Online Square Packing

Sándor P. Fekete, Tom Kamphans*, and Nils Schweer

Braunschweig University of Technology
Department of Computer Science, Algorithms Group
Mühlenpfordtstrasse 23, 38106 Braunschweig, Germany

Abstract. We analyze the problem of packing squares in an online fashion: Given a semi-infinite strip of width 1 and an unknown sequence of squares of side length in $[0, 1]$ that arrive from above, one at a time. The objective is to pack these items as they arrive, minimizing the resulting height. Just like in the classical game of Tetris, each square must be moved along a collision-free path to its final destination. In addition, we account for gravity in both motion and position. We apply a geometric analysis to establish a competitive factor of 3.5 for the bottom-left heuristic and present a $\frac{34}{13} \approx 2.6154$-competitive algorithm.

1 Introduction

In this paper, we consider online *strip packing* of squares. Squares arrive from above in an online fashion, one at a time, and have to be moved to their final positions in a semi-infinite, vertical strip of unit width. On its path, a square may move only through unoccupied space; in allusion to the well-known computer game, this is called the *Tetris constraint*. In addition, an item is not allowed to move upwards and has to be supported from below when reaching its final position (i.e., the bottom side of the square touches either another square or the bottom side of the strip). These conditions are called *gravity constraints*. Note that the gravity constraints make the problem harder, because we are not allowed to "hang squares in the air". The objective is to minimize the total height of the packing. Applications of this problem arise whenever physical access to the packed items is required. For example, objects stored in a warehouse need to be packed in a way such that the final positions can be accessed. Moreover, gravity is—obviously—a natural constraint in real-world packing applications.

Related Work. The strip packing problem was first considered by Baker et al. [1]. They showed that for online packing of *rectangles*, the bottom left heuristic does not necessarily guarantee a constant competitive ratio. For the offline case they proved an upper bound of 3 for a sequence of rectangles, and of 2 for squares. Kenyon and Rémila designed a FPTAS [2] for packing rectangles. For the online case, Csirik and Woeginger [3] gave a lower bound of 1.69103 on rectangle packings and an algorithm whose asymptotic worst-case ratio comes arbitrarily close to this value.

* Supported by DFG grant FE 407/8-3, project "ReCoNodes".

F. Dehne et al. (Eds.): WADS 2009, LNCS 5664, pp. 302–314, 2009.

For the one-dimensional online version of bin packing (pack a set of items into a minimum number of unit-capacity bins), the current best algorithm is $1,58889$-competitive [4]. The best known lower bound is $1,54014$ [5]. Coppersmith and Raghavan considered this problem in higher dimensions $d \geq 2$ [6]. They introduced a 2.6875-competitive algorithm and a lower bound of 4/3 for $d = 2$. Epstein and van Stee improved both bounds by designing an optimal, bounded-space, 2.3722-competitive algorithm [7].

Life gets considerably harder if items cannot be packed at arbitrary positions, but must be placed from above avoiding previously packed objects as obstacles—just like in the classical game of Tetris. In this setting, no item can ever move upward, no collisions between objects must occur, an item will come to a stop if and only if it is supported from below, and each placement has to be fixed before the next item arrives. Tetris is PSPACE-hard, even for the original game with a limited set of different objects; see Breukelaar et al. [8].

Azar and Epstein [9] considered tetris-constraint online packing of rectangles into a strip. For the case without rotation, they showed that no constant competitive ratio is possible, unless there is a fixed-size lower bound of ε on the side length of the objects, in which case there is an upper bound of $O(\log \frac{1}{\varepsilon})$. For the case with rotation, they showed a 4-competitive strategy, based on shelf-packing methods; until now, this is also the best deterministic upper bound for squares. Observe that their strategy does not take the gravity constraints into account, as items are allowed to be placed at appropriate levels, even if they are unsupported. Coffmann et al. [10] considered probabilistic aspects of online rectangle packing with Tetris constraint, without allowing rotations. If rectangle side lengths are chosen uniformly at random from the interval $[0, 1]$, they showed that there is a lower bound of $(0.3138...)n$ on the expected height of the strip. Using another strategy, which arises from the bin-packing–inspired *Next Fit Level*, they established an upper bound of $(0.3697...)n$ on the expected height.

Our Results. In this paper, we demonstrate that it pays off to take a closer look at the geometry of packings. We analyze a natural and simple heuristic called *BottomLeft*, similar to the one introduced by Baker et al. [1]. We show that it is possible to give a better competitive guarantee than 4 (as achieved by Azar and Epstein), even in the presence of gravity. We obtain an asymptotic competitive ratio of 3.5 for *BottomLeft*, implying an asymptotic density of at least 0.2857... Improving this ratio even further, we introduce the strategy *SlotAlgorithm* and establish a competitive ratio $34/13 = 2.6154...$

2 Preliminaries

We are given a vertical strip, S, of width 1 and a sequence, $\mathcal{A} = (A_1, \dots, A_n)$, of squares with side lengths $a_i \leq 1$. Our goal is to find a non-overlapping, axis-parallel placement of squares in the strip that keeps the height of the strip as low as possible. A packing has to fulfill the Tetris and the gravity constraints. Moreover, we consider the online problem.

We denote the bottom (left, right) side of the strip by B_S (R_S, L_S; respectively), and the sides of a square, A_i, by B_{A_i}, T_{A_i}, R_{A_i}, L_{A_i} (bottom, top, right, left; respectively). The x-coordinates of the left and right side of A_i in a packing are l_{A_i} and r_{A_i}; the y-coordinates of the top and bottom side are t_{A_i} and b_{A_i}, respectively. Let the *left neighborhood*, $N_L(A_i)$, be the set of squares that touch the left side of A_i. In the same way we define the bottom, top, and right neighborhoods, denoted by $N_B(A_i), N_T(A_i)$, and $N_R(A_i)$, respectively.

A packing may leave areas of the strip empty. We call a maximal connected component of the strip's empty area a *hole*, denoted by H_h, $h \in \mathbb{N}$. A point, P, is called *unsupported*, if there is a vertical line segment from P downwards whose interior lies inside a hole. Otherwise, P is *supported*. A section of a line segment is supported, if every point in this section is supported. For an object ξ we refer to the boundary by $\partial\xi$, to the interior by ξ°, and to its area by $|\xi|$.

3 The Strategy *BottomLeft*

In this section, we analyze the packing generated by the strategy *BottomLeft*, which works as follows: We place the current square as close as possible to the bottom of the strip (provided that there is a collision-free path from the top of the strip to the desired position that never moves in positive y-direction). We break ties by choosing the leftmost among all possible bottommost positions.[1]

For a simplified analysis, we finish the packing with an additional square, A_{n+1}, of side length 1. This implies that all holes have a closed boundary. Let $H_1, \dots, H_s$ be the holes in the packing. Then the height of the packing produced by *BottomLeft* is $BL = \sum_{i=1}^{n} a_i^2 + \sum_{h=1}^{s} |H_h|$. In the following sections, we prove $\sum_{h=1}^{s} |H_h| \leq 2.5 \cdot \sum_{i=1}^{n+1} a_i^2$. Because any strategy needs at least a height of $\sum_{i=1}^{n} a_i^2$, our bound implies that asymptotically $BL \leq 3.5 \cdot OPT$.

We proceed as follows. First, we state some properties of the generated packing (Section 3.1). In Section 3.2 we simplify the shape of the holes by partitioning a hole into several disjoint new parts.[2] In the packing, these new holes are open at their top side, so we introduce *virtual lids* that close these holes. Afterwards, we estimate the area of a hole in terms of the squares that enclose the hole (Section 3.3). Summing up the charges to a single square (Table 1) we get

Theorem 1. BottomLeft *is (asymptotically) 3.5-competitive.*

3.1 Basic Properties of the Generated Packing

In this section, we analyze structural properties of the boundary of a hole. We say that a square, A_i, *contributes* to the boundary of a hole, H_h, iff ∂A_i and ∂H_h intersect in more than one point. Let $\tilde{A}_1, \dots, \tilde{A}_k$ denote the squares on the

[1] To implement the strategy, we can use robot-motion-planning techniques. For k placed squares, this can be done in time $O(k \log^2 k)$; see de Berg et al. [11].

[2] Let the new parts replace the original hole, so that we do not have to distinguish between 'holes' and 'parts of a hole'.

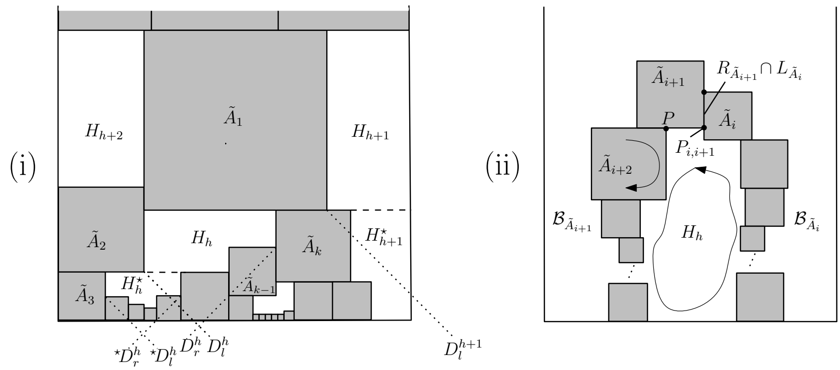

Fig. 1. (i) A packing produced by *BottomLeft*. The squares $\tilde{A}_1, \ldots, \tilde{A}_k$ contribute to the boundary of the hole H_h, which is split into a number of subholes. In the shown example one new subhole $H^{\star}_h$ is created. Note that the square $\tilde{A}_1$ also contributes to the holes H_{h+1} and H_{h+2} and serves as a virtual lid for $H^{\star}_{h+1}$. (ii) The hole H_h with the two squares $\tilde{A}_i$ and $\tilde{A}_{i+1}$ and their bottom sequences. Here, $\tilde{A}_{i+1}$ is $\tilde{A}_1$. If ∂H_h is traversed in ccw order, then $\partial H_h \cap \partial \tilde{A}_{i+2}$ is traversed in cw order w.r.t. to $\partial \tilde{A}_{i+2}$.

boundary of H_h in counterclockwise order starting with the upper left square.[3] We call $\tilde{A}_1$ the *lid* of H_h and define $\tilde{A}_{k+1} = \tilde{A}_1$, $\tilde{A}_{k+2} = \tilde{A}_2$ and so on. By $P_{i,i+1}$ we denote the point where ∂H_h leaves $\partial \tilde{A}_i$ and enters $\partial \tilde{A}_{i+1}$.

Let A_i be a packed square. We define the *left (bottom) sequence*, $\mathcal{L}_{A_i}$, $(\mathcal{B}_{A_i})$, of A_i, as follows: The first element of $\mathcal{L}_{A_i}$ $(\mathcal{B}_{A_i})$ is A_i. The next element is chosen as an arbitrary left (bottom) neighbor of the previous element. The sequence ends if no such neighbor exists. We call the polygonal chain from the upper right corner of the first element of $\mathcal{L}_{A_i}$ to the upper left corner of the last element while traversing the boundary of the sequence in counterclockwise order the *skyline*, $\mathcal{S}_{A_i}$, of A_i. Obviously, $\mathcal{S}_{A_i}$ has an endpoint on L_S. Further, $\mathcal{S}^{\circ}_{A_i} \cap H^{\circ}_h = \emptyset$.

Lemma 1. *Let $\tilde{A}_i$ be a square that contributes to ∂H_h. Then,*
(i) $\partial H_h \cap \partial \tilde{A}_i$ is a single curve, and
(ii) if ∂H_h is traversed in counterclockwise (clockwise) order, then $\partial H_h \cap \partial \tilde{A}_i$ is traversed in clockwise (counterclockwise) order w.r.t. $\partial \tilde{A}_i$; see Fig. 1(ii).

Proof. (i) Assume that $\partial H_h \cap \partial \tilde{A}_i$ consists of (at least) two curves, c_1 and c_2. Consider a simple curve, C, that lies inside H_h and has one endpoint in c_1 and the other one in c_2. We add the straight line between the endpoints to C and obtain a simple closed curve C'. As c_1 and c_2 are not connected, there is a square $\tilde{A}_j$ inside C' that is a neighbor of $\tilde{A}_i$. If $\tilde{A}_j$ is a left, right or bottom neighbor of $\tilde{A}_i$ this contradicts the existence of $\mathcal{B}_{\tilde{A}_j}$; if it is a top neighbor this contradicts the existence of $\mathcal{L}_{\tilde{A}_j}$. Hence, $\partial H_h \cap \partial \tilde{A}_i$ is a single curve.

[3] It is always clear from the context which hole defines this sequence of squares. Thus, we chose not to introduce an additional superscript referring to the hole.

(ii) Imagine that we walk along ∂H_h in ccw order: The interior of H_h lies on our left-hand side and all squares that contribute to ∂H_h lie on our right-hand side. Hence, their boundaries are traversed in cw order w.r.t. their interior. □

Let P and Q be the left and right endpoint, respectively, of the line segment $\partial \tilde{A}_1 \cap \partial H_h$. The next lemma restricts the relative position of two squares:

Lemma 2. *Let $\tilde{A}_i$, $\tilde{A}_{i+1}$ contribute to the boundary of a hole H_h.*
(i) If $\tilde{A}_{i+1} \in N_L(\tilde{A}_i)$ then either $\tilde{A}_{i+1} = \tilde{A}_1$ or $\tilde{A}_i = \tilde{A}_1$.
(ii) If $\tilde{A}_{i+1} \in N_T(\tilde{A}_i)$ then $\tilde{A}_{i+1} = \tilde{A}_1$ or $\tilde{A}_{i+2} = \tilde{A}_1$.
(iii) There are two types of holes: Type I with $\tilde{A}_k \in N_R(\tilde{A}_{k-1})$, and Type II with $\tilde{A}_k \in N_T(\tilde{A}_{k-1})$; see Fig. 3.

Proof. (i) Let $\tilde{A}_{i+1} \in N_L(\tilde{A}_i)$. Consider the endpoints of the vertical line $R_{\tilde{A}_{i+1}} \cap L_{\tilde{A}_i}$; see Fig. 1(ii). We traverse ∂H_h in counterclockwise order starting in P. By Lemma 1, we traverse $\partial \tilde{A}_i$ in clockwise order and, therefore, $P_{i,i+1}$ is the lower endpoint of $R_{\tilde{A}_{i+1}} \cap L_{\tilde{A}_i}$. Now, $\mathcal{B}_{\tilde{A}_i}$, $\mathcal{B}_{\tilde{A}_{i+1}}$, and the segment of B_S completely enclose an area that completely contains the hole, H_h. If the sequences share a square, A_j, we consider the area enclosed up to the first intersection. Therefore, if $b_{\tilde{A}_{i+1}} \geq b_{\tilde{A}_i}$ then $\tilde{A}_{i+1} = \tilde{A}_1$ else $\tilde{A}_i = \tilde{A}_1$ by the definition of $\overline{PQ}$.

The proof of (ii) follows almost directly from (i). Let $\tilde{A}_{i+1} \in N_T(\tilde{A}_i)$. We know that $\partial \tilde{A}_{i+1}$ is traversed in clockwise order and we know that $\tilde{A}_{i+1}$ has to be supported to the left. Therefore, $\tilde{A}_{i+2} \in N_L(\tilde{A}_{i+1}) \cup N_B(\tilde{A}_{i+1})$ and the result follows from (i). For (iii) we traverse ∂H_h from P in clockwise order. From the definition of $\overline{PQ}$ and Lemma 1 we know that $P_{k,1}$ is a point on $L_{\tilde{A}_k}$. If $P_{k-1,k} \in L_{\tilde{A}_k}$, then $\tilde{A}_k \in N_R(\tilde{A}_{k-1})$; if $P_{k-1,k} \in B_{\tilde{A}_k}$, then $\tilde{A}_k \in N_T(\tilde{A}_{k-1})$. In any other case $\tilde{A}_k$ does not have a lower neighbor. □

3.2 Splitting Holes

Let H_h be a hole whose boundary does not touch the boundary of the strip. We define two lines: The *left diagonal*, D_l^h, is defined as the straight line with slope -1 starting in $P_{2,3}$ if $P_{2,3} \in R_{\tilde{A}_2}$ or, otherwise, in the lower right corner of $\tilde{A}_2$; see Fig. 3. We denote the point in which D_l^h starts by P'. The *right diagonal*, D_r^h, is defined as the line with slope 1 starting in $P_{k-1,k}$ if $\tilde{A}_k \in N_R(\tilde{A}_{k-1})$ (Type I) or in $P_{k-2,k-1}$, otherwise (Type II). Note that $P_{k-2,k-1}$ lies on $L_{\tilde{A}_{k-1}}$, otherwise there would not be a left neighbor of $\tilde{A}_{k-1}$. We denote the point in which D_r^h starts by Q'. If h is clear or does not matter, we omit the superscript.

Lemma 3. *Let H_h be a hole, D_r its right diagonal. Then $D_r \cap H_h^\circ = \emptyset$ holds.*

Proof. Consider the left sequence, $\mathcal{L}_{\tilde{A}_k} = (\tilde{A}_k = \alpha_1, \alpha_2, \ldots)$ or $\mathcal{L}_{\tilde{A}_{k-1}} = (\tilde{A}_{k-1} = \alpha_1, \alpha_2, \ldots)$ for H_h being of Type I or II, respectively. By induction, the upper left corners of the α_i's lie above D_r: If D_r intersects $\partial \alpha_i$ at all, the first intersection is on R_{α_i}, the second on B_{α_i}. Thus, at least the skyline separates D_r and H_h. □

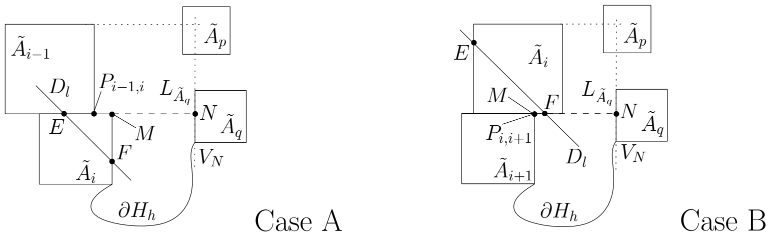

Fig. 2. D_l can intersect $\tilde{A}_i$ (for the second time) in two different ways: on the right side or on the bottom side. In Case A, the square $\tilde{A}_{i-1}$ is on top of $\tilde{A}_i$; in Case B, $\tilde{A}_i$ is on top of $\tilde{A}_{i+1}$.

It is a simple observation that if D_l intersects a square $\tilde{A}_i$ in a nontrivial way[4] then either $F \in R_{\tilde{A}_i}$ and $E \in T_{\tilde{A}_i}$ or $F \in B_{\tilde{A}_i}$ and $E \in L_{\tilde{A}_i}$. To break ties, we define that an intersection in the lower right corner of $\tilde{A}_i$ belongs to $B_{\tilde{A}_i}$ (Fig. 1(i) and 2). Unfortunately, Lemma 3 does not hold for D_l. Therefore, we split our hole, H_h, into two new holes, $H_h^{(1)}$ and $H_h^\star$, as follows: Let F be the first nontrivial intersection point of ∂H_h and D_l while traversing ∂H_h in counterclockwise fashion, starting in P. We consider two cases, $F \in R_{\tilde{A}_i} \setminus B_{\tilde{A}_i}$ (Case A) and $F \in B_{\tilde{A}_i}$ (Case B); see Fig. 2.

Let E be the other intersection point of D_l and $\partial\tilde{A}_i$. In Case A, let $\tilde{A}_{\text{up}} := \tilde{A}_{i-1}$ and $\tilde{A}_{\text{low}} := \tilde{A}_i$, in Case B $\tilde{A}_{\text{up}} := \tilde{A}_i$ and $\tilde{A}_{\text{low}} := \tilde{A}_{i+1}$. The horizontal ray that emanates from the upper right corner of $\tilde{A}_{\text{low}}$ to the right is subdivided into supported and unsupported sections. Let $U = \overline{MN}$ be the leftmost unsupported section. Now we split H_h into two parts, $H_h^\star$ below $\overline{MN}$ and $H_h^{(1)} := H_h \backslash H_h^\star$.

We split $H_h^{(1)}$ into $H_h^{(2)}$ and $H_h^{\star\star}$ etc., until there is no further intersection between the boundary of $H_h^{(z)}$ and D_l^h. Every split is caused by a pair of squares. It can be shown that $\overline{MN} < \tilde{a}_{up}$ and, therefore, a copy of $\tilde{A}_{\text{up}}$, denoted by $\tilde{A}'_{\text{up}}$, placed on $\overline{MN}$ can serve as a *virtual lid* for the hole below. Moreover, a square serves as a virtual lid for at most one hole. Regarding the holes, they are either of Type I or Type II and, thus, can be analyzed in the same way as original holes. See the full version of this paper for a rigorous proof.

3.3 Computing the Area of a Hole

We eliminated all intersections of D_l^h with the boundary of the hole $H_h^{(z)}$ by splitting the hole. Thus, we have a set of holes $\hat{H}_h$, $h = 1, \ldots, s^\star$, that fulfill $\partial\hat{H}_h \cap D_l^h = \emptyset$ and have either a non-virtual or a virtual lid.

Our aim is to bound $|\hat{H}_h|$ by the areas of the squares that contribute to $\partial\hat{H}_h$. A square A_i may contribute to more than one hole. It is too expensive to use

[4] An intersection, $p \in \Gamma \cap \Delta$, of a Jordan curve Δ and a line, ray, or line segment Γ is called *nontrivial*, iff there is a line segment ℓ of length $\varepsilon > 0$ on the line through Γ such that p is in the interior of ℓ and the endpoints of ℓ lie on different sides of Δ.

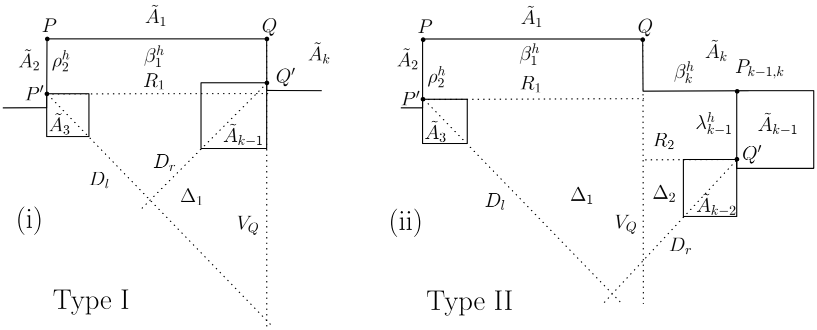

Fig. 3. Holes of Type I and Type II with their left and right diagonals

its total area a_i^2 in the bound for a single hole. Instead, we charge only fractions of a_i^2 per hole. Moreover, we charge every edge of A_i separately. By Lemma 1, $\partial\hat{H}_h \cap \partial A_i$ is connected. In particular, every side of A_i contributes at most one line segment to $\partial\hat{H}_h$. For the left (bottom, right) side of a square A_i, we denote the length of the line segment contributed to $\partial\hat{H}_h$ by λ_i^h (β_i^h, ρ_i^h; respectively).[5]

Let $c_{h,i}^{\{\lambda,\beta,\rho\}}$ be appropriate coefficients, such that the area of a hole can be charged against the area of the adjacent squares; i.e., $|\hat{H}_h| \leq \sum_{i=1}^{n+1} c_{h,i}^{\lambda}(\lambda_i^h)^2 + c_{h,i}^{\beta}(\beta_i^h)^2 + c_{h,i}^{\rho}(\rho_i^h)^2$. As each point on ∂A_i is on the boundary of at most one hole, the line segments are pairwise disjoint. Thus, for the left side of A_i, the two squares inside A_i induced by the line segments λ_i^h and λ_i^g of two different holes, $\hat{H}_h$ and $\hat{H}_g$, do not overlap. Therefore, we obtain $\sum_{h=1}^{s^\star} c_{h,i}^{\lambda}\cdot(\lambda_i^h)^2 \leq c_i^{\lambda}\cdot a_i^2$, where $c_i^{\lambda} := \max_h c_{h,i}^{\lambda}$. We call c_i^{λ} the *charge of* L_{A_i} and define c_i^{β} and c_i^{ρ} analogously.

We use virtual copies of some squares as lids. However, for every square, A_i, there is at most one copy, A_i'. We denote the line segments and charges corresponding to A_i' by $\lambda_{i'}^h$, $c_{h,i'}^{\lambda}$ and so on. The *total charge of* A_i is given by $c_i = c_i^{\lambda}+c_i^{\beta}+c_i^{\rho}+c_{i'}^{\lambda}+c_{i'}^{\beta}+c_{i'}^{\rho}$. Altogether, we bound $\sum_{h=1}^{s^\star}|\hat{H}_h| \leq \sum_{i=1}^{n+1} c_i\cdot a_i^2 \leq \sum_{i=1}^{n+1} c\cdot a_i^2$, with $c = \max_i c_i$. Next, we want to find an upper bound on c.

Holes with a Non-Virtual Lid. We removed all intersections of $\hat{H}_h$ with its diagonal D_l^h. Therefore, $\hat{H}_h$ lies completely inside the polygon formed by D_l^h, D_r^h and the part of $\partial\hat{H}_h$ that is clockwise between P' and Q'; see Fig. 3. If $\hat{H}_h$ is of Type I, we consider the rectangle, R_1, of area $\rho_2^h\cdot\beta_1^h$ induced by P, P' and Q. Let Δ_1 be the triangle below R_1 formed by the bottom side of R_1, D_l^h, and the vertical line V_Q passing through Q; see Fig. 3(i). Obviously, $|\hat{H}_h| \leq |R_1|+|\Delta_1|$. As D_l^h has slope -1, we get $|\Delta_1| = \frac{1}{2}(\beta_1^h)^2$. We have $|R_1| = \rho_2^h\cdot\beta_1^h \leq \frac{1}{2}(\rho_2^h)^2+\frac{1}{2}(\beta_1^h)^2$. Thus, for a Type I-hole we get $|\hat{H}_h| \leq (\beta_1^h)^2 + \frac{1}{2}(\rho_2^h)^2$, i.e., we charge the bottom side of $\tilde{A}_1$ with 1 and the right side of $\tilde{A}_2$ with $\frac{1}{2}$. In this case, we get $c_{h,1}^{\beta} = 1$

[5] If a side of a square does not contribute to a hole, the corresponding length of the line segment is defined to be zero.

and $c^{\rho}_{h,2} = \frac{1}{2}$. For a Type II hole, we additionally get a rectangle R_2 and a triangle, Δ_2, as in Fig. 3(ii). Using similar arguments as above we get charges $c^{\beta}_{h,1} = c^{\beta}_{h,k} = 1$ and $c^{\rho}_{h,2} = \frac{1}{2} = c^{\lambda}_{h,k-1} = \frac{1}{2}$.

Holes with a Virtual Lid. Let $\hat{H}_h$ be a hole with a virtual lid, $\hat{H}_g$ be immediately above $\hat{H}_h$, $\tilde{A}_{\text{up}}$ be the square whose copy, $\tilde{A}'_{\text{up}}$, becomes a new lid, and $\tilde{A}_{\text{low}}$ the bottom neighbor of $\tilde{A}_{\text{up}}$. We show that $\tilde{A}'_{\text{up}}$ increases the charge of $\tilde{A}_{\text{up}}$ by at most $\frac{1}{2}$: If $\tilde{A}_{\text{up}}$ does not exceed $\tilde{A}_{\text{low}}$ to the left, it cannot serve as a lid for any other hole (Fig. 4). Hence, the charge of the bottom side of $\tilde{A}_{\text{up}}$ is 0; like in the preceding section, we obtain a charge ≤ 1 to the bottom of $\tilde{A}'_{\text{up}}$. If it exceeds $\tilde{A}_{\text{low}}$ to the left, we know that the part $B_{\tilde{A}_{\text{up}}} \cap T_{\tilde{A}_{\text{low}}}$ of $B_{\tilde{A}_{\text{up}}}$ is not charged by another hole, because it does not belong to a hole and the lid is defined uniquely.

We define points P and P' for $\hat{H}_h$ in the same way as in the preceding section. Independent of $\hat{H}_h$'s type, $\tilde{A}'_{\text{up}}$ gets charged only for the rectangle R_1 induced by P, P' and N, as well as for the triangle below R_1 (Fig. 3). Now we show that we do not have to charge $\tilde{A}'_{\text{up}}$ for R_1, since the part of R_1 above D^g_l is already included in the bound for $\hat{H}_g$, and the remaining part can be charged to $B_{\tilde{A}_{\text{up}}}$ and $R_{\tilde{A}_{\text{low}}}$. $\tilde{A}'_{\text{up}}$ gets charged $\frac{1}{2}$ for the triangle.

D^g_l splits R_1 into a part that is above this line, and a part that is below this line. The latter part of R_1 is not included in the bound for $\hat{H}_g$. Let F be the intersection of $\partial\hat{H}_g$ and D^g_l that caused the creation of $\hat{H}_h$. If $F \in R_{\tilde{A}_{\text{low}}}$, this part is at most $\frac{1}{2}(\rho^h_{\text{low}})^2$, where ρ^h_{low} is the length of $\overline{P'F}$. We charge $\frac{1}{2}$ to $R_{\tilde{A}_{\text{low}}}$.

If $F \in B_{\tilde{A}_{\text{up}}}$, the part of R_1 below D^g_l can be split into a rectangular part of area $\rho^h_{\text{low}} \cdot \beta^h_{\text{up}}$, and a triangular part of area $\frac{1}{2}(\rho^h_{\text{low}})^2$. Here β^h_{up} is the length of $\overline{PF}$. The cost of the triangle is charged to $R_{\tilde{A}_{\text{low}}}$. Note that the part of $B_{\tilde{A}_{\text{up}}}$ that exceeds $\tilde{A}_{\text{low}}$ to the right is not charged and ρ^h_{low} is not larger than $B_{\tilde{A}_{\text{up}}} \cap T_{\tilde{A}_{\text{low}}}$ (i.e., the part of $B_{\tilde{A}_{\text{up}}}$ that was not charged before). Thus, we can charge the rectangular part completely to $B_{\tilde{A}_{\text{up}}}$. Hence, $\tilde{A}'_{\text{up}}$ is charged $\frac{1}{2}$ in total.

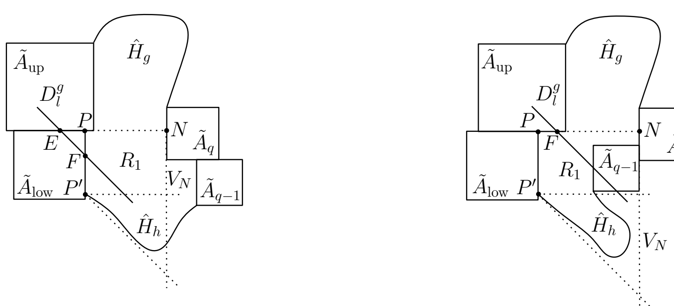

Fig. 4. The holes $\hat{H}_g$ and $\hat{H}_h$ and the rectangle R_1 which is divided into two parts by D^g_l. The upper part is already included in the bound for $\hat{H}_g$. The lower part is charged completely to $R_{\tilde{A}_{\text{low}}}$ and $B_{\tilde{A}'_{\text{up}}}$. Here, P and P' are defined w.r.t. $\hat{H}_h$.

Holes Containing Parts of ∂S. We show in this section that holes that touch ∂S are just special cases of the ones discussed in the preceding sections.

Because the top side of a square never gets charged for a hole, it does not matter whether a part of B_S belongs to the boundary. Moreover, for any hole $\hat{H}_h$ either L_S or R_S can be a part of $\partial\hat{H}_h$, because otherwise there exits a curve with one endpoint on L_S and the other endpoint on R_S, with the property that this curve lies completely inside of $\hat{H}_h$. This contradicts the existence of the bottom sequence of a square lying above the curve.

For a hole $\hat{H}_h$ touching L_S, $L_S \cap \partial\hat{H}_h$ is a single line segment (similar to Lemma 1). Let P be the topmost point of this line segment and $\tilde{A}_1$ be the square containing P. The existence of $\mathcal{B}_{\tilde{A}_1}$ implies that $\tilde{A}_1$ is the lid of $\hat{H}_h$. As $\tilde{A}_1$ must have a bottom neighbor, $\tilde{A}_k$, and $\tilde{A}_k$ must have a right neighbor, $\tilde{A}_{k-1}$, we get $P_{k,1} \in B_{\tilde{A}_1}$ and $P_{k-1,k} \in L_{\tilde{A}_k}$, respectively. We define the right diagonal D_r and the point Q' as above and conclude that $\hat{H}_h$ lies completely inside the polygon formed by $L_S \cap \partial\hat{H}_h$, D_r and the part of $\partial\hat{H}_h$ that is between P and Q' (in clockwise order). We split this polygon into a rectangle and a triangle in order to obtain charges of 1 to $B_{\tilde{A}_1}$ and $\frac{1}{2}$ to $L_{\tilde{A}_k}$.

Now consider a hole where a part of R_S belongs to $\partial\hat{H}_h$. We denote the topmost point on $R_S \cap \partial\hat{H}_h$ by Q, and the square containing Q by $\tilde{A}_1$. $\tilde{A}_1$ is the lid of this hole. As above, we eliminate the intersections of D_l and $\partial\hat{H}_h$ by creating new holes. After this, the modified hole $\hat{H}_h^{(z)}$ can be viewed as a hole of Type II, for which the part on the right side of V_Q has been cut off. We obtain charges of 1 to $B_{\tilde{A}_1}$, $\frac{1}{2}$ to $R_{\tilde{A}_2}$, and $\frac{1}{2}$ to the bottom of a virtual lid.

Table 1. Charges to different sides of a single square. The charges depend on the type of the adjacent hole (Type I, II, touching or not touching the strip's boundary), but the maximal charge dominates the other one. Moreover, the square may also serve as a virtual lid. These charges sum up to a total charge of 2.5 per square.

	Non-virtual Lid					Virtual Lid				Total
	Type I	Type II	L_S	R_S	Max.	Type I	Type II	R_S	Max.	
Left side	0	0.5	0.5	0	0.5	0	0	0	0	0.5
Bottom side	1	1	1	1	1	0.5	0.5	0.5	0.5	1.5
Right Side	0.5	0.5	0	0.5	0.5	0	0	0	0	0.5
Total					2				0.5	2.5

4 The Strategy *SlotAlgorithm*

Consider two vertical lines going upward from the bottom side of S and parallel to sides of S. We call the area between these lines a *slot*, the lines the slot's *left* and *right boundary*, and the distance between the lines the *width* of the slot.

Our strategy *SlotAlgorithm* works as follows: We divide the strip S of width 1 into one slot of width 1, two slots of width 1/2, four slots of width 1/4 etc.

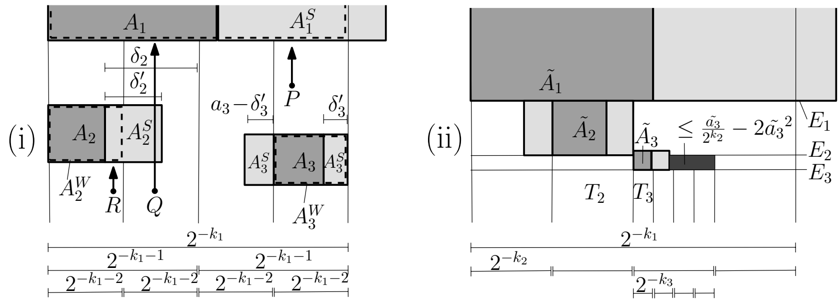

Fig. 5. (i) Squares A_i with *shadows* A_i^S and *widenings* A_i^W. $\delta'_2 = a_2$ and $\delta'_3 = \delta_3$. P and Q are charged to A_1. R is charged to A_2. (ii) The first three squares of the sequence. Here, $\tilde{A}_2$ is the smallest square that bounds $\tilde{A}_1$ from below. $\tilde{A}_3$ is the smallest one that intersects E_2 in an active slot (w.r.t. E_2) of width $1/2^{k_2}$. T_2 is nonactive (w.r.t. E_2) and also w.r.t. all E_j, $j \geq 3$. The part of $F_{\tilde{A}_1}$ (darkest gray) between E_2 and E_3 in an active slot of width 2^{-k_2} is $\leq \tilde{a}_3/2^{k_2} - 2\tilde{a}_3^2$ as points in $\tilde{A}_3^W$ are not charged to $\tilde{A}_1$.

(i.e., creating 2^j of width 2^{-j}). Note that a slot of width 2^{-j} contains 2 slots of width 2^{-j-1}; see Fig. 5(i). For every square A_i we round the side length a_i to the smallest number $1/2^{k_i}$ that is larger than or equal to a_i. We place A_i in the slot of width 2^{-k_i} that allows A_i to be placed as near to the bottom of S as possible by moving A_i down along the left boundary of the chosen slot until another square is reached. *SlotAlgorithm* satisfies the Tetris and the Gravity constraints.

Theorem 2. SlotAlgorithm *is (asymptotically) 2.6154-competitive.*

Proof. Let A_i be a square placed by *SlotAlgorithm* in a slot T_i of width 2^{-k_i}. Let δ_i be the distance between the right side of A_i and the right boundary of the slot of width 2^{-k_i+1} that contains A_i and $\delta'_i := \min\{a_i, \delta_i\}$. We call the area obtained by enlarging A_i by δ'_i to the right and by $a_i - \delta'_i$ to the left the *shadow* of A_i and denote it by A_i^S. Thus, A_i^S is an area of the same size as A_i and lies completely inside a slot of twice the width of A_i's slot. Moreover, we define the *widening* of A_i as $A_i^W = (A_i \cup A_i^S) \cap T_i$; see Fig. 5(i).

Now, consider a point P in T_i that is not inside an A_j^W for any square A_j. We charge P to the square A_i if A_i^W is the first widening that intersects the vertical line going upwards from P. Let F_{A_i} be the set of all points charged to A_i. For the analysis, we place a closing square, A_{n+1}, of side length 1 on top of the packing. Therefore, every point in the packing that does not lie inside an A_j^W is charged to a square. Because A_i and A_i^S have the same area, we can bound the height of the packing by $2\sum_{i=1}^{n} a_i^2 + \sum_{i=1}^{n+1} |F_{A_i}|$.

The height of an optimal packing is at least $\sum_{i=1}^{n} a_i^2$ and, therefore, it suffices to show $|F_{A_i}| \leq 0.6154 \cdot a_i^2$. We construct for every A_i a sequence of squares $\tilde{A}_1^i, \tilde{A}_2^i, \ldots, \tilde{A}_m^i$ with $\tilde{A}_1^i = A_i$. (To ease notation, we omit the superscript i in the following.) We denote by E_j the extension of the bottom side of $\tilde{A}_j$ to the left and to the right; see Fig. 5(ii). We will show that by an appropriate choice of the sequence we can bound the area of the part of $F_{\tilde{A}_1}$ that lies between a consecutive pair of extensions, E_j and E_{j+1}, in terms of $\tilde{A}_{j+1}$ and the slot widths. From this we will derive the upper bound on the area of $F_{\tilde{A}_1}$. We assume throughout the proof that the square $\tilde{A}_j$, $j \geq 1$, is placed in a slot, T_j, of width 2^{-k_j}. Note that $F_{\tilde{A}_1}$ is completely contained in T_1. A slot is called *active (w.r.t. E_j and $\tilde{A}_1$)* if there is a point in the slot that lies below E_j and that is charged to $\tilde{A}_1$ and *nonactive* otherwise. If it is clear from the context we leave out the $\tilde{A}_1$.

The sequence of squares is chosen as follows: $\tilde{A}_1$ is the first square and $\tilde{A}_{j+1}$, $j = 1, \ldots, m-1$ is chosen as the smallest one that intersects or touches E_j in an active slot (w.r.t. E_j and $\tilde{A}_1$) of width 2^{-k_j} and that is not equal to $\tilde{A}_j$. The sequence ends if all slots are nonactive w.r.t. to an extension E_m. We claim:

(i) $\tilde{A}_{j+1}$ exists for $j+1 \leq m$ and $\tilde{a}_{j+1} \leq 2^{-k_j-1}$ for $j+1 \leq m-1$.
(ii) The number of active slots (w.r.t. E_j) of width 2^{-k_j} is at most 1 for $j=1$ and $\prod_{i=2}^{j}(\frac{1}{2^{k_{i-1}}}2^{k_i}-1)$ for $j \geq 2$.
(iii) The area of the part of $F_{\tilde{A}_1}$ that lies in an active slot of width 2^{-k_j} between E_j and E_{j+1} is at most $2^{-k_j}\tilde{a}_{j+1} - 2\tilde{a}_{j+1}^2$.

We prove the claims by induction. Assume that the $(j+1)$st element does not exist for $j+1 \leq m$. Let T' be an active slot in T_1 (w.r.t. E_j) of width 2^{-k_j} for which E_j is not intersected by a square in T'. If there is a rectangle of height ε below $T' \cap E_j$ for which every point is charged to $\tilde{A}_1$, SlotAlgorithm would have chosen this slot for $\tilde{A}_j$. Hence, at least one point, Q, below E_j is not charged to $\tilde{A}_1$. Consider the bottom sequence (see Section 3.1) of the square, $\hat{A}$, to which Q is charged. This sequence has to intersect E_j outside of T' (by choice of T'). But then one of its elements has to intersect the left or the right boundary of T' and we can conclude that this square has at least the width of T', because (by the algorithm) a square with rounded side length $2^{-\ell}$ cannot cross a slot's boundary of width larger than $2^{-\ell}$. Hence, a setting as shown in the figure is not possible. In turn, a square larger than T' completely covers T' and T' cannot be active w.r.t. to E_j and $\tilde{A}_1$. Thus, all points in T' below E_j are charged to this square; a contradiction. This proves the existence of $\tilde{A}_{j+1}$. Because we chose $\tilde{A}_{j+1}$ to be of minimal side length, $\tilde{a}_{j+1} \geq 2^{-k_j}$ would imply that all slots inside T are nonactive (w.r.t. E_j). Therefore, if $\tilde{A}_{j+1}$ is not the last element of the sequence, $\tilde{a}_{j+1} \leq 2^{-k_j-1}$ holds.

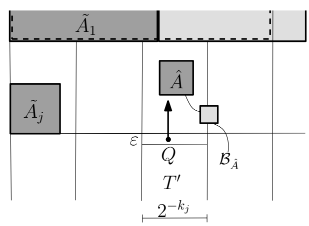

By the induction hypothesis there are at most $(2^{-k_1}2^{k_2}-1) \cdot (2^{-k_2}2^{k_3}-1) \cdot \ldots \cdot (2^{-k_{j-2}}2^{k_{j-1}}-1)$ active slots of width $2^{-k_{j-1}}$ (w.r.t. E_{j-1}). Each of these slots contains $2^{k_j-k_{j-1}}$ slots of width 2^{-k_j} and in every active slot of width $2^{-k_{j-1}}$ at least one slot of width 2^{-k_j} is nonactive because we chose $\tilde{A}_j$ to be of minimum side length. Hence, the number of active slots (w.r.t. E_j) is a factor of $(\frac{1}{2^{k_{j-1}}}2^{k_j}-1)$ larger than the number of active slots (w.r.t. E_{j-1}).

By the choice of $\tilde{A}_{j+1}$ and the fact that in every active slot of width 2^{-k_j} there is at least one square that intersects E_j (points below its widening are not charged to $\tilde{A}_1$) we conclude that the area of $F_{\tilde{A}_1}$ between E_j and E_{j+1} is at most $2^{-k_j}\tilde{a}_{j+1} - 2\tilde{a}_{j+1}^2$ in every active slot of width 2^{-k_j} (Fig. 5). We get $|F_{\tilde{A}_1}| \leq \frac{\tilde{a}_2}{2^{k_1}} - 2\tilde{a}_2^2 + \sum_{j=2}^{m}\left[\left(\frac{\tilde{a}_{j+1}}{2^{k_j}} - 2\tilde{a}_{j+1}^2\right)\prod_{i=1}^{j-1}\left(\frac{2^{k_{i+1}}}{2^{k_i}} - 1\right)\right]$. This is maximized for $\tilde{a}_{i+1} = 1/2^{k_i+2}$, $1, \leq i \leq m$ implying $k_i = k_1 + 2(i-1)$. We get $|F_{\tilde{A}_1}| \leq \sum_{i=0}^{\infty}\frac{3^i}{2^{2k_1+4i+3}}$. $|F_{\tilde{A}_1}|/\tilde{a}_1^2$ is maximized for $\tilde{a}_1$ as small as possible; i.e., $\tilde{A}_1 = 2^{-(k_1+1)} + \varepsilon$. We get: $\frac{|F_{\tilde{A}_1}|}{\tilde{a}_1^2} \leq \sum_{i=0}^{\infty}\frac{2^{2k_1+2}.3^i}{2^{2k_1+4i+3}} = \frac{1}{2}\sum_{i=0}^{\infty}\left(\frac{3}{16}\right)^i = \frac{8}{13} = 0.6154...$ □

5 Conclusion

We have demonstrated that geometric analysis improves the best competitive guarantee for online square packing. We believe that this is not the end of the line: It should be possible to combine this type of analysis with more sophisticated, shelf-based algorithms. Our best lower bound for *BottomLeft* is a competitive factor of 5/4: Consider a sequence of small items of total width 1/3, followed by two items slightly larger than 1/3. Asymptotically, this yields a lower bound of 5/4 by taking turns with unit squares.

The bottleneck in our analysis are squares that have large holes at their right, left, and bottom side and also serve as a virtual lid; see Fig. 1(i). This worst case can happen to only a few squares, but never to all of them, so it may be possible to transfer charges between squares. It may also be possible to apply better lower bounds than the total area, e.g., similar to [12].

We also presented an algorithm that is 2.6154-competitive. We believe that our algorithm can be improved (as the best known lower bound is only 1.2). Moreover, we believe that our approach can be extended to higher dimensions. Rectangles may require a slightly different analysis. These topics will be the subject of future research. It is an open question whether our analysis is tight or can be improved. The best lower bound for *SlotAlgorithm* known to us is 2.

References

1. Baker, B.S., Coffman Jr., E.G., Rivest, R.L.: Orthogonal packings in two dimensions. SIAM Journal on Computing 9(4), 846–855 (1980)
2. Kenyon, C., Rémila, E.: Approximate strip packing. In: Proc. 37th Annu. IEEE Sympos. Found. Comput. Sci., pp. 31–36 (1996)
3. Csirik, J., Woeginger, G.J.: Shelf algorithms for on-line strip packing. Information Processing Letters 63, 171–175 (1997)
4. Seiden: On the online bin packing problem. JACM: Journal of the ACM 49 (2002)
5. van Vliet, A.: An improved lower bound for online bin packing algorithms. Information Processing Letters 43, 277–284 (1992)
6. Coppersmith, D., Raghavan, P.: Multidimensional online bin packing: Algorithms and worst case analysis. orl 8, 17–20 (1989)
7. Epstein, L., van Stee, R.: Optimal online bounded space multidimensional packing. In: Proc. 15th Symp. on Discrete Algorithms (SODA), pp. 214–223. ACM/SIAM (2004)

8. Breukelaar, R., Demaine, E.D., Hohenberger, S., Hoogeboom, H.J., Kosters, W.A., Liben-Nowell, D.: Tetris is hard, even to approximate. Internat. J. Comput. Geom. Appl. 14, 41–68 (2004)
9. Azar, Y., Epstein, L.: On two dimensional packing. J. Algor. 25, 290–310 (1997)
10. Coffman Jr., E., Downey, P.J., Winkler, P.: Packing rectangles in a strip. Acta Inform. 38, 673–693 (2002)
11. de Berg, M., van Kreveld, M., Overmars, M., Schwarzkopf, O.: Computational Geometry: Algorithms and Applications. Springer, Berlin (2000)
12. Fekete, S.P., Schepers, J.: New classes of lower bounds for the bin packing problem. Math. Progr. 91, 11–31 (2001)