

Online Bin Covering with Advice

Boyar, Joan; Favrholt, Lene M.; Kamali, Shahin; Larsen, Kim S.

Published in:
Algorithmica

DOI:
10.1007/s00453-020-00728-0

Publication date:
2021

Document version:
Accepted manuscript

Citation for published version (APA):
Boyar, J., Favrholt, L. M., Kamali, S., & Larsen, K. S. (2021). Online Bin Covering with Advice. *Algorithmica*, 83(3), 795-821. <https://doi.org/10.1007/s00453-020-00728-0>

Go to publication entry in University of Southern Denmark's Research Portal

Terms of use

This work is brought to you by the University of Southern Denmark.
Unless otherwise specified it has been shared according to the terms for self-archiving.
If no other license is stated, these terms apply:

- You may download this work for personal use only.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying this open access version

If you believe that this document breaches copyright please contact us providing details and we will investigate your claim.
Please direct all enquiries to puresupport@bib.sdu.dk

Online Bin Covering with Advice

Joan Boyar · Lene M. Favrholt ·
Shahin Kamali · Kim S. Larsen *

Received: date / Accepted: date

Abstract The bin covering problem asks for covering a maximum number of bins with an online sequence of n items of different sizes in the range $(0, 1]$; a bin is said to be covered if it receives items of total size at least 1. We study this problem in the advice setting and provide asymptotically tight bounds of $\Theta(n \log \text{OPT})$ on the size of advice required to achieve optimal solutions.

The first, second, and fourth authors were supported in part by the Danish Council for Independent Research, Natural Sciences, grant DFF-1323-00247. A preliminary version of this paper appeared in the 16th International Algorithms and Data Structures Symposium (WADS), volume 11646 of Lecture Notes in Computer Science, Springer 2019.

* Corresponding author. E-mail: kslarsen@imada.sdu.dk

J. Boyar
University of Southern Denmark, Odense
Department of Mathematics and Computer Science, Campusvej 55, 5230 Odense M, Denmark
E-mail: joan@imada.sdu.dk

L. M. Favrholt
University of Southern Denmark, Odense
Department of Mathematics and Computer Science, Campusvej 55, 5230 Odense M, Denmark
E-mail: lenem@imada.sdu.dk

S. Kamali
University of Manitoba
Department of Computer Science, E2-EITC, Winnipeg, Manitoba R3T 2N2, Canada
E-mail: shahin.kamali@umanitoba.ca

K. S. Larsen
University of Southern Denmark, Odense
Department of Mathematics and Computer Science, Campusvej 55, 5230 Odense M, Denmark
Tel.: +45 6550 2328
E-mail: kslarsen@imada.sdu.dk
ORCID: 0000-0003-0560-3794

Moreover, we show that any algorithm with advice of size $o(\log \log n)$ has a competitive ratio of at most 0.5. In other words, advice of size $o(\log \log n)$ is useless for improving the competitive ratio of 0.5, attainable by an online algorithm without advice. This result highlights a difference between the bin covering and the bin packing problems in the advice model: for the bin packing problem, there are several algorithms with advice of constant size that outperform online algorithms without advice. Furthermore, we show that advice of size $O(\log \log n)$ is sufficient to achieve an asymptotic competitive ratio of $0.5\bar{3}$ which is strictly better than the best ratio 0.5 attainable by purely online algorithms. The technicalities involved in introducing and analyzing this algorithm are quite different from the existing results for the bin packing problem and confirm the different nature of these two problems.

Finally, we show that a linear number of advice bits is necessary to achieve any competitive ratio better than $15/16$ for the online bin covering problem.

Keywords Online Algorithms · Bin Covering · Advice Complexity · Competitive Analysis.

1 Introduction

In the bin covering problem [3], the input is a multi-set of items of different sizes in the range $(0, 1]$ which need to be placed into a set of bins. A bin is said to be covered if the total size of items in it is at least 1. The goal of the bin covering problem is to place items into bins so that a maximum number of bins is covered. In the online setting, items form a sequence which is revealed in a piece-by-piece manner; that is, at each given time, one item of the sequence is revealed and an online algorithm has to place the item into a bin without any information about the forthcoming items. The decisions of the algorithm are irrevocable.

Bin covering is closely related to the classic bin packing problem and is sometimes called the dual bin packing problem¹. The input to both problems is the same. In the bin packing problem, however, the goal is to place items into a minimum number of bins so that the total size of items in each bin is at most 1. Online algorithms for bin packing can naturally be extended to bin covering. For example, Next-Fit is a bin packing algorithm which keeps one “open” bin at any time: To place an incoming item x , if the size of x is smaller than the remaining capacity of the open bin, x is placed in the open bin; otherwise, the bin is closed (never used again) and a new bin is opened. Dual-Next-Fit [3] is a bin covering algorithm that behaves similarly, except that it closes the bin when the total size of items in it becomes at least 1.

¹ There is another problem, also sometimes referred to as “dual bin packing”, which asks for maximizing the number of items packed into a fixed number of bins; for the advice complexity of that dual bin packing problem, see [22, 10].

1.1 Offline algorithms

In the offline setting, the bin packing and bin covering problems are NP-hard. There is an asymptotic fully polynomial-time approximation scheme (AFP-TAS) for bin covering [18]. There are also bin packing algorithms which open $\text{OPT} + o(\text{OPT})$ bins [19, 24, 16], where OPT is the number of bins in the optimal packing. The additive term was improved in [24] and further, to $O(\log \text{OPT})$, in [16].

1.2 Online algorithms

Online algorithms are often compared under the framework of competitive analysis. Roughly speaking, the *competitive ratio* of a bin covering (respectively bin packing) algorithm is the minimum (respectively maximum) ratio between the number of bins covered (respectively opened) by the algorithm and that of an optimal offline algorithm on the same input.

Despite similarities between bin covering and bin packing, the status of these problems are different in the online setting. In the case of bin covering, it is known that no online algorithm can achieve a competitive ratio better than 0.5 [13], while bin covering algorithms such as Dual-Next-Fit [3] have the best possible competitive ratio of 0.5. Hence, we have a clear picture of the complexity of deterministic bin covering under competitive analysis. The situation is more complicated for the bin packing problem. It is known that no deterministic algorithm can achieve a competitive ratio of 1.54278 [5] while the best existing deterministic algorithm has a competitive ratio of 1.5783 [4]. Note there is a gap between the best known upper and lower bounds.

1.3 Online algorithms with advice

Advice complexity is a formalized way of measuring how much knowledge of the future is required for an online algorithm to obtain a certain level of performance, as measured by the competitive ratio. When such advice is available, algorithms with advice could lead to semi-online algorithms. Unlike related approaches such as “lookahead” [15] (in which some forthcoming items are revealed to the algorithm) and “closed bin packing” [2] (where the length of the input is revealed), *any* information can be encoded and sent to the algorithm under the advice setting. This generality means that lower bound results under the advice model also imply strong lower bound results on semi-online algorithms, where one can infer impossibility results simply from the length of an encoding of the information a semi-online algorithm is provided with. Advice complexity is also closely related to randomization; complexity bounds from advice complexity can be transferred to the randomization case and vice versa [8, 6, 21, 14].

We use the advice on tape model defined in [17, 9]: The advice is generated by a benevolent oracle with unlimited computational power. The advice is

written on a tape and the algorithm knows its meaning. This general approach has been studied for many problems (we refer the reader to a recent survey on advice complexity of online problems [11]). In particular, bin packing has been studied under the advice model [12, 23, 1].

For *bin covering with advice*, the decision of where to pack the i th item is based on the content of the advice tape and the sizes of the first i items. For any bin covering algorithm, \mathbb{A} , and any input sequence, σ , $\mathbb{A}(\sigma)$ and $\text{OPT}(\sigma)$ denote the number of bins covered by \mathbb{A} and an optimal offline algorithm, OPT , respectively, when given the input sequence σ . A bin covering algorithm, \mathbb{A} , is *c-competitive with advice complexity $s(n)$* if there exists a constant b such that, for all n and for all input sequences σ of length at most n , there exists some advice Φ such that $\mathbb{A}(\sigma) \geq c \cdot \text{OPT}(\sigma) - b$ and at most $s(n)$ bits of Φ are accessed by the algorithm. If $c = 1$ and $b = 0$, the algorithm is *optimal*. For a given algorithm, \mathbb{A} , with a given advice complexity, $s(n)$, the *competitive ratio* is $\sup\{c \mid \mathbb{A} \text{ is } c\text{-competitive}\}$. Thus, competitive ratios are fractions, with 1 being the best possible. For the *asymptotic* competitive ratio, the b term is allowed to be non-constant, as long as it is $o(\text{OPT})$. Thus, the *asymptotic competitive ratio* of \mathbb{A} is $\limsup_{\text{OPT}(\sigma) \rightarrow \infty} \{\mathbb{A}(\sigma) / \text{OPT}(\sigma)\}$.

Note that bin covering is a maximization problem. For minimization problems, such as bin packing, the competitive ratio is defined analogously, except that the inequality is $\mathbb{A}(\sigma) \leq c \cdot \text{OPT}(\sigma) + b$. Similarly, the competitive ratio is the *infimum* over all c such that algorithm is c -competitive, giving values at least 1.

For bin packing, it is known that $\Theta(n \log(\text{OPT}))$ advice bits are necessary and sufficient to produce optimal solutions [12], but a constant number of advice bits are sufficient to obtain a competitive ratio close to 1.47, beating the best possible online algorithm without advice [1]. Furthermore, $2n + o(n)$ advice bits suffice to get arbitrarily close to a competitive ratio of $4/3$ [12], and getting below 1.17 requires at least a linear number of bits [21]. In [23], $(1 + \varepsilon)$ -competitive online algorithms using $O(n \cdot \frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ advice bits are designed based on round and group techniques known from offline algorithms.

1.4 Contributions

In this article, we provide the first results with respect to the advice complexity of the bin covering problem.

To obtain an optimal result, advice essentially corresponding to an encoding of an entire optimal solution is necessary and sufficient. Not surprisingly, this follows from a similar proof for bin packing, since for both problems, bins filled to size one in an optimal solution are at the core of the proof.

However, unlike the bin packing problem, advice of constant size cannot improve the competitive ratio of algorithms. We establish this result by showing that any algorithm with advice of size $o(\log \log n)$ has a competitive ratio of at most 0.5, which is the optimal competitive ratio of online algorithms without advice.

We prove a tight result that advice of size $O(\log \log n)$ suffices to achieve a competitive ratio arbitrarily close to $0.5\bar{3}$. Some techniques that we develop for this result are quite different from the existing results for bin packing and are likely helpful for future analysis of bin covering with advice. The idea is to let the advice communicate the number of bins of certain “types” in an optimal packing. However, to get down to $O(\log \log n)$ bits of advice, only approximate values are given. This idea is similar to that for bin packing in [1], except that in [1] only a constant number of bits are used, approximating the ratio of the number of bins in two different sets. Another difference between our algorithm and previous solutions for bin packing is that, we employ the “dual-worst-fit” scheme for placing certain items for bin covering. In contrast, the worst-fit algorithm for bin packing has not been used previously and does not seem to be helpful.

Finally, using a reduction from the binary string guessing problem [7], we show that advice of linear size is necessary to achieve any competitive ratio larger than $15/16$. This is similar to, but more intricate, than the corresponding result for bin packing.

1.5 Techniques

We provide an intuitive explanation of the difference between bin packing and bin covering under the advice model, and use this explanation to describe our techniques in designing algorithms and impossibility results for the bin covering problem with advice. Online bin packing is relatively “easy” when items are relatively large (close to 1) or small (close to 0). Online algorithms can place large and small items separately, and this gives relatively good competitive ratios because an optimal algorithm has to open a bin for every large item and the online algorithm can fill any bin almost completely with small items. In contrast, in the bin packing problem, the “tricky” items are those that are close to $1/2$ (a bit more or less than $1/2$), and other items can be handled without wasting too much space. For inputs formed by the tricky items, a bin packing algorithm acts like a “matching algorithm”, where items smaller than $1/2$ can match with themselves or some items larger than $1/2$. Advice can help by encoding the number of items slightly larger than $1/2$. This is consistent with reserving some space for “critical” items, which is the main technique used in some results for bin packing [12, 1]. The resulting bins reflect the matching of large and small items or small items among themselves. It turns out that if we know the ratio between these two “types” of bins (approximated by a constant number of bits), we can do better than any deterministic online algorithm (see [1] for details). For bin covering, however, the tricky items are those that are either very large (close to 1) or very small (close to 0). Inputs formed by such tricky items are used in [13] to derive a lower bound on the competitive ratio of purely online algorithms. We also use them in Section 3 to derive an impossibility result when the advice size is $o(\log \log n)$. Similarly to bin packing, for inputs formed by tricky items, a bin

covering algorithm becomes a “matching algorithm”. The matching process is albeit a bit harder than it is for bin packing. This is because, unlike bin packing where matching two large items is not possible (as they do not fit in the bin), in bin covering, matching two large items is possible and sometimes necessary. To be a bit more precise, in bin covering, we cannot afford having a bin with an unmatched large item; such a bin will be almost full, but not covered, which implies that the item placed in the bin is wasted. When designing algorithms for bin covering in Section 4, we exploit advice in order to form packings that ensure that large items are always matched with other items, preferably with a set of small items, and if not possible with other large items. This ensures that items are not wasted in our packings, that is, all bins that receive items will be covered (except potentially $o(\text{OPT})$ of them). This family of packings is formally defined in Section 4 as “ (α, ε) -desirable packings”, in which large items are matched with each other, except for a fraction α of them that are placed in bins that are covered by small items. The value of ε is a parameter determined by the accuracy of the encodings of the “approximate values” of a few numbers passed to the algorithm in the form of advice.

In Section 5, we use a reduction from the binary separation problem (closely related to the binary string guessing problem) to derive an impossibility result when the advice size is sub-linear. While reductions of this form have been used for the bin packing problem, there are major differences between these reductions. These differences are again rooted in the fact that hard sequences for bin covering are formed by items that are close to 0 or 1, and that there are more possibilities for placing items in a bin covering instance compared to bin packing.

Throughout the paper, we let the *level* of a bin denote the total size of items packed in that bin at the current time.

2 Optimal covering and advice

It is not hard to see that advice of size $O(n \log(\text{OPT}(\sigma)))$ is sufficient to achieve an optimal covering for an input σ of length n ; note that $\text{OPT}(\sigma)$ denotes the number of bins in an optimal covering of σ . Provided with $O(\log(\text{OPT}(\sigma)))$ bits of advice for each item, the offline oracle can indicate in which bin the item is placed in the optimal packing. Provided with this advice, the online algorithm just needs to pack each item in the bin indicated by the advice. Clearly, the size of the advice is $O(n \log(\text{OPT}(\sigma)))$ and the outcome is an optimal packing. Note that it is always assumed that the oracle that generates the advice has unbounded computational power. However, if the time complexity of the oracle is a concern, we can use the AFPTAS of [18] to generate an almost-optimal packing and encode it in the advice. Similarly, if the input is assumed to have only m distinct known sizes, one can encode the entire request sequence, specifying for each distinct size how many of that size occur in the sequence. This only requires $O(m \log(n))$ bits of advice. The following theorem shows that the above naive solutions are asymptotically tight.

Theorem 1 *For online bin covering on sequences σ of length n , advice of size $\Theta(n \log \text{OPT}(\sigma))$ is required and sufficient to achieve an optimal solution, assuming $2 \text{OPT}(\sigma) \leq (1 - \varepsilon)n$ for some positive value of ε . When the input is formed by n items with $m \in o(n)$ distinct, known item sizes, advice of size $\Theta(m \log n)$ is required and sufficient to achieve an optimal solution.*

Proof The lower bounds follow immediately from the corresponding results for bin packing [12, Theorems 1, 3]. Since the optimal result in those proofs have all bins filled to level 1, any non-optimal bin packing would also lead to a non-optimal bin covering. \square

3 Advice of size $o(\log \log n)$ is not helpful

In this section, we show that advice of size $o(\log \log n)$ does not help for improving the competitive ratio of bin covering algorithms. This result is in contrast to bin packing where advice of constant size can improve the competitive ratio. Our lower bound sequence is similar to the one in [13], where the authors proved a lower bound on the competitive ratio of purely online algorithms.

Theorem 2 *There is no algorithm with advice of size $o(\log \log n)$ and competitive ratio better than $1/2$.*

Proof Consider a family of sequences formed as follows:

$$\sigma_j = \underbrace{(\varepsilon, \varepsilon, \dots, \varepsilon)}_{n \text{ items}}, \underbrace{(1 - j\varepsilon, 1 - j\varepsilon, \dots, 1 - j\varepsilon)}_{\lfloor n/j \rfloor \text{ items}}$$

Here, j ranges from 1 through $\lfloor \sqrt{n} \rfloor$, giving rise to $\lfloor \sqrt{n} \rfloor$ sequences in the family. All sequences start with the same prefix of n items of size ε . We assume that $\varepsilon < \frac{1}{2n}$ to ensure that, even if all these items are placed in the same bin, the level of that bin is still less than $1/2$. Note that the suffix, formed by items of size $1 - j\varepsilon$, has length $O(n)$, so the length of any of the sequences is $\Theta(n)$.

Clearly, for packing σ_j , an optimal algorithm places j items of size ε in each bin and covers $\lfloor n/j \rfloor$ bins. So we have $\text{OPT}(\sigma_j) = \lfloor n/j \rfloor$.

The proof is by contradiction, so assume there is an algorithm, \mathbb{A} , using $o(\log \log n)$ advice bits and having competitive ratio $1/2 + \mu$ for some constant $0 < \mu \leq \frac{1}{2}$. Then, there exists a fixed constant d such that for any sequence σ_j , we have

$$\mathbb{A}(\sigma_j) \geq (1/2 + \mu) \text{OPT}(\sigma_j) - d = \frac{1}{2} \left\lfloor \frac{n}{j} \right\rfloor + \mu \left\lfloor \frac{n}{j} \right\rfloor - d \quad (1)$$

We say two sequences belong to the same *sub-family* if they receive the same advice string. Since the advice has size $o(\log \log n)$, there are $o(\log n)$ sub-families. Let $\sigma_{a_1}, \dots, \sigma_{a_w}$ be the sequences in one sub-family. Since the advice and the first n items (of size ε) are the same for any two members of this sub-family, \mathbb{A} will place these n items identically. Let m_i denote the

number of bins receiving at least i items in such a placement. So, we have $\sum_{i=1}^n m_i = n$ (a bin with x items is counted x times, since it receives at least one item, at least two items, etc.). Moreover, for any σ_j , we have

$$\mathbb{A}(\sigma_j) \leq m_j + (\lfloor n/j \rfloor - m_j)/2 = \frac{1}{2} \left\lfloor \frac{n}{j} \right\rfloor + \frac{m_j}{2} \quad (2)$$

This follows since any bin with at least j items of size ε can be covered using only one item of size $1 - j\varepsilon$, while the other bins require two such items. From Equations (1) and (2), we get $\mu \left\lfloor \frac{n}{j} \right\rfloor \leq \frac{m_j}{2} + d$. Summing over $j \in \{a_1, \dots, a_w\}$, we get that

$$\mu n \left(\frac{1}{a_1} + \frac{1}{a_2} + \dots + \frac{1}{a_w} \right) - \mu w \leq \frac{1}{2} (m_{a_1} + m_{a_2} + \dots + m_{a_w}) + wd,$$

where μw is subtracted to compensate for removing the floor. Since $\frac{1}{2} (m_{a_1} + m_{a_2} + \dots + m_{a_w}) + (d + \mu)w \leq \frac{1}{2} \cdot \sum_{i=1}^n m_i + (d + \frac{1}{2})n = (d + 1)n$, we have

$$\frac{1}{a_1} + \frac{1}{a_2} + \dots + \frac{1}{a_w} \in O(1)$$

Summing the left-hand side over all families, we include every sequence and therefore every fraction, $\frac{1}{i}$, once and obtain $\sum_{i=1}^{\lfloor \sqrt{n} \rfloor} \frac{1}{i}$. Since there are $o(\log n)$ sub-families, and the contribution to the sum of fractions from each sub-family has been proven constant, it follows that $\sum_{i=1}^{\lfloor \sqrt{n} \rfloor} \frac{1}{i} \in o(\log n)$. This is a contradiction, since the Harmonic number $\sum_{i=1}^{\lfloor \sqrt{n} \rfloor} \frac{1}{i} \in \Theta(\log \sqrt{n}) = \Theta(\log n)$.

Thus, our initial assumption is wrong and with advice of size $o(\log \log n)$, no algorithm with competitive ratio strictly better than $1/2$ can exist.

For clarity in the exposition, we have ignored the issue of the sequences in the family having different lengths, all of which are larger than n . However, all sequences have length at most $2n$, so we can take that to be our n , and pad all sequences with items small enough that even the sum of them cannot fill up the missing space in any bin. For the asymptotic result proven here, changing n by a factor of at most 2 is immaterial. \square

4 An algorithm with advice of size $O(\log \log n)$

In this section, we show that advice of size $O(\log \log n)$ is sufficient to achieve an asymptotic competitive ratio of $8/15 = 0.5\bar{3}$.

Throughout this section, we call an item *small* if it has size less than $1/2$ and *large* otherwise.

Consider a packing of the input sequence σ . We partition the bins in this packing into three groups. A *large-small* (LS) bin includes one large item and some small items, a *large-large* (LL) bin includes only two large items, and a *small* (S) bin includes only small items. We use m_S , m_{LS} , and m_{LL} ,

respectively, to denote the number of S-bins, LS-bins, and LL-bins in the optimal packing. For $m_{LS} \geq 1$, we let $\beta \geq 1$ satisfy $m_{LS} + m_{LL} = \beta m_{LS}$. See Table 1 for a summary of notation used in this section.

4.1 A simple algorithm for the case $m_{LS} = 0$ or $\beta \geq 15/14$

The following lemma shows that sequences, where there are not too many LS-bins compared to LL-bins in an optimal packing, are “easy” instances. More specifically, when there are at most 14 times as many LS-bins as LL-bins, there is a simple $8/15$ -competitive algorithm.

Lemma 1 *When $m_{LS} = 0$ or $\beta \geq 15/14$, there is an online bin covering algorithm with competitive ratio at least $8/15$.*

Proof Consider a simple algorithm, \mathbb{A} , that places large and small items separately. Each pair of large items cover one bin and small items are placed using the Dual-Next-Fit strategy, that is, they are placed in the same bin until the bin is covered (and then a new bin is started). Let S denote the total size of small items. Note that the number of large items is $m_{LS} + 2m_{LL}$.

Table 1: Notation used in Section 4

Notation	Meaning
n	The length of the input.
m_{LS}	The number of LS-bins in the optimal packing.
m_{LL}	The number of LL-bins in the optimal packing.
m_S	The number of S-bins in the optimal packing.
β	The value of $\frac{m_{LS} + m_{LL}}{m_{LS}}$. The algorithm behaves differently when $\beta \geq 15/14$ compared to when $\beta < 15/14$.
α	A parameter of the algorithm when $\beta < 15/14$. Approximately $\lfloor \alpha m_{LS} \rfloor$ of covered bins include exactly one large item.
ε	$\varepsilon \in O\left(\frac{1}{\log n}\right)$.
n_G	$n_G = \lfloor \alpha m_{LS} / 3 \rfloor$ is the number of good items.
d	$d = 1 - s_G$, where s_G is the size of the smallest good item.
s_B	The size of the n_B 'th smallest black item.
m_R	The number of reserved bins.
m_B	Among the reserved bins, m_B bins are black bins.
m_W	Among the reserved bins, m_W bins are white bins.

The number of bins covered by \mathbb{A} is at least $\lfloor (m_{\text{LS}} + 2m_{\text{LL}})/2 \rfloor + \lfloor 2S/3 \rfloor \geq (m_{\text{LS}} + 2m_{\text{LL}})/2 + 2S/3 - 2$. The number of bins covered by OPT is at most $m_{\text{LS}} + m_{\text{LL}} + \lfloor S \rfloor$. Thus, for any input sequence, σ ,

$$\mathbb{A}(\sigma) \geq \frac{(m_{\text{LS}} + 2m_{\text{LL}})/2 + 2S/3}{m_{\text{LS}} + m_{\text{LL}} + S} \cdot \text{OPT}(\sigma) - 2$$

For $m_{\text{LS}} = 0$, this proves a ratio of at least $2/3$ which is larger than $8/15$. For $\beta \geq 15/14$, this gives a competitive ratio of at least $\min\{2/3, \frac{2\beta-1}{2\beta}\}$, since

$$\frac{(m_{\text{LS}} + 2m_{\text{LL}})/2}{m_{\text{LS}} + m_{\text{LL}}} = \frac{2m_{\text{LS}} + 2m_{\text{LL}} - m_{\text{LS}}}{2m_{\text{LS}} + 2m_{\text{LL}}} = \frac{2 \frac{m_{\text{LS}} + m_{\text{LL}}}{m_{\text{LS}}} - 1}{2 \frac{m_{\text{LS}} + m_{\text{LL}}}{m_{\text{LS}}}} = \frac{2\beta - 1}{2\beta}.$$

For $\beta \geq 15/14$, $(2\beta - 1)/(2\beta) \geq 8/15$. This proves the lemma. \square

4.2 A more complicated algorithm for the case $m_{\text{LS}} \geq 1$ and $\beta < 15/14$

In what follows, we define (α, ε) -desirable packings, which act as reference packings for our algorithm. Here α and ε are parameters of the algorithm that we will introduce later. Note that in an optimal packing, there are m_{LS} bins covered by one large item and some small items. Consider a covering in which large and small items are placed in separate bins. Clearly, there are 0 bins that are covered by one large and some small items and, in the worst case, the level of all bins covered by small bins is almost 1.5 (when all small items have size a bit less than 0.5). Intuitively speaking, an (α, ε) -desirable packing is an intermediate solution between this packing and the optimal packing, in which there are almost αm_{LS} bins covered by one large item and some small items.

For the following definition, it may be helpful to confer with Figure 1.

Definition 1 A packing is (α, ε) -desirable, where α is a real number in the range $(0, 1]$ and ε is a (small) positive real, if and only if all the following hold:

- I The packing has at least $\lfloor \alpha m_{\text{LS}} \rfloor$ LS-bins. All LS-bins, except possibly $O(\varepsilon m_{\text{LS}})$ of them, are covered.
- II The large items that are not packed in LS-bins appear in pairs, with each pair covering one bin (except one item when there are an odd number of such large items).
- III The packing has at least $\frac{2m_{\text{S}}}{3} - O(\varepsilon m_{\text{LS}})$ covered S-bins.

Lemma 2 For any input sequence, σ , the number of bins covered in an (α, ε) -desirable packing is at least $\min\{\frac{\alpha+2\beta-1}{2\beta}, \frac{2}{3}\} \cdot \text{OPT}(\sigma) - O(\varepsilon m_{\text{LS}})$.

Proof The number of bins covered in the optimal packing is $\text{OPT}(\sigma) = m_{\text{LL}} + m_{\text{LS}} + m_{\text{S}} = \beta m_{\text{LS}} + m_{\text{S}}$.

Now consider the (α, ε) -desirable packing and let x_{LS} be the number of LS-bins in the packing. Then by Definition 1, $x_{\text{LS}} \geq \lfloor \alpha m_{\text{LS}} \rfloor$, and the number of covered LS-bins is $x_{\text{LS}} - O(\varepsilon m_{\text{LS}})$. The number of covered LL-bins is $\lfloor (2m_{\text{LL}} +$




Fig. 1: (top) an optimal packing with $m_{LS} = 4$, $m_{LL} = 2$, and $m_S = 3$; (bottom) an (α, ε) -desirable packing with $\alpha = 1/2$ and $\varepsilon = 1/64$.

$m_{LS} - x_{LS})/2] = \lfloor ((2\beta - 1)m_{LS} - x_{LS})/2 \rfloor$. By Definition 1, the number of covered S-bins is at least $\frac{2m_S}{3} - O(\varepsilon m_{LS})$. Hence, the number of bins covered in the (α, ε) -desirable packing of σ will be at least

$$\begin{aligned}
& x_{LS} - O(\varepsilon m_{LS}) + \left\lfloor \frac{(2\beta - 1)m_{LS} - x_{LS}}{2} \right\rfloor + \frac{2m_S}{3} - O(\varepsilon m_{LS}) \\
&= \frac{1}{2}x_{LS} + \frac{(2\beta - 1)m_{LS}}{2} + \frac{2m_S}{3} - O(\varepsilon m_{LS}) \\
&\geq \frac{\alpha m_{LS}}{2} + \frac{(2\beta - 1)m_{LS}}{2} + \frac{2m_S}{3} - O(\varepsilon m_{LS}) \\
&= \frac{\alpha m_{LS}/2 + (2\beta - 1)m_{LS}/2 + 2m_S/3}{\beta m_{LS} + m_S} \cdot \text{OPT}(\sigma) - O(\varepsilon m_{LS}) \\
&\geq \min \left\{ \frac{\alpha + 2\beta - 1}{2\beta}, \frac{2}{3} \right\} \cdot \text{OPT}(\sigma) - O(\varepsilon m_{LS}).
\end{aligned}$$

In the first equality, we simply remove terms that are anyway absorbed in $O(\varepsilon m_{\text{LS}})$; similar for the first inequality, using the bound $x_{\text{LS}} \geq \lfloor \alpha m_{\text{LS}} \rfloor$, where the constant coming from lifting the floor is also absorbed. In the last equality, we have gathered the first three terms and multiplied with the identity $\text{OPT}(\sigma)/(\beta m_{\text{LS}} + m_{\text{S}})$. Finally, for the last inequality, we have used the arithmetic inequality, $\frac{a+b}{c+d} \geq \min(\frac{a}{c}, \frac{b}{d})$, to collect the m_{LS} and m_{S} terms, respectively. \square

In the remainder of this section, we describe an algorithm that achieves an (α, ε) -desirable packing for certain values of α and ε . Here, ε is a parameter determined by the accuracy of the encodings of approximate values of a few numbers passed to the algorithm. Before describing these numbers, we explain how the approximate encodings work.

Our algorithm requires approximate encoding of integers (associated with the number of certain bins), as well as positive values smaller than 1 (associated with size of items or levels of bins). Given a positive integer x , we can write the length of the binary encoding of x in $O(\log \log x)$ bits, using self-delimited encoding as in [20]. The approximate value of x will be represented by the binary encoding of the length of x , plus the k most significant bits of x after the high-order 1, with $k = \lceil \log \log x \rceil$. Thus we use $O(\log \log n)$ bits in total for encoding numbers no larger than n . Setting the unknown lower order bits to zero gives an approximation to x which we denote by x^- . We can bound x^- as follows: If $x^- = y \cdot 2^\ell$ for some y represented by $k+1$ bits, where the high-order bit is a one, then $2^k \leq y < 2^{k+1}$. Given x^- , the largest x could be is $y \cdot 2^\ell + (2^\ell - 1)$. Thus, $(1 - \frac{1}{2^k})x < x^- \leq x$. Similarly, setting the unknown lower order bits to one gives an approximation to x which we denote by x^+ . We can bound x^+ as follows: $x \leq x^+ < (1 + \frac{1}{2^k})x$. For a real-valued, positive $x < 1$, we consider its binary representation as the sum of powers of $1/2$. To represent x^- , we write the first $k = O(\log \log n)$ powers of $1/2$, ignoring any unknown remaining bits. If these unknown bits are all 1, we get x^+ , the value of which is no more than $x + \frac{1}{2^k}$.

We will let $\varepsilon = \frac{1}{2^k} \leq \frac{1}{\log n}$.

In the remaining more technical part of the section, it may be beneficial to consider that if we had had $O(\log n)$ bits of advice instead of $O(\log \log n)$, many arguments would be simplified, and it could be helpful on a first reading to ignore multiplicative terms such as $1 - \varepsilon$ that are there because we know only approximate as opposed to exact values of the parameters we receive information about in the advice.

4.2.1 Outlining the algorithm

We call the $n_{\text{G}} = \lfloor m_{\text{LS}}/3 \rfloor$ largest items in the input sequence *good items* and let $d = 1 - s_{\text{G}}$, where s_{G} denotes the size of the smallest good item. Thus, packing a good item in a bin with level at least d ensures that the bin is covered.

Our algorithm aims to place a subset of small items into n_G *reserved bins* in a way that the level of each reserved bin is at least d , while the remaining small items will cover almost $2m_S/3$ separate bins. As we will describe later, some of the reserved bins will receive an additional good item and hence become covered.

At first glance, placing small items to provide the above guarantee does not seem too hard: an optimal packing covers m_S small bins while we need to cover roughly $2m_S/3$ bins, and the remaining small items have total size at least $2m_{LS}/3 \times d$ (because $2/3$ of OPT's LS-bins have large items of size at most $1-d$) while we want to cover a level $n_G \times d \approx (m_{LS}/3) \times d$ (in a sense, the total volume of available small items is at least twice the required volume for achieving the desired level). Ideally, if we can balance the level of small items in the reserved bins, then all levels are close to the average level (roughly $2d$). This can be done via the Dual-Worst-Fit algorithm that places each item in the bin that currently has the minimum level among a predefined number of bins (here n_G , the number of reserved bins). Unfortunately however, the discrepancy between the sizes of the small items can create situations where simply using Dual-Worst-Fit is not sufficient to achieve the desired level for all reserved bins. Consequently, the algorithm is more delicate as we will describe in what follows.

The algorithm receives the approximation d^+ as advice. Small items of size at least d^+ are called *black items*, and items smaller than d^+ are called *white items*. Let m_{LSB} denote the number of LS-bins in the optimal packing that contain a black item.

The following offline scheme illustrates the idea behind our algorithm. See Figure 2 for an illustration.

- Open n_G bins called *reserved bins*.
- Place the $n_B = \min\{m_{LSB}, n_G\}$ smallest black items in n_B of these bins, one item per bin. These bins are called *black bins*.
- Place white items of total size at least d , but less than $d + d^+$, in each of the remaining $m_W = n_G - n_B$ reserved bins. These bins are called *white bins*. This is done using Dual-Worst-Fit.
- The remaining small items are packed in S-bins using Dual-Next-Fit.
- Place a large item in each of the reserved bins, ensuring that at least αm_{LS} of them are good items. Place an additional large item in each reserved bin without a good item. Place the remaining large items pairwise in new bins.

The above algorithm is strictly better than 0.5-competitive. Here is an informal argument for this claim. The total size of small items placed in S-bins cannot be much smaller than the total size of small items placed in S-bins in an optimal packing. This is due to the following:

- The smallest black items are placed in the reserved bins; the total size of these items is no more than the total size of black items placed with a good item in the LS-bins of OPT.




Fig. 2: (top) An optimal packing with 12 LS-bins ($m_{LS} = 12$ and hence $n_G = 12/3 = 4$). The four good items are colored dark gray. The smallest good item has size 0.77 and hence $d = 1 - 0.77 = 0.23$. The approximate value of d is encoded in 6 bits, i.e., $d^+ = 1/8 + 1/16 + 1/32 + 1/64 = 0.234375$. Black items are those no smaller than 0.234375 and they are colored dark green (they are 0.24, 0.25, and 0.26). There are $m_{LSB} = 2$ black items in LS-bins. (bottom) A packing that covers more bins than half of the optimal packing. There are $n_G = 4$ reserved bins. We have $n_B = \min\{m_{LSB} = 2, n_G = 4\} = 2$. So, the two smallest black items (0.24 and 0.25) are placed in the reserved bins (black bins). The other 2 reserved bins (white bins) are covered with small items up to an approximate level d using the Dual-Worst-Fit algorithm. The remaining small items are placed using Dual-Next-Fit in separate S-bins. Assuming that $\alpha = 1/6$, there are $\alpha m_{LS} = 2$ reserved bins that receive good items. The remaining large items are all paired in their bins.

- Each white bin receives white items of total size approximately d (assuming $n_G > n_B$, i.e., a white bin exists). There are enough white items in the input sequence to ensure this, because OPT covers $m_{LS} - m_{LSB}$ LS-bins with

white items together with a large item. By the assumption that $n_G > n_B$, $m_{LSB} = n_B$. Thus, OPT has $m_{LS} - n_B$ LS-bins without a black item. In at most $m_{LS}/3$ of these bins, the total size of white items is less than d (these are bins with a good item). The total size of white items in any of the remaining $2m_{LS}/3 - n_B$ bins is at least d , that is, the total size of white items in all bins is at least $(2m_{LS}/3 - n_B) \times d \geq (m_{LS}/3 - n_B) \times 2d$. In the above scheme, white items are placed in $n_G - n_B \approx m_{LS}/3 - n_B$ bins, that is, on average, a size of at least $2d$ is available for each white bin. Since Dual-Worst-Fit tends to balance the load of white items among the bins, the level of each bin will not be less than d , which is half of the average available size (it will be formally proved later).

For bins where we use Dual-Next-Fit, each bin has a level less than $3/2$. Thus, the number of the non-reserved bins covered by small items is roughly lower-bounded by $2m_S/3$. Reserved bins that received a good item are all covered; so, at least αm_{LS} bins are covered by a large and some small items. Finally, other large items are placed pairwise in new bins. We conclude that the resulting packing is (α, ε) -desirable and hence (for $\alpha > 0$ and $\varepsilon < \frac{1}{4}$) strictly better than 0.5-competitive by Lemma 2.

Unfortunately however, using only $O(\log \log n)$ bits of advice, the oracle cannot give the exact values of d and m_{LSB} nor identify the m_{LSB} smallest black items. Thus, we make a number of adjustments to the above algorithm as described in the two following subsections, resulting in a value of α smaller than $1/15$.

4.2.2 Placing the small items

Let s_B denote the size of the n_B 'th smallest black item in the input sequence. Let e_B denote the number of items among the n_B smallest black items which have sizes in the range $(s_B^-, s_B^+]$. Let x_B denote the number of black items no larger than s_B^- and let $m_B = x_B + e_B$. The numbers s_B^- , e_B^- , and m_B^- are given as advice, and the algorithm opens m_B^- reserved black bins. This ensures that the number of black items used in the algorithm's reserved bins is close to n_B , but not larger than n_B . Each black item of size at most s_B^- is packed in a black bin. These black items are among the n_B smallest black items. Each black item with a size in $(s_B^-, s_B^+]$ is also packed in a black bin, if less than e_B^- black items in the size range $(s_B^-, s_B^+]$ have already been placed in black bins. These black items might not be among the n_B smallest black items but are at most a factor $(1 + \varepsilon)$ larger. So, m_B^- reserved bins are covered by black items that are not much larger than the smallest n_B black items.

Let $m_W = n_G - n_B$. The number m_W^- is given as advice, and the algorithm opens m_W^- reserved white bins. As long as the total size of white items received so far is at most $2d^+ m_W^-$, white items are packed in the white bins using Dual-Worst-Fit. The remaining small items are packed in S-bins using Dual-Next-Fit.

We let $m_R = m_B^- + m_W^-$ denote the total number of reserved bins. Note that

$$\begin{aligned}
m_R &= m_B^- + m_W^- \\
&= (x_B + e_B^-)^- + (n_G - n_B)^- \\
&> (1 - \varepsilon)(x_B + e_B^-) + (1 - \varepsilon)(n_G - n_B) \\
&> (1 - \varepsilon)^2 n_B + (1 - \varepsilon)(n_G - n_B) \\
&> (1 - \varepsilon)^2 n_G
\end{aligned} \tag{3}$$

Lemma 3 *A reserved bin that includes any good item will be covered in the final packing of the algorithm.*

Proof Since any good item has size at least $1 - d$, we just need to prove that each reserved bin receives small items of total size at least d . Since $d^+ \geq d$, the claim is true for all black bins. Thus, we only need to consider the white bins. For each white bin, we let the *white level* of the bin denote the total size of white items packed in the bin.

Recall that Dual-Worst-Fit places each item in a bin with minimum white level. This means that, if a white item is ever placed in a bin that already has a white level of at least d , all white bins have a white level of at least d . Assume for the sake of contradiction that no white item is placed in a bin that already has a white level of at least d .

The algorithm stops packing white items in the white bins when

- (a) we have reached the upper bound of $2d^+ m_W^-$ total size, or
- (b) there are no more white items in the sequence,

whichever happens first. In both cases, we arrive at a contradiction.

We first consider (a), i.e., we assume that white items of total size at least $2d^+ m_W^-$ are packed in the white bins. By the assumption that no white item is placed in a bin that already has a white level of at least d , the final white level of each of the m_W^- white bins is less than $d + d^+$, a contradiction.

Now, consider (b), i.e., assume that all white items of the input sequence are packed in white bins. Assume that exactly ℓ white items packed in white bins are larger than d , and let $d + e_1, \dots, d + e_\ell$ be the sizes of these items. By the assumption that no white item is placed in a bin with a white level of at least d , these ℓ items are packed in bins with a final white level less than $2d + e_i$, $1 \leq i \leq \ell$, respectively. By the same assumption, the remaining $m_W^- - \ell$ white bins each have a white level less than $2d$. Thus, the total size of white items is

$$S_W < (m_W^- - \ell) \cdot 2d + \sum_{i=1}^{\ell} (2d + e_i) = m_W^- \cdot 2d + \sum_{i=1}^{\ell} e_i \tag{4}$$

OPT covers m_{LS} bins, each with a large item and some small items. Out of the m_{LS} large items used for these bins, fewer than n_G items are larger than $1 - d$ (by the definition of d). Thus, at least $m_{LS} - n_G \geq 2n_G$ of these large

items have size at most $1 - d$, and at most n_B of them are combined with a black item in the optimal packing, as $n_B \geq m_{LSB}$. Therefore OPT has at least $2n_G - n_B \geq 2m_{\bar{W}}$ bins, each with a total size of white items of at least d . Hence, the total size of white items is

$$S_W \geq (2m_{\bar{W}} - \ell) \cdot d + \sum_{i=1}^{\ell} (d + e_i) = 2m_{\bar{W}} \cdot d + \sum_{i=1}^{\ell} e_i,$$

contradicting Equation (4). \square

We now prove that the algorithm fulfills Property III of an (α, ε) -desirable packing:

Lemma 4 *The algorithm covers at least $\frac{2}{3}m_S - O(\varepsilon m_{LS})$ S-bins.*

Proof We first show that the total size of small items packed in S-bins in the algorithm's packing is at least $m_S - O(\varepsilon m_{LS})$. The total size of items that OPT packs in S-bins is at least m_S . Thus, since all small items that are not packed in reserved bins are packed in S-bins, it is sufficient to argue that the total size of small items packed in the reserved bins is at most a factor of $O(1 + \varepsilon)$ larger than the total size of the small items packed in the LS-bins of OPT.

As argued in the proof of Lemma 3, OPT packs white items of total size at least $2dm_{\bar{W}}$ in its LS-bins. Since our algorithm stops packing white items after total size at least $2d^+m_{\bar{W}}$ has been reached, it packs white items of less than $2d^+m_{\bar{W}} + d^+$ total size in its reserved bins. Thus, the algorithm uses less than

$$\begin{aligned} 2d^+m_{\bar{W}} + d^+ - 2dm_{\bar{W}} &< 2(1 + \varepsilon)dm_{\bar{W}} - 2dm_{\bar{W}} + d^+ \\ &= 2\varepsilon dm_{\bar{W}} + d^+ \\ &< 2\varepsilon m_{\bar{W}} + d^+ \\ &\leq 2\varepsilon m_{LS} + d^+ \end{aligned}$$

more total size of white items for its reserved bins as does OPT for its LS-bins.

The algorithm packs no more black items in its reserved bins than does OPT in its LS-bins. Furthermore, these items are not much larger than those OPT uses in LS-bins since at most e_B^- black items in the size interval $(s_B^-, s_B^+]$ are used in those reserved bins. Thus, since the black items no larger than s_B^- are the smallest black items in the input, the total size of black items packed in the reserved bins is at most $1 + \varepsilon$ times the total size of items packed in OPT's LS-bins.

This proves that the algorithm packs small items of total size at least $m_S - O(\varepsilon m_{LS})$ in its S-bins. Since no items larger than $1/2$ are packed in the S-bins, no bin gets filled to more than $3/2$. This means that at least $2m_S/3 - O(\varepsilon m_{LS})$ bins are covered. \square

4.2.3 Placing the large items

In order to achieve an (α, ε) -desirable packing, our algorithm needs to select approximately $\alpha m_{\text{LS}} \leq m_{\text{R}}$ good items, and place them in the reserved bins; Lemma 3 guarantees that these bins will be covered (Property I holds). Note that Lemma 4 establishes Property III. Meanwhile, as we will describe, the algorithm ensures that other large items are paired and hence each pair of them covers a bin (Property II holds).

In order to pack at least $\alpha m_{\text{LS}} - O(\varepsilon m_{\text{LS}})$ good items in the reserved bins, the algorithm considers three cases depending on the location of good items in the sequence. Advice will be used to select the correct case. Furthermore, m_{LS}^- will be given as advice.

Lemma 5 *For $\beta < 15/14$, there exists an online algorithm using $O(\log \log n)$ advice bits and producing an (α, ε) -desirable packing for $\alpha \leq \frac{(1-\varepsilon)^2(\frac{7}{6}-\beta)}{1+\frac{3}{2}(1-\varepsilon)^2(1+\varepsilon)}$ and $\varepsilon \in O(1/(\log n))$.*

Proof By Lemma 4, Property III holds.

To prove that Property I holds, we first argue that the value of αm_{LS} is not larger than the number of reserved bins. By Equation (3), the number of reserved bins is

$$m_{\text{R}} > (1 - \varepsilon)^2 \left\lfloor \frac{m_{\text{LS}}}{3} \right\rfloor > \frac{3}{4} \left\lfloor \frac{m_{\text{LS}}}{3} \right\rfloor, \text{ for } \varepsilon \leq \frac{1}{8}.$$

Since $\beta \geq 1$, we get

$$\alpha \leq \frac{(1 - \varepsilon)^2(\frac{7}{6} - \beta)}{1 + \frac{3}{2}(1 - \varepsilon)^2(1 + \varepsilon)} < \frac{7 - 6\beta}{6/(1 - \varepsilon)^2 + 9(1 + \varepsilon)} < \frac{7 - 6\beta}{15} < \frac{1}{15}.$$

Large items are placed in a way to establish Properties I and II, following a case analysis based on where good items appear in the request sequence (the appropriate case is encoded using advice).

Case 1: Assume there are at least $\lfloor \alpha m_{\text{LS}}^- \rfloor$ good items among the first m_{R} large items in the sequence.

In this case, the algorithm places the first m_{R} large items into the reserved bins. After seeing all these m_{R} items, the algorithm chooses the largest $\lfloor \alpha m_{\text{LS}}^- \rfloor$ of them and declares them to be good items, which by Lemma 3 are guaranteed to be covered. Thus, Property I holds.

The remaining $m_{\text{R}} - \lfloor \alpha m_{\text{LS}}^- \rfloor$ large items in the reserved bins will be paired with forthcoming large items. Since there are at least $m_{\text{LS}} - m_{\text{R}} \geq 2m_{\text{LS}}/3$ forthcoming large items and fewer than $m_{\text{LS}}/3$ large items in the reserved bins waiting to be paired, all these large items (except possibly one) can be paired (Property II holds). In summary, in the final covering, there are $\lfloor \alpha m_{\text{LS}}^- \rfloor$ bins covered by a large item (and some small items) while the remaining large items are paired (except possibly one). Hence, Property II holds.

Case 2: Assume there are fewer than $\lfloor \alpha m_{\text{LS}}^- \rfloor$ good items among the first m_{R} large items in the sequence (Case 1 does not apply). Furthermore, assume

there are at least $\lfloor \alpha m_{\text{LS}}^- \rfloor$ good items among the m_{R} large items that follow the first m_{R} large items.

In this case, the algorithm places the first m_{R} large items pairwise in $\lceil m_{\text{R}}/2 \rceil$ bins that are not reserved. The m_{R} large items that follow are placed in the reserved bins. After placing the last of these items in the reserved bins, the algorithm considers these m_{R} items and declares the $\lfloor \alpha m_{\text{LS}}^- \rfloor$ largest to be good items, which by Lemma 3 are guaranteed to be covered. Thus, Property I holds.

The remaining $m_{\text{R}} - \lfloor \alpha m_{\text{LS}}^- \rfloor$ reserved bins (with large items) will need to be covered by forthcoming large items. We know there are at least $m_{\text{LS}} - 2m_{\text{R}} \geq \lfloor m_{\text{LS}}/3 \rfloor$ forthcoming large items and fewer than $\lfloor m_{\text{LS}}/3 \rfloor$ large items in reserved bins waiting to be paired, so all these large items (except possibly one) can be paired. Thus, Property II holds.

Case 3: Assume there are fewer than $\lfloor \alpha m_{\text{LS}}^- \rfloor$ good items among the first m_{R} large items and also fewer than $\lfloor \alpha m_{\text{LS}}^- \rfloor$ good items among the following m_{R} items (Cases 1 and 2 do not apply).

In this case, the algorithm places the first $2m_{\text{R}}$ large items in pairs. There are $L = 2m_{\text{LL}} + m_{\text{LS}} - 2m_{\text{R}}$ remaining large items. The algorithm places the first $F = m_{\text{LL}}^+ + \lfloor m_{\text{R}}/2 \rfloor + \lfloor \alpha m_{\text{LS}}^-/2 \rfloor - 1$ of the last L large items in the reserved bins (note that this is roughly half of the last L large items when α is small). For this to be possible, we show first that the number of reserved bins is at least F , i.e., $F - m_{\text{R}} \leq 0$.

$$\begin{aligned}
F - m_{\text{R}} &= m_{\text{LL}}^+ + \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor + \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor - 1 - m_{\text{R}} \\
&\leq m_{\text{LL}}^+ - \frac{m_{\text{R}}}{2} + \frac{\alpha m_{\text{LS}}^-}{2} - 1 \\
&< m_{\text{LL}}^+ - (1 - \varepsilon)^2 \frac{n_{\text{G}}}{2} + \frac{\alpha m_{\text{LS}}^-}{2} - 1, \text{ by Equation (3)} \\
&< m_{\text{LL}}^+ - (1 - \varepsilon)^2 \frac{m_{\text{LS}}}{6} + \frac{\alpha m_{\text{LS}}^-}{2} \\
&< (1 + \varepsilon)m_{\text{LL}} - (1 - \varepsilon)^2 \frac{m_{\text{LS}}}{6} + \frac{\alpha m_{\text{LS}}}{2} \\
&= (1 + \varepsilon)(\beta - 1)m_{\text{LS}} - (1 - \varepsilon)^2 \frac{m_{\text{LS}}}{6} + \frac{\alpha m_{\text{LS}}}{2} \\
&= \left((1 + \varepsilon)(\beta - 1) - \frac{(1 - \varepsilon)^2}{6} + \frac{\alpha}{2} \right) m_{\text{LS}} \\
&< \left(\frac{1 + \varepsilon}{14} - \frac{(1 - \varepsilon)^2}{6} + \frac{1}{30} \right) m_{\text{LS}} \\
&< 0, \text{ for } \varepsilon \leq \frac{1}{8},
\end{aligned}$$

The second to last equality follows since $m_{\text{LL}} = (\beta - 1)m_{\text{LS}}$, and the second to last inequality follows from $\beta < 15/14$ and $\alpha < 1/15$.

This establishes that there are enough reserved bins.

Next, we show that at least $\alpha m_{LS} - 6$ of the F items placed in the reserved bins are good items. There are fewer than $2 \lfloor \alpha m_{LS}^- \rfloor$ good items among the first $2m_R$ large items. Thus, the number of good items among the last L large items is more than $n_G - 2 \lfloor \alpha m_{LS}^- \rfloor$. Even if all items among the last $L - F$ large items are good, there are still more than $G = n_G - 2 \lfloor \alpha m_{LS}^- \rfloor - (L - F)$ good items that are placed in the reserved bins. In the following we deduce that $G > \alpha m_{LS} - 6$:

$$\begin{aligned}
G &= n_G - 2 \lfloor \alpha m_{LS}^- \rfloor - L + F \\
&= n_G - 2 \lfloor \alpha m_{LS}^- \rfloor - (2m_{LL} + m_{LS} - 2m_R) + m_{LL}^+ + \left\lfloor \frac{m_R}{2} \right\rfloor + \left\lfloor \frac{\alpha m_{LS}^-}{2} \right\rfloor - 1 \\
&> n_G - 2\alpha m_{LS}^- - (2m_{LL} + m_{LS} - 2m_R) + m_{LL}^+ + \frac{m_R}{2} + \frac{\alpha m_{LS}^-}{2} - 3 \\
&= n_G - \frac{3}{2}\alpha m_{LS}^- - 2m_{LL} - m_{LS} + \frac{5}{2}m_R + m_{LL}^+ - 3 \\
&> n_G - \frac{3}{2}\alpha m_{LS}^- - 2m_{LL} - m_{LS} + \frac{5}{2}(1 - \varepsilon)^2 n_G + m_{LL}^+ - 3 \\
&> \frac{7}{2}(1 - \varepsilon)^2 n_G - \frac{3}{2}\alpha m_{LS}^- - m_{LL} - m_{LS} - 3 \\
&> \frac{7}{6}(1 - \varepsilon)^2 m_{LS} - \frac{3}{2}\alpha m_{LS} - m_{LL} - m_{LS} - 6 \\
&= \frac{7}{6}(1 - \varepsilon)^2 m_{LS} - \frac{3}{2}\alpha m_{LS} - \beta m_{LS} - 6 \\
&> (1 - \varepsilon)^2 \left(\frac{7}{6} - \frac{3}{2}\alpha - \beta \right) m_{LS} - 6
\end{aligned}$$

The second to last inequality follows since $n_G = \lfloor \frac{m_{LS}}{3} \rfloor \geq \frac{m_{LS}}{3} - \frac{2}{3}$.

Since we assumed $\alpha \leq \frac{(1 - \varepsilon)^2 (\frac{7}{6} - \beta)}{1 + \frac{3}{2}(1 - \varepsilon)^2 (1 + \varepsilon)}$, we can write

$$\beta \leq \frac{7}{6} - \frac{1 + \frac{3}{2}(1 - \varepsilon)^2 (1 + \varepsilon)}{(1 - \varepsilon)^2} \alpha = \frac{7}{6} - \frac{\alpha}{(1 - \varepsilon)^2} - \frac{3}{2}(1 + \varepsilon)\alpha,$$

and

$$\begin{aligned}
G &> (1 - \varepsilon)^2 \left(\frac{7}{6} - \frac{3}{2}\alpha - \frac{7}{6} + \frac{\alpha}{(1 - \varepsilon)^2} + \frac{3}{2}(1 + \varepsilon)\alpha \right) m_{LS} - 6 \\
&> \alpha m_{LS} - 6.
\end{aligned}$$

Hence, at least $\lfloor \alpha m_{LS} \rfloor - 6$ of the F items placed in the reserved bins are good items. After placing these F items, the algorithm declares the largest $\lfloor \alpha m_{LS}^- \rfloor - 6$ among them to be good items. By Lemma 3, these items (along with small items in the reserved bins) will cover their respective bins. Thus, Property I holds.

There are $F - \lfloor \alpha m_{LS}^- \rfloor + 6$ large items in reserved bins which have not been declared good, and the $L - F$ large items which have not arrived at this point will be paired with them. We prove that, among the $F - \lfloor \alpha m_{LS}^- \rfloor + 6$ large

items that are not declared good, only $O(\varepsilon m_{\text{LS}})$ items might not be paired with one of the remaining $L - F$ large items:

$$\begin{aligned}
& F - \lfloor \alpha m_{\text{LS}}^- \rfloor + 6 \\
&= m_{\text{LL}}^+ + \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor + \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor - 1 - \lfloor \alpha m_{\text{LS}}^- \rfloor + 6 \\
&\leq m_{\text{LL}}^+ + \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor - \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor + 6 \\
&\leq (2m_{\text{LL}}^+ - m_{\text{LL}}^+) + \left(3m_{\text{R}} - 2m_{\text{R}} - \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor \right) - \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor + 6 \\
&\leq (2m_{\text{LL}} - m_{\text{LL}}^+) + \left(m_{\text{LS}} - 2m_{\text{R}} - \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor \right) - \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor + O(\varepsilon m_{\text{LL}}) \\
&< (2m_{\text{LL}} + m_{\text{LS}} - 2m_{\text{R}}) - \left(m_{\text{LL}}^+ + \left\lfloor \frac{m_{\text{R}}}{2} \right\rfloor + \left\lfloor \frac{\alpha m_{\text{LS}}^-}{2} \right\rfloor - 1 \right) + O(\varepsilon m_{\text{LL}}) \\
&= L - F + O(\varepsilon m_{\text{LL}}) \\
&= L - F + O(\varepsilon m_{\text{LS}}), \text{ since } m_{\text{LL}} = (\beta - 1)m_{\text{LS}} < \frac{1}{14}m_{\text{LS}}
\end{aligned}$$

So, the number of large items in the reserved bins which are not paired will be $O(\varepsilon m_{\text{LS}})$. The bins in which these unpaired items are placed, along with the $\lfloor \alpha m_{\text{LS}}^- \rfloor - 6$ bins that include good items, will be the LS-bins in the final (α, ε) -desirable packing. All large items placed in bins other than LS-bins are paired and hence, Property II also holds. \square

4.3 Wrapping it up

Theorem 3 *There is an algorithm that, provided with $O(\log \log n)$ bits of advice, achieves an asymptotic competitive ratio of at least $\frac{8}{15}$.*

Proof The advice indicates the values of m_{LS}^- , m_{LL}^+ , d^+ , s_{B}^- , s_{B}^+ , e_{B}^- , m_{B}^- , and m_{W}^- . These values can all be encoded in $O(\log \log n)$ bits of advice. Note that one cannot calculate β exactly, since m_{LS} and m_{LL} are not known exactly. Thus, the advice also includes 1 bit to indicate if Lemma 1 should be used because β is larger than $15/14$, or $m_{\text{LS}} = 0$. If not, the advice also indicates one of the three cases described above; this requires two more bits. Thus, the advice is $O(\log \log n)$ bits.

If Lemma 1 is used, the competitive ratio is at least $8/15$. Otherwise, provided with this advice and a sufficiently small ε , the algorithm can create an (α, ε) -desirable packing of the input sequence for any $\alpha \leq \frac{(1-\varepsilon)^2(\frac{7}{6}-\beta)}{1+\frac{3}{2}(1-\varepsilon)^2(1+\varepsilon)}$. By Lemma 2, the resulting packing has an asymptotic competitive ratio of at least $\frac{\alpha+2\beta-1}{2\beta}$, for $\varepsilon \in o(1)$. Choosing $\alpha = \frac{(1-\varepsilon)^2(\frac{7}{6}-\beta)}{1+\frac{3}{2}(1-\varepsilon)^2(1+\varepsilon)}$ and $\varepsilon \leq 1/\log n$ gives an algorithm with asymptotic competitive ratio at least $\frac{12\beta-4}{15\beta} \geq \frac{8}{15}$. \square

5 Impossibility result for advice of sub-linear size

This section uses what is normally referred to as lower bound techniques, but since our ratios are smaller than 1, an upper bound is a negative result, and we refer to such results as negative or impossibility results. In what follows, we show that, in order to achieve any competitive ratio larger than $15/16$, advice of linear size is necessary. We use a reduction from the binary separation problem, introduced in [12]:

Definition 2 The *Binary Separation Problem* is the following online problem. The input $I = (n_1, \sigma = \langle y_1, y_2, \dots, y_n \rangle)$ consists of $n = n_1 + n_2$ positive values which are revealed one by one. There is a fixed partitioning of the set of values into a subset of n_1 *large* values and a subset of n_2 *small* values, so that all large values are larger than all small values. Upon receiving a value y_i , an online algorithm must guess if y belongs to the set of small or large values. After the algorithm has made a guess, it is revealed to the algorithm which class y_i belongs to. The number n_1 is known to the algorithm from the beginning.

A reduction from a closely related problem named “binary string guessing with known history” shows that, in order to guess more than half of the values correctly, advice of linear size is required:

Lemma 6 ([12]) *For any fixed $\beta > 0$, any deterministic algorithm for the Binary Separation Problem that is guaranteed to guess correctly on more than $(1/2 + \beta)n$ input values on an input of length n needs at least $\Omega(n)$ bits of advice.*

The following lemma provides the actual reduction from the Binary Separation Problem to bin covering.

Lemma 7 *Consider the bin covering problem on sequences of length $2n$ for which OPT covers n bins. Assume that there is an online algorithm \mathbb{A} that solves the problem on these instances using $b(n)$ bits of advice and covers at least $n - r(n)/8$ bins. Then there is also an algorithm BSA that solves the Binary Separation Problem on sequences of length n using $b(n) + O(\log n)$ bits of advice and guessing incorrectly at most $r(n)$ times.*

Proof In the reduction, we encode requests for the algorithm BSA as items for bin covering, which will be given to the algorithm \mathbb{A} . Assume we are given an instance $I = (n_1, \sigma = \langle y_1, y_2, \dots, y_n \rangle)$ of the Binary Separation Problem, in which n_1 is the number of large values ($n_1 + n_2 = n$), and the values y_t are revealed in an online manner ($1 \leq t \leq n$). We create an instance of the bin covering problem for \mathbb{A} which has length $2n$.

The bin covering sequence starts with n_1 “huge” items of size $1 - \varepsilon$ for some $\varepsilon < \frac{1}{2n}$ (this will ensure that at least one item of size at least $1/2$ is in every covered bin). The value n is given as advice, using $O(\log n)$ bits. In the optimal covering, each of these huge items will be placed in a different bin. The next n items are created in an online manner, so that we can use the result

of their packing to guess the requests for the Binary Separation Problem. Let $\tau = y_t$ ($1 \leq t \leq n$) be a requested value from the Binary Separation Problem, and choose an increasing function $f : \mathbb{R}^+ \rightarrow (\varepsilon, 2\varepsilon)$. When τ is presented to BSA, $f(\tau)$ is presented to \mathbb{A} , and the item of size $f(\tau)$ is said to be associated with the request τ . If the algorithm places the item $f(\tau)$ in one of the bins opened by the huge items, BSA will answer that τ is small; otherwise, BSA will answer that τ is large.


The last n_2 items of the bin covering instance are defined as complements of the items in the bin covering instance associated with large values in the binary separation instance (the complement of item x is $1-x$). We do not need to give the last items complementing the small items in order to implement the algorithm BSA, but we need them for the proof of the quality of the correspondence that we are proving.

Call an item in the bin covering sequence “large” if it is associated with a large value in the Binary Separation Problem, and “small” if associated with a small value. For the bin covering sequence produced by the reduction, an optimal algorithm pairs each of the small items with a huge item, placing it in one of the first n_1 bins. OPT pairs the large items with their complements, starting one of the next n_2 bins with each of these large items. Hence, the number of bins in an optimal covering is $n_1 + n_2 = n$.


Let a_1 and a_2 denote the number of the two types of mistakes that the bin covering algorithm \mathbb{A} makes, causing incorrect answers to be given by BSA. Thus, we let a_1 denote the number of small items which do not get placed with a huge item, and a_2 denote the number of large items which are placed with a huge item. Clearly, the number of errors in the answers provided by the binary separation algorithm BSA is $a_1 + a_2$. We claim that the number of bins covered by the bin covering algorithm \mathbb{A} is at most $n_1 + n_2 - (a_1 + a_2)/8$. \mathbb{A} covers at most $(n_1 - a_1) + (n_2 - a_2)$ bins in the same way that OPT does; these bins include either a huge item with a small item or a large item and its complement. Moreover, at most $\min\{a_1, a_2\}$ bins are covered by huge items that are placed with large items, but no small item. Similarly, at most $\min\{a_1/2, a_2\}$ bins are covered with two small items and a complement of a large item. Finally, there are $p = \max\{a_1 - a_2, 0\}$ bins with huge items that are not covered with any small or large items; similarly, there are $q = \max\{a_2 - a_1/2, 0\}$ complements of large items that are not covered by large or pairs-of-small items. These $p+q$ items can be paired to cover at most $(p+q)/2$ bins. In total, the number of bins that are covered by \mathbb{A} is

$$\begin{aligned} \mathbb{A}(I) \leq & (n_1 - a_1) + (n_2 - a_2) + \min\{a_1, a_2\} + \min\{a_1/2, a_2\} \\ & + \frac{1}{2}(\max\{a_1 - a_2, 0\} + \max\{a_2 - a_1/2, 0\}) \end{aligned}$$

If $a_1 \leq a_2$, the above value becomes $n_1 + n_2 - a_2/2 + a_1/4$ which is indeed at most $n_1 + n_2 - (a_1 + a_2)/8$. If $a_1/2 \leq a_2 \leq a_1$, the above value becomes $n_1 + n_2 - a_1/4$ which is at most $n_1 + n_2 - (a_1 + a_2)/8$. Finally, if $a_2 < a_1/2$, the above value becomes $n_1 + n_2 - (a_1 + a_2)/2$ which is less than $n_1 + n_2 - (a_1 + a_2)/8$. In summary, when the algorithm BSA makes $a_1 + a_2$ errors in partitioning small



(a) The optimal covering



(b) A covering with 8 mistakes

Fig. 3: Two coverings for the sequences used in reduction from binary separation to bin covering. Grey items show huge items that arrive first. Red and green items are respectively large and small items that need to be separated. The sequence ends with complements of large items. (a) shows an optimal packing and (b) shows a covering with eight mistakes; as a result of the eight mistakes, one fewer bin is covered.

and large items, \mathbb{A} covers at most $n - (a_1 + a_2)/8$ bins (intuitively speaking, each eight mistakes in binary guessing causes at least one fewer bin to be covered; see Figure 3). In other words, if the number of covered bins is at least $n - r(n)/8$, then the number of binary separation errors must be at most $r(n)$. \square

It turns out that reducing the Binary Separation Problem to bin covering (the above lemma) is more involved than a similar reduction to the bin packing problem [12]. The difference roots in the fact that there are more ways to place items into bins in the bin covering problem compared to bin packing; this is because many arrangements of items are not allowed in bin packing due to the capacity constraint.

Theorem 4 *Consider the bin covering problem on sequences of length n . To achieve a competitive ratio of $15/16 + \delta$, in which δ is a small, but fixed positive constant, an online algorithm needs to receive $\Omega(n)$ bits of advice.*

Proof Consider a bin covering algorithm \mathbb{A} with competitive ratio $15/16 + \delta$, and sequences of length $2n$ for which OPT covers n bins. \mathbb{A} covers $(15/16 + \delta)n = n - r(n)/8$ bins for $r(n) = (1/2 - 8\delta)n$. Applying Lemma 7, we conclude that there is an algorithm that solves the Binary Separation Problem on sequences of length n using at most $O(\log n)$ bits of additional advice, while making at most $(1/2 - 8\delta)n$ errors. By Lemma 6, we know that such an algorithm requires $\Omega(n)$ bits of advice. Thus, \mathbb{A} requires $\Omega(n)$ bits of advice as well. \square

6 Concluding remarks

We have established that for bin covering $\Theta(\log \log n)$ bits of advice are necessary and sufficient to improve the competitive ratio obtainable by purely online algorithms. This differs significantly from the results from bin packing, where a constant number of bits of advice are sufficient.

Obvious questions are: How much better than our bound of $8/15 = 0.53\bar{3}$ can one do with $O(\log \log n)$ bits of advice? Can one do better with $O(\log n)$ bits of advice?

Acknowledgements We thank anonymous referees for their valuable comments.

References

1. Angelopoulos, S., Dürr, C., Kamali, S., Renault, M.P., Rosén, A.: Online bin packing with advice of small size. *Theory of Computing Systems* **62**(8), 2006–2034 (2018)
2. Ásgeirsson, E.I., Ayesta, U., Jr., E.G.C., Etra, J., Momcilovic, P., Phillips, D.J., Vokshoori, V., Wang, Z., Wolfe, J.: Closed on-line bin packing. *Acta Cybernetica* **15**(3), 361–367 (2002)
3. Assmann, S.F., Johnson, D.S., Kleitman, D.J., Leung, J.Y.: On a dual version of the one-dimensional bin packing problem. *Journal of Algorithms* **5**(4), 502–525 (1984)
4. Balogh, J., Békési, J., Dósa, G., Epstein, L., Levin, A.: A new and improved algorithm for online bin packing. In: 26th Annual European Symposium on Algorithms (ESA), *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 112, pp. 5:1–5:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2018)
5. Balogh, J., Békési, J., Dósa, G., Epstein, L., Levin, A.: A new lower bound for classic online bin packing. *ArXiv* **1807.05554** [cs:DS] (2018)
6. Böckenhauer, H., Hromkovic, J., Komm, D.: A technique to obtain hardness results for randomized online algorithms – A survey. In: Computing with New Resources – Essays Dedicated to Jozef Gruska on the Occasion of His 80th Birthday, *Lecture Notes in Computer Science*, vol. 8808, pp. 264–276. Springer (2014)
7. Böckenhauer, H., Hromkovič, J., Komm, D., Krug, S., Smula, J., Sprock, A.: The string guessing problem as a method to prove lower bounds on the advice complexity. *Theoretical Computer Science* **554**, 95–108 (2014)
8. Böckenhauer, H., Komm, D., Královic, R., Královic, R.: On the advice complexity of the k-server problem. *Journal of Computer and System Sciences* **86**, 159–170 (2017)

9. Böckenhauer, H., Komm, D., Královic, R., Královic, R., Mömke, T.: Online algorithms with advice: The tape model. *Information and Computation* **254**, 59–83 (2017)
10. Borodin, A., Pankratov, D., Salehi-Abari, A.: A simple PTAS for the dual bin packing problem and advice complexity of its online version. In: 1st Symposium on Simplicity in Algorithms (SOSA), LIPIcs, pp. 8:1–8:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2018)
11. Boyar, J., Favrholt, L.M., Kudahl, C., Larsen, K.S., Mikkelsen, J.W.: Online algorithms with advice: A survey. *ACM Computing Surveys* **50**(2), 19:1–19:34 (2017)
12. Boyar, J., Kamali, S., Larsen, K.S., López-Ortiz, A.: Online bin packing with advice. *Algorithmica* **74**(1), 507–527 (2016)
13. Csirik, J., Totik, V.: Online algorithms for a dual version of bin packing. *Discrete Applied Mathematics* **21**(2), 163–167 (1988)
14. Dürr, C., Konrad, C., Renault, M.P.: On the power of advice and randomization for online bipartite matching. In: 24th Annual European Symposium on Algorithms (ESA), *LIPIcs*, vol. 57, pp. 37:1–37:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)
15. Grove, E.F.: Online bin packing with lookahead. In: 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 430–436. SIAM (1995)
16. Hoberg, R., Rothvoss, T.: A logarithmic additive integrality gap for bin packing. In: 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2616–2625. SIAM (2017)
17. Hromkovic, J., Královic, R., Královic, R.: Information complexity of online problems. In: 35th International Symposium on Mathematical Foundations of Computer Science (MFCS) 2010, *Lecture Notes in Computer Science*, vol. 6281, pp. 24–36. Springer (2010)
18. Jansen, K., Solis-Oba, R.: An asymptotic fully polynomial time approximation scheme for bin covering. *Theoretical Computer Science* **306**(1–3), 543–551 (2003)
19. Karmarkar, N., Karp, R.M.: An efficient approximation scheme for the one-dimensional bin-packing problem. In: 23rd Annual Symposium on Foundations of Computer Science (FOCS), pp. 312–320. IEEE Computer Society (1982)
20. Komm, D.: An Introduction to Online Computation – Determinism, Randomization, Advice. *Texts in Theoretical Computer Science. An EATCS Series*. Springer (2016)
21. Mikkelsen, J.W.: Randomization can be as helpful as a glimpse of the future in online computation. In: 43rd International Colloquium on Automata, Languages, and Programming (ICALP), pp. 39:1–39:14. Springer (2016)
22. Renault, M.P.: Online algorithms with advice for the dual bin packing problem. *Central European Journal of Operations Research* **25**(4), 953–966 (2017)
23. Renault, M.P., Rosén, A., van Stee, R.: Online algorithms with advice for bin packing and scheduling problems. *Theoretical Computer Science* **600**, 155–170 (2015)
24. Rothvoss, T.: Approximating bin packing within $o(\log \text{OPT} * \log \log \text{OPT})$ bins. In: 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 20–29. IEEE Computer Society (2013)