MATHEMATICS OF OPERATIONS RESEARCH
Vol. 31, No. 1, February 2006, pp. 31–49
ISSN 0364-765X | EISSN 1526-5471 | 06 | 3101 | 0031

DOI 10.1287/moor.1050.0168
© 2006 INFORMS

Bin Packing in Multiple Dimensions: Inapproximability Results and Approximation Schemes

Nikhil Bansal
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598,
nikhil@us.ibm.com, http://www.research.ibm.com/people/n/nikhil/

José R. Correa
School of Business, Univesidad Adolfo Ibáñez, Avenida Presidente Errázuriz 3485, Las Condes, Santiago, Chile,
correa@uai.cl, http://jcorrea.uai.cl/

Claire Kenyon
Department of Computer Science, Brown University, Providence, Rhode Island 02912,
claire@cs.brown.edu, http://www.cs.brown.edu/~claire/

Maxim Sviridenko
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598,
sviri@us.ibm.com, http://www.research.ibm.com/people/s/sviri/sviridenko.html

We study the following packing problem: Given a collection of d-dimensional rectangles of specified sizes, pack them into the minimum number of unit cubes. We show that unlike the one-dimensional case, the two-dimensional packing problem cannot have an asymptotic polynomial time approximation scheme (APTAS), unless $P = NP$. On the positive side, we give an APTAS for the special case of packing d-dimensional cubes into the minimum number of unit cubes. Second, we give a polynomial time algorithm for packing arbitrary two-dimensional rectangles into at most OPT square bins with sides of length $1+\varepsilon$, where OPT denotes the minimum number of unit bins required to pack these rectangles. Interestingly, this result has no additive constant term, i.e., is not an asymptotic result. As a corollary, we obtain the first approximation scheme for the problem of placing a collection of rectangles in a minimum-area encasing rectangle.

Key words: bin packing; approximation algorithms; polynomial time approximation schemes; hardness of approximation
MSC2000 subject classification: Primary: 90B35, 68W25; secondary: 60K30
OR/MS subject classification: Primary: Production/scheduling; secondary: cutting stock/trim, approximations/heuristic
History: Received February 21, 2004; revised December 7, 2004, and May 17, 2005.

1. Introduction. In the *two-dimensional bin packing problem*, rectangles of specified size (width, height) have to be packed into larger squares (bins). The most interesting and well-studied version of this problem is the so-called *orthogonal packing without rotation*, where each rectangle must be packed parallel to the edges of a bin. The goal is to find a feasible packing, i.e., a packing where rectangles do not overlap, using the smallest number of bins.

Bin packing and its d-dimensional variants have been extensively studied since the 1960s both in the context of offline approximation algorithms and online algorithms (see, e.g., Baker et al. [1], Caprara [3], Caprara et al. [4], Chung et al. [6], Coffman et al. [8], Epstein and van Stee [11], Ferreira et al. [13], Kenyon and Rémila [18], Kohayakawa et al. [20], Leung et al. [23], Seiden and van Stee [29]). Detailed surveys can be found in Coffman et al. [7] and Csirik and Woeginger [10]. Throughout this paper we only consider offline algorithms.

Clearly the NP-hardness of two-dimensional bin packing follows from that of one-dimensional bin packing (which is a special case, when all the heights are exactly 1). Furthermore, it is known that given a collection of rectangles (in fact, squares), it is NP-hard to decide in polynomial time whether these can be packed in a single bin or require two bins (Leung et al. [23]). Hence, no $(2-\varepsilon)$ approximation algorithm for the problem can exist unless $P = NP$. These kind of hardness of approximation results are typical in bin packing problems; nevertheless, one can usually prove much stronger performance guarantees when considering asymptotic analysis. Therefore, the standard measure used to analyze the performance of a packing Algorithm A is the asymptotic approximation ratio R_A^∞ defined as

$$R_A^n = \max_L \left\{ \frac{A(L)}{\mathrm{OPT}(L)} \colon \mathrm{OPT}(L) = n \right\},$$
$$R_A^\infty = \lim_{n\to\infty} \sup R_A^n,$$

where L ranges over the set of all problem instances and $A(L)$ (resp. $\mathrm{OPT}(L)$) denotes the number of bins used by A (resp. the optimal solution). Thus, the notion of asymptotic approximation ratio allows us to ignore the anomalous behavior of the algorithm for small instances. A problem admits an asymptotic approximation scheme, APTAS for short, if for any $\varepsilon > 0$, there is a polynomial time algorithm with an asymptotic approximation ratio of $1+\varepsilon$.

For one-dimensional bin packing, Fernandez de La Vega and Lueker [12] gave the first asymptotic approximation scheme. Later, this was improved by Karmarkar and Karp [17] to give an algorithm that uses $\mathrm{OPT} + O(\log^2(\mathrm{OPT}))$ bins. For the two-dimensional case, the first results were obtained by Chung et al. [6] who gave a 2.125-approximation algorithm. For a long time this was the best known until a $2+\varepsilon$ (for any $\varepsilon > 0$) approximation was obtained (implicitly) by Kenyon and Rémila [18]. The recent breakthrough is an elegant 1.691-approximation algorithm due to Caprara [3]. Interestingly, many more results are known for some special cases in which there is a restriction on how the rectangles can be packed in a bin. Two particular cases that are widely studied are *strip packing* and *shelf packing* (details about these can be found in Caprara [3] and the references therein). While clever asymptotic approximation schemes are known for some of these special cases, it was unclear whether the general two-dimensional bin packing problem admits an approximation scheme.

For the special case of packing squares in square bins (which is also NP-hard by Leung et al. [23]), only algorithms with constant factor approximation ratios were known prior to our work. The first guarantee better than 2.125 was obtained by Ferreira et al. [13], who gave a 1.988-approximation algorithm. This was later improved to $14/9 + \varepsilon = 1.55 + \varepsilon$ by Kohayakawa et al. [20] and by Seiden and van Stee [29]. Recently, Caprara [3] gave an algorithm and showed that it has approximation ratio in the interval $[1.490, 1.507]$, provided a conjecture is true. The most recent (and best) known guarantee prior to our work is due to Epstein and van Stee [11], who give a $16/11 + \varepsilon = 1.454 + \varepsilon$ algorithm for the two-dimensional case. The algorithm in Kohayakawa et al. [20] also works in the d-dimensional case, and its asymptotic approximation ratio of $2-(2/3)^d$ was the best-known approximation ratio for $d > 2$ dimensions prior to our work.

Another relevant problem in the context of two-dimensional packing is the *minimum rectangle placement problem*: Find the minimum area rectangle in which a given set of rectangles can be placed. This problem has many applications in scheduling and maybe the most prominent one, in very large-scale integration (VLSI) design. Despite the significant number of approximate and exact heuristic methods that have been proposed for this problem (see, e.g., Korf [21], Murata et al. [26], and references therein), very few results that prove worst-case approximation guarantees are known. Exceptions are the work of Kleitman and Krieger [19] and the work by Novotny [27] that consider the special case of packing squares into a minimum area rectangle.

A related multidimensional packing problem is *vector bin packing*, described as follows: Given a set of n rational, d-dimensional vectors $p_1, \ldots, p_n$ from $[0,1]^d$, find a partition of the vectors into sets $A_1, \ldots, A_m$ such that $\|\bar{A}_i\|_\infty \le 1$ for $1 \le i \le m$, where $\bar{A}_i = \sum_{j\in A_i} p_j$ is the sum of the vectors in A_i. The objective is to minimize m, the size of the partition. For $d = 1$, the vector bin packing problem is identical to the classical one-dimensional bin packing, but this is not true for $d > 1$. Chekuri and Khanna [5] showed an interesting connection between d-dimensional vector bin packing (for arbitrary d) and graph coloring, which implies that vector bin packing is hard to approximate within $O(d^{1/2-\varepsilon})$ for any $\varepsilon > 0$. Woeginger [33] showed that the problem is APX hard even for the case of $d = 2$. The best-known result for this problem is a $(1+\varepsilon d + O(\ln \varepsilon^{-1}))$-approximation for any fixed $\varepsilon > 0$, by Chekuri and Khanna [5].

Another closely related problem is the packing problem where we are given a set of rectangles with nonnegative profits. The goal is to maximize the total profit of rectangles that can be packed into a larger rectangle (or square by scaling). The most recent paper on this problem is due to Jansen and Zhang [15]. They describe a few constant factor approximation algorithms for this problem and provide references on the state of the art for it.

1.1. Our results.

We now state the main result in each section of the paper.

1. In §2 we prove that the two-dimensional bin packing problem does not admit an asymptotic approximation scheme. This directly implies the nonexistence of an asymptotic PTAS for all $d \ge 2$.

2. We give the first asymptotic approximation scheme for packing squares into square bins in §3. More generally, we give an APTAS for packing d-dimensional hypercubes into bins for any fixed d. Specifically, we prove the following:

THEOREM 1.1. *There exists an Algorithm A that, given a list I of n d-dimensional cubes and a positive number ε, produces a packing of I into $A(I)$ copies of $[0,1]^d$ such that:*

$$A(I) \le \lceil (1+\varepsilon)\,\mathrm{OPT}(I) \rceil + 1.$$

The running time of A is $O(n\log(n))$ for fixed d and ε.

We will also show that a slight variation of the algorithm achieves an absolute approximation guarantee of 2, matching a recent result by van Stee [31]. Note that this algorithm achieves the best possible nonasymptotic ratio. This follows directly due to the impossibility of distinguishing in polynomial time whether a collection of squares can be packed in a single bin or requires two bins Ferreira et al. [13].

3. In §4 we design a resource-augmented approximation scheme for the general two-dimensional rectangle packing problem, i.e., an algorithm which, given $\varepsilon > 0$, packs rectangles into bins of size $(1+\varepsilon)\times(1+\varepsilon)$. Our result is found in Theorem 1.2.

Theorem 1.2. *There exists an Algorithm A that, given a list I of n rectangles and a positive number ε, produces a packing of I into $A(I)$ copies of $[0, 1+\varepsilon]^2$ such that:*

$$A(I) \le \mathrm{OPT}(I),$$

where $\mathrm{OPT}(I)$ is the minimum number of unit cubes in which I can be packed. The running time of A is polynomial in n for fixed ε.

Unfortunately, it turns out that the running time of our algorithm is $O(n^{2^{2^{\tilde{O}(1/\varepsilon)}}})$, which is rather impractical.

4. Observe that the result in Theorem 1.2 above is not asymptotic (there is no additive term in the expression for the number of bins used by the algorithm). We use this observation to give a polynomial time approximation scheme (PTAS) for the minimum rectangle placement problem. This is described in §5.

Theorem 1.3. *There exists an Algorithm A that, given a list I of n rectangles and a positive number ε, finds a rectangle R, in which all rectangles in I fit and such that*

$$\mathrm{Vol}(R) \le (1+\varepsilon)\,\mathrm{OPT}(I).$$

Here $\mathrm{Vol}(R)$ is the area of rectangle R and $\mathrm{OPT}(I)$ is the area of the minimum area rectangle in which I can be packed. The running time of A is polynomial in n for fixed ε.

1.2. Techniques. Our result for the APX hardness of two-dimensional bin packing has ideas similar to those used by Woeginger [33] to show the APX hardness of two-dimensional vector packing. However, our construction is much more involved in Woeginger [33], as there is much less structure in how the rectangles can be packed in a bin for two-dimensional bin packing, as compared with that for vector packing. For example, given a collection of rectangles (or equivalently vectors) $p_1, \ldots, p_n$, it is NP-hard to decide whether these rectangles can be packed in a single bin, whereas this is trivial for vector bin packing (to do this, simply sum up the vectors and verify that all coordinates are at most 1).

For hypercube packing, most previous approaches that obtain a constant factor approximation Ferreira et al. [13], Seiden and van Stee [29], Kohayakawa et al. [20], and Epstein and van Stee [11] use the classical techniques used for one-dimensional bin packing. That is, classify the objects into large and small. Find a packing of the large objects using rounding and exhaustive search and then pack the remaining small objects. In the case when $d > 1$, the above approach does not directly yield an approximation scheme for the following reason: The gaps left in the bin after packing the large objects can have arbitrary structure, and it is not clear how to pack the small objects in these gaps without wasting a constant fraction of the space. Our approach extends the ideas of Fernandez de la Vega and Lueker [12] and is based on a technique originally used by Sevastianov and Woeginger [30] in the context of some problems in shop scheduling. We partition the objects into three sets, large, medium, and small, such that the medium objects do make up a significant portion of the input instance, and at the same time this gives us a sufficient gap between the sizes of the large objects and the small objects. We pack the medium objects separately and then show how to pack the large and the small objects together.

We now consider the case of packing rectangles using resource augmentation. The main difficulty of obtaining a PTAS in this case is that there is no natural total order on rectangles. Hence, the rounding techniques used by Fernandez de la Vega and Lueker [12] and for the square packing above do not seem to extend here. We adopt the following approach. Because we are allowed to enlarge the bin sizes from $[0,1]^2$ to $[0,1+\varepsilon]^2$, we show that it is not a problem if we round the sides of rectangles that are both tall and wide (i.e., whose height and width are both larger than a small constant δ). This enables us to reduce the rectangles that are both wide and tall to a constant number of types. The main problem is to deal with the rectangles that are either very wide and flat or very thin and tall. For either kind of these rectangles, we show that we can essentially use the strip-packing algorithm of Kenyon and Rémila [18]. Our algorithm thus relies on an appropriate partition of rectangles into "large," "small," "horizontal," "vertical," and "medium," in such a way that the medium rectangles are negligible. It only remains to mix gracefully the various kinds of rectangles; this requires discretizing a near-optimal solution appropriately so as to be able to "guess" not only where the large rectangles go but also the areas used to pack horizontal rectangles and the areas used to pack vertical rectangles.

Our result for the minimum encasing rectangle is obtained by using the previous algorithm for packing rectangles using resource augmentation. Specifically, we use the previous algorithm to decide whether a list of rectangles does not fit in a given rectangle or they fit in a slightly larger rectangle. We then do exhaustive search over all possible rectangles and select the one of minimum area. This procedure takes care of all instances having at least one "wide" and one "tall" rectangle. The rest of the instances are solved by a slight variation of a strip-packing type result in Coffman et al. [8].

2. APX hardness of two-dimensional bin packing.

We give an approximation preserving reduction from the maximum bounded three-dimensional matching problem (MAX-3-DM). The MAX-3-DM problem is defined as follows.

INPUT. Three sets $X = \{x_1, \ldots, x_q\}$, $Y = \{y_1, \ldots, y_q\}$, and $Z = \{z_1, \ldots, z_q\}$. A subset $T \subseteq X \times Y \times Z$ such that any element in X, Y, Z occurs in one, two, or three triples in T. Note that this implies that $q \leq |T| \leq 3q$.

GOAL. Find a maximum cardinality subset T' of T such that no two triples in T' agree in any coordinate.

MEASURE. Cardinality of T'.

Kann [16] was the first who proved the MAX SNP hardness of the MAX-3-DM problem. Recently, explicit lower bounds on the approximation ratio of MAX-3-DM have also been obtained Berman and Karpinski [2] and Hazan et al. [14]. Petrank [28] (Theorem 4.4) proved a refined hardness result, where he showed that it is NP-hard to distinguish between instances where $|T'| = q$ and instances where $|T'| \leq (1-\varepsilon)q$ for some constant $\varepsilon > 0$.

Given an instance I of MAX-3-DM we will construct an instance of two-dimensional bin packing. The following is a high-level description of the steps involved. First, we create a numerical version of the MAX-3-DM instance by associating an integer with each element in X, Y, Z, and T. These integers will have the property that any four of them sum up exactly to a number B, if and only if the integers correspond to a triple $t_l = (x_i, y_j, z_k)$. Let x'_i (resp. y'_j, z'_k, t'_l) denote the integer corresponding to x_i (resp. x_j, z_k, t_l). Next, with each integer we will associate two rectangles. Each rectangle has width close to $1/4$ and height close to $1/2$. Each of the rectangles will have an area of about $1/8$, so no more than eight rectangles can fit in a bin. The exact width and height of a rectangle is a small perturbation around $1/4$ and $1/2$, respectively, depending of the value of the integer associated to it. Typically, one of the rectangles will be "thin" and "tall," and the other rectangle will be "wide" and "short." These perturbations are chosen such that the following property is satisfied. If $t_l = (x_i, y_j, z_k)$ or equivalently $t'_l + x'_i + y'_j + z'_k = B$, then the eight rectangles corresponding to x'_i, y'_j, z'_k, and t'_l can fit in a bin. Otherwise, each bin is suboptimal in the sense that it either contains at most seven rectangles or contains at most three "tall" and "thin" rectangles. We begin by describing the construction and then prove some useful structural properties of this construction. We then show how these properties, together with the APX hardness of MAX-3-DM, imply the asymptotic hardness of two-dimensional bin packing.

We begin by defining the integers corresponding to x_i, y_j, z_k, t_l. Let $r = 32q$. Define

$$x'_i = ir^3 + i^2 r + 1, \quad \text{for } 1 \leq i \leq q,$$
$$y'_j = jr^6 + j^2 r^4 + 2, \quad \text{for } 1 \leq j \leq q,$$
$$z'_k = kr^9 + k^2 r^7 + 4, \quad \text{for } 1 \leq k \leq q.$$

For every triple $t_l = (x_i, y_j, z_k)$ in T, we define

$$t'_l = r^{10} - x'_i - y'_j - z'_k + 15 = r^{10} - kr^9 - k^2 r^7 - jr^6 - j^2 r^4 - ir^3 - i^2 r + 8.$$

Let $\delta = 1/500$ and let $c = (r^{10} + 15)/\delta$. Observe that $0 < x'_i, y'_j, z'_k < \delta c/10$ and $t'_l < \delta c$ hold for all i, j, k, l. Intuitively, think of c as a scaling factor much larger than x'_i, y'_j, z'_k, and b'_l.

We now describe the rectangles in our instance. A rectangle of width w and height h will be denoted by (w, h). For each element $x_i \in X$ (resp. $y_i \in Y$ and $z_i \in Z$), we define a pair of rectangles $a_{x,i}$, $a'_{x,i}$ (resp. $a_{y,i}$, $a'_{y,i}$, and $a_{z,i}$, $a'_{z,i}$) as follows (the reader is encouraged to look at Figure 1):

$$a_{x,i} = \left(\frac{1}{4} - 4\delta + \frac{x'_i}{c}, \frac{1}{2} + 4\delta - \frac{x'_i}{c}\right) \quad \text{and} \quad a'_{x,i} = \left(\frac{1}{4} + 4\delta - \frac{x'_i}{c}, \frac{1}{2} - 4\delta + \frac{x'_i}{c}\right),$$

$$a_{y,i} = \left(\frac{1}{4} - 3\delta + \frac{y_i'}{c}, \frac{1}{2} + 3\delta - \frac{y_i'}{c}\right) \quad \text{and} \quad a_{y,i}' = \left(\frac{1}{4} + 3\delta - \frac{y_i'}{c}, \frac{1}{2} - 3\delta + \frac{y_i'}{c}\right),$$

$$a_{z,i} = \left(\frac{1}{4} - 2\delta + \frac{z_i'}{c}, \frac{1}{2} + 2\delta - \frac{z_i'}{c}\right) \quad \text{and} \quad a_{z,i}' = \left(\frac{1}{4} + 2\delta - \frac{z_i'}{c}, \frac{1}{2} - 2\delta + \frac{z_i'}{c}\right).$$

As x_i'/c, y_i'/c, and z_i'/c are all no more than $\delta/10$, observe that for each pair of rectangles, the first rectangle is "thin" and "tall" and the second rectangle is "wide" and "short." Also observe how the perturbations in height and width are related to each other (this fact is made precise in Lemmas 2.1 and 2.2 below).

Next for each $t_l \in T$ we define two rectangles b_l and b_l' as

$$b_l = \left(\frac{1}{4} + 8\delta + \frac{t_l'}{c}, \frac{1}{2} + \delta - \frac{t_l'}{c}\right) \quad \text{and} \quad b_l' = \left(\frac{1}{4} - 8\delta - \frac{t_l'}{c}, \frac{1}{2} - \delta + \frac{t_l'}{c}\right).$$

Finally, we define D to be a collection of $|T| - q$ dummy rectangles d_i such that $d_i = (3/4 - 10\delta, 1)$. These dummy rectangles will not play a significant role and will just serve as "garbage collectors."

Formally, we say that two rectangles a and a' are buddies iff $\{a, a'\}$ is $\{a_{x,i}, a_{x,i}'\}$ or $\{a_{y,j}, a_{y,j}'\}$ or $\{a_{z,k}, a_{z,k}'\}$ for $1 \le i, j, k \le q$ or $\{a, a'\} = \{b_l, b_l'\}$ for some $1 \le l \le |T|$.

Let $A_x = \{a_{x,1}, \ldots, a_{x,q}\}$ and $A_x' = \{a_{x,1}', \ldots, a_{x,q}'\}$. A_y, A_y', A_z, and A_z' are defined similarly to be the set of rectangles $a_{y,i}$, $a_{y,i}'$, $a_{z,i}$, and $a_{z,i}'$ respectively. Let A denote the collection $A_x \cup A_y \cup A_z$ and $A' = A_x' \cup A_y' \cup A_z'$ and, finally, let $B = \{b_1, \ldots, b_{|T|}\}$ and $B' = \{b_1', \ldots, b_{|T|}'\}$. Having all the notation in place, we first list some simple observations about these rectangles that follow directly from the choice of sizes of these rectangles.

Lemma 2.1. *For each rectangle $a \in A \cup A'$, $w(a) + h(a) = 3/4$. For any $b \in B$, $w(b) + h(b) = 3/4 + 9\delta$ and for any $b' \in B'$, $w(b') + h(b') = 3/4 - 9\delta$. In particular, this implies that for any $b \in B$ and $b' \in B'$, $h(b) + h(b') + w(b) + w(b') = 3/2$.*

Lemma 2.2. *For any two rectangles a, a' in $A \cup A' \cup B \cup B'$, $h(a) + h(a') = 1$ iff a and a' are buddies.*

Lemma 2.3. *The width (resp. height) of any rectangle in the instance is at least $1/4 - 10\delta$ (resp. at least $1/2 - 5\delta$) and, hence, the area of any rectangle is at least $1/8 - 25/4\delta > 1/9$. Thus, no bin can have more than eight rectangles.*

Lemma 2.4. *For any packing of squares in a bin, every vertical line (i.e., of the form $x = c$), intersects at most one rectangle from $A \cup B$. This follows as the height of each rectangle in $A \cup B$ is strictly greater than $1/2$.*

This observation implies that any bin can contain at most four rectangles in $A \cup B$. Moreover, any bin can contain at most three rectangles in B, as the width of each rectangle in $A \cup B$ is more than $1/5$, and the width of each rectangle in B is more than $1/4$. Finally, it is easy to see that:

Lemma 2.5. *If a rectangle $d_i \in D$ lies in some bin S, then S contains at most two other rectangles and at most one of them is a rectangle from $A \cup B$.*

Next, we give two less obvious consequences (Lemmas 2.6 and 2.7) of the choices of sizes of these rectangles. The following definition is needed only for the next two lemmas. For a rectangle a, we now define $\Delta(a)$ which allows us to relate the rectangle back to the integers x_i', y_j', z_k', or b_l'. For a rectangle $a_{x,i} \in A_x$, let $\Delta(a_{x,i}) = x_i'$ and for $a_{x_i}' \in A_x'$, $\Delta(a_{x,i}') = -x_i'$. Similarly, $\Delta(a_{y,j}) = y_j'$, $\Delta(a_{z,k}) = z_k'$, $\Delta(b_l) = t_l'$, and $\Delta(a_{y,j}') = -y_j'$, $\Delta(a_{z,k}') = -z_k'$, $\Delta(b_l') = -t_l'$.

Lemma 2.6. *For any three rectangles $a_1, a_2, a_3 \in A$ and $b \in B$, we have that $w(a_1) + w(a_2) + w(a_3) + w(b) = 1$ iff there is a triple $t_l = (x_i, y_j, z_k)$ in the instance I of MAX-3-DM and $\{a_1, a_2, a_3, b\} = \{a_{x,i}, a_{y,j}, a_{z,k}, b_l\}$.*

Proof. The "if" part follows directly from the choice of the sizes of the rectangles. In particular, if $t_l = (x_i, y_j, z_k)$, then

$$w(a_{x,i}) + w(a_{y,j}) + w(a_{z,k}) + w(b_l) = 1 - \delta + \frac{(x_i' + y_j' + z_k' + t_l')}{c} = 1 - \delta + \frac{r^{10} + 15}{c} = 1.$$

For the "only if" part of the lemma, we first observe that

$$w(a_1) + w(a_2) + w(a_3) + w(b) = 1 - \delta + \frac{\Delta(a_1) + \Delta(a_2) + \Delta(a_3) + \Delta(b)}{c}.$$

Thus, if $w(a_1)+w(a_2)+w(a_3)+w(b)=1$, it must be that

$$\sum_{i=1}^{3}\Delta(a_i)+\Delta(b)=\delta c=r^{10}+15.$$

Consider the quantity $\sum_{i=1}^{3}\Delta(a_i)+\Delta(b)$ modulo r. Since r is a multiple of 32, it follows that there is exactly one rectangle each from A_x, A_y, A_z, and B since $15=1+2+4+8$ and this is the only way to represent 15 as a multisum (with possible repetitions) of four numbers from the set $\{1,2,4,8\}$. Next, considering the sum $\sum_{i=1}^{3}\Delta(a_i)+\Delta(b)$ modulo r^2 it follows that i and l are such that $x_i\in t_l$. Similarly, considering modulo r^4 and r^7, it follows that $y_j\in t_l$ and $z_k\in t_l$. As x_i, y_j, and z_k determine t_l uniquely, the result follows. □

Lemma 2.7. *Let $\mathcal{R}=\{a_1,a_2,a_3,a_4\}$ be a set of any four rectangles lying in $A\cup A'$, such that no two of them are buddies. Then $\sum_{i=1}^{4}w(a_i)\neq 1$.*

Proof. Suppose for the sake of contradiction that $\sum_{i=1}^{4}w(a_i)=1$. As $0\le|\Delta(a_i)|\le\delta c/10$, it must be that $\sum_{i=1}^{4}\Delta(a_i)=0$. We first show that there cannot be more than one rectangle from $A_x\cup A'_x$ in the set $\mathcal{R}$ (later we will show that the same argument works for $A_y\cup A'_y$ and $A_z\cup A'_z$), which contradicts that $\mathcal{R}$ contains four rectangles in $A\cup A'$. Consider the coefficient of r and r^3 in $\sum_{i=1}^{4}\Delta(a_i)$. These coefficients depend on rectangles in $(A_x\cup A'_x)\cap\mathcal{R}$ only. Let i_1, i_2, i_3, i_4 denote the indices of rectangles from $(A_x\cup A'_x)\cap\mathcal{R}$ (where an index is 0, if fewer than 4 occur). Since no two rectangles are buddies, we cannot have both a_{x,i_k} and a'_{x,i_k} in $\mathcal{R}$. So, for each i_k, $1\le k\le 4$, we associate a variable $\alpha(i_k)$, where $\alpha(i_k)=1$ if $a_{x,i_k}\in\mathcal{R}$ and -1 if $a'_{x,i_k}\in\mathcal{R}$.

Since the coefficients of r and r^3 sum to 0, we must have that $\sum_{k=1}^{4}\alpha(i_k)i_k=0$ and $\sum_{k=1}^{4}\alpha(i_k)i_k^2=0$, with the constraints that all nonzero i_ks are distinct (since no two rectangles are buddies). We now claim that the only feasible solution to this system is $i_k=0$ for all $1\le k\le 4$. As all the i_ks are distinct, we need at least three of the i_ks to be nonzero, else we cannot have $\sum_k\alpha(i_k)i_k=0$. Also, it is not the case that all $\alpha(i_k)$ are either -1 or that all are $+1$. Under these constraints, when there are exactly three nonzero i_ks, without loss of generality, we have the following system of equations: $i_1+i_2=i_3$ and $i_1^2+i_2^2=i_3^2$ such that $i_1,i_2,i_3>0$. However, squaring the first equation and subtracting the second from it implies that $2i_1i_2=0$, which contradicts that all of i_1, i_2, and i_3 are nonzero. Thus, there is no solution to this system of equations.

Finally, we consider the case when all the i_ks are nonzero. As not all $\alpha(i_k)$ have the same sign, without loss of generality we have only the following two cases.

(i) $i_1+i_2=i_3+i_4$ and $i_1^2+i_2^2=i_3^2+i_4^2$,

(ii) $i_1+i_2+i_3=i_4$ and $i_1^2+i_2^2+i_3^2=i_4^2$.

As all i_1, i_2, i_3, and i_4 are nonzero, the last case clearly does not have a solution. For the first case, rewrite the equations as $i_1-i_3=i_4-i_2$ and $i_1^2-i_3^2=i_4^2-i_2^2$. As i_1, i_2, i_3, and i_4 are all pairwise distinct, this implies that $i_1+i_3=i_4+i_2$. Together with $i_1+i_2=i_3+i_4$, this implies that $i_1=i_4$, which gives a contradiction. Repeating a similar argument for $A_y\cup A'_y$ and $A_z\cup A'_z$ by considering the coefficients of r^4 and r^6 and those of r^7 and r^9, respectively, it follows that there can be at most one rectangle each from $A_y\cup A'_y$ and $A_z\cup A'_z$ in $\mathcal{R}$, which contradicts the assumption that $\mathcal{R}$ contains four rectangles. □

We now have enough structural results about how rectangles can be packed in a bin. This will allow us to show that if a bin does not contain eight rectangles corresponding to a triple, then it is suboptimally packed. We make the notion of "suboptimally packed" precise with the following definition of a good bin.

Definition 2.1. Given a packing of the bins, call a bin *good* if it contains exactly eight rectangles and additionally it has exactly four rectangles from $A\cup B$.

The following crucial lemma characterizes the structure of good bins.

Lemma 2.8. *A bin is good if and only if it contains the rectangles $a_{x,i}$, $a_{y,j}$, $a_{z,k}$, b_l and the corresponding rectangles $a'_{x,i}$, $a'_{y,j}$, $a'_{z,k}$, b'_l such that $t_l=(x_i,y_j,z_k)$ corresponds to a triple in the MAX-3-DM instance.*

Proof. We first show that the rectangles corresponding to a triple can be packed in a bin. Starting from the bottom left corner of the bin and moving towards the right, we pack the rectangles $a_{x,i}$, $a_{y,j}$, $a_{z,k}$, and b_l. Each of these rectangles is placed such that it touches the bottom of the bin.

Figure 1 shows the packing. It is easy to verify that these four rectangles can be packed as described, as

$$w(a_{x,i})+w(a_{y,j})+w(a_{z,k})+w(b_l)=1-\delta+\frac{(x'_i+y'_j+z'_k+t'_l)}{c}=1.$$

Next we observe that each of the rectangles $a'_{x,i}$, $a'_{y,j}$, $a'_{z,k}$, and b'_l can be placed in the remaining gaps (as shown in Figure 1). Clearly, $a'_{x,i}$ can be placed on top of $a_{x,i}$ because $h(a_{x,i})+h(a'_{x,i})=1$ and as $h(a_{x,i})>$

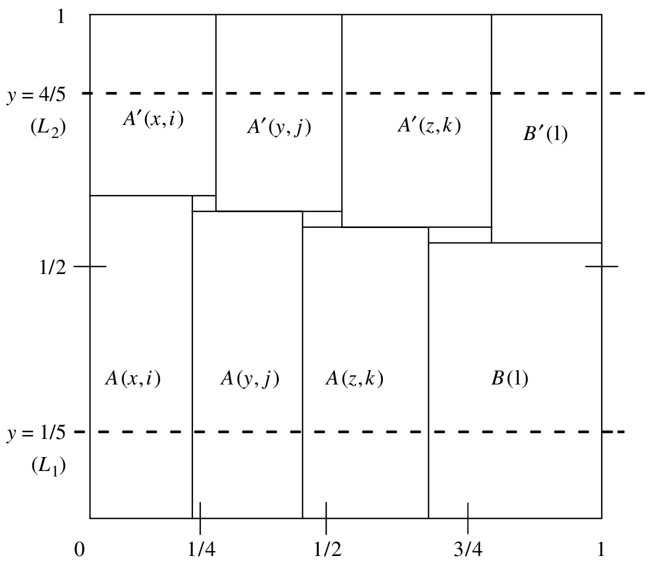

FIGURE 1. Packing of the rectangles corresponding to a triple.

$h(a_{y,j}) > h(a_{z,k}) > h(b_l)$, this allows $a'_{x,i}$ to extend horizontally beyond $a_{x,i}$. Arguing similarly and observing that $w(a'_{x,i}) + w(a'_{y,j}) + w(a'_{z,k}) + w(b'_l) = 1$, it is easy to see that rectangles $a'_{y,j}$, $a'_{z,k}$, and b'_l also fit.

We now show that any good bin must correspond to a triple. We first give a way for labeling the eight rectangles in a good bin. Consider the lines $L_1 = \{y = 1/5\}$ and $L_2 = \{y = 4/5\}$. It is easy to see that in any packing of a bin with eight rectangles, each rectangle must intersect exactly one of L_1 or L_2. This follows because any rectangle has height at most $1/2 + 1/50 < 3/5$ and at least $1/2 - 1/50 > 2/5$. Moreover, as any rectangle has width strictly larger than 1/5, it follows that each L_1 and L_2 intersects exactly four rectangles. Let $\{a_1, a_2, a_3, a_4\}$ denote the rectangles that intersect L_1 such that a_i is to the left of a_j for $i < j$. Similarly, let $\{a_5, a_6, a_7, a_8\}$ denote the rectangles that intersect L_2 in the left-to-right order. Thus, we have that

$$w(a_1) + w(a_2) + w(a_3) + w(a_4) \le 1, \tag{1}$$

and that

$$w(a_5) + w(a_6) + w(a_7) + w(a_8) \le 1. \tag{2}$$

Finally, as the width of each rectangle is at least 1/5, for each $1 \le i \le 4$, it must be that rectangle a_i overlaps with a_{i+4} in an x-coordinate (i.e., there is a vertical line $x = c$ for some c that intersects both a_i and a_{i+4}), implying that

$$h(a_i) + h(a_{i+4}) \le 1 \quad \text{for } 1 \le i \le 4. \tag{3}$$

We now show that any good bin must contain exactly one rectangle from B.

(i) Suppose that at least two rectangles in B lie in a good bin. As each rectangle in B has width at least $1/4 + 8\delta$, two such rectangles use at least $1/2 + 16\delta$. Thus, the width left for rectangles from A is at most $1/2 - 16\delta$ (we cannot put rectangles from A on top of the rectangles from B), and, hence, at most one rectangle from A can fit. Similarly, if exactly three rectangles from B lie in the bin, there cannot be any rectangle from A. Thus, in either case there are at most three rectangles from $A \cup B$, which contradicts the fact that the bin is good.

(ii) Suppose that no rectangle in B lies in a good bin. We claim the bin cannot have more than three rectangles from A. For the sake of contradiction suppose $r_1, r_2, r_3, r_4 \in A$ lie in the bin. Then no rectangle from B' can lie in the bin because any rectangle in A has height at least $1/2 + 2\delta$ while the height of any rectangle in B' is at least $1/2 - \delta$ and, moreover, any rectangle in B' must overlap in some x-coordinate with some r_i for $1 \le i \le 4$. Thus, all the eight rectangles lie in $A \cup A'$.

Adding (1), (2), and (3) we have that $\sum_{i=1}^{8}(w(a_i) + h(a_i)) \le 6$. Moreover from Lemma 2.1, as $w(a) + h(a) = 3/4$ for each rectangle $a \in A \cup A'$, we have that $\sum_{i=1}^{8}(w(a_i) + h(a_i)) = 6$; thus it must be the case that each of the inequalities (1), (2), and (3) must hold with equality. By Lemma 2.2, this implies that a_i and a_{i+4} are buddies for each $i = 1, \ldots, 4$. This in particular implies that no two rectangles are buddies among the rectangles a_1, a_2, a_3, and a_4. Therefore, by Lemma 2.7, it is impossible that $\sum_{i=1}^{4} w(a_i) = 1$ and, hence, we have a contradiction.

Thus, any good bin can contain exactly one rectangle from B. To finish the proof, we show that if b_l lies in a good bin, then it must contain rectangles corresponding to the triple $t_l = (x_i, y_j, z_k)$.

First, by definition, any good bin with exactly one rectangle b_l from B must contain exactly three other rectangles from A. Because any rectangle from B' cannot overlap in an x-coordinate with any rectangle in A, it follows that there can be *at most* one rectangle from B', which must overlap with b_l. Next, we show that there has to be *at least* one rectangle from B'. Suppose there are no rectangles from B'; then we claim that the total width of all the rectangles must be strictly larger than 2, which is not possible. To see this, because there is no rectangle from B', there must be four rectangles from A'. As the width of any rectangle in A' (resp. B) is strictly more than $1/4+\delta$ (resp. $1/4+8\delta$), any four rectangles from A', together with any rectangle from B, take up width strictly larger than $5/4+12\delta$. However, as the width of any rectangle in A is at least $1/4-4\delta$, we do not have sufficient total width to pack three rectangles from A, which contradicts the fact that the bin is good.

Hence, there is exactly one rectangle from B'. Call it $b'_{l'}$. By Lemma 2.1, we have that $h(b_l)+w(b_l)+h(b'_{l'})+w(b'_{l'}) = 3/2$. Moreover, observing that $w(a)+h(a) = 3/4$ for each $a \in A\cup A'$, it follows that each of the inequalities (1), (2), and (3) is satisfied with equality. To complete the argument, suppose b_l intersects line L_1, let a_1, a_2, a_3 denote the rectangles in $A\cup A'$, which also intersect L_1. Thus, we have that $w(a_1)+w(a_2)+w(a_3)+w(b_l) = 1$. None of the a_i, $1\le i\le 3$ can lie in A', because otherwise it is always the case that $w(a_1)+w(a_2)+w(a_3)+w(b_l) > 1$. Because all a_1, a_2, and a_3 lie in A, by Lemma 2.6 it follows that these rectangles correspond to a triple $t_l = (x_i, y_j, z_k)$. □

THEOREM 2.1. *There is no asymptotic PTAS for the two-dimensional bin-packing problem unless $P = NP$.*

PROOF. If the MAX-3-DM problem has a matching consisting of q triples, then we can get a bin-packing solution that uses $|T|$ bins as follows. For each of the triples in the matching, create a good bin as described in Lemma 2.8. For each t_l not in the matching, we put b_l and b'_l along with a dummy rectangle and, hence, we use $q+(|T|-q) = |T| \le 3q$ bins.

Assume now that every feasible solution of the MAX-3-DM problem has at most $(1-\varepsilon)q$ triples. We will show that any solution to the corresponding bin packing problem uses at least $(1+\varepsilon/33)|T|$ bins. Consider any feasible solution to the bin-packing instance. There will be exactly $n_d = |T|-q$ bins with dummy objects. Let n_g denote the number of good bins. Because the set of good bins corresponds to some feasible solution by Lemma 2.8 we have $n_g \le (1-\varepsilon)q$. By Lemmas 2.3 and 2.4, and Lemma 2.8, it follows that if a bin is not good, then it either has at most seven rectangles or else it has at most three rectangles from $A\cap B$. Among the bins that are not good let n_{b_1} denote the number of bins (other than the bins with dummy objects) that contain at most seven rectangles and let n_{b_2} denote bins that have eight rectangles but three and fewer rectangles from $A\cup B$.

Because any solution must cover all of the rectangles in $A\cup B$ and any bin with a dummy rectangle can have at most one rectangle from $A\cup B$, we have that

$$4n_g + 4n_{b_1} + 3n_{b_2} + n_d \ge 3q + |T|.$$

Equivalently,

$$4n_g + 4n_{b_1} + 3n_{b_2} \ge 4q.$$

Finally, since all the rectangles in $A\cup A'\cup B\cup B'$ must be covered, we have that

$$8n_g + 7n_{b_1} + 8n_{b_2} + 2n_d \ge 6q + 2|T|.$$

Equivalently,

$$8n_g + 7n_{b_1} + 8n_{b_2} \ge 8q.$$

Adding the inequalities above, $12n_g + 11n_{b_1} + 11n_{b_2} \ge 12q$. Equivalently, $n_g + n_{b_1} + n_{b_2} \ge 12q/11 - n_g/11$. Adding the bins with dummy objects, this implies that the total number of bins used is at least $|T| - q + 12q/11 - n_g/11 = |T| + (q-n_g)/11 \ge |T| + \varepsilon q/11 \ge |T|(1+\varepsilon/33)$.

Now if there is an APTAS for two-dimensional bin packing, then for every $\varepsilon > 0$, there exists an algorithm A_ε and a constant c_ε, such that for instances I if $|\mathrm{OPT}(I)| > c_\varepsilon$, then $A_\varepsilon \le (1+2\varepsilon)|\mathrm{OPT}(I)|$. Thus, for any $\varepsilon > 0$, if $q > c_\varepsilon$ we can distinguish between two instances of the MAX-3-DM problem with $|T'| = q$ and $|T'| \le (1-66\varepsilon)q$, which is an NP-hard problem by Petrank [28]. □

3. Hypercube packing. In this section we describe our asymptotic approximation scheme for packing a collection of d-dimensional cubes into a minimum number of unit cubes. Our result will consist of three parts. In §3.2 we show how to pack small cubes into rectangular regions without wasting too much space. In §3.3 we show how to compute an almost optimum packing of large cubes into unit bins. Finally, in §3.4 we show how to combine these two ideas together to get a close-to-optimum packing of large and small cubes simultaneously. We begin with some preliminaries.

3.1. Definitions and preliminaries. We adopt two standard assumptions in multidimensional bin packing: first, items are allowed to "touch," i.e., they can intersect in a face; second, items cannot be rotated (*orthogonal packing without rotation*).

A d-dimensional cube is given by a positive number a, representing the length of its side. Since we will pack squares into unit squares, we need to assume $a \leq 1$. Given an input list I of n d-dimensional cubes with sides of length $a_i \in (0, 1]$ for $i = 1, \ldots, n$, the volume of the list is defined as: $\mathrm{Vol}(I) = \sum_{i=1}^{n} a_i^d$. Given two cubes and their positions in the space, we say that they are *nonoverlapping* if their interiors are disjoint.

A packing of I into k bins is a positioning of the cubes into k copies $H_1, \ldots, H_k$ of the unit hypercube $[0, 1]^d$, so that no two cubes overlap. The d-dimensional cube-packing problem consists of finding a packing of I into the minimum number of bins, $\mathrm{OPT}(I)$. We denote by $A(I)$ the number of bins algorithm A uses to pack I.

We now establish a simple result that will be needed later.

Lemma 3.1. *Let P be a packing of m d-dimensional cubes in $[0, 1]^d$. Then the unused space in the bin, denoted by $[0, 1]^d \setminus P$, can be divided into at most $(2m)^d$ nonoverlapping d-dimensional rectangles.*

Proof. Extending each facet of each cube in P, we obtain a grid in $[0, 1]^d$ consisting of $(2m+1)^d$ cells. By pushing cubes toward the origin, in Lemma 3.1 we can assume without loss of generality that in the packing P there is a cube touching each of the hyperplanes $x_i = 0$, for $i = 1, \ldots, d$. Thus, $[0, 1]^d \setminus P$ can be seen as the union of no more than $(2m)^d$ nonoverlapping d-dimensional rectangles. □

We can make a stronger statement for the two-dimensional case.

Lemma 3.2 (The Two-Dimensional Case). *Let P be a packing of m squares in $[0, 1]^2$. Then, $[0, 1]^2 \setminus P$ can be divided into at most $3m$ nonoverlapping rectangles.*

Proof. For each square in P draw a horizontal line at its top and bottom until it intersects some other square, as in Figure 2. Assuming the bottom-most square is at height 0, we have drawn $2m - 1$ horizontal lines. This partitions the unit square into at most $3m$ rectangles (two to the side of each square and another to the top of each square in P) plus the original m squares. □

3.2. Packing small cubes. In this section we look at multidimensional cube packing in case all cubes are small. In this setting we are given an input list $S = (a_1, \ldots, a_n)$ of cubes with sides $a_i \leq \delta$ for some positive (small) constant δ. We will use a multidimensional version of the next-fit-decreasing-height shelf heuristic (NFDH). Coffman et al. [8] analyzed this heuristic in the context of strip packing, pointing out properties closely related to what we extend here to higher dimensions.

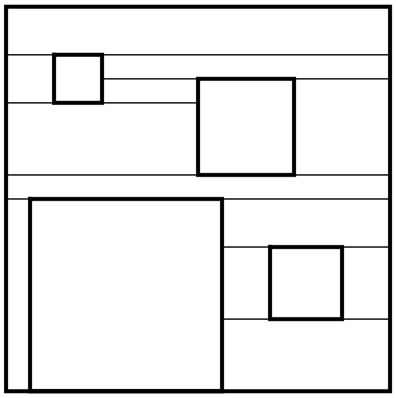

Figure 2. Decomposing the used space in the bin into rectangular regions.

Let us assume that the cubes are sorted according to nonincreasing order of sizes. In two dimensions, the NFDH algorithm is a level algorithm that uses the next-fit approach to pack the sorted list of cubes. The cubes are packed, left-justified on a shelf, until the next rectangle will not fit. This rectangle is used to define a new shelf, and the packing continues on this shelf. The earlier shelves are not revisited. In general, for higher dimensions, we define the NFDH heuristic inductively. Assume we know how to perform NFDH in $d-1$ dimensions. The d-dimensional NFDH heuristic will consider a facet $\mathscr{F}$ of the bin and pack cubes on it using its $(d-1)$-dimensional version. For that, one of the dimensions (say, the height) of the d-dimensional cubes is ignored, and these cubes are seen as $(d-1)$-dimensional cubes. After it finishes packing in this facet, it will cut off from the bin a d-dimensional rectangle (*shelf*) with base $\mathscr{F}$ and height a_1. It will proceed to pack the remaining list in the rest of the bin. The following lemma illustrates a fundamental property of NFDH. It is a consequence of a result of Meir and Moser [24], but we include a proof for completeness.

Lemma 3.3. *Let C be a set of d-dimensional cubes (where $d \geq 2$) with sides smaller than δ. Consider the NFDH heuristic applied to C. If NFDH cannot place any other cube in a rectangle R of size $r_1 \times r_2 \times \cdots \times r_d$ (with $r_i \leq 1$), the total wasted (unfilled) volume in that bin is at most:*

$$\delta \sum_{i=1}^{d} r_i.$$

Proof. Let $h_1, h_2, \ldots, h_l$ denote the height of the shelves generated by the algorithm. Similarly, let c_i and d_i denote the sides of the largest and smallest cube, respectively, in shelf i. Clearly, since the cubes are packed according to nonincreasing sizes, we have that $c_{i+1} \leq d_i$ for all $i = 1, \ldots, l-1$. Similarly, we have that $\sum_{i=1}^{l} h_i \geq r_d - c_{l+1} \geq r_d - d_l$; otherwise, the algorithm would open another shelf.

We will prove the required result by induction on the number of dimensions. For $d = 2$, from the result of Coffman et al. [8], we know that the wasted volume is at most $\delta(r_1 + r_2)$. Suppose the result is true for $d-1$; hence, for each $(d-1)$-dimensional facet, the wasted volume along the first $d-1$ dimensions is at most $\delta \sum_{i=1}^{d-1} r_i$. We now bound the wasted volume in d dimensions. First, the volume wasted in the topmost shelf (that cannot be packed with any cube) is at most $d_l \prod_{i=1}^{d-1} r_i$. Second, for each shelf $i \leq l$, the wasted volume in the $(d-1)$-dimensional facet contributes at most $h_i \cdot \delta(\sum_{j=1}^{d-1} r_j)$, and the additional volume wasted along the dth dimension can be bounded by $(c_i - d_i) \cdot \prod_{j=1}^{d-1} r_j$. Adding all of the contributions of all the shelves and the topmost unpacked shelf, we get that the total wasted volume is at most $c_1 \cdot \prod_{j=1}^{d-1} r_j + \delta(\sum_{j=1}^{d-1} r_j)$. Without loss of generality, suppose that $r_1 \leq r_2 \leq \cdots \leq r_d$. Because $c_1 \leq \delta$ and $\prod_{j=1}^{d-1} r_j \leq r_d$, the total waste is bounded by $\delta(\sum_{i=1}^{d} r_i)$, which implies the result. □

3.3. Packing large cubes. We now concentrate on the case of packing large cubes. In this case we are given an input list L of n cubes whose sides are at least δ. For simplicity, we split the analysis into two steps (as in Fernandez de la Vega [12] and Vazirani [32]).

Lemma 3.4. *Suppose the input list L contains large cubes only, say, $a_i \geq \delta$ for all $i = 1, \ldots, n$, and there are only K different cube sizes for some constant K. Then, we can solve the cube-packing problem optimally in time $O(\text{polylog}(n))$.*

Proof. Clearly, the number of cubes that fit in a bin is bounded by $M = \lfloor 1/\delta^d \rfloor$. For a cube of type i, we denote its side length by α_i. Given a packing of cubes in a bin, its bin type will be described by the number of cubes of each type in that bin. This will be denoted by a vector $(x_1, x_2, \ldots, x_K)$, which means that there are x_1 cubes of type 1, x_2 cubes of type 2, up to x_K cubes of type K in that bin. Clearly, there are at most M^K (a constant) bin types. We now show that, given a vector $(x_1, x_2, \ldots, x_K)$ such that $x_i \leq M$ for each $1 \leq i \leq K$, we can test in constant time whether it denotes a valid bin type. To see this, we show that there are only a constant number of possible positions for each cube in a bin at which it can be placed. We can then test in constant time whether a vector is valid, by exhaustively trying all possible positions for the cubes in that vector.

If a given vector of cubes fit in a bin, without loss of generality we can assume that their vertices lie at coordinates belonging to the set:

$$\mathscr{C} = \left\{ \sum_{i=1}^{K} \lambda_i \alpha_i \colon \lambda_i \in \mathbb{Z}_+ \text{ and } 0 \leq \sum_{i=1}^{K} \lambda_i \alpha_i \leq 1 \right\}.$$

Indeed, given a general placement of the cubes in the bin, fix a dimension $l = 1, \ldots, d$ and sort the cubes by their leftmost point in their lth coordinate. Now, consider each cube in this sorted order and push them one-by-one to the left until the cube touches another cube, which always happens at a coordinate in $\mathscr{C}$. By doing this for all coordinates we obtain the claimed result. Let C denote the cardinality of the set $\mathscr{C}$. Since $\alpha_i \geq \delta$, any point $\sum_{i=1}^{K} \lambda_i \alpha_i$ in $\mathscr{C}$ will satisfy $\lambda_i \leq 1/\delta$ for all $i = 1, \ldots, K$. Thus $C \leq (1/\delta)^K$, which is a (huge) constant. Thus there are at most $(C^d)^M$ ways to specify the position of each possible cube in a bin.

As mentioned above, we can thus check in constant time if a given vector $(x_1, \ldots, x_K)$ is a bin type or not. Thus, by exhaustive enumeration we can generate in constant time a list $\mathscr{T} = \{T_1, \ldots, T_Q\}$ consisting of all valid bin types.

It only remains to see how to find an optimal packing. To this end we will use an integer program of size $O(\text{polylog}(n))$ with a fixed number of variables and constraints. For $1 \leq j \leq K$, let T_{ij} denote the number of type j cubes in bin type T_i and let n_j denote the number of cubes of type j in the problem instance. Let x_i correspond to the number of bins that have configuration corresponding to the type T_i. We can formulate an integer program (denoted by IP) in the x variables as follows:

$$\min \sum_{i=1}^{Q} x_i, \tag{4}$$

$$\begin{aligned} \sum_{i=1}^{Q} T_{ij} x_i &\geq n_j \quad \text{for } j = 1, \ldots, K, \\ x_i &\geq 0 \quad \text{for } i = 1, \ldots, Q, \\ x_i &\in \mathbb{Z}_+ \quad \text{for } i = 1, \ldots, Q. \end{aligned}$$

Clearly the size of IP is $O(\log(n))$. Moreover, it only has a constant number of variables. Therefore, by Lenstra's algorithm for integer programming in fixed dimension Lenstra [22], we can find an optimum integer solution to IP in time $O(\text{polylog}(n))$. It is straightforward to see that such a solution represents an optimal packing. □

LEMMA 3.5. *Suppose the input list L contains large cubes ($a_i \geq \delta$ for all $i = 1, \ldots, n$) only. Then, for any $\varepsilon > 0$, there is a $(1+\varepsilon)$-approximation algorithm for packing cubes into a minimum number of unit cubes that runs in time $O(n \log(n))$.*

PROOF. Let L denote the given instance list. The algorithm is as follows:

(i) Sort the n cubes in nonincreasing order of sizes and partition them in $K = \lceil 1/(\varepsilon\delta^d) \rceil$ groups, each containing at most $Q = \lceil n/K \rceil$ cubes.

(ii) Construct an instance J by rounding up all items in each group to the largest cube in the group. By construction, the cubes in J have at most K distinct sizes.

(iii) Apply the algorithm of Lemma 3.4 to solve the rounded instance and output the packing found by it.

To analyze the algorithm, we use the argument of Fernandez de la Vega and Lueker [12]. Consider two instances J and J' derived from the sorted input list L. As in the algorithm, J is constructed by rounding up all items in each group to the largest cube in the group. On the other hand, J' is constructed by rounding down all items in each group to the smallest cube in the group. Clearly,

$$\text{OPT}(J') \leq \text{OPT}(L) \leq \text{OPT}(J).$$

Now, each cube in group g of J' has size at least that of the cubes in group $g+1$ of J. Thus, possibly placing each of the largest Q cubes from J in a bin by itself, it follows that

$$\text{OPT}(J) \leq \text{OPT}(J') + Q \leq \text{OPT}(L) + Q \leq \lceil (1+\varepsilon)\,\text{OPT}(L) \rceil.$$

The last inequality follows directly from $\text{OPT}(L) \geq n\delta^d$ together with $Q = \lceil n/K \rceil \leq \lceil n\varepsilon\delta^d \rceil \leq \lceil \varepsilon\,\text{OPT}(L) \rceil$. □

3.4. Packing small and large cubes together. In this section we prove Theorem 1.1. Although we already know how to construct almost-optimal packings of only large and only small cubes separately, we cannot directly apply the algorithms in the previous sections to construct a packing in the general case. One extra step will be required to allow the combination of the previously seen algorithms.

For the sake of simplicity in the notation we work in this section only in the two-dimensional case. The extension to higher dimensions is straightforward and is discussed at the end of the section.

The general asymptotic polynomial time approximation scheme for square packing is as follows:

ALGORITHM FOR SQUARE PACKING.

(1) Let I be the input list and $\varepsilon > 0$ fixed. Let $r = \lceil 1/\varepsilon \rceil$. Consider the sequence $\varepsilon, \varepsilon^3, \ldots, \varepsilon^{2^i - 1}, \ldots$ for $i = 1, \ldots, r+1$. Let $M_i = \{j: a_j \in [\varepsilon^{2^{i+1}-1}, \varepsilon^{2^i - 1})\}$, for $i = 1, \ldots, r$.

(2) Take $M := M_i$ for some index $1 \le i \le r$ satisfying $\mathrm{Vol}(M_i) \le \varepsilon \, \mathrm{Vol}(I)$. Define the set of large items as $L = \{j\colon a_j \ge \varepsilon^{2^i - 1}\}$ and the set of small items as $S = \{j\colon a_j < \varepsilon^{2^{i+1}-1}\}$.

(3) Find an almost-optimal packing of L as in Lemma 3.5.

(4) Partition the remaining space in the opened bins into rectangular regions using Lemma 3.2. Use NFDH to pack as many squares in S as possible into the rectangular free space. Let $S' \subset S$ denote the subset of the small items that could not be packed (S' could possibly be empty).

(5) Open new bins and use NFDH to pack $M \cup S'$.

From now on, we refer to this algorithm as Algorithm A. We show that it achieves the desired approximation ratio and running time, as given in Theorem 1.1.

Analysis of the running time. First, it requires time $O(n \log n)$ to sort the cubes according to their sizes. Next, it is easy to see that Steps (1) and (2) take at most $O(n)$ time. By Lemma 3.5, Step (3) requires $O(\mathrm{polylog}(n))$ time. Finally, the running time in Steps (4) and (5) is dominated by the running time of the NFDH shelf heuristic, which is also dominated by the time to sort the list of cubes and, hence, is at most $O(n \log n)$.

Analysis of the algorithm. First observe that since we have $\lceil 1/\varepsilon \rceil$ groups M_i, there exists an i such that M_i satisfies $\mathrm{Vol}(M_i) \le \varepsilon \, \mathrm{Vol}(I)$.

To prove that Algorithm A satisfies the result in Theorem 1.1, we need to distinguish two cases depending on what the set S' was after Step (4). In both cases we will show that the bound in the theorem holds.

(i) *After Step* (4), *S' is empty.* In this case, by Lemma 3.5, the number of bins Algorithm A has opened by Step (4) is bounded by

$$A(L \cup S) = A(L) \le \lceil (1+\varepsilon) \, \mathrm{OPT}(L) \rceil \le \lceil (1+\varepsilon) \, \mathrm{OPT}(I) \rceil.$$

It only remains to account for the bins used to pack the cubes in M. Since the size of each cube in M is at most ε, and these are packed using NFDH, the number of bins required by Algorithm A to pack M is at most

$$\lceil \mathrm{Vol}(M)/(1-2\varepsilon) \rceil \le \lceil \varepsilon \, \mathrm{OPT}(I)/(1-2\varepsilon) \rceil \le \lceil 2\varepsilon \, \mathrm{OPT}(I) \rceil$$

new bins. The last step of the inequality follows by assuming that $\varepsilon < 1/4$ (clearly, if the result holds for $\varepsilon < 1/4$, then it also holds for any $\varepsilon' \ge 1/4$). It follows that the total number of bins used by Algorithm A is no more than

$$\lceil (1+3\varepsilon) \, \mathrm{OPT}(I) \rceil + 1.$$

(ii) *After Step* (4), *S' is nonempty*. In this case we derive a volume argument for the bins opened to pack L. Consider a bin obtained after Step (4) and let $L' \subseteq L$ denote the set of large cubes packed in this bin. Clearly, $|L'| \le (1/\varepsilon^{(2^i-1)})^2 = 1/\varepsilon^{2(2^i-1)}$. From Lemma 3.2, $[0,1]^2 \backslash L'$ can be decomposed into no more than $3/\varepsilon^{(2^i-1)2}$ rectangles; these are filled in Step (4) of the algorithm (by NFDH) with cubes from the set S. Lemma 3.3 tells us that for each rectangle in the partition we will cover everything but a volume of at most $2\varepsilon^{2^{i+1}-1}$ (since each rectangle in the partition has side with length no more than 1). Adding over all rectangles, the total wasted volume can be bounded by $2\varepsilon^{2^{i+1}-1} \times 3/\varepsilon^{(2^i-1)2} \le 6\varepsilon$. This last bound implies that for each bin containing cubes in L, at least a fraction $(1-6\varepsilon)$ of its volume is filled with cubes from $L \cup S$. Moreover, by the argument in the previous case, the cubes in $M \cup (S')$ can be packed such that $(1-2\varepsilon)$ fraction of volume is used, except possibly for the last bin,

$$A(I) \le \left\lceil \frac{\mathrm{Vol}(I)}{1-6\varepsilon} \right\rceil \le (1+12\varepsilon) \, \mathrm{OPT}(I) + 1.$$

The last step follows by assuming without loss of generality that $\varepsilon < 1/12$. This proves Theorem 1.1 for the two-dimensional case.

The higher-dimensional analysis. The asymptotic approximation scheme in the d-dimensional case is almost the same as the one described above. The only difference lies in the sequence we need to consider in Step (1) used to define L, M, and S (the large, medium, and small cubes). In d dimensions, we need to use the decomposition in Lemma 3.1. Hence, Steps (1) and (2) should be replaced by:

(1′) Let I be the input list and $\varepsilon > 0$. Consider the sequence (α_i) satisfying $\alpha_0 = \varepsilon$ and $\alpha_{i+1} = \alpha_i^{2d}\varepsilon$. Such sequence is:

$$\alpha_i = \varepsilon^{((2d)^{(i+1)}-1)/(2d-1)} \quad \text{for } i \ge 0.$$

Let $M_i = \{j\colon a_j \in [\alpha_{i+1}, \alpha_i)\}$.

(2′) Take $M := M_i$ for some i such that $\text{Vol}(M_i) \le \varepsilon\,\text{OPT}(I)$. Define the set of large items as $L = \{j\colon a_j \ge \alpha_i\}$ and the set of small items as $S = \{j\colon a_j < \alpha_{i+1}\}$.

Clearly, the running time of the algorithm is again polynomial in the input size if d and ε are fixed constants.

Finally, the analysis of the algorithm is exactly the same in Case (i) and very similar in Case (ii) above. The only difference is that we will need to use the bounds from Lemma 3.1 and the definition of large and small cubes in the d-dimensional setting.

An exact 2-approximation algorithm. Recently, van Stee [31] gave a (nonasymptotic) 2-approximation algorithm for square packing. Due to the impossibility of distinguishing in polynomial time whether a collection of squares can be packed in a single bin or requires two bins Ferreira et al. [13], this constant is the best that can be obtained. We describe in what follows how a slight variation of our algorithm also achieves such an approximation result, even for hypercubes.

Take ε small (in particular $\varepsilon < 1/12$ so that Theorem 1.1 holds). Let I be the instance and $\text{Vol}(I)$ be the total volume of the instance. The new algorithm should check in Step (1) whether $\text{Vol}(I) > 1$ or $\text{Vol}(I) \le 1$. In the former case the algorithm continues exactly as described before. In the latter case, Step (3) in the algorithm is replaced by: (3′) *Find an optimum packing of* L. Note that, from Lemma 3.4, this new Step (3′) can be executed in polynomial time since L only contains a constant number of items.

Clearly if $\text{Vol}(I) > 1$, then $\text{OPT}(I) \ge 2$ implying that

$$\lceil (1+\varepsilon)\,\text{OPT}(I)\rceil + 1 \le 2\cdot \text{OPT}(I).$$

Thus, in this case, our algorithm is a 2-approximation. On the other hand, if $\text{Vol}(I) \le 1$, the modified algorithm will find a packing using no more than $\text{OPT}(I) + 1 \le 2\cdot\text{OPT}(I)$ bins (specifically, it will use at most $\text{OPT}(I) + \lceil \text{Vol}(M)/(1-2\varepsilon)\rceil \le \text{OPT}(I) + \lceil \varepsilon/(1-2\varepsilon)\rceil = \text{OPT}(I)+1$ bins in case $S' = \varnothing$, and at most $\lceil \text{Vol}(I)/(1-6\varepsilon)\rceil \le \lceil 1/(1-6\varepsilon)\rceil \le 2$ bins in case $S' \ne \varnothing$).

4. Packing rectangles.

By our results in §2 we know that we cannot hope to get an asymptotic approximation scheme for the general problem of packing rectangles. Hence, we now consider a relaxed version of the problem, where we are allowed to pack in slightly larger square bins of size $1+\varepsilon$. We will give an algorithm that uses no additional bins than those required by the optimum possible packing using unit size bins. We begin by describing our algorithm.

4.1. The algorithm.

The following algorithm packs any list I of rectangles into no more than $\text{OPT}(I)$ bins of size $(1+15\varepsilon)\times(1+15\varepsilon)$, where $\text{OPT}(I)$ is the minimum number of unit squares in which I can be packed. Of course, to obtain the exact result in Theorem 1.2 it is enough to reassign $\varepsilon \leftarrow \varepsilon/15$ and apply the algorithm to this newly defined ε.

The algorithm will first decompose the input into "large," "horizontal," "vertical," "medium," and "small," rectangles in such a way that medium rectangles are negligible. So we will pack medium rectangles in thin strips that will be added to the bins along their sides. Large, horizontal, and vertical rectangles will be rounded so that there are few different large rectangles, few different widths of horizontal rectangles, and few different heights of vertical rectangles. Then, a bin of size $(1+O(\varepsilon))\times(1+O(\varepsilon))$ will be seen as a grid consisting of cells of size $\varepsilon\varepsilon' \times \varepsilon\varepsilon'$. We will guess labelings of these cells and attempt to pack all rectangles of each type (large, horizontal, and vertical) into cells having the corresponding label. After that, small rectangles will be packed in the remaining space using the NFDH heuristic. Eventually, a correct labeling together with a desired packing will be found.

Algorithm for Rectangle Packing.

Input. Let I denote the input list consisting of n rectangles to be packed and let $\varepsilon > 0$. Assume that the ith rectangle has width a_i and height b_i, with $0 \le a_i, b_i \le 1$. Denote also by $\text{Vol}(I) = \sum_{i=1}^{n} a_i b_i$, the total area of the input.

Partitioning the input. For $j \in \{1, 2, \ldots, 2/\varepsilon\}$, let M_j denote the set of rectangles in I such that $a_i \in (\varepsilon^{2(j+1)+1}, \varepsilon^{2j+1}]$ or $b_i \in (\varepsilon^{2(j+1)+1}, \varepsilon^{2j+1}]$. Let $j_0 \in \{1, 2, \ldots, 2/\varepsilon\}$ be such that the total area of the rectangles of M_{j_0} is minimum. Let $\varepsilon' = \varepsilon^{2j_0+1}$, and define the partition $I = M_{j_0} \cup L \cup H \cup V \cup S$, where:

- $L = \{i\colon a_i > \varepsilon' \text{ and } b_i > \varepsilon'\}$,
- $S = \{i\colon a_i < \varepsilon'\varepsilon^2 \text{ and } b_i < \varepsilon'\varepsilon^2\}$,
- $H = \{i\colon a_i > \varepsilon' \text{ and } b_i < \varepsilon'\varepsilon^2\}$,
- $V = \{i\colon a_i < \varepsilon'\varepsilon^2 \text{ and } b_i > \varepsilon'\}$.

Rounding the input. For each rectangle in I, we round every side greater than ε' up to the nearest multiple of $\varepsilon'\varepsilon$. Denote by I' the instance containing only the rounded rectangles, i.e., $L \cup H \cup V$.

Let C denote the number of distinct rectangles in L: thus, L contains ℓ_i rectangles of type i, for $1 \le i \le C$. Also, note that the rectangles in H have a constant number of widths and the rectangles in V have a constant number of heights.

Rounding and partitioning the output. We define bin types by decomposing squares of size $(1+2\varepsilon) \times (1+2\varepsilon)$ as follows:

• We consider all possible packings of (large) rectangles of type i, $1 \le i \le C$, into a $(1+2\varepsilon) \times (1+2\varepsilon)$ square, such that the corners of the rectangles are placed at coordinates that are integer multiples of $\varepsilon'\varepsilon$.

• For each possible packing into a $(1+2\varepsilon) \times (1+2\varepsilon)$ square, we decompose the area of the bin that is still uncovered as a union of small square cells of side length $\varepsilon'\varepsilon$ (where each square is positioned at integer multiples of $\varepsilon'\varepsilon$) and consider all possible labelings of each cell with labels H or V.

Each packing of large rectangles together with a labeling of the uncovered area defines a bin type. Let K denote the total number of bin types (which is a constant that we will estimate later).

Main loop. For each $(n_1, \ldots, n_K)$ such that $\sum_j n_j \le n$, we attempt to construct a packing of $I' \cup S \cup M_{j_0}$ using n_j bins of type j, such that the rectangles from L are packed in the spaces reserved for them in the bin type, the cells labeled H are only used for rectangles from $H \cup S$ and the cells labeled V are only used for rectangles from $V \cup S$. Almost all rectangles in I' are packed in Steps (1) to (3), rectangles in S are packed in Step (4), while the rest is packed in Step (5). This is done as follows.

(1) To decide whether the rectangles from L can be placed, we check that for every rectangle type i, $1 \le i \le C$, the number ℓ_i of type i rectangles in I' is less than or equal to the total space available for them:

$$\ell_i \le \sum_{1\le j\le K} n_j \cdot \begin{pmatrix} \text{number of type } i \text{ rectangles} \\ \text{positioned in type } j \text{ bins} \end{pmatrix}.$$

(2) We use the following algorithm for packing the rectangles from H.

(a) For each bin type j, consider the union U of the cells labeled H. Drawing horizontal lines at y-coordinates integer multiples of $\varepsilon'\varepsilon$, we can interpret U as a union of horizontal strips of height $\varepsilon'\varepsilon$ and width multiple of $\varepsilon'\varepsilon$. For each integer multiple ℓ of $\varepsilon'\varepsilon$, let $h_\ell^{(j)}$ denote the sum of the heights of the strips of width ℓ in a type j bin. Let h_ℓ denote the total height of the strips of width ℓ in the packing that we are currently constructing:

$$h_\ell = \sum_{1\le j\le K} n_j h_\ell^{(j)}.$$

(b) We now consider the problem of packing rectangles from H with rounded widths into two-dimensional bins of height h_ℓ and width ℓ; moreover, the number of different width types for rectangles in H is bounded above by a constant. To do this we will solve the following fractional strip-packing problem. Consider all *configurations* $(w_1, w_2, \ldots)$ of widths (including empty configurations) which are multiples of $\varepsilon'\varepsilon$ and sum to at most $1+\varepsilon$. Note that the number of such configurations is constant. Let $A_{i,r}$ denote the number of occurrences of the width $i\varepsilon'\varepsilon$ in configuration r. Let B_i denote the sum of all heights of the rectangles of H whose width equals $i\varepsilon'\varepsilon$. We define one variable $x_r^{(\ell)}$ for each strip width ℓ and for each configuration r whose widths sum to at most ℓ. We find, in polynomial time, a basic feasible solution to the following system of linear constraints, if it exists.

$$\begin{cases} (\forall\, i) & \displaystyle\sum_{\ell,\, r} A_{i,r} x_r^{(\ell)} \ge B_i, \\ (\forall\, \ell) & \displaystyle\sum_{r} x_r^{(\ell)} \le h_\ell, \\ (\forall\, r, \ell) & x_r^{(\ell)} \ge 0. \end{cases}$$

(c) We place rectangles from H in the configurations thus defined, proceeding in a greedy fashion.

(d) We cut back bins of height h_l into strips of height $\varepsilon'\varepsilon$ and place them back into the bins. If a rectangle is cut we throw it away from the packing and pack it later.

(3) Similarly, we pack the rectangles of V into the parts of the bins labeled V.

Let $M_H \subseteq H$ denote the set of rectangles which either did not fit in the fractional packing after (c) or are cut in the process (d). Analogously, we define the set of rectangles $M_V \subseteq V$ which remained unpacked after Step (3).

(4) We pack the rectangles of S into all the $\varepsilon'\varepsilon \times \varepsilon'\varepsilon$ cells which have available space, using the next-fit-decreasing-height (NFDH) algorithm.

(5) We expand each bin by adding 13 thin $1 \times \varepsilon$ horizontal strips and also 13 thin $\varepsilon \times 1$ vertical strips and use them to pack the rectangles from $M_{j_0} \cup M_H \cup M_V$ using an $O(1)$-approximation algorithm such as NFDH in the horizontal strips and next-fit-decreasing-width (NFDW) in the vertical strips.

Output. We output the best packing among all feasible packings of I thus constructed.

4.2. Analysis of the running time. The running time is relatively easy to analyze. As $j_0 \leq 2/\varepsilon$, we have that $\varepsilon' \geq \varepsilon^{4/\varepsilon} = \Omega(1)$. The number C of large rectangle types is at most $(1/\varepsilon'\varepsilon)^2 = O(1)$. A bin type can be defined by labeling each $\varepsilon'\varepsilon \times \varepsilon'\varepsilon$ cell by H, V, or $i \leq C$; thus the number K of bin types is at most $(C+2)^{((1+2\varepsilon)/(\varepsilon'\varepsilon))^2} = O(1)$, and so the number of iterations through the main loop is at most n^K, which is polynomial in n. Note that since $\log(1/\varepsilon') \leq (4/\varepsilon)\log(1/\varepsilon)$, then $((1+2\varepsilon)/(\varepsilon'\varepsilon))^2 \log(C+2) = \tilde{O}(1/(\varepsilon')^2)$. Thus, we estimate the number of iterations as

$$n^K = n^{(C+2)^{((1+2\varepsilon)/(\varepsilon'\varepsilon))^2}} = n^{2^{((1+2\varepsilon)/(\varepsilon'\varepsilon))^2 \log(C+2)}} = n^{2^{\tilde{O}(1/(\varepsilon')^2)}} = n^{2^{2^{\tilde{O}(1/\varepsilon)}}}.$$

The number of strip widths is at most $1/\varepsilon'\varepsilon = O(1)$. The number of configurations is at most $2^{2/(\varepsilon'\varepsilon)} = O(1)$. Multiplying, the number of variables in the linear program is $O(1)$. The number of constraints is at most $2/(\varepsilon'\varepsilon) = O(1)$. The coefficients $A_{i,r}$ are bounded by $O(1)$ and B_i are written on at most $O(\log(n))$ bits; hence, the linear program can be solved efficiently in polylogarithmic time.

The next-fit-decreasing-height and next-fit-decreasing-width algorithms run in time $O(n\log(n))$.

Overall, we conclude that the algorithm thus runs in time:

$$O\Big(n\log(n)\cdot \text{polylog}(n)\cdot n^{2^{2^{\tilde{O}(1/\varepsilon)}}}\Big) = n^{2^{2^{\tilde{O}(1/\varepsilon)}}}.$$

4.3. Analysis of correctness.

Lemma 4.1. *The area of M_{j_0} satisfies*

$$\text{Vol}(M_{j_0}) \leq \varepsilon\,\text{Vol}(I).$$

Proof. Each rectangle of I belongs to at most two sets M_j; thus $\sum_{1\leq j\leq 2/\varepsilon} \text{Vol}(M_j) \leq 2\,\text{Vol}(I)$. The minimum area is less than the average area, which is bounded by $\varepsilon\,\text{Vol}(I)$. □

Lemma 4.2. *Let I' denote the rounded input. If I can be packed into* OPT *unit size bins, then $I' \cup S \cup M_{j_0}$ can be packed into* OPT *bins of size $(1+2\varepsilon)\times(1+2\varepsilon)$, in such a way that any rectangle with $a_i' \geq \varepsilon'$ is positioned at an x-coordinate that is an integer multiple of $\varepsilon'\varepsilon$, and any rectangle with $b_i' \geq \varepsilon'$ is positioned at a y-coordinate that is an integer multiple of $\varepsilon'\varepsilon$.*

Proof. Consider the optimal packing of I. Define a partial order $\prec_H$ on rectangles as the transitive closure of the relation: i is in relation with i' if rectangle i, when translated horizontally to the right by $2\varepsilon'\varepsilon$, intersects i'. Note that any chain in $\prec_H$ contains at most $1/\varepsilon'$ rectangles with $a_i \geq \varepsilon'$.

Consider a linear order $\leq$ that extends $\prec_H$. We take the rectangles (such that $a_i > \varepsilon'$) one by one in that order and deal with i as follows: we extend i horizontally to the right so that its width becomes a_i'; we translate it horizontally to the right so that it is positioned at an integer multiple of $\varepsilon'\varepsilon$; and then, for each $i' \geq i$ in increasing order, if i' intersects some rectangle $i'' \leq i'$, then we translate i' to the right by $2\varepsilon'\varepsilon$. This construction produces a feasible packing because rectangles are shifted every time infeasibility occurs.

By induction we can show that the number of times a rectangle i is shifted is upper bounded by the number of rectangles in a longest chain in the partially ordered set $\{j \prec_H i\colon a_j > \varepsilon'\}$ under the partial order $\prec_H$. Since such chains contain at most $1/\varepsilon'$ rectangles, each rectangle is translated at most $1/\varepsilon'$ times, hence in total by at most $(2\varepsilon'\varepsilon)/\varepsilon' = 2\varepsilon$. Thus, the final packing fits in bins of size $(1+2\varepsilon)\times 1$.

We then proceed similarly for rounding the b_is into b_i'. □

Lemma 4.3. *Consider a packing of I' into bins of size $(1+2\varepsilon)\times(1+2\varepsilon)$, satisfying the condition of Lemma 4.2. Consider any cell $\mathscr{C} = [m\varepsilon'\varepsilon, (m+1)\varepsilon'\varepsilon] \times [p\varepsilon'\varepsilon, (p+1)\varepsilon'\varepsilon]$ in any bin. Then either $H \cap \mathscr{C} = \varnothing$ or $V \cap \mathscr{C} = \varnothing$; i.e., if a rectangle in H intersects cell $\mathscr{C}$, then no rectangle in V can intersect it.*

Proof. Assume, for a contradiction, that there are rectangles $i \in H \cap \mathscr{C} \neq \varnothing$ and $i' \in V \cap \mathscr{C} \neq \varnothing$. Because i has width and starting point integer multiples of $\varepsilon'\varepsilon$, it must be that i spans the whole width of $\mathscr{C}$. Similarly, i' must span the whole height of $\mathscr{C}$. However, in this case i and i' must intersect, which gives a contradiction. □

Lemma 4.4. *The areas of M_H and M_V can be estimated as*:

$$\mathrm{Vol}(M_H) = O(\varepsilon)\,\mathrm{Vol}(I) \qquad \textit{and} \qquad \mathrm{Vol}(M_V) = O(\varepsilon)\,\mathrm{Vol}(I).$$

Proof. Let us only prove the first equation; the second is analogous. Consider the sets M_c and M_d denoting the rectangles that were discarded in Steps 2(c) and 2(d), respectively. Then $M_c \cup M_d = M_H$. We prove that both $\mathrm{Vol}(M_c) = O(\varepsilon)\,\mathrm{Vol}(I)$ and $\mathrm{Vol}(M_d) = O(\varepsilon)\,\mathrm{Vol}(I)$.

To see the first equality, we recall the strip-packing analysis by Kenyon and Rémila [18]. Consider a fractional strip packing (i.e., a solution to the linear program in Step (2)(b)) $x_r^{(\ell)}$ for all r and ℓ. Clearly, such a solution has at most $(2/\varepsilon'\varepsilon)$ nonzero coordinates. Fix ℓ and let $x_1^{(\ell)}, \ldots, x_k^{(\ell)}$ be the nonzero variables corresponding to that ℓ. To construct an integer strip packing we proceed as follows: Let $x_j^{(\ell)} > 0$ be the variable corresponding to the current configuration. This configuration will be used between levels $l_j^\ell = (x_1^{(\ell)} + \varepsilon'\varepsilon^2) + \cdots + (x_{j-1}^{(\ell)} + \varepsilon'\varepsilon^2)$ and $l_{j+1}^\ell = l_j^\ell + x_j^{(\ell)} + \varepsilon'\varepsilon^2$. For each i such that $A_{ij} \neq 0$ we draw A_{ij} columns of width $i\varepsilon'\varepsilon$ going from level l_j^ℓ to level l_{j+1}^ℓ. After this is done for all ℓ and all configurations, we take all columns of width $i\varepsilon'\varepsilon$ and start filling them up with the corresponding rectangles of the same width in a greedy manner; as the maximum height of a rectangle in H is $\varepsilon'\varepsilon$, it is not difficult to see that all rectangles fit. Now, we take all rectangles whose top end belongs to the interval $(l_j^\ell - \varepsilon'\varepsilon^2, l_j^\ell]$, for any ℓ and j, and set them aside (they are put in M_c). The total area of the removed rectangles is clearly no more than

$$2 \cdot (\text{number of constraints}) \cdot (\text{max item height}) = 2 \cdot \frac{2}{\varepsilon'\varepsilon} \cdot \varepsilon'\varepsilon^2 = 4\varepsilon.$$

This proves that a fractional strip packing where strips of width ℓ are of height h_ℓ (the linear program in Step (2)(b)), can be turned into an integer strip packing of almost all rectangles where strips of width ℓ are of height no more than h_ℓ. We can conclude that the total area of rectangles in H that may not fit after Step (2)(c) is bounded by 4ε. In other words, $\mathrm{Vol}(M_c) \leq 4\varepsilon\,\mathrm{Vol}(I)$.

Noting that rectangles in H have height smaller than $\varepsilon'\varepsilon^2$ and the cut strips of Step (2)(d) are of height $\varepsilon'\varepsilon$, the area of the rectangles put aside in Step (2)(d) is no more than a fraction ε of the area of H. Therefore, $\mathrm{Vol}(M_d) = \varepsilon\,\mathrm{Vol}(I)$ and the total area of M_H is no more than a fraction $O(\varepsilon)$ (indeed 5ε) of the total area of the instance. □

Lemma 4.5. *The rectangles in $M_{j_0} \cup M_H \cup M_V$ fit into the thin strips added along the bins in Step* (5).

Proof. Clearly, the whole area of $M_{j_0} \cup M_H \cup M_V$ is no more than $O(\varepsilon)\,\mathrm{Vol}(I)$. Moreover, we can partition $M_{j_0} \cup M_H \cup M_V$ into two sets A and B such that:

- A contains only rectangles with $a_i < \varepsilon' < \varepsilon^2$ and $\mathrm{Vol}(A) \leq \mathrm{Vol}(M_V) + \mathrm{Vol}(M_{j_0}) \leq 6\varepsilon\,\mathrm{Vol}(I)$ (A contains M_V and part of M_{j_0}).
- B contains only rectangles with $b_i < \varepsilon' < \varepsilon^2$ and $\mathrm{Vol}(B) \leq \mathrm{Vol}(M_H) + \mathrm{Vol}(M_{j_0}) \leq 6\varepsilon\,\mathrm{Vol}(I)$ (B contains M_H and part of M_{j_0}).

As all rectangles in A have width smaller than ε^2, for small enough ε, the NFDW heuristic packs A into the $13 \cdot \mathrm{OPT}(I)$ added strips of size $\varepsilon \times 1$. On the other hand, NFDH does the work for the rectangles in B. Note that the number of thin strips that we need to add to the $(1+2\varepsilon) \times (1+2\varepsilon)$ bins depends on the algorithm used to pack rectangles in $M_{j_0} \cup M_H \cup M_V$. □

Lemma 4.6. *Consider the two-dimensional NFDH heuristic applied to rectangles in S. When NFDH cannot place any other rectangle in a bin of size $a \times b$, then the total unused area in that bin is no more than*

$$\varepsilon'\varepsilon^2 \cdot (a+b).$$

Proof. The result is very similar to Lemma 3.3 and to a result in Coffman et al. [8]. □

We are now ready to give the overall analysis of the algorithm.

Proof of Theorem 1.2. Consider an input list I and let I' be the rounded input. From Lemma 4.2 we know that there is a packing of $I' \cup S \cup M_{j_0}$ into no more than $\mathrm{OPT}(I)$ bins of size $(1+2\varepsilon) \times (1+2\varepsilon)$. Consider then the optimal packing in such bins for $I' \cup S \cup M_{j_0}$ satisfying the conditions of Lemma 4.2. By Lemma 4.3 we know that in such packing all cells intersect either a rectangle in L, H, or V but not two of them. In other words in a packing of $I' \cup S \cup M_{j_0}$ satisfying the conditions of Lemma 4.2 each cell is labeled either V, H, or i for $i = 1, \ldots, C$ (where C was the number of distinct large rectangles); or will have no label at all.

Clearly, our algorithm will eventually guess a labeling of the cells that coincides with the labeling in the optimal packing. At that point the algorithm will find a feasible packing of all rectangles in L and almost all rectangles in H and V. By Lemmas 4.4 and 4.5 the unpacked rectangles M_{j_0}, M_H, and M_V are of small area and they can be packed in the extra space added in Step (5) of the algorithm.

It only remains to see that the small rectangles S will be successfully packed by NFDH. We prove that is not possible that in Step (4) a new bin is opened if already $\mathrm{OPT}(I)$ bins of size $(1+2\varepsilon)\times(1+2\varepsilon)$ have been used. We do this by a volume argument. Suppose by contradiction that such a new bin is opened in Step (4). At this step we distinguish four types of $\varepsilon'\varepsilon\times\varepsilon'\varepsilon$ cells: the ones completely filled with a rectangle in L, the ones filled with only rectangles in S, the ones filled only with rectangles in H or V, and the ones partly filled with rectangles in H or V and partly with rectangles in S. By the arguments in Lemmas 4.4 and 4.6 all cells are almost filled. Namely, a fraction $(1-2\varepsilon)$ of their area is filled. Overall this implies that a fraction $(1-2\varepsilon)$ of the first $\mathrm{OPT}(I)$ bins is filled, and then the total area that has been filled in the first $\mathrm{OPT}(I)$ bins of size $(1+2\varepsilon)\times(1+2\varepsilon)$ is at least

$$(1-2\varepsilon)\,\mathrm{OPT}(I)(1+2\varepsilon)^2>(1-2\varepsilon)(1+4\varepsilon)\,\mathrm{OPT}(I)>\mathrm{OPT}(I).$$

The inequality follows because $\varepsilon<1/4$.

5. A related rectangle packing problem. The algorithm presented in §4 is closely related to the optimization version of the following question posed by Moser [25]: *Determine the smallest number x such that any system of squares with total area* 1 *may be packed in parallel into a rectangle of area x.*

Bounds for this problem have been obtained by Kleitman and Krieger [19] and by Novotny [27]. Indeed, the first authors proved that any such system of squares can be packed in a rectangle of size $\sqrt{2}\times\sqrt{4/3}$, while the latter author proved that they can be packed in a rectangle of area no more than 1.53.

Although we do not obtain absolute bounds (i.e., independent of the input) for this problem, our algorithm can easily be adapted to obtain approximate solutions for a related optimization problem: the minimum rectangle placement problem. The problem can be formally stated as follows: *Given a system of n rectangles with total area* 1, *determine the smallest number x such that all n rectangles may be packed in parallel into a rectangle R of area x.*

We now proceed to prove Theorem 1.3; i.e., we present a PTAS for the minimum rectangle placement problem.

Proof of Theorem 1.3. Let I denote the instance. For each rectangle, let $a_i\in(0,1]$ denote its width and $b_i\in(0,1]$ its height; the total area of the instance, denoted by $\mathrm{Vol}(I)$, equals 1. To solve this problem note first that the techniques in §4 can be easily adapted to solve the underlying approximate decision problem: Given n rectangles with sides $a_i, b_i\in(0,1]$, given a rectangle R of size $a\times b$, and given $\delta>0$, determine whether all n rectangles can be packed in a rectangle of size $[(1+\delta)a+\delta]\times[(1+\delta)b+\delta]$ or they cannot be packed in R.

Consider a given $\varepsilon>0$, and let $\alpha=\min\{\max_{i=1,\ldots,n}a_i, \max_{i=1,\ldots,n}b_i\}$. We distinguish two cases:

(i) $\alpha>\varepsilon$. In this case, we ensure that the sides of the minimum area rectangle in which all n items can be packed are at least ε. It is then enough to solve the decision problem, taking $\delta=\varepsilon^2$, for all rectangles R such that:

—Their lower-left corner is $(0,0)$ and their upper-right corner is on or below the curve $xy=4\,\mathrm{Vol}(I)=4$ (because the NFDH shelf packing heuristic can always pack I in a square of area $4\,\mathrm{Vol}(I)$).

—The coordinates of their upper-right corner are at least ε.

—Their side lengths are integer multiples of ε^2.

The number of such rectangles is clearly no more than the total number of points with coordinates multiples of ε^2 contained in the square $[\varepsilon,4/\varepsilon]\times[\varepsilon,4/\varepsilon]$, which is less than $4/\varepsilon^3\cdot 4/\varepsilon^3=16/\varepsilon^6$ (a slightly more careful analysis shows that the number of rectangles with the above properties is less than $\int_{\varepsilon}^{1/\varepsilon}(4/x)\,dx\cdot(1/\varepsilon^4)=(4/\varepsilon^4)\ln(1/\varepsilon^2)\le(8/\varepsilon^4)\ln(1/\varepsilon)$). Since $16/\varepsilon^6$ is constant, it suffices to solve the decision problem for a constant number of rectangles: we can then choose the one with minimum area. Since in this case a and b are guaranteed to be large, $[(1+\delta)a+\delta]\le(1+2\varepsilon)a$ and $[(1+\delta)b+\delta]\le(1+2\varepsilon)b$. Therefore, the best rectangle is guaranteed to have an area within a factor of $(1+O(\varepsilon))$ that of the optimal one.

(ii) $\alpha\le\varepsilon$. In this case either all rectangles are flat or all are narrow. Without loss of generality, assume they are all narrow. We will pack all rectangles using the FFDH heuristic for strip packing in a strip of width $\sqrt{\varepsilon}$. A slight adaptation of the following theorem—essentially shown by Coffman et al. [8]—proves that all items can be placed in a rectangle of area roughly equal to $\mathrm{Vol}(I)$.

Theorem 5.1. *If all rectangles in I have width less than or equal to δ, then the height $FFDH(I)$ achieved by the FFDH heuristic* (*on a strip of width* 1) *satisfies*:

$$FFDH(I)\le(1+2\delta)\,\mathrm{Vol}(I)+1.$$

In our case the width of the strip is $\sqrt{\varepsilon}$, and then FFDH finds a packing in a rectangle of size:

$$\sqrt{\varepsilon}\times\left[(1+2\sqrt{\varepsilon})\frac{\mathrm{Vol}(I)}{\sqrt{\varepsilon}}+1\right].$$

The area of such a rectangle is then

$$(1+2\sqrt{\varepsilon})\,\mathrm{Vol}(I)+\sqrt{\varepsilon}=(1+3\sqrt{\varepsilon}).$$

Again the result follows in this case.

6. Open problems.

Even though the result on cube packing in this paper holds in general (fixed) dimension, this is not the case for the results on rectangle packing in §§4 and 5. An interesting open question is, therefore, to generalize Theorems 1.2 and 1.3 to higher-dimensional rectangle packing. We have not been able to do so because the interaction of, say, 3-dimensional rectangles, which are long in one direction but short in the other two, seems much more complicated than in the two-dimensional case. Another interesting open problem is to find explicit asymptotic inapproximability gaps for the two-dimensional rectangle packing problem. For the d-dimensional rectangle packing problem, the best-known approximation algorithms, due to Csirik and van Vliet [9], have an asymptotic approximation ratio of $(1.691\ldots)^d$ (which is exponential in d). However, the only lower bound on the achievable approximation ratio follows from our hardness result in §2. It would be very interesting to close this gap, in particular, to find out whether the lower bound construction can be extended to higher dimensions to yield a bound exponential in d, or whether an algorithm for d-dimensional rectangle packing with an approximation ratio subexponential in d can be obtained.

Acknowledgments. The authors thank two anonymous referees for many valuable comments and suggestions regarding the presentation of the paper. The authors also thank David Johnson for his detailed and informative comments in his review of the original SODA submission, pointing out several minor errors in the proof of Theorem 1.2. The research of the third author was partially completed while she was at Ecole Polytechnique. The research of the second author was partially completed while he was visiting Ecole Polytechnique, supported by Institut Universitaire de France and by an AMAFC Franco-Chilean cooperation grant.

References

[1] Baker, B. S., D. J. Brown, H. P. Kasteff. 1981. A 5/4 algorithm for two-dimensional packing. *J. Algorithms* **2** 348–368.
[2] Berman, P., M. Karpinski. 2003. Improved approximation lower bounds on small occurrence optimization.
[3] Caprara, A. 2002. Packing two-dimensional bins in harmony. *Proc. 43rd IEEE Sympos. Foundations Comput. Sci.* (*FOCS*). IEEE, Vancouver, Canada, 490–499.
[4] Caprara, A., A. Lodi, M. Monaci. 2005. Fast approximation schemes for two-stage, two-dimensional bin packing. *Math. Oper. Res.* **30** 136–156.
[5] Chekuri, C., S. Khanna. 1999. On multi-dimensional packing problems. *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). ACM-SIAM, Baltimore, MD, 185–194.
[6] Chung, F. R. K., M. R. Garey, D. S. Johnson. 1982. On packing two-dimensional bins. *SIAM J. Algebraic Discrete Methods* **3** 66–76.
[7] Coffman, E. G., M. R. Garey, D. S. Johnson. 1996. Approximation algorithms for bin packing: A survey. D. Hochbaum, ed. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing, Boston, MA, 46–93.
[8] Coffman, E. G., M. R. Garey, D. S. Johnson, R. E. Tarjan. 1980. Performance bounds for level-oriented two-dimensional packing algorithms. *SIAM J. Comput.* **9** 808–826.
[9] Csirik, J., A. van Vliet. 1993. An on-line algorithm for multidimensional bin packing. *Oper. Res. Lett.* **13** 149–158.
[10] Csirik, J., G. Woeginger. 1998. On-line packing and covering problems. A. Fiat, G. Woeginger, eds. *Online Algorithms: The State of the Art*, *Lecture Notes in Computer Science*, Vol. 1442. Springer, 147–177.
[11] Epstein, L., R. van Stee. 2004. Optimal online bounded space multidimensional packing. *Proc. 15th ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). ACM-SIAM, New Orleans, LA, 207–216.
[12] Fernandez de la Vega, W., G. Lueker. 1981. Bin packing can be solved within $1+\varepsilon$ in linear time. *Combinatorica* **1** 349–355.
[13] Ferreira, C. E., F. K. Miyazawa, Y. Wakabayashi. 1999. Packing squares into squares. *Pesquisa Operacional* **19** 223–237.
[14] Hazan, E., S. Safra, O. Schwartz. 2003. On the complexity of approximating k-dimensional matching. *Proc. 6th Internat. Workshop on Approximation Algorithms for Combin. Optim. Problems and of the 7th Internat. Workshop on Randomization and Comput.* (*RANDOM-APPROX*), *Lecture Notes in Computer Science*, Vol. 2764. Springer, Berlin, Germany, 83–97.
[15] Jansen, K., G. Zhang. 2004. On rectangle packing: Maximizing benefits. *Proc. 15th ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). ACM-SIAM, New Orleans, LA, 197–206.
[16] Kann, V. 1991. Maximum bounded 3-dimensional matching is MAX SNP-complete. *Inform. Processing Lett.* **37** 27–35.
[17] Karmarkar, N., R. M. Karp. 1982. An efficient approximation scheme for the one-dimensional bin-packing problem. *Proc. 23rd IEEE Sympos. Foundations Comput. Sci.* (*FOCS*). IEEE, New York, 312–320.
[18] Kenyon, C., E. Rémila. 2000. A near-optimal solution to a two-dimensional cutting stock problem. *Math. Oper. Res.* **25** 645–656.

[19] Kleitman, D. J., M. Krieger. 1975. An optimal bound for two-dimensional packing. *Proc. 16th IEEE Sympos. Foundations Comput. Sci.* (*FOCS*). IEEE, Berkeley, CA, 163–168.
[20] Kohayakawa, Y., F. K. Miyazawa, P. Raghavan, Y. Wakabayashi. 2004. Multidimensional cube packing. *Algorithmica* **40** 173–187.
[21] Korf, R. 2003. Optimal rectangle packing: Initial results. *Proc. 13th Internat. Conf. on Automated Planning and Scheduling* (*ICAPS*). AAAI, Trento, Italy, 287–295.
[22] Lenstra, H. W. 1983. Integer programming with a fixed number of variables. *Math. Oper. Res.* **8** 538–548.
[23] Leung, J. Y. T., T. W. Tam, C. S. Wong, G. H. Young, F. Y. L. Chin. 1990. Packing squares into a square. *J. Parallel Distributed Comput.* **10** 271–275.
[24] Meir, A., L. Moser. 1968. On packing of squares and cubes. *J. Combin. Theory. Ser. A* **5** 126–134.
[25] Moser, L. 1965. Poorly formulated unsolved problems of combinatorial geometry. Mimeographed.
[26] Murata, H., K. Fujiyoshi, S. Nakatake, Y. Kajitani. 1996. VLSI module placement based on rectangle-packing by the sequence-pair. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **15** 1518–1524.
[27] Novotny, P. 1996. On packing of squares into a rectangle. *Archivum Mathematicum* **32** 75–83.
[28] Petrank, E. 1994. The hardness of approximation: Gap location. *Comput. Complexity* **4** 133–157.
[29] Seiden, S., R. van Stee. 2003. New bounds for multi-dimensional packing. *Algorithmica* **36** 261–293.
[30] Sevastianov, S., G. Woeginger. 1998. Makespan minimization in open shops: A polynomial time approximation scheme. *Math. Programming Ser. B* **82** 191–198.
[31] van Stee, R. 2004. An approximation algorithm for square packing. *Oper. Res. Lett.* **32** 535–539.
[32] Vazirani, V. 2001. *Approximation Algorithms*. Springer, Berlin, Germany.
[33] Woeginger, G. 1997. There is no asymptotic PTAS for two-dimensional vector packing. *Inform. Processing Lett.* **64** 293–297.