

Online Bin Packing with Advice

Joan Boyar · Shahin Kamali · Kim S.
Larsen · Alejandro López-Ortiz

Received: date / Accepted: date

Abstract We consider the online bin packing problem under the advice complexity model where the “online constraint” is relaxed and an algorithm receives partial information about the future items. We provide tight upper and lower bounds for the amount of advice an algorithm needs to achieve an optimal packing. We also introduce an algorithm that, when provided with $\log n + o(\log n)$ bits of advice, achieves a competitive ratio of $3/2$ for the general problem. This algorithm is simple and is expected to find real-world applications. We introduce another algorithm that receives $2n + o(n)$ bits of advice and achieves a competitive ratio of $4/3 + \varepsilon$. Finally, we provide a lower bound argument that implies that advice of linear size is required for an algorithm to achieve a competitive ratio better than $9/8$.

The work of the first and third author was partially supported by the Danish Council for Independent Research, Natural Sciences, and the Villum Foundation, and most of the work was carried out while these authors were visiting the University of Waterloo.

Joan Boyar

Department of Mathematics and Computer Science, University of Southern Denmark,
Odense, Denmark

E-mail: joan@imada.sdu.dk

Shahin Kamali

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada

E-mail: s3kamali@cs.uwaterloo.ca

Kim S. Larsen

Department of Mathematics and Computer Science, University of Southern Denmark,
Odense, Denmark

E-mail: kslarsen@imada.sdu.dk

Alejandro López-Ortiz

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada

E-mail: alopez-o@uwaterloo.ca

1 Introduction

In the classical one-dimensional bin packing problem the goal is to pack a given sequence of *items* into a minimum number of *bins* with fixed and equal capacities. For convenience, it is assumed that items sizes are in the range $(0, 1]$ and the capacities of bins are 1. In the *online* version of the problem, the items are revealed one by one, and an algorithm must pack each item without any knowledge about future items. The decisions of an online algorithm are irrevocable, i.e., it is not possible to move an item from one bin to another after it is *packed* in a bin.

The online bin packing problem has many applications in practice, from cutting stock applications to creating file backups in removable media [15, 16]. Heuristics that have been proposed for the problem include Next-Fit (NF), First-Fit (FF), Best-Fit (BF), and the Harmonic-based class of algorithms. NF maintains a single *open* bin and places an item in that bin; in the case the item does not fit, it *closes* the bin and opens a new one. FF keeps a list of bins in the order they are opened, packs an item in the first bin that has enough space, and opens a new bin if necessary. BF performs similarly to FF, except that the bins are ordered in increasing order of their remaining capacity. Harmonic-based algorithms are based on the idea of packing items of similar sizes together in a bin. For Harmonic_K , an item has type i ($1 \leq i \leq K-1$) if it is in the range $(\frac{1}{i+1}, \frac{1}{i}]$, and type K if it is in the range $(0, \frac{1}{K}]$. The algorithm applies the NF strategy for items of each type separately.

As for other online problems, the standard method for comparing bin packing algorithms is competitive analysis. Under competitive analysis, the performance of an algorithm \mathbb{A} is compared to that of OPT , which is the optimal offline algorithm. More precisely, the competitive ratio of an algorithm \mathbb{A} is the asymptotically maximum ratio of the number of bins used by \mathbb{A} to that of OPT for packing the same sequence σ . FF and BF have the same competitive ratio of 1.7, while the best Harmonic-based algorithm has a competitive ratio of at most 1.58889 [33]. It is also known that no online algorithm can have a competitive ratio better than 1.54037 [4].

The total lack of information about the future is unrealistic in many real-world scenarios [19]. A natural approach for addressing this issue is to relax the problem by providing extra information about the input sequence. For the online bin packing problem, such relaxations have been studied in the contexts of *lookahead*, in which the online algorithm can look at the items arriving in the near future [25], and *closed bin packing*, in which the length of the input sequence is known to the online algorithm [2]. In both cases, the average performance of the online algorithm improves, compared to the online algorithm with no information about the future. In [20], bin packing is studied in an almost-online setting in which the online algorithm knows the value of OPT in advance. It is proven that any online algorithm has a competitive ratio of at least $47/36 = 1.30\bar{5}$ under this setting. This lower bound was later improved to 1.3231 [3]. Other almost-online settings allow the online algorithms to repack some items or arrange for some pre-ordering

of the items [22, 27, 23]. Along similar lines, the almost-online setting is studied for the related problem of scheduling; see [14, 1, 34], for example.

Under the advice complexity model for online algorithms, the “no knowledge assumption” behind online algorithms is relaxed by letting the online algorithm receive some bits of *advice* about the future items. The advice can be any information about the input sequence and is generated by an offline oracle which has unbounded computational power. Provided with the appropriate advice, the online algorithms are expected to achieve improved competitive ratios. The advice model has received significant attention since its introduction [12, 26, 19, 11, 29, 30, 9, 6, 17, 21, 28, 10, 7, 32].

The advice model provides a framework to measure the information content of online problems [26]. There is generally a trade-off between the number of advice bits and the quality of solutions. For some problems, including list updating and a variant of knapsack, a constant number of bits of advice is known to be sufficient to outperform all online algorithms [13, 9]. On the other hand, for many online problems, at least a linear number of bits is required to achieve optimal solutions. In this paper, we investigate the trade-off between the size of advice and the competitive ratio of online bin packing algorithms. This leads to a better understanding of the problem when compared to other problems in terms of the information content.

Our interest in studying the problem under this setting is mostly theoretical. Nevertheless, in many practical scenarios, it can be justified to allow a fast offline oracle to take a “quick look” at the input sequence and send some advice to the online algorithm. For example, it may be possible to take a quick look and count the number of items which are larger than $1/2$ and smaller than $2/3$ of the bin capacity. We show that this form of advice can be used to achieve an algorithm which outperforms all online algorithms.

1.1 Model

In the last few years, slightly different models of advice complexity have been proposed for online problems. All these models assume that there is an offline oracle with infinite computational power, which provides the online algorithm with some bits of advice. How these bits of advice are given to the algorithm is the source of difference between the models. In the first model, presented in [18], an online algorithm poses a series of questions which are answered by the offline oracle in *blocks of answers*. The total size of the answers, measured in the number of bits, defines the advice complexity. The problem with this model is that a lot of information can be encoded in the individual length of each block. To address this issue, another model is proposed in [19] which assumes that online algorithms receive a fixed number of bits of advice (per item). We call this model the *advice-with-request model*. This model is studied for problems, such as metrical task systems and k -server, and the results tend to use at least a constant number of bits of advice per request [19, 30]. Nevertheless, there are many online problems for which a sublinear and even a

constant number of bits of advice in total is sufficient to achieve good competitive ratios. However, under the advice-with-request model, the possibility of sending a sublinear number of advice bits to the algorithm is not well defined. In [12, 11] another model of advice complexity is presented which assumes that the online algorithm has access to an *advice tape*, written by the offline oracle. At any time step, the algorithm may refer to the tape and read any number of advice bits. The advice complexity is the number of bits on the tape accessed by the algorithm. We refer to this model as *advice-on-tape model*. Since its introduction, the advice-on-tape model has been used to analyze the advice complexity of many online problems including paging [12, 26, 29], disjoint path allocation [12], job shop scheduling [12, 29], k -server [11, 30], knapsack [9], various coloring problems [5, 21, 7, 32], set cover [28, 10], maximum clique [10], and graph exploration [17].

Under the advice-on-tape model, we require a mechanism to infer how many bits of advice the algorithm should read at each time step. This could be implicitly derived during the execution of the algorithm or explicitly encoded in the advice string itself. For example, we may use a *self-delimited* encoding as used in [11], in which the value of a non-negative integer X is encoded by writing the value of $\lceil \log(\lceil \log(X + 1) \rceil + 1) \rceil$ in unary (a string of 1's followed by a zero), the value of $\lceil \log(X + 1) \rceil$ in binary, and the value of X in binary (throughout, we use $\log n$ to denote $\log_2(n)$). These codes respectively require $\lceil \log(\lceil \log(X + 1) \rceil + 1) \rceil + 1$, $\lceil \log(\lceil \log(X + 1) \rceil + 1) \rceil$, and $\lceil \log(X + 1) \rceil$ bits. Thus, the self-delimited encoding of X requires

$$e(X) = \lceil \log(X + 1) \rceil + 2\lceil \log(\lceil \log(X + 1) \rceil + 1) \rceil + 1$$

bits. The existence of self-delimited encodings at the beginning of the tape usually adds a lower-order term to the number of advice bits required by an algorithm.

Regarding notation, we use $\mathbb{A}(\sigma)$ to denote the number of bins used by \mathbb{A} to pack a sequence σ . When σ follows from the context, we simply use \mathbb{A} to denote this number. We use similar notation for all algorithms, including OPT .

We consider the bin packing problem under the advice-on-tape model, which is formally defined as follows, based on the definition of the advice model in [11]:

Definition 1 In the *online bin packing problem with advice*, the input is a sequence of items $\sigma = \langle x_1, \dots, x_n \rangle$, revealed to the algorithm in an online manner ($0 < x_i \leq 1$). The goal is to pack these items in the minimum number of bins of unit size. At time step t , an online algorithm should pack item x_t into a bin. The decision of the algorithm to select the target bin is a function of $\Phi, x_1, \dots, x_{t-1}$, where Φ is the content of the advice tape. An algorithm \mathbb{A} is *c -competitive with advice complexity $s(n)$* if there exists a constant c_0 such that, for all n and for all input sequences σ of length at most n , there exists some advice Φ such that $\mathbb{A}(\sigma) \leq c \text{OPT}(\sigma) + c_0$, and at most the first $s(n)$ bits of Φ have been accessed by the algorithm. If $c = 1$ and $c_0 = 0$, then \mathbb{A} is *optimal*.

1.2 Contribution

We answer different questions about the advice complexity of the online bin packing problem. First, we study how many bits of advice are required to achieve an optimal solution. We consider two different settings of the problem. When there is no restriction on the number of distinct items or their sizes, we present the easy result that $n \lceil \log \text{OPT}(\sigma) \rceil$ bits of advice are sufficient to achieve an optimal solution, where $\text{OPT}(\sigma)$ is the number of bins in an optimal packing. We also prove that at least $(n - 2 \text{OPT}(\sigma)) \log \text{OPT}(\sigma)$ bits of advice are required to achieve an optimal solution.

When there are m distinct items in the sequence, we prove that at least $(m - 3) \log n - 2m \log m$ bits of advice are required to achieve an optimal solution. If m is a constant, there is a polynomial time online algorithm that receives $m \log n + o(\log n)$ bits of advice and achieves an optimal solution. We also show that, even if m is not a constant, there is a polynomial time online algorithm that receives $m \lceil \log(n + 1) \rceil + o(\log n)$ bits of advice and achieves a packing with $(1 + \varepsilon) \text{OPT}(\sigma) + 1$ bins.

We also study a relevant question that asks how many bits of advice are required to perform strictly better than all online algorithms. We bound this by providing an algorithm which receives $\log n + o(\log n)$ bits of advice and achieves a competitive ratio of $3/2$. Recall that any online bin packing algorithm has a competitive ratio of at least 1.54037 [4]. Hence, our algorithm outperforms all online algorithms.

Moreover, we introduce an algorithm that receives $2n + o(n)$ bits of advice and achieves a competitive ratio of $4/3 + \varepsilon$, for arbitrarily small constant values of $\varepsilon > 0$. We also prove a lower bound that implies that a linear number of bits of advice are required to achieve a competitive ratio of $9/8 - \delta$ for any fixed value of $\delta > 0$.

In the advice model, the offline oracle has unlimited computational power. However, except for optimal online algorithms, the algorithms that we introduce in this paper can be implemented to run in polynomial time.

1.2.1 Comparison to other results

Online bin packing with advice has also been studied by Renault et al. in [31]. They present an algorithm for online bin packing which is $(1 + 3\delta)$ -competitive (or asymptotically $(1 + 2\delta)$ -competitive), using $s = \frac{1}{\delta} \log \frac{2}{\delta^2} + \log \frac{2}{\delta^2} + 3$ bits of advice per item. This is a nice theoretical result, showing that techniques for designing polynomial approximation schemes can also be useful when considering online algorithms with advice. Unfortunately, applying their theorem for reasonable length input sequences and c -competitiveness, where $c \leq 4/3$, the result is not better than the $\log \text{OPT}(\sigma)$ bits per item, which are sufficient for optimality. To put this in context, we present an algorithm below which is $4/3 + \varepsilon$ -competitive and uses only two bits of advice per item, plus advice of $o(n)$ at the beginning. In total, the amortized size will be $2 + o(1)$ bits per item.

To achieve a ratio of $4/3$, Renault et al.'s algorithm has $\delta = 1/9$ (for the asymptotic result, $\delta = 1/6$), so the number of bits is at least $9 \log 162 + \log 162 + 3 > 76$ (for the asymptotic result, $6 \log 72 + \log 72 + 2 > 45$). Thus, there must be more than 2^{76} (respectively 2^{45}) bins before the result improves on the naive approach of using $\log \text{OPT}(\sigma)$ bits per item.

2 Optimal Algorithms with Advice

In this section we study the amount of advice required to achieve an optimal solution. We first investigate the theoretical setting in which there is no restriction on the number of distinct items or on their sizes. We observe that there is a simple algorithm that receives $n \lceil \log \text{OPT}(\sigma) \rceil$ bits of advice and achieves an optimal solution. Such an algorithm basically reads $\lceil \log \text{OPT}(\sigma) \rceil$ bits for each item, encoding the index of the bin that includes the item in an optimal packing. We show that the upper bound given by this algorithm is tight up to lower order terms, when $n - 2 \text{OPT}(\sigma) \in \Theta(n)$.

Theorem 1 *To achieve an optimal packing for a sequence of size n which can be optimally packed in $\text{OPT}(\sigma)$ bins, it is sufficient to receive $n \lceil \log \text{OPT}(\sigma) \rceil$ bits of advice, and any deterministic online algorithm requires at least $(n - 2 \text{OPT}(\sigma)) \log \text{OPT}(\sigma)$ bits of advice to achieve an optimal packing.*

Proof Upper Bound: Consider an offline oracle that knows an optimal packing (note that such an oracle has unbounded computational power). This oracle simply writes on the advice tape, for each item x , except for the last two, the index of the bin in an optimal packing that x is packed in. To pack any item x , the online algorithm simply reads the index of the bin that x should be packed in and packs x accordingly. For the last two items, the algorithm simply uses Best-Fit. Since the packing is the same as one for an optimal algorithm up to that point, if it is impossible to fit both of the remaining items in the bins already used, Best-Fit will ensure that at least one fits if that is possible. If both of the remaining items fit in the same already open bin, it is fine to put the first one of the last two items anywhere it fits, since there will still be space remaining for the last. If both of the remaining items fit in open bins, but should be in different bins, using Best-Fit will ensure that they are both placed there. This requires $\lceil \log \text{OPT}(\sigma) \rceil$ bits of advice per item which sums up to $(n - 2) \lceil \log \text{OPT}(\sigma) \rceil$ bits of advice. The algorithm should also know the value of $X = \lceil \log \text{OPT}(\sigma) \rceil$ in order to read the appropriate number of bits on each request. This can be done by encoding X in unary and terminating with a zero. This uses no more than $2 \lceil \log \text{OPT}(\sigma) \rceil$ bits. Consequently the number of advice bits used by the algorithm is $n \lceil \log \text{OPT}(\sigma) \rceil$ as stated by the theorem.

Lower Bound: Consider a set $S = \{\sigma_1, \dots, \sigma_N\}$ of sequences, so that each σ_r has length n for $1 \leq r \leq N$. Let $1 \leq k \leq n - 1$. Each sequence σ_r in the set has the form

$$\left\langle \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \dots, \frac{1}{2^{n-k+1}}, u_1^r, u_2^r, \dots, u_k^r \right\rangle$$

in which u_1^r, \dots, u_k^r are defined as follows. Consider a set V of vectors of the form

$$V_r = (v_1^r = 1, v_2^r = 2, \dots, v_k^r = k, v_{k+1}^r, v_{k+2}^r, \dots, v_{n-k}^r)$$

such that also for each $h \in \{k+1, \dots, n-k\}$, we have $v_h^r \in \{1, \dots, k\}$.

An example, when $n = 8$ and $k = 3$, the vector $(1, 2, 3, 2, 1)$ is a vector in V .

We associate with each vector $V_r \in V$ a sequence $\sigma_r \in S$. For a vector $V_r \in V$ and bin j , define

$$u_j^r = 1 - \sum_{\substack{1 \leq i \leq n-k \\ v_i^r = j}} a_i,$$

where a_i is the i th item in the sequence σ_r , i.e., $a_i = \frac{1}{2^{i+1}}$. Note that all u_j s are strictly larger than 0.5. Clearly, $\text{OPT}(\sigma_r) = k$ for all r . We refer to the first $n - \text{OPT}(\sigma)$ items as *small* items and the last $\text{OPT}(\sigma)$ items as *large* items.

As an example, assume $n = 8$ and $\text{OPT}(\sigma) = 3$. For a vector $V_r = (1, 2, 3, 2, 1)$, we have $u_1^r = 1 - (\frac{1}{4} + \frac{1}{64}) = 0.734375$, $u_2^r = 1 - (\frac{1}{8} + \frac{1}{32}) = 0.84375$, and $u_3^r = 1 - \frac{1}{16} = 0.9375$. Hence, the sequence σ_r associated with V_r is $(\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{64}, 0.734375, 0.84375, 0.9375)$.

In fact, V_r indicates in which bin each of the first $n - \text{OPT}(\sigma)$ items of σ_r should be packed, and at the end, u_j^r fills the empty space of the j th bin to capacity to achieve an optimal packing P for a given sequence (it is optimal since all bins are fully packed). The restriction that the sequence starts with k distinct items ensures that we do not need to consider permutations of the bins in P as additional optimal packings. We claim that P is the unique optimal packing. Suppose there is another optimal packing P' . Observe that each bin includes at most one large item, and indeed exactly one since we assume it is also optimal. Let $a_i (1 \leq i \leq n - \text{OPT}(\sigma))$ be the first item which is packed in some other bin in P' than the one prescribed by P . Consider the bin B that a_i is packed into in P . This bin cannot be fully packed in P' since a_i is strictly larger than the total size of all remaining small items, i.e., even if we put all of them in the empty space of a_i , there is still some empty space in B . As a result P' cannot be optimal. Hence there is unique solution to pack each sequence in the set S .

Note that there are $N = \text{OPT}(\sigma)^{n-2\text{OPT}(\sigma)}$ sequences S . We claim that these sequences need separate advice strings. Assume to the contrary that this is not the case, and let $\sigma_r, \sigma_{r'} \in S$ ($r \neq r'$) be two different sequences with the same advice string. Note that the first $n - \text{OPT}(\sigma)$ items in these sequences are the same. Since the online algorithm performs deterministically and we assume it receives the same advice for both σ_r and $\sigma_{r'}$, the partial packings of the algorithms after packing the first $n - \text{OPT}(\sigma)$ items are the same for both sequences. However, as discussed earlier, this implies that the final packing of the algorithm is different from the optimal packing prescribed by $V_{r''}$ for at least one of the sequences. As discussed, such a packing is the unique optimal packing and deviating from that increases the number of bins

used by the algorithm by at least one unit. As a result, the algorithm performs non-optimally for at least one of σ_r or $\sigma_{r'}$. We conclude that the sequences in the set S need separate advice strings. Since there are $N = \text{OPT}(\sigma)^{n-2\text{OPT}(\sigma)}$ sequences in S , at least $\log(\text{OPT}(\sigma)^{n-2\text{OPT}(\sigma)}) = (n - 2\text{OPT}(\sigma)) \log \text{OPT}(\sigma)$ bits of advice are required to get that many distinct advice strings. \square

Next, we consider a more realistic scenario where there are $m \in o(n)$ distinct item sizes and these are known to the algorithm. Assume that the advice tape specifies the number of items of each size. If we are not concerned about the running time of the online algorithm, there is enough information to obtain an optimal solution. If we are concerned, we can use known results for solving the offline problem [8, 36, 35]. We formalize this in what follows.

Lemma 1 ([8]) *Consider the restriction of the bin packing problem to instances in which the number of distinct item sizes is a constant non-negative integer m . There is a polynomial time algorithm that optimally solves this restricted problem.*

If there are more than a constant number of distinct items sizes, we can solve the problem almost optimally if the item sizes are lower bounded by a fixed value ε .

Lemma 2 ([36, 35]) *There is a polynomial algorithm for the bin packing problem which opens at most $(1 + \varepsilon)\text{OPT}(\sigma) + 1$ bin, in which ε is any small but fixed value.*

We use the above results to obtain the following:

Theorem 2 *Consider the online bin packing problem in which there are m distinct item sizes. The item sizes are assumed to be known. If m is a constant, there is a (linear time) optimal online algorithm that receives $m \log n + o(\log n)$ bits of advice. If m is not a constant, there is a (polynomial time) online algorithm that reads $m \lceil \log(n + 1) \rceil + o(\log n)$ bits of advice and achieves an almost optimal packing with at most $(1 + \varepsilon)\text{OPT}(\sigma) + 1$ bins, for any small but constant value of ε .*

Proof The offline oracle simply encodes the input sequence, considered as a multi-set, in $m \lceil \log(n + 1) \rceil$ bits of advice. In order to do that, it writes the number of occurrences of each of the m distinct items on the tape. The online algorithm uses the algorithms of Lemma 1 (for constant values of m) or that of Lemma 2 (for non-constant m) to compute an (almost) optimal packing. Then it packs the items in an online manner according to such an (almost) optimal packing. The algorithm reads frequencies of items in chunks of $X = \lceil \log(n + 1) \rceil$ bits and consequently needs to know the value of X . So, we add a self-delimited encoding of X at the beginning of the tape using $e(X)$ bits. The number of advice bits used by the algorithm is thus $m \lceil \log(n + 1) \rceil + O(\log \log n)$, which is $m \lceil \log(n + 1) \rceil + o(\log n)$ as $m \in o(n)$. \square

We show that the above upper bound is asymptotically tight. We start with the following simple lemma.

Lemma 3 *Consider the equation $x_1 + 2x_2 + \dots + \alpha x_\alpha = X$ in which the x_i s ($i \leq \alpha$) and X are non-negative integers. If X is sufficiently large, then this equation has at least $\left(1 + \frac{2X}{\alpha(\alpha+1)}\right)^{\alpha-1}$ solutions.*

Proof Define $A = \sum_{i=1}^{\alpha} i$. Assign arbitrary values in the range $[0, X/A]$ to all x_i s for $2 \leq i \leq \alpha$ (for simplicity assume X/A is an integer). There are $(1 + X/A)^{\alpha-1}$ different such assignments. Any of these assignments defines a valid solution for the equation since by definition of A we have $\sum_{i=2}^{\alpha} ix_i \leq X$, and we can assign $x_1 = X - \sum_{i=2}^{\alpha} ix_i$. Replacing A with $\alpha(\alpha+1)/2$ completes the proof. \square

Theorem 3 *At least $(m-3)\log n - 2m\log m$ bits of advice are required to achieve an optimal solution for the online bin packing problem on sequences of length n with m distinct item sizes.*

Proof We define a family of sequences of length n and containing m distinct item sizes and show that the sequences in this family need separate advice strings to be optimally packed by an online algorithm. To define the family, we fix m item sizes as being $\{\frac{1}{2m}, \frac{m+2}{2m}, \frac{m+3}{2m}, \dots, \frac{2m-1}{2m}, 1\}$. To simplify the argument, we scale up the sizes of bins and items by a factor of $2m$. So, we assume the item sizes are $\{1, m+2, m+3, \dots, 2m-1, 2m\}$, and the bins have capacity $2m$. Each sequence in the family starts with $n/2$ items of size 1. Consider any packing of these items in which all bins have level at most equal to $m-2$. Such a packing includes a_1 bins of level 1 (one item of size 1 in each), a_2 bins of level 2 (two items of size 1 in each), etc., such that the a_i s are non-negative integers and $a_1 + 2a_2 + \dots + (m-2)a_{m-2} = n/2$. By Lemma 3, there are at least $\left(1 + \frac{n}{(m-1)(m-2)}\right)^{m-3}$ distinct packings with the desired property. For any of these packings, we define a sequence in our family. Such a sequence starts with $n/2$ items of size 1 and is followed by another $n/2$ items. Let B denote the number of bins in a given packing of the first $n/2$ items, so that $B \leq n/2$. The sequence associated with the packing is followed by B items of size larger than $m+1$ which completely fit these bins (in non-increasing order of their sizes). Finally, we include another $n/2 - B$ items of size $2m$ in the sequence to achieve a sequence of length n .

We claim that any of the sequences in the family has a unique optimal packing of size $n/2$. This is because there are exactly $n/2$ large items of size strictly greater than m (more than half the capacity of the bin), and the other $n/2$ items have small size 1 (which fit the empty space of all bins). So each bin is fully packed with one large item of size x and $2m - x$ items of size 1 (see Figure 1).

The unique optimal packing of each sequence is defined by the partial packing of the first $n/2$ small items. Consider a deterministic online algorithm \mathbb{A} receiving the same advice string for two sequence σ_1 and σ_2 . Since \mathbb{A} is





Figure 1 The optimal packings for two sequences of the family when $n = 30$ and $m = 6$ (item sizes and bin capacities are scaled by $2m = 12$).

deterministic and both sequences start with the same sub-sequence of small items, the partial packing of the algorithm after packing the first $n/2$ items is the same for both σ_1 and σ_2 . As a result, the final packing of \mathbb{A} is sub-optimal for at least one them. We conclude that any deterministic online algorithm should receive distinct advice strings for each sequence in the family. Since there are at least $\left(1 + \frac{n}{(m-1)(m-2)}\right)^{m-3}$ sequences in the family, at least $(m-3) \log \left(1 + \frac{n}{(m-1)(m-2)}\right) > (m-3) \log n - 2m \log m$ bits of advice are required. \square

3 An Algorithm with Sublinear Advice

In what follows we introduce an algorithm that receives $\log n + o(\log n)$ bits of advice and achieves a competitive ratio of $\frac{3}{2}$, for any instance of the online bin packing problem. An offline oracle can compute and write the advice on the tape in linear time, and the online algorithm runs as fast as First-Fit. Thus, the algorithm might be applied in practical scenarios in which it is allowed to have a “quick look” at the input sequence. Note that the offline oracle runs in linear time, and the time complexity of the algorithm is polynomial ($O(n \log n)$).

We call items *tiny*, *small*, *medium*, and *large* if their sizes lie in the intervals $(0, 1/3]$, $(1/3, 1/2]$, $(1/2, 2/3]$, and $(2/3, 1]$, respectively. The advice that the algorithm receives is the number of medium items, which we denote by α .

The algorithm reads the advice tape, obtains α , opens α bins, called *critical bins*, and reserves $2/3$ of the space in each of them. This reserved space will be used to pack a medium item in each of the critical bins, and these bins have a *virtual level* of size $2/3$ at the beginning. All other bins have virtual level zero when they are opened. The algorithm packs an item x in the following manner:

- If x is a large item, open a new bin for it. Set the virtual level to its size.
- If x is a medium item, put it in the reserved space of a critical bin B . Update the virtual level to the actual level. (B will not have any reserved space now.)
- If x is small or tiny, use the First Fit (FF) strategy to put it into any of the open bins, based on virtual levels (open a new bin if required). Add the size of the item to the virtual level.

Note that the critical bins appear first in the ordering maintained by the algorithm as they are opened before other bins.

Theorem 4 *There is an online algorithm which receives $\log n + o(\log n)$ bits of advice and opens at most $3/2 \text{OPT}(\sigma) + 2$ bins to pack any sequence σ of size n .*

Proof We prove that the algorithm described above has the desired property. The value of α is encoded in $X = \lceil \log(n+1) \rceil$ bits of advice. In order to read this properly from the tape, the algorithm needs to know the value of X . This can be done by adding the self-delimited encoding of X in $e(X) = \lceil \log X \rceil + 2\lceil \log \log(X) \rceil + 2$ bits at the beginning of the tape. Consequently the number of advice bits used by the algorithm is $X + O(\log X)$, which is $\log n + o(\log n)$ as stated by the theorem.

Consider the packing of the algorithm for a sequence σ . There are two cases. In the first case, there is a critical bin B so that no other item, except a medium item, is packed in it. Since all tiny items are smaller than $1/3$ and can fit in B , all the non-critical bins that are opened after B include small and large items only. More precisely, they include either a single large item or two small items (except the last one which might have a single small item). Let L , M , and S denote the number of large, medium, and small items. The number of bins used by the algorithm is at most $L + M + S/2 + 1$. Now, if $S \leq M$, this would be at most $L + 3/2M + 1$. Since $L + M$ is a lower bound for the number of bins in an optimal solution, the number of bins used by the algorithm is at most $3/2 \text{OPT}(\sigma) + 1$ and we are done. If $S > M$, OPT should open $L + M$ bins for large and medium items, and in the best case, it packs M small items together with medium ones. For the other $S - M$ small items, OPT has to open at least $(S - M)/2$ bins. Hence the number of bins of OPT is at least $L + M + (S - M)/2 = L + M/2 + S/2$, and we have $3/2 \text{OPT}(\sigma) \geq 3L/2 + 3M/4 + 3S/4 > L + M + S/2$. Thus, the number of bins used by the algorithm is at most $3/2 \text{OPT}(\sigma) + 1$.

In the second case, we assume that all critical bins include another item in addition to the medium item. We claim that, at the end of packing a sequence,

all bins, except possibly two, have level at least $2/3$. First, we verify this for non-critical bins (bins without medium items). If a non-critical bin is opened by a large item, it clearly has level higher than $2/3$. All other non-critical bins only include items of size at most $1/2$. Hence, these bins, except possibly the last one, include at least two items. Among the non-critical bins that include two items, consider two bins b_i and b_j ($i < j$) that have levels smaller than $2/3$. Since b_j contains at least two items, at least one of them has size smaller than $1/3$. This item could fit in b_i by the FF property. We conclude that all non-critical bins, except possibly two, have level at least $2/3$. Now, suppose two critical bins b_i and b_j have levels smaller than $2/3$. Consider the first non-medium item x which is packed in b_j (in the second case, such an item exists). Since a medium item is packed in the bin, x should be either tiny or small. If x is small, then the level of b_j is at least $1/2 + 1/3$, which contradicts the level of b_j being smaller than $2/3$. Similarly, x cannot be a tiny item of size larger than $1/6$ (since $1/2 + 1/6 \geq 2/3$). Hence, x is a tiny item of size at most $1/6$. This implies that at the time the online algorithm packs x , bin b_i has a virtual level of at least $5/6$. The virtual level is at most $1/6$ larger than the actual level (the final level). Hence, the actual level of b_i is at least $5/6 - 1/6 = 2/3$. We conclude that at most one critical bin has level smaller than $2/3$.

To summarize, there are at most two non-critical bins and one critical bin which have level smaller than $2/3$. On the other hand, it is not possible to have one critical bin with content less than $2/3$ and one non-critical bin (other than the last) with such content. This is because the rightmost bin will have an item of size at most $1/3$, which should have gone in the leftmost. To conclude, in the final packing, there are at most two bins with content less than $2/3$ (in which case both will be non-critical). Hence, the number of bins used by the algorithm is at most $3/2 \text{OPT}(\sigma) + 2$. \square

4 An Algorithm with Linear Advice

In this section, we present an algorithm that receives $2n + o(n)$ bits of advice and achieves a competitive ratio of $4/3 + \varepsilon$ for any sequence of size n , and arbitrarily small constant values of ε . Consider an algorithm that receives an *approximate size* for each sufficiently large item x encoded using k bits. The approximate size of x would be larger than its *actual size* by at most an additive term of $1/2^k$. The algorithm can optimally pack items by their approximate sizes and achieve an *approximate packing* which includes a reserved space of size $x + \varepsilon$ ($\varepsilon \leq 1/2^k$) for each item. Precisely, for each sufficiently large item x , the approximate packing includes a reserved space of size $x + \varepsilon$ ($\varepsilon \leq 1/2^k$) for x . This enables the algorithm to place x in the reserved space for it in the approximate packing. Note that, since the number of approximated sizes is bounded (it is $1/\varepsilon$), the approximate packing can be created in polynomial time using the algorithm of [24]. Smaller items are treated differently and the algorithm does not reserve any space for them in the approximate

packing. In the remainder of this section, we elaborate on this idea to achieve a $4/3$ -competitive algorithm.

Notice that the number of bins in an approximate packing can be as large as $\frac{3}{2}$ times that of OPT. To see that, consider a sequence which is a permutation of

$$\langle \frac{1}{2} + \varepsilon_1, \frac{1}{2} - \varepsilon_1, \frac{1}{2} + \varepsilon_2, \frac{1}{2} - \varepsilon_2, \dots, \frac{1}{2} + \varepsilon_{n/2}, \frac{1}{2} - \varepsilon_{n/2} \rangle,$$

where $\varepsilon_i < 1/2^n$ ($1 \leq i \leq n/2$). Since OPT packs all bins tightly, an increase in the sizes of items by a constant (small) ε results in opening a new bin for each two bins OPT uses. Hence, the number of bins in an optimal approximate packing can be as bad as $\frac{3}{2}$ OPT. This example suggests that using approximate packings is not good for the bins in which a small number of large items are tightly packed. To address this issue we divide the bins of OPT into two groups:

Definition 2 Consider an optimal packing of a sequence σ . Given a small parameter $\varepsilon' < 1/60$, define *good bins* to be those where the total size of the items smaller than $1/4$ in the bin is at least $5\varepsilon'$. Define all other bins to be *bad bins*.

A part of the advice received for each item x indicates if x is packed by OPT in a good bin or in a bad bin. This enables us to treat items packed in these two groups separately.

Lemma 4 Consider sequences for which all bins in the optimal packing are good (as defined above). There is an online algorithm that receives $o(n)$ bits of advice and achieves a competitive ratio of $4/3$.

Proof Call an item *small* if it is smaller than or equal to $1/6$ and *large* otherwise. The advice bits define the approximate sizes of all large items with a precision of ε' , i.e., it gives the count for each possible rounded item size. The amount of advice will be roughly $\frac{1}{\varepsilon'} \log n$, which is $o(n)$ for constant values of ε' . The online algorithm \mathbb{A} can build the optimal approximate packing of large items. In such a packing, there is a reserved space of size at most $x + \varepsilon'$ for any large item of size x . The algorithm considers this packing as a partial packing and initializes the level of each bin to be the total sizes of approximated items in that bin. To pack an item x , if x is large, \mathbb{A} places it in the space reserved for it in the approximate packing. It also updates the level of the bin to reflect the actual size of x . If x is small, \mathbb{A} simply applies the First-Fit strategy to pack x in a bin of the partial packing (and opens a new bin for it if necessary). We prove that \mathbb{A} is $4/3$ -competitive. In the final packing by \mathbb{A} , call a bin “red” if all items packed in it are small items and call it “blue” otherwise (the blue bins constitute the approximated packing at the beginning). There are two cases to consider.

In the first case, there is no red bin in the final packing of \mathbb{A} , i.e., all small items fit in the remaining space of the bins in the approximate packing of large items. Let σ' be a copy of the input sequence in which the sizes of large items are approximated, i.e., increased by at most ε' ; also let X be the number of bins

for the optimal packing of σ' . Since there is no red bin in the final packing of \mathbb{A} , the number of bins used by \mathbb{A} is equal to X . Consider the optimal packing of the actual input sequence σ . By definition, good bins contain a set S of items, each of size at most $1/4$, and with a total size of at least $5\varepsilon'$. We claim that one can transfer a subset $T \subset S$ of items to provide an available space of size at least $5\varepsilon'$ in each bin so that the total size of items in T is no more than $1/4$. If S includes an item x with size in the range $[5\varepsilon', 1/4]$, we define $T = \{x\}$, and we are done. Otherwise (when all items in S are smaller than $5\varepsilon'$), starting from an empty set, we add items of S to T , one by one, until their total size exceeds $5\varepsilon'$. At this point, since the total size is increased by at most $5\varepsilon'$ after adding each item, the total size of items in T is no more than $5\varepsilon' + 5\varepsilon' < 1/6$. Consequently, a subset T with the desired property exists. After transferring items into T , we can increase the sizes of large items to their approximate sizes. Since there are at most 5 large items in each bin and also available space of size at least $5\varepsilon'$, the packing constructed this way is a valid packing for the sequence σ' . Since the size of the transferred items for each bin is at most $1/4$, the transferred items from each group of four bins can fit in one new bin. Consequently the number of bins in the new packing is at most $5/4 \text{OPT}(\sigma)$. We know that the final packing by \mathbb{A} is the optimal packing for σ' (with X bins), and in particular not worse than the packing constructed above. Hence, the number of bins used by \mathbb{A} is no more than $5/4 \text{OPT}(\sigma)$.

In the second case, there is at least one red bin in the final packing of \mathbb{A} . We claim that all bins in the final packing of \mathbb{A} , except possibly the last, have levels larger than $3/4$. The claim obviously holds for the red bins since the levels of all these bins (excluding the last one) are larger than $5/6$. Moreover, since there is a bin which is opened by a small item, all blue bins have levels larger than $5/6$, i.e., the total size of packed items and reserved space for the large items is larger than $5/6$. Since there are at most 5 large items in each bin, the actual level of each bin in the final packing of \mathbb{A} is at least $5/6 - 5\varepsilon'$, which is not smaller than $3/4$ for $\varepsilon' \leq 1/60$. So, all bins, except possibly one, have levels larger than $3/4$. Consequently, the algorithm is $4/3$ -competitive. \square

It remains to address how to deal with bad bins. The next three lemmas do this.

Lemma 5 *Consider sequences for which all bins in the optimal packing include precisely two items. There is an algorithm that receives 1 bit of advice per item and achieves an optimal packing.*

Proof The single bit of advice for an item x determines whether or not the partner of x appeared as a previous item, where the partner of x is the item which is packed in the same bin as x in OPT 's packing. Consider an algorithm \mathbb{A} that works as follows: If the partner of x has not appeared yet, \mathbb{A} opens a new bin for x . Otherwise, it uses the BF strategy to pack x in one of the open bins. We claim that \mathbb{A} achieves an optimal packing.

Assume that initially we have a mapping that maps the last item to go into a bin to the item it goes on top of in the optimal packing, i.e., it maps

the second item of each bin to the first item. We update this mapping when necessary and maintain the invariant that we can always pack optimally according to the mapping. For packing an item x , if BF does not pack according to this mapping, it packs x on top of y' , while, according to the mapping, it was supposed to pack x on top of y , and a later x' is supposed to go on top of y' . Due to the BF strategy, $y' \geq y$, so we can update the mapping to map the currently unprocessed x' to y , and, of course, x to y' . \square

Lemma 6 *Consider a sequence σ for which all items have sizes larger than $1/4$ and for which each bin in OPT's packing includes precisely three items. The Harmonic algorithm uses at most $4/3 \text{OPT}(\sigma) + 3$ bins to pack such a sequence.*

Proof The proof is based on a simple weighting function. Call an item x *large* if $1/3 < x < 1/2$ and *small* otherwise ($1/4 < x \leq 1/3$). Note that, since there are three items in each bin and all are larger than $1/4$, no item can have size $1/2$ or larger. Define the weight of x to be $1/2$ if x is large and $1/3$ if it is small. Consider a bin B in the packing of σ by OPT. Since there are three items in B , its weight is maximized when there are two large items and one small item in it (three large items do not fit in the same bin). Hence, the weight of each bin in the OPT packing is at most $2(1/2) + 1/3 = 4/3$. Consequently, we have $\text{OPT}(\sigma) \geq 3/4W$, where W is the total weights of all items.

The Harmonic algorithm (HA) simply packs small and large items in separate collections of bins. So, each of the algorithm's bins, except possibly two bins, contains either three small items or two large items. In both cases, the weight of each bin is at least 1 and we have $\text{HA}(\sigma) \leq W + 2$. As a conclusion $\text{HA}(\sigma) \leq 4/3 \text{OPT}(\sigma) + 2$ which completes the proof. \square

Lemma 7 *Consider a sequence σ for which all bins in the optimal packing are bad bins (as defined earlier). There is an algorithm that receives two bits of advice per item, and opens at most $(4/3 + \frac{5\varepsilon'}{1-5\varepsilon'}) \text{OPT}(\sigma) + 4$ bins.*

Proof By the definition of bad bins, for any bin in the optimal packing, all items are either smaller than $5\varepsilon'$ or larger than $1/4$. We call the former group of items *tiny* items and pack them separately using the FF strategy. We refer to other items as *normal items*. Consider an offline packing P which is the same as OPT's packing, except that all tiny items are removed from their bins and packed separately in new bins using the FF strategy. By the definition of bad bins, the total size of the tiny items in each bin is at most $5\varepsilon'$, so the total size of the tiny items in all bins is at most $5\varepsilon' \text{OPT}(\sigma)$. In P , the level of each bin (except the last) opened by these items is at least $1 - 5\varepsilon'$. Consequently, the number of bins in P opened by tiny items is at most $\frac{5\varepsilon'}{1-5\varepsilon'} \text{OPT}(\sigma) + 1$. Let Q be the optimal packing for normal items. Since all normal items are larger than $1/4$, each bin of Q contains at most three items. We say a bin of Q has type i ($i \in \{1, 2, 3\}$), if it contains i normal items. Similarly, we say an item x has type i if it is packed in a type i bin. All items in type 3 bins have sizes smaller than $1/2$ (otherwise one will have size at most $1/4$ which

contradicts the assumption). Moreover, the sizes of the items in all type 1 bins (except possibly the last one) are larger than $1/2$ (otherwise a better packing is achieved by pairing two of them). With two bits of advice, we can detect the type of an item as follows: Let b denote the two bits of advice with item x . If b is “01” and $x > 1/2$, then x has type 1; if b is “01” and $x \leq 1/2$, then x has type 3; and if b is “10” or b is “11”, then x has type 2. Note that the code “00” is not used at this point (this is used later on), and the use of “10” and “11” is still to be detailed.

Let X_i denote the number of bins of type i ($1 \leq i \leq 3$). Hence, the number of bins used by Q is $X_1 + X_2 + X_3$, and consequently the number of bins used by P is at least $X_1 + X_2 + X_3 + X'$, where X' is the number of bins filled by tiny items. Consider an algorithm A that performs as follows. If an item x has type 1, A simply opens a new bin for x . If x has type 2, A applies the strategy of Lemma 5 to place it in one of the bins maintained for items of type 2. Recall that the advice in this case is either “10” or “11”, so the second bit provides the advice required by Lemma 5. If x has type 3, A applies the Harmonic strategy to pack the item in a set of bins maintained for type 3 items. By Lemma 6, the number of bins used by A for these items is at most $4/3X_3 + 3$. Finally, A uses the FF strategy to pack tiny items in separate bins. Consequently, the number of bins used by the algorithm is at most $X_1 + X_2 + 4/3X_3 + X' + 3 \leq (\frac{5\varepsilon'}{1-5\varepsilon'}) \text{OPT}(\sigma) + X_3/3 + 4 \leq (4/3 + \frac{5\varepsilon'}{1-5\varepsilon'}) \text{OPT}(\sigma) + 4$. \square

Provided with the above results, we arrive at the following result:

Theorem 5 *There is an online algorithm which receives two bits of advice per item, plus an additive lower order term, and achieves a competitive ratio of $4/3 + \varepsilon$, for any positive value of ε .*

Proof Define ε' to be $\frac{11\varepsilon}{60}$. For $\varepsilon < 1/11$, we have $\varepsilon' < 1/60$. Moreover, we have $\frac{5\varepsilon'}{1-5\varepsilon'} \leq \frac{5\varepsilon'}{1-1/12} = \frac{60\varepsilon'}{11} = \varepsilon$. In an optimal packing, divide bins into good and bad bins using Definition 2. Also, let Gd and Bd respectively denote the number of good and bad bins. Use advice bits to distinguish items which are packed in good and bad bins, and pack them in separate lists of bins. More precisely, let the two bits of advice for an item x be “00” if it is packed by OPT in a good bin, and apply Lemma 4 to pack these items in at most $4/3\text{Gd}$ bins. Similarly, apply Lemma 7 to pack items from bad bins in at most $(4/3 + \frac{5\varepsilon'}{1-5\varepsilon'})\text{Bd} + 4 \leq (4/3 + \varepsilon)\text{Bd} + 4$ bins, using bits of advice of the form “01”, “10”, or “11”, as discussed in the proof of Lemma 7. Consequently, the number of bins used by the algorithm will be at most $4/3\text{Gd} + (4/3 + \varepsilon)\text{Bd} + 4 \leq (4/3 + \varepsilon)\text{OPT}(\sigma) + 4$. \square

5 A Lower Bound for Linear Advice

The *GMP problem* [19] and the *String Guessing Problem* [10] both contain a core special case of guessing a binary sequence. We use their results to show

that an online algorithm needs a linear number of bits of advice to achieve a competitive ratio better than $9/8$ for bin packing.

Definition 3 ([19,10]) The *Binary String Guessing Problem with known history (2-SGKH)* is the following online problem. The input is of the form $(n, \sigma = \langle x_1, x_2, \dots, x_n \rangle)$. It consists of n items that are either “0” or “1” and that are revealed one by one. For each item x_t , the online algorithm \mathbb{A} must guess if it is a “0” or a “1”. After the algorithm has made a guess, the value of x_t is revealed to the algorithm.

Lemma 8 ([10]) *On any input of length n , any deterministic algorithm for 2-SGKH that is guaranteed to guess correctly on more than αn bits, for $1/2 \leq \alpha < 1$, needs to read at least $(1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n$ bits of advice.*

Since the number of bits needed to express the number of “0”s in the input is at most $\lceil \log(n + 1) \rceil \leq \log n + 1$, and this number can be given as advice by an oracle, if it is not given to the algorithm otherwise, we easily obtain the following lemma. Recall that the definition of e , the length of the encoding function, is given in Section 1.1.

Lemma 9 *Consider instances of size n of the 2-SGKH problem in which the number of “0”s is given to the algorithm as part of the input. For these instances, any deterministic algorithm that is guaranteed to guess correctly on more than αn bits, for $1/2 \leq \alpha < 1$, needs to read at least $(1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n - e(n)$ bits of advice.*

Proof Assume to the contrary that the statement is not true. Hence, there is an algorithm, BSGA, that knows the number of “0”s and receives fewer than $(1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n - e(n)$ bits of advice while guessing correctly on more than αn bits. This algorithm can be used to serve arbitrary instances of the 2-SGKH problem (in which the number of “0”s is not known). Modify the advice tape used by the algorithm BSGA so that it contains at most $e(n)$ additional bits at the beginning specifying the number of “0”s. (This can be done with the self-delimited encoding of the number of “0”s.) The algorithm for 2-SGKH reads this number and gives it to BSGA. Then it asks BSGA for its guess for each bit in the sequence and answers the same as BSGA. It also informs BSGA of when it is correct and when it is wrong, with the same information it is given. The algorithm is correct exactly when BSGA is correct. The total number of advice bits will be less than $e(n) + (1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n - e(n) = (1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n$. However, Lemma 8 implies that no algorithm can guess correctly on more than αn bits with this many bits of advice. In conclusion, the initial assumption is incorrect and the statement holds. \square

In order to relate the Binary String Guessing Problem to the online bin packing problem, we introduce another problem called the Binary Separation Problem.

Algorithm 1 Implementing Binary String Guessing via Binary Separation.

The Binary Guessing algorithm knows the number of “0”s (n_1) and passes it as a parameter (the number of large items) to the Binary Separation algorithm

```

1: small = 0; large = 1
2: repeat
3:   mid = (large - small) / 2
4:   class_guess = SeparationAlgorithm.ClassifyThis(mid)
5:   if class_guess = “large” then
6:     bit_guess = 0
7:   else
8:     bit_guess = 1
9:   actual_bit = Guess(bit_guess) {The actual value is received after guessing (2-SGKH).}

10:  if actual_bit = 0 then
11:    large = mid {We let “large” be the correct decision.}
12:  else
13:    small = mid {We let “small” be the correct decision.}
14: until end of sequence

```

Definition 4 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 items into a subset of n_1 *large* items and a subset of n_2 *small* items, so that all large items are larger than all small items. Upon receiving an item y_i , an online algorithm for the problem must guess if y belongs to the set of small or large items. After the algorithm has made a guess, it is revealed to the algorithm whether y_i actually belongs to class of small or large items.

We provide reductions from the modified Binary String Guessing Problem to the Binary Separation Problem, and from the Binary Separation Problem to the online bin packing problem. In order to reduce a problem P_1 to another problem P_2 , given an instance of P_1 defined by a sequence σ_1 and a set of parameters η_1 (such as the length of σ_1 or the number of “0”s in it), we create an instance of P_2 which is defined by a sequence σ_2 and also a set of parameters η_2 . In our reductions, we assume η_2 is derived from η_1 , and since σ_1 is revealed in an online manner, σ_2 is created in an online manner by looking only at η_1 and the revealed items of σ_1 .

Lemma 10 *Assume that there is an online algorithm that solves the Binary Separation Problem on sequences of length n with $b(n)$ bits of advice, and makes at most $r(n)$ mistakes. Then there is also an algorithm that solves the Binary String Guessing Problem on sequences of length n , assuming the number of “0”s is given as a part of input, so that the algorithm receives $b(n)$ bits of advice and makes at most $r(n)$ errors.*

Proof We assume that we have an algorithm BSA that solves the Binary Separation Problem under the conditions of the lemma statement. Using that algorithm, we define the number n_1 of large items to be the number of “0”s in the instance of the Binary String Guessing Problem. Then, we implement our

algorithm BSGA for the Binary String Guessing Problem as outlined in Algorithm 1, which defines the reduction. This BSGA implementation, defined in Algorithm 1, functions as an adversary for BSA, e.g., in Line 4, BSGA gives BSA its next request. Notice that we ensure that the BSGA makes a correct guess if and only if BSA makes a correct guess. The advice tape is filled with bits of advice for this combined algorithm. The BSGA uses the BSA as a sub-routine, but all the questions are effectively coming from the BSA.

The set-up, reminiscent of binary search, is carried out as specified in the algorithm with the purpose of ensuring that when the BSA is informed of the actual class of the item it considered, no result can contradict information already obtained. Specifically, the next item for the BSA to consider is always in between the largest item which has previously been deemed “small” and the smallest item which has previously been deemed “large”. The fact that we give the middle item from that interval is unimportant; any value chosen from the open interval would work. \square

Now, we prove that if we can solve a special case of the bin packing problem, we can also solve the Binary Separation Problem.

Lemma 11 *Consider the bin packing problem on sequences of length $2n$ for which OPT opens n bins. Assume that there is an online algorithm \mathbb{A} that solves the problem on these instances with $b(n)$ bits of advice and opens at most $n + r(n)/4$ bins. Then there is also an algorithm BSA that solves the Binary Separation Problem on sequences of length n with $b(n)$ bits of advice and makes at most $r(n)$ errors.*

Proof In the reduction, we encode requests for the BSA as items for bin packing. 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 items ($n_1 + n_2 = n$), and the values of y_t s are revealed in an online manner ($1 \leq t \leq n$). We create an instance of the bin packing problem which has length $2n$. Algorithm 2 shows the details of the reduction. The bin packing sequence starts with n_1 items of size $\frac{1}{2} + \varepsilon_{min}$ (in Algorithm 2, the variable “NumberOfLargeItems” is n_1 from the Binary Separation Problem). Any algorithm needs to open a bin for each of these n_1 items. We create the next n items 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 item of the Binary Separation Problem; we ask the bin packing algorithm to pack an item whose size is an increasing function of τ , and slightly less than $\frac{1}{2}$. Depending on the decision of the bin packing algorithm for opening a new bin or placing the item in one of the existing bins, we decide the type of τ as being consecutively small or large. The last n_2 items of the bin packing instance are defined as complements of the items in the bin packing instance associated with small items 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, but we need them for the proof of the quality of the correspondence that we are proving.

Algorithm 2 Implementing Binary Separation via Special Case Bin Packing.

```

1: Choose  $\varepsilon_{min}$  and  $\varepsilon_{max}$  so that  $0 < \varepsilon_{min} < \varepsilon_{max} < \frac{1}{6}$ 
2: Choose a decreasing function  $f: \mathbb{R} \rightarrow (\varepsilon_{min}, \varepsilon_{max})$ 
3: for  $i = 1$  to  $\text{NumberOfLargeItems}$  do
4:    $\text{BinPacking.Treat}(\frac{1}{2} + \varepsilon_{min})$  {The decision can only be to open a bin.}
5: repeat
6:   Let  $\tau$  be the next request
7:    $\text{decision} = \text{BinPacking.Treat}(\frac{1}{2} - f(\tau))$  {Placing an item of size  $\frac{1}{2} - f(\tau)$ .}
8:   if  $\text{decision} = \text{"packed with an } \frac{1}{2} + \varepsilon_{min} \text{ item"}$  then
9:      $\text{class\_guess} = \text{"large"}$ 
10:  else
11:     $\text{class\_guess} = \text{"small"}$ 
12:     $\text{actual\_class} = \text{Guess}(\text{class\_guess})$ 
13:    if  $\text{actual\_class} = \text{"small"}$  then
14:       $\text{SmallItems.append}(\frac{1}{2} - f(\tau))$  {Collecting small items for later.}
15:  until end of request sequence
16: for  $i = 1$  to  $\text{length}(\text{SmallItems})$  do
17:    $\text{BinPacking.Treat}(1 - \text{SmallItems}[i])$  {The decision is not used.}

```

Call an item in the bin packing sequence “large” if it is associated with large items in the Binary Separation Problem, and “small” otherwise. For the bin packing sequence produced by the reduction, an optimal algorithm pairs each of the large items with one of the first n_1 items (those with size $\frac{1}{2} + \varepsilon_{min}$), placing them in the first n_1 bins. OPT pairs the small items with their complements, starting one of the next n_2 bins with each of these small items. Hence, the number of bins in an optimal packing is $n_1 + n_2 = n$. The values ε_{min} and ε_{max} in Algorithm 2 must be small enough so that no more than two of any of the items given in the algorithm can fit together in a bin. No other restriction is necessary.

We claim that each extra bin used by the bin packing algorithm, but not by OPT, results in at most four mistakes made by the derived algorithm on the given instance of the Binary Separation Problem. Consider an extra bin in the final packing of \mathbb{A} . This bin is opened by a large item which is incorrectly guessed as being small (bins which are opened by small items also appear in OPT’s packing). Note that large items do not fit in the same bins as complements of small items. The extra bin has enough space for another large item. Moreover, there are at most two small items which are incorrectly guessed as being large and placed in the space dedicated to the large items of the extra bin. Hence, there is an overhead of at least one for four mistakes. To summarize, \mathbb{A} has to decide if a given item is small or large and performs accordingly, and it pays a cost of at least $1/4$ for each incorrect decision. If \mathbb{A} opens at most $n + r(n)/4$ bins, the algorithm derived from \mathbb{A} for the Binary Separation Problem makes at most $r(n)$ mistakes. \square

Theorem 6 *Consider the online bin packing problem on sequences of length n . To achieve a competitive ratio of c ($1 < c < 9/8$), an online algorithm needs to receive at least $(n(1 + (4c - 4)\log(4c - 4) + (5 - 4c)\log(5 - 4c)) - (\lceil \log(n + 1) \rceil + 2\lceil \log(\lceil \log(n + 1) \rceil + 1) \rceil + 1))/2$ bits of advice.*

Proof Consider a bin packing algorithm \mathbb{A} that receives $b(n)$ bits of advice and achieves a competitive ratio of c . This algorithm opens at most $(c-1)\text{OPT}(\sigma)$ bins more than OPT , so when $\text{OPT}(\sigma) = n/2$, it opens at most $(c-1)n/2$ more bins. By Lemma 11, the existence of such an algorithm implies that there is an algorithm \mathbb{A} that solves the Binary Separation Problem on sequences of length $n/2$ with b bits of advice and makes at most $2(c-1)n$ errors. By Lemma 10, this implies that there is an algorithm \mathbb{B} that solves the Binary String Guessing Problem on sequences of length $n/2$ with b bits of advice and makes at most $2(c-1)n$ mistakes, i.e., it correctly guesses the other $n/2 - 2(c-1)n = (5-4c)n/2$ items. Let $\alpha = 5 - 4c$, and note that α is in the range $[1/2, 1)$ when c is in the range $(1, 9/8]$. Lemma 9 implies that in order to correctly guess more than $\alpha n/2$ of the items in the binary sequence, we must have $b(n)$ larger than or equal to $((1 + (1 - \alpha) \log(1 - \alpha) + \alpha \log \alpha)n - e(n))/2$. Replacing α with $5 - 4c$ completes the proof. \square

Thus, to obtain a competitive ratio strictly better than $9/8$, a linear number of bits of advice is required. For example, to achieve a competitive ratio of $17/16$, at least $0.188n$ bits of advice are required asymptotically.

Corollary 1 *Consider the bin packing problem of packing sequences of length n . To achieve a competitive ratio of $9/8 - \delta$, in which δ is a small, but fixed positive number, an online algorithm needs to receive $\Omega(n)$ bits of advice.*

6 Concluding Remarks

We conjecture that a sublinear number of bits of advice is enough to achieve competitive ratios smaller than $4/3$. Note that our results imply that we cannot hope for ratios smaller than $9/8$ with sublinear advice.

The lower bound presented here does not give a result better than half a bit per item asymptotically, regardless of how close to a competitive ratio of 1 we wish to obtain. It would be interesting to strengthen this result.

References

1. Albers, S., Hellwig, M.: Semi-online scheduling revisited. *Theoretical Computer Science* **443**, 1–9 (2012)
2. Ásgeirsson, E.I., Ayesta, U., Coffman, E.G., Etra, J., Momcilovic, P., Phillips, D.J., Vokhshoori, V., Wang, Z., Wolfe, J.: Closed on-line bin packing. *Acta Cybernetica* **15**(3), 361–367 (2002)
3. Balogh, J., Békési, J.: Semi-on-line bin packing: a short overview and a new lower bound. *Central European Journal of Operations Research* **21**(4), 685–698 (2013)
4. Balogh, J., Békési, J., Galambos, G.: New lower bounds for certain classes of bin packing algorithms. *Theoretical Computer Science* **440–441**, 1–13 (2012)
5. Bianchi, M.P., Böckenhauer, H., Hromkovic, J., Keller, L.: Online coloring of bipartite graphs with and without advice. *Algorithmica* **70**(1), 92–111 (2014)
6. Bianchi, M.P., Böckenhauer, H.J., Hromkovic, J., Keller, L.: Online coloring of bipartite graphs with and without advice. In: J. Gudmundsson, J. Mestre, T. Viglas (eds.) *Proc. 18th Computing and Combinatorics Conference (COCOON)*, *Lecture Notes in Computer Science*, Springer, vol. 7434, pp. 519–530 (2012)

7. Bianchi, M.P., Böckenhauer, H.J., Hromkovic, J., Krug, S., Steffen, B.: On the advice complexity of the online $L(2,1)$ -coloring problem on paths and cycles. In: D. Du, G. Zhang (eds.) Proc. 19th Computing and Combinatorics Conference (COCOON), *Lecture Notes in Computer Science, Springer*, vol. 7936, pp. 53–64 (2013)
8. Błażewicz, J., Ecker, K.: A linear time algorithm for restricted bin packing and scheduling problems. *Operations Research Letters* **2**(2), 80–83 (1983)
9. Böckenhauer, H., Komm, D., Královic, R., Rossmann, P.: The online knapsack problem: Advice and randomization. *Theoretical Computer Science* **527**, 61–72 (2014)
10. Böckenhauer, H.J., Hromkovic, J., Komm, D., Krug, S., Smula, J., Sprock, A.: The string guessing problem as a method to prove lower bounds on the advice complexity. In: D. Du, G. Zhang (eds.) Proc. 19th Computing and Combinatorics Conference (COCOON), *Lecture Notes in Computer Science, Springer*, vol. 7936, pp. 493–505 (2013)
11. Böckenhauer, H.J., Komm, D., Královic, R., Královic, R.: On the advice complexity of the k -server problem. In: L. Aceto, M. Henzinger, J. Sgall (eds.) Proc. 38th International Colloquium on Automata, Languages, and Programming (ICALP), *Lecture Notes in Computer Science, Springer*, vol. 6755, pp. 207–218 (2011)
12. Böckenhauer, H.J., Komm, D., Královic, R., Královic, R., Mömke, T.: On the advice complexity of online problems. In: Y. Dong, D. Du, O.H. Ibarra (eds.) Proc. 20th International Symposium on Algorithms and Computation (ISAAC), *Lecture Notes in Computer Science, Springer*, vol. 5878, pp. 331–340 (2009)
13. Boyar, J., Kamali, S., Larsen, K.S., López-Ortiz, A.: On the list update problem with advice. In: A.H. Dediu, C. Martín-Vide, J.L. Sierra-Rodríguez, B. Truthe (eds.) Proc. 8th International Conference on Language and Automata Theory and Applications (LATA), pp. 210–221 (2014)
14. Cheng, T.C.E., Kellerer, H., Kotov, V.: Semi-on-line multiprocessor scheduling with given total processing time. *Theoretical Computer Science* **337**(1-3), 134–146 (2005)
15. Coffman, E.G., Garey, M.R., Johnson, D.S.: Approximation algorithms for bin packing: A survey. In: D. Hochbaum (ed.) *Approximation algorithms for NP-hard Problems*. PWS Publishing Co. (1997)
16. Coffman Jr., E.G., Csirik, J., Galambos, G., Martello, S., Vigo, D.: Bin packing approximation algorithms: survey and classification. In: P.M. Pardalos, D.Z. Du, R.L. Graham (eds.) *Handbook of Combinatorial Optimization*, pp. 455–531. Springer (2013)
17. Dobrev, S., Královic, R., Markou, E.: Online graph exploration with advice. In: G. Even, M.M. Halldórsson (eds.) Proc. 19th International Colloquium on Structural Information and Communication Complexity (SIROCCO), *Lecture Notes in Computer Science, Springer*, vol. 7355, pp. 267–278 (2012)
18. Dobrev, S., Královic, R., Pardubská, D.: Measuring the problem-relevant information in input. *RAIRO – Theoretical Informatics and Applications* **43**(3), 585–613 (2009)
19. Emek, Y., Fraigniaud, P., Korman, A., Rosén, A.: Online computation with advice. *Theoretical Computer Science* **412**(24), 2642–2656 (2011)
20. Epstein, L., Levin, A.: On bin packing with conflicts. *SIAM Journal on Optimization* **19**(3), 1270–1298 (2008)
21. Forišek, M., Keller, L., Steinová, M.: Advice complexity of online coloring for paths. In: A.H. Dediu, C. Martín-Vide (eds.) Proc. 6th International Conference on Language and Automata Theory and Applications (LATA), *Lecture Notes in Computer Science, Springer*, vol. 7183, pp. 228–239 (2012)
22. Galambos, G., Woeginger, G.J.: Repacking helps in bounded space online bin packing. *Computing* **49**, 329–338 (1993)
23. Gambosi, G., Postiglione, A., Talamo, M.: Algorithms for the relaxed online bin-packing model. *SIAM Journal on Computing* **30**(5), 1532–1551 (2000)
24. Goemans, M.X., Rothvoß, T.: Polynomiality for bin packing with a constant number of item types. In: C. Chekuri (ed.) Proc. 25th Symposium on Discrete Algorithms (SODA), pp. 830–839 (2014)
25. Grove, E.F.: Online binpacking with lookahead. In: K.L. Clarkson (ed.) Proc. 6th Symposium on Discrete Algorithms (SODA), pp. 430–436 (1995)
26. Hromkovič, J., Královic, R., Královic, R.: Information complexity of online problems. In: P. Hliněný, A. Kucera (eds.) Proc. 35th Symposium on Mathematical Foundations of Computer Science (MFCS), *Lecture Notes in Computer Science, Springer*, vol. 6281, pp. 24–36 (2010)

27. Ivkovic, Z., Lloyd, E.: A fundamental restriction on fully dynamic maintenance of bin packing. *Information Processing Letters* **59**, 229–232 (1996)
28. Komm, D., Královic, R., Mömke, T.: On the advice complexity of the set cover problem. In: E.A. Hirsch, J. Karhumäki, A. Lepistö, M. Prilutskii (eds.) *Proc. 7th International Computer Science Symposium in Russia (CSR)*, *Lecture Notes in Computer Science*, Springer, vol. 7353, pp. 241–252 (2012)
29. Komm, D., Královic, R.: Advice complexity and barely random algorithms. *RAIRO – Theoretical Informatics and Applications* **45**(2), 249–267 (2011)
30. Renault, M.P., Rosén, A.: On online algorithms with advice for the k -server problem. In: R. Solis-Oba, G. Persiano (eds.) *Proc. 9th International Workshop in Approximation and Online Algorithms (WAOA)*, *Lecture Notes in Computer Science*, Springer, vol. 7164, pp. 198–210 (2011)
31. Renault, M.P., Rosén, A., van Stee, R.: Online algorithms with advice for bin packing and scheduling problems. *CoRR* **abs/1311.7589** (2013)
32. Seibert, S., Sprock, A., Unger, W.: Advice complexity of the online coloring problem. In: P.G. Spirakis, M.J. Serna (eds.) *Proc. 8th International Conference on Algorithms and Complexity (CIAC)*, *Lecture Notes in Computer Science*, Springer, vol. 7878, pp. 345–357 (2013)
33. Seiden, S.S.: On the online bin packing problem. *Journal of the ACM* **49**, 640–671 (2002)
34. Tan, Z., Zhang, A.: Online and semi-online scheduling. In: P.M. Pardalos (ed.) *Handbook of Combinatorial Optimization*, pp. 2191–2252. Springer (2013)
35. Vazirani, V.V.: *Approximation Algorithms*. Springer (2004)
36. de la Vega, W.F., Lueker, G.S.: Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica* **1**, 349–355 (1981)