The Relative Worst Order Ratio Applied to Paging*

Joan Boyar, Lene M. Favrholdt**, and Kim S. Larsen

Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark,
{joan,lenem,kslarsen}@imada.sdu.dk

Abstract. The relative worst order ratio, a new measure for the quality of on-line algorithms, was recently defined and applied to two bin packing problems. Here, we apply it to the paging problem. Work in progress by various researchers shows that the measure gives interesting results and new separations for bin coloring, scheduling, and seat reservation problems as well.

Using the relative worst order ratio, we obtain the following results: We devise a new deterministic paging algorithm, Retrospective-LRU, and show that it performs better than LRU. This is supported by experimental results, but contrasts with the competitive ratio. All deterministic marking algorithms have the same competitive ratio, but here we find that LRU is better than FWF. No deterministic marking algorithm can be significantly better than LRU, but the randomized algorithm MARK is better than LRU. Finally, look-ahead is shown to be a significant advantage, in contrast to the competitive ratio, which does not reflect that look-ahead can be helpful.

1 Introduction

The standard measure for the quality of on-line algorithms is the competitive ratio [17, 28, 20], which is, roughly speaking, the worst-case ratio, over all possible input sequences, of the on-line performance to the optimal off-line performance. The definition of the competitive ratio is essentially identical to that of the approximation ratio. This seems natural in that on-line algorithms can be viewed as a special class of approximation algorithms. However, for approximation algorithms, the comparison to an optimal off-line algorithm, OPT, is natural, since the approximation algorithm is compared to another algorithm of the same general type, just with more computing power, while for on-line algorithms, the comparison to OPT is to a different type of algorithm.

Although the competitive ratio has been an extremely useful notion, in many cases, and particularly for the paging problem, it has appeared inadequate at differentiating between on-line algorithms. In a few cases (bin coloring [25] and dual bin packing [9]), one algorithm $\mathbb{A}$ even has a better competitive ratio than another algorithm $\mathbb{B}$, though intuitively, $\mathbb{B}$ is clearly better than $\mathbb{A}$.

When differentiating between on-line algorithms is the goal, performing a direct comparison between the algorithms, instead of involving an intermediate comparison to OPT, seems the obvious choice. A direct comparison on exactly the same sequences will produce the result that many algorithms are not comparable, because one algorithm does well on one type of permutation, while the other does well on another type. With the relative worst order ratio, on-line algorithms are compared directly to each other on their respective worst permutations of sequences. In this way, the relative worst order ratio [7] combines some of the desirable properties of the Max/Max ratio [5] and the random order ratio [22].

The Max/Max Ratio The Max/Max ratio [5] allows direct comparison of two on-line algorithms for an optimization problem, without the intermediate comparison to OPT. Rather than comparing two algorithms on the same sequence, they are compared on their respective worst-case sequences of the same length. The Max/Max Ratio applies only when the length of an input sequence yields a bound on the profit/cost of an optimal solution. Technically, it applies to the paging problem, but the Max/Max ratio of any paging algorithm (deterministic or randomized) approaches 1 as the size of the slow memory approaches infinity.

The Random Order Ratio The random order ratio [22] gives the possibility of considering some randomness of the request sequences without specifying a complete probability distribution. For an on-line algorithm $\mathbb{A}$, the random order ratio is the worst-case ratio, over all input sequences, of the expected performance of $\mathbb{A}$ on a random permutation of the sequence, compared with an optimal solution. If, for all possible input sequences, any permutation of the sequence is equally likely, this ratio gives a meaningful worst-case measure of how well an algorithm can do. Unfortunately, the random order ratio seems to be difficult to compute.

The Relative Worst Order Ratio With the relative worst order ratio, one considers the worst-case performance over all permutations instead of the average-case performance as with the random order ratio. Thus, when comparing two on-line algorithms, one considers a worst-case sequence and takes the ratio of how the two algorithms perform on their respective worst permutations of that sequence. Note that the two algorithms may have different worst permutations for the same sequence. The relative worst order ratio is formally defined in Section 2.

The relative worst order ratio can be viewed as a worst case version of Kenyon's random order ratio, with the modification that on-line algorithms are compared directly, rather than indirectly through OPT. It can also be viewed as a modification of the Max/Max ratio, where a finer partition of the request sequences is used; instead of finding the worst sequence among those having the same length, one finds the worst sequence among those which are permutations of each other. This particular finer partition was inspired by the random order ratio.

The Paging Problem We consider the well studied paging problem. The input sequence consists of requests for pages in a slow memory, which contains N pages. There is a fast memory, the cache, which has space for $k < N$ pages. A request for a page currently in cache is a *hit*, while a request for a page not in cache is a *page fault*. When a page fault occurs, the requested page must be brought into cache. If the cache already contains k pages when this happens, at least one of these must be evicted. A paging algorithm decides which page to evict on a fault. This decision must usually be made on-line, i.e., without any knowledge about future requests. The goal is to minimize the number of faults.

Paging Algorithms Two major classes of deterministic algorithms for the paging problem are conservative algorithms [32] and marking algorithms [6].

A paging algorithm $\mathbb{A}$ is called *conservative*, if no request sequence has a consecutive subsequence with requests to at most k distinct pages causing $\mathbb{A}$ to fault more than k times. The algorithms, Least-Recently-Used (LRU) and First-In/First-Out (FIFO) are examples of conservative algorithms. On a page fault, LRU evicts the least recently used page in cache and FIFO evicts the page which has been in cache longest.

Marking algorithms work in phases. Each time a page is requested, this page is marked (implicitly in the analysis or explicitly by the algorithm). When a page must be evicted, one of the unmarked pages is chosen, if one exists. Otherwise all marks are erased, and the

requested page is marked. This request starts a new phase. Note that LRU is a marking algorithm, whereas FIFO is not. Another example of a marking algorithm is Flush-When-Full (FWF), the algorithm which evicts all pages in cache at the end of each phase. The randomized marking algorithm MARK chooses the unmarked page to be evicted uniformly at random.

Previous Results All conservative and marking algorithms have competitive ratio k [31, 29] and this is optimal among deterministic algorithms [28]. However, in practice, these algorithms do not all have the same performance: LRU is better than FIFO and much better than FWF [32]. Moreover, results from [15] suggest there may be algorithms that perform even better than LRU.

In [3] an alternative model, the Max-/Average-Model, for the paging problem capturing locality of reference was suggested. It was proven that, in this model, LRU is slightly better than FIFO, but LRU is still best possible among deterministic algorithms. Using access graphs, it has been proven that LRU is better than FIFO [12] and algorithms have been designed that are better than LRU [6]. Hence, these alternative ways of measuring the quality of paging algorithms give more satisfactory results. However, they are only defined for paging and paging-like problems.

In contrast to deterministic algorithms, MARK [14] has a competitive ratio of $2H_k - 1$[1], where H_k is the kth harmonic number, i.e., $H_k = \sum_{i=1}^{k} \frac{1}{i} \approx \ln k$. Other randomized algorithms have been shown to have the optimal competitive ratio for randomized algorithms of H_k [26, 1].

Look-Ahead. Look-ahead, where the algorithm deciding which page to evict is allowed to see the next ℓ page requests before making that decision, is a model which intuitively lies between on-line and off-line. It is well known that look-ahead cannot reduce the competitive ratio of any algorithm, but clearly it can be useful when it can be implemented.

Previously, alternative definitions of look-ahead have led to results showing that look-ahead helps. In each case, the algorithm is allowed to see a sequence of future requests satisfying some property. Young [31] proposed *resource-bounded look-ahead*, where the sequence is a maximal sequence of future requests for which it would incur ℓ page faults, Albers [2] proposed *strong look-ahead*, where the sequence of future requests contains ℓ distinct pages different from the current request, and Breslauer [11] proposed *natural look-ahead*, where the sequence of future requests contains ℓ pages not currently in cache. Here we retain the original definition, so the algorithm is only allowed to see the next ℓ pages, regardless of what they are.

The Max/Max Ratio [5] has been somewhat successfully applied to the standard definition of look-ahead, showing that a greedy strategy achieves a Max/Max ratio of $\frac{N-1}{\ell}$ for $N - k < \ell \leq N - 1$ (recall that N is the size of the slow memory). *Comparative analysis* [24] is more successful, showing that look-ahead gives a result which is a factor $\min\{k, \ell + 1\}$ better than without look-ahead. This is the same result we obtain with the relative worst order ratio.

Other Measures. Many alternatives to or variations on the competitive ratio have been proposed. We have already mentioned the Max/Max ratio, the random order ratio, access graphs, the Max-/Average-Model, and comparative analysis. Other alternatives are Markov paging [21], diffuse adversaries [24], extra resource analysis [19, 28], the accommodating function [9], and statistical adversaries [27]. Most of these techniques have been applied to only a few closely related problems. So far, the techniques which have been applied to a broader range of problems, extra resource analysis and the accommodating function, for instance, have given new separation results for only a limited number of different types of problems.

The Relative Worst Order Ratio. The relative worst order ratio has already been applied quite successfully to two very different problem types: bin packing [7] and now paging. For Classical Bin Packing, Worst-Fit is better than Next-Fit according to the relative worst order ratio, even though they have the same competitive ratio [18]. Thus, the advantage of keeping all bins open, instead of just one, is reflected by the relative worst order ratio. For Dual Bin Packing, the relative worst order ratio shows that First-Fit is better than Worst-Fit, while the competitive ratio indicates the opposite [9].

Other New Results on the Relative Worst Order Ratio The wide applicability of the relative worst order ratio has been confirmed by other new results. Recently, various researchers have applied the relative worst order ratio to other problems and obtained separations not given by the competitive ratio, but consistent with intuition and/or practice.

A few simple examples are given in [13]. For instance, for the problem of minimizing makespan on two related machines with speed ratio s, the optimal competitive ratio of $\frac{s+1}{s}$ for $s \geq \Phi \approx 1.618$ is obtained both by the post-greedy algorithm, which schedules each job on the machine where it will finish earliest, and by the algorithm which simply schedules all jobs on the fast machine. In contrast, the relative worst order ratio shows that the post-greedy algorithm is better. A similar result is obtained for the problem of minimizing makespan on $m \geq 2$ identical machines with preemption.

The relative worst order ratio was also found by [23] to give the intuitively correct result for the bin coloring problem, where the competitive ratio gives the opposite result [25]: a trivial algorithm using only one open bin has a better competitive ratio than a natural greedy-type algorithm.

The proportional price version of the seat reservation problem has largely been ignored due to very negative impossibility results using competitive analysis [8]. However, algorithms for the problem were compared and separated using the relative worst order ratio in [10].

Our Results First, we propose a new algorithm, Retrospective-LRU (RLRU), which is a variation on LRU that takes into account which pages would be in the cache of the optimal off-line algorithm, LFD, if it were given the subsequence of page requests seen so far. We show that, according to the relative worst order ratio, RLRU is better than LRU. This is interesting, since it contrasts with results on the competitive ratio and with results in [3] where a new model of locality of reference is studied.

It is easily shown that RLRU does not belong to either of the common classes of algorithms, conservative and marking algorithms which all have the optimal competitive ratio k. In fact, the competitive ratio of RLRU is $k + 1$ and thus slightly worse than that of LRU. Initial testing of RLRU shows that it performs better than LRU.

Analyzing paging algorithms with the relative worst order ratio, we obtain more detailed information than with competitive analysis. With the relative worst order ratio, LRU is better than FWF, so not all marking algorithms are equivalent, but no marking algorithm is significantly better than LRU. All conservative algorithms are equivalent, so LRU and FIFO have the same performance.

Look-ahead is shown to help significantly with respect to the relative worst order ratio. Compared to the competitive ratio which does not reflect that look-ahead can be of any use, this is a very nice property of the relative worst order ratio.

The definition of the relative worst order ratio is extended to randomized algorithms, and MARK and LRU are compared, giving that MARK is the better algorithm, as with the competitive ratio.

Most proofs have been moved to the appendix.

	minimization	maximization
$\mathbb{A}$ better than $\mathbb{B}$	< 1	> 1
$\mathbb{B}$ better than $\mathbb{A}$	> 1	< 1

Table 1. Ratio values for minimization and maximization problems

2   The Relative Worst Order Ratio.

The definition of the relative worst order ratio uses $\mathbb{A}_W(I)$, the performance of an algorithm $\mathbb{A}$ on the worst permutation of the input sequence I, formally defined in the following way.

Definition 1. *Consider an optimization problem P, let I be any input sequence, and let n be the length of I. If σ is a permutation on n elements, then $\sigma(I)$ denotes I permuted by σ. Let $\mathbb{A}$ be any algorithm for P.*

If P is a maximization *problem, $\mathbb{A}(I)$ is the profit of running $\mathbb{A}$ on I, and $\mathbb{A}_W(I) = \min_\sigma \mathbb{A}(\sigma(I))$.*

If P is a minimization *problem, $\mathbb{A}(I)$ is the cost of running $\mathbb{A}$ on I, and $\mathbb{A}_W(I) = \max_\sigma \mathbb{A}(\sigma(I))$.*

For many on-line problems, some algorithms perform well on particular types of permutations of the input, while other algorithms perform well on other types of permutations. The purpose of comparing on the worst permutation of sequences, rather than on each sequence independently, is to be able to differentiate between such pairs of algorithms, rather than just concluding that they are incomparable. Sequences with the same "content" are considered together, but the measure is worst case, so the algorithms are compared on their respective worst permutations. This was originally motivated by problems where all permutations are equally likely, but appears to be applicable to other problems as well.

Definition 2. *Let S_1 and S_2 be statements about algorithms $\mathbb{A}$ and $\mathbb{B}$ defined in the following way.*

$S_1(c)$: There exists a constant b such that $\mathbb{A}_W(I) \leq c \cdot \mathbb{B}_W(I) + b$ for all I.

$S_2(c)$: There exists a constant b such that $\mathbb{A}_W(I) \geq c \cdot \mathbb{B}_W(I) - b$ for all I.

The relative worst order ratio *$WR_{\mathbb{A},\mathbb{B}}$ of on-line algorithm $\mathbb{A}$ to algorithm $\mathbb{B}$ is defined if $S_1(1)$ or $S_2(1)$ holds. In this case, $\mathbb{A}$ and $\mathbb{B}$ are said to be* comparable.

If $S_1(1)$ holds, then $WR_{\mathbb{A},\mathbb{B}} = \sup\{r \mid S_2(r)\}$, and if $S_2(1)$ holds, then $WR_{\mathbb{A},\mathbb{B}} = \inf\{r \mid S_1(r)\}$.

The statements $S_1(1)$ and $S_2(1)$ check that the one algorithm is always at least as good as the other on every sequence (on their respective worst permutations). When one of them holds, the relative worst order ratio is a bound on how much better the one algorithm can be. Note that if $S_1(1)$ holds, the supremum involves S_2 rather than S_1, and vice versa. A ratio of 1 means that the two algorithms perform identically with respect to this quality measure; the further away from 1, the greater the difference in performance. The ratio may be greater than or less than one, depending on whether the problem is a minimization or a maximization problem and on which of the two algorithms is better. These possibilities are illustrated in Table 1.

It is easily shown [7] that the relative worst order ratio is a *transitive measure*, i.e., for any three algorithms $\mathbb{A}$, $\mathbb{B}$, and $\mathbb{C}$, $WR_{\mathbb{A},\mathbb{B}} \geq 1$ and $WR_{\mathbb{B},\mathbb{C}} \geq 1$ implies $WR_{\mathbb{A},\mathbb{C}} \geq 1$. The proof of transitivity shows that when $WR_{\mathbb{A},\mathbb{B}} \geq 1$, $WR_{\mathbb{B},\mathbb{C}} \geq 1$, and both are bounded above by some constant, then $\max\{WR_{\mathbb{A},\mathbb{B}}, WR_{\mathbb{B},\mathbb{C}}\} \leq WR_{\mathbb{A},\mathbb{C}} \leq WR_{\mathbb{A},\mathbb{B}} \cdot WR_{\mathbb{B},\mathbb{C}}$. Thus, when a new algorithm is analyzed, it need not be compared to all other algorithms.

Although one of the goals in defining the relative worst order ratio was to avoid the intermediate comparison of any on-line algorithm, $\mathbb{A}$, to the optimal off-line algorithm, OPT, it is still possible to compare on-line algorithms to OPT. In this case, the measure is called the *worst order ratio* of $\mathbb{A}$ and denoted $WR_{\mathbb{A}}$.

3 A Better Algorithm than LRU

In this section, we introduce an algorithm which turns out to be better than LRU according to the relative worst order ratio. This is supported by initial experimental results but is in contrast to the competitive ratio which says that LRU is best possible among deterministic algorithms. The algorithm, called Retrospective-LRU (RLRU), is defined in Figure 1. The name comes from the algorithm's marking policy. When evicting pages, RLRU uses the LRU policy, but it chooses only from the unmarked pages in cache, unless they are all marked. Marks are set according to what the optimal off-line algorithm, LFD [4], would have in cache, if given the part of the sequence seen so far. LFD is the algorithm that, on a fault, evicts the page that will be requested farthest in the future.

If RLRU has a fault and LFD does not, RLRU marks the page requested. If RLRU has a hit, the page p requested is marked if it is different from the page of the previous request. Requiring the page to be different from the previous page ensures that at least one other page has been requested since p was brought into the cache. A phase of the execution starts with the removal of all marks and this occurs whenever there would otherwise be a second fault on the same page within the current phase.

Theorem 1. $WR_{LRU,RLRU} = \frac{k+1}{2}$.

```
The first phase begins with the first request.

On request r to page p:
    Update p's timestamp
    if p is not in cache then
        if there is no unmarked page then
            evict the least recently used page in cache
        else
            evict the least recently used unmarked page
        if this is the second fault on p
          since the start of the current phase then
            unmark all pages
            start a new phase with r
        if p was in LFD's cache just before this request then
            mark p
    else
        if p is different from the previous page then
            mark p
```

Fig. 1. Retrospective-LRU (RLRU)

The proof of Theorem 1 relies on a few basic properties of RLRU. Modifications to the algorithm which do not change these basic properties will result in other algorithms which, according to the relative worst order ratio, are also better than LRU. One example of this is the test as to whether or not the current page is the same as the previous. This test could

be removed and the page marked unconditionally or never marked, and the proofs still hold. Another example is the decision when to end a phase. The most important property is that each phase consists of requests to at least $k + 1$ distinct pages and there is at most one fault on each of them. This leaves room for experimentally testing a number of variations, and it could lead to algorithms which are even better in practice than the one we present here.

Note that RLRU is neither a conservative nor a marking algorithm. This can be seen from the sequence $\langle p_1, p_2, p_3, p_4, p_1, p_2, p_3, p_4, p_3 \rangle$ for $k = 3$, where RLRU faults on every request.

In contrast to Theorem 1, the competitive ratio of RLRU is slightly worse than that of LRU:

Theorem 2. *The competitive ratio of RLRU is $k + 1$.*

When LRU and RLRU are compared to OPT using the worst order ratio, instead of the competitive ratio, one finds that they have the same ratio k, confirming the intuition that information is lost in an indirect comparison to OPT.

Theorem 3. $WR_{LRU} = WR_{RLRU} = k$.

Implementation As outlined in the appendix, RLRU can be implemented to run in time $O(\log N)$ and space $O(N)$, where N is the number of different pages requested. The question is whether or not these time and space bounds are good enough in practice. We believe there are at least two interesting scenarios to consider. One is the interaction between two high speed storage media, the speed of which differ by only a small multiplicative constant, such as primary versus secondary cache. Here, a paging algorithm must be very efficient, which also implies that it cannot be allowed much working space. In such a scenario, even LRU is most often too time and space consuming. Another scenario is the interaction of storage media, the speed of which differ by orders of magnitude. This could be the buffer pool versus the disk in database systems or local file caching of Internet files. In those situations, we can use substantial space, and time logarithmic in either the number of different pages or just in the cache size would both be insignificant compared with almost any small improvement in cache behavior. A similar point is made in [15]. If, in some special application, space is a problem, then it could possibly be reduced to a function of k using the techniques of [1]. In summary, a comparison between LRU and RLRU is interesting because the circumstances under which they can reasonably be applied are close to identical.

Empirical Analysis To see if the positive theoretical results are also reflected in practice, we have investigated the behavior of LRU and RLRU on traces[1] collected from very different applications, including key words searches in text files, selections and joins in the Postgres database system, external sorting, and various kernel operations. We have used all ten data files from the site.

In Table 2, we list the results for each data file, and for cache sizes of $8, 16, \ldots, 1024$. Each entry shows the percent-wise improvement of RLRU over LRU. If ℓ and r denote the number of faults by LRU and RLRU, respectively, then the improvement is computed as $100\frac{\ell - r}{\ell}$. This number is negative if LRU performs best. In the appendix, the same table is given, but each entry in the table shows also the number of page faults of each of the three algorithms LFD, LRU, and RLRU. Similar results are also presented there for a variant of RLRU.

Out of the 80 tests, 16 are negative. The largest negative result of -0.74% is from a short sequence and is due to a difference of only one page fault. The remaining negative results lie

---

Cache	File names and lengths									
	bigsort	j1	j2	j3	j4	j5	j6	pjoin	pq7	xds
Size	40167	18533	25881	38112	59744	95723	20709	41558	32989	88558
8	9.76	0.61	0.44	-0.33	0.00	0.11	10.88	10.16	-0.51	0.41
16	14.92	0.43	0.53	-0.28	-0.28	0.02	1.54	9.26	0.19	0.49
32	1.70	8.21	0.74	-0.21	-0.08	0.02	0.07	6.93	0.50	0.55
64	1.75	-0.74	0.78	-0.21	-0.07	-0.01	-0.02	0.60	0.89	0.55
128	1.48	0.00	0.78	-0.26	-0.08	0.18	-0.04	0.10	0.07	0.51
256	0.70	0.00	0.75	-0.19	0.01	1.17	0.16	0.19	0.06	0.28
512	1.35	0.00	0.71	-0.05	0.36	1.83	0.11	0.43	0.39	0.39
1024	0.89	0.00	17.85	6.35	0.52	1.38	0.13	0.03	0.77	0.07

Table 2. Empirical comparison of LRU and RLRU

between zero and approximately half a per cent. RLRU beats LRU with more than half a per cent in 32 cases, more than 1% in 17 cases, and more than 5% in 9 cases. This is illustrated in Figure 2.


Fig. 2. Percentages with which RLRU is better than LRU

4 Conservative and Marking Algorithms

The following technical lemma is used extensively in the proofs of the results in this section and in the proof of Theorem 1.

Lemma 1. *For any sequence I of page requests, there exists a worst permutation of I for LRU with all faults appearing before all hits.*

It is easy to see that both LRU and FIFO are conservative algorithms [32]: between any two faults on the same page there must be requests to at least k other pages. Using Lemma 1, we can prove that for any sequence I, there exists a permutation $I_{\mathbb{C}}$ of I which is worst possible for any conservative algorithm and that all conservative algorithms behave exactly the same when given $I_{\mathbb{C}}$.

Theorem 4. *For any pair of conservative algorithms $\mathbb{C}_1$ and $\mathbb{C}_2$, $WR_{\mathbb{C}_1,\mathbb{C}_2} = 1$.*

Therefore, by Theorem 1 and the transitivity of the relative worst order ratio, we have the following:

Corollary 1. *For any conservative algorithm $\mathbb{C}$, $WR_{\mathbb{C},RLRU} = \frac{k+1}{2}$.*

In contrast to the competitive ratio, the relative worst order ratio distinguishes between different marking algorithms. In particular, LRU is better than FWF, as it is in practice:

Theorem 5. *For any conservative algorithm $\mathbb{C}$, $WR_{FWF,\mathbb{C}} = \frac{2k}{k+1}$.*

Furthermore, LRU is close to being a best possible marking algorithm:

Theorem 6. *For any marking algorithm $\mathbb{M}$ and any conservative algorithm $\mathbb{C}$ with $WR_{\mathbb{C},\mathbb{M}}$ defined, $WR_{\mathbb{C},\mathbb{M}} \leq \frac{k+1}{k}$.*

5   Look-Ahead

In the standard on-line model, requests arrive one by one. A model in which the algorithm is informed of the next $\ell \geq 1$ page requests before servicing the current one, is a look-ahead model. This model is in-between the standard on-line model and the off-line model.

It is well known that using standard competitive analysis one cannot show that knowing the next ℓ requests is any advantage for any fixed ℓ; for any input sequence, an adversary can "fill up" the look-ahead by using $\ell+1$ consecutive copies of each request, adding no cost to the optimal off-line solution. In contrast, results on the relative worst order ratio, indicate that look-ahead helps significantly. Here we only look at a modification of LRU, using look-ahead, though the technique can be applied to other algorithms as well.

Define $\mathrm{LRU}(\ell)$ to be the algorithm which on a fault evicts the least recently used page in cache which is not among the next ℓ requests. If $\ell \geq k$, all pages in cache may be among the next ℓ requests. In this case, the page whose next request is farthest in the future is evicted. Note that by transitivity, Theorem 7 shows that for any conservative algorithm, $\mathbb{C}$, $\mathrm{WR}_{\mathbb{C},\mathrm{LRU}(\ell)} = \min\{k, \ell+1\}$.

Theorem 7. $WR_{LRU,LRU(\ell)} = \min\{k, \ell+1\}$.

6   Randomized Algorithms

The relative worst order ratio can also be applied to randomized algorithms. The only change to the definition is that an algorithm's expected profit/cost on a worst permutation of a sequence is used in place of the profit/cost obtained by a deterministic algorithm.

Definition 3. *Consider an optimization problem P and let I be any input sequence of length n. If σ is a permutation on n elements, then $\sigma(I)$ denotes I permuted by σ. Let $\mathbb{A}$ be any randomized algorithm for P.*

If P is a maximization *problem, $E[\mathbb{A}(I)]$ is the expected profit of running $\mathbb{A}$ on I, and $\mathbb{A}_W(I) = \min_\sigma E[\mathbb{A}(\sigma(I))]$. If P is a* minimization *problem, $E[\mathbb{A}(I)]$ is the expected cost of running $\mathbb{A}$ on I, and $\mathbb{A}_W(I) = \max_\sigma E[\mathbb{A}(\sigma(I))]$.*

Using the above definition, the relative worst order ratio is now defined as in the deterministic case.

Consider the randomized paging algorithm MARK [14]. On a fault, MARK chooses the unmarked page to be evicted uniformly at random. In this section, we show that $\mathrm{WR}_{\mathrm{LRU,MARK}} = k/H_k$. This is consistent with the results one obtains with the competitive ratio where MARK has ratio $2H_k - 1$ [1], while LRU has ratio k.

Recall that marking algorithms, such as MARK, work in phases. In each phase (except possibly the last), exactly k distinct pages are requested, and the first page requested within a phase was not requested in the previous phase. Thus, the subsequence processed within a phase (except possibly the last) is a maximal subsequence containing requests to exactly k distinct pages. A subsequence processed within one marking phase is called a *k-phase*. Note that the partitioning of a sequence into k-phases is independent of the particular marking algorithm.

For Lemma 2 and Theorem 8 below, we need the fact that MARK's expected number of faults in the ith k-phase is $m_i(H_k - H_{m_i} + 1)$ [14], where m_i is the number of *new* pages in the ith phase, i.e., the number of pages that are requested in the ith phase and not in the $(i-1)$st phase.

Lemma 2. *There exists a sequence I, which is a worst permutation for both MARK and LRU, where MARK's expected number of faults is H_k per k-phase, while LRU faults k times per k-phase.*

Proof. Consider the sequence with $k+1$ pages repeated cyclicly. □

Lemma 3. *For any sequence I of page requests, there exists a worst permutation I_{MARK} of I with respect to MARK, such that all k-phases, except possibly the last, have the following properties:*

1. *The first page requested in a k-phase did not appear in the previous k-phase.*
2. *There are exactly k requests, all to distinct pages.*

Lemma 4 below uses the following definition of rare and frequent pages and blocks.

Definition 4. *Consider a sequence S consisting of $s \geq 2$ consecutive k-phases. Call pages requested in every phase of S* frequent *pages, and the others* rare *pages. The sequence S is called a* block, *if it has the following properties.*

1. *Each k-phase in S contains exactly $s-1$ rare pages.*
2. *There is no $r < s$ such that the first $r \geq 2$ k-phases of S contain exactly $r-1$ rare pages.*

Note that any sequence with m k-phases contains at least $\lfloor \frac{m}{k+1} \rfloor$ consecutive blocks.

Lemma 4. *There exists a constant b such that, for any sequence I, $LRU_W(I) \geq MARK_W(I) - b$.*

Proof. For any request sequence I, consider a worst permutation I_{MARK} of I with respect to MARK, satisfying the conditions of Lemma 3. Partition the sequence I_{MARK} in blocks. Each block in the partition will be analyzed separately, and it will be shown that the sequence can be permuted so that LRU faults at least as many times as the expected number of faults by MARK on the requests of that block.

Consider a block, S, containing $s+1$ k-phases and thus s rare pages and $k-s$ frequent pages in each k-phase. Clearly, no frequent page is a new page in any of the last s k-phases of S. Therefore, if the first k-phase, P_1, in the block has at most s new pages, then MARK's expected number of faults is at most $s(s+1)(1 + H_k - H_s)$.

Since each rare page occurs at most s times in S, one can permute the block into s groups of $k+1$ distinct pages, plus $(s+1)k - s(k+1) = k - s$ extra pages. Thus, LRU can be forced to fault $s(k+1)$ times. If P_1 has at most s new pages, MARK's expected number of faults on this block is at most $s(s+1)(1 + H_k - H_s) \leq s(s+1)(1 + \frac{k-s}{s+1}) = s(k+1)$, so in this case the result holds.

Now assume that the first k-phase, P_1, in the block, S, has $s+i$ new pages, where $0 < i \leq k-s$. Then, some frequent page in P_1 is also a new page. MARK's expected number of faults is at most $s^2(1 + H_k - H_s) + (s+i)(1 + H_k - H_{s+i})$.

Let S' be the block immediately preceding S. Assume that it contains $s'+1$ k-phases and thus s' rare pages in each k-phase. Consider any frequent, new page, p, in P_1. It is clearly not a frequent page in S'. Assume for a moment that p occurs in all but the last k-phase of

10

S'. In this case, the first s' k-phases of S' have at least one more frequent page than all of S' does. Generally, removing k-phases from the end of a block cannot decrease the number of frequent pages, and the first two k-phases have at most $k-1$ frequent pages. Thus, removing k-phases from the end of S', we would eventually end up with $2 \leq r < s+1$ consecutive k-phases with $r-1$ rare pages. This contradicts the fact that S' is a block, so p occurs at most $s'-2$ times in S'. Hence, one can choose i requests to frequent, new pages in P_1 which can be moved back into the previous block, S', permuting S' such that LRU faults on these i pages, in addition to the $s'(k+1)$ pages originally in S' which it faults on. After removing these i requests from S, there are still s requests to rare pages in each k-phase, and a total of at least $s(k+1)$ requests in S, so the remaining requests can still be permuted to give LRU $s(k+1)$ faults. Thus, one can count $s(k+1)+i$ requests from block S which LRU will fault on. The lemma now follows by the following claim, which is proven in the appendix. $\qquad \square$

Claim 1. For all integers s, i, k, such that $1 \leq s \leq k$, $i \geq 0$, and $s+i \leq k$,

$$s^2(1 + H_k - H_s) + (s+i)(1 + H_k - H_{s+i}) \leq s(k+1) + i.$$

Theorem 8. $WR_{LRU,MARK} = k/H_k$.

Proof. The lower bound follows from Lemmas 2 and 4. To see that the ratio cannot be higher than k/H_k, consider any k-phase in LRU's worst permutation. LRU never faults more than k times on any k-phase, and MARK never has an expected number of faults less than H_k on any complete k-phase [14]. MARK would fault at least as many times on its own worst ordering. Thus, the result is tight. $\qquad \square$

7 Conclusion and Open Problems

This second problem area, paging, studied using the relative worst order ratio gives even more convincing evidence than the first, bin packing, that this new performance measure could become an important tool for analyzing on-line algorithms. Comparing algorithms directly to each other, rather than doing it indirectly through a comparison to OPT, appears to give more meaningful results, both for paging and for bin packing. Previous measures and models, proposed as alternatives or supplements to the competitive ratio, have been more limited as to applicability, usually to very few problems. Further study is needed to determine how widely applicable the relative worst order ratio is, but paging and bin packing are very different problems. Together with the results on the bin coloring, scheduling, and seat reservation problems mentioned in the introduction, this gives a convincing basis for further investigation.

For paging, many algorithms with widely varying performance in practice all have a competitive ratio of k. The relative worst order ratio is able to distinguish between some of these algorithms. Most notably, LRU is found to be better than FWF, and look-ahead is shown to help. It is also promising that this new performance measure is leading to the discovery of new algorithms. Further testing is needed to determine which variant of RLRU is best in practice and how much better it is than LRU.

Theorem 6 shows that no marking algorithm can be much better than LRU. It would be interesting to know if LRU is in fact the best marking algorithm according to the relative worst order ratio.

References

1. D. Achlioptas, M. Chrobak, and J. Noga. Competitive analysis of randomized paging algorithms. *Theoretical Computer Science*, 234(1–2):203–218, 2000.

2. S. Albers. On the influence of lookahead in competitive paging algorithms. *Algorithmica*, 18:283–305, 1997.

3. S. Albers, L. M. Favrholdt, and O. Giel. On paging with locality of reference. In *34th Annual ACM Symposium on the Theory of Computing*, pages 258–267, 2002.

4. L. A. Belady. A study of replacement algorithms for virtual storage computers. *IBM Systems Journal*, 5:78–101, 1966.

5. S. Ben-David and A. Borodin. A new measure for the study of on-line algorithms. *Algorithmica*, 11(1):73–91, 1994.

6. A. Borodin, S. Irani, P. Raghavan, and B. Schieber. Competitive paging with locality of reference. *Journal of Computer and System Sciences*, 50(2):244–258, 1995.

7. J. Boyar and L. M. Favrholdt. The relative worst order ratio for on-line algorithms. In *5th Italian Conference on Algorithms and Complexity*, volume 2653 of *LNCS*, pages 58–69, 2003.

8. J. Boyar and K. S. Larsen. The seat reservation problem. *Algorithmica*, 25:403–417, 1999.

9. J. Boyar, K. S. Larsen, and M. N. Nielsen. The accommodating function — a generalization of the competitive ratio. *SIAM Journal of Computation*, 31(1):233–258, 2001.

10. J. Boyar and P. Medvedev. The relative worst order ratio applied to seat reservation. In *Proceedings of the Ninth Scandinavian Workshop on Algorithm Theory*, 2004. To appear.

11. D. Breslauer. On competitive on-line paging with lookahead. *Theoretical Computer Science*, 209(1–2):365–375, 1998.

12. M. Chrobak and J. Noga. LRU Is Better than FIFO. *Algorithmica*, 23:180–185, 1999.

13. L. Epstein, L. M. Favrholdt, and J. S. Kohrt. The relative worst order ratio applied to scheduling problems. Work in progress, 2003.

14. A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. *Journal of Algorithms*, 12:685–699, 1991.

15. A. Fiat and Z. Rosen. Experimental studies of access graph based heuristics: Beating the LRU standard? In *8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 63 – 72, 1997.

16. J. S. Frederiksen and K. S. Larsen. On-line seat reservations via off-line seating arrangements. In *Eighth International Workshop on Algorithms and Data Structures*, volume 2748 of *LNCS*, pages 174–185. Springer-Verlag, 2003.

17. R. L. Graham. Bounds for certain multiprocessing anomalies. *Bell Systems Technical Journal*, 45:1563–1581, 1966.

18. D. S. Johnson. Fast algorithms for bin packing. *Journal of Computer and System Sciences*, 8:272–314, 1974.

19. B. Kalyanasundaram and K. Pruhs. Speed is as powerful as clairvoyance. *Journal of the ACM*, 47:617–643, 2000.

20. A. R. Karlin, M. S. Manasse, L. Rudolph, and D. D. Sleator. Competitive snoopy caching. *Algorithmica*, 3(1):79–119, 1988.

21. A. R. Karlin, S. J. Phillips, and P. Raghavan. Markov paging. *SIAM Journal of Computation*, 30:906–922, 2000.

22. C. Kenyon. Best-fit bin-packing with random order. In *7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 359–364, 1996.

23. J. S. Kohrt. The relative worst order ratio applied to bin coloring. Work in progress, 2003.

24. E. Koutsoupias and C. H. Papadimitriou. Beyond competitive analysis. In *35th Annual Symposium on Foundations of Computer Science*, pages 394–400, 1994.

25. S. O. Krumke, W. E. de Paepe, J. Rambau, and L. Stougie. Online bin coloring. In *9th Annual European Symposium on Algorithms*, volume 2161 of *LNCS*, pages 74–85, 2001.

26. L. A. McGeoch and D. D. Sleator. A strongly competitive randomized paging algorithm. *Algorithmica*, 6:816–825, 1991.

27. P. Raghavan. A statistical adversary for on-line algorithms. In D. D. Sleator L. A. McGeoch, editor, *On-Line Algorithms*, volume 7 of *Series in Discrete Mathematics and Theoretical Computer Science*, pages 79–83. American Mathematical Society, 1992.

28. D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

29. E. Torng. A unified analysis of paging and caching. *Algorithmica*, 20:175–200, 1998.

30. M. J. van Kreveld and M. H. Overmars. Union-copy structures and dynamic segment trees. *Journal of the ACM*, 40(3):635–652, 1993.

31. N. Young. Competitive paging and dual-guided algorithms for weighted caching and matching (thesis). Technical Report CS-TR-348-91, Computer Science Department, Princeton University, 1991.

32. N. Young. The k-server dual and loose competitiveness for paging. *Algorithmica*, 11:525–541, 1994.

A Additional Comments on an Implementation of RLRU

In addition to the administration required to evict least recently used pages, which is similar to the administration necessary for LRU, RLRU needs to be able to perform the following operations:

1. Check if it faults on a page for the second time in a phase.
2. Mark a page, and unmark all pages.
3. Find the least recently used page, possibly just among unmarked pages.
4. Check for a page in LFD's cache.

These can all be implemented in time $O(\log N)$ and space $O(N)$, where N is the number of different pages requested.

1. We use a balanced binary search tree over all the different pages on which RLRU has faulted during the phase.
2. Using a balanced binary search tree over all the different pages which have been requested, we mark a page by associating the current phase number with the page. Thus, by incrementing the phase number, we can unmark all pages in constant time.
3. Using a balanced binary search tree ordered on timestamp, the least recently used page can be found in logarithmic time. If the timestamp is also associated with pages in cache, then old timestamp entries can be found and updated when a page is requested. By adding information to the nodes in the tree regarding the last phase in which the page stored in the node was marked and information regarding the least recent phase of any node in the subtree, it is also possible in logarithmic time to find the least recently used page among those which are unmarked, i.e., not marked in the current phase. In an actual implementation, points 1, 2, and 3 can be combined.
4. At any given point in time, it is of course impossible to compute the entire contents of LFD's cache, since this depends on future requests. It is, however, possible, given a request and the request sequence up to that point, to compute whether or not LFD would have that particular page in cache. Using techniques [16] inspired by geometric algorithms [30], this can be done by registering the known time intervals of pages in LFD's cache in a balanced binary search tree. Also here, time $O(\log N)$ and space $O(N)$ can be obtained.

B Further Empirical Analysis of RLRU

In Table 3, we give the full table from which the percentages were shown in Table 2. As described before Table 2, we list the results for each data file, and for cache sizes of $8, 16, \dots, 1024$. Each entry shows the percent-wise improvement of RLRU over LRU. If ℓ and r denote the number of faults by LRU and RLRU, respectively, then the improvement is computed as $100\frac{\ell - r}{\ell}$. This number is negative if LRU performs best. Here in the appendix, in addition to the percentages, each entry shows the number of page faults of each of the three algorithms LFD, LRU, and RLRU, in that order. For convenience, a copy of Figure 2 is shown in Figure 3 below Table 3.

We also consider another variant, RLRU$'$, of RLRU. The only difference is that RLRU$'$ never marks pages that are already in cache. Thus, RLRU$'$ is defined as in Figure 1 with the else-statement deleted. For this variant, we obtain the results displayed in Table 4. For this algorithm, only 8 out of the 80 tests are negative. Except for the result of -1.01%, all results are larger than $-\frac{1}{3}\%$. RLRU$'$ beats LRU with more than $\frac{1}{3}\%$ in 39 cases, more than 1% in 13 cases, and more than 5% in 6 cases. The distribution of the percentages can be seen in Figure 4.

Cache Size	File names and lengths									
	bigsort	j1	j2	j3	j4	j5	j6	pjoin	pq7	xds
	40167	18533	25881	38112	59744	95723	20709	41558	32989	88558
8	*11080*	*418*	*8110*	*4194*	*7064*	*25169*	*4490*	*6906*	*9046*	*10665*
	14632	494	8233	4262	7278	25412	5100	8014	9371	10768
	13204	491	8197	4276	7278	25385	4545	7200	9419	10724
	9.76	**0.61**	**0.44**	**-0.33**	**0.00**	**0.11**	**10.88**	**10.16**	**-0.51**	**0.41**
16	*10346*	*331*	*8016*	*4143*	*6945*	*25014*	*4461*	*6760*	*8887*	*10630*
	12619	470	8177	4243	7201	25332	4596	7718	9277	10762
	10736	468	8134	4255	7221	25326	4525	7003	9259	10709
	14.92	**0.43**	**0.53**	**-0.28**	**-0.28**	**0.02**	**1.54**	**9.26**	**0.19**	**0.49**
32	*10054*	*205*	*7882*	*4076*	*6795*	*24773*	*4425*	*6594*	*8718*	*10566*
	10744	463	8138	4239	7180	25307	4516	7401	9216	10756
	10561	425	8078	4248	7186	25303	4513	6888	9170	10697
	1.70	**8.21**	**0.74**	**-0.21**	**-0.08**	**0.02**	**0.07**	**6.93**	**0.50**	**0.55**
64	*9757*	*126*	*7658*	*3974*	*6586*	*24325*	*4386*	*6363*	*8514*	*10438*
	10587	136	8120	4230	7135	25276	4505	6879	9185	10754
	10402	137	8057	4239	7140	25278	4506	6838	9103	10695
	1.75	**-0.74**	**0.78**	**-0.21**	**-0.07**	**-0.01**	**-0.02**	**0.60**	**0.89**	**0.55**
128	*9440*	*126*	*7210*	*3782*	*6370*	*23477*	*4322*	*6026*	*8141*	*10182*
	10466	126	8120	4223	7087	25256	4503	6815	9075	10749
	10311	126	8057	4234	7093	25211	4505	6808	9069	10694
	1.48	**0.00**	**0.78**	**-0.26**	**-0.08**	**0.18**	**-0.04**	**0.10**	**0.07**	**0.51**
256	*8928*	*126*	*6314*	*3398*	*5986*	*21813*	*4194*	*5474*	*7501*	*9768*
	10238	126	8118	4213	7039	25209	4499	6793	8989	10564
	10166	126	8057	4221	7038	24913	4492	6780	8984	10534
	0.70	**0.00**	**0.75**	**-0.19**	**0.01**	**1.17**	**0.16**	**0.19**	**0.06**	**0.28**
512	*8139*	*126*	*4522*	*2630*	*5218*	*18771*	*3938*	*4796*	*6221*	*9236*
	10016	126	8115	4171	6933	24470	4491	6782	8870	10272
	9881	126	8057	4173	6908	24021	4486	6753	8835	10232
	1.35	**0.00**	**0.71**	**-0.05**	**0.36**	**1.83**	**0.11**	**0.43**	**0.39**	**0.39**
1024	*6744*	*126*	*1288*	*1180*	*3682*	*14032*	*3426*	*4098*	*4571*	*8571*
	9618	126	5060	1921	6709	24024	4476	6042	8674	10190
	9532	126	4157	1799	6674	23693	4470	6040	8607	10183
	0.89	**0.00**	**17.85**	**6.35**	**0.52**	**1.38**	**0.13**	**0.03**	**0.77**	**0.07**

Table 3. Empirical comparison of LRU and RLRU

Fig. 3. Percentages with which RLRU is better than LRU

Cache Size	File names and lengths									
	bigsort	j1	j2	j3	j4	j5	j6	pjoin	pq7	xds
	40167	18533	25881	38112	59744	95723	20709	41558	32989	88558
8	*11080*	*418*	*8110*	*4194*	*7064*	*25169*	*4490*	*6906*	*9046*	*10665*
	14632	494	8233	4262	7278	25412	5100	8014	9371	10768
	13451	499	8208	4265	7254	25368	5033	7205	9357	10724
	8.07	**-1.01**	**0.30**	**-0.07**	**0.33**	**0.17**	**1.31**	**10.09**	**0.15**	**0.41**
16	*10346*	*331*	*8016*	*4143*	*6945*	*25014*	*4461*	*6760*	*8887*	*10630*
	12619	470	8177	4243	7201	25332	4596	7718	9277	10762
	10862	470	8149	4220	7181	25315	4570	6986	9220	10714
	13.92	**0.00**	**0.34**	**0.54**	**0.28**	**0.07**	**0.57**	**9.48**	**0.61**	**0.45**
32	*10054*	*205*	*7882*	*4076*	*6795*	*24773*	*4425*	*6594*	*8718*	*10566*
	10744	463	8138	4239	7180	25307	4516	7401	9216	10756
	10620	426	8097	4217	7131	25299	4516	6927	9152	10703
	1.15	**7.99**	**0.50**	**0.52**	**0.68**	**0.03**	**0.00**	**6.40**	**0.69**	**0.49**
64	*9757*	*126*	*7658*	*3974*	*6586*	*24325*	*4386*	*6363*	*8514*	*10438*
	10587	136	8120	4230	7135	25276	4505	6879	9185	10754
	10521	136	8079	4199	7051	25250	4507	6895	9122	10703
	0.62	**0.00**	**0.50**	**0.73**	**1.18**	**0.10**	**-0.04**	**-0.23**	**0.69**	**0.47**
128	*9440*	*126*	*7210*	*3782*	*6370*	*23477*	*4322*	*6026*	*8141*	*10182*
	10466	126	8120	4223	7087	25256	4503	6815	9075	10749
	10422	126	8079	4199	6958	25149	4505	6836	9075	10703
	0.42	**0.00**	**0.50**	**0.57**	**1.82**	**0.42**	**-0.04**	**-0.31**	**0.00**	**0.43**
256	*8928*	*126*	*6314*	*3398*	*5986*	*21813*	*4194*	*5474*	*7501*	*9768*
	10238	126	8118	4213	7039	25209	4499	6793	8989	10564
	10226	126	8079	4189	6920	25127	4498	6783	8984	10541
	0.12	**0.00**	**0.48**	**0.57**	**1.69**	**0.33**	**0.02**	**0.15**	**0.06**	**0.22**
512	*8139*	*126*	*4522*	*2630*	*5218*	*18771*	*3938*	*4796*	*6221*	*9236*
	10016	126	8115	4171	6933	24470	4491	6782	8870	10272
	9934	126	8077	4045	6866	24409	4487	6772	8842	10262
	0.82	**0.00**	**0.47**	**3.02**	**0.97**	**0.25**	**0.09**	**0.15**	**0.32**	**0.10**
1024	*6744*	*126*	*1288*	*1180*	*3682*	*14032*	*3426*	*4098*	*4571*	*8571*
	9618	126	5060	1921	6709	24024	4476	6042	8674	10190
	9617	126	5074	1921	6723	23564	4471	6041	8658	10188
	0.01	**0.00**	**-0.28**	**0.00**	**-0.21**	**1.91**	**0.11**	**0.02**	**0.18**	**0.02**

Table 4. Empirical comparison of LRU and RLRU′

−2%　　0%　　　　　　　　　　　　　　　　　　　　　　　　　18%

Fig. 4. Percentages with which RLRU′ is better than LRU

C Omitted Proofs

Retrospective-LRU

Lemma 5. *For any request sequence, each complete phase defined by RLRU contains requests to at least $k + 1$ distinct pages.*

Proof. Consider any phase P and the page p which starts the next phase. Page p was requested in phase P, and was later evicted, also within phase P. At that point, all other pages in the cache must either be marked or have been requested since the last request to p, so every page in cache at that point has been requested in phase P. The page requested when p is evicted must be different from the k pages in cache at that point. Thus, there must be at least $k + 1$ different pages requested in phase P. □

Lemma 6. *For any sequence I of page requests, $RLRU_W(I) \leq LRU_W(I)$.*

Proof. Consider a worst permutation I_{RLRU} of I with respect to RLRU. By definition, RLRU never faults twice on the same page within any single phase of I_{RLRU}.

   Move the last, possibly incomplete, phase of I_{RLRU} to the beginning and call the resulting sequence I_{LRU}. Process the requests in this sequence phase by phase (the phases are the original RLRU phases), starting at the beginning. LRU faults on each distinct page in the first phase. Since, by Lemma 5, there are at least $k + 1$ distinct pages in each of the later phases, all of the distinct pages in a phase can be ordered so that there will be a fault by LRU on each of them. Hence, in each phase, LRU faults at least as many times as RLRU, i.e., LRU has at least as many faults on I_{LRU} as RLRU on I_{RLRU}. □

   The previous lemma established that $\mathrm{WR}_{\mathrm{LRU,RLRU}} \geq 1$. To find the exact relative worst order ratio for the two algorithms, the following technical lemma for LRU is proven. This lemma is also used extensively in the section on conservative and marking algorithms.

Lemma 1. For any sequence I of page requests, there exists a worst permutation of I for LRU with all faults appearing before all hits.

Proof. We describe how any permutation I' of I can be transformed, step by step, to a permutation I_{LRU} with all hits appearing at the end of the sequence, without decreasing the number of faults LRU will have on the sequence. Let I' consist of the requests $r_1, r_2, \ldots, r_n$, in that order.

   Consider the first hit r_i in I' with respect to LRU. We construct a new sequence I'' by moving r_i later in I'. Let p denote the page requested by r_i.

   First, we remove r_i from the sequence. If p is evicted at some point after r_{i-1} in this shorter sequence, and is not evicted at the same point in I', r_i is placed just after the first request r_j, $j > i$, causing p to be evicted (see Figure 5). Otherwise, r_i is inserted after r_n. In this case, let $j = n$.

$$I' : r_1, \ldots, r_{i-1}, r_i, r_{i+1}, \ldots, r_j, r_{j+1}, \ldots, r_n$$
$$I'' : r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_j, r_i, r_{j+1}, \ldots, r_n$$

Fig. 5. The two sequences I' and I'' in the case where p is evicted at r_j.

LRU maintains a queue of the pages in cache, and, on a fault, evicts the first page in the queue. Moving r_i within the sequence affects the position of p in the queue, but the mutual order of the other pages stays the same. Just before r_{i+1}, the cache contents are the same for both sequences. Therefore, for I'', the behavior of LRU is the same as for I' until p is evicted. Just after this eviction, p is requested by r_i in I''. Thus, just before r_{j+1}, the cache contents are again the same for both sequences, but for I'', p is at the end of queue. This means that all pages that are in cache just before r_{j+1}, except p, are evicted no later for I'' than for I'. The first request to p after the jth request may be a fault in I' and a hit in I''. On the other hand, r_i is a hit in I' and a fault in I''.

Let r_ℓ be the first request after r_i in I'', where p is either requested or evicted. After r_ℓ, the state of LRU is the same for both sequences.

By moving r_i, the number of faults among the first j requests is increased by at least one, and the total number of faults is not decreased. Thus, continuing in this way, we obtain I_{LRU} in a finite number of steps. $\qquad\square$

Theorem 1. $\text{WR}_{\text{LRU,RLRU}} = \frac{k+1}{2}$.

Proof. Since Lemma 6 shows that $\text{WR}_{\text{LRU,RLRU}} \geq 1$, for the lower bound, it is sufficient to find a family of sequences I_n with $\lim_{n\to\infty} \text{LRU}(I_n) = \infty$, where there exists a constant b such that for all I_n,

$$\text{LRU}_W(I_n) \geq \frac{k+1}{2}\text{RLRU}_W(I_n) - b.$$

Let I_n consist of n phases, where, in each phase, the first $k-1$ requests are to the $k-1$ pages $p_1, p_2, ..., p_{k-1}$, always in that order, and the last two requests are to completely new pages. LRU will fault on every page, so it will fault $n(k+1)$ times.

Regardless of the order this sequence is given in, LFD will never evict the pages $p_1, p_2, ..., p_{k-1}$ from cache, so RLRU will mark them the first time they are requested in each phase, if they have ever been requested before. Since there are never more than $k-1$ marked pages in cache, none of these pages is ever evicted in a phase in which it is marked. Thus, for each of these pages p', at most one phase is ended because of a fault on p', and the requests to the pages which only occur once cannot end phases. This gives at most $k-1$ phases, each containing at most one fault on each of the pages $p_1, p_2, \ldots, p_{k-1}$, which limits the number of faults RLRU has on these $k-1$ pages to a constant (dependent on k, but not n), so RLRU faults at most $2n + c$ times for some constant c. Asymptotically, the ratio is $\frac{k+1}{2}$.

For the upper bound, suppose there exists a sequence I, where LRU faults s times on its worst permutation, I_{LRU}, RLRU faults s' times on its worst permutation, I_{RLRU}, and $s > \frac{k+1}{2} \cdot s'$. Then, $s > \frac{k+1}{2} \cdot s''$, where s'' is the number of times RLRU faults on I_{LRU}. Assume by Lemma 1, that I_{LRU} is such that LRU faults on each request of a prefix I_1 of I_{LRU} and on no request after I_1. Then there must exist a subsequence, $J = \langle r_1, r_2, ..., r_{k+1}\rangle$, of consecutive requests in I_1, where RLRU faults at most once. Since LRU faults on every request, they must be to $k+1$ different pages. One may assume that r_1 is not the first request, since then RLRU would fault on all the requests in J. Let p be the page requested immediately before J. Clearly, p must be in RLRU's cache when it begins processing J. If r_{k+1} is not a request to p, then the fact that LRU faulted on every request in J means that J contains $k+1$ pages different from p, but at most $k-1$ of them could be in RLRU's cache when it begins processing J. Thus, RLRU must fault at least twice on the requests in J. On the other hand, if r_{k+1} is a request to p, there are exactly k requests in J which are different from p. At least one of them must cause a fault, since at most $k-1$ of them could have been in the cache when RLRU began processing J. If no others caused faults, then they must have all

been marked. In this case RLRU evicts the least recently used page in cache, which cannot be a page requested in J before this fault, so it must be a later page in J, causing a second fault. This is a contradiction. ☐

Theorem 2. The competitive ratio of RLRU is $k + 1$.

Proof. The upper bound of $k + 1$ follows since each phase of the algorithm contains requests to at least $k + 1$ different pages, and RLRU faults at most once on each page within a phase. If there are $s > k$ different pages in a phase, OPT must fault at least $s - k$ times in that phase. The worst ratio is obtained when there are exactly $k + 1$ different pages in a phase, giving a ratio of $k + 1$.

The lower bound follows from a sequence with $k + 1$ distinct pages $p_1, p_2, \ldots, p_{k+1}$, where each request is to the page not in RLRU's cache. This sequence is $\langle p_1, p_2, \ldots, p_{k+1} \rangle^2 \langle p_k, p_1, p_2, \ldots, p_{k-1}, p_{k+1} \mid p_{k-1}, p_k, p_1, p_2, \ldots, p_{k-2}, p_{k+1} \mid p_{k-2}, p_{k-1}, p_k, p_1, p_2, \ldots, p_{k-3}, p_{k+1} \mid \ldots \mid p_1, p_2, \ldots, p_{k+1} \rangle^n$, where $\mid$ marks the beginning of a new phase. The part of the sequence which is repeated n times is called a superphase and consists of k phases, the ith phase consisting of the sequence $\langle p_{k+1-i}, \ldots, p_k, p_1, \ldots, p_{k-i}, p_{k+1} \rangle$, for $1 \leq i \leq k - 1$, and $\langle p_1, p_2, \ldots, p_{k+1} \rangle$, for $i = k$. The optimal strategy is to evict page p_{k-1-i} in the ith phase of a superphase for $1 \leq i \leq k - 2$, p_k for $i = k - 1$, and p_{k-1} for $i = k$. Hence, an optimal off-line algorithm faults $k + 1$ times on the initial $2k + 1$ requests and then exactly once per phase, while RLRU faults on all $k + 1$ requests of each phase. ☐

Theorem 3. $\mathrm{WR}_{\mathrm{LRU}} = \mathrm{WR}_{\mathrm{RLRU}} = k$.

Proof. Consider any sequence I. Since no algorithm is better than OPT, on OPT's worst permutation of I, LRU will fault at least as many times as OPT, so it also will on its own worst permutation. The sequence consisting of n copies of $k + 1$ pages repeated cyclicly is a worst permutation of the underlying multi-set for both LRU and OPT. LRU faults k times for every time that OPT faults. Since the worst order ratio cannot be larger than the competitive ratio, and LRU's competitive ratio is k, $\mathrm{WR}_{\mathrm{LRU}} = k$.

Consider any sequence I. As above, on OPT's worst permutation of I, RLRU will fault at least as many times as OPT, so it also will on its own worst permutation. By Lemma 6, for any sequence I, $\mathrm{RLRU}_W(I) \leq \mathrm{LRU}_W(I)$. Thus, since $\mathrm{WR}_{\mathrm{LRU}} = k$, $\mathrm{WR}_{\mathrm{RLRU}} \leq k$. The sequence $\langle p_1, p_2, \ldots, p_{k+1} \rangle^2 \langle p_k, p_1, p_2, \ldots, p_{k-1}, p_{k+1} \mid p_{k-1}, p_k, p_1, p_2, \ldots, p_{k-2}, p_{k+1} \mid p_{k-2}, p_{k-1}, p_k, p_1, p_2, \ldots, p_{k-3}, p_{k+1} \mid \ldots \mid p_1, p_2, \ldots, p_{k+1} \rangle^n$, where $\mid$ marks the beginning of a new phase, will cause RLRU to fault every time. A worst permutation for OPT will repeat the $k + 1$ pages in a cyclic manner and OPT will fault once on every k pages, giving the ratio k. ☐

Conservative and Marking Algorithms

We first prove that LRU is best possible among conservative algorithms.

Lemma 7. $\mathrm{WR}_{\mathbb{C}, \mathrm{LRU}} \geq 1$, for any conservative paging algorithm $\mathbb{C}$.

Proof. By Lemma 1, we can consider a sequence I where all faults by LRU occur before all hits. Let I_1 denote the subsequence consisting of the faults. We prove by induction on the lengths of prefixes of I_1 that, on any request in I_1, any conservative algorithm $\mathbb{C}$ evicts the same page as LRU, and hence has as many faults on I as LRU.

For the base case, consider the first $k + 1$ requests in the sequence. Since LRU faults on each request, these $k + 1$ requests are all to different pages (ignoring the trivial case with at most k pages in I). Hence, on the $(k + 1)$st request, any algorithm must evict a page. Since $\mathbb{C}$ is conservative it evicts p_1 (if it evicted some page $p_i \neq p_1$, requesting p_i after p_{k+1} would yield a sequence with a subsequence $\langle p_2, \ldots, p_{k+1}, p_i \rangle$ with requests to only k distinct pages, but with $k + 1$ faults).

The induction step is similar to the base case. By the induction hypothesis, $\mathbb{C}$ has the same pages in cache as LRU. For each request r to some page p, the previous k requests were all to different pages different from p. Hence, $\mathbb{C}$ must evict the first of these k pages, as LRU does. □

In addition, LRU is a worst possible conservative algorithm.

Lemma 8. $WR_{LRU,\mathbb{C}} \geq 1$, *for any conservative paging algorithm* $\mathbb{C}$.

Proof. Consider any conservative algorithm $\mathbb{C}$ and any request sequence I. Divide I into phases, so that $\mathbb{C}$ faults exactly $k + 1$ times per phase, starting the next phase with a fault (the last phase may have fewer than $k + 1$ faults). Since $\mathbb{C}$ is conservative, each phase, except possibly the last, contains at least $k+1$ distinct pages. These can be reordered, phase by phase, so that LRU faults once on each distinct page within a phase. Thus, with this new permutation LRU faults at least as many times as $\mathbb{C}$ on I, except possibly in the last phase. Since $\mathbb{C}$ faults at least once in the last phase, LRU faults as many times on the new permutation of I as $\mathbb{C}$ on I, except for at most k faults. □

Thus, all conservative algorithms are equivalent under the relative worst order ratio:

Theorem 4. For any pair of conservative algorithms $\mathbb{C}_1$ and $\mathbb{C}_2$, $WR_{\mathbb{C}_1,\mathbb{C}_2} = 1$.

Proof. By Lemmas 7 and 8, $WR_{\mathbb{C}_1,LRU} \geq 1$ and $WR_{LRU,\mathbb{C}_2} \geq 1$. Since the relative worst order ratio is easily seen to be a transitive measure [7], this means that $WR_{\mathbb{C}_1,\mathbb{C}_2} \geq 1$. Similarly, $WR_{\mathbb{C}_1,LRU} \leq 1$ and $WR_{LRU,\mathbb{C}_2} \leq 1$, which implies $WR_{\mathbb{C}_1,\mathbb{C}_2} \leq 1$. □

The marking algorithm Flush-When-Full is strictly worse than any conservative algorithm:

Lemma 9. *For any conservative algorithm* $\mathbb{C}$, $WR_{FWF,\mathbb{C}} \geq \frac{2k}{k+1}$.

Proof. By Theorem 4 and transitivity, it is sufficient to show that $WR_{FWF,LRU} \geq \frac{2k}{k+1}$. Consider any sequence I. If LRU faults on request r in I to page p, then p was not among the last k different pages that were requested. Thus, p could not be in FWF's cache when request r occurs and FWF will also fault. Hence, on any sequence, FWF will fault at least as many times on its worst permutation as LRU will on its. This shows that $WR_{FWF,LRU} \geq 1$.

It is now sufficient to find a family of sequences I_n with $\lim_{n\to\infty} LRU(I_n) = \infty$, where there exists a constant b such that for all I_n,

$$FWF_W(I_n) \geq \frac{2k}{k+1} LRU_W(I_n) - b.$$

Let $I_n = \langle p_1, p_2, \ldots, p_k, p_{k+1}, p_k, \ldots, p_2 \rangle^n$. FWF will fault on every page, so it will fault $n(2k)$ times.

By Lemma 1, a worst permutation for LRU consists of some permutation of the $k + 1$ pages, with p_1 and p_{k+1} at the end, repeated n times, followed by the first $k - 1$ pages in the permutation repeated n times, i.e., $\langle p_2, p_3, \ldots, p_{k+1}, p_1 \rangle^n \langle p_2, p_3, \ldots, p_k \rangle^n$ is a worst permutation of I_n with respect to LRU. LRU will fault $n(k+1) + k - 1$ times. Asymptotically, the ratio is $\frac{2k}{k+1}$. □

Combining the previous lemma and transitivity with the following lemma, we find that Flush-When-Full is worst possible among marking algorithms.

Lemma 10. *For any marking algorithm $\mathbb{M}$ and any request sequence I, $\mathbb{M}_W(I) \leq \frac{2k}{k+1} LRU_W(I) + k$.*

Proof. For any sequence with n complete phases, $\mathbb{M}$ faults at most kn times. Since for any request sequence, every pair of consecutive phases contains requests to at least $k+1$ distinct pages, for the first $2 \cdot \lfloor \frac{n}{2} \rfloor$ complete phases, there is a permutation such that LRU faults at least $(k+1)\lfloor \frac{n}{2} \rfloor$ times. If the remaining requests consist of a complete phase, plus a partial phase, then LRU will also fault on $k+1$ of those requests if given in the correct order. Thus, the additive constant is bounded by the number of faults $\mathbb{M}$ makes on the last partial or complete phase and is thus at most k. $\square$

Combining the above two lemmas, Theorem 4, and the the transitivity of the relative worst order ratio gives:

Theorem 5. For any conservative algorithm $\mathbb{C}$, $\mathrm{WR}_{\mathrm{FWF},\mathbb{C}} = \frac{2k}{k+1}$.

LRU is not far from being a best possible marking algorithm:

Lemma 11. *For any marking algorithm $\mathbb{M}$ and any sequence I of page requests,*

$$LRU_W(I) \leq \frac{k+1}{k} \mathbb{M}_W(I).$$

Proof. Consider any sequence I of requests. By Lemma 1, there is a worst permutation I_{LRU} such that LRU faults on each request of a prefix I_1 of I_{LRU} and on no request after I_1. Partition I_1 into consecutive subsequences, each consisting of exactly $k+1$ requests (the last subsequence may contain fewer). Since LRU faults on all requests in I_1, each subsequence, except possibly the last, must contain $k+1$ distinct pages. Hence, for each subsequence with pages $p_1, p_2, \ldots, p_{k+1}$, an adversary can create a marking phase, by choosing k of the pages $p_1, p_2, \ldots, p_{k+1}$, such that the marking algorithm faults on all k pages. This is easily seen in the following way. Pages requested within a phase stay in cache throughout the phase. Therefore, when x of the pages $p_1, p_2, \ldots, p_{k+1}$ have been requested, the remaining $k+1-x$ pages cannot all be in the cache. $\square$

This, combined with Theorem 4 and the transitivity of the relative worst order ratio, gives the following:

Theorem 6. For any marking algorithm $\mathbb{M}$ and any conservative algorithm $\mathbb{C}$ with $\mathrm{WR}_{\mathbb{C},\mathbb{M}}$ defined,

$$\mathrm{WR}_{\mathbb{C},\mathbb{M}} \leq \frac{k+1}{k}.$$

Look-Ahead

One can see that $\mathrm{LRU}(\ell)$ is at least as good as LRU on any sequence by noting that $\mathrm{LRU}(\ell)$ is conservative.

Lemma 12. *$LRU(\ell)$ is a conservative algorithm.*

Proof. Let I be a request sequence, and assume that there is an interval, I', in I, containing only k distinct pages, on which LRU(ℓ) faults at least $k+1$ times. Then it must fault on some page, p, twice in I'. Between these two faults, say at request r, page p must be evicted. First assume that $\ell < k$. At this point, p is the least recently used page which is not among the next ℓ. Clearly the second request causing a fault on p must be beyond these next ℓ. So the other $k-1$ pages in cache, when request r occurs, must all have been requested between the two faults on p. In addition, the request r cannot be for p or any of the other pages in cache at that time. Thus, there must be at least $k+1$ distinct pages in I', giving a contradiction. Now assume that $\ell \geq k$. If p is not among the next ℓ requests when r occurs, the previous argument holds, so assume that it is. In this case p must have been the page in cache which was requested furthest in the future, so the other $k-1$ pages are requested between request r and the second fault on p. Again, counting the request r and p, there must be at least $k+1$ distinct pages in I', which is a contradiction. Thus, LRU(ℓ) is conservative. $\qquad\square$

Observe that Lemma 8 holds for algorithms using look-ahead, though Lemma 7 does not.

Theorem 7. $\mathrm{WR}_{\mathrm{LRU,LRU}(\ell)} = \min\{k, \ell+1\}$

Proof. Since the previous lemma combined with Lemma 8 show that $\mathrm{WR}_{\mathrm{LRU,LRU}(\ell)} \geq 1$, to prove the lower bound, it is sufficient to find a family of sequences I_n with $\lim_{n\to\infty} \mathrm{LRU}(I_n) = \infty$, where there exists a constant b such that for all I_n,

$$\mathrm{LRU}_W(I_n) \geq \min\{k, \ell+1\}\mathrm{LRU}(\ell)_W(I_n) - b.$$

Let I_n consist of n phases, each containing the pages $p_1, p_2, ..., p_k, p_{k+1}$, in that order. LRU will fault $n(k+1)$ times. However, if $\ell \leq k-1$, after the first k faults, LRU(ℓ) never faults on any of the next ℓ pages after a fault. Thus, regardless of the order, LRU(ℓ) faults on at most $k + \lfloor \frac{n(k+1)-k}{\ell+1} \rfloor$ pages. Asymptotically, this gives a ratio of $\ell+1$. If $\ell \geq k$, then LRU(ℓ) faults on at most one out of every k pages.

For the upper bound, suppose there exists a sequence I, where LRU faults s times on its worst permutation, I_{LRU}, LRU(ℓ) faults s' times on its worst permutation, $I_{\mathrm{LRU}(\ell)}$, and $s > \min(k, \ell+1) \cdot s'$. Then, $s > \min(k, \ell+1) \cdot s''$, where s'' is the number of times LRU(ℓ) faults on I_{LRU}. One cannot have $\ell \geq k$, since then LRU(ℓ) faults fewer times than OPT. So suppose $\ell < k$, and assume by Lemma 1, that I_{LRU} is such that LRU faults on each request of a prefix I_1 of I_{LRU} and on no request after I_1. Then there must exist a request r in I_{LRU} where LRU(ℓ) faults, but it does not fault on any of the next $\ell+1$ requests, all of which are in I_1. The last of these $\ell+1$ requests caused LRU to fault, so it was not among the last k distinct requests at that point. Since $l < k$, it was not in any of the requests in the look-ahead when LRU(ℓ) processed request r, and all of the pages in the look-ahead were in cache then since LRU(ℓ) did not fault on any of them. Hence, this $\ell+1$st page was evicted by LRU(ℓ) when r was requested, and there must have been a fault the next time it was requested after that, giving a contradiction. $\qquad\square$

MARK

Lemma 3. For any sequence I of page requests, there exists a worst permutation I_{MARK} of I with respect to MARK, such that all k-phases, except possibly the last, have the following properties:

1. The first page requested in a k-phase did not appear in the previous k-phase.

2. There are exactly k requests, all to distinct pages.

Proof. Consider a worst permutation I_{MARK} of I for MARK and consider its k-phase partition. The first property follows by the definition of a k-phase partition. Within a phase, the first occurrence of each page is the only one MARK has any chance of faulting on. Thus, moving extra occurrences of a page within a phase to the end of the sequence will never decrease the probability of MARK faulting on any page. After this has been completed, each phase (except possibly the last) consists of exactly k requests, all to distinct pages. $\square$

Claim 1. For all integers s, i, k, such that $1 \leq s \leq k$, $i \geq 0$, and $s + i \leq k$,

$$s^2(1 + H_k - H_s) + (s + i)(1 + H_k - H_{s+i}) \leq s(k + 1) + i.$$

Proof. One can check by hand that the claim holds for all $k \leq 5$.
For $s > \frac{k}{2}$, the claim holds by the following calculations.

$$
\begin{aligned}
s^2(1 + H_k - H_s) + (s + i)(1 + H_k - H_{s+i}) &= s(s+1)\left(1 + \sum_{j=s+1}^{k} \frac{1}{j}\right) + i\left(1 + \sum_{j=s+i+1}^{k} \frac{1}{j}\right) - s\left(\sum_{j=s+1}^{s+i} \frac{1}{j}\right) \\
&\leq s(s+1)\left(1 + \frac{k-s}{s+1}\right) + i\left(1 + \frac{k-s-i}{s+i+1}\right) - s\frac{i}{s+1} \\
&= s(k+1) + i + \left(\frac{i(k-s-i)}{s+i+1} - \frac{si}{s+1}\right) \\
&< s(k+1) + i + \left(\frac{i\left(\frac{k}{2} - i\right)}{s+i+1} - \frac{\frac{k}{2}i}{s+1}\right), \text{ for } s > \frac{k}{2} \\
&\leq s(k+1) + i, \text{ since } i \geq 0
\end{aligned}
$$

Note that approximating by integrals, one gets that $H_x - H_y \leq \ln(x) - \ln(y)$ for $x > y$. Thus, $s^2(1 + H_k - H_s) + (s+i)(1 + H_k - H_{s+i}) \leq s^2(1 + \ln(k) - \ln(s)) + (s+i)(1 + \ln(k) - \ln(s+i))$, so it is sufficient to prove that

$$f(s, i, k) = s^2(1 + \ln(k) - \ln(s)) + (s + i)(1 + \ln(k) - \ln(s + i)) - s(k + 1) - i \leq 0$$

for $k \geq 6$ and $s \leq k/2$.
Taking the derivative of $f(s, i, k)$ with respect to i gives $f'_i(s, i, k) = \ln(k) - \ln(s + i) - 1$, which is zero at $i = \frac{k-se}{e}$, positive for smaller i and negative for larger. Thus, for fixed s and k, $f(s, i, k)$ has its maximum at $i = \frac{k-se}{e}$.
Assume now that $\frac{k}{e} \leq s \leq \frac{k}{2}$. In this case, $i = \frac{k-se}{e}$ is negative, and hence outside the specified range for i. Since $f(s, i, k)$ is decreasing for $i > \frac{k-se}{e}$, $f(s, i, k)$ is maximum at its smallest allowable value, $i = 0$. Hence,

$$f(s, i, k) \leq f(s, 0, k) = s(s + 1)(1 + \ln(k) - \ln(s)) - s(k + 1).$$

The derivative of this with respect to s is $-s - k - 2 + (1 + \ln(k) - \ln(s))(2s + 1)$ which is zero where $(s + k + 2)/(2s + 1) = (1 + \ln(k) - \ln(s))$. At this point, $f(s, 0, k) = s(s + 1)(s + k + 2)/(2s + 1) - sk - s = \frac{s^3 + s^2 + s - s^2k}{2s+1}$. This is equal to zero, where $s^2 + s + 1 - sk = 0$, which has no solution in the range $\frac{k}{e} \leq s \leq \frac{k}{2}$ for $k \geq 4$. Looking at the endpoints of this range, we

see that $f(\frac{k}{e}, 0, k)$ is negative for $k \geq 4$, and $f(\frac{k}{2}, 0, k)$ is negative for $k \geq 6$. Thus, $f(s, i, k)$ is negative for $\frac{k}{e} \leq s \leq \frac{k}{2}$ and $k \geq 6$.

Finally, consider $s \leq \frac{k}{e}$. In this range, the maximum value of f is

$$f\left(s, \frac{k - se}{e}, k\right) = s^2(1 + \ln(k) - \ln(s)) + 2s + \frac{k - se}{e} - s(k + 1).$$

Taking the derivative with respect to s gives $s - k + 2s(\ln(k) - \ln(s))$, which is negative for s small enough and then positive, so the function has a local minimum where the derivative is zero. Thus, the maximum values are at the endpoints. For $s = 1$, one gets that $f(1, \frac{k}{e} - 1, k) = 1 + \ln(k) + \frac{k}{e} - k$, which is negative for $k \geq 4$. For $s = \frac{k}{e}$, one gets that $f(\frac{k}{e}, 0, k) = (\frac{2}{e} - 1)sk + s$, which is negative for $k \geq 4$.

Thus, for $1 \leq s \leq k$, $i \geq 0$, and $s + i \leq k$, $\quad s^2(1 + H_k - H_s) + (s + i)(1 + H_k - H_{s+i}) \leq s(k + 1) + i$. $\qquad \square$