0, w] \times [0, h]$ below the hyperbola arc $f(x) = u^2/x$. (c) The decomposition of t_i into vertical sectors. The tips of t_i are shaded.

Lemma 3 *The width of the right tip of t_i is at least $|a'_i|$, hence $2|a'_i| \le |a_i|$. Similarly, the height of the upper tip of t_i is at least $|b'_i|$, hence $2|b'_i| \le |b_i|$.*

Proof. By symmetry, it is enough to prove the first claim. Let r' be the maximum-area axis-aligned rectangle in t_i whose lower side is a'_i. Since t_i is a β-tile, the area of r' is less than $\frac{1}{\beta}\text{area}(t_i)$. Recall that the area of each tip is at least $\frac{1}{\beta}\text{area}(t_i)$, thus area$(r')$ is less than the area of the right tip of t_i, and so area(r') is less than the area of the bounding box of the right tip of t_i. However, the height of r' is strictly greater than the height of the right tip (and its bounding box). Therefore, the width of r', which is $|a'_i|$, is less than the width of the right tip of t_i. Thus $|a'_i| \le |a_i| - |a'_i|$, or $2|a'_i| \le |a_i|$, as required. □

Recall that the areas of the right and upper tips of t_i are each less than $\frac{2}{\beta}\text{area}(t_i)$. Hence $\text{area}(t_i) < \frac{\beta}{\beta-4}\text{area}(t'_i)$. Since t_i is a β-tile and $t'_i \subset t_i$, the area of every axis-aligned rectangle contained in t'_i is less than $\text{area}(t_i)/\beta < \text{area}(t'_i)/(\beta - 4)$. Applying Lemma 2 for the main body t'_i yields

$$\text{area}(t_i) < \frac{\beta}{\beta - 4} \cdot \text{area}(t'_i) < \frac{\beta}{\beta - 4} \cdot \frac{\beta - 4}{e^{\beta-5}} \cdot |a'_i| \cdot |b'_i| = \frac{\beta}{e^{\beta-5}} \cdot |a'_i| \cdot |b'_i|. \quad (3)$$

Tall and wide tiles. We distinguish two types of β-tiles based on the height and width of their main body. A β-tile t_i is *tall* if $|a'_i| < |b'_i|$; and it is *wide* if $|a'_i| \ge |b'_i|$. For wide β-tiles, we have $\max(|a'_i|, |b'_i|) = |a'_i|$, and (3) implies

$$\text{area}(t_i) < \frac{\beta}{e^{\beta-5}}|a'_i|^2. \quad (4)$$

Similarly, if a β-tile t_i is tall, then $\text{area}(t_i) < \frac{\beta}{e^{\beta-5}}|b'_i|^2$.

3.2 Upper bound on the total area of β-tiles

In this section we give an upper bound $F(\beta, \lambda)$ (see equation (10) further bellow) on the total area of all β-tiles for every $\beta \ge 5$ and $0 < \lambda < 1$. It is enough to bound the total area of wide β-tiles by $\frac{1}{2}F(\beta, \lambda)$. By symmetry, the same upper bound holds for the total area of tall β-tiles. Let $W \subset \{1, \dots n\}$ be the set of indices of the wide β-tiles.

To begin, for every tile t_i we define two adjacent triangles. Let Δ_i be the isosceles right triangle bounded by a_i, the line of slope -1 through s_i, and a vertical line through the right endpoint of

a_i (see Fig. 4(a)). Similarly, let Γ_i be isosceles right triangle adjacent to b_i that lies left of t_i. The two triangles Δ_i and Γ_i are not part of the tiling $\{t_i : i = 1 \dots n\}$, and they may intersect several tiles. A key fact is that these two triangles are empty of points from S in their interior, regardless whether the tile t_i is a β-tile or not.

Lemma 4 *For every $i = 1, \dots, n$, the interior of Δ_i (resp., Γ_i) is disjoint from S.*

Proof. Suppose to the contrary, that there exist points in S in the interior of Δ_i, and let s_j be the first such point processed by the algorithm. Then s_j is processed before s_i, and so the left side b_j of the tile t_j would cut through the horizontal segment a_i, which is a contradiction. Similarly, if s_j lies in the interior of Γ_i, then the lower side a_j would cut through the left side b_i. □

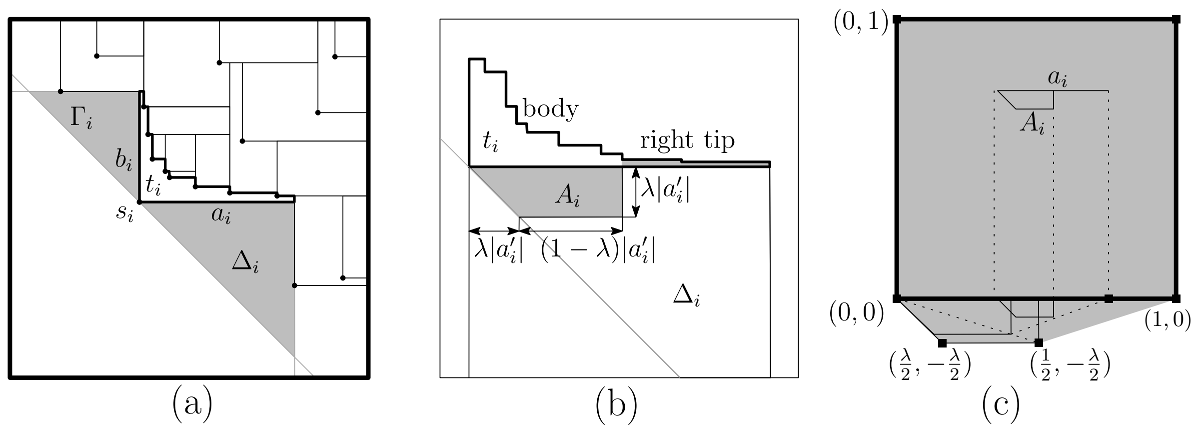

Figure 4: (a) Tiles $t_1, \dots, t_i$, and the triangles Δ_i and Γ_i. (b) The trapezoid $A_i \subset \Delta_i$. (c) Every trapezoid A_i is contained in the shaded (hexagonal) region.

We charge the area of each wide β-tile t_i to the trapezoid $A_i \subset \Delta_i$ defined below. Let A_i be the set of points in Δ_i that lie vertically below the segment a'_i at distance at most $\lambda|a'_i|$ from it. See Fig. 4(c). Using Inequality (4), the area of A_i can be bounded from below as follows in terms of the area of t_i:

$$\text{area}(A_i) = \frac{|a'_i| + (1-\lambda)|a'_i|}{2} \cdot \lambda|a'_i| = \frac{\lambda(2-\lambda)}{2} \cdot |a'_i|^2 > \frac{\lambda(2-\lambda) \cdot e^{\beta-5}}{2\beta} \cdot \text{area}(t_i). \tag{5}$$

The trapezoids A_i, $i \in W$, are homothetic copies of each other, and their area depends only on $|a'_i|$. Note that the triangles Δ_i (and also the trapezoids A_i) may extend beyond the boundary of $U = [0,1]^2$ (*e.g.*, in Fig. 4(b), Δ_i extends below U). We show that all trapezoids A_i, $i \in W$, lie in a polygon whose area is at most $1 + \lambda(3-\lambda)/8$.

Lemma 5 *Every trapezoid A_i, $i \in W$, lies in a polygon of area $1 + \lambda(3-\lambda)/8$.*

Proof. Every trapezoid A_i lies vertically below a segment a'_i, which is part of the lower side of tile t_i. Therefore, A_i cannot extend beyond the left, right, and upper sides of the unit square U. Moreover, we show that A_i is contained in the shaded polygon in Fig. 4(c).

Consider a trapezoid A_i, and the corresponding lower side a_i of a β-tile t_i, $i \in W$. Translate them vertically down until a_i lies on the x-axis. Then apply a dilation centered at the right endpoint of a_i (now on the x-axis), such that the left endpoint of a_i becomes $(0,0)$. Finally, apply a dilation centered at $(0,0)$ such that the right endpoint of a_i becomes $(1,0)$. Observe that through all three transformations, the trapezoid A_i remains in the shaded polygon. The shaded polygon is the union of U and a trapezoid of area $\frac{1}{2}(1 + \frac{1-\lambda}{2})\frac{\lambda}{2} = \lambda(3-\lambda)/8$. □

The case of pairwise disjoint trapezoids. If the trapezoids A_i, for all $i \in W$, are pairwise disjoint, then we can deduce an upper bound for the total area of wide β-tiles: By Lemma 5, we have

$$\sum_{i\in W} \text{area}(A_i) \leq 1 + \frac{\lambda(3-\lambda)}{8} = \frac{8+3\lambda-\lambda^2}{8}.$$

Using Inequalities (4) and (5), this implies

$$\sum_{i\in W} \text{area}(t_i) < \frac{8+3\lambda-\lambda^2}{8} \cdot \frac{2\beta}{\lambda(2-\lambda)\cdot e^{\beta-5}} = \frac{(8+3\lambda-\lambda^2)}{4\lambda(2-\lambda)} \cdot \frac{\beta}{e^{\beta-5}}. \tag{6}$$

The general case of overlapping trapezoids. However, it is possible that the trapezoids A_i, $i \in W$, are not disjoint. To take care of this possibility, we set up a charging scheme, in which we choose a set of "large" pairwise disjoint trapezoids. For every trapezoid A_i, $i \in W$, denote by ℓ_i the supporting line of a_i. We say that A_i is *above* A_j (and A_j is *below* A_i) if A_i and A_j, $i \neq j$ and ℓ_i is above ℓ_j. We next show that if A_i and A_j overlap, then the trapezoid below the other is significantly larger.

Lemma 6 *Assume that $A_i \cap A_j \neq \emptyset$, for some $i, j \in W$, $i \neq j$; and A_i is above A_j. Then $\ell_j \cap \Delta_i \subseteq a'_j$ and $(2-\lambda)|a'_i| \leq |a'_j|$.*

Proof. Refer to Fig. 5(a). If $A_i \cap A_j \neq \emptyset$, and line ℓ_i is above line ℓ_j, then the segment a'_j has to intersect A_i. Note that the left endpoint of a'_j is $s_j \in S$, and there is some point from S on the rightmost edge of t'_j. By Lemma 4, however, there is no point from S in the interior of Δ_i. Therefore, a'_j has to traverse both A_i and Δ_i, hence $\ell_j \cap \Delta_i \subseteq a'_j$.

The minimum horizontal cross-section of A_i is $(1-\lambda)|a'_i|$, and the width of the right tip of t_i is at least $|a'_i|$ by Lemma 3. It follows that $|a'_j| \geq |\ell_j \cap \Delta_i| \geq (1-\lambda)|a'_i| + |a'_i| = (2-\lambda)|a'_i|$. □

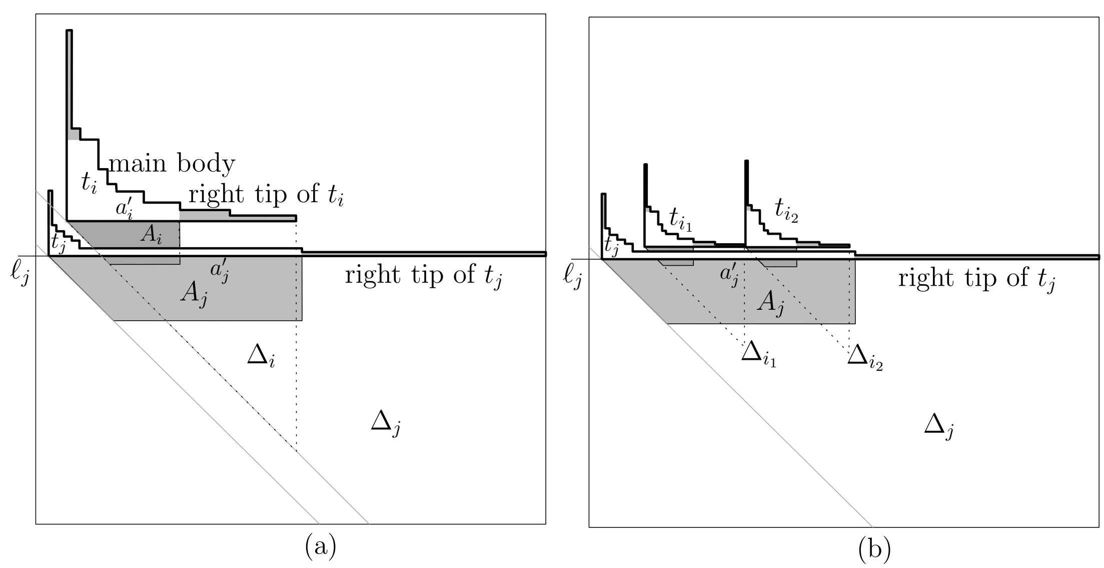

Figure 5: (a) If A_i and A_j overlap, and a_j lies below a_i, then $|a'_j| \geq (2-\lambda)|a'_i|$. (b) Trapezoids A_{i_1} and A_{i_2} intersect A_j from above such that the intervals $a'_j \cap \Delta_{i_1}$ and $a'_j \cap \Delta_{i_2}$ are disjoint.

Charging scheme. We introduce a charging scheme among the trapezoids A_i, $i \in W$. Initially, each trapezoid A_i has a charge of $\text{area}(A_i)$. We transfer the charges to a subset of pairwise disjoint trapezoids. The transfer of charges is represented by a directed acyclic graph G defined as follows. The nodes of G correspond to the trapezoids A_i, $i \in W$. If A_i intersects some other trapezoid below, we add a unique outgoing edge from A_i to the trapezoid A_j, $j \in W$, whose top side a'_j is the highest below a'_i. Observe that all edges of G are oriented downwards, thus G is acyclic. By construction, the out-degree of G is at most one. However, the in-degree of a node in G may be higher than one.

Lemma 7 *For every trapezoid A_j, the total area of all trapezoids A_i, $i \neq j$, with a directed path in G to A_j is at most $\frac{1}{(1-\lambda)(2-\lambda)}\text{area}(A_j)$.*

Proof. Fix a trapezoid A_j, $j \in W$, and denote by ℓ_j be the supporting line of a_j. Refer to Fig. 6(a). For $k \geq 1$, let $W_j^k \subseteq W$ be the set of indices i of trapezoids A_i that have a directed path of length exactly k to A_j in G. We say that the trapezoids A_i, $i \in W_j^k$ are *on level k*. In particular, the trapezoids A_i, $i \in W_j^1$, on level 1 are connected to A_j with a directed edge $(A_i, A_j) \in G$. Let $W_j = \bigcup_{k \geq 1} W_j^k$ be the set of indices of *all* trapezoids A_i, $i \neq j$, with a directed path in G to A_j. For each A_i, $i \in W_j$, denote by A_i^* the unique trapezoid with $(A_i, A_i^*) \in G$.

By Lemma 6, every trapezoid A_i, $i \in W_j^1$, has width at most $|a'_i| \leq |a'_j|/(2-\lambda)$, and height at most $\frac{\lambda}{2-\lambda}|a'_j|$. Equality is possible if the lower left corner of A_i coincides with the upper left corner s_j of A_j (see A_{i_1} in Fig. 6(a)). The gray triangle in Fig. 6(a) is the minimum triangle with base a'_j that contains the maximal possible trapezoid intersecting A_j from above. By triangle similarity, the height of this triangle is $\frac{\lambda}{1-\lambda}|a'_j|$, hence its area is (using Equation (5))

$$\frac{1}{2} \cdot |a'_j| \cdot \frac{\lambda}{1-\lambda}|a'_j| = \frac{\lambda}{2(1-\lambda)}|a'_j|^2 = \frac{\text{area}(A_j)}{(1-\lambda)(2-\lambda)}. \tag{7}$$

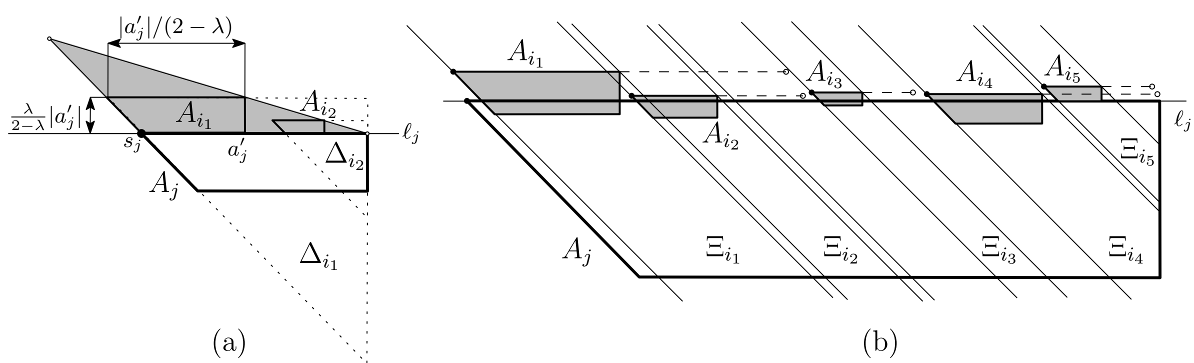

Figure 6: (a) A_{i_1} is the largest possible trapezoid that can intersect A_j from above. All trapezoids with a directed path in G to A_j can be translated (without overlaps) into the gray triangle. (b) Disjoint trapezoids $A_{i_1}\ldots,A_{i_5}$ intersect A_j from above. Each of these trapezoids induces a parallel strip Ξ_i. If the strips Ξ_{i_1} and Ξ_{i_2} intersect, and A_{i_1} is above A_{i_2}, then s_{i_2} lies in the interior of Δ_{i_1}. Dashed lines indicate segments $a_i \setminus a'_i$.

We claim that $\sum_{i \in W_j} \text{area}(A_i)$ is at most the area of the gray triangle in Figure 6(a). To verify the claim, we translate every trapezoid A_i, $i \in W_j$, into the gray triangle region such that they remain pairwise disjoint. Each trapezoid will be translated in the same direction, $(-1, 1)$, but at different distances. In order to control the possible location of the translates, we enclose each A_i,

$i \in W_j$, in a parallel strip. For every $i \in W_j$, draw lines of slope -1 through the two endpoints of a'_i, and denote by Ξ_i the strip bounded by the two lines. Refer to Fig. 6(b).

First consider the trapezoids A_i, $i \in W_j^1$, which are connected to A_j by a directed edge in G. By the definition of G, the trapezoids A_i, $i \in W_j^1$, are pairwise disjoint. Label their strips in increasing order from left to right. We show that the strips Ξ_i, $i \in W_j^1$, are pairwise interior-disjoint. Suppose to the contrary that Ξ_{i_1} and Ξ_{i_2} intersect, where $i_1 < i_2$, and such that A_{i_1} is above A_{i_2} (if A_{i_1} is below A_{i_2}, the strips are obviously disjoint). Since the trapezoids are disjoint, the left endpoint of a'_{i_2}, $s_{i_2} \in S$, lies in the interior of the isosceles right triangle bounded by the right side of A_{i_1}, line ℓ_j, and the line of slope -1 bounding the strip Ξ_{i_1} from the right. Since $\lambda < 1$ and $|a_{i_1}| - |a'_{i_1}| \geq |a'_{i_1}|$ by Lemma 3, this triangle is contained in Δ_{i_1}. However, the triangle Δ_{i_1} is empty of points from S, and we reached a contradiction.

Applying the above argument for the trapezoids on level k, $k = 1, 2, \ldots$, we conclude that the trapezoids on level $k+1$ are pairwise disjoint, and they induce pairwise disjoint parallel strips.

We are now ready to describe the translation of the trapezoids A_i, $i \in W_j$. We translate the trapezoids in direction $(-1, 1)$ in phases $k = 1, 2, \ldots$. In phase k, consider each A_i, $i \in W_j^k$, independently. Translate A_i together with all other trapezoids that have a directed path to A_i by the same vector in direction $(-1, 1)$ until the lower side of A_i becomes collinear with the upper side of A_i^*. Each trapezoid A_i remains in its parallel strip Ξ_i, therefore the trapezoids on the same level remain pairwise disjoint. After phase k, there is no overlap between a trapezoid A_i, $i \in W_j^k$ of level k and trapezoids in lower levels. When all phases are complete, all trapezoids are pairwise interior-disjoint.

It remains to show that after the translation, all trapezoids lie in the gray triangle in Fig. 6(a). From Lemma 6 and since all translations were done in direction $(-1, 1)$, the trapezoids are on or to the right of the line of slope -1 passing through s_j. Also from Lemma 6, the right endpoint of a'_j is to the right of the right side of every triangle Δ_i, $i \in W_j^1$, hence to the right of the right side of every triangle Δ_i, $i \in W_j$. Also, if $i \in W_j^1$, by Lemma 3 we have $2|a'_i| \leq |a_i|$, and so $\frac{\lambda |a'_i|}{|a_i| - |a'_i|} \leq \frac{\lambda |a'_i|}{|a'_i|} = \lambda$. Hence the upper right corner of every A_i, $i \in W_j^1$, is below the line of slope $-\lambda$ passing through the right endpoint of a'_j, and consequently, the upper right corner of every A_i, $i \in W_j$, is below the line of slope $-\lambda$ passing through the right endpoint of a'_j. Therefore, after the above translation, all trapezoids A_i, $i \in W_j$, are contained in the gray triangle in Figure 5(c). This verifies the above claim and completes the proof of the lemma. □

Transfer the charges from all nodes to the sinks in G along directed paths. By Lemma 7, the total area charged to a sink A_j is less than

$$\text{area}(A_j) + \frac{1}{(1-\lambda)(2-\lambda)}\text{area}(A_j) = \frac{3 - 3\lambda + \lambda^2}{(1-\lambda)(2-\lambda)}\text{area}(A_j) \tag{8}$$

The area of every trapezoid A_i, $i \in W$, is charged to some sink in G, and the sinks correspond to pairwise disjoint trapezoids. We can now adjust Inequality (6) to obtain

$$\sum_{i \in W} \text{area}(t_i) < \frac{3 - 3\lambda + \lambda^2}{(1-\lambda)(2-\lambda)} \cdot \frac{(8 + 3\lambda - \lambda^2)}{4\lambda(2-\lambda)} \cdot \frac{\beta}{e^{\beta-5}} = \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{4\lambda(1-\lambda)(2-\lambda)^2} \cdot \frac{\beta}{e^{\beta-5}}. \tag{9}$$

The area of *all* β-tiles is less than twice the right hand-side of (9), namely we can set

$$F(\beta, \lambda) = \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1-\lambda)(2-\lambda)^2} \cdot \frac{\beta}{e^{\beta-5}}. \tag{10}$$

From (1), it follows that the total area of all rectangles in R is

$$\rho \geq \frac{1-F(\beta,\lambda)}{\beta} = \frac{1}{\beta}\left(1 - \frac{(3-3\lambda+\lambda^2)(8+3\lambda-\lambda^2)}{2\lambda(1-\lambda)(2-\lambda)^2}\cdot\frac{\beta}{e^{\beta-5}}\right)$$
$$= \frac{1}{\beta} - \frac{(3-3\lambda+\lambda^2)(8+3\lambda-\lambda^2)}{2\lambda(1-\lambda)(2-\lambda)^2}\cdot\frac{1}{e^{\beta-5}}.$$

Whenever $F(\beta,\lambda) < 1$, this already gives a lower bound of $\rho = \Omega(1)$. We have optimized the parameters β and λ with numerical methods. With the choice of $\beta = 12.75$ and $\lambda = 0.45$, we obtain an initial lower bound of $\rho \geq 0.07229$.

3.3 Making β a continuous variable

In this section we further improve the lower bound on the covered area to 0.09121 by making β a continuous variable and using integration. We define the *contribution* of each point $p \in t_i \subseteq U$, as $u(p) = \text{area}(r_i)/\text{area}(t_i)$. With this definition, we have

$$\sum_{i=1}^{n} \text{area}(r_i) = \iint_{p\in U} u(p)\ \mathrm{d}A.$$

Let $\beta_0 \geq 5$ be a parameter to be optimized later (we will choose $\beta_0 = 9.955$). Partition the interval $[\beta_0,\infty)$ into subintervals of length $\varepsilon > 0$: $[\beta_0,\infty) = \bigcup_{j=1}^{\infty}[\beta_{j-1},\beta_j)$, where $\beta_j = \beta_0 + j\varepsilon$. Denote by $B_j \subset U$ the union of all β_j-tiles, and let $\overline{B}_j = U \setminus B_j$. By definition, we have $u(P) \geq 1/\beta_j$ for every $p \in \overline{B}_j$, and so $\iint_{p\in\overline{B}_j} u(p)\ \mathrm{d}A \geq \text{area}(\overline{B}_j)/\beta_j$.

Observe that the sets $\overline{B}_j$ form a nested sequence $\overline{B}_0 \subseteq \overline{B}_1 \subseteq \overline{B}_2 \subseteq \ldots \subseteq U$. The total contribution of all points in U can be written as

$$\begin{aligned}
\sum_{j=1}^{n} \text{area}(r_i) &= \iint_{p\in U} u(p)\ \mathrm{d}A \\
&= \iint_{p\in\overline{B}_0} u(p)\ \mathrm{d}A + \sum_{j=1}^{\infty}\iint_{p\in\overline{B}_j\setminus\overline{B}_{j-1}} u(p)\ \mathrm{d}A \\
&\geq \frac{\text{area}(\overline{B}_0)}{\beta_0} + \sum_{j=1}^{\infty}\frac{\text{area}(\overline{B}_j\setminus\overline{B}_{j-1})}{\beta_j} \\
&= \frac{\text{area}(\overline{B}_0)}{\beta_0} + \sum_{j=1}^{\infty}\frac{\text{area}(\overline{B}_j)-\text{area}(\overline{B}_{j-1})}{\beta_j} \\
&= \sum_{j=0}^{\infty}\text{area}(\overline{B}_j)\left(\frac{1}{\beta_j}-\frac{1}{\beta_{j+1}}\right).
\end{aligned}$$

In Section 3.2, we showed that for any $j \geq 0$, $\text{area}(B_j) < F(\beta_j,\lambda)$, where $F(\beta,\lambda)$ is given by (10).

It follows that $\text{area}(\overline{B}_j) > 1 - F(\beta_j, \lambda)$, and therefore,

$$\begin{aligned}
\rho &\geq \sum_{j=0}^{\infty}(1 - F(\beta_j, \lambda))\left(\frac{1}{\beta_j} - \frac{1}{\beta_{j+1}}\right) \\
&= \frac{1 - F(\beta_0, \lambda)}{\beta_0} + \sum_{j=0}^{\infty} \frac{(1 - F(\beta_{j+1}, \lambda)) - (1 - F(\beta_j, \lambda))}{\beta_{j+1}} \\
&= \frac{1 - F(\beta_0, \lambda)}{\beta_0} + \sum_{j=0}^{\infty} \frac{F(\beta_j, \lambda) - F(\beta_{j+1}, \lambda)}{\beta_{j+1}} \\
&= \frac{1 - F(\beta_0, \lambda)}{\beta_0} - \sum_{j=0}^{\infty} \frac{1}{\beta_{j+1}} \cdot \frac{F(\beta_{j+1}, \lambda) - F(\beta_j, \lambda)}{\beta_{j+1} - \beta_j} \cdot (\beta_{j+1} - \beta_j).
\end{aligned}$$

Letting ε go to 0 yields

$$\begin{aligned}
\rho &\geq \frac{1 - F(\beta_0, \lambda)}{\beta_0} - \int_{\beta_0}^{\infty} \left(\frac{1}{\beta} \cdot \frac{\partial}{\partial \beta} F(\beta, \lambda)\right) \mathrm{d}\beta \\
&= \frac{1}{\beta_0} + \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1 - \lambda)(2 - \lambda)^2} e^5 \left(-\frac{1}{e^{\beta_0}} - \int_{\beta_0}^{\infty} \left(\frac{1}{\beta} \cdot \frac{\mathrm{d}}{\mathrm{d}\beta} \frac{\beta}{e^{\beta}}\right) \mathrm{d}\beta\right) \\
&= \frac{1}{\beta_0} + \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1 - \lambda)(2 - \lambda)^2} e^5 \left(-\frac{1}{e^{\beta_0}} + \int_{\beta_0}^{\infty} \left(\frac{1}{\beta} \cdot \frac{\beta - 1}{e^{\beta}}\right) \mathrm{d}\beta\right) \\
&= \frac{1}{\beta_0} + \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1 - \lambda)(2 - \lambda)^2} e^5 \left(-\frac{1}{e^{\beta_0}} + \int_{\beta_0}^{\infty} \frac{1}{e^{\beta}} - \frac{1}{\beta e^{\beta}} \mathrm{d}\beta\right) \\
&= \frac{1}{\beta_0} - \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1 - \lambda)(2 - \lambda)^2} e^5 \int_{\beta_0}^{\infty} \frac{1}{\beta e^{\beta}} \mathrm{d}\beta \\
&= \frac{1}{\beta_0} - \frac{(3 - 3\lambda + \lambda^2)(8 + 3\lambda - \lambda^2)}{2\lambda(1 - \lambda)(2 - \lambda)^2} e^5 E_1(\beta_0),
\end{aligned}$$

where $E_1(x)$ is the exponential integral

$$\mathrm{E}_1(x) = \int_x^{\infty} \frac{1}{te^t} \mathrm{d}t.$$

For every $x > 0$, this exponential integral can be approximated by the initial terms of the convergent series

$$\mathrm{E}_1(x) = -\gamma - \ln x - \sum_{k=1}^{\infty} \frac{(-1)^k x^k}{k \cdot k!},$$

where $\gamma = 0.57721566\ldots$ is Euler's constant; see [2]. With the choice of $\beta_0 = 9.955$ and $\lambda = 0.452$, we obtain $\rho \geq 0.09121$.

Taking into account Lemma 1, we summarize our main result in the following theorem.

Theorem 8 *For any finite point set $S \subset U$, $(0,0) \in S$, the algorithm* TilePacking *chooses a set of rectangles of total area $\rho \geq 0.09121$. Consequently, the same guarantee holds for the algorithm* GreedyPacking.

3.4 Runtime analysis

It is not difficult to show that TilePacking can be implemented in $O(n \log n)$ time and $O(n)$ space in the RAM model of computation. The input is a set S of n points in the unit square $U = [0,1]^2$. Clearly, S can be sorted in $O(n \log n)$ time in non-increasing order of the sum of coordinates. Assume that the points are labeled $s_1, s_2, \ldots, s_n$ in this order.

We compute the tiles sequentially in n steps. Let $P_i = U \setminus \bigcup_{j<i} t_j$ be the staircase polygon left from U after deleting the first $i-1$ tiles. We maintain the x- and y-coordinates of the vertices of P_i, respectively, in two binary search trees. In step i, we compute the tile t_i by shooting a vertical (resp., horizontal) ray from s_i until it hits the boundary of P_i. The point hit by an axis-parallel ray can be found with a simple binary search in $O(\log n)$ time. Once the sides a_i and b_i have been determined, we insert the x- and y-coordinates of s_i into the search trees, and delete the points that are not vertices of P_i subsequently. Each point in S is inserted and deleted at most once, so the search trees can be maintained in $O(n \log n)$ total time.

Recall that the generated tiles are staircase polygons, and their reflex vertices are points in S. Since every point in S is a reflex vertex in at most one tile, the total complexity of the n tiles is $O(n)$. A tile with k reflex vertices contains exactly $k+1$ maximal axis-aligned rectangles, and one with the largest area can be selected in $O(k)$ time. Altogether, we can pick an axis-aligned rectangle of maximum area from each of the n tiles in $O(n)$ total time. Hence TilePacking runs in $O(n \log n)$ time.

4 Conclusion

We have shown that in the 1-round rectangle packing game, no matter how Alice chooses a finite set of points in S, where $(0,0) \in S$, Bob can always construct a rectangle packing with rectangles anchored at the points in S that cover at least a constant area. Allen Freedman [19] asked whether Bob can always cover at least $1/2$ of the unit square. Bill Pulleyblank and Peter Winkler conjectured that this is true. While we cannot confirm this at the moment, we believe that the performance of our GreedyPacking and TilePacking algorithms is significantly better than what we proved here.

We suspect that the problem of finding the rectangles with maximum total area anchored at the given points is NP-hard, but this remains to be shown. Our algorithms certainly achieve a constant approximation ratio 0.09121. No efficient exact algorithm or good approximation was previously known.

Special cases and variants. It is easily seen that the conjecture holds for "permutation point sets", namely integer n-element point sets from the $\{0,1,\ldots,n-1\} \times \{0,1,\ldots,n-1\}$ grid, with exactly one grid point in each row and column, and containing $(0,0)$, as required. The unit square is now $[0,n] \times [0,n]$. This is in fact the only family of sets for which we could verify the conjecture. Indeed, for each point, say (i,j), select the rectangle $[i,n] \times [j,j+1]$; then the average width of the chosen rectangles is $(n+1)/2$, each rectangle has unit height, and so the corresponding covered area ratio is $\frac{1}{2} + \frac{1}{2n}$ for each of the $(n-1)!$ point sets. If the anchoring condition is relaxed so that the anchor point of a rectangle can be either of its leftmost two vertices, then it is easy to cover an area of at least $1/2$ as well.

Higher-dimensional version. The d-dimensional generalization of the 1-round rectangle packing game is also very interesting, and almost nothing is known about it. If S is a set of equally spaced points along the main diagonal of a d-dimensional unit cube $U = [0,1]^d$, where $(0,\ldots,0) \in S$, then the total volume covered by any anchored d-dimensional axis-parallel rectangle packing is roughly $1/d$. Our GreedyPacking and TilePacking algorithms readily generalize to d dimensions, but the performance analysis does not seem to be easily extendible. In particular, the definition of β-tiles extends to arbitrary dimensions. Lemma 2 also carries over (i.e., the volume of a β-tile is exponentially smaller than the volume of its bounding box); and there are large empty convex polytopes along the edges of a β-tile (in d-space, these are the d edges incident to the anchor point) similarly to the empty triangles Δ_i and Γ_i in the plane. However, it is not clear what could be the analogues of the trapezoids A_i in higher dimensions and whether any charging scheme could be set up.

Multi-round versions. A natural generalization of the problem is the multi-round rectangle packing game. One can consider two versions, depending on whether the number of rounds is known in advance. In the *n-round rectangle packing game*, both Alice and Bob know the number of rounds. In round i, first Alice places a point $s_i \in [0,1]^2$ somewhere outside of Bob's rectangles, and then Bob chooses an axis-aligned rectangle $r_i \subseteq [0,1]^2$ with lower left corner at s_i and interior-disjoint from his previous rectangles. Alice has to choose the origin in one of the n rounds. In the *unlimited rectangle packing game*, the number of rounds (or points) is not known in advance. Each round goes exactly as in the n-round version, but the game terminates when Alice decides to put a point at the origin and Bob chooses his last rectangle incident to the origin.

For both versions of the multi-round game, Bob could employ a greedy strategy: for each point s_i, let r_i be an axis-aligned rectangle of maximum area with lower left corner at s_i that is interior-disjoint from all previous rectangles. However, our analysis does not extend to these versions of the game. In fact, we can show that the greedy strategy cannot guarantee any constant area for Bob. Whether Bob can secure a constant fraction of the area by other means in any of the multi-round versions of the game remains open.

Theorem 9 *In both multi-round versions of the rectangle packing game, Bob cannot always cover* $\Omega(1)$ *area with a greedy strategy.*

Proof. We show that for every $\varepsilon > 0$, Alice can construct a finite sequence of n points $s_1, \ldots, s_n$, such that Bob can cover at most ε area with a greedy strategy. Essentially, Alice can force Bob to choose a rectangle from a $\frac{\varepsilon}{2}$-tile (using at most $\frac{\varepsilon}{2}$ of the tile's area) and then fence off the remainder of the tile so that it cannot be covered later. Alice can make sure that the total area of these $\frac{\varepsilon}{2}$-tiles is arbitrarily close to 1, say $1-\frac{\varepsilon}{2}$. Then Bob can cover at most $\frac{\varepsilon}{2} + (1-\frac{\varepsilon}{2}) \cdot \frac{\varepsilon}{2} < \varepsilon$ area. We proceed with the details.

We say that a staircase polygon P is a β-staircase, for some $\beta \geq 1$, if the area of every axis-aligned rectangle contained in P is at most $\text{area}(P)/\beta$. By Lemma 2, for every $\beta \geq 1$, $h > 0$, and $w > 0$, one can construct a β-staircase of height h and width w whose area is roughly $\beta e^{1-\beta} hw$.

Alice first computes a packing of the unit square $U = [0,1]^2$ with $\frac{2}{\varepsilon}$-staircases, by successively choosing interior-disjoint $\frac{2}{\varepsilon}$-staircases of smaller and smaller sizes until their total area is at least $1-\frac{\varepsilon}{2}$ (Fig. 7(a,b)). Specifically, in the current step, given an axis aligned rectangle R of height h and width w, a $\frac{2}{\varepsilon}$ staircase polygon P with the same height and width is anchored at the lower left corner of R, and the remaining space $R \setminus P$ is partitioned into vertical or horizontal sectors to be processed (see Section 3).

Then she slightly shrinks these staircases, to make them pairwise disjoint, and perturbs them to ensure that each $\frac{2}{\varepsilon}$-staircase P_i contains a *unique* axis-aligned rectangle r_i of maximum area, and r_i has the same width as P_i. The point set S contains, for every $\frac{2}{\varepsilon}$-staircase in this packing, all vertices of P_i, including the lower left corner s_i. In addition, for every $\frac{2}{\varepsilon}$-staircase whose lower left

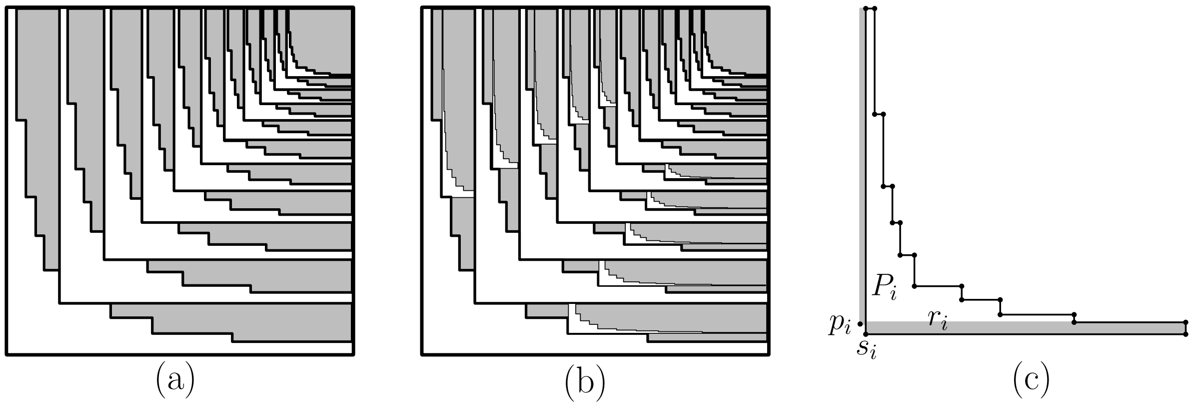

Figure 7: (a) An initial $\frac{2}{\varepsilon}$-staircase packing. (b) In any empty region (shaded gray), Alice can place an additional $\frac{2}{\varepsilon}$-staircase. (c) The placement of an extra point p_i near the left side of the maximal rectangle $r_i \subset P_i$.

corner is not on the left side of U, S also contains a point p_i very close to the left side of rectangle r_i, as shown in Fig. 7(c).

It remains to determine the order in which Alice reveals the points to Bob. The points associated with each $\frac{2}{\varepsilon}$-staircase P_i are revealed in a contiguous sequence such that the last two points in each sequence are the lower left corner s_i followed by the extra point p_i. For the lower left corner s_i, Bob has to choose the unique rectangle r_i of maximum area in P_i, which is adjacent to the lower side of P_i. For the extra point p_i, Bob has to choose a tall rectangle of negligible area, which covers the left side of P_i. These two rectangles guarantee that no subsequent rectangle can cover any additional part of P_i, while $\text{area}(r_i) \leq \frac{\varepsilon}{2} \cdot \text{area}(P_i)$.

To determine the order of sequences of points associated with the staircases, we define a partial order over the staircase polygons. Note that in the initial staircase packing, each tip of every P_i is adjacent to the left or lower side of another staircase, or the right or upper side of U. This defines a partial order: let $P_i \prec P_j$, if the tip of P_j is adjacent to the left or lower side of P_i; or if the lower left corner of P_i is a reflex vertex of P_j. Order the staircases in any linear extension of this partial order. This ensures that Bob cannot choose a rectangle intersecting the interior of P_i before Alice reveals the lower left corner s_i. □

Best versus worst greedy strategy. For a finite set $S \subset [0,1]^2$, and a permutation (ordering) π of S, we can select anchored rectangles greedily (ties are broken arbitrarily) in the order prescribed by π. One can ask which permutation gives the best or the worst performance for a greedy strategy. Our main result, Theorem 8, says that for every n-element point set S, $(0,0) \in S$, we can find in $O(n \log n)$ time a permutation π for which the greedy strategy covers $\Theta(1)$ area. In the worst case, greedy covers only $o(1)$ area by Theorem 9. In the best case, however, we will show (Lemma 10) that greedy is always optimal for some permutation π. We say that an anchored rectangle packing is *Pareto optimal* if each rectangle $r(s)$ has maximum area assuming that all other rectangles are fixed. It is clear that every optimal solution is Pareto optimal.

In particular, each rectangle in an optimal solution is bounded by the two rays (going up and to the right) from its anchor point, and two other such rays (from other points) that limit it from the right and from the top. This immediately implies the existence of an exact algorithm for the optimization problem running in exponential time, based on brute force enumeration. The next lemma also shows that the greedy algorithm and brute force enumeration of permutations yields yet another exact algorithm for the optimization problem.

Lemma 10 *For every finite point set $S \subset [0,1]^2$ and every Pareto optimal anchored packing $R = \{r(s) : s \in S\}$, there is a permutation π for which a greedy algorithm (with some tie breaking) computes R.*

Proof. Let $S \subset [0,1]^2$ be a finite set, which may not contain the origin. Since every $r(s)$ is Pareto optimal, it is a greedy choice assuming that s is the last point in the order π. Suppose that $s \in S$ is the last point in a permutation π. If the smaller problem $S \setminus \{s\}$ with $R \setminus \{r(s)\}$ is not Pareto optimal, then there is a point $s' \in S$ for which we could choose a larger rectangle anchored at s', which intersects $r(s)$ only. In this case, either s dominates s' or the ray shot from s' vertically up (resp., horizontally right) hits the lower (resp., left) side of rectangle $r(s)$. This motivates the definition of a binary relation over S, which is an extension of the dominance order. Let $s' \prec s$ if either s dominates s' or an axis-aligned ray shot from s' hits the lower or left side of rectangle $r(s)$. It is not difficult to see that this is a partial order over S. If $s \in S$ is a minimal element in the poset $(S, \prec)$, then the rectangles in $R \setminus \{r(s)\}$ are still Pareto optimal for the anchors $S \setminus \{s\}$. Now let π be the reverse order of any linear extension of this partial order. □

We have shown (our main result) that for any set of n points in the unit square $U = [0,1]^2$, one can find a set of disjoint empty rectangles anchored at the given points and covering more than 9% of U. The same conclusion holds for points in any axis-aligned rectangle V instead of U, since it is straightforward to use an affine transformation to map the input into the unit square. Concerning the bound obtained, a sizable gap to the conjectured 50% remains. While certainly small adjustments in our proof can lead to improvements in the bound, obtaining substantial improvements probably requires new ideas.

Acknowledgement. The authors thank Richard Guy for tracing back the origins of this problem and Jan Kynčl for comments and remarks.

References

[1] Ponder this challenge: puzzle for June 2004; `http://domino.research.ibm.com/comm/wwwr_ponder.nsf/Challenges/June2004.html`.

[2] M. Abramovitz and I. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover, New York, 1964.

[3] N. Bansal, A. Caprara, and M. Sviridenko, A new approximation method for set covering problems, with applications to multidimensional bin packing, *SIAM J. Comput.* **39** (2009), 1256–1278.

[4] N. Bansal, J. Correa, C. Kenyon, and M. Sviridenko, Bin packing in multiple dimensions: inapproximability results and approximation schemes, *Math. Operat. Research* **31** (2006), 31–49.

[5] N. Bansal, X. Han, K. Iwama, M. Sviridenko, and G. Zhang, Harmonic algorithm for 3-dimensional strip packing problem, in *Proc. 18th SODA*, ACM-SIAM, 2007, pp. 1197–1206.

[6] A. Caprara, Packing 2-dimensional bins in harmony, in *Proc. 43rd FOCS*, IEEE, 2002, pp. 490–499.

[7] A. Caprara, Packing d-dimensional bins in d stages, *Math. Oper. Res.* **33** (2008), 203–215.

[8] M. Chlebík and J. Chlebíková, Hardness of approximation for orthogonal rectangle packing and covering problems, *J. Discrete Alg.* **7** (3) (2009), 291–305.

[9] T. Christ, A. Francke, H. Gebauer, J. Matoušek, and T. Uno, A doubly exponentially crumbled cake, manuscript, arXiv:1104.0122v1, 2011.

[10] H. T. Croft, K. J. Falconer, and R. K. Guy, *Unsolved Problems in Geometry*, volume II of Unsolved Problems in Intuitive Mathematics, Springer, New York, 1991.

[11] P. Erdős and R. L. Graham, Note on packing squares with equal squares, *J. Combin. Theory Ser. A* **19** (1975), 119–123.

[12] R. Harren, K. Jansen, L. Prädel, and R. van Stee, A $(5/3+\varepsilon)$-approximation for strip packing, in *Proc. 12th WADS*, vol. 6844 of LNCS, Springer, 2011, to appear.

[13] K. Jansen , L. Prädel, and U. M. Schwarz, Two for one: tight approximation of 2D bin packing, in *Proc. 11th WADS*, vol. 5664 of LNCS, Springer, 2009, pp. 399–410.

[14] K. Jansen and R. Solis-Oba, A polynomial time approximation scheme for the square packing problem, in *Proc. 13th IPCO*, vol. 5035 of LNCS, Springer, 2008, pp. 184–198.

[15] K. Jansen and G. Zhang, Maximizing the total profit of rectangles packed into a rectangle, *Algorithmica* **47** (3) (2007), 323–342.

[16] L. Prädel, Approximation algorithms for two-dimensional geometrical knapsack, Masters thesis, Department of Computer Science, University of Kiel, 2008.

[17] I. Schiermeyer, Reverse-fit: a 2-optimal algorithm for packing rectangles, in *Proc. 2nd ESA*, vol. 855 of LNCS, Springer, 1994, pp. 290–299.

[18] A. Steinberg, A strip-packing algorithm with absolute performance bound 2, *SIAM J. Comput.* **26** (2) (1997), 401–409.

[19] W. Tutte, *Recent Progress in Combinatorics: Proceedings of the 3rd Waterloo Conference on Combinatorics*, May 1968, Academic Press, New York, 1969.

[20] P. Winkler, Packing rectangles, pp. 133-134, in *Mathematical Mind-Benders*, A. K. Peters Ltd., Wellesley, MA, 2007.

[21] P. Winkler, Puzzled: rectangles galore, *Communications of the ACM*, **53** (11) (2010), 112.

[22] P. Winkler, Puzzled: solutions and sources, *Communications of the ACM*, **53** (12) (2010), 128.