

Placing your Coins on a Shelf*

Helmut Alt¹, Kevin Buchin², Steven Chaplick³, Otfried Cheong^{†4},
Philipp Kindermann⁵, Christian Knauer⁶, and Fabian Stehn⁷

- 1 Freie Universität Berlin, Germany
alt@fu-berlin.de
- 2 Technische Universiteit Eindhoven, Netherlands
k.a.buchin@tue.nl
- 3 Universität Würzburg, Germany
steven.chaplick@uni-wuerzburg.de
- 4 KAIST, Daejeon, Republic of Korea
otfried@kaist.airpost.net
- 5 FernUniversität in Hagen, Germany
philipp.kindermann@fernuni-hagen.de
- 6 Universität Bayreuth, Germany
christian.knauer@uni-bayreuth.de
- 7 Universität Bayreuth, Germany
fabian.stehn@uni-bayreuth.de

Abstract

We consider the problem of packing a family of disks “on a shelf,” that is, such that each disk touches the x -axis from above and such that no two disks overlap. We prove that the problem of minimizing the distance between the leftmost point and the rightmost point of any disk is NP-hard. On the positive side, we show how to approximate this problem within a factor of $4/3$ in $O(n \log n)$ time, and provide an $O(n \log n)$ -time exact algorithm for a special case, in particular when the ratio between the largest and smallest radius is at most four.

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases packing problems, approximation algorithms, NP-hardness

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2017.4

1 Introduction

Packing problems have a long history and abundant literature. Circular disks and spherical balls, because of their symmetry and simplicity, are of particular interest from a theoretical point of view. Historically, Johannes Kepler conjectured that an optimal packing of unit spheres into the Euclidean three-space cannot have greater density than the face-centered cubic packing [8]. The conjecture was first proven to be correct by Hales and Ferguson [7]. A more recent treatment of the proof is given by Hales et al. [6]. The proof of the 2-dimensional version of Kepler’s conjecture, that is, packing unit disks into the Euclidean two-space, is elementary and attributed to Lagrange (1773).

Packing unit disks into 2-dimensional shapes in the plane is a well studied problem in recreational mathematics. Croft et al. [2] give an overview of packing geometrical objects in finite-sized containers, for instance finding the smallest square (circle, isosceles triangle, etc.)

* A full version of this result is available on the ArXiv [1], <https://arxiv.org/abs/1707.01239>.

† OC is supported by NRF grant 2011-0030044 (SRC-GAIA) funded by the government of Korea.



© Helmut Alt, Kevin Buchin, Steven Chaplick, Otfried Cheong, Philipp Kindermann,
Christian Knauer, and Fabian Stehn;
licensed under Creative Commons License CC-BY


28th International Symposium on Algorithms and Computation (ISAAC 2017).

Editors: Yoshio Okamoto and Takeshi Tokuyama; Article No. 4; pp. 4:1–4:12



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Illustration of the span w of a valid (but not optimal) placement of five discs.

such that a given number of n unit disks can be packed into it. Specht [10] presents the best known packings of up to 10,000 disks into various containers.

Algorithmically, many packing problems are NP-hard, some are not even known to be in NP. Demaine, Fekete, and Lang showed that the problems whether a given set of circular disks of arbitrary radii can be packed into a given square, rectangle, or triangle are all NP-hard problems [3].

We will discuss a particular “nearly” one-dimensional packing problem for disks from an algorithmic perspective. We are given a family of disks that we wish to arrange “on a shelf,” that is, such that each disk touches the x -axis from above and such that no two disks overlap; see Figure 1. The goal is to minimize the *span* of the resulting configuration, that is, to minimize the horizontal distance between the leftmost point and the rightmost point of any disk. In other words, we want to minimize the required width of the shelf. Obviously, this problem is trivial for unit disks, so we allow the disks to have different sizes.

Related work. Independently from us, Dürr et al. [4] have studied the same problem, but for an isosceles, right-angled triangle. Given n sizes of this triangle, they ask for the shortest horizontal span in which the triangles can be arranged so that their lowest point lies on the x -axis, while the triangles do not overlap. Their entirely independent results are quite similar to ours: an NP-hardness proof by reduction from 3-PARTITION, a fast algorithm for a special case, and a $3/2$ -approximation algorithm.

Klemz et al. [9] show that it is NP-hard to decide if n given disks fit around a large center disk, such that each disk is in contact with the center disk while all disks are disjoint. Their proof is by reduction from 3-PARTITION as well.

Stoyan and Yaskov [11] introduce the problem of packing disks of unequal sizes into a strip of given height and minimizing the required width which is known as the *circular open dimension problem*.

Our results. We first give some useful definitions and properties for touching disks in Section 2. The hardness of the problem arises from the fact that disks can sometimes “hide” in the holes formed by larger disks, as in Figure 2b. For this reason, in Section 3, we consider the special case where, for any ordering of the disks, each disk can touch only its left and its right neighbor (where the two walls bounding the span count as neighbors as well). In particular, this implies that no disk will ever fit in a gap between two other disks. We call this the *linear case*, see Figure 2a. It turns out that for this (linear) case the optimal configuration depends *only* on the relative order of the disk sizes,¹ so it suffices to sort the disks in $O(n \log n)$ time to determine the optimal sequence.

¹ The median disk for an odd number of disks is the only exception, it can be on either end, depending on its actual size.




Figure 2 Illustration of different instances of the problem.




Figure 3 The footpoint distance of two touching disks.

In Section 4, we show that in its general form, the problem is NP-hard. More precisely, we show that given n disk sizes and a number $\delta > 0$, it is NP-hard to decide if a non-overlapping arrangement of the disks with horizontal span at most δ exists. Our NP-hardness proof is by a reduction from 3-PARTITION, and exploits the fact that disks can “hide” in the holes formed by larger disks.

Finally, in Section 5, we give an approximation algorithm that runs in $O(n \log n)$ time and guarantees a span at most $4/3$ times the optimal span.

2 Preliminaries

For reasons that will become obvious shortly, it will be convenient to define the *size* of a disk as the *square root* of its radius. We will denote disks by capital letters, and their size by the corresponding lower-case letter. Namely, disk A has size a , radius a^2 , and diameter $2a^2$.

In a valid placement, each disk A touches the x -axis in its lowest point. We will call this point the *footpoint* of the disk and denote it A . All of our arguments are based on calculations involving the distances between footpoints, so we start with the following lemma.

► **Lemma 1.** *If A and B touch, then their footpoint distance AB is $2ab$.*

Proof. The statement holds for $a = b$, so we assume $a > b$ and consider the right-angled triangle with edge lengths AB , $a^2 + b^2$, and $a^2 - b^2$, see Figure 3. We obtain $(AB)^2 = (a^2 + b^2)^2 - (a^2 - b^2)^2 = 4a^2b^2$. ◀

► **Lemma 2.** *Let G be the largest disk that fits in the gap formed by two touching disks A and B . Then $1/g = 1/a + 1/b$.*

Proof. Since G is the largest disk that fits in the gap, it must touch both A and B . By Lemma 1 we have $2ab = AB = AG + GB = 2ag + 2gb$, proving the lemma. ◀

► **Lemma 3.** *Let G be the largest disk that fits in the gap between a disk A and the vertical wall through A 's rightmost point. Then $g = (\sqrt{2} - 1) \cdot a$.*

Proof. Again, G must touch both A and the wall, so we have $a^2 = AG + g^2 = 2ag + g^2$. The positive solution to $g^2 + 2ag - a^2 = 0$ is $(\sqrt{2} - 1) \cdot a$. ◀

In any valid placement of the disks, their footpoints are distinct. Thus, the footpoints induce a linear left-to-right order on the disks. We refer to this linear order as the *footpoint sequence* of a valid placement. Further, disks are called *consecutive* or *neighbors* when their footpoints are consecutive in the footpoint sequence.

3 The Linear Case

In this section, we consider *linear case instances*, that is, instances where in any valid placement only consecutive pairs of disks can touch, only the first disk (with the leftmost footpoint) touches the left wall, and only the last disk touches the right wall.

By Lemmas 2 and 3, this is true if and only if the following condition holds: Let A be the largest disk, B the second largest, and Z the smallest disk in the collection. Then $1/z < 1/a + 1/b$, and $z > (\sqrt{2} - 1)a$. The condition holds in particular if the ratio between the largest and smallest disk size is less than two (that is, if the ratio of diameters is less than four), since then we have $1/z < 2/a \leq 1/a + 1/b$ and $z > a/2 > (\sqrt{2} - 1)a$.

In an optimal placement of a linear case instance, each disk must touch both its neighbors. Thus, the ordering of the disks uniquely determines the exact placement of every disk in any layout of minimal span. From now on, we represent placements by the *ordering* of the disks, with the understanding that the placement minimizes the span for this ordering. It remains to determine the optimal ordering. We will first give a lemma that allows us to improve a given ordering.

► **Lemma 4.** *Let \mathcal{D} be a left-to-right or right-to-left ordering of the disks in a linear case instance. Let A, B, Z be three disks that appear in this order in \mathcal{D} such that AB is a consecutive pair. Let \mathcal{D}' be the ordering obtained from \mathcal{D} by reversing the subsequence from B to Z . Then \mathcal{D}' has smaller span than \mathcal{D} if one of the following is true:*


1. Z is the last disk and $a > b > z$;
2. Z is the last disk and $a < b < z$;
3. $a > y$ and $b > z$, where Y is the disk after Z in \mathcal{D} ;
4. $a < y$ and $b < z$, where Y is the disk after Z in \mathcal{D} .

Proof. First, suppose that Z is the last disk in \mathcal{D} . Then, except for AB being replaced by AZ , each consecutive footpoint distance in \mathcal{D}' is the same as in \mathcal{D} . So, since the last disk in \mathcal{D}' is B , the change in span is $AZ + b^2 - AB - z^2 = 2az + b^2 - 2ab - z^2 = (b + z - 2a)(b - z)$. For both $a < b < z$ and $a > b > z$, this is negative, and so \mathcal{D}' has smaller span than \mathcal{D} .

Now suppose Z is not the last disk, and let Y be the disk after Z . Here, except for AB being replaced by AZ and ZY being replaced by BY , each consecutive footpoint distance in \mathcal{D}' is the same as in \mathcal{D} . Thus, the change in span is $AZ + BY - AB - ZY = 2(az + by - ab - zy) = 2(a - y)(z - b)$. For $a > y$ and $b > z$ or $a < y$ and $b < z$, this is negative. So, again \mathcal{D}' has smaller span than \mathcal{D} . ◀

We label a given family of n disks in order of decreasing size as $D_1, D_2, D_3, \dots, D_n$, and in order of increasing size as $S_1, S_2, S_3, \dots, S_n$. In other words, $d_1 \geq d_2 \geq d_3 \geq \dots \geq d_n$ and $s_1 \leq s_2 \leq s_3 \leq \dots \leq s_n$. Thus, each disk has two names, and we have $D_1 = S_n, D_2 = S_{n-1}$, and so on until $D_n = S_1$.

We now prove our claim about the structure of the optimal ordering (see also Figure 4):



■ **Figure 4** An optimal placement in the linear case. For instance for $k = 2$, the disks in $\{S_1, S_2, D_1, D_2\}$ form the consecutive subsequence starting with S_2 and ending with D_2 .

► **Lemma 5.** *Let k be an integer with $1 \leq k \leq n/2$. In any optimal placement of n disks with distinct sizes in a linear case instance, the k largest disks D_1, \dots, D_k and the k smallest disks S_1, \dots, S_k appear as a consecutive subsequence terminated by the disks S_k and D_k . If $k > 1$, then $D_k S_{k-1}$ and $S_k D_{k-1}$ are consecutive pairs.*

Proof. We use induction over k . For $k = 1$, it suffices to prove that S_1 and D_1 are consecutive, so assume for a contradiction that this is not the case. Let $A = D_1$, $Z = S_1$, assume A is to the left of Z , and let B be the right neighbor of A . By Lemma 4 (Case 1 or 3), the sequence can now be improved by reversing the subsequence from B up to Z .

Assume now that $k > 1$ and that the statement holds for $k - 1$. This means that there is a consecutive subsequence of the disks $\{S_1, \dots, S_{k-1}, D_1, \dots, D_{k-1}\}$, terminated by disk S_{k-1} at the, say, right end and disk D_{k-1} at the left end, as in the example of Figure 4.

We first show that the right neighbor of S_{k-1} is D_k . Assume this is not the case. We distinguish four cases:

1. If D_k appears to the right of S_{k-1} (but not immediately adjacent), then we apply Lemma 4 (Case 2 or 4) with $A = S_{k-1}$, B the right neighbor of S_{k-1} , and $Z = D_k$.
2. If D_k appears to the left of S_{k-1} , then it must appear to the left of D_{k-1} . If D_k is not the left neighbor of D_{k-1} , then apply Lemma 4 (Case 1 or 3) with $A = D_k$, B the right neighbor of D_k , and $Z = S_{k-1}$.
3. If D_k is the left neighbor of D_{k-1} and S_{k-1} is not the rightmost disk, then apply Lemma 4 (Case 3) with $A = D_k$, $B = D_{k-1}$, and $Z = S_{k-1}$.
4. If D_k is the left neighbor of D_{k-1} and S_{k-1} is the rightmost disk, then S_k appears somewhere to the left of D_k . We apply Lemma 4 (Case 1 or 3) with $A = D_{k-1}$, $B = D_k$, and $Z = S_k$.

We next show that the left neighbor of D_{k-1} is S_k . Assume this is not the case. If S_k appears somewhere to the left of D_{k-1} , apply Lemma 4 (Case 1 or 3) with $A = D_{k-1}$, B the left neighbor of D_{k-1} , and $Z = S_k$. If, on the other hand, S_k appears to the right of D_k , apply Lemma 4 (Case 2 or 4) with $A = S_k$, B the left neighbor of S_k , and $Z = D_{k-1}$. (Note that in this case B might be D_k .) ◀

► **Theorem 6.** *Let \mathcal{D} be a linear case instance of n disks D_1, \dots, D_n of sizes $d_1 \geq d_2 \geq \dots \geq d_n$. If n is even, then the following ordering is optimal:*

$$\dots, D_{n-5}, D_5, D_{n-3}, D_3, D_{n-1}, D_1, D_n, D_2, D_{n-2}, D_4, D_{n-4}, D_6, D_{n-6} \dots$$

For odd n , the median disk needs to be appended at the end of the sequence with the larger size difference.

Proof. Let \mathcal{D} be in the given ordering, and assume a better ordering \mathcal{D}' exists. We can modify the disk sizes slightly so as to make them unique while keeping \mathcal{D}' better than \mathcal{D} . But then we have a contradiction to Lemma 5. If n is odd, then the only possible placements of the median disk are the left end and the right end, so choosing the end with the larger size difference gives the optimal solution. ◀

4 NP-Hardness of the General Case

Let us denote the decision version of our problem as COINSONASHELF. Its input is a set of disks with rational radii and a rational number $\delta > 0$, the question is whether there is a feasible placement of the disks with span at most δ .

► **Theorem 7.** *COINSONASHELF is NP-hard, even when the ratio of the largest and smallest disk size is bounded by six and when all numbers are given in unary notation.*

Our proof is by reduction from 3-PARTITION [5, Problem SP15]. An instance of 3-PARTITION consists of $3m$ integers $\mathcal{A} = a_1, \dots, a_{3m}$ and another integer B , with $\sum_{i=1}^{3m} a_i = mB$ and $B/4 < a_i < B/2$ for all i . 3-PARTITION decides if there is a partition of \mathcal{A} into m three-element groups A_1, \dots, A_m such that $\sum_{a \in A_i} a = B$ for each group A_i .

Given a 3-PARTITION instance (\mathcal{A}, B) , we construct a family \mathcal{D} of $12m + 11$ disks, as follows:

- $m + 1$ disks of size 1, we will refer to these disks as *outer frame disks*;
- $4(m + 1)$ disks of size $s_0 = 33/100 = 0.33$, we will refer to these disks as *inner frame disks*;
- $2(m + 1)$ disks of size $s_1 = s_0/1+s_0 = 33/133 (\approx 0.24812)$, we will refer to these disks as *large filler disks*;
- $2(m + 1)$ disks of size $s_2 = s_1/1+s_1 = 33/166 (\approx 0.198795)$, we will refer to these disks as *small filler disks*;
- 2 disks of size $s_3 = \frac{1-s_0^2-2s_0}{4s_0} = 2311/13200 (\approx 0.175076)$, referred to as *end disks*;
- $3m$ disks D_1, \dots, D_{3m} , referred to as *partition disks*, where $d_i = \frac{17}{99} \left(\frac{3}{100} \frac{a_i}{B} + \frac{99}{100} \right)$.

In the following, we will identify disks by their size or type. We observe that all disk sizes are rational, where numerator and denominator can be computed in time polynomial in the input size. The radius of a disk is obtained by squaring its size. Note that, if we multiply all radii by the product of the denominators, then we obtain in polynomial time an instance of our problem with integer radii.

► **Lemma 8.** *Each end disk and partition disk has size at least $s_4 = 2261/13200 > 0.17128$.*

Proof. Since $s_3 > s_4$, the statement is trivial for end disks. Let $a_i \in \mathcal{A}$. From $a_i > B/4$ follows that the size d_i of the corresponding partition disk is $d_i \geq \frac{17}{99} (3/400 + 99/100) = \frac{17}{99} \cdot \frac{399}{400} = 2261/13200$. ◀

Equivalence of the problem instances. We show that \mathcal{D} has a placement with span $2(m+1)$ if and only if (\mathcal{A}, B) is a Yes-instance of 3-PARTITION, implying the NP-hardness of COINSONASHELF.

The $m + 1$ outer frame disks alone already require a span of $2(m + 1)$, so no better span is possible. A placement of all disks of \mathcal{D} with span $2(m + 1)$ therefore implies that consecutive outer frame disks touch, and that all remaining disks fit into the space under these outer frame disks.

Let's call the m spaces between two consecutive (and touching) outer frame disks *gaps*. The space to the left of the leftmost outer frame disk is called the *left end*, the *right end* is defined symmetrically.

► **Lemma 9.** *There is only one pattern of frame and filler disks (ignoring end disks and partition disks) that has span $2(m + 1)$.*




Figure 5 The unique pattern of span $2(m + 1)$ in Lemma 9.

The proof can be found in the full paper [1], here we only show the pattern in Figure 5a. Each gap contains eight disks of sizes $s_2, s_1, s_0, s_0, s_0, s_0, s_1, s_2$; see Figure 5b. The left end contains four disks of sizes s_0, s_0, s_1, s_2 , the right end contains disks of sizes s_2, s_1, s_0, s_0 .

► **Lemma 10.** *Three end/partition disks $X, Y,$ and Z fit in the three gaps formed by the three pairs of consecutive inner frame disks in a common gap if and only if $x + y + z \leq 17/33$.*

Proof. By Lemma 2, the largest disk that fits in the space between two touching disks of size s_0 has size $s_0/2$. By Lemma 8, an end/partition disk has size at least $s_4 > s_0/2$, so it does not fit entirely in this space. It follows that the total footprint distance of the sequence $1, s_0, x, s_0, y, s_0, z, s_0, 1$ is at least $4s_0 + 4s_0x + 4s_0y + 4s_0z = 4s_0(x + y + z + 1)$. $X, Y,$ and Z fit in the prescribed manner if and only if this total footprint distance is at most two, proving the lemma. ◀

► **Lemma 11.** *Placing a disk X in the space between the two consecutive inner frame disks in the left end or the right end causes the total span to increase if and only if $x > s_3$.*

Proof. If $x \leq s_0/2 < s_3$, the statement follows from Lemma 2, so assume $x > s_0/2$. Then the total width of the sequence $1, s_0, x, s_0$ is $2s_0 + 4s_0x + s_0^2$. The span increases if and only if this is larger than one, proving the lemma. ◀

A 3-partition implies small span. Assume that \mathcal{A} can be partitioned into m groups A_i such that $\sum_{a \in A_i} a = B$. Consider a group $A_i = (a_{i1}, a_{i2}, a_{i3})$ and let $X, Y,$ and Z be the partition disks corresponding to a_{i1}, a_{i2}, a_{i3} . Then we have

$$x + y + z = \frac{17}{99} \left(\frac{3}{100} \frac{a_{i1} + a_{i2} + a_{i3}}{B} + 3 \cdot \frac{99}{100} \right) = \frac{17}{33}.$$

By Lemma 10 this implies that $X, Y,$ and Z can be placed in a common gap in the pattern of Figure 5 without increasing the total span. Since there are m gaps, we can place all partition disks into the m gaps. Finally, by Lemma 11, we can place the two end disks inside the left end and the right end.

Small span implies a 3-partition. We assume now that a placement of the disks \mathcal{D} with span $2(m + 1)$ exists. By Lemma 9, the frame and filler disks must be placed in the pattern of Figure 5. It remains to discuss the possible locations of the end disks and the partition disks. We need a number of observations about a placement of span $2(m + 1)$:

1. The left end and right end can contain at most one end disk or partition disk, and only between the two inner frame disks or between the outer frame disk and the small filler disk, see top of Table 1.
2. A gap can contain at most three partition disks or end disks. If a gap contains three such disks, each has to appear between two inner frame disks, see bottom of Table 1.

■ **Table 1** Impossible placements of end/partition disks...

... in the right end	
sequence	width
1 $s_0 s_0 s_4$	$2s_0 + 2s_0^2 + 2s_0s_4 + s_4^2 > 1.0201$
1 $s_1 s_4 s_0 s_0$	$2s_1 + 2s_1s_4 + 2s_0s_4 + 3s_0^2 > 1.0209$
1 $s_2 s_4 s_1 s_0 s_0$	$2s_2 + 2s_2s_4 + 2s_1s_4 + 2s_1s_0 + 3s_0^2 > 1.0411$
1 $s_0 s_4 s_4 s_0$	$2s_0 + 4s_0s_4 + 2s_4^2 + s_0^2 > 1.0536$
1 $s_4 s_4 s_2 s_1 s_0 s_0$	$2s_4 + 2s_4^2 + 2s_2s_4 + 2s_1s_2 + 2s_1s_0 + 3s_0^2 > 1.0584$
1 $s_4 s_2 s_1 s_0 s_4 s_0$	$2s_4 + 2s_2s_4 + 2s_1s_2 + 2s_1s_0 + 4s_0s_4 + s_0^2 > 1.0080$

... in a gap	
sequence	total footprint distance
1 $s_1 s_4 s_0 s_0 s_0 s_0 1$	$2s_1 + 2s_1s_4 + 2s_0s_4 + 6s_0^2 + 2s_0 > 2.0076$
1 $s_2 s_4 s_1 s_0 s_0 s_0 s_0 1$	$2s_2 + 2s_2s_4 + 2s_4s_1 + 2s_1s_0 + 6s_0^2 + 2s_0 > 2.0278$
1 $s_0 s_4 s_4 s_0 s_0 s_0 1$	$4s_0 + 4s_0s_4 + 2s_4^2 + 4s_0^2 > 2.0403$
1 $s_4 s_4 s_2 s_1 s_0 s_0 s_0 s_0 1$	$2s_4 + 2s_4^2 + 2s_4s_2 + 2s_2s_1 + 2s_1s_0 + 6s_0^2 + 2s_0 > 2.0451$
1 $s_4 s_2 s_1 s_0 s_4 s_0 s_4 s_0 s_0 1$	$2s_4 + 2s_4s_2 + 2s_2s_1 + 2s_1s_0 + 8s_0s_4 + 2s_0^2 + 2s_0 > 2.0030$
1 $s_4 s_2 s_1 s_0 s_4 s_0 s_0 s_0 s_1 s_2 s_4 1$	$4s_4 + 4s_4s_2 + 4s_2s_1 + 4s_1s_0 + 4s_0s_4 + 4s_0^2 > 2.0078$

- Since there are $3m + 2$ end and partition disks, (1) and (2) imply that each gap contains three such disks, while the left end and right end each contain one.
- By (1) and Lemma 11, the left end and the right end can contain only disks of size at most s_3 . We can assume that these are the two end disks (otherwise, swap them with an end disk).
- Consider a gap. It contains exactly three partition disks X, Y , and Z . By Lemma 10, we have $x + y + z \leq 17/33$. Let a, b, c be the elements of \mathcal{A} corresponding to X, Y , and Z . Then we have

$$x + y + z = \frac{17}{99} \left(\frac{3}{100} \frac{a + b + c}{B} + 3 \cdot \frac{99}{100} \right) \leq \frac{17}{33},$$

which implies $a + b + c \leq B$. It follows that we have partitioned the elements of \mathcal{A} into m groups A_1, A_2, \dots, A_m with $\sum_{a \in A_i} a \leq B$. Since $\sum_{a \in \mathcal{A}} a = mB$, we must have $\sum_{a \in A_i} a = B$ for each i , so (\mathcal{A}, B) is a Yes-instance of 3-PARTITION.


This concludes the proof of Theorem 7, noting that by Lemma 8 all disks have size at least $s_4 > 1/6$.

5 A $4/3$ -Approximation

In this section, we give a *greedy algorithm* and prove that it computes a $4/3$ -approximation to the problem.

Our algorithm starts by sorting the disks D_1, D_2, \dots, D_n by decreasing size, such that $d_1 \geq d_2 \geq \dots \geq d_n$. It then considers the disks one by one, in this order, maintaining a placement of the disks considered so far. Each disk D is placed as follows:

- If there is a gap between two consecutive disks A and B in the current placement that is large enough to contain D , then we place D in this gap, touching the *smaller* one of the two disks A and B .
- Otherwise, let A be the leftmost disk in the current placement (that is, the disk with the leftmost footprint—this is not necessarily the disk defining the left end of the current



■ **Figure 6** Support of three disks of radius 1, 2 and 3 respectively.

span), and let Z be the rightmost disk. Since $d \leq a$, we can place D so that it touches A from the left (candidate placement D_A), and since $d \leq z$, we can place D so that it touches Z from the right (candidate placement D_Z).

3. If one of the candidate placements D_A or D_Z does not increase the span, we place D in this way.
4. Otherwise, we place D at D_A if $a > z$ and at D_Z otherwise.

The algorithm can be implemented to run in time $O(n \log n)$ as follows: We maintain a priority queue that stores, for each pair of consecutive disks, the size of the largest disk that will fit between them. Since we are placing disks in order of decreasing size, a newly placed disk can only touch its two neighbors, and so it will fit into the gap if and only if its size is at most the stored gap size.

For the analysis of the approximation factor, we will assume, without loss of generality, that the final disk D_n is placed using the last rule (as otherwise it does not contribute to the final span and can be ignored in the analysis). We also assume that $d_n = 1$.

Next, let's call a disk D *large* if $d \geq 2$, and *small* otherwise. We have the following:

► **Lemma 12.** *Any two consecutive small disks placed by the algorithm touch.*

Proof. Assume, for a contradiction, that D is the *first* small disk whose placement causes two small disks to be consecutive but non-touching.

If D was placed by the third or fourth rule (at the left or right end of the sequence), it is touching its only neighbor. Therefore, D must have been placed in a gap between two disks A and B . If both A and B are small, they must be touching (since D is the first small disk that will not touch a neighboring small disk). But by Lemma 2 that implies that the gap between A and B is too small to contain a disk of size $d \geq 1$. It follows that at most one of A and B is small, say B . But then the algorithm will place D such that it touches B , a contradiction. ◀

We now associate with each disk a *support interval*. The support interval of a disk A is the interval $[A - 2a + 1, A + 2a - 1]$. Since $0 \leq (a - 1)^2 = a^2 - 2a + 1$, we have $2a - 1 \leq a^2$, and so the support interval of a disk lies within the disk's span, see Figure 6.

► **Lemma 13.** *In any feasible placement of disks of size at least one, the open support intervals of the disks are disjoint.*

Proof. Consider the function $f(a, b) = (a + b - 1)/ab$ for $a, b \geq 1$. Since $f(1, \cdot) = f(\cdot, 1) = 1$ and the partial derivatives of f are negative for $a, b > 1$, we have $f(a, b) \leq 1$.

Consider two consecutive touching disks of size a and b . Their footpoints are at distance $2ab$. The support intervals cover $2a + 2b - 2$ of this distance. From $f(a, b) \leq 1$ it follows that $2a + 2b - 2 \leq 2ab$, and so the support intervals do not overlap. ◀

4:10 Placing your Coins on a Shelf

Lemma 13 implies that the total length of the support intervals is a lower bound for the span of a family of disks. We will show that our greedy algorithm computes a solution where the support intervals cover at least $3/4$ of the span, implying approximation factor $4/3$.

Consider a pair of two consecutive disks A and B placed by the algorithm, and let G be the (imaginary) largest disk that can be placed in the gap between A and B . Since D_n was not placed in this gap, we have $g < 1$. By Lemma 1, we have $\overline{AB} = \overline{AG} + \overline{GB} = 2ag + 2gb = 2g(a + b)$.

Consider first the case where A and B touch. Lemma 2 gives $1/g = 1/a + 1/b$ or $g = ab/(a + b)$. The support intervals cover $2a + 2b - 2$ of the footpoint distance $2ab$, so the ratio is $1/a + 1/b - 1/ab$. For $1 \leq a, b$ under the constraint $1/a + 1/b > 1$ this is minimized at $a = b = 2$ and we have $1/a + 1/b - 1/ab \geq 3/4$, so the claim holds for this interval.

Now suppose that A and B do not touch. By Lemma 12, this means at least one of the disks is large, say A , that is $a \geq 2$. The footpoint distance \overline{AB} is $2g(a + b) \leq 2(a + b)$, and the support intervals cover $2a + 2b - 2$ of this distance, so the ratio is

$$\frac{2a + 2b - 2}{2g(a + b)} \geq \frac{a + b - 1}{a + b} = 1 - \frac{1}{a + b}.$$

If $a \geq 3$ or $b \geq 2$, we already have $1 - 1/(a + b) \geq 3/4$, and this bound is good enough.

It remains to consider the situation when $2 \leq a < 3$ and $1 \leq b \leq 2$. Without loss of generality, we assume that B is to the right of A . We denote the first disk to the right of A that is touching A as D . By the nature of our algorithm, when B was placed, it was placed inside the space between A and D (possibly, other disks were already present in this space at that time). Since B does not touch A , the disk D must be smaller than A , that is $1 \leq d \leq a < 3$.

We analyze the entire interval $[A, D]$ as a whole. Since A and D touch, the length of this interval is $2ad$. In between A and D , some $k \geq 1$ disks have been placed, with B being the leftmost of these.

We first consider the case $k \geq 2$. If two disks X and Y of size one fit between A and D , then we have

$$2ad = \overline{AD} = \overline{AX} + \overline{XY} + \overline{YD} \geq 2a + 2 + 2d,$$

and from $a < 3$ follows

$$d \geq \frac{a + 1}{a - 1} = 1 + \frac{2}{a - 1} > 2.$$

The total length of the support intervals in the interval \overline{AD} is at least $2a - 1 + 2d - 1 + 2k \geq 2(a + d + 1)$. The distance \overline{AD} is $2ad$. For $2 \leq a, d \leq 3$, the ratio $(a + d + 1)/ad$ is at least $7/9 > 3/4$, implying the claim.

In the second case, B is the only disk between A and D . This means that B touches D . The total support interval length in the interval \overline{AD} is

$$2a - 1 + 4b - 2 + 2d - 1 = 2a + 4b + 2d - 4.$$

Let G be the largest disk that fits in the gap between A and B . Its size is determined by the equality $2ag + 2gb + 2bd = 2ad$, so $g = (a - b)d/(a + b)$. Since D_n was not placed in this gap, we have $g < 1$, and so $(a - b)d < a + b$. Minimizing the expression

$$\frac{a + 2b + d - 2}{ad}$$

under the constraints $2 \leq a \leq 3$, $1 \leq d \leq a$, $1 \leq b \leq 2$, and $(a - b)d < a + b$ leads to the minimum $7/9 > 3/4$ for $a = d = 3$ and $b = 3/2$.

To complete the proof, we need to argue about the part of the span that does not lie between two footpoints, in other words, the two intervals between the left wall (defined by the leftmost point on any disk) and the leftmost footpoint, and between the rightmost footpoint and the right wall. Recall that we assumed that placing D_n increased the total span. This implies that D_n was placed using the algorithm's last rule and therefore touches one of the two walls, let's say the right wall. Let A and B be the leftmost two disks (in footpoint order), and let Y and Z be the rightmost two disks (in footpoint order). By assumption, $Z = D_n$ and so $z = 1$. Since D_n was placed using the last rule, we have $y \geq a$, and Z touches Y . Let us call G the (imaginary) largest disk that would fit into the space between the left wall and A . Since D_n was not placed in this position, we have $g < 1$. Note that the left wall is at coordinate $G - g^2$, the right wall at coordinate $Z + 1$. We now distinguish two cases.

We first consider the case where $a \geq 3/2$. We then analyze the two intervals $[G - g^2, A]$ and $[Y, Z + 1]$ together. Their total length is $g^2 + 2ga + 2y + 1 < 2y + 2a + 2$, and the support intervals of A , Y , and Z cover $2a - 1 + 2y - 1 + 2 = 2y + 2a$ of this. The ratio is

$$\frac{2y + 2a}{2y + 2a + 2} = 1 - \frac{1}{y + a + 1} \geq 1 - \frac{1}{4} = \frac{3}{4} \quad \text{since } y \geq a \geq 3/2.$$

In the second case we have $a < 3/2$. Then B must be touching A . This is true if $b \geq a$, because then A was placed later than B using the third rule. When $b < a$, then it follows from Lemma 12. The distance between $G - g^2$ and B is then $g^2 + 2ag + 2ab \leq 2ab + 2a + 1 \leq 3b + 4$. Since B fits inside the span, we must have $b^2 \leq 3b + 4$, which solves to $-1 \leq b \leq 4$.

We now analyze the intervals $[G - g^2, B]$ and $[Y, Z + 1]$ together. Their total length is

$$g^2 + 2ga + 2ab + 2y + 1 < 2y + 2a + 2ab + 2,$$

while the support intervals of A , B , Y , and Z cover

$$4a - 2 + 2b - 1 + 2y - 1 + 2 = 2y + 4a + 2b - 2.$$

Since $y \geq a$, we can lower-bound the ratio

$$\frac{2y + 4a + 2b - 2}{2y + 2a + 2ab + 2} \geq \frac{6a + 2b - 2}{4a + 2ab + 2} = \frac{3a + b - 1}{2a + ab + 1}.$$

Consider the function $h(a, b) = 3a + 2b - 3ab/2$ over the domain $1 \leq a \leq 3/2$ and $1 \leq b \leq 4$. For fixed b , the function $h(a, b)$ is linear in a , so $h(a, b) \geq \min \{h(1, b), h(3/2, b)\}$. We have $h(1, b) = 3 + 2b - 3b/2 = 3 + b/2 \geq 7/2$ and $h(3/2, b) = 9/2 + 2b - 9b/4 = 9/2 - b/4 \geq 7/2$.

It follows that $\frac{3}{2}a + b - \frac{3}{4}ab \geq \frac{7}{4}$, and so

$$3a + b - 1 \geq \frac{3}{2}a + \frac{3}{4}ab + \frac{3}{4} = \frac{3}{4}(2a + ab + 1).$$

Note that in this second case we have used the interval $[A, B]$ to help bound the coverage of the two end intervals. This could be a problem if the same interval was also needed to help bound a larger interval of the form $[A, C]$, where A and C touch and B was inserted into this interval later. But note that we needed to analyze $[A, C]$ as a whole only if $c < 3$. Since $a < 3/2$, no disk of size one would then fit into the gap between A and C , so this situation cannot occur.

This completes the proof of the following theorem.

► **Theorem 14.** *The greedy algorithm computes a $4/3$ -approximation in time $O(n \log n)$.*

6 Conclusions

Our best approximation algorithm achieves an approximation factor of $4/3$. We were unable to find a polynomial time approximation scheme, so it would be natural to try to prove that the problem is APX-hard. This, however, seems unlikely to be true, for the same reasons as outlined by Dürr et al. [4]: The ideas they present appear to transfer to our problem, and would lead to an $2^{O(\log^{O(1)} n)}$ algorithm with approximation factor $(1 + \varepsilon)$. APX-hardness, on the other hand, would imply that for some $\varepsilon > 0$ this approximation problem is NP-hard, implying subexponential algorithms for NP.

Acknowledgments. We thank Peyman Afshani and Ingo van Duin for helpful discussions during O.C.'s visit to Madalgo in 2016. We also thank all participants at the Korean Workshop on Computational Geometry in Würzburg 2016.

References

- 1 Helmut Alt, Kevin Buchin, Steven Chaplick, Otfried Cheong, Philipp Kindermann, Christian Knauer, and Fabian Stehn. Placing your coins on a shelf, 2017. [arXiv:1707.01239](https://arxiv.org/abs/1707.01239).
- 2 Hallard T. Croft, Kenneth J. Falconer, and Richard K. Guy. *Unsolved Problems in Geometry*. Springer-Verlag, 1991. pp. 108–110.
- 3 Erik D. Demaine, Sándor P. Fekete, and Robert J. Lang. Circle packing for origami design is hard. In A. K. Peters, editor, *Origami⁵: Proceedings of the 5th International Conference on Origami in Science, Mathematics and Education (OSME 2010)*, pages 609–626, 2010.
- 4 Christoph Dürr, Zdeněk Hanzálek, Christian Konrad, Yasmina Seddik, René Sitters, Óscar C. Vásquez, and Gerhard Woeginger. The triangle scheduling problem. *Journal of Scheduling*, pages 1–8, 2017. doi:10.1007/s10951-017-0533-1.
- 5 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.
- 6 Thomas C. Hales, Mark Adams, Gertrud Bauer, Dat Tat Dang, John Harrison, Truong Le Hoang, Cezary Kaliszzyk, Victor Magron, Sean McLaughlin, Thang Tat Nguyen, Truong Quang Nguyen, Tobias Nipkow, Steven Obua, Joseph Pleso, Jason Rute, Alexey Solov'yev, An Hoai Thi Ta, Trung Nam Tran, Diep Thi Trieu, Josef Urban, Ky Khac Vu, and Roland Zumkeller. A formal proof of the Kepler conjecture. *Forum of Mathematics, Pi*, 5(e2):1–29, 2017. doi:10.1017/fmp.2017.1.
- 7 Thomas C. Hales and Samuel P. Ferguson. A formulation of the Kepler conjecture. *Discrete & Computational Geometry*, 36(1):21–69, 2006.
- 8 Johannes Kepler. *Strena seu de nive sexangula (The six-cornered snowflake)*. 1611.
- 9 Boris Klenz, Martin Nöllenburg, and Roman Prutkin. Recognizing weighted disk contact graphs. In *Graph Drawing and Network Visualization*, volume 9411 of *LNCS*, pages 433–446. Springer, 2015.
- 10 Eckehard Specht. www.packomania.com. Accessed 2017-06-28. URL: <http://www.packomania.com/>.
- 11 Yu. G. Stoyan and Georgiy Yaskov. A mathematical model and a solution method for the problem of placing various-sized circles into a strip. *European Journal of Operational Research*, 156(3):590–600, 2004.