Best Fit Bin Packing with Random Order Revisited

Susanne Albers
Technische Universität München, Germany
albers@in.tum.de

Arindam Khan
Indian Institute of Science, Bangalore, India
arindamkhan@iisc.ac.in

Leon Ladewig
Technische Universität München, Germany
ladewig@in.tum.de

Abstract

Best Fit is a well known online algorithm for the bin packing problem, where a collection of one-dimensional items has to be packed into a minimum number of unit-sized bins. In a seminal work, Kenyon [SODA 1996] introduced the (asymptotic) *random order ratio* as an alternative performance measure for online algorithms. Here, an adversary specifies the items, but the order of arrival is drawn uniformly at random. Kenyon's result establishes lower and upper bounds of 1.08 and 1.5, respectively, for the random order ratio of Best Fit. Although this type of analysis model became increasingly popular in the field of online algorithms, no progress has been made for the Best Fit algorithm after the result of Kenyon.

We study the random order ratio of Best Fit and tighten the long-standing gap by establishing an improved lower bound of 1.10. For the case where all items are larger than 1/3, we show that the random order ratio converges quickly to 1.25. It is the existence of such large items that crucially determines the performance of Best Fit in the general case. Moreover, this case is closely related to the classical maximum-cardinality matching problem in the fully online model. As a side product, we show that Best Fit satisfies a monotonicity property on such instances, unlike in the general case.

In addition, we initiate the study of the *absolute* random order ratio for this problem. In contrast to asymptotic ratios, absolute ratios must hold even for instances that can be packed into a small number of bins. We show that the absolute random order ratio of Best Fit is at least 1.3. For the case where all items are larger than 1/3, we derive upper and lower bounds of 21/16 and 1.2, respectively.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases Online bin packing, random arrival order, probabilistic analysis

Digital Object Identifier 10.4230/LIPIcs.MFCS.2020.7

Funding Work supported by the European Research Council, Grant Agreement No. 691672.

1 Introduction

One of the fundamental problems in combinatorial optimization is *bin packing*. Given a list $I = (x_1, \dots, x_n)$ of n items with sizes from $(0, 1]$ and an infinite number of unit-sized bins, the goal is to pack all items into the minimum number of bins. Formally, a *packing* is an assignment of items to bins such that for any bin, the sum of assigned items is at most 1. While an offline algorithm has complete information about the items in advance, in the online variant, items are revealed one by one. Therefore, an online algorithm must pack x_i without knowing future items $x_{i+1}, \dots, x_n$ and without modifying the packing of previous items.

45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020).
Editors: Javier Esparza and Daniel Kráľ; Article No. 7; pp. 7:1–7:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

As the problem is strongly NP-complete [15], research mainly focuses on efficient approximation algorithms. The offline problem is well understood and admits even approximation schemes [8, 18, 26]. The online variant is still a very active field in the community [6], as the asymptotic approximation ratio of the best online algorithm is still unknown [2, 3]. As one of the first algorithms for the problem, Garey et al. proposed the algorithms Best Fit and First Fit [14]. Johnson published the Next Fit algorithm briefly afterwards [22]. All of these algorithms work in the online setting and attract by their simplicity: Suppose that x_i is the current item to pack. The algorithms work as follows:

Best Fit (BF) Pack x_i into the fullest bin possible, open a new bin if necessary.

First Fit (FF) Maintain a list of bins ordered by the time at which they were opened. Pack x_i into the first possible bin in this list, open a new bin if necessary.

Next Fit (NF) Pack x_i into the bin opened most recently if possible; open a new bin if necessary.

Another important branch of online algorithms is based on the HARMONIC algorithm [29]. This approach has been massively tuned and generalized in a sequence of papers [2, 35, 36].

To measure the performance of an algorithm, different metrics exist. For an algorithm $\mathcal{A}$, let $\mathcal{A}(I)$ and $\mathrm{OPT}(I)$ denote the number of bins used by $\mathcal{A}$ and an optimal offline algorithm, respectively, to pack the items in I. Let $\mathcal{I}$ denote the set of all item lists. The most common metric for bin packing algorithms is the *asymptotic (approximation) ratio* defined as

$$R^{\infty}_{\mathcal{A}} = \limsup_{k\to\infty} \sup_{I\in\mathcal{I}} \{\mathcal{A}(I)/\operatorname{OPT}(I) \mid \operatorname{OPT}(I) = k\} .$$

Note that $R^{\infty}_{\mathcal{A}}$ focuses on instances where $\mathrm{OPT}(I)$ is large. This avoids anomalies typically occurring on lists that can be packed optimally into few bins. However, many bin packing algorithms are also studied in terms of the stronger *absolute (approximation) ratio*

$$R_{\mathcal{A}} = \sup_{I\in\mathcal{I}} \{\mathcal{A}(I)/\operatorname{OPT}(I)\} .$$

Here, the approximation ratio $R_{\mathcal{A}}$ must hold for each possible input. An online algorithm with (absolute or asymptotic) ratio α is also called *α-competitive*.

Table 1 shows the asymptotic and absolute approximation ratios of the three heuristics Best Fit, First Fit, and Next Fit. Interestingly, for these algorithms both metrics coincide. While the asymptotic ratios of Best Fit and Next Fit were established already in early work [23], the absolute ratios have been settled rather recently [9, 10].

Note that the above performance measures are clearly worst-case orientated. An adversary can choose items and present them in an order that forces the algorithm into its worst possible behavior. In the case of Best Fit, hardness examples are typically based on lists where small items occur before large items [14]. In contrast, it is known that Best Fit performs significantly better if items appear in non-increasing order [23]. For real-world instances, it seems overly pessimistic to assume adversarial order of input. Moreover, sometimes worst-case ratios hide interesting properties of algorithms that occur in average cases. This led to the development of alternative measures.

A natural approach that goes beyond worst-case was introduced by Kenyon [28] in 1996. In the model of random order arrivals, the adversary can still specify the items, but the arrival order is permuted randomly. The performance measure described in [28] is based on the asymptotic ratio, but can be applied to absolute ratios likewise. In the resulting performance metrics, an algorithm must satisfy its performance guarantee in expectation

Table 1 Approximation ratios in different metrics of common bin packing heuristics. In R_{NF}, the symbol γ refers to the total size of items in the list.

Algorithm $\mathcal{A}$	Abs. ratio $R_{\mathcal{A}}$	Asym. ratio $R_{\mathcal{A}}^{\infty}$	Asym. random order ratio $RR_{\mathcal{A}}^{\infty}$
Best Fit	1.7 [10]	1.7 [23]	$1.08 \leq RR_{BF}^{\infty} \leq 1.5$ [28]
First Fit	1.7 [9]	1.7 [23]	–
Next Fit	$2 - 1/\lceil\gamma\rceil$ [4]	2 [22]	2 [24]

over all permutations. We define

$$RR_{\mathcal{A}}^{\infty} = \limsup_{k\to\infty} \sup_{I\in\mathcal{I}} \{\mathrm{E}[\mathcal{A}(I^{\sigma})]/\operatorname{OPT}(I) \mid \operatorname{OPT}(I) = k\} \quad \text{and}$$
$$RR_{\mathcal{A}} = \sup_{I\in\mathcal{I}} \{\mathrm{E}[\mathcal{A}(I^{\sigma})]/\operatorname{OPT}(I)\}$$

as the *asymptotic random order ratio* and the *absolute random order ratio* of algorithm $\mathcal{A}$, respectively. Here, σ is drawn uniformly at random from $\mathcal{S}_n$, the set of permutations of n elements, and $I^{\sigma} = (x_{\sigma(1)}, \ldots, x_{\sigma(n)})$ is the permuted list.

1.1 Related work

The following literature review only covers results that are most relevant to our work. We refer the reader to the article [7] by Coffman et al. for an extensive survey on (online) bin packing. For further problems studied in the random order model, see [17].

Bin packing. Kenyon introduced the notion of asymptotic random order ratio $RR_{\mathcal{A}}^{\infty}$ for online bin packing algorithms in [28]. For the Best Fit algorithm, Kenyon proves an upper bound of 1.5 on $RR_{\mathrm{BF}}^{\infty}$, demonstrating that random order significantly improves upon $R_{\mathrm{BF}}^{\infty} = 1.7$. However, it is conjectured in [7,28] that the actual random order ratio is close to 1.15. The proof of the upper bound crucially relies on the following scaling property: With high probability, the first t items of a random permutation can be packed optimally into $\frac{t}{n}\operatorname{OPT}(I) + o(n)$ bins. On the other side, Kenyon proves that $RR_{\mathrm{BF}}^{\infty} \geq 1.08$. This lower bound is obtained from the weaker i.i.d-model, where item sizes are drawn independently and identically distributed according to a fixed probability distribution.

Coffman et al. [24] analyzed next-fit in the random order model and showed that $RR_{\mathrm{NF}}^{\infty} = 2$, matching the asymptotic approximation ratio $RR_{\mathrm{NF}}^{\infty} = 2$ (see Table 1). Fischer and Röglin [12] obtained analogous results for worst-fit [22] and smart next-fit [34]. Therefore, all three algorithms fail to perform better in the random order model than in the adversarial model.

A natural property of bin packing algorithms is monotonicity, which holds if an algorithm never uses fewer bins to pack I' than for I, where I' is obtained from I by increasing item sizes. Murgolo [33] showed that next-fit is monotone, while Best Fit and First Fit are not monotone in general. The concept of monotonicity also arises in related optimization problems, such as scheduling [16] and bin covering [12].

Bin covering. The dual problem of bin packing is bin covering, where the goal is to cover as many bins as possible. A bin is covered if it receives items of total size at least 1. Here, a well-studied and natural algorithm is Dual Next Fit (DNF). In the adversarial setting, DNF has asymptotic ratio $R_{\mathrm{DNF}}^{\infty} = 1/2$ which is best possible for any online algorithm [5]. Under random arrival order, Christ et al. [5] showed that $RR_{\mathrm{DNF}}^{\infty} \leq 4/5$. This upper bound was improved later by Fischer and Röglin [11] to $RR^{\infty}_{\mathrm{DNF}} \leq 2/3$. The same group of authors further showed that $RR^{\infty}_{\mathrm{DNF}} \geq 0.501$, i.e., DNF performs strictly better under random order than in the adversarial setting [12].

Matching. Online matching can be seen as the key problem in the field of online algorithms [32]. Inspired by the seminal work of Karp, Vazirani, and Vazirani [27], who introduced the online bipartite matching problem with one-sided arrivals, the problem has been studied in many generalizations. Extensions include fully online models [13, 19, 20], vertex-weighted versions [1, 21] and, most relevant to our work, random arrival order [21, 31].

1.2 Our results

While several natural algorithms fail to perform better in the random order model, Best Fit emerges as a strong candidate in this model. The existing gap between 1.08 and 1.5 clearly leaves room for improvement; closing (or even narrowing) this gap has been reported as challenging and interesting open problem in several papers [5, 17, 24].

To the best of our knowledge, our work provides the first new results on the problem since the seminal work by Kenyon. Below we describe our results in detail. In the following theorems, the expectation is over the permutation σ drawn uniformly at random.

Case of 1/3-large items

If all items are strictly larger than 1/3, the objective is to maximize the number of bins containing two items. This problem is closely related to finding a maximum-cardinality matching in a vertex-weighted graph; our setting corresponds with the fully online model studied in [1] under random order arrival. Also in the analysis from [28], this special case arises. There, it is sufficient to argue that $\mathrm{BF}(I) \leq \frac{3}{2}\,\mathrm{OPT}(I) + 1$ under adversarial order. We show that Best Fit performs significantly better under random arrival order:

► **Theorem 1.1.** *For any list I of items larger than $1/3$, we have $\mathrm{E}[\mathrm{BF}(I^{\sigma})] \leq \frac{5}{4}\,\mathrm{OPT}(I) + \frac{1}{4}$.*

The proof of Theorem 1.1 is developed in Section 3 and based on several pillars. First, we show that Best Fit is monotone in this case (Proposition 3.2), unlike in the general case [33]. This property can be used to restrict the analysis to instances with well-structured optimal packing. The main technical ingredient is introduced in Section 3.3 with Lemma 3.5 as the key lemma. Here, we show that Best Fit maintains some parts of the optimal packing, depending on certain structures of the input sequence. We identify these structures and show that they occur with constant probability for a random permutation. It seems likely that this property can be used in a similar form to improve the bound $RR^{\infty}_{\mathrm{BF}} \leq 1.5$ for the general case: Under adversarial order, much hardness comes from relatively large items of size more than 1/3; in fact, if all items have size at most 1/3, an easy argument shows 4/3-competitiveness even for adversarial arrival order [23].

Moreover, it is natural to ask for the performance in terms of absolute random order ratio. It is a surprising and rather recent result that for Best Fit, absolute and asymptotic ratios coincide. The result of [28] has vast additive terms and it seems that new techniques are required for insights into the absolute random order ratio. In Section 3.4, we show an upper bound of 21/16 for 1/3-large items, which is complemented by a lower bound of 6/5.

► **Proposition 1.2.** *For any list I of items larger than $1/3$, we have $\mathrm{E}[\mathrm{BF}(I^{\sigma})] \leq \frac{21}{16}\,\mathrm{OPT}(I)$.*

► **Proposition 1.3.** *There is a list I of items larger than $1/3$ with $\mathrm{E}[\mathrm{BF}(I^{\sigma})] > \frac{6}{5}\,\mathrm{OPT}(I)$.*

A proof sketch of Proposition 1.3 is presented in Section 4.2.

Lower bounds

We also make progress on the hardness side, which is presented in Section 4. First, we show that the asymptotic random order ratio is larger than 1.10, improving the previous lower bound of 1.08 from [28].

▶ **Theorem 1.4.** *The asymptotic random order ratio of Best Fit is* $RR^{\infty}_{\mathrm{BF}} > 1.10$*.*

As it is typically challenging to obtain lower bounds in the random order model, we exploit the connection to the i.i.d-model. Here, items are drawn independently and identically distributed according to a fixed probability distribution. By defining an appropriate distribution, the problem can be analyzed using Markov chain techniques. Moreover, we present the first lower bound on the absolute random order ratio:

▶ **Theorem 1.5.** *The absolute random order ratio of Best Fit is* $RR_{\mathrm{BF}} \geq 1.30$*.*

Interestingly, our lower bound on the absolute random order ratio is notably larger than in the asymptotic case (see [28] and Theorem 1.4). This suggests either

- a significant discrepancy between RR_{BF} and $RR^{\infty}_{\mathrm{BF}}$, which is in contrast to the adversarial setting ($R_{\mathrm{BF}} = R^{\infty}_{\mathrm{BF}}$, see Table 1), or
- a disproof of the conjecture $RR^{\infty}_{\mathrm{BF}} \approx 1.15$ mentioned in [7, 28].

2 Notation

We consider a list $I = (x_1, \ldots, x_n)$ of n items throughout the paper. Due to the online setting, I is revealed in *rounds* $1, \ldots, n$. In round t, item x_t arrives and in total, the prefix list $I(t) := (x_1, \ldots, x_t)$ is revealed to the algorithm. The items in $I(t)$ are called the *visible* items of round t. We use the symbol x_t for the item itself and its size $x_t \in (0, 1]$ interchangeably. An item x_t is called *large* (L) if $x_t > 1/2$, *medium* (M) if $x_t \in \left(\frac{1}{3}, \frac{1}{2}\right]$, and *small* (S) if $x_t \leq 1/3$. We also say that x_t is α-large if $x_t > \alpha$.

Bins contain items and therefore can be represented as sets. As a bin usually can receive further items in later rounds, the following terms refer always to a fixed round. We define the *load* of a bin $\mathcal{B}$ as $\sum_{x_i \in \mathcal{B}} x_i$. Sometimes, we classify bins by their internal structure. We say $\mathcal{B}$ is of *configuration LM* (or $\mathcal{B}$ is an *LM-bin*) if it contains one large and one medium item. The configurations L, MM, etc. are defined analogously. Moreover, we call $\mathcal{B}$ a k-bin if it contains exactly k items. If a bin cannot receive further items in the future, it is called *closed*; otherwise, it is called *open*.

The number of bins which Best Fit uses to pack a list I is denoted by $\mathrm{BF}(I)$. We slightly abuse the notation and refer to the corresponding packing by $\mathrm{BF}(I)$ as well whenever the exact meaning is clear from the context. Similarly, we denote by $\mathrm{OPT}(I)$ the number of bins and the corresponding packing of an optimal offline solution.

Finally, for any natural number n we define $[n] := \{1, \ldots, n\}$. Let $\mathcal{S}_n$ be the set of permutations in $[n]$. If not stated otherwise, σ refers to a permutation drawn uniformly at random from $\mathcal{S}_n$.

3 Upper bound for 1/3-large items

In this section, we consider the case where I contains no small items, i.e., where all items are 1/3-large. In Sections 3.1 to 3.3 we develop the technical foundations. The final proofs of Theorem 1.1 and Proposition 1.2 are presented in Section 3.4.

3.1 Monotonicity

We first define the notion of monotone algorithms.

▶ **Definition 3.1.** *We call an algorithm monotone if increasing the size of one or more items cannot decrease the number of bins used by the algorithm.*

One might suspect that any reasonable algorithm is monotone. While this property holds for an optimal offline algorithm and some online algorithms as ext-fit [25], Best Fit is not monotone in general [33]. As a counterexample, consider the lists

$$I = (0.36,\ 0.65,\ \mathbf{0.34},\ 0.38,\ 0.28,\ 0.35,\ 0.62) \ \text{ and}$$
$$I' = (0.36,\ 0.65,\ \mathbf{0.36},\ 0.38,\ 0.28,\ 0.35,\ 0.62)\,.$$

Before arrival of the fifth item, $\mathrm{BF}(I(4))$ uses two bins $\{0.36, 0.38\}$ and $\{0.65, 0.34\}$, while $\mathrm{BF}(I'(4))$ uses three bins $\{0.36, 0.36\}$, $\{0.65\}$, and $\{0.38\}$. Now, the last three items fill up the existing bins in $\mathrm{BF}(I'(4))$ exactly. In contrast, these items open two further bins in the packing of $\mathrm{BF}(I(4))$. Therefore, $\mathrm{BF}(I) = 4 > 3 = \mathrm{BF}(I')$.

However, we can show that Best Fit is monotone for the case of 1/3-large items. Interestingly, 1/3 seems to be the threshold for the monotonicity of Best Fit: As shown in the counterexample from the beginning of this section, it is sufficient to have one item $x \in \left(\frac{1}{4}, \frac{1}{3}\right]$ to force Best Fit into anomalous behavior.

▶ **Proposition 3.2.** *Given a list I of items larger than $1/3$ and a list I' obtained from I by increasing the sizes of one or more items, we have $\mathrm{BF}(I) \le \mathrm{BF}(I')$.*

Sketch of proof. For simplicity, first assume that both lists differ only in the i-th element. All bins in any packing of I or I' contain at most two items. We call two 1-bins of $\mathrm{BF}(I)$ and $\mathrm{BF}(I')$ *pairwise-identical* if they contain items of the same size. Moreover, we call any two 2-bins of $\mathrm{BF}(I)$ and $\mathrm{BF}(I')$ *pairwise-closed*, as neither of the two bins can receive a further item. For ease of notation, let $I_t = I(t)$ and $I'_t = I'(t)$. We can show that at any time t, the packings $\mathrm{BF}(I_t)$ and $\mathrm{BF}(I'_t)$ are related in one of three ways:

(1) All bins are pairwise-identical or pairwise-closed.

(2) All bins are pairwise-identical or pairwise-closed, except for two 1-bins $B = \{b\}$ and $B' = \{b'\}$ in $\mathrm{BF}(I_t)$ and $\mathrm{BF}(I'_t)$, respectively, where $b < b'$.

(3) All bins are pairwise-identical or pairwise-closed, except for a 2-bin $C = \{c_1, c_2\}$ in $\mathrm{BF}(I_t)$ which does not exist in $\mathrm{BF}(I'_t)$, and two 1-bins $B'_1 = \{b'_1\}$, $B'_2 = \{b'_2\}$ in $\mathrm{BF}(I'_t)$ which do not exist in $\mathrm{BF}(I_t)$.

Note in all three cases, $\mathrm{BF}(I_t) \le \mathrm{BF}(I'_t)$. As this property is maintained until $t = n$, it implies the lemma. ◀

The entire proof of Proposition 3.2 will be given in the full version of this paper.

3.2 Simplifying the instance

Let I be a list of items larger than 1/3. Note that both the optimal and the Best Fit packing use only bins of configurations L, LM, MM, and possibly one M-bin. However, we can assume a simpler structure without substantial implications on the competitiveness of Best Fit.

▶ **Lemma 3.3.** *Let I be any list that can be packed optimally into $\mathrm{OPT}(I)$ LM-bins. If Best Fit has (asymptotic or absolute) approximation ratio α for I, then it has (asymptotic or absolute) approximation ratio α for any list of items larger than $1/3$ as well.*

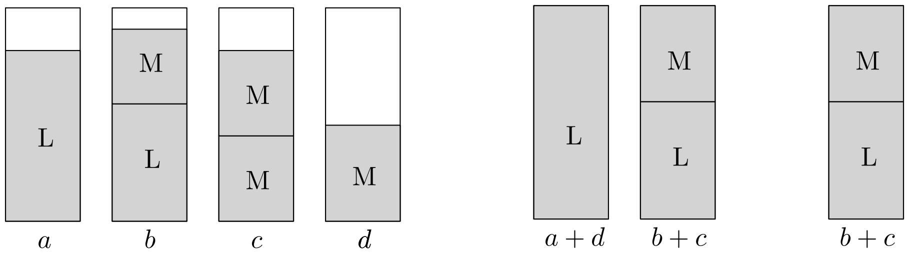

(a) $\mathrm{OPT}(I_0) = \mathrm{OPT}(I_1)$. **(b)** $\mathrm{OPT}(I_2)$. **(c)** $\mathrm{OPT}(I_3)$.

Figure 1 Construction from Lemma 3.3 to eliminate L-, MM-, and M-bins in the optimal packing.

Proof. Let I_0 be a list of items larger than 1/3 and let a, b, c, and $d \leq 1$ be the number of bins in $\mathrm{OPT}(I_0)$ with configurations L, LM, MM, and M, respectively (see Figure 1a). In several steps, we eliminate L-,MM-, and M-bins from $\mathrm{OPT}(I_0)$ while making the instance only harder for Best Fit.

First, we obtain I_1 from I_0 by replacing items of size 1/2 by items of size $1/2 - \varepsilon$. By choosing $\varepsilon > 0$ small enough, i.e., $\varepsilon < \min\{\delta^+ - 1/2, 1/2 - \delta^-\}$, where $\delta^+ = \min\{x_i \mid x_i > 1/2\}$ and $\delta^- = \max\{x_i \mid x_i < 1/2\}$, it is ensured that Best Fit packs all items in the same bins than before the modification. Further, the modification does not decrease the number of bins in an optimal packing, so we have $\mathrm{BF}(I_0) = \mathrm{BF}(I_1)$ and $\mathrm{OPT}(I_0) = \mathrm{OPT}(I_1)$. Now, we obtain I_2 from I_1 by increasing item sizes: We replace each of the $a+d$ items packed in 1-bins in $\mathrm{OPT}(I_1)$ by large items of size 1. Moreover, any 2-bin (MM or LM) in $\mathrm{OPT}(I_1)$ contains at least one item smaller than 1/2. These items are enlarged such that they fill their respective bin completely. Therefore, $\mathrm{OPT}(I_2)$ has $a + d$ L-bins and $b + c$ LM-bins (see Figure 1b). We have $\mathrm{OPT}(I_2) = \mathrm{OPT}(I_1)$ and, by Proposition 3.2, $\mathrm{BF}(I_2) \geq \mathrm{BF}(I_1)$. Finally, we obtain I_3 from I_2 by deleting the $a + d$ items of size 1. As size-1 items are packed separately in any feasible packing, $\mathrm{OPT}(I_3) = \mathrm{OPT}(I_2) - (a + d)$ and $\mathrm{BF}(I_3) = \mathrm{BF}(I_2) - (a + d)$. Note that $\mathrm{OPT}(I_3)$ contains only LM-bins (see Figure 1c) and, by assumption, Best Fit has (asymptotic or absolute) approximation ratio α for such lists. Therefore, in general we have a factor $\alpha \geq 1$ and an additive term β such that $\mathrm{BF}(I_3) \leq \alpha\,\mathrm{OPT}(I_3) + \beta$. It follows that

$$\mathrm{BF}(I_0) \leq \mathrm{BF}(I_2) = \mathrm{BF}(I_3) + (a + d) \leq \alpha\,\mathrm{OPT}(I_3) + (a + d) + \beta \leq \alpha\,\mathrm{OPT}(I_0) + \beta\,. \quad ◀$$

By Lemma 3.3, we can impose the following constraints on I without loss of generality.

Assumption. For the remainder of the section, we assume that the optimal packing of I has $k = \mathrm{OPT}(I)$ LM-bins. For $i \in [k]$, let l_i and m_i denote the large item and the medium item in the i-th bin, respectively. We call $\{l_i, m_i\}$ an *LM-pair*.

3.3 Good order pairs

If the adversary could control the order of items, he would send all medium items first, followed by all large items. This way, Best Fit opens $k/2$ MM-bins and k L-bins and therefore is 1.5-competitive. In a random permutation, we can identify structures with a positive impact on the Best Fit packing. This is formalized in the following random event.

▶ **Definition 3.4.** *Consider a fixed permutation $\pi \in \mathcal{S}_n$. We say that an LM-pair $\{l_i, m_i\}$ arrives in good order (or is a good order pair) if l_i arrives before m_i in π.*

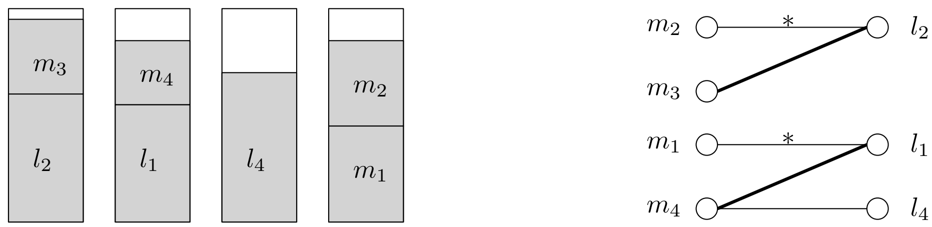

(a) Best Fit packing $\mathrm{BF}(I(7))$. **(b)** Graph G_7.

Figure 2 Visualization of Example 3.6. In Figure 2b, BF-edges are solid, while OPT-edges are thin. An asterisk indicates an OPT-edge in good order.

Note that in the adversarial setting no LM-pair arrives in good order, while in a random permutation, this holds for any LM-pair independently with probability 1/2. The next lemma is central for the proof of Theorem 1.1. It shows that the number of LM-pairs in good order bound the number of LM-bins in the final Best Fit packing from below.

▶ **Lemma 3.5.** *Let $\pi \in \mathcal{S}_n$ be any permutation and let X be the number of LM-pairs arriving in good order in I^π. The packing $\mathrm{BF}(I^\pi)$ has at least X LM-bins.*

To prove Lemma 3.5, we model the Best Fit packing by the following bipartite graph: Let $G_t = (\mathcal{M}_t \cup \mathcal{L}_t, E_t^{\mathrm{BF}} \cup E_t^{\mathrm{OPT}})$, where $\mathcal{M}_t$ and $\mathcal{L}_t$ are the sets of medium and large items in $I^\pi(t)$, respectively. The sets of edges represent the LM-matchings in the Best Fit packing and in the optimal packing at time t, i.e.,

$$E_t^{\mathrm{BF}} = \big\{\{m, l\} \in (\mathcal{M}_t \times \mathcal{L}_t) \mid m \text{ and } l \text{ are packed into the same bin in } \mathrm{BF}(I^\pi(t))\big\}$$
$$E_t^{\mathrm{OPT}} = \big\{\{m_i, l_i\} \in (\mathcal{M}_t \times \mathcal{L}_t) \mid i \in [k]\big\}\,.$$

We distinguish OPT-edges in good and bad order, according to the corresponding LM-pair. Note that G_t is not necessarily connected and may contain parallel edges. We illustrate the graph representation by a small example.

▶ **Example 3.6** (see Figure 2)**.** Let $\varepsilon > 0$ be sufficiently small and define for $i \in [4]$ large items $l_i = 1/2 + i\varepsilon$ and medium items $m_i = 1/2 - i\varepsilon$. Consider the list $I^\pi = (l_2, l_1, m_3, m_4, l_4, m_1, m_2, l_3)$. Figures 2a and 2b show the Best Fit packing and the corresponding graph G_7 before arrival of the last item. Note that I^π has two good order pairs ($\{l_1, m_1\}$ and $\{l_2, m_2\}$) and, according to Lemma 3.5, the packing has two LM-bins.

The proof of Lemma 3.5 essentially boils down to the following claim:

▷ Claim 3.7. In each round t and in each connected component C of G_t, the number of BF-edges in C is at least the number of OPT-edges in good order in C.

We first show how Lemma 3.5 follows from Claim 3.7. Then, we work towards the proof of Claim 3.7.

Proof of Lemma 3.5. Claim 3.7 implies that in G_n, the total number of BF-edges (summed over all connected components) is at least X. Therefore, the packing has at least X LM-bins and thus not less than the number of good order pairs X. ◀

Before proving Claim 3.7, it is reasonable to observe the following property of G_t.

▷ Claim 3.8. Consider the graph G_t for some $t \in [n]$. Let $Q = (b_w, a_{w-1}, b_{w-1}, \ldots, a_1, b_1)$ with $w \geq 1$ be a maximal alternating path such that $\{a_j, b_j\}$ is an OPT-edge in good order and $\{a_j, b_{j+1}\}$ is a BF-edge for any $j \in [w-1]$ (i.e., a-items and b-items represent medium and large items, respectively). It holds that $b_w \geq b_1$.

Proof. We show the claim by induction on w. Note that the items' indices only reflect the position along the path, not the arrival order. For $w = 1$, we have $Q = (b_w) = (b_1)$ and thus, the claim holds trivially.

Now, fix $w \geq 2$ and suppose that the claim holds for all paths Q' with $w' \leq w - 1$. We next prove $b_w \geq b_1$. Let $t' \leq t$ be the arrival time of the a-item a_d that arrived latest among all a-items in Q. We consider the graph $G_{t'-1}$, i.e., the graph immediately before arrival of a_d and its incident edges. Note that in $G_{t'-1}$, all items a_i with $i \in [w-1] \setminus \{d\}$ and b_i with $i \in [w-1]$ are visible. Let $Q' = (b_w, \ldots, a_{d+1}, b_{d+1})$ and $Q'' = (b_d, \ldots, a_1, b_1)$ be the connected components of b_w and b_1 in $G_{t'-1}$. As Q' and Q'' are maximal alternating paths shorter than Q, we obtain from the induction hypothesis $b_w \geq b_{d+1}$ and $b_d \geq b_1$. Note that b_{d+1} and b_1 were visible and packed into L-bins on arrival of a_d. Further, a_d and b_1 would fit together, as $a_d + b_1 \leq a_d + b_d \leq 1$. However, Best Fit packed a_d with b_{d+1}, implying $b_{d+1} \geq b_1$. Combining the inequalities yields $b_w \geq b_{d+1} \geq b_1$, which concludes the proof. ◁

Now, we are able to prove the remaining technical claim.

Proof of Claim 3.7. Note that the number of OPT-edges in good order can only increase on arrival of a medium item m_i where $\{m_i, l_i\}$ is an LM-pair in good order. Therefore, it is sufficient to verify Claim 3.7 in rounds $t_1 < \ldots < t_j$ such that in round t_i, item m_i arrives and l_i arrived previously.

Induction base. In round t_1, there is one OPT-edge $\{m_1, l_1\}$ in good order. We need to show that there exists at least one BF-edge in G_{t_1}, or, alternatively, at least one LM-bin in the packing. If the bin of l_1 contains a medium item different from m_1, we identified one LM-bin. Otherwise, Best Fit packs m_1 together with l_1 or some other large item, again creating an LM-bin.

Induction hypothesis. Fix $i \geq 2$ and assume that Claim 3.7 holds up to round t_{i-1}.

Induction step. We only consider the connected component of m_i, as by the induction hypothesis, the claim holds for all remaining connected components. If m_i is packed into an LM-bin, the number of BF-edges increases by one and the claim holds for round t_i. Therefore, assume that m_i is packed by Best Fit in an M- or MM-bin. This means that in G_{t_i}, vertex m_i is incident to an OPT-edge in good order, but not incident to any BF-edge. Let $P = (m_i, l_i, \ldots, v)$ be the maximal path starting from m_i alternating between OPT-edges and BF-edges.

Case 1: v is a medium item For illustration, consider Figure 2b with $m_i = m_2$ and $v = m_3$. Since P begins with an OPT-edge and ends with a BF-edge, the number of BF-edges in P equals the number of OPT-edges in P. The latter number is clearly at least the number of OPT-edges in good order in P.

Case 2: v is a large item For illustration, consider Figure 2b with $m_i = m_1$ and $v = l_4$. We consider two cases. If P contains at least one OPT-edge which is not in good order, the claim follows for the same argument as in Case 1.

Now, suppose that all OPT-edges in P are in good order. Let P' be the path obtained from P by removing the item m_i. As P' satisfies the premises of Claim 3.8, we obtain $l_i \geq v$. This implies that m_i and v would fit together, as $m_i + v \leq m_i + l_i \leq 1$. However, m_i is packed in an M- or MM-bin by assumption, although v is a feasible option on arrival of m_i. As this contradicts the Best Fit rule, we conclude that case 2 cannot happen. ◁

3.4 Final proofs

Finally, we prove the main result of this section.

Proof of Theorem 1.1. Let X be the number of good order pairs in I^σ and let Y be the number of LM-bins in the packing $\mathrm{BF}(I^\sigma)$. We have $Y \geq X$ by Lemma 3.5. For the remaining large and medium items, Best Fit uses $(k-Y)$ L-bins and $\lceil (k-Y)/2 \rceil$ MM-bins (including possibly one M-bin), respectively. Therefore,

$$\mathrm{BF}(I^\sigma) = Y + (k-Y) + \left\lceil \frac{k-Y}{2} \right\rceil \leq k + \left\lceil \frac{k-X}{2} \right\rceil = \frac{3k}{2} - \frac{X}{2} + \frac{\xi(X)}{2}\,, \tag{1}$$

where $\xi(X) = (k-X) \bmod 2$. Using linearity and monotonicity of expectation, we obtain

$$\mathrm{E}[\mathrm{BF}(I^\sigma)] \leq \frac{3k}{2} - \frac{\mathrm{E}[X]}{2} + \frac{\Pr[\xi(X)=1]}{2}\,. \tag{2}$$

Since σ is uniformly distributed on $\mathcal{S}_n$, each LM-pair arrives in good order with probability 1/2, independently of all other pairs. Therefore, X follows a binomial distribution with parameters k and 1/2, implying $\mathrm{E}[X] = k/2$ and $\Pr[\xi(X)=1] = 1/2$. Hence,

$$\mathrm{E}[\mathrm{BF}(I^\sigma)] \leq \frac{3k}{2} - \frac{k/2}{2} + \frac{1/2}{2} = \frac{5k}{4} + \frac{1}{4} = \frac{5}{4}\,\mathrm{OPT}(I) + \frac{1}{4}\,, \tag{3}$$

where we used $k = \mathrm{OPT}(I)$. This concludes the proof. ◀

To obtain a slightly weaker bound on the absolute random order ratio (Proposition 1.2), we analyze some special cases more carefully.

Proof of Proposition 1.2. For $k \geq 4$ the claim follows immediately from Equation (3):

$$\frac{\mathrm{E}[\mathrm{BF}(I^\pi)]}{\mathrm{OPT}(I)} = \frac{(5k)/4 + 1/4}{k} = \frac{5}{4} + \frac{1}{4k} \leq \frac{21}{16}\,.$$

Since Best Fit is clearly optimal for $k = 1$, it remains to verify the cases $k \in \{2,3\}$.

$k = 2$ It is easily verified that there are 16 out of $4! = 24$ permutations where Best Fit is optimal and that it opens at most 3 bins otherwise. Therefore,

$$\mathrm{E}[\mathrm{BF}(I^\sigma)] = \frac{1}{4!} \cdot \left(16\,\mathrm{OPT}(I) + 8 \cdot \frac{3}{2}\,\mathrm{OPT}(I)\right) = \frac{7}{6}\,\mathrm{OPT}(I) < \frac{21}{16}\,\mathrm{OPT}(I)\,.$$

$k = 3$ When k is odd, there must be at least one LM-bin in the Best Fit packing: Suppose for contradiction that all M-items are packed in MM- or M-bins. As k is odd, there must be an item m_i packed in an M-bin. If l_i arrives before m_i, item l_i is packed in an L-bin, as there is no LM-bin. Therefore, Best Fit would pack m_i with l_i or some other L-item instead of opening a new bin. If l_i arrives after m_i, Best Fit would pack l_i with m_i or some other M-item. We have a contradiction in both cases.
Therefore, for $k = 3$ we have at least one LM-bin, even if no LM-pair arrives in good order. Consider the proof of Theorem 1.1. Instead of $Y \geq X$, we can use the stronger bound

$Y \geq X'$ with $X' := \max\{1, X\}$ on the number of LM-bins. The new random variable satisfies $\mathrm{E}[X'] = k/2 + 1/2^k$ and $\Pr[\xi(X') = 1] = 1/2 - 1/2^k$. Adapting Equations (1) and (2) appropriately, we obtain

$$\frac{\mathrm{E}[\mathrm{BF}(I^\sigma)]}{\mathrm{OPT}(I)} = \frac{1}{k} \cdot \left(\frac{3k}{2} - \frac{k/2 + 1/2^k}{2} + \frac{1/2 - 1/2^k}{2} \right) = \frac{5}{4} + \frac{1}{4k} - \frac{1}{k2^k} = \frac{31}{24} < \frac{21}{16} . \blacktriangleleft$$

4 Lower bounds

In this section, we present the improved lower bound on RR^∞_{BF} (Theorem 1.4) and the first lower bound on the absolute random order ratio RR_{BF}.

4.1 Asymptotic random order ratio

Consider the i.i.d.-model, where the input is a sequence of independent and identically distributed (i.i.d) random variables. Here, the performance measure for an algorithm $\mathcal{A}$ is $\mathrm{E}[\mathcal{A}(I_n(F))]/\mathrm{E}[\mathrm{OPT}(I_n(F))]$, where $I_n(F) := (X_1, \ldots, X_n)$ is a list of n random variables drawn i.i.d according to F. This model is in general weaker than the random order model, which is why lower bounds in the random order model can be obtained from the i.i.d. model. This is formalized in the following lemma.

▶ **Lemma 4.1.** *Consider any online bin packing algorithm $\mathcal{A}$. Let F be a discrete distribution and $I_n(F) = (X_1, \ldots, X_n)$ be a list of i.i.d. samples. There exists a list I of n items such that for $n \to \infty$,*

$$\frac{\mathrm{E}[\mathcal{A}(I^\sigma)]}{\mathrm{OPT}(I)} \geq \frac{\mathrm{E}[\mathcal{A}(I_n(F))]}{\mathrm{E}[\mathrm{OPT}(I_n(F))]} .$$

Moreover, if there exists a constant $c > 0$ such that $X_i \geq c$ for all $i \in [n]$, we have $\mathrm{OPT}(I) \geq cn$.

This technique has already been used in [28] to establish the previous bound of 1.08, however, without a formal proof. Apparently, the only published proofs of this reduction technique address bin covering [5, 11]. We will provide a constructive proof of Lemma 4.1 in the full version of this paper. Theorem 1.4 follows by combining Lemmas 4.1 and 4.2.

▶ **Lemma 4.2.** *There exists a discrete distribution F such that for $n \to \infty$, we have $\mathrm{E}[\mathcal{A}(I_n(F))] > \frac{11}{10} \mathrm{E}[\mathrm{OPT}(I_n(F))]$ and each sample X_i satisfies $X_i \geq 1/4$.*

Proof. Let F be the discrete distribution which gives an item of size $1/4$ with probability p and an item of size $1/3$ with probability $q := 1 - p$. First, we analyze the optimal packing. Let N_4 and N_3 be the number of items with size $1/4$ and $1/3$ in $I_n(F)$, respectively. We have

$$\mathrm{E}[\mathrm{OPT}(I_n(F))] \leq \mathrm{E}\left[\frac{N_4}{4} + \frac{N_3}{3} + 2\right] = \frac{np}{4} + \frac{nq}{3} + 2 = n\left(\frac{1}{3} - \frac{p}{12} + \frac{2}{n}\right) .$$

Now, we analyze the expected behavior of Best Fit for $I_n(F)$. As the only possible item sizes are $1/4$ and $1/3$, we can consider each bin of load more than $3/4$ as closed. Moreover, the number of possible loads for open bins is small and Best Fit maintains at most two open bins at any time. Therefore, we can model the Best Fit packing by a Markov chain as follows. Let the nine states $\mathsf{A}, \mathsf{B}, \ldots, \mathsf{I}$ be defined as in Figure 3b. The corresponding

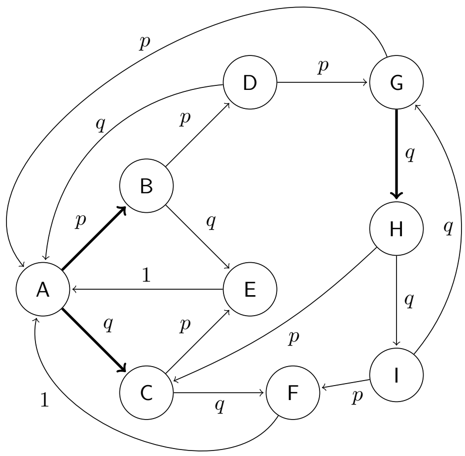

(a) Transition diagram.

State	Load of open bin(s)
A	–
B	1/4
C	1/3
D	2/4
E	7/12
F	2/3
G	3/4
H	3/4, 1/3
I	3/4, 2/3

(b) Description of states.

Figure 3 Markov chain from Lemma 4.2. Bold arcs in Figure 3a indicate transitions where Best Fit opens a new bin.

transition diagram is depicted in Figure 3a. This Markov chain converges to the stationary distribution

$$\omega = (\omega_{\mathsf{A}}, \dots, \omega_{\mathsf{I}}) = \frac{1}{\lambda}\left(1,\ p,\ q + pq\vartheta,\ p^2,\ 2pq + p^2q\vartheta,\ q^2 + 2pq^2\vartheta,\ \vartheta,\ q\vartheta,\ q^2\vartheta\right), \tag{4}$$

where we defined $\vartheta = \frac{p^3}{1-q^3}$ and $\lambda = \vartheta q\left(3 - q^2\right) + \vartheta + 3$. A formal proof of this fact will appear in the full version of this paper.

Let $V_{\mathsf{S}}(t)$ denote the number of visits to state $\mathsf{S} \in \{\mathsf{A}, \dots, \mathsf{I}\}$ up to time t. By a basic result from the theory of ergodic Markov chains (see [30, Sec. 4.7]), it holds that $\lim_{t\to\infty} V_{\mathsf{S}}(t) = n\omega_S$. In other words, the proportion of time spent in state S approaches its probability ω_{S} in the stationary distribution. This fact can be used to bound the total number of opened bins over time. Note that Best Fit opens a new bin on the transitions $A \to B$, $A \to C$, and $G \to H$ (see Figure 3a). Hence, $\mathrm{E}[\mathrm{BF}(I_n(F))] = n\left(\omega_{\mathsf{A}} + q\omega_{\mathsf{G}}\right)$. Setting $p = 0.60$, we obtain finally

$$\lim_{n\to\infty} \frac{\mathrm{E}[\mathrm{BF}(I_n(F))]}{\mathrm{E}[\mathrm{OPT}(I_n(F))]} \geq \lim_{n\to\infty} \frac{\omega_{\mathsf{A}} + q\omega_{\mathsf{G}}}{\frac{1}{3} - \frac{p}{12} + \frac{2}{n}} = \frac{1 + q\vartheta}{\lambda \cdot \left(\frac{1}{3} - \frac{p}{12}\right)} > \frac{11}{10}\,. \qquad ◀$$

4.2 Absolute random order ratio

Theorem 1.5 follows from the following lemma.

▶ **Lemma 4.3.** *There exists a list I such that* $\mathrm{E}[\mathrm{BF}(I^\sigma)] = \frac{13}{10}\,\mathrm{OPT}(I)$.

Proof. Let $\varepsilon > 0$ be sufficiently small and let $I := (a_1, a_2, b_1, b_2, c)$ where

$$a_1 = a_2 = \frac{1}{3} + 4\varepsilon, \quad b_1 = b_2 = \frac{1}{3} + 16\varepsilon, \quad c = \frac{1}{3} - 8\varepsilon\,.$$

An optimal packing of I needs two bins $\{a_1, a_2, c\}$ and $\{b_1, b_2\}$, thus $\mathrm{OPT}(I) = 2$. Best Fit needs two or three bins depending on the order of arrival.

Let E be the event that exactly one b-item arrives within the first two rounds. After the second item, the first bin is closed, as its load is at least $\frac{1}{3} + 16\varepsilon + \frac{1}{3} - 8\varepsilon = \frac{2}{3} + 8\varepsilon$. Among the remaining three items, there is a b-item of size $\frac{1}{3} + 16\varepsilon$ and at least one a-item of size $\frac{1}{3} + 4\varepsilon$. This implies that a third bin needs to be opened for the last item. As there are exactly $2 \cdot 3 \cdot 2! \cdot 3! = 72$ permutations where E happens, we have $\Pr[E] = \frac{72}{5!} = \frac{3}{5}$.

On the other side, Best Fit needs only two bins if one of the events F and G, defined in the following, happen. Let F be the event that both b-items arrive in the first two rounds. Then, the remaining three items fit into one additional bin. Moreover, let G be the event that the set of the first two items is a subset of $\{a_1, a_2, c\}$. Then, the first bin has load at least $\frac{2}{3} - 4\varepsilon$, thus no b-item can be packed there. Again, this ensures a packing into two bins. By counting permutations, we obtain $\Pr[F] = \frac{2!\cdot 3!}{5!} = \frac{1}{10}$ and $\Pr[G] = \frac{3\cdot 2!\cdot 3!}{5!} = \frac{3}{10}$.

As the events E, F, and G partition the probability space, we obtain

$$\frac{\mathrm{E}[\mathrm{BF}(I^\sigma)]}{\mathrm{OPT}(I)} = \frac{\Pr[E]\cdot 3 + (\Pr[F] + \Pr[G])\cdot 2}{2} = \frac{\frac{3}{5}\cdot 3 + \left(\frac{1}{10} + \frac{3}{10}\right)\cdot 2}{2} = \frac{13}{10}. \qquad \blacktriangleleft$$

The construction from the above proof is used in [23] to prove that Best Fit is 1.5-competitive under adversarial arrival order if all item sizes are close to 1/3. Interestingly, it gives a strong lower bound on the absolute random order ratio as well.

Finally, we revisit the case of 1/3-large items. To prove Proposition 1.3, we need to construct a list I with 1/3-large items and $\mathrm{E}[\mathrm{BF}(I^\sigma)] > \frac{6}{5}\,\mathrm{OPT}(I)$. Due to space restrictions, we only sketch the construction here and will provide the entire analysis in the full version of this paper.

Proof sketch of Proposition 1.3. We construct a list of $k = 3$ LM-pairs. For sufficiently small $\varepsilon > 0$ and $i \in [k]$ define $l_i = \frac{1}{2} + i\varepsilon$ and $m_i = \frac{1}{2} - i\varepsilon$. This way, $l_1 < l_2 < l_3$ and $m_1 > m_2 > m_3$. Clearly, $\mathrm{OPT}(I) = 3$. We can show that Best Fit uses 4 instead of 3 bins in at least 440 permutations. Therefore,

$$\frac{\mathrm{E}[\mathrm{BF}(I^\sigma)]}{\mathrm{OPT}(I)} \geq \frac{\frac{1}{6!}\cdot(440\cdot 4 + (6! - 440)\cdot 3)}{3} = \frac{65}{54} > \frac{6}{5}. \qquad \blacktriangleleft$$

References

1 A. Al-Herz and A. Pothen. A 2/3-approximation algorithm for vertex-weighted matching. *Discrete Applied Mathematics*, 2019. (in press).

2 J. Balogh, J. Békési, G. Dósa, L. Epstein, and A. Levin. A new and improved algorithm for online bin packing. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*, volume 112 of *LIPIcs*, pages 5:1–5:14, 2018.

3 J. Balogh, J. Békési, G. Dósa, L. Epstein, and A. Levin. A new lower bound for classic online bin packing. In *Approximation and Online Algorithms - 17th International Workshop, (WAOA)*, volume 11926 of *Lecture Notes in Computer Science*, pages 18–28. Springer, 2019.

4 J. Boyar, G. Dósa, and L. Epstein. On the absolute approximation ratio for first fit and related results. *Discret. Appl. Math.*, 160(13-14):1914–1923, 2012.

5 M. G. Christ, L. M. Favrholdt, and K. S. Larsen. Online bin covering: Expectations vs. guarantees. *Theor. Comput. Sci.*, 556:71–84, 2014.

6 H. I. Christensen, A. Khan, S. Pokutta, and P. Tetali. Approximation and online algorithms for multidimensional bin packing: A survey. *Comput. Sci. Rev.*, 24:63–79, 2017.

7 E. G. Coffman, J. Csirik, G. Galambos, S. Martello, and D. Vigo. Bin packing approximation algorithms: survey and classification. In *Handbook of combinatorial optimization*, pages 455–531. Springer New York, 2013.

8 W. F. de la Vega and G. S. Lueker. Bin packing can be solved within 1+epsilon in linear time. *Combinatorica*, 1(4):349–355, 1981.

9 G. Dósa and J. Sgall. First fit bin packing: A tight analysis. In *30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 20 of *LIPIcs*, pages 538–549, 2013.

10 G. Dósa and J. Sgall. Optimal analysis of best fit bin packing. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 429–441, 2014.

11 C. Fischer and H. Röglin. Probabilistic analysis of the dual next-fit algorithm for bin covering. In *LATIN 2016: Theoretical Informatics - 12th Latin American Symposium*, pages 469–482, 2016.

12 C. Fischer and H. Röglin. Probabilistic analysis of online (class-constrained) bin packing and bin covering. In *LATIN 2018: Theoretical Informatics - 13th Latin American Symposium*, volume 10807 of *Lecture Notes in Computer Science*, pages 461–474. Springer, 2018.

13 B. Gamlath, M. Kapralov, A. Maggiori, O. Svensson, and D. Wajc. Online matching with general arrivals. In *60th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 26–37, 2019.

14 M. R. Garey, R. L. Graham, and J. D. Ullman. Worst-case analysis of memory allocation algorithms. In *Proceedings of the 4th Annual ACM Symposium on Theory of Computing (STOC)*, pages 143–150, 1972.

15 M. R. Garey and D. S. Johnson. "Strong" NP-completeness results: Motivation, examples, and implications. *J. ACM*, 25(3):499–508, 1978.

16 R. L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal of Applied Mathematics*, 17(2):416–429, 1969.

17 A. Gupta and S. Singla. Random-order models. *CoRR*, abs/2002.12159, 2020. `arXiv:2002.12159`.

18 R. Hoberg and T. Rothvoss. A logarithmic additive integrality gap for bin packing. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2616–2625, 2017.

19 Z. Huang, N. Kang, Z. G. Tang, X. Wu, Y. Zhang, and X. Zhu. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 17–29, 2018.

20 Z. Huang, B. Peng, Z. Gavin Tang, R. Tao, X. Wu, and Y. Zhang. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2875–2886, 2019.

21 Z. Huang, Z. G. Tang, X. Wu, and Y. Zhang. Online vertex-weighted bipartite matching: Beating 1-1/e with random arrivals. *ACM Trans. Algorithms*, 15(3):38:1–38:15, 2019.

22 D. S. Johnson. Fast algorithms for bin packing. *J. Comput. Syst. Sci.*, 8(3):272–314, 1974.

23 D. S. Johnson, A. J. Demers, J. D. Ullman, M. R. Garey, and R. L. Graham. Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM J. Comput.*, 3(4):299–325, 1974.

24 E. G. Coffman Jr., J. Csirik, L. Rónyai, and A. Zsbán. Random-order bin packing. *Discret. Appl. Math.*, 156(14):2810–2816, 2008.

25 E. G. Coffman Jr. and G. S. Lueker. *Probabilistic analysis of packing and partitioning algorithms*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1991.

26 N. Karmarkar and R. M. Karp. An efficient approximation scheme for the one-dimensional bin-packing problem. In *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 312–320, 1982.

27 R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 352–358, 1990.

28 C. Kenyon. Best-fit bin-packing with random order. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 359–364, 1996.

29 C. C. Lee and D. T. Lee. A simple on-line bin-packing algorithm. *J. ACM*, 32(3):562–572, 1985.

30 D. A. Levin and Y. Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

31 M. Mahdian and Q. Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 597–606, 2011.

32 A. Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.

33 F. D. Murgolo. Anomalous behavior in bin packing algorithms. *Discret. Appl. Math.*, 21(3):229–243, 1988.

34 P. V. Ramanan. Average-case analysis of the smart next fit algorithm. *Inf. Process. Lett.*, 31(5):221–225, 1989.

35 P. V. Ramanan, D. J. Brown, C. C. Lee, and D. T. Lee. On-line bin packing in linear time. *J. Algorithms*, 10(3):305–326, 1989.

36 S. S. Seiden. On the online bin packing problem. *J. ACM*, 49(5):640–671, 2002.