Alcom-FT Technical Report Series
ALCOMFT-TR-02-4

Extending the Accommodating Function*

Joan Boyar, Lene M. Favrholdt, Kim S. Larsen, and Morten N. Nielsen
University of Southern Denmark, Odense, Denmark
{joan,lenem,kslarsen,nyhave}@imada.sdu.dk

Abstract

The applicability of the accommodating function, a relatively new measure for the quality of on-line algorithms, is extended. If a limited amount n of some resource is available, the accommodating function $\mathcal{A}(\alpha)$ is the competitive ratio when input sequences are restricted to those for which the amount αn of resources suffices for an optimal off-line algorithm. The standard competitive ratio is $\lim_{\alpha\to\infty} \mathcal{A}(\alpha)$. The accommodating function was originally used only for $\alpha \geq 1$. We focus on $\alpha < 1$, observe that the function now appears interesting for a greater variety of problems, and use it to make new distinctions between known algorithms and to find new ones.

1 Introduction

1.1 The Accommodating Function

The accommodating function $\mathcal{A}$ is a new performance measure for on-line optimization problems with a limited amount n of some resource. Informally, $\mathcal{A}(\alpha)$ is the competitive ratio when input sequences are restricted to those for which an optimal off-line algorithm does not benefit from having more than the amount αn of resources. These sequences are called α-sequences. Note that $\lim_{\alpha\to\infty} \mathcal{A}(\alpha)$ is the standard competitive ratio. The accommodating function was recently defined in [6], and it was applied to various problems in [6] and [1], but only for $\alpha \geq 1$. In this paper, values of $\alpha < 1$ are considered for the first time. The accommodating function is formally defined in Section 2.

1.2 Motivation and Background

The original motivation for considering this type of restriction of request sequences is from the Seat Reservation Problem [5], the problem of assigning seats to passengers in a train on-line, in a "fair" manner, to maximize earnings. For the unit price version, the competitive

*Supported in part by the Danish Natural Science Research Council (SNF) and in part by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

ratio for this problem is $\Theta(\frac{1}{k})$, where k is the number of stations where the train stops. This very discouraging performance cannot occur, however, for realistic request sequences. Since the management is often able to judge approximately how many cars are necessary to accommodate all requests, it is more realistic to consider only request sequences that can be fully accommodated by an optimal off-line algorithm. Such sequences, corresponding to $\alpha = 1$, are called accommodating sequences. For the unit price problem, $\mathcal{A}(1) \geq \frac{1}{2}$ [5].

The idea of restricting the adversary to only giving accommodating sequences carries over to any optimization problem with some limited resource, such as the seats in a train. Thus, for instance, it can be used on the k-Server Problem, where the servers constitute a limited resource since there are only k of them, or scheduling problems, where there is only a fixed number of machines.

In addition to giving rise to new interesting algorithmic and analytical problems, the accommodating function, compared to just one ratio, contains more information about the on-line algorithms. This information can be exploited in several ways. The shape of the function, for instance, can be used to warn against critical scenarios, where the performance of the on-line algorithm compared to the off-line can suddenly drop rapidly when fewer resources are available.

In [6] a variant of bin packing is investigated, for which the number of bins is fixed and the goal is to maximize the number of items packed. The bins have height k and the items are integer sized. It is shown that, in general, Worst-Fit has a strictly better competitive ratio than First-Fit, while, in the special case of accommodating sequences, First-Fit has a strictly better competitive ratio than Worst-Fit. In this case, the competitive ratio on accommodating sequences seems the more appropriate measure, since it is constant while, in the general case, the competitive ratio is $\Theta(\frac{1}{k})$, basically due to some sequences which seem very contrived. This shows that in addition to giving more realistic performance measures for some problems, the competitive ratio on accommodating sequences can be used to distinguish between algorithms, showing, not surprisingly, that the decision as to which algorithm should be used depends on what sort of request sequence is expected. The obvious question at this point was: Where is the cross-over point? When does First-Fit become better than Worst-Fit? This is another motivation for considering the accommodating function.

Main Contribution In [6], the accommodating function was investigated for $\alpha \geq 1$, thus spanning from the one known interesting case (accommodating sequences) to the other (the standard competitive ratio). In this paper we extend the definition of the accommodating function to include $\alpha < 1$. As example problems, we first consider two maximization problems. For Unrestricted Bin Packing, we show upper and lower bounds for First-Fit and investigate a new variant of Unfair-First-Fit [1], called Unfair-First-Fit$_\alpha$, which turns out to be better than First-Fit for α close to 1. For Seat Reservation, we consider three deterministic algorithms, which asymptotically have the same competitive ratio on accommodating sequences, and show that these algorithms can be distinguished using the proposed extension of the accommodating function. Finally, we consider well known on-

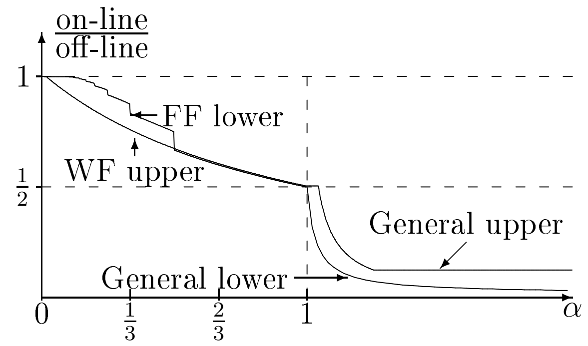

Figure 1: Bounds on $\mathcal{A}(\alpha)$ for the Seat Reservation Problem.

line minimization problems to emphasize the broad applicability of the accommodating function for $\alpha < 1$.

1.3 The Accommodating Function for Two Maximization Problems

For the Seat Reservation Problem, considering the accommodating function for $\alpha < 1$ corresponds to the situation where the management has provided more cars than needed by an optimal off-line algorithm to accommodate all passenger requests. This seems to be a desirable situation, since the only way the train company can hope to accommodate all requests on-line is by having more resources than would be necessary if the problem was solved optimally off-line. In this paper, we obtain positive and negative results on the accommodating function for the problem in general and for First-Fit and Worst-Fit in particular. The results for First-Fit and Worst-Fit are depicted in Figure 1 to the left of the line $\alpha = 1$. To the right of this line, general results from [6] are shown. Although the competitive ratio on accommodating sequences fails to distinguish between fair algorithms for the Seat Reservation Problem (they all have ratio $\frac{1}{2}$ in the limit), we investigate the accommodating function at $\alpha = \frac{1}{3}$ for three different algorithms and discover that Worst-Fit is the worst there, First-Fit is better, and Kierstead-Trotter has competitive ratio 1 at $\alpha = \frac{1}{3}$.

We also consider another maximization problem, a variant of bin packing called Unrestricted Bin Packing. We obtain positive and negative results on the accommodating function for the problem in general and for First-Fit in particular. The results for First-Fit are depicted in Figure 2 to the left of the line $\alpha = 1$. The results to the right of the line are from [6]. Moreover, we show how positive results for accommodating sequences can be extended in a fairly general manner to better bounds on the competitive ratio for more restricted sets of input sequences. Thus, accommodating sequences play a more important role than originally anticipated. One can also see from the graphs in Figures 1 and 2 that accommodating sequences have unique properties; the shape of the curve seems to often change significantly at $\alpha = 1$. In addition, for Unrestricted Bin Packing, we investi-

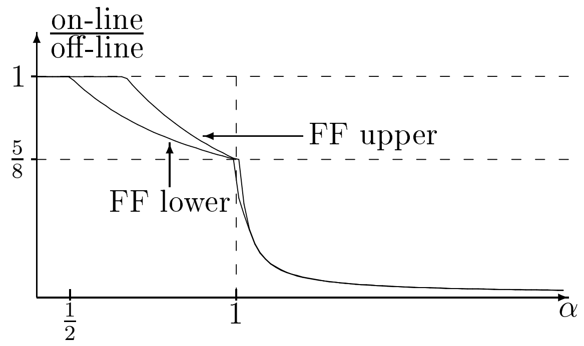

Figure 2: Bounds on $\mathcal{A}_{\mathrm{FF}}(\alpha)$ for Unrestricted Bin Packing.

gate a variant called Unfair-First-Fit$_\alpha$, which is shown to perform significantly better than First-Fit given that a good bound on α is known in advance.

1.4 The Connection to Resource Augmentation

Resource augmentation is another technique which is used to give more realistic results when the standard competitive ratio seems too negative [11]. With resource augmentation analysis, the on-line algorithm is given more resources than the off-line algorithm, but the performance ratio is still the worst case over all request sequences. There is clearly a similarity between resource augmentation and the accommodating function with $\alpha < 1$. However, for some problems there is a tremendous difference between the results one gets with the two techniques. In these cases, the accommodating function gives the more positive results, because there can exist sequences where the off-line algorithm could have done somewhat better with more resources (and thus would not be considered in the accommodating function analysis), while the on-line algorithm is unable to fully take advantage of its extra resources. As an example of this phenomenon, for $\alpha \leq 1$, we prove a performance guarantee for the algorithm First-Fit for the Seat Reservation Problem, $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq 1 - \frac{1}{2^l}$, where $l = \lfloor \frac{1}{\alpha} \rfloor$, but even with $\frac{1}{\alpha}$ times as many seats as the off-line algorithm, First-Fit's competitive ratio can be as bad as $\frac{1+\alpha}{(\alpha-2/n)(k-1)}$; the limit is zero as the number of stations, k, approaches infinity. For Unrestricted Bin Packing, one can also show that resource augmentation analysis gives results which are much more pessimistic than the corresponding accommodating function results. We prove a performance guarantee for First-Fit for $\frac{1}{2} \leq \alpha \leq 1$ of $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq \frac{3+2\alpha}{8\alpha}$, while the competitive ratio of First-Fit is at most $\frac{1}{\alpha k}$, even when First-Fit has $\frac{1}{\alpha}$ times as many bins as the off-line algorithm (again k is the height of the bins).

On the other hand, there are many cases where results or proofs from resource augmentation can be used to calculate the accommodating function for $\alpha < 1$. Positive results about algorithms analyzed using resource augmentation directly give positive results for the accommodating function. Negative results can also be applied if the original adversary only used accommodating sequences. We illustrate this with three minimization problems:

the Paging Problem, the k-Server Problem, and Scheduling (minimizing the makespan). The clear intuitive applicability of the accommodating function when the actual request sequences are expected to have the proper form makes these results and proofs from resource augmentation even more interesting than previously.

2 The Accommodating Function

If $\mathbb{A}$ is an on-line algorithm for a maximization (minimization) problem, then, for any input sequence I for the problem, $\mathbb{A}(I)$ denotes the value (cost) of running $\mathbb{A}$ on I. For any maximization (minimization) problem, OPT denotes an optimal off-line algorithm for the problem, and $\mathrm{OPT}(I)$ denotes the value (cost) of OPT when run on I. For a problem with some limited resource, OPT_n denotes the value (cost) of OPT when the amount n of the limited resource is available.

Definition 1 Let P be an on-line problem with a fixed amount n of resources. For any $\alpha > 0$, an input sequence I is said to be an *α-sequence*, if $\mathrm{OPT}_{\alpha n}(I) = \mathrm{OPT}_{n'}(I)$, for all $n' \geq \alpha n$. 1-sequences are also called *accommodating sequences.* □

If an input sequence is an α-sequence, then OPT does not benefit from having more than the amount αn of resources. Thus, if an input sequence is an accommodating sequence, then OPT does not benefit from having more resources than the amount already available. For many problems, the amount n of the resource is always a natural number. In that case, one should only consider values of α such that $\alpha n \in \mathbb{N}$. If not, we define $\mathcal{A}(\alpha)$ to have the same value as $\mathcal{A}(\frac{\lfloor \alpha n \rfloor}{n})$. Thus, $\mathcal{A}$ becomes a step function.

Definition 2 Let $\mathbb{A}$ be an on-line algorithm for a maximization (minimization) problem. Then $\mathbb{A}$ is *c-competitive on α-sequences*, if there exists a constant b, such that

$$\mathbb{A}(I) \geq c \cdot \mathrm{OPT}(I) - b \quad \big(\mathbb{A}(I) \leq c \cdot \mathrm{OPT}(I) + b\big),$$

for any α-sequence I. The *accommodating function* is defined as

$$\mathcal{A}_{\mathbb{A}}(\alpha) = \sup\{c \mid \mathbb{A} \text{ is } c\text{-comp on } \alpha\text{-seq}\} \ (\mathcal{A}_{\mathbb{A}}(\alpha) = \inf\{c \mid \mathbb{A} \text{ is } c\text{-comp on } \alpha\text{-seq}\}).$$

□

Note that $\lim_{\alpha \to \infty} \mathcal{A}(\alpha)$ is the standard competitive ratio.

3 Unrestricted Bin Packing

In this section we consider a maximization variant of Bin Packing, called Unrestricted Bin Packing, in which the objective is to maximize the number of items accepted within n bins all of the same size $k \in \mathbb{N}$. An input sequence consists of integer-sized items, and an α-sequence can be packed (by an optimal off-line algorithm) within αn bins.

We use the following notation for a given $\alpha \leq 1$. Assume that we have a numbering of the bins from 1 to n. The first αn bins, denoted B_A, are called *internal* bins, and the remaining $n - \alpha n$ bins, denoted B_X, are called *external* bins. Whenever we consider a fixed request sequence, we let A denote the set of items accepted within internal bins, and let X denote the set of items accepted within the external bins. The set R contains the remaining (rejected) items. The bounds on First-Fit found in this section are depicted in Figure 2.

3.1 Performance guarantees

Definition 3 An algorithm for Unrestricted Bin Packing is called *fair* if it never rejects an item that it is able to pack. □

In this section, we give a lower bound for the accommodating function for the class of fair algorithms and better bounds for three specific algorithms, two of which are fair. In fact, the bounds for the three algorithms are corollaries of a more general theorem extending results for accommodating sequences ($\alpha = 1$) to the α-sequences where α is less than 1.

Theorem 1 Any fair algorithm $\mathbb{A}$ for Unrestricted Bin Packing has $\mathcal{A}_{\mathbb{A}}(\alpha) \geq \frac{1}{1+\alpha-\frac{1}{k}}$, when $\alpha \leq 1$.

Proof Given an $\alpha \leq 1$, let I be an α-sequence. We assume that R is non-empty, since otherwise the result of $\mathbb{A}$ is optimal. Thus, $|A \cup X| \geq n$ since $\mathbb{A}$ is fair. Let s denote the size of a smallest item in R. Since $\mathbb{A}$ is fair, the empty space in each bin in B_A and B_X is at most $s - 1$, meaning that the items accepted have a total size of at least $n(k - s + 1)$. Since I is an α-sequence, the total size of the items in I is at most αnk. Thus, $|R| \leq \frac{1}{s}(\alpha nk - n(k - s + 1)) \leq \frac{1}{s}(\alpha nk - \alpha n(k - s) - n) = \alpha n - \frac{n}{s} \leq \alpha n - \frac{n}{k}$, since $\alpha \leq 1$ and $s \leq k$. □

In Section 3.2 this performance guarantee is shown to be asymptotically tight due to Worst-Fit's behavior.

The proof of Theorem 1 shows that fair algorithms will reject less than αn items. When considering First-Fit, a stronger fact can be shown, namely that αn is an upper bound on the number of items not packed in the internal bins, i.e., $|R| + |X| < \alpha n$. Similar results can be shown for a broader class of algorithms.

Definition 4 Consider any α-sequence I with $\alpha \leq 1$ and an algorithm $\mathbb{A}$. Let X' be the set of items in I which $\mathbb{A}$ places as the first in some external bin, so $X' \subset X$. Define $f(I)$ to be the subsequence of I consisting of those requests in $A \cup X' \cup R$. An algorithm $\mathbb{A}$ is said to be *uniform* if, for any sequence I, when $\mathbb{A}$ processes $f(I)$ using αn bins (and any α the algorithm may use internally set to 1), it accepts exactly those items in A. □

Lemma 1 For any fair, uniform algorithm $\mathbb{A}$ and any α-sequence with $\alpha \leq 1$, $|R| + |X'| < \alpha n$ and whenever $|R| > 0$, $|X'| = n - \alpha n$.

Proof Consider $\mathbb{A}$'s processing of $f(I)$ with αn bins. Using the argument from the proof of Theorem 1, one gets that $|R|+|X'| < \alpha n$. Since $\mathbb{A}$ is fair, if $|R| > 0$, then $|X'| = n-\alpha n$. $\square$

Note that, by Lemma 1, any fair, uniform algorithm is optimal when $\alpha \leq \frac{1}{2}$. This can be seen in the following way. For $\alpha \leq \frac{1}{2}$, $|R| + |X'| < \alpha n$ and $|X'| = n - \alpha n$ cannot both be fulfilled. Thus, for any fair, uniform algorithm, $|R| = 0$ when $\alpha \leq \frac{1}{2}$.

First-Fit is a fair, uniform algorithm, as are other algorithms, such as Best-Fit. For such algorithms, the theorem below extends good performance guarantees obtained on accommodating sequences to results for the accommodating function for $\alpha < 1$. The theorem can also be applied to some algorithms which are not fair.

Theorem 2 Suppose the following four conditions hold for an on-line algorithm $\mathbb{A}$:

1. $\mathcal{A}_{\mathbb{A}}(1) \geq \beta$,
2. $\mathbb{A}$ is uniform,
3. $|R| + |X'| \leq \alpha n$,
4. whenever $|R| > 0$, $|X'| = n - \alpha n$.

Then for $\alpha \leq 1$, we have $\mathcal{A}_{\mathbb{A}}(\alpha) \geq \min\left(\frac{1-\alpha-\beta+2\alpha\beta}{\alpha}, 1\right)$.

Proof Let $\alpha \leq 1$ be given. Since we consider an α-sequence and $\mathbb{A}$ is uniform, the first condition implies that $\frac{|A|}{|A|+|X'|+|R|} \geq \beta$. Assume that $|R| > 0$, since otherwise $\mathbb{A}$'s packing is optimal. We have $\frac{|A|+|X|}{|A|+|X|+|R|} \geq \frac{|A|+|X'|}{|A|+|X'|+|R|} \geq \beta + \frac{|X'|}{|A|+|X'|+|R|} \geq \beta + \frac{n-\alpha n}{|A|+\alpha n}$; the last inequality holds by conditions 3 and 4. If $|A| \leq \frac{\beta\alpha n}{1-\beta}$, we get

$$\begin{aligned}\frac{|A|+|X|}{|A|+|X|+|R|} &\geq \beta + \frac{(1-\beta)(n-\alpha n)}{\beta\alpha n + (1-\beta)\alpha n} = \beta + \frac{(1-\beta)(n-\alpha n)}{\alpha n} \\ &= \frac{1-\alpha-\beta+2\alpha\beta}{\alpha}.\end{aligned}$$

If $|A| > \frac{\beta\alpha n}{1-\beta}$, we get

$$\begin{aligned}\frac{|A|+|X|}{|A|+|X|+|R|} &\geq \frac{|A|+|X'|}{|A|+|X'|+|R|} \geq \frac{|A|+n-\alpha n}{|A|+\alpha n} \\ &> \frac{\beta\alpha n + (1-\beta)(n-\alpha n)}{\beta\alpha n + (1-\beta)\alpha n} = \frac{1-\alpha-\beta+2\alpha\beta}{\alpha}.\end{aligned}$$

$\square$

The first corollaries to this theorem concern the fair algorithms First-Fit and Best-Fit.

Corollary 1 For $\alpha \leq 1$, $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq \min\left(\frac{3+2\alpha}{8\alpha}, 1\right)$.

Proof We have to verify the four conditions in Theorem 2. The first condition holds with $\beta \geq \frac{5}{8}$ [6]. It is not hard to see that First-Fit is uniform since items from $f(I)$ are put in the external bins only if they do not fit in the internal bins. The last two conditions follow from Lemma 1. $\square$

Corollary 2 For $\alpha \leq 1$, $\mathcal{A}_{\mathrm{BF}}(\alpha) \geq \min\left(\frac{3+2\alpha}{8\alpha}, 1\right)$.

Proof The first condition in Theorem 2 holds with $\beta \geq \frac{5}{8}$ [6]. The other properties are easily verified. Let I be an α-sequence and let A, B_A, X', R, and $f(I)$ be defined as earlier. Assume that when Best-Fit is given the sequence $f(I)$ and only αn bins to pack it in, one of the following two situations occur. 1.) Some item $i \in A$ is rejected or put in a bin that does not correspond to the bin it was put in in B_A. 2.) Some item $i \in X' \cup R$ is accepted. In either situation, we say that i is treated nonuniformly.

If Best-Fit is not uniform, there must exist an α-sequence I and an item $u \in I$ that is treated nonuniformly. Assume that u is the first such item in the sequence. At the time the item arrives, the configuration in the αn bins is the same as it was in B_A at the corresponding time. Thus, if $u \in A$, u is packed in the bin corresponding to the bin it was packed in in B_A, and if $u \in X' \cup R$, there was not space for u in B_A, so u will be rejected. Thus, u cannot exist.

The last two conditions follow from Lemma 1. $\square$

The third corollary to the theorem above concerns a slight variation of the algorithm Unfair-First-Fit [1]. Unfair-First-Fit behaves just like First-Fit unless the given item has size larger than $\frac{k}{2}$. In this case, the item is rejected on purpose, if the number of currently accepted items is at least $\frac{2}{3}$ of the number of items in the entire sequence seen so far. The new algorithm given in Figure 3, called Unfair-First-Fit$_\alpha$ (UFF$_\alpha$), assumes that α is known in advance. Using this knowledge, the algorithm divides the bins into B_A and B_X. Note that the ratio considered in the if-statement of the algorithm is the number of items in the internal bins, compared to all items given. If this ratio is at least $\frac{2}{3}$, the item is rejected if it does not fit in an external bin.

The actions "Accept" and "Reject" should be understood with respect to the internal bins, in that they distinguish between whether the internal bins are considered or not. In the "Accept" case, the item is placed in the internal bins according to the First-Fit packing rule, if there is enough space. If there is not enough space in any of the internal bins, the First-Fit rule is tried against the external bins. If there is not enough space in either an internal or an external bin, the item is rejected (added to R). The "Reject" case is similar, but the internal bins are not considered.

If a good upper bound α' on α is known, UFF$_{\alpha'}$ can be used, giving a lower bound of $\mathcal{A}_{\mathrm{UFF}_{\alpha'}}(\alpha')$ on the performance.

Input: $S = \langle o_1, o_2, \ldots, o_n \rangle$
while $S \neq \langle\rangle$
  $o := \mathrm{hd}(S); S := \mathrm{tail}(S)$
  if $\mathrm{size}(o) \leq \frac{k}{2}$ **or** $\frac{|A|}{|A|+|X|+|R|+1} < \frac{2}{3}$
    "Accept" - Try to pack o using First-Fit
  else
    "Reject" - Try to pack o in B_X using First-Fit
  Update A, X, R accordingly

Figure 3: The algorithm Unfair-First-Fit$_\alpha$

Corollary 3 $\mathcal{A}_{\mathrm{UFF}_\alpha}(\alpha) \geq \min\left(\frac{1+\alpha}{3\alpha} - \epsilon, 1\right)$, for $\alpha \leq 1$ and $\alpha n \geq 9$, where $\epsilon = \frac{4\alpha-2}{(4n+1)\alpha} < \frac{1}{n}$.

Proof From [1], we know that the first property from Theorem 2 holds for $\beta = \frac{2}{3} - \frac{2}{4n+1}$, for $\alpha n \geq 9$. The second property holds by the order in which the sets B_A, B_X, and R are tried. If $R = \emptyset$, the result is optimal, so we can assume that R is non-empty. Then there is at least one item in each of the external bins. This justifies property (4) of Theorem 2. Recall the definition of X' from Definition 4. To show the third property, let s denote the size of a smallest item in $X' \cup R$. If this item was packed by an "Accept" attempt, every bin in B_A has empty space less than s, and by the same calculations as in Theorem 1, we are done. If the item was packed by a "Reject" attempt, this item has size larger than $\frac{k}{2}$, which then all items in $X' \cup R$ must have. Since no two items of this size can be packed together, there are at most αn such items, since we consider an α-sequence. Finally, assuming $\mathcal{A}_{\mathrm{UFF}_\alpha}(\alpha) < 1$, we derive the bound

$$\begin{aligned}\mathcal{A}_{\mathrm{UFF}_\alpha}(\alpha) \geq \frac{1-\alpha-(\frac{2}{3}-\frac{2}{4n+1})+2\alpha(\frac{2}{3}-\frac{2}{4n+1})}{\alpha} &= \frac{1-\alpha-\frac{2}{3}+\frac{2}{3}2\alpha-\frac{2}{4n+1}(2\alpha-1)}{\alpha} \\ &= \frac{1+\alpha}{3\alpha} - \frac{4\alpha-2}{\alpha(4n+1)}.\end{aligned}$$

□

3.2 Impossibility Results

The theorem below extends a result in [6] and it shows that any (even randomized) on-line algorithm $\mathbb{A}$ for Unrestricted Bin Packing has $\mathcal{A}_{\mathbb{A}}(\alpha) < 1$, for $\alpha > \frac{4}{5}$.

Theorem 3 Any on-line algorithm $\mathbb{A}$ for Unrestricted Bin Packing has $\mathcal{A}_{\mathbb{A}}(\alpha) \leq \frac{2\alpha+4}{7\alpha}$, when $\frac{4}{5} \leq \alpha \leq 1$.

Proof Assume that $\frac{4}{5} \leq \alpha \leq 1$. The adversary gives two phases of items. In the first phase, αn items of size $\frac{k}{2} - 1$ are given. Let q denote the fraction of bins containing two items in $\mathbb{A}$'s packing. Depending on the expected value of q, $E[q]$, the next phase is different.

If $E[q] \leq \frac{3\alpha-1}{7}$, the adversary gives $\frac{\alpha n}{2}$ items of size k. By linearity of expectation, the expected ratio is at most $\frac{n+E[q]n}{\frac{3}{2}\alpha n} = \frac{2+2E[q]}{3\alpha} \leq \frac{14+2(3\alpha-1)}{21\alpha} = \frac{4+2\alpha}{7\alpha}$.

If $E[q] > \frac{3\alpha-1}{7}$, the adversary gives αn items of size $\frac{k}{2}+1$. In this case $\mathbb{A}$ accepts at most $\alpha n + n - qn$. Thus, the expected ratio is at most $\frac{\alpha n+n-E[q]n}{2\alpha n} < \frac{7\alpha+7-(3\alpha-1)}{14\alpha} = \frac{2\alpha+4}{7\alpha}$. $\square$

The following theorem extends a hardness result for First-Fit on accommodating sequences in [1] and shows that, when n is sufficiently large, $\mathcal{A}_{\mathrm{UFF}}(\alpha) \geq \mathcal{A}_{\mathrm{FF}}(\alpha)$ for $\frac{7+\sqrt{85}}{18} < \alpha \leq 1$.

Theorem 4 When $\frac{2}{3}+\frac{1}{3n} \leq \alpha \leq 1$, the accommodating function for First-Fit for Unrestricted Bin Packing is at most $\mathcal{A}_{\mathrm{FF}}(\alpha) \leq \frac{5}{9\alpha-1} + O(\frac{1}{\sqrt{n}})$. If n is of the form $n = 9 \cdot 2^i - 5$, $i \in \mathbb{N}_0$, then $\mathcal{A}_{\mathrm{FF}}(\alpha) \leq \frac{5}{9\alpha-1} + O(\frac{1}{n})$.

Proof We first show the bound for the special values of n. Then we extend to all values of n. Assume that, for some $i \in \mathbb{N}_0$, n is of the form $n = 9 \cdot 2^i - 5$ and that k is at least $12 \cdot 2^{3i}$ and divisible by 3. An adversary can give the following request sequence, divided into $i+3$ phases:

Phase 0. 3 items of size $A = \frac{k}{3} - 2^{3i}$.
Phase $j = 1 \ldots i$. $3 \cdot 2^j$ pairs, each with one item of size $B_j = \frac{k}{3} + 2^{3i-3j+2}$ followed by one of size $C_j = \frac{k}{3} - 2^{3i-3j}$.
Phase $i+1$. $3 \cdot 2^i$ items of size $D = \frac{2k}{3} + 1$.
Phase $i+2$. $3 \cdot (\alpha n - (6 \cdot 2^i - 3))$ items of size $E = \frac{k}{3}$.

The number of items given in Phase $i+2$ is non-negative, since $\alpha \geq \frac{2}{3} + \frac{1}{3n}$. First-Fit will pack the items of Phase 0 in one bin. The assumption on k ensures that these three items will remain alone, since the smallest item given after Phase 0 has size $\frac{k}{3} - 2^{3i-3}$ and therefore the requirement is that $3 \cdot 2^{3i} < \frac{k}{3} - 2^{3i-3}$ giving that $k > (9 + \frac{3}{8})2^{3i}$. For each phase j, $1 \leq j \leq i$, First-Fit will pack one pair consisting of an item of size B_j and an item of size C_j in each bin, using $3 \cdot 2^j$ bins. After such a pair is packed, all future items are too large to join a pair. The number of bins used in the first $i+1$ phases is $1 + \sum_{j=1}^{i} 3 \cdot 2^j = 6 \cdot 2^i - 5$. In the next phase, each item will be placed in its own bin, using the last $3 \cdot 2^i$ bins. There will be no space for items from the last phase. Adding up the number of items from the first $i+2$ phases, First-Fit will accept $15 \cdot 2^i - 9$ items.

OPT can pack each item from Phase 1 with two of the items of size B_1 from Phase 2, using a total of 3 bins for this. Then, for $1 \leq j \leq i-1$, it can place every pair of items of size B_{j+1} together with one item of size C_j. This occupies $3 \cdot 2^j$ bins. Then it can pack one item of size C_i together with one item of size D using a total of $3 \cdot 2^i$ bins for this. The number of bins used is $3 + 3\sum_{j=1}^{i-1} 2^j + 3 \cdot 2^i = 6 \cdot 2^i - 3$. From the last phase, three items can be packed together in each bin, and the entire sequence will in this way occupy αn bins.

The ratio is thus

$$\frac{15\cdot 2^i-9}{15\cdot 2^i-9+3\cdot(\alpha n-(6\cdot 2^i-3))}=\frac{5\cdot 2^i-3}{\alpha n-2^i}=\frac{5-\frac{2}{n}}{9\alpha-1-\frac{5}{n}}=\frac{5}{9\alpha-1}+O\left(\frac{1}{n}\right).$$

For n general, we choose integers $n'\in\Theta(\sqrt{n}), c\leq n'$, and $s\in\Theta(\sqrt{n})$ satisfying $n=sn'+c$ and $n'=9\cdot 2^i-5$. By giving c items of size k followed by s times the sequence described above (with n' substituted for n), we get the following performance ratio

$$\frac{c+s(15\cdot 2^i-9)}{c+s(3\alpha n'-3\cdot 2^i)}\leq\frac{c+s(5\cdot 2^i-3)}{s(\alpha n'-2^i)}=\frac{5}{9\alpha-1}+O\left(\frac{c}{n}+\frac{1}{n'}\right)=\frac{5}{9\alpha-1}+O\left(\frac{1}{\sqrt{n}}\right).$$

□

The following theorem establishes that Theorem 1, showing that any fair algorithm has accommodating function at least $\frac{1}{1+\alpha-1/k}$, is asymptotically tight.

Theorem 5 For Unrestricted Bin Packing, the accommodating function for Worst-Fit is at most $\mathcal{A}_{\mathrm{WF}}(\alpha)<\frac{1}{1+\alpha-\frac{1}{k}-\frac{1}{n}}$, if $\alpha\leq 1$.

Proof The adversary gives two phases of items

- n items of size 1
- $\alpha n-\lceil\frac{n}{k}\rceil$ items of size k.

Worst-Fit will accept only the first phase, whereas the optimal algorithm can behave like First-Fit and pack all items in αn bins. The ratio is

$$\mathcal{A}_{\mathrm{WF}}(\alpha)\leq\frac{n}{n+(\alpha n-\lceil\frac{n}{k}\rceil)}<\frac{n}{n+(\alpha n-\frac{n+k}{k})}=\frac{1}{1+\alpha-\frac{1}{k}-\frac{1}{n}}.$$

□

4 Unit Price Seat Reservation

The other maximization problem we consider in detail is the Unit Price Seat Reservation Problem. A train with n seats travels from a start station to an end station, stopping at $k\geq 2$ stations, including the first and last. Reservations can be made for any trip from a station s to a station t. The passenger is given a single seat number when the ticket is purchased, which can be any time before departure. For political reasons, the problem must be solved in a *fair* manner, i.e., the ticket agent may not refuse a passenger if it is possible to accommodate him when he attempts to make his reservation. For this problem an α-sequence can be packed by an optimal off-line algorithm using αn seats. The algorithms (ticket agents) attempt to maximize income, i.e., the sum of the prices of the tickets sold. In this paper, we consider only the pricing policy in which all tickets have the same price, the *unit price problem*, since the results in [5] for the proportional price problem show that its competitive ratio, even on accommodating sequences, is $\Theta(\frac{1}{k})$.

The seat reservation problem is similar to the problem of coloring an interval graph on-line (the seats corresponding to the colors), which has been well studied because of applications to dynamic storage allocation. The difference is that with graph coloring, all vertices must be given a color and the goal is to minimize the number of colors. With the seat reservation problem, there is a fixed number n of colors, and the goal is to maximize the number of vertices that are colored. We use the following result from interval graph theory: Interval graphs are *perfect* [10], so the size of the largest clique is exactly the number of colors needed. Thus, when there is no pair of stations $(s, s+1)$ such that the number of people who want to be on the train between stations s and $s+1$ is greater than n, the optimal off-line algorithm will be able to accommodate all requests. The contrapositive is clearly also true; if there is a pair of stations such that the number of people who want to be on the train between those stations is greater than n, the optimal off-line algorithm will be unable to accommodate all requests. We will refer to the number of people who want to be on the train between two stations as the *density* between those stations. Some results for First-Fit and Worst-Fit found in this section are depicted in Figure 1.

4.1 Performance Guarantees

The performance guarantee for fair algorithms on accommodating sequences found in [5] can be extended to a general performance guarantee for $\alpha \leq 1$. Consider any α-sequence and any fair on-line algorithm $\mathbb{A}$. Let A be the set of requests which are accepted and R be the set rejected. For any subset S of the n seats, let A_S denote the subset of the intervals in A which $\mathbb{A}$ places on seats in S. The following lemma shows that if many requests are rejected, then the number of intervals in A_S must be large. The lemma is used to prove a performance guarantee on all fair algorithms. The main idea of the proof is to focus on some specific subset, consisting of αn of the n seats.

Lemma 2 Let S be any subset of the seats, consisting of αn seats, $\alpha \leq 1$. For any fair algorithm $\mathbb{A}$, $|A_S| \geq |R|$.

Proof Since $\mathbb{A}$ is fair, none of the requests in R can be placed on any of the seats in S. Consider the following request sequence I', consisting of all the intervals in $A_S \cup R$, with the intervals in A_S coming before those in R. Clearly, there exists a fair algorithm $\mathbb{A}'$ which given I' and a train with αn seats places the intervals in A_S exactly as $\mathbb{A}$ did, except that it uses all αn seats instead of αn selected seats out n. Algorithm $\mathbb{A}'$ accepts all the intervals in A_S, and, as a consequence, rejects all intervals from R.

The request sequence I' is an α-sequence, since it is constructed from a subset of requests from an α-sequence, so every request could be accommodated in those αn seats by an optimal off-line algorithm. By [5], the competitive ratio of any fair algorithm on accommodating sequences is at least $\frac{1}{2}$, so $|A_S| \geq |R|$. $\square$

Theorem 6 Any fair algorithm for Unit Price Seat Reservation is $\frac{1}{1+\alpha}$-competitive on α-sequences, for $\alpha \leq 1$.

Proof Choose S to be a subset containing αn seats which minimizes $|A_S|$. By Lemma 2, $|A_S| \geq |R|$. Since S was chosen to minimize $|A_S|$, $\mathbb{A}$ places at least as many intervals on each seat not in S as it does on any seat in S. Thus, $|A| \geq n\frac{|A_S|}{\alpha n} \geq \frac{|R|}{\alpha}$. This gives a performance ratio of $\frac{|A|}{|A|+|R|} \geq \frac{\frac{1}{\alpha}|R|}{\frac{1}{\alpha}|R|+|R|} = \frac{1}{1+\alpha}$. □

This result is asymptotically tight due to Worst-Fit's behavior, as shown in the next subsection.

It was shown in [5] that First-Fit's competitive ratio on accommodating sequences is not strictly better than $\frac{1}{2}$, the performance guarantee for any fair algorithm. It is interesting that for $\alpha \leq \frac{1}{2}$, it is possible to prove that First-Fit's performance is better than Worst-Fit's.

Theorem 7 $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq \frac{2^{l-1}(1-l\alpha)+(2^l-1)\alpha}{2^{l-1}(1-l\alpha)+(2^l)\alpha} \geq 1 - \frac{1}{2^l}$, where $l = \lfloor \frac{1}{\alpha} \rfloor$ and $\alpha \leq 1$.

Proof Consider an α-sequence. Let A be the set of requests which are accepted by First-Fit and R be the set rejected.

Divide the seats up into l blocks, where $l = \lfloor \frac{1}{\alpha} \rfloor$, such that the first αn seats are in block 1, the next αn seats in block 2, etc. Any leftover seats will be added to the lth block, so that it has $n - l\alpha n + \alpha n$ seats. Let A_i denote the set of intervals which First-Fit places in the ith block.

We consider the last block first and compute a lower bound on $|A_l|$, the number of intervals First-Fit places in it. Using the techniques from the proof of Theorem 6, one can choose the αn seats from among the last $n - l\alpha n + \alpha n$ which contain the fewest intervals. Those seats must contain at least $|R|$ intervals, so $|A_l| \geq \frac{1-l\alpha+\alpha}{\alpha}|R|$.

First-Fit could not place any of the intervals it placed in the last block in any previous block. Thus, one can look at the intervals in the second to last block, in the last block, and in R, and note that they must form an α-sequence. Again using the result that any fair algorithm is at least $\frac{1}{2}$-competitive on accommodating sequences, we have that $|A_{l-1}| \geq |R| + |A_l|$. Similarly, $|A_i| \geq |R| + \sum_{j=i+1}^{l} |A_j|$. If C_i denotes the total number of intervals First-Fit places in the last i blocks, then $C_i \geq 2C_{i-1} + |R|$, so $C_l \geq \frac{2^{l-1}(1-l\alpha)+(2^l-1)\alpha}{\alpha}|R|$. Thus, the performance ratio

$$\mathcal{A}_{\mathrm{FF}}(\alpha) \geq \frac{|A|}{|A|+|R|} = \frac{C_l}{C_l+|R|} \geq \frac{2^{l-1}(1-l\alpha)+(2^l-1)\alpha}{2^{l-1}(1-l\alpha)+(2^l)\alpha}.$$

□

This performance guarantee is not tight. Kierstead and Qin [12] have shown that First-Fit colors every interval graph using at most 25.72 times as many colors as necessary, so the competitive ratio on α-sequences is 1 when $\alpha \leq \frac{1}{25.72} \approx 0.039$.

4.2 Impossibility Results

Our first impossibility result for the Unit Price Seat Reservation Problem shows that the general lower bound of Theorem 6 is asymptotically tight by giving a matching upper bound on the performance of Worst-Fit.

Theorem 8 For $\alpha \leq 1$, $\mathcal{A}_{\mathrm{WF}}(\alpha) < \frac{1}{1+\alpha-\frac{1}{k-1}-\frac{1}{n}}$.

Proof Since this upper bound is at least 1 for $\alpha n < \lceil \frac{n}{k-1} \rceil$, we assume that $\alpha n \geq \lceil \frac{n}{k-1} \rceil$. Consider the sequence consisting of n intervals each of size 1 "evenly" distributed among the intervals of the form $[i, i+1)$, for $1 \leq i \leq k-1$, i.e., the number of intervals for each i is either $\lfloor \frac{n}{k-1} \rfloor$ or $\lceil \frac{n}{k-1} \rceil$, followed by $\alpha n - \lceil \frac{n}{k-1} \rceil$ copies of the interval $[1, k)$. The unit-length intervals can all be placed on $\lceil \frac{n}{k-1} \rceil$ seats, so the sequence is an α-sequence. Worst-Fit will put the unit-length intervals on separate seats and reject all of the long intervals. OPT will accept all of the intervals since $\alpha \leq 1$, so the ratio is

$$\mathcal{A}_{\mathrm{WF}}(\alpha) \leq \frac{n}{n+\alpha n - \lceil \frac{n}{k-1} \rceil} < \frac{n}{n + \alpha n - \frac{n+k-1}{k-1}} = \frac{1}{1+\alpha - \frac{1}{k-1} - \frac{1}{n}}.$$

□

Theorem 9 below says that if k can be arbitrarily large, no fair on-line deterministic or randomized algorithm for the Unit Price Seat Reservation Problem is better than $\frac{8-\alpha}{9\alpha}$-competitive on α-sequences. Thus, for $\alpha > \frac{4}{5}$, $\mathcal{A}_{\mathbb{A}}(\alpha) < 1$ for any algorithm $\mathbb{A}$ for the problem. The proof of this is based on Theorem 3.1 in [3], which in turn is based on a proof from [5].

Theorem 9 For $k \geq 9$, $\alpha \leq 1$, and αn even, no fair on-line algorithm for the Unit Price Seat Reservation Problem is more than $\frac{(8-\alpha)2k'+33\alpha-48}{9\alpha \cdot 2k'-27}$-competitive, even on α-sequences, where $k' \leq k$ is largest possible satisfying that 6 divides $k'-3$.

Proof Assume that αn is even, $\alpha \leq 1$ and $k \geq 9$. Let $k' \leq k$ be chosen largest possible such that 6 divides $k'-3$. The adversary begins with $\frac{\alpha n}{2}$ requests for each of the intervals $[3s+1, 3s+3)$ for $s = 0, 1, \cdots, \frac{k'-3}{3}$. We refer to these intervals as the *original* intervals. Any fair on-line algorithm is able to accept this set of $\frac{k'}{3} \cdot \frac{\alpha n}{2}$ requests. Suppose that after these requests are accepted, there are q_i seats which are empty from station $3i+2$ to station $3i+5$. Let $p_i = q_{2i} + q_{2i+1}$.

If $E[p_i] \leq \frac{16-11\alpha}{9} n$, the adversary proceeds with $\frac{\alpha n}{2}$ requests for the interval $[6i+2, 6i+5)$ and $\frac{\alpha n}{2}$ requests for the interval $[6i+5, 6i+8)$. For these αn additional requests, the on-line algorithm can accommodate exactly p_i of them. Thus, for the total of $2\alpha n$ requests, the starting station of which is in $[6i+1, 6i+6)$, the on-line algorithm accommodates $\alpha n + p_i$.

If $E[p_i] > \frac{16-11\alpha}{9} n$, the adversary proceeds with $\frac{\alpha n}{2}$ requests for the interval $[6i+3, 6i+7)$, followed by $\frac{\alpha n}{2}$ requests for the interval $[6i+2, 6i+4)$ and $\frac{\alpha n}{2}$ requests for the interval $[6i+6, 6i+8)$. The $[6i+3, 6i+7)$ requests are all accepted. Out of the next αn requests, the on-line algorithm can accommodate exactly $(4-3\alpha)n - p_i$ of them. This is seen as follows. Consider the $[6i+2, 6i+4)$ intervals. The on-line algorithm can place $n - q_{2i} - \frac{\alpha n}{2}$ of these intervals on the same seat as a $[6i+4, 6i+6)$ interval, namely where there is no $[6i+1, 6i+3)$ interval at the same time. On the $n - \frac{\alpha n}{2}$ seats where there are no $[6i+4, 6i+6)$ intervals, one of the $\frac{\alpha n}{2}$ intervals of type $[6i+3, 6i+7)$ would prevent us from placing a $[6i+2, 6i+4)$ interval. Thus, the on-line algorithm can place at most $n - \frac{\alpha n}{2} - \frac{\alpha n}{2}$ intervals of type $[6i+2, 6i+4)$ on seats with no $[6i+4, 6i+6)$ interval. Adding up, the total number of accepted requests are at most $(2 - \frac{3\alpha}{2})n - q_{2i}$. Similarly, the on-line algorithm accepts at most $(2 - \frac{3\alpha}{2})n - q_{2i+1}$ of the $[6i+6, 6i+8)$ intervals, totaling $(4-3\alpha)n - p_i$. Thus, of the $\frac{5\alpha n}{2}$ requests, the starting station of which is in $[6i+1, 6i+6)$, the on-line algorithm accommodates at most $\frac{3\alpha n}{2} + (4-3\alpha)n - p_i = (4 - \frac{3}{2}\alpha)n - p_i$ of them.

Clearly for both cases, the density is nowhere more than αn, so the optimal off-line algorithm can accommodate all requests on αn seats.

In this way, the requests are partitioned into $\frac{k'-3}{6} + 1$ groups; each of the first $\frac{k'-3}{6}$ groups consists of either $2\alpha n$ or $\frac{5\alpha n}{2}$ requests and the last group consists of $\frac{\alpha n}{2}$ requests. For each of the first $\frac{k'-3}{6}$ groups, the on-line algorithm can accommodate up to a fraction $\frac{8-\alpha}{9\alpha}$ of the requests in the group. This leads to the theorem.

More precisely, let S denote the set of indices for which the first case happens, and let $\bar{S}$ denote the set of indices for which the second case happens. By linearity of expectation, the expected ratio of the number of requests accepted by the on-line algorithm to the number of requests accepted by an optimal off-line algorithm is bounded by

$$
\begin{aligned}
& \frac{\frac{\alpha n}{2} + \sum_{i \in S}(\alpha n + E[p_i]) + \sum_{i \in \bar{S}}((4 - \frac{3}{2}\alpha)n - E[p_i])}{\frac{\alpha n}{2} + \sum_{i \in S} 2\alpha n + \sum_{i \in \bar{S}} \frac{5\alpha n}{2}} \\
\le\ & \frac{\frac{\alpha}{2} + \sum_{i \in S}(\alpha + \frac{16-11\alpha}{9}) + \sum_{i \in \bar{S}}(4 - \frac{3}{2}\alpha - \frac{16-11\alpha}{9})}{\frac{\alpha}{2} + \sum_{i \in S} 2\alpha + \sum_{i \in \bar{S}} \frac{5\alpha}{2}} \\
\le\ & \frac{\frac{\alpha}{2} + (\alpha + \frac{16-11\alpha}{9}) \cdot \frac{k'-3}{6}}{\frac{\alpha}{2} + 2\alpha \cdot \frac{k'-3}{6}} \\
=\ & \frac{(8-\alpha)2k' + 33\alpha - 48}{9\alpha \cdot 2k' - 27\alpha},
\end{aligned}
$$

where the last inequality holds because in general $\frac{a}{b} = \frac{c}{d} < 1$ and $a < c$ imply that $\frac{e+ax+cy}{e+bx+dy} \le \frac{e+a(x+y)}{e+b(x+y)}$. This completes the proof. $\square$

We consider a class of fair algorithms called Any-Fit which will use an empty seat, only if there is not enough space on partially used seats. Any-Fit includes both First-Fit and Best-Fit. The following theorem shows that, if k can be arbitrarily large, the competitive ratio of an Any-Fit algorithm on α-sequences is no better than $\frac{1}{3\alpha-1}$. Thus, for $\alpha > \frac{2}{3}$, the competitive ratio of an Any-Fit algorithm on α-sequences is strictly worse than 1.

Theorem 10 Let $\mathbb{A}$ be an Any-Fit algorithm and let $\frac{1}{3} < \alpha \le 1$. If 3 divides n and 6 divides $k-4$, then $\mathcal{A}_{\mathbb{A}}(\alpha) \le \frac{k+4}{(3\alpha-1)k+8-6\alpha}$. Otherwise, $\mathcal{A}_{\mathbb{A}}(\alpha) \le \frac{k+4}{(3\alpha-1)k+8-6\alpha} + O(\frac{1}{n})$.

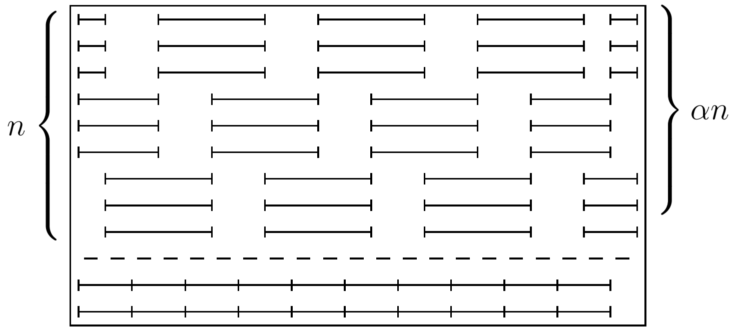

Figure 4: Placement for $k = 22$ and $n = 9$.

Proof Since the upper bound is at least 1 for $\alpha \leq \frac{2}{3}$, we assume that $\alpha \geq \frac{2}{3}$. Assume first that 3 divides n and 6 divides $k-4$. Give $\frac{n}{3}$ requests of each of the following types (see Figure 4, where we have illustrated the placement by $\mathbb{A}$):

- $[1,2)$, $[k-1,k)$, and for $i \in \{1,\ldots,\frac{k-4}{6}\}$: $[6i-2, 6i+2)$.
- $[1,4)$, $[k-4,k-1)$, and for $i \in \{1,\ldots,\frac{k-4}{6}-1\}$: $[6i, 6i+4)$.
- $[k-2,k)$, and for $i \in \{1,\ldots,\frac{k-4}{6}\}$: $[6i-4, 6i)$.

Since $\mathbb{A}$ is fair, it accepts all these requests, amounting to $\frac{n}{3}(3\frac{k-4}{6}+4) = \frac{k+4}{6}n$.

OPT uses $\frac{2n}{3}$ seats to accommodate all these requests, since the density between any two stations is $\frac{2n}{3}$. OPT can use the space from $\frac{2n}{3}$ up to αn to accept the following intervals, which $\mathbb{A}$ will be unable to accommodate:

- For $i \in \{1,\ldots,\frac{k}{2}-1\}$: $[2i-1, 2i+1)$.

We get the following ratio:

$$\mathcal{A}_{\mathbb{A}}(\alpha) \leq \frac{\frac{k+4}{6}n}{\frac{k+4}{6}n + (\alpha-\frac{2}{3})n(\frac{k}{2}-1)} = \frac{k+4}{(3\alpha-1)k+8-6\alpha}.$$

Let $c = n \bmod 3$ and $d = (k-4) \bmod 6$. If $c = 0$ and $d > 0$, the sequence above with $\frac{k-4}{6}$ replaced by $\lfloor\frac{k-4}{6}\rfloor$ can be used. If $c > 0$, the adversary first gives $c \leq 2$ intervals $[1,k)$ and then gives the sequence above with $\frac{n}{3}$ replaced by $\lfloor\frac{n}{3}\rfloor$ and $\frac{k-4}{6}$ replaced by $\lfloor\frac{k-4}{6}\rfloor$. The ratio is then

$$
\begin{aligned}
\mathcal{A}_{\mathbb{A}}(\alpha) &\leq \frac{c + \frac{k-d+4}{6}(n-c)}{c + \frac{k-d+4}{6}(n-c) + (\alpha - \frac{2}{3})(n-c)(\frac{k-d}{2} - 1)} \\
&= \frac{k + 4 - \frac{k-2}{n}c - (1 - \frac{c}{n})d}{(3\alpha - 1)k + 8 - 6\alpha - \frac{(3\alpha-1)k+2-6\alpha}{n}c - (1 - \frac{c}{n})(3\alpha - 1)d} \\
&\leq \frac{k+4}{(3\alpha - 1)k + 8 - 6\alpha - \frac{(6\alpha-2)k+4-12\alpha}{n} - 15\alpha + 5} \\
&= \frac{(k+4)\left(1 + \frac{\left((6\alpha-2)k+4-12\alpha\right)/n+15\alpha-5}{(3\alpha-1)k+8-6\alpha-\left((6\alpha-2)k+4-12\alpha\right)/n-15\alpha+5}\right)}{(3\alpha-1)k+8-6\alpha} \\
&= \frac{k+4}{(3\alpha-1)k+8-6\alpha} + O\left(\frac{1}{n}\right).
\end{aligned}
$$

□

4.3 Distinguishing Between Algorithms

As mentioned above, First-Fit has a competitive ratio on accommodating sequences which is not strictly better than $\frac{1}{2}$. This result is from [5]. Later, in [3], it was shown that $\frac{1}{2}$ is an asymptotic (for k much larger than n) upper bound for any fair algorithm. Similarly, its competitive ratio is $\frac{2-\frac{1}{k-1}}{k-1}$, while the performance guarantee for any fair algorithm is $\frac{2}{k}$ [5]. Theorems 7 and 8 together show that in contrast to these results for $\alpha \geq 1$, when $\alpha < \frac{1}{2}$, First-Fit is not as bad as any other possible fair algorithm. Specifically, Worst-Fit has an asymptotic competitive ratio of at most $\frac{3}{4}$ on $\frac{1}{3}$-sequences, while First-Fit has a ratio of at least $\frac{7}{8}$.

Kierstead and Trotter [13] have shown that every interval graph with clique number $\omega(G)$ can be $(3\omega(G)-2)$-colored on-line. Since colors correspond to seat assignments, this means that their algorithm has a competitive ratio of 1 on $\frac{1}{3}$-sequences. Their algorithm solves the problem of minimizing the number of seats (colors) used; it is undefined for values of $\alpha > \frac{1}{3}$ when applied to the maximization problem where the number of seats is limited. However, their algorithm can obviously be extended in many ways so that it solves this maximization problem. For example, one could use their algorithm whenever it assigns a seat number which exists and use First-Fit otherwise.

Chrobak and Ślusarek [8] have shown that there exists a sequence where First-Fit uses more than 4.4 times as many seats (colors) as OPT, and thus is not 1-competitive on $\frac{1}{3}$-sequences. Thus, Kierstead and Trotter's algorithm has a better competitive ratio on $\frac{1}{3}$-sequences than First-Fit.

Kierstead and Trotter's algorithm is not, however, better than First-Fit on 1-sequences, since it has been shown [5] that no fair algorithm has a competitive ratio on accommodating sequences which is strictly better than $\frac{1}{2}$ asymptotically.

In conclusion, we have that the competitive ratio on α-sequences, where $\alpha < 1$, can be useful in distinguishing between algorithms.

Theorem 11 Kierstead and Trotter's algorithm, First-Fit, and Worst-Fit, all have asymptotic competitive ratio $\frac{1}{2}$ on accommodating sequences, but have different competitive ratios on $\frac{1}{3}$-sequences.

Note 1 It is tempting to try to use the techniques in Chrobak and Ślusarek [8] to improve the hardness result on First-Fit or the techniques in Kierstead and Trotter [13] to improve the general hardness result. Unfortunately, the number of requests accepted is so large compared to the number rejected that the hardness results which can be obtained are essentially 1. □

5 Resource augmentation in comparison with the accommodating function

The concept of the accommodating function should not be confused with resource augmentation introduced in [11]. Resource augmentation analysis gives the on-line algorithm more resources than the optimal off-line algorithm that it is compared to, but there is no restriction on the input sequences.

Note 2 In the resource augmentation setting, the on-line algorithm has the amount m of resources and the optimal off-line algorithm has the amount $n \leq m$ of resources. Positive results (upper bounds for minimization problems and lower bounds for maximization problems) obtained in this setting are also valid for the case where all input sequences are $\frac{n}{m}$-sequences. This is because, even though the optimal off-line algorithm in our setting has the same amount of resources as the on-line algorithm, the result it obtains cannot be improved with extra resources beyond n when considering $\frac{n}{m}$-sequences. □

The contrapositive of the observation above gives that when considering negative results, it is the other way around: negative results for the accommodating function carry over to the resource augmentation setting.

Resource augmentation does not always give as realistic or optimistic results as the accommodating function for $\alpha < 1$. In order to obtain the same results, one needs to restrict to accommodating sequences while also doing resource augmentation. To see that resource augmentation alone can be insufficient, consider the algorithm First-Fit for the Seat Reservation Problem. Theorem 7 gives a lower bound on First-Fit's accommodating function for $\alpha \leq 1$: $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq 1 - \frac{1}{2^l}$, where $l = \lfloor \frac{1}{\alpha} \rfloor$. In contrast, resource augmentation cannot give a constant lower bound on First-Fit's accommodating function.

Theorem 12 For $\alpha > \frac{2}{n}$, the competitive ratio of First-Fit for the Seat Reservation Problem is at most $\frac{1+\alpha}{(\alpha - \frac{2}{n})(k-1)}$, when First-Fit has at most $\frac{1}{\alpha}$ times as many seats as OPT.

Proof Let $\frac{2}{n} \leq \alpha < 1$ be given. We allow First-Fit to have n seats, while OPT has $\lfloor \alpha n \rfloor$ seats. Consider the sequence beginning with $\lfloor \frac{\alpha n}{k-1} \rfloor$ requests for $[i, i+1)$ for $1 \leq i \leq k-1$, one request for $[i, i+1)$ for $1 \leq i \leq \lfloor \alpha n \rfloor - (k-1)\lfloor \frac{\alpha n}{k-1} \rfloor$, $n - 1 - \lfloor \frac{\alpha n}{k-1} \rfloor$ requests for $[1, k)$, and finally the request $[\lfloor \alpha n \rfloor + 1 - (k-1)\lfloor \frac{\alpha n}{k-1} \rfloor, k)$. First-Fit will accept all of these and the train will be full. The sequence continues with $\lfloor \alpha n \rfloor - 1 - \lfloor \frac{n}{k-1} \rfloor$ requests for $[i, i+1)$ for $1 \leq i \leq k-1$, and First-Fit must reject them all. Note that OPT can place one of the first short intervals on each of its $\lfloor \alpha n \rfloor$ seats and thus reject all of the long intervals, except possibly the last "long" interval. Thus, it can accept all of the unit-length intervals at the end. The performance ratio is at most $\frac{\lfloor \alpha n \rfloor + n - \lfloor \frac{\alpha n}{k-1} \rfloor}{\lfloor \alpha n - 1 \rfloor (k-1)} \leq \frac{1+\alpha}{(\alpha - \frac{2}{n})(k-1)}$. □

For Unrestricted Bin Packing, one can also show that resource augmentation analysis alone gives results which are much more pessimistic than the corresponding accommodating function results. Recall that the performance guarantee from Corollary 1 for First-Fit for $\alpha \leq 1$ is $\mathcal{A}_{\mathrm{FF}}(\alpha) \geq \frac{3+2\alpha}{8\alpha}$. We let n denote the resources available for the on-line algorithm and let n' denote the resources available for the off-line algorithm.

Theorem 13 For Unrestricted Bin Packing, the competitive ratio of First-Fit is at most $\frac{1}{\alpha k}$, when First-Fit has at most $\frac{1}{\alpha}$ times as many bins as OPT.

Proof The adversary can give n items of size k followed by $\alpha n k$ items of size 1. First-Fit will accept the first n items, but none of the small items. OPT will reject the first n items and accept all the small items, giving a ratio of $\frac{n}{\alpha n k} = \frac{1}{\alpha k}$. □

Next, we consider three concrete minimization problems to demonstrate how the accommodating function is applicable to them. These problems show the subtleties in the definition of the accommodating sequences. They also demonstrate how resource augmentation results can be applied to give accommodating function results.

5.1 Paging

We consider the Paging Problem in the page fault model, i.e., the algorithm maintains a fast memory (cache) consisting of k pages of memory and the input is a sequence of page requests. If a page in cache is requested, no cost is incurred; otherwise the requested page must be transferred from the slow memory at a cost of 1 and another page must be evicted from the cache. The goal is to choose an eviction strategy which minimizes cost. Before the first page request is served, the cache is empty. In this problem, the limited resource is the cache.

When an optimal off-line algorithm serves an α-sequence, it will have the same cost for every cache size $k' \geq \alpha k$. To see why such a sequence can consist of more than αk different pages, consider the following example: For $k = 6$, the following sequence is a $\frac{1}{2}$-sequence, $S = \langle 1, 2, 3, 4, 5, 6, 7, 1, 2, 8, 9, 10, 1, 2, 11, 12 \rangle$. Keeping the two pages 1 and 2 in the cache will give a cost of 12 with a cache of size k' for all $k' \geq 3$.

In the (h, k)-Paging Problem, the on-line algorithm is assumed to have a cache of size k as in the usual Paging Problem, but the performance of the on-line algorithm is measured relative to an optimal off-line algorithm whose cache has size $h \leq k$. (See [4] for a survey on (h, k)-Paging.)

Using Note 2 and a result from [16], we find that FIFO (First-In/First-Out) and LRU (Least-Recently-Used) are $\frac{k}{(1-\alpha)k+1}$-competitive. Sleator and Tarjan [16] have shown that their performance guarantee is tight, and a very similar proof shows that it is also tight for α-sequences. The proof is given in detail here to give further intuition for α-sequences.

Theorem 14 Any deterministic Paging algorithm $\mathbb{A}$ has $\mathcal{A}_{\mathbb{A}}(\alpha) \geq \frac{k}{(1-\alpha)k+1}$, when $\alpha \leq 1$.

Proof Let $\mathbb{A}$ be a deterministic Paging algorithm. Assume that OPT has a cache of size αk. The requests are given in phases $P_0, P_1, \ldots, P_n$. In phase P_0, αk different pages are requested once. The rest of the phases consist of k requests each. For each phase P_i, $1 \leq i \leq n$, the first $k - \alpha k + 1$ requests are to pages that have not been requested before. Let S_i be the set consisting of these $k - \alpha k + 1$ pages and the αk pages that were in OPT's cache at the beginning of P_i. Then, $|S_i| = k + 1$. Each of the remaining $\alpha k - 1$ requests in phase P_i is to a page in S_i that is currently not in $\mathbb{A}$'s cache.

The first $k - \alpha k + 1$ pages of phase P_i , $1 \leq i \leq n$, cost both algorithms $k - \alpha k + 1$, since they have not been requested earlier. During these $k - \alpha k + 1$ requests, OPT keeps the pages to be requested during the last $\alpha k - 1$ requests of the phase in its cache. Thus, the optimal cost of each phase P_i, $1 \leq i \leq n$, is $(1 - \alpha)k + 1$, while the on-line cost is k. Since OPT has faults only on requests to pages that have not been requested earlier, the sequence is an α-sequence. □

Using Note 2 and a result from [17], we find that when $\alpha < 1$, the randomized algorithm MARK is 2-competitive when $\frac{1}{1-\alpha} < e$ and $2(\ln \frac{1}{1-\alpha} - \ln\ln \frac{1}{1-\alpha} + \frac{1}{e-1})$-competitive when $\frac{1}{1-\alpha} \geq e$. Young's hardness result [17] for any randomized algorithm can be used directly for the competitive ratio on α-sequences for $\alpha < 1$, since his sequences are α-sequences. Thus, any randomized algorithm has a competitive ratio of at least $\ln \frac{1}{1-\alpha} - \ln\ln \frac{1}{1-\alpha} - \frac{3}{k(1-\alpha)}$ when $\frac{1}{1-\alpha} \geq e$, so MARK is within a factor of approximately 2 of optimal.

5.2 The k-Server Problem

The k-Server Problem [15] is a natural generalization of the Paging Problem which was discussed in the previous subsection. For the k-Server Problem [15], the algorithm controls k servers, which are located on points of a metric space (M, d), where M is a set of points with $|M| > k$ and d is a metric for M. The request sequence is a sequence of points in M, and the algorithm must service all requests sequentially, by having a server at the point requested or moving one there. The goal is to minimize the total distance the servers move. When considering the accommodating function for the k-Server Problem, we use the servers as the resource; this is the natural generalization from the last section where the pages in cache were the resource.

5.2.1 A Lower Bound

Since the k-Server Problem is a generalization of the Paging Problem, the lower bound of $\frac{k}{(1-\alpha)k+1}$ found in Section 5.1 is also valid for the k-Server Problem.

5.2.2 An Upper Bound for WFA

In [14], the Work Function Algorithm for the k-Server Problem is investigated in the case where the on-line algorithm has k servers and the off-line algorithm has $h \leq k$ servers. Using Note 2 and a result in [14], the upper bounds found in [14] are also valid for the k-Server Problem on α-sequences. One of these upper bounds is

$$\mathrm{WFA}_k(I) \leq k \cdot OPT_{\alpha k}(I) + (\alpha k - 1) \cdot OPT_k(I) + \mathrm{const},$$

which reduces to $\mathrm{WFA}_k(I) \leq (k + \alpha k - 1) \cdot OPT_k(I) + \mathrm{const}$, in the case of α-sequences. Thus, we arrive at the following theorem.

Theorem 15 For $\alpha \leq 1$, $\mathcal{A}_{\mathrm{WFA}}(\alpha) \leq (1+\alpha)k - 1$.

5.2.3 Balance

The algorithm Balance serves each request r with a server s_i such that $D_i + d_i(r)$ is minimized, where D_i is the total distance traveled by the server s_i before the request r and $d_i(r)$ is the distance it will have to travel to serve r. In general, Balance is not competitive. This is normally shown using only two servers and a metric space with only 4 points forming a rectangle. For instance in $\mathbb{R}^2$, using the Euclidean distance, we could have that $a_1 = (M, \epsilon), a_2 = (M, -\epsilon), b_1 = (2M, \epsilon)$, and $b_2 = (2M, -\epsilon)$. Let all servers start at $(0,0)$ and consider the request sequence

$$S = \langle a_1, a_2, b_1, b_2, a_1, a_2, b_1, b_2, \ldots, a_1, a_2, b_1, b_2 \rangle.$$

Balance will let one server travel back and forth between a_1 and b_1 and the other server will travel between a_2 and b_2. OPT on the other hand will let one server service a_1 and a_2, while the other will take care of b_1 and b_2. Thus, the ratio of the cost of Balance to the optimal cost is asymptotically $\frac{d(a_1,b_1)}{d(a_1,a_2)} \geq \frac{M}{2\epsilon}$, which can be arbitrarily large. Assume that the subsequence a_1, a_2, b_1, b_2 is repeated n times. Then, if $M > n \cdot \epsilon$, then $OPT_2(S) = OPT_{k'}(S)$, for any $k' \geq 2$. Thus, by choosing an appropriate rectangle, it can be shown that the sequence is a $\frac{2}{k}$-sequence. This example can easily be extended to larger values of k, using $a_1, a_2, \ldots, a_k$ and $b_1, b_2, \ldots, b_k$.

For $\alpha = \frac{1}{k}$, consider the following request sequence in $\mathbb{R}^2$: For $i = 1, 2, \ldots, n$ give $(i, 0), (i, \epsilon), \ldots, (i, (k-1)\epsilon)$. The cost of Balance is larger than kn and the cost of OPT is less than $n + 2(k-1)n\epsilon$. Thus, choosing ϵ sufficiently small, the ratio of the cost of Balance to the optimal cost can be arbitrarily close to k. On the other hand, because of the triangle inequality, no Balance server moves a longer distance than OPT's server.

This proves the following theorem.

Theorem 16 For $\alpha \geq \frac{2}{k}$, Balance is not competitive, even on α-sequences. The competitive ratio of Balance on $\frac{1}{k}$-sequences is k.

5.2.4 Greedy

The algorithm Greedy serves each request r with a server s_i, such that the distance moved, $d_i(r)$, is minimized. Like Balance, Greedy is not competitive on α-sequences, as long as $\alpha \geq \frac{2}{k}$. This can be shown using only three points a, b, c. Let a be the starting point of the servers and let $d(b,c) < d(a,b)$. Greedy will service the sequence $\langle b, c, b, c, \ldots, b, c\rangle$ using only one server, whereas OPT will put one server on b and another server on c. It is clear that the sequence is a $\frac{2}{k}$-sequence. In contrast to Balance, Greedy is optimal on $\frac{1}{k}$-sequences. This can be seen in the following way. Let a be the starting point of the servers. For any two consecutive requests r_{i-1}, r_i, $d(r_{i-1}, r_i) \leq d(a, r_i)$, since otherwise the sequence is not a $\frac{1}{k}$-sequence. Thus, in each step, the cost of Greedy is the same as the cost of OPT.

Theorem 17 For $\alpha \geq \frac{2}{k}$, Greedy is not competitive, even on α-sequences. On $\frac{1}{k}$-sequences, Greedy is 1-competitive.

5.3 Machine Scheduling Minimizing Makespan

We study the problem of assigning jobs to machines such that the makespan is minimized, i.e., the completion time of the latest completing job is minimized. The jobs are characterized by their length, and there are m identical machines.

5.3.1 List Scheduling

The List Scheduling algorithm (LS) [9] schedules each new job on a least loaded machine.

Theorem 18 For $\alpha \leq 1$, $\mathcal{A}_{\mathrm{LS}}(\alpha) = 1 + \alpha - \frac{1}{m}$.

Proof In [9] it is shown that the competitive ratio of List Scheduling is at most $1+\alpha-\frac{1}{m}$, if LS has m machines and the off-line algorithm has only αm machines. Thus, according to Note 2 in Section 2, $\mathcal{A}_{\mathrm{LS}}(\alpha) \leq 1 + \alpha - \frac{1}{m}$.

The sequence used in [7] to match the upper bound consists of $(\alpha m - 1)m$ jobs of length 1, followed by one job of length m. Clearly, this sequence is an α-sequence. Thus, $\mathcal{A}_{\mathrm{LS}}(\alpha) \geq 1 + \alpha - \frac{1}{m}$. □

5.3.2 A General Lower Bound

In [2], the scheduling problem is investigated in the resource augmentation setting. A general hardness result of $1+e^{-\frac{1}{\alpha}(1-o(1))}$, valid for both deterministic and randomized algorithms, is given using α-sequences. They also give a hardness result of $\frac{5}{4}$ for deterministic algorithms for $\alpha = \frac{1}{2}$. The following hardness result is a generalization of this result. In [2], the scheduling problem is investigated in the resource augmentation setting. A general lower bound of $1 + e^{-\frac{1}{\alpha}(1-o(1))}$ valid for both deterministic and randomized algorithms is given using α-sequences. The following lower bound is a better bound, but it is only valid for deterministic algorithms.

Theorem 19 Let $\alpha \le 1$ and let $l \le \alpha m$ be an integer such that $i \mid \alpha m$, for all $i \in \{1, 2, \ldots, l\}$. Then, for any deterministic algorithm $\mathbb{A}$ for the scheduling problem, if $\alpha \ge \frac{m+1}{m}\left(\sum_{i=1}^{l} \frac{1}{i}\right)^{-1}$, then $\mathcal{A}_{\mathbb{A}}(\alpha) \ge \frac{l+1}{l}$.

Proof The jobs are given in phases starting with Phase 1. In the ith phase, $\frac{\alpha m}{i}$ jobs of size i are given. The sequence is ended as soon as two jobs have been assigned the same machine. This happens at the latest in Phase l, since $\sum_{i=1}^{l} \frac{\alpha m}{i} \ge m + 1$.

Note that, for $1 \le i \le l$, the total length of jobs of length i is αm and the total length of all of the jobs is $\alpha m l$. Thus, to prove that the sequence is an α-sequence, we must prove that the jobs can be scheduled on αm machines such that the last job on each machine has completion time l. For the purpose of the proof, we now schedule the jobs (without idle time) in non-increasing order of length. Assume to the contrary that a job of length i cannot be scheduled such that its completion time is no later than l. This means that all machines are busy until time $l - i + 1$. Thus, a total volume of $(l - i + 1)\alpha m$ has been given. However, this is a contradiction, since this volume will not be reached until all the jobs of length i have been processed. Now, let $r \le l$ be the number of phases. The makespan of $\mathbb{A}$ is at least $r + 1$, and the optimal makespan is r. Thus, $\mathcal{A}_{\mathbb{A}}(\alpha) \ge \frac{r+1}{r} \ge \frac{l+1}{l}$. $\square$

Example:

$$
\begin{aligned}
\alpha &\ge \frac{2}{3}\frac{m+1}{m} &\Rightarrow \mathcal{A}(\alpha) \ge \frac{3}{2} \\
\alpha &\ge \frac{6}{11}\frac{m+1}{m} &\Rightarrow \mathcal{A}(\alpha) \ge \frac{4}{3} \\
\alpha &\ge \frac{12}{25}\frac{m+1}{m} &\Rightarrow \mathcal{A}(\alpha) \ge \frac{5}{4}
\end{aligned}
$$

For $\alpha \ge \frac{4}{5} + \frac{2}{m}$, the following theorem supplies a stronger hardness result than Theorem 19. Specifically, for $i = 2$ and $p = 1 + \sqrt{2}$, it says that $\alpha \ge \frac{8}{9} + \frac{16}{9m}$ gives $\mathcal{A}_{\mathbb{A}}(\alpha) \ge 1 + \frac{1}{\sqrt{2}} \approx 1.707$, and for $i = 3$ and $p = \frac{1}{2}(1 + \sqrt{3})$, $\alpha \ge \frac{24}{29} + \frac{52}{29m}$ gives $\mathcal{A}_{\mathbb{A}}(\alpha) \ge 1 + \frac{1}{\sqrt{3}} \approx 1.577$.

Theorem 20 If ($i = 2$ and $p \ge 1 + \sqrt{2}$) or ($i = 3$ and $p \ge \frac{1}{2}(1 + \sqrt{3})$), then $\alpha \ge \frac{2+\frac{5}{m}-\frac{2}{mi}}{3-\frac{1}{i}-\frac{1}{\lfloor ip \rfloor}}$ gives $\mathcal{A}_{\mathbb{A}}(\alpha) \ge \frac{i+1}{i} + \frac{1}{ip}$ for any algorithm $\mathbb{A}$ for the scheduling problem.

Proof Consider the following sequence consisting of three phases.

1. αm jobs of length 1
2. $\alpha m - 2$ jobs of length p

3. 1 job of length $p+1$, or $\alpha m - \lceil \frac{\alpha m-2}{i} \rceil - \lceil \frac{\alpha m}{\lfloor ip \rfloor} \rceil$ jobs of length ip

The jobs given in the first phase form an α-sequence with optimal makespan 1. Thus, if two of these jobs are scheduled on the same machine, $\mathcal{A}_{\mathbb{A}}(\alpha) \geq 2$.

Assume now that no two jobs from Phase 1 are placed on the same machine. If two jobs from Phase 2 are scheduled together with a job from Phase 1, or if three Phase 2 jobs are scheduled together, the on-line makespan is at least $2p+1$. In this case, the adversary gives one job of length $p+1$ in Phase 3. The sequence is now an α-sequence with optimal makespan $p+1$, giving a competitive ratio of at least $\frac{2p+1}{p+1}$. Assuming that ($i = 2$ and $p \geq 1+\sqrt{2}$) or ($i = 3$ and $p \geq \frac{1}{2}(1+\sqrt{3})$), this ratio is at least $\frac{i+1}{i} + \frac{1}{ip}$.

Assume now that no two jobs from Phase 2 are scheduled together with a job from Phase 1, and no three jobs from Phase 2 are scheduled together. In this case, the adversary gives $\alpha m - \lceil \frac{\alpha m-2}{i} \rceil - \lceil \frac{\alpha m}{\lfloor ip \rfloor} \rceil$ jobs of length ip in Phase 3. At most $m - \alpha m$ machines could receive two jobs from Phase 2, so at least $(\alpha m - 2) - (m - \alpha m)$ machines receive at least one job from Phase 2. Thus, some job from Phase 3 must be placed with a job from Phase 2 plus another job from either Phase 1 or Phase 2. Assuming that $\alpha \geq \frac{2+\frac{5}{m}-\frac{2}{mi}}{3-\frac{1}{i}-\frac{1}{\lfloor ip \rfloor}}$, we get that $(\alpha m - 2) - (m - \alpha m) + (\alpha m - \lceil \frac{\alpha m-2}{i} \rceil - \lceil \frac{\alpha m}{\lfloor ip \rfloor} \rceil) \geq m+1$. Thus, the on-line makespan is at least $1+p+ip$. Since the jobs from Phase 1 can be scheduled on $\lceil \frac{\alpha m}{\lfloor ip \rfloor} \rceil$ machines, the Phase 2 jobs on $\lceil \frac{\alpha m-2}{i} \rceil$ machines, and the Phase 3 jobs on $\alpha m - \lceil \frac{\alpha m-2}{i} \rceil - \lceil \frac{\alpha m}{\lfloor ip \rfloor} \rceil$ machines, giving a makespan of ip and using only αm machines, the sequence is an α-sequence. Thus, $\mathcal{A}_{\mathbb{A}}(\alpha) \geq \frac{1+p+ip}{ip} = \frac{i+1}{i} + \frac{1}{ip}$. □

6 Conclusion

The accommodating function for $\alpha \leq 1$ seems to be interesting for a variety of on-line problems, possibly for a greater variety than when $\alpha > 1$. For example, the accommodating function for deterministic algorithms for the Paging Problem has a very uninteresting shape for $\alpha > 1$; the value is constant at k, while for $\alpha \leq 1$, $\mathcal{A}(\alpha) = \frac{k}{(1-\alpha)k+1}$.

The study of the accommodating function in general and in particular for $\alpha \leq 1$ has given rise to new algorithms, Unfair-First-Fit and Unfair-First-Fit$_\alpha$.

In addition, the Seat Reservation Problem demonstrates the utility of the accommodating function with $\alpha < 1$ in distinguishing between different algorithms. Three algorithms, which all have competitive ratio close to $\frac{1}{2}$ on accommodating sequences, have different competitive ratios on $\frac{1}{3}$-sequences.

The proofs of positive results seem, in general, to be more interesting than the proofs of negative results. This appears less true for those problems where positive results concerning resource augmentation can be used directly, giving positive results for the accommodating function with $\alpha \leq 1$. However, the original proofs in the resource augmentation setting tend to be interesting, and those results become even more interesting given their application to this setting where the set of request sequences is limited. On the other hand, the accommodating function with $\alpha \leq 1$ sometimes gives much more useful information than resource augmentation. Examples of this were given for the Seat Reservation Problem and Unrestricted Bin Packing, using First-Fit.

References

[1] Y. Azar, J. Boyar, L. Epstein, L. M. Favrholdt, K. S. Larsen, and M. N. Nielsen. Fair versus Unrestricted Bin Packing. *Algorithmica*. To appear.

[2] Y. Azar, L. Epstein, and R. van Stee. Resource Augmentation in Load Balancing. In *Proceedings of the 7th Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 189–199. Springer-Verlag, 2000.

[3] E. Bach, J. Boyar, L. Epstein, L. M. Favrholdt, T. Jiang, K. S. Larsen, G.-H. Lin, and R. van Stee. Tight Bounds on the Competitive Ratio on Accommodating Sequences for the Seat Reservation Problem. *Journal of Scheduling*. To appear.

[4] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[5] J. Boyar and K. S. Larsen. The Seat Reservation Problem. *Algorithmica*, 25:403–417, 1999.

[6] J. Boyar, K. S. Larsen, and M. N. Nielsen. The Accommodating Function: A Generalization of the Competitive Ratio. *SIAM Journal on Computing*, 31(1):233–258, 2001.

[7] M. Brehop, E. Torng, and P. Uthaisombut. Applying Extra Resource Analysis to Load Balancing. *Journal of Scheduling*, 3:273–288, 2000.

[8] M. Chrobak and M. Ślusarek. On Some Packing Problems Related to Dynamic Storage Allocation. *RAIRO Informatique Théoretique et Applications*, 22:487–499, 1988.

[9] R. L. Graham. Bounds for Certain Multiprocessing Anomalies. *Bell Systems Technical Journal*, 45:1563–1581, 1966.

[10] T. R. Jensen and B. Toft. *Graph Coloring Problems*. John Wiley & Sons, 1995.

[11] B. Kalyanasundaram and K. Pruhs. Speed is as Powerful as Clairvoyance. In *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, pages 214–221, 1995.

[12] H. A. Kierstead and J. Qin. Coloring Interval Graphs with First-Fit. *Discrete Mathematics*, 144:47–57, 1995.

[13] H. A. Kierstead and W. T. Trotter. An Extremal Problem in Recursive Combinatorics. *Congressus Numerantium*, 33:143–153, 1981.

[14] E. Koutsoupias. Weak Adversaries for the k-Server Problem. In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, pages 444–449, 1999.

[15] M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive Algorithms for Server Problems. *Journal of Algorithms*, 11(2):208–230, 1990.

[16] D. D. Sleator and R. E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of the ACM*, 28(2):202–208, 1985.

[17] N. Young. On-Line Caching as Cache Size Varies. In *Proceedings of the 2nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 241–250, 1991.