Best Fit Bin Packing with Random Order Revisited

Susanne Albers¹ · Arindam Khan² · Leon Ladewig¹

Received: 13 September 2020 / Accepted: 9 June 2021 / Published online: 1 July 2021
© The Author(s) 2021

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.

Keywords Online bin packing · Random arrival order · Probabilistic analysis

Work supported by the European Research Council, Grant Agreement No. 691672. A preliminary version of this paper received the best paper award at *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)* and was published in the conference proceedings [2].

✉ Leon Ladewig
leon.ladewig@tum.de

Extended author information available on the last page of the article

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. An online algorithm must pack x_i without knowing the items following x_i and without modifying the packing of previous items.

Bin packing was mentioned first by Ullman [38]. As the problem is strongly NP-complete [17], research mainly focuses on efficient approximation algorithms. The offline problem is well understood and admits even approximation schemes [20, 26, 39]. The online variant is still a very active field in the community [7], as the asymptotic approximation ratio of the best online algorithm is still unknown [3, 4]. The first approximation algorithms for the problem, First Fit and Best Fit, have been analyzed in [38] and a subsequent work by Garey et al. [16]. Johnson published the Next Fit algorithm briefly afterwards [24]. 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 [3, 35, 36].

To measure the performance of an algorithm, different metrics exist. For an algorithm $\mathcal{A}$, let $\mathcal{A}(I)$ and $\text{OPT}(I)$ 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

$$R_{\mathcal{A}}^{\infty} = \limsup_{k \rightarrow \infty} \sup_{I \in \mathcal{I}} \{ \mathcal{A}(I) / \text{OPT}(I) \mid \text{OPT}(I) = k \}.$$

Note that $R_{\mathcal{A}}^{\infty}$ focuses on instances where $\text{OPT}(I)$ 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*

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

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 [25], the absolute ratios have been settled rather recently [11, 12].

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 [16]. In contrast, it is known that Best Fit performs significantly better if items appear in non-increasing order [25]. 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 over all permutations. We define

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

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 $I^{\sigma} = (x_{\sigma(1)}, \dots, x_{\sigma(n)})$ 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 [8] by Coffman et al. for an extensive survey

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 [12]	1.7 [25]	$1.08 \leq RR_{\text{BF}}^{\infty} \leq 1.5$ [28]
First Fit	1.7 [11]	1.7 [25]	–
Next Fit	$2 - 1/\lceil \gamma \rceil$ [5]	2 [24]	2 [9]

on (online) bin packing. For further problems studied in the random order model, see [19].

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_{BF}^{∞} , demonstrating that random order significantly improves upon $R_{\text{BF}}^{\infty} = 1.7$. However, it is conjectured in [8, 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 $\frac{t}{n}\text{OPT}(I) + o(n)$ bins. On the other side, Kenyon proves that $RR_{\text{BF}}^{\infty} \geq 1.08$. 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. [9] analyzed next-fit in the random order model and showed that $RR_{\text{NF}}^{\infty} = 2$, matching the asymptotic approximation ratio $RR_{\text{NF}}^{\infty} = 2$ (see Table 1). Fischer and Röglin [14] obtained analogous results for Worst Fit [24] 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 [18] and bin covering [14].

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 $R_{\text{DNF}}^{\infty} = 1/2$ which is best possible for any online algorithm [6]. Under random arrival order, Christ et al. [6] showed that $RR_{\text{DNF}}^{\infty} \leq 4/5$. This upper bound was improved later by Fischer and Röglin [13] to $RR_{\text{DNF}}^{\infty} \leq 2/3$. The same group of authors further showed that $RR_{\text{DNF}}^{\infty} \geq 0.501$, i.e., DNF performs strictly better under random order than in the adversarial setting [14].

Matching Online matching can be seen as the key problem in the field of online algorithms [32]. Inspired by the seminal work of Karp et al. [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 [15, 21, 22], vertex-weighted versions [1, 23] and, most relevant to our work, random arrival order [23, 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 [6, 9, 19]. 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.

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 $\text{BF}(I) \leq \frac{3}{2}\text{OPT}(I) + 1$ under adversarial order. We show that Best Fit performs significantly better under random arrival order:

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

The proof of Theorem 1 is developed in Sect. 3 and based on several pillars. First, we show that Best Fit is monotone in this case (Proposition 3), 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 Sect. 3.3 with Lemma 2 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 $RR_{\text{BF}}^\infty \leq 1.5$ 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 [25].

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 Sect. 3.4, we investigate the absolute random order ratio for items larger than $1/3$ and obtain the following result.

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

The upper bound of $21/16$ is complemented by the following lower bound.

Proposition 2 *There is a list I of items larger than $1/3$ with $\mathbb{E}[\text{BF}(I^\sigma)] > \frac{6}{5}\text{OPT}(I)$.*

The proof of Proposition 2 is given in Appendix B.

We also make progress on the hardness side in the general case, which is presented in Sect. 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 2 *The asymptotic random order ratio of Best Fit is $RR_{BF}^{\infty} > 1.10$.*

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 3 *The absolute random order ratio of Best Fit is $RR_{BF} \geq 1.30$.*

Interestingly, our lower bound on the absolute random order ratio is notably larger than in the asymptotic case (see [28] and Theorem 2). This suggests either

- a significant discrepancy between RR_{BF} and RR_{BF}^{∞} , which is in contrast to the adversarial setting ($R_{BF} = R_{BF}^{\infty}$, see Table 1), or
- a disproof of the conjecture $RR_{BF}^{\infty} \approx 1.15$ mentioned in [8, 28].

2 Notation

We consider a list $I = (x_1, \dots, x_n)$ of n items throughout the paper. Due to the online setting, I is revealed in rounds $1, \dots, n$. In round t , item x_t arrives and in total, the prefix list $I(t) := (x_1, \dots, x_t)$ 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 $x_t \in (1/3, 1/2]$, 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 $BF(I)$. We slightly abuse the notation and refer to the corresponding packing by $BF(I)$ as well whenever the exact meaning is clear from the context. Similarly, we denote by $OPT(I)$ the number of bins and the corresponding packing of an optimal offline solution.

Finally, for any natural number n we define $[n] := \{1, \dots, n\}$. 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. We develop the technical foundations in Sects. 3.1 to 3.3. The final proofs of Theorem 1 and Proposition 1 are presented in Sect. 3.4.

3.1 Monotonicity

We first define the notion of monotone algorithms.

Definition 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 Next Fit [10], Best Fit is not monotone in general [33]. As a counterexample, consider the lists

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

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

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 $x \in (1/4, 1/3]$ to force Best Fit into anomalous behavior. Anyway, we have the following proposition.

Proposition 3 *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 $\text{BF}(I) \leq \text{BF}(I')$.*

We provide the proof of Proposition 3 in Appendix A. Enabled by the monotonicity, we can reduce an instance of $1/3$ -large items to an instance of easier structure. This construction is described in the following.

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 1 *Let I be any list that can be packed optimally into $\text{OPT}(I)$ 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.*

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 $\text{OPT}(I_0)$ with configurations L, LM, MM, and M, respectively (see Fig. 1a). In several steps, we eliminate L-, MM-, and M-bins from $\text{OPT}(I_0)$ 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 as before the modification. Further, the modification does not decrease the number of bins in an optimal packing, so we have $\text{BF}(I_0) = \text{BF}(I_1)$ and $\text{OPT}(I_0) = \text{OPT}(I_1)$.

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 $\text{OPT}(I_1)$ by large items of size 1. Moreover, any 2-bin (MM or LM) in $\text{OPT}(I_1)$ contains at least one item smaller than $1/2$. These items are enlarged such that they fill their respective bin completely. Therefore, $\text{OPT}(I_2)$ has $a + d$ L-bins and $b + c$ LM-bins (see Fig. 1b). We have $\text{OPT}(I_2) = \text{OPT}(I_1)$ and, by Proposition 3, $\text{BF}(I_2) \geq \text{BF}(I_1)$.

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, $\text{OPT}(I_3) = \text{OPT}(I_2) - (a + d)$ and $\text{BF}(I_3) = \text{BF}(I_2) - (a + d)$.

Note that $\text{OPT}(I_3)$ contains only LM-bins (see Fig. 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 $\text{BF}(I_3) \leq \alpha \text{OPT}(I_3) + \beta$. It follows that

$$\text{BF}(I_0) \leq \text{BF}(I_2) = \text{BF}(I_3) + (a + d) \leq \alpha \text{OPT}(I_3) + (a + d) + \beta \leq \alpha \text{OPT}(I_0) + \beta,$$

which concludes the proof. $\square$

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

By Lemma 1, 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 $k = \text{OPT}(I)$ 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 2 Consider a fixed permutation $\pi \in \mathcal{S}_n$. We say that the 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 .

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. It shows that the number of LM-pairs in good order bounds the number of LM-bins in the final Best Fit packing from below.

Lemma 2 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 $\text{BF}(I^\pi)$ has at least X LM-bins.

To prove Lemma 2, we model the Best Fit packing by the following bipartite graph: Let $G_t = (\mathcal{M}_t \cup \mathcal{L}_t, E_t^{\text{BF}} \cup E_t^{\text{OPT}})$, 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.,

$$E_t^{\text{BF}} = \{ \{m, l\} \in (\mathcal{M}_t \times \mathcal{L}_t) \mid m \text{ and } l \text{ are packed into the same bin in } \text{BF}(I^\pi(t)) \}$$

$$E_t^{\text{OPT}} = \{ \{m_i, l_i\} \in (\mathcal{M}_t \times \mathcal{L}_t) \mid i \in [k] \}.$$

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 1 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 $I^\pi = (l_2, l_1, m_3, m_4, l_4, m_1, m_2, l_3)$. Figure 2a, b 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 2, the packing has two LM-bins.

Fig. 2 Visualization of Example 1. In Fig. 2, BF-edges are solid, while OPT-edges are thin. An asterisk indicates an OPT-edge in good order

The proof of Lemma 2 essentially boils down to the following claim:

Claim 1 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 2 follows from Claim 1. Then, we work towards the proof of Claim 1.

Proof of Lemma 2 Claim 1 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 . $\square$

Before proving Claim 1, we show the following property of G_t .

Claim 2 Consider the graph G_t for some $t \in [n]$. Let $Q = (b_w, a_{w-1}, b_{w-1}, \dots, a_1, b_1)$ 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 $Q' = (b_w, \dots, a_{d+1}, b_{d+1})$ and $Q'' = (b_d, \dots, a_1, b_1)$ 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 $b_w \geq b_{d+1} \geq b_1$, concluding the proof. $\square$

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

Proof of Claim 1 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 1 in rounds $t_1 < \dots < t_j$ 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 1 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 $P = (m_i, l_i, \dots, v)$ 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 Fig. 2 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 Fig. 2 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 by 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 2, 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. $\square$

3.4 Final Proofs

Finally, we prove the main result of this section.

Proof of Theorem 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 $\text{BF}(I^\sigma)$. We have $Y \geq X$ by Lemma 2. 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,

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

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

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

Since σ is uniformly distributed on $\mathcal{S}_n$, each LM-pair arrives in good order with probability $1/2$. Therefore, $\mathbb{E}[X] = k/2$ and $\Pr[\xi(X) = 1] = 1/2$. Hence,

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

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

To obtain the upper bound of $21/16$ on the absolute random order ratio (Proposition 1), we analyze a few special cases more carefully.

Proof of Proposition 1 For $k \geq 4$, the claim follows immediately from Equation (3):

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

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,

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

$k = 3$ 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 packs 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 packs 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. 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 $\mathbb{E}[X'] = k/2 + 1/2^k$

and $\Pr[\xi(X') = 1] = 1/2 - 1/2^k$. Adapting Equations (1) and (2) appropriately, we obtain

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

□

4 Lower Bounds

In this section, we present the improved lower bound on RR_{BF}^∞ (Theorem 2) and the first lower bound on the absolute random order ratio RR_{BF} (Theorem 3).

4.1 Asymptotic Random Order Ratio

Another model of probabilistic analysis is the i.i.d.-model, where the input of the algorithm is a sequence of independent and identically distributed (i.i.d.) random variables. Here, the performance measure of algorithm $\mathcal{A}$ is $E[\mathcal{A}(I_n(F))]/E[\text{OPT}(I_n(F))]$, where $I_n(F) := (X_1, \dots, X_n)$ 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 3 *Consider any online bin packing algorithm $\mathcal{A}$. Let F be a discrete distribution and $I_n(F) = (X_1, \dots, X_n)$ be a list of i.i.d. samples. For $n \rightarrow \infty$, there exists a list I of n items such that*

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

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

This technique has already been used in [28] to establish the lower bound of 1.08, however, without a formal proof. Apparently, the only published proofs of this connection address bin covering [6, 13]. We provide a constructive proof of Lemma 3 in Appendix C for completeness. The improved lower bound from Theorem 2 now follows by combining Lemma 3 with the next lemma.

Lemma 4 *There exists a discrete distribution F such that for $n \rightarrow \infty$, we have $E[\mathcal{A}(I_n(F))] > \frac{11}{10} E[\text{OPT}(I_n(F))]$ 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

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

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 A, B, ..., I be defined as in Fig. 3b. The corresponding transition diagram is depicted in Fig. 3. This Markov chain converges to the stationary distribution

$$\begin{aligned} \omega &= (\omega_A, \dots, \omega_I) \\ &= \frac{1}{\lambda} (1, p, q + pq\vartheta, p^2, 2pq + p^2q\vartheta, q^2 + 2pq^2\vartheta, \vartheta, q\vartheta, q^2\vartheta), \end{aligned}$$

where we defined $\vartheta := \frac{p^3}{1-q^3}$ and $\lambda := \vartheta q(3 - q^2) + \vartheta + 3$. A formal proof of this fact can be found in Appendix C.2.

Fig. 3 Markov chain from Lemma 4. Bold arcs in Figure 3a indicate transitions where Best Fit opens a new bin

Let $V_S(t)$ denote the number of visits to state $S \in \{A, \dots, I\}$ up to time t . By a basic result from the theory of ergodic Markov chains (see [30, Sect. 4.7]), it holds that $\lim_{t \rightarrow \infty} \frac{1}{t} \cdot V_S(t) = \omega_S$. 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 $A \rightarrow B$, $A \rightarrow C$, and $G \rightarrow H$ (see Fig. 3a).

Hence, $E[\text{BF}(I_n(F))] = V_A(n) + qV_G(n)$. Setting $p = 0.60$, we obtain finally

$$\lim_{n \rightarrow \infty} \frac{E[\text{BF}(I_n(F))]}{E[\text{OPT}(I_n(F))]} \geq \lim_{n \rightarrow \infty} \frac{n\omega_A + nq\omega_G}{n\left(\frac{1}{3} - \frac{p}{12} + \frac{2}{n}\right)} = \frac{1 + q\vartheta}{\lambda \cdot \left(\frac{1}{3} - \frac{p}{12}\right)} > \frac{11}{10}. \quad (4)$$

□

4.2 Absolute Random Order Ratio

Theorem 3 follows from the following lemma.

Lemma 5 *There exists a list I such that $E[\text{BF}(I^\sigma)] = \frac{13}{10}\text{OPT}(I)$.*

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 has two bins $\{a_1, a_2, c\}$ and $\{b_1, b_2\}$, thus $\text{OPT}(I) = 2$. Subsequently, we argue that 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

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

□

The construction from the above proof is used in [25] 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.

Appendix A: Monotonicity

Proposition 3 follows by applying the following lemma iteratively. A technically similar proof appeared in [37], where Shor showed that the MBF algorithm from [37] is monotone under removal of items.

Lemma 6 *Let $I = (x_1, \dots, x_n)$ be any list of items larger than $1/3$. Let $I' = (x'_1, \dots, x'_n)$ with $x'_i > x_i$ for a single $i \in [n]$ and $x'_j = x_j$ for all $j \neq i$. We have $\text{BF}(I) \leq \text{BF}(I')$.*

Proof All bins in any packing of I or I' contain at most two items. We call two 1-bins of $\text{BF}(I)$ and $\text{BF}(I')$ *pairwise-identical* if they contain items of the same size. Moreover, we call any two 2-bins of $\text{BF}(I)$ and $\text{BF}(I')$ *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 show that at any time t , the packings $\text{BF}(I_t)$ and $\text{BF}(I'_t)$ are related in one of three ways (see Fig. 4).

Fig. 4 Relating the packings $\text{BF}(I_t)$ and $\text{BF}(I'_t)$ in the proof of Lemma 6

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

Note in all three cases, $\text{BF}(I_t) \leq \text{BF}(I'_t)$. As this property is maintained until $t = n$, it implies the lemma. Subsequently, we show the claim inductively for each round.

Before round i , both lists contain items of identical size and thus, the packings are clearly related by (* 1). In round i , both packings deviate possibly, since $x_i < x'_i$. Here, three cases can occur.

- (a) Both items go into new bins. Then, (* 2) holds.
- (b) Both items go into identical 1-bins. Then, (* 1) holds.
- (c) Best-fit packs x_i into an existing 1-bin B_1 in $\text{BF}(I_i)$, while it opens a new bin B'_2 for x'_i in $\text{BF}(I'_i)$. Before packing x_i into B_1 , there was a 1-bin B'_1 in $\text{BF}(I'_i)$ such that B_1 and B'_1 were pairwise identical. By packing x_i , bin B_1 becomes a 2-bin. Therefore, (* 3) holds with $C = B_1$ as 2-bin and B'_1, B'_2 as 1-bins.

Note that these three cases are exhaustive: The algorithm packs either both items into new bins (a), both items into existing 1-bins (b), or x_i into an existing 1-bin and x'_i into a new bin (c). In case (b), the 1-bins must be pairwise-identical as this holds for all bins up to time i . The inverse situation of (c) cannot occur since $x_i < x'_i$.

For the induction step, we consider round $t \geq i + 1$ and suppose that $\text{BF}(I_{t-1})$ and $\text{BF}(I'_{t-1})$ are related either by (* 1), (* 2), or (* 3). Note that in round t , in both lists the current item has equal size $x_t = x'_t$ again.

Case 1 If (* 1) holds at time $t - 1$, then (* 1) is maintained in round t .

Case 2 Suppose (* 2) holds at time $t - 1$. If x'_t is packed in $\text{BF}(I'_t)$ into an existing 1-bin, we know that Best Fit does not open a new bin for x_t in $\text{BF}(I_t)$ as well. Depending on the chosen bins, either (* 2) is maintained, or (* 1) holds.

Now, assume that x'_t is packed into a new bin in $\text{BF}(I'_t)$. Either, a new bin is opened in $\text{BF}(I_t)$ as well, or $\text{BF}(I_t)$ packs x_t into bin B . In the first case, (* 2) is maintained, in the second case (* 3) holds. Note that $\text{BF}(I_t)$ cannot pack x_t into any existing 1-bin other than B , as otherwise $\text{BF}(I'_t)$ would pack x'_t into the corresponding identical bin instead of opening a new bin.

Case 3 Suppose (* 3) holds at time $t - 1$. If Best Fit packs x_t into an existing 1-bin, we know that it packs x'_t into an existing 1-bin as well. Hence, in both packings a 1-bin becomes a 2-bin and (* 3) is maintained.

If Best Fit packs x_t into a new bin, we know that none of the existing 1-bins in $\text{BF}(I'_t)$, except for possibly B'_1 or B'_2 , are suitable for x'_t . Therefore, Best Fit either opens a new bin to pack x'_t as well (then, (* 3) is maintained), or it packs x'_t into $B' \in \{B'_1, B'_2\}$. Suppose that $B' = B'_1$. Then, all bins in the packings are pairwise-identical or closed, except for $B'_2 = \{b'_2\}$ in $\text{BF}(I'_t)$ and the new bin $\{x_t\}$ in $\text{BF}(I_t)$. Hence, (* 2) holds if $b'_2 > x_t$. To see this, observe that $b'_1 + b'_2 > 1$ since b'_1 and b'_2

have been packed into separate bins previously. Moreover, since Best Fit packed x'_i with b'_1 , we have $x_i + b'_1 \leq 1$. Therefore, $b'_2 > 1 - b'_1 \geq 1 - (1 - x_i) = x_i$. The case $B' = B'_2$ is analogous. $\square$

Appendix B: Lower Bound for $1/3$ -Large Items

Here, we prove the existence of a list I with $1/3$ -large items and $E[\text{BF}(I^\sigma)] > \frac{6}{5} \text{OPT}(I)$.

Proof of Proposition 2 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, $\text{OPT}(I) = 3$. We can show that Best Fit uses 4 instead of 3 bins in at least 440 permutations. Therefore,

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

We call the event that Best Fit packs two items different from $\{l_i, m_i\}$ into the same bin a *non-optimal match*. This event occurs if Best Fit packs either l_i and m_j for $1 \leq i < j \leq 3$, or two medium items m_i and m_j together into the same bin. It is easy to see that any packing with a non-optimal match needs at least 4 bins. However, Best Fit never uses more than 4 bins.

In the following, we partition the set of all permutations with a non-optimal match according to the first time where this happens. Note that the first possibility is after arrival of the second item. Moreover, if no non-optimal match happened before arrival of the fifth item, the fifth and sixth item build an LM-pair and therefore, no non-optimal match occurs at all.

Let $a_1, \dots, a_4$ denote the first four items in a fixed permutation.

Case 1: First non-optimal match after the second item

We have either $a_1, a_2 \in \{l_i, m_j\}$ with $i < j$, or $a_1, a_2 \in \{m_i, m_j\}$ with $i \neq j$. In both cases, we can choose the set $\{a_1, a_2\}$ in three ways, can arrange a_1 and a_2 in $2!$ ways and the remaining items in $4!$ ways. The total number of such permutations is

$$2 \cdot 3 \cdot 2! \cdot 4! = 288.$$

Case 2: First non-optimal match after the third item

Before a_3 arrives and is matched non-optimally, a_1 and a_2 need to be packed into separate bins and a_3 is matched non-optimally. This happens in the following cases.

Case 2a $a_1, a_2 \in \{l_1, l_2\}$ and $a_3 = m_3$.

Case 2b $a_1, a_2 \in \{l_1, l_3\}$ and $a_3 = m_2$.

Case 2c $a_1, a_2 \in \{l_2, m_1\}$ and $a_3 = m_3$.

Case 2d $a_1, a_2 \in \{l_3, m_1\}$ and $a_3 = m_2$.

Case 2e $a_1, a_2 \in \{m_2, l_3\}$ and $a_3 \in \{m_1, l_1\}$.

Note that for each case 2a-2d, we have $2! \cdot 3!$ permutations, while in case 2e, in addition, we can choose a_3 in one of two ways. Therefore, we have a total number of permutations for case 2 of

$$4 \cdot (2! \cdot 3!) + 2! \cdot 2 \cdot 3! = 48 + 24 = 72.$$

Case 3: First non-optimal match after the fourth item

Here, we must have an optimal match either after the second item (Case 3a) or after the third item (Case 3b).

Case 3a

We have $a_1, a_2 \in \{l_i, m_i\}$ for some $i \in [3]$. Now, $\{a_3, a_4\}$ can contain either the remaining two medium items, or the larger of the remaining two large items and the smaller of the remaining two medium items. The fifth and sixth item can be arranged in $2!$ ways. The total number of such permutations is

$$3 \cdot 2! \cdot 2 \cdot 2! \cdot 2! = 48.$$

Case 3b

We have an optimal match after the third item followed by a non-optimal match after the fourth item. This happens in the following cases.

- $a_1, a_2 \in \{l_3, m_1\}, a_3 = m_3, a_4 = m_2$.
- $a_1, a_2 \in \{l_1, l_3\}, a_3 = m_3, a_4 = m_2$.
- $a_1, a_2 \in \{l_2, m_1\}, a_3 \in \{m_2, l_1\}, a_4 = m_3$.
- $a_1, a_2 \in \{l_1, l_2\}, a_3 \in \{m_1, m_2\}, a_4 = m_3$.
- $a_1, a_2 \in \{l_3, m_2\}, a_3 = m_3, a_4 = \{m_1, l_1\}$.

In each of the first two cases, we get $2! \cdot 2!$ permutations. In each of the remaining three cases, we have $2! \cdot 2 \cdot 2!$ permutations, since we can choose one additional item among two elements. The total number of permutations in case 3b is thus

$$2 \cdot (2! \cdot 2!) + 3 \cdot (2! \cdot 2 \cdot 2!) = 8 + 24 = 32.$$

□

Appendix C: Lower Bound on the Asymptotic Random Order Ratio

Appendix C.1: Proof of Lemma 3

Let $\mathcal{I} = \{I \mid \Pr[I_n(F) = I] > 0\}$ be the set of possible outcomes of $I_n(F)$. We say that two lists $I_1, I_2 \in \mathcal{I}$ are *similar* ($I_1 \sim I_2$) if there is a permutation $\pi \in \mathcal{S}_n$ such that $I_1 = I_2^\pi$. Note that $\sim$ defines an equivalence relation on $\mathcal{I}$. Let $\mathcal{H}$ be a complete set of representatives of $\sim$. This way, $\mathcal{I} = \bigsqcup_{H \in \mathcal{H}} \{H^\pi \mid \pi \in \mathcal{S}_n\}$.

We will use the following two technical claims which we will prove later.

Claim 3 Let $ALG \in \{\mathcal{A}, \text{OPT}\}$. For each $H \in \mathcal{H}$, there exists $\lambda_H > 0$ such that

$$\mathbb{E}[ALG(I_n(F))] = \sum_{H \in \mathcal{H}} \left(\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} ALG(H^\pi) \right).$$

Claim 4 For any two nonnegative sequences $(a_i)_{i \in [m]}$ and $(b_i)_{i \in [m]}$, we have

$$\frac{\sum_{j=1}^m a_j}{\sum_{j=1}^m b_j} \leq \max_{j \in [m]} \frac{a_j}{b_j}.$$

Proof of Lemma 3 Using Claims 3 and 4, Lemma 3 follows from the following reasoning. Let

$$I = \arg \max_{H \in \mathcal{H}} \frac{\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} \mathcal{A}(H^\pi)}{\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} \text{OPT}(H^\pi)}.$$

Now, it holds that

$$\begin{aligned} \frac{\mathbb{E}[\mathcal{A}(I_n(F))]}{\mathbb{E}[\text{OPT}(I_n(F))]} &= \frac{\sum_{H \in \mathcal{H}} \left(\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} \mathcal{A}(H^\pi) \right)}{\sum_{H \in \mathcal{H}} \left(\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} \text{OPT}(H^\pi) \right)} && \text{Claim 3} \\ &\leq \frac{\lambda_I \cdot \sum_{\pi \in \mathcal{S}_n} \mathcal{A}(I^\pi)}{\lambda_I \cdot \sum_{\pi \in \mathcal{S}_n} \text{OPT}(I^\pi)} && \text{Claim 4} \\ &= \frac{\frac{1}{n!} \cdot \sum_{\pi \in \mathcal{S}_n} \mathcal{A}(I^\pi)}{\frac{1}{n!} \cdot \sum_{\pi \in \mathcal{S}_n} \text{OPT}(I^\pi)} \\ &= \frac{\mathbb{E}[\mathcal{A}(I^\sigma)]}{\text{OPT}(I)}. \end{aligned}$$

The second claim follows immediately from $\text{OPT}(I) \geq \sum_{i=1}^n X_i \geq cn$. $\square$

It remains to show Claims 3 and 4.

Proof of Claim 3 We first show that all lists of the same equivalence class have the same probability to be an outcome of $I_n(F) = (X_1, \dots, X_n)$.

Let $I_1 = (a_1, \dots, a_n)$ and $I_2 = (b_1, \dots, b_n)$ with $I_1 \sim I_2$. Note that both lists contain the same multiset of items. Let m be the number of different item sizes and let z_i be the multiplicity of size y_i for $i \in [m]$ in I_1 (and thus in I_2). As the random variables X_i are drawn i.i.d., we have

$$\begin{aligned} \Pr[I_n(F) = I_1] &= \prod_{i=1}^n \Pr[X_i = a_i] \\ &= \prod_{j=1}^m (\Pr[X_1 = y_j])^{z_j} \\ &= \prod_{i=1}^n \Pr[X_i = b_i] \\ &= \Pr[I_n(F) = I_2]. \end{aligned}$$

Now, consider any equivalence class $[H]$ for $H \in \mathcal{H}$. All lists $I \in [H]$ have the same probability $\Pr[I_n(F) = I] =: \lambda_H$ to be an outcome of $I_n(F)$. Further, we have $[H] = \{H^\pi \mid \pi \in \mathcal{S}_n\}$ and thus

$$\begin{aligned}
\mathbb{E}[ALG(I_n(F))] &= \sum_{I \in \mathcal{I}} ALG(I) \cdot \Pr[I_n(F) = I] \\
&= \sum_{H \in \mathcal{H}} \sum_{I \in [H]} ALG(I) \cdot \Pr[I_n(F) = I] \\
&= \sum_{H \in \mathcal{H}} \left(\lambda_H \cdot \sum_{I \in [H]} ALG(I) \right) \\
&= \sum_{H \in \mathcal{H}} \left(\lambda_H \cdot \sum_{\pi \in \mathcal{S}_n} ALG(H^\pi) \right).
\end{aligned}$$

□

Claim 4 is easy to see:

Proof of Claim 4 Define $M := \max_{j \in [m]} (a_j/b_j)$. We have

$$\sum_{j=1}^m a_j = \sum_{j=1}^m \frac{a_j}{b_j} \cdot b_j \leq \sum_{j=1}^m M b_j = M \cdot \sum_{j=1}^m b_j.$$

□

Appendix C.2: Markov Chain from Fig. 3

Lemma 7 The Markov chain from Fig. 3 converges to the stationary distribution $\omega = (\omega_A, \dots, \omega_I)$ with

$$\omega = \frac{1}{\lambda} (1, p, q + pq\vartheta, p^2, 2pq + p^2q\vartheta, q^2 + 2pq^2\vartheta, \vartheta, q\vartheta, q^2\vartheta),$$

where $\vartheta = \frac{p^3}{1-q^3}$ and $\lambda = \vartheta q(3 - q^2) + \vartheta + 3$.

Proof As the Markov chain from Fig. 3 is irreducible and aperiodic, it converges to a unique stationary distribution. Let $\omega = (\omega_A, \dots, \omega_I)$ be this distribution, then we have the following system of equations.

$$\omega_A = \omega_E + \omega_F + p\omega_G + q\omega_D \quad (Q1)$$

$$\omega_B = p\omega_A \quad (Q2)$$

$$\omega_C = q\omega_A + p\omega_H \quad (Q3)$$

$$\omega_D = p\omega_B \quad (Q4)$$

$$\omega_E = q\omega_B + p\omega_C \quad (Q5)$$

$$\omega_F = q\omega_C + p\omega_I \quad (\text{Q6})$$

$$\omega_G = p\omega_D + q\omega_I \quad (\text{Q7})$$

$$\omega_H = q\omega_G \quad (\text{Q8})$$

$$\omega_I = q\omega_H \quad (\text{Q9})$$

$$1 = \omega_A + \omega_B + \omega_C + \omega_D + \omega_E + \omega_F + \omega_G + \omega_H + \omega_I. \quad (\text{Q10})$$

We define $\vartheta := \frac{p^3}{1-q^3}$ and $\lambda := \vartheta q(3 - q^2) + \vartheta + 3$ and claim that

$$\omega = \frac{1}{\lambda} (1, p, q + pq\vartheta, p^2, 2pq + p^2q\vartheta, q^2 + 2pq^2\vartheta, \vartheta, q\vartheta, q^2\vartheta)$$

satisfies (Q1) to (Q10). In fact, the validity of equations (Q2) to (Q9) can be seen immediately. To verify (Q1) and (Q10), we first observe that for any $i \geq 0$ and $1 \leq j \leq 3$, it holds that

$$p^i q^j = \begin{cases} p^i - p^{i+1} & j = 1 \\ p^i - 2p^{i+1} + p^{i+2} & j = 2 \\ p^i - 3p^{i+1} + 3p^{i+2} - p^{i+3} & j = 3. \end{cases} \quad (5)$$

Now, we prove (Q1). We have

$$\begin{aligned} & \lambda \cdot (\omega_E + \omega_F + p\omega_G + q\omega_D) \\ &= (2pq + p^2q\vartheta) + (q^2 + 2pq^2\vartheta) + p\vartheta + qp^2 \\ &= \vartheta \cdot (p^2q + 2pq^2 + p) + 2pq + q^2 + p^2q \\ &\stackrel{\text{eq.5}}{=} \vartheta \cdot (p^2 - p^3 + 2 \cdot (p - 2p^2 + p^3) + p) \\ &\quad + 2 \cdot (p - p^2) + (1 - 2p + p^2) + (p^2 - p^3) \\ &= \vartheta(3p - 3p^2 + p^3) + 1 - p^3 \\ &= 1 \\ &= \lambda \cdot \omega_A, \end{aligned}$$

which implies (Q1). Equation (Q10) follows from the following calculation.

$$\begin{aligned}
& \lambda \cdot (\omega_A + \omega_B + \omega_C + \omega_D + \omega_E + \omega_F + \omega_G + \omega_H + \omega_I) \\
&= 1 + p + (q + pq\vartheta) + p^2 + (2pq + p^2q\vartheta) + (q^2 + 2pq^2\vartheta) + \vartheta + q\vartheta + q^2\vartheta \\
&= q\vartheta \cdot (p + p^2 + 2pq + 1 + q) + (1 + p + q + p^2 + 2pq + q^2) + \vartheta \\
&\stackrel{\text{eq.5}}{=} q\vartheta \cdot (p + p^2 + 2 \cdot (p - p^2) + 1 + 1 - p) \\
&\quad + (1 + p + 1 - p + p^2 + 2 \cdot (p - p^2) + 1 - 2p + p^2) + \vartheta \\
&= q\vartheta \cdot (3 - q^2) + 3 + \vartheta \\
&= \lambda.
\end{aligned}$$

□

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Al-Herz, A., Pothén, A.: A 2/3-approximation algorithm for vertex-weighted matching. In: Discrete Applied Mathematics (2019). ((in press))
2. Albers, S., Khan, A., Ladewig, L.: Best fit bin packing with random order revisited. In: 45th International Symposium on Mathematical Foundations of Computer Science (MFCS), pp. 7:1–7:15 (2020)
3. Balogh, J., Békési, J., Dósa, G., Epstein, L., Levin, A.: A new and improved algorithm for online bin packing. In: Proceedings of the 26th Annual European Symposium on Algorithms (ESA), LIPIcs, vol. 112, pp. 5:1–5:14 (2018)
4. Balogh, J., Békési, J., Dósa, G., Epstein, L., Levin, A.: A new lower bound for classic online bin packing. In: Approximation and Online Algorithms—17th International Workshop, (WAOA), Lecture Notes in Computer Science, vol. 11926, pp. 18–28. Springer (2019)
5. Boyar, J., Dósa, G., Epstein, L.: On the absolute approximation ratio for first fit and related results. Discret. Appl. Math. **160**(13–14), 1914–1923 (2012)
6. Christ, M.G., Favrholt, L.M., Larsen, K.S.: Online bin covering: expectations vs. guarantees. Theor. Comput. Sci **556**, 71–84 (2014)
7. Christensen, H.I., Khan, A., Pokutta, S., Tetali, P.: Approximation and online algorithms for multidimensional bin packing: a survey. Comput. Sci. Rev. **24**, 63–79 (2017)

8. Coffman Jr., E.G., Csirik, J., Galambos, G., Martello, S., Vigo, D.: Bin packing approximation algorithms: survey and classification. In: *Handbook of Combinatorial Optimization*, pp. 455–531. Springer, New York (2013)
9. Coffman, E.G., Jr., Csirik, J., Rónyai, L., Zsbán, A.: Random-order bin packing. *Discret. Appl. Math.* **156**(14), 2810–2816 (2008)
10. Coffman, E.G., Jr., Lueker, G.S.: *Probabilistic Analysis of Packing and Partitioning Algorithms*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley, Hoboken (1991)
11. Dósa, G., Sgall, J.: First fit bin packing: a tight analysis. In: *30th International Symposium on Theoretical Aspects of Computer Science (STACS), LIPIcs*, vol. 20, pp. 538–549 (2013)
12. Dósa, G., Sgall, J.: Optimal analysis of best fit bin packing. In: *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 429–441 (2014)
13. Fischer, C., Röglin, H.: Probabilistic analysis of the dual next-fit algorithm for bin covering. In: *LATIN 2016: Theoretical Informatics—12th Latin American Symposium*, pp. 469–482 (2016)
14. Fischer, C., Röglin, H.: Probabilistic analysis of online (class-constrained) bin packing and bin covering. In: *LATIN 2018: Theoretical Informatics—13th Latin American Symposium, Lecture Notes in Computer Science*, vol. 10807, pp. 461–474. Springer (2018)
15. Gamlath, B., Kapralov, M., Maggiori, A., Svensson, O., Wajc, D.: Online matching with general arrivals. In: *60th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 26–37 (2019)
16. Garey, M.R., Graham, R.L., Ullman, J.D.: Worst-case analysis of memory allocation algorithms. In: *Proceedings of the 4th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 143–150 (1972)
17. Garey, M.R., Johnson, D.S.: Strong NP-completeness results: motivation, examples, and implications. *J. ACM* **25**(3), 499–508 (1978)
18. Graham, R.L.: Bounds on multiprocessing timing anomalies. *SIAM J. Appl. Math.* **17**(2), 416–429 (1969)
19. Gupta, A., Singla, S.: Random-order models. *CoRR*, [arXiv:abs/2002.12159](https://arxiv.org/abs/2002.12159) (2020)
20. Hoberg, R., Rothvoss, T.: A logarithmic additive integrality gap for bin packing. In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2616–2625 (2017)
21. Huang, Z., Kang, N., Tang, Z.G., Wu, X., Zhang, Y., Zhu, X.: How to match when all vertices arrive online. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 17–29 (2018)
22. Huang, Z., Peng, B., Tang, Z.G., Tao, R., Wu, X., Zhang, Y.: 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)*, pp. 2875–2886 (2019)
23. Huang, Z., Tang, Z.G., Wu, X., Zhang, Y.: Online vertex-weighted bipartite matching: beating $1-1/e$ with random arrivals. *ACM Trans. Algorithms* **15**(3), 38:1–38:15 (2019)
24. Johnson, D.S.: Fast algorithms for bin packing. *J. Comput. Syst. Sci.* **8**(3), 272–314 (1974)
25. Johnson, D.S., Demers, A.J., Ullman, J.D., Garey, M.R., Graham, R.L.: Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM J. Comput.* **3**(4), 299–325 (1974)
26. Karmarkar, N., Karp, R.M.: An efficient approximation scheme for the one-dimensional bin-packing problem. In: *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 312–320 (1982)
27. Karp, R.M., Vazirani, U.V., Vazirani, V.V.: An optimal algorithm for on-line bipartite matching. In: Ortiz, H. (ed.) *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pp. 352–358 (1990)
28. Kenyon, C.: Best-fit bin-packing with random order. In: *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 359–364 (1996)
29. Lee, C.C., Lee, D.T.: A simple on-line bin-packing algorithm. *J. ACM* **32**(3), 562–572 (1985)
30. Levin, D.A., Peres, Y.: *Markov Chains and Mixing Times*, vol. 107. American Mathematical Society, Providence (2017)
31. Mahdian, M., Yan, Q.: 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)*, pp. 597–606 (2011)
32. Mehta, A.: Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.* **8**(4), 265–368 (2013)
33. Murgolo, F.D.: Anomalous behavior in bin packing algorithms. *Discret. Appl. Math.* **21**(3), 229–243 (1988)

34. Ramanan, P.V.: Average-case analysis of the smart next fit algorithm. *Inf. Process. Lett.* **31**(5), 221–225 (1989)
35. Ramanan, P.V., Brown, D.J., Lee, C.C., Lee, D.T.: On-line bin packing in linear time. *J. Algorithms* **10**(3), 305–326 (1989)
36. Seiden, S.S.: On the online bin packing problem. *J. ACM* **49**(5), 640–671 (2002)
37. Shor, P.W.: The average-case analysis of some on-line algorithms for bin packing. *Combinatorica* **6**(2), 179–200 (1986)
38. Ullman, J.: The Performance of a Memory Allocation Algorithm, vol. 47. Department of Electrical Engineering, Computer Science Laboratory, Princeton University, Princeton (1971)
39. de la Vega, W.F., Lueker, G.S.: Bin packing can be solved within $1+\epsilon$ in linear time. *Combinatorica* **1**(4), 349–355 (1981)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Susanne Albers¹ · Arindam Khan² · Leon Ladewig¹

Susanne Albers
albers@in.tum.de

Arindam Khan
arindamkhan@iisc.ac.in

¹ Department of Computer Science, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany

² Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560012, India