Almost Tight Bounds for Reordering Buffer Management *

Anna Adamaszek
A.M.Adamaszek@warwick.ac.uk

Artur Czumaj
A.Czumaj@warwick.ac.uk

Matthias Englert
M.Englert@warwick.ac.uk

Harald Räcke
H.Raecke@warwick.ac.uk

Department of Computer Science and
Centre for Discrete Mathematics and its Applications (DIMAP)
University of Warwick, Coventry, UK

ABSTRACT

We give almost tight bounds for the online reordering buffer management problem on the uniform metric. Specifically, we present the first non-trivial lower bounds for this problem by showing that deterministic online algorithms have a competitive ratio of at least $\Omega(\sqrt{\log k / \log \log k})$ and randomized online algorithms have a competitive ratio of at least $\Omega(\log \log k)$, where k denotes the size of the buffer.

We complement this by presenting a deterministic online algorithm for the reordering buffer management problem that obtains a competitive ratio of $O(\sqrt{\log k})$, almost matching the lower bound. This improves upon an algorithm by Avigdor-Elgrabli and Rabani (SODA 2010) that achieves a competitive ratio of $O(\log k / \log \log k)$.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Sequencing and scheduling*

General Terms

Theory, Algorithms

Keywords

Online Algorithms, Reordering Buffer, Sorting Buffer

1. INTRODUCTION

In the reordering buffer management problem a stream of colored items arrives at a service station and has to be processed. The cost for servicing the items heavily depends on the processing order: servicing an item with color c, when the most recently serviced item had color $c' \neq c$, incurs a context switching cost w_c.

In order to reduce the total processing cost, the servicing station is equipped with a reordering buffer able to store k items. This buffer can be used to reorder the input sequence in a restricted fashion to construct an output sequence with lower processing cost. At each point in time, the buffer contains the first k items of the input sequence that have not yet been processed. A scheduling strategy has to decide which item to service next. Upon its decision, the corresponding item is removed from the buffer and serviced, while the next item from the input sequence takes its place in the buffer.

This simple and versatile framework has many important applications in areas like production engineering, computer graphics, storage systems, and information retrieval, among others [3, 5, 10, 13, 14]. We give two examples.

In the paint shop of a car manufacturing plant, switching colors between two consecutive cars induces non-neglible cleaning and set-up costs. Therefore, paint shops are preceded by a reordering buffer (see [10]) to reorder the stream of incoming cars into a stream with a lower number of color-changes.

In a 3D graphic rendering engine [13], a change in attributes between two consecutively rendered polygons slows down the GPU, as, for instance, the shader program needs to be replaced. A reordering buffer can be included between application and graphics hardware in order to reduce such state changes.

In this paper we focus on the *online version* of the reordering buffer problem, in which when the buffer becomes full, one has to decide which item to service next, without knowing the rest of the input sequence. The cost of an online algorithm is compared to the cost of an optimal offline strategy that knows all items in the input sequence in advance and may use the buffer of size k to reorder these items. The worst case ratio between the cost of the online algorithm and the cost of an optimal offline algorithm is called the *competitive ratio*.

1.1 Related work

The reordering buffer problem was introduced by Räcke et al. [14], who developed an $O(\log^2 k)$-competitive online algorithm for the version with *uniform costs* where $w_c = 1$ for all colors c. Englert and Westermann [8] improved the competitive ratio to $O(\log k)$, and their algorithm is also able to handle non-uniform costs with the same bound. Their

proof works in two steps. First, it is shown that an online algorithm with a buffer of size k is constant competitive w.r.t. an optimal offline algorithm with a buffer of size $k/4$. Then, it is shown that an optimum algorithm with a buffer of size $k/4$ only looses a logarithmic factor compared to an optimum algorithm with a buffer of size k.

It was shown in [1] that with this proof technique it is not possible to derive online algorithms with a competitive ratio $o(\log k)$, by presenting an input sequence where the gap between an optimum algorithm with a buffer of size $k/4$ and an optimum algorithm with buffer size k is $\Omega(\log k)$. Nevertheless, Avigdor-Elgrabli and Rabani [3] were able to go beyond the logarithmic threshold by presenting an online algorithm with competitive ratio $O(\log k/\log\log k)$ using linear programming based techniques. This is currently the best known algorithm.

For the offline problem it was recently shown by Chan et al. [6] and independently by Asahiro, Kawahara, and Miyano [2] that the problem is NP-hard even for uniform costs. To the best of our knowledge, no approximation algorithms that achieve better bounds than the best online algorithm are known.

There also exists a more general version of the problem. Instead of letting the context switching cost for switching from an item with color c' to an item with color c only depend on c, it is sometimes desirable to let it depend on c' and c. Khandekar and Pandit [11], and Gamzu and Segev [9] study the problem where the colors correspond to points in a line-metric. Colors c' and c are integer points on the line and the cost for switching from c' to c is $|c'-c|$. This version of the problem is motivated by disc scheduling. Khandekar and Pandit [11] give a randomized $O(\log^2 n)$-competitive online algorithm for n uniformly spaced points on a line, and a constant factor offline approximation in quasi-polynomial time. Gamzu and Segev [9] improve the first result to $O(\log n)$. They also shows a lower bound of 2.1547 on the competitive ratio of deterministic online algorithms on the line. This is the only non-trivial lower bound that exists for any variant (i.e. metric) of the problem.

Englert et al. [7] consider the more general version where colors correspond to arbitrary points in a metric space (C,d) and the cost for switching from color c' to color c is the distance $d(c',c)$ between the corresponding points in the metric space. They obtain a competitive ratio of $O(\log^2 k \log|C|)$.

Research has also been done on the maximization version of the problem, where the cost-measure is the number of color-changes that the output sequence saved over the unordered input sequence. For this version there exist constant factor approximation algorithms due to Kohrt and Pruhs [12] and Bar-Yehuda and Laserson [4].

1.2 Our results

In this work, we improve the best known upper bound, as well as the best known lower bounds. We start by presenting the first non-trivial lower bound on the competitive ratio of online algorithms for the problem. We show in Theorem 2.6 that any deterministic online algorithm for the reordering buffer management problem has a competitive ratio of at least $\Omega(\sqrt{\log k/\log\log k})$ even in the uniform case. For randomized algorithms we are able to construct a lower bound of $\Omega(\log\log k)$ (Theorem 2.9). Before this work, no lower bounds were known.

We complement the lower bounds with a deterministic online algorithm whose competitive ratio nearly matches it. We present a deterministic online algorithm that obtains a competitive ratio of $O(\sqrt{\log k})$ for the non-uniform case in which the ratio between the smallest and the largest weight of a color is polynomially bounded in k. This improves upon the result of Avigdor-Elgrabli and Rabani [3] who obtained a competitive ratio of $O(\log k/\log\log k)$.

All previous results [3, 8, 14] for the reordering buffer management problem used very similar algorithms with only subtle differences between them. The differences between the results were mostly based on the analysis. In contrast, our new result relies on an important modification in the algorithm. In addition to techniques similar to those used in [3, 8, 14], our algorithm also relies on classifying colors according to the number of items of the color in the buffer. Then, the algorithm tries to evict items of a color class that currently occupy a large fraction of the buffer. This algorithmic ingredient plays a crucial role in reducing the competitive ratio to $O(\sqrt{\log k})$ in our analysis.

Finally, it is worth mentioning that we are not aware of any candidate for even a randomized algorithm that possibly could give a competitive ratio better than $O(\sqrt{\log k})$. It is possible to use the techniques developed in this paper to show better lower bounds for specific algorithms as, e.g., for the algorithm proposed in [13], which, in the case when a color-change is required, chooses a random item in the buffer and switches to the corresponding color. This shows that, e.g., this algorithm does not have a competitive ratio better than $\Omega(\sqrt{\log k})$.

2. LOWER BOUNDS

In this section we give lower bounds on the competitive ratio of online algorithms for the reordering buffer problem. We do this by carefully constructing an input sequence σ for which any online algorithm ONL exhibits a large cost while the optimum algorithm OPT can process σ with a significantly lower number of color-changes.

2.1 Preliminaries

We first describe the general scheme for constructing σ. For this we introduce quite a few parameters whose precise values will be fixed later. For simplicity of notation we assume that k is sufficiently large, and chosen in such a way that no rounding issues occur.

2.1.1 The general scheme for constructing σ

The input sequence σ that we construct has the property that an optimum algorithm OPT′ with a buffer of size $(1+\alpha)k$ can process σ in such a way that the number of color-changes is equal to the total number of different colors contained in σ. This means OPT′ is truly optimal for the sequence, and even increasing the size of the buffer further cannot reduce the cost for processing σ.

To specify σ further, we will view the buffer of OPT′ as partitioned into d classes $C_1, \ldots, C_d$. Each class can store $(1+\alpha)k/d$ elements (a $1/d$-fraction of the total buffer size of OPT′). We further partition the buffer of a class C_i into N_i slots, where each slot can store $s_i = \frac{(1+\alpha)k}{dN_i}$ elements (a $1/N_i$-fraction of the buffer space of class C_i). The number of slots in a class will be decreasing with the number of the class, i.e., for $i < j$ we have $N_i > N_j$.

The main property of the input sequence σ will be as follows:

> Whenever OPT′ has to make a color-change while processing σ, the elements in its buffer that share a particular color fill one complete slot.

This means, e.g., that if n_c denotes the number of elements with color c at a certain time (i.e., right before a color-change by OPT′), then $n_c = s_i$ for some $i \in \{1, \ldots, d\}$. Because of this property it makes sense to refer to the *color of a slot* as the color that every element in the slot has (at a specified time).

We obtain the above property by constructing σ as follows: The initial $(1 + \alpha)k$ elements of σ fill up the buffer of OPT′. They are chosen such that the invariant is satisfied right before OPT′ outputs its first element. The exact order in which these first $(1 + \alpha)k$ elements appear in σ is not important but we assume that elements among them that share the same color appear consecutively.

Further elements in σ are chosen in rounds in the following way. For a round we choose a slot z_i, $i \in \{1, \ldots, d\}$ from every class.

Then we first add $k + 1$ elements of the color of z_d to σ, which will force any algorithm to switch to this color at this point. Then, for every i, starting from $d - 1$ down to 1, we add $s_{i+1} - s_i$ elements of the color of slot z_i to σ. Finally, we add s_1 elements of a completely new color to σ. This finishes the round.

The algorithm OPT′ works as follows. In the beginning of a round it switches to the color of slot z_d. This frees up space s_d in the buffer, and, hence, OPT′ can hold all further elements appearing in the round without requiring any further color-changes.

In order to maintain the main property, we do the following. For $i \leq d - 1$, all s_i elements in slot z_i plus the $s_{i+1} - s_i$ elements with the same color arriving in the round are moved to slot z_{i+1} (and completely fill this slot).[1] We say that we *promote slot z_i in class C_i*. Finally, slot z_1 holds the s_1 elements of the new color.

This finishes the description of the construction of σ up to the selection of the d slots for each round. Note that regardless of how we choose these slots it follows from the above discussion that the cost of OPT′ for processing σ is equal to the number of different colors in the sequence.

2.1.2 A sketch of the analysis

From the above description it is clear that the cost of OPT′ for processing σ is equal to the number of different colors in the sequence.

However, the online algorithm ONL and the optimum algorithm OPT only have a buffer of size k. Hence, at any time, these algorithms have already removed at least αk elements that are still held by OPT′. Suppose for the time being that these algorithms only remove whole slots (remember that we can simply view a slot as a set of elements that share the same color).

Clearly, if e.g. ONL removed all elements of a slot z_i, and this slot is promoted, then ONL will have an additional color-change that OPT′ doesn't have. Now, if our aim is to

maximize the ratio between the cost of ONL and the cost of OPT′ it turns into a caching problem, where in each round the adversary has to pick a slot from every class in such a way that she "hits" many slots that ONL has removed.

For making this idea work we need to show

- that it is more or less optimal for ONL to remove whole slots,

- that an adversary always can promote a large number of slots that have been removed by ONL, and

- that OPT can handle the resulting input sequence fairly well (as we are not interested in the ratio between OPT′ and ONL but in the ratio between OPT and ONL).

2.1.3 The caching framework

In this section we make a formal connection of our problem to the caching-related problem hinted at in the above sketch.

We say that an algorithm ALG *cleared* slot z before round r if right before the time ALG reads the first element of round r, there are no elements from slot z in ALG's buffer. This implies that ALG does not hold any elements of the color of slot z. Define the *cost* $\text{cost}^r_{\text{ALG}}$ of an algorithm ALG in round r as the number of slots promoted in round r that are cleared by ALG (i.e. $\text{cost}^r_{\text{ALG}}$ corresponds to the number of "cache misses" in round r).

A lower bound for ONL

The following lemma gives a lower bound on the cost of the online algorithm in terms of $\text{cost}^r_{\text{ONL}}$.

LEMMA 2.1. *The total cost of* ONL *on the generated input sequence σ is at least $\sum_{i=1}^{d} N_i + \sum_r \text{cost}^r_{\text{ONL}}$.*

PROOF. First observe that we can compute the cost of an algorithm by increasing its cost by 1 whenever an element of a color arrives whose color is different from the last color that was appended to the output sequence (the *active output color*) and also different from all colors present in the buffer.

Initially, the first $(1 + \alpha)k$ elements of the input sequence consist of $\sum_{i=1}^{d} N_i$ different colors, each of them contributing 1 to the cost of any algorithm.

Then, consider a slot z in class C_i for $i \leq d - 1$ that is cleared by ONL and promoted in round r. Since ONL cleared z before round r, ONL does not store any elements of the color of slot z in its buffer at the beginning of the round. On the other hand, z is promoted, which means that elements of the color of z appear in σ in round r.

The first $k + 1$ elements of round r belong to the color c of a slot in class C_d. After they arrived, ONL's active output color is c. Since ONL cannot switch to a color not currently stored in the buffer and z was cleared by ONL before round r, ONL's active output color has to be different from the color of z at the time the first element of the color of z appears in σ in round r.

Therefore, each such slot z contributes 1 to the total cost of ONL and there are at least $\text{cost}^r_{\text{ONL}} - 1$ such slots, where the -1 accounts for the slot in class C_d. The cost of ONL is further increased by one in every round r, because of the single element of a completely new color appearing in σ.

By summing over all rounds and taking the first $(1 + \alpha)k$ elements into account, it follows that the total cost of ONL is bounded from below by $\sum_{i=1}^{d} N_i + \sum_r \text{cost}^r_{\text{ONL}}$. $\qquad\square$

---

[1] Observe that the notion of a slot has only been introduced for illustration. Since, it is irrelevant where in the buffer something is stored it is also possible to simply view a slot as a set of elements (all with the same color and currently stored by OPT′).

A lower bound for OPT

In order to give an analogous upper bound on the cost of OPT, we will present specific (offline) algorithms and analyze their cost. In order to do this analysis in a round-based manner, we further restrict these offline algorithms in the following way. We require that right before a new round, the algorithm has at most $k - s_d$ elements stored in the buffer. To be more specific, we consider algorithms that process a round as follows.

1. Right before the first element of the round appears in σ, the number of elements stored in the buffer is reduced to less than $k - s_d$. This is done by selecting $\alpha N_i + 1$ slots from each class C_i and removing all of the elements in these slots from the buffer.

2. Let $z_1, \ldots, z_d$ denote the slots that are promoted during the round. The algorithm outputs every element of the color of slot z_d. This includes the first $k + 1$ elements arriving in the round and also elements of the same color that may be stored in the buffer.

3. Finally, the algorithm stores all other elements arriving during the round in the buffer. This is possible since the number of these elements is s_d.

In order to satisfy the buffer constraint for the first $(1 + \alpha)k$ elements of the input sequence, we assume that the algorithm immediately outputs elements from slots that get cleared from the buffer before the first round.

LEMMA 2.2. *Given any offline algorithm* OFF *with the property above, the total cost of* OPT *on the generated input sequence σ is at most $\sum_{i=1}^d N_i + \sum_r (\text{cost}_{\text{OFF}}^r + 1)$.*

PROOF. The proof is similar to the proof of Lemma 2.1. Again, we observe that we can compute the cost of an algorithm by increasing its cost by 1 whenever an element of a color arrives whose color is different from the last color that was appended to the output sequence and also different from all colors present in the buffer.

Initially, the first $(1 + \alpha)k$ elements of the input sequence consist of $\sum_{i=1}^d N_i$ different colors, each of them contributing 1 to the cost of any algorithm. OFF's cost also increases by 1 in each round due to the single element of a completely new color that appears in every round. The sum of all these cost is $\sum_{i=1}^d N_i + \sum_r 1$.

All further increases of OFF's cost are caused by a sequence of elements arriving in some round r that have a color that is not currently present in OFF's buffer but present in the buffer of OPT$'$. Such a sequence of elements corresponds to the promotion of a slot z_i, where at the time of the promotion OFF does not store any elements of the slot in the buffer. This implies that OFF does not store any elements of the slot at the time the first element of the round is read from the input. To see this for a slot in class C_i, with $i < d$, observe that OFF does not remove any elements from class C_i from the buffer until the end of a round. Therefore, all the elements that are not in the buffer at some point in time, were already missing from the buffer when the round started. For a slot from class C_d, the statement is also true, since this slot is the first one to be promoted in the round.

In other words, each such increase of OFF's cost is caused by a promotion of a cleared slot in some round r and therefore also contributes 1 to cost$_{\text{OFF}}^r$. Hence, $\sum_r \text{cost}_{\text{OFF}}^r$ is an upper

bound on these increases to the cost of OFF. Summing up everything and observing that the cost of OPT is bounded by the cost of OFF completes the proof. □

2.1.4 Choosing parameters

For the remainder we fix the number of classes as $d := \log k / (2 \log \log k)$ and the size of a slot in class C_i as $s_i := \log^{i-1} k$. The parameter α will be chosen differently depending on whether we want to derive lower bounds for deterministic or for randomized online algorithms. If we deal with deterministic algorithms we choose $\alpha := \sqrt{\log \log k / \log k}$, otherwise we set $\alpha := (\log \log k)^2 / \log k$.

2.1.5 An important lemma

We now prove a lemma that shows that in the beginning of a round the online algorithm has many *cleared slots* and that these lie in different classes (so that the adversary can choose many of them). This lemma forms the basis of our analysis.

We first require a technical claim that essentially states that for our specific choice of s_i's the online algorithm either has a slot cleared or has stored nearly all elements of the slot.

CLAIM 2.3. *For a round r and a slot z in class C_i at least one of the following is true:*

(a) ONL *cleared slot z before round r, or*

(b) *The color of slot z is equal to the color of the element that* ONL *appended to the output sequence, right before reading the first element of round r,*

(c) ONL *holds at least $\log^{i-1} k - \log^{i-2} k$ of the $\log^{i-1} k$ elements* OPT$'$ *stores in slot z, right before* ONL *reads the first element of round r.*

PROOF. Consider a slot z in class i. There are only two possible reasons that z is not cleared (i.e., we are not in Case a), but ONL does not store all $\log^{i-1} k$ elements of z. Either z is the active output color, which means that ONL is in the process of removing elements from z (Case b), or some elements of z have been previously removed by ONL.

In the latter case some elements have arrived after the removal, as otherwise z would be cleared. However, the last sequence of elements with the color of slot z was the sequence of length $\log^{i-1} k - \log^{i-2} k$ that filled the slot z in the buffer of OPT$'$. All these elements must still be in ONL's buffer. This means we are in Case c. This proves the claim. □

Using the claim, we can now prove the desired bounds on the number of slots cleared by ONL.

LEMMA 2.4. *Let ℓ_i be the number of slots from class C_i cleared by* ONL *before round r. The following holds:*

(a) $\sum_{i=1}^d \ell_i \log^{i-1} k \geq \frac{\alpha k}{2}$.

(b) *At least $\alpha d / 4$ of the ℓ_i's are not 0. In other words, at least $\alpha d / 4$ different classes contain at least one cleared slot.*

PROOF. At the beginning of a round there must exist at least αk elements that ONL has already removed from its buffer while they are still held by OPT$'$.

Due to Claim 2.3, every slot of class C_i that is not cleared by ONL before round r and whose color is not ONL's active

output color, contains at least $\log^{i-1} k - \log^{i-2} k$ elements. Hence, the number of elements that are held by OPT' but not by ONL is at most $\log^{d-1} k + \sum_{i=1}^{d}(N_i - \ell_i)\log^{i-2} k + \sum_{i=1}^{d}\ell_i \log^{i-1} k$, where the first term accounts for the active output color of ONL and the second term accounts for the possible excess elements of OPT' in slots that are not cleared. This, however, has to be at least αk. We get

$$\alpha k \leq \log^{d-1} k + \sum_{i=1}^{d}(N_i - \ell_i)\log^{i-2} k + \sum_{i=1}^{d}\ell_i \log^{i-1} k$$

$$\leq \frac{k}{\log k} + \frac{(1+\alpha)k}{\log k} + \sum_{i=1}^{d}\ell_i \log^{i-1} k$$

$$\leq \frac{\alpha k}{2} + \sum_{i=1}^{d}\ell_i \log^{i-1} k \ ,$$

where the last step follows since $(2+\alpha)k/\log k \leq \alpha k/2$ for sufficiently large k. This implies the first claim.

For the second claim assume that less than $\alpha d/4$ of the ℓ_i's are greater than 0. Then we obtain $\sum_{i=1}^{d}\ell_i \log^{i-1} k < (1+\alpha)k/d \cdot \alpha d/4 = \alpha(1+\alpha)k/4 \leq \alpha k/2$, which is a contradiction to the first claim. $\square$

In the following sections we describe how to choose the slots to be promoted in a round in such a way that for ONL many cleared slots are promoted while for OPT this happens very rarely. This choice depends on whether we want to derive lower bounds for deterministic or for randomized algorithms.

2.2 Lower bound for deterministic algorithms

In this section we present a lower bound of $\Omega(\sqrt{\log k/\log\log k})$ on the competitive ratio of any deterministic online algorithm for the reordering buffer problem. For this section, we define α to be $\sqrt{\log\log k/\log k}$.

For every class C_i, we choose a slot for promotion as follows:

> If in class C_i there exists a slot cleared by ONL, we choose an arbitrary such slot to be promoted. Otherwise, we promote the first slot of class C_i.

We present a randomized algorithm RND that processes σ with a buffer of size k and has small expected cost compared to ONL. As in the general outline for our offline algorithms in the previous section, the algorithm ensures that at the beginning of a round at least $\alpha N_i + 1$ slots from class C_i are cleared. More precisely, for each class C_i, RND chooses $\alpha N_i + 1$ slots uniformly at random from all but the first slot in the class. At the beginning of each round, RND removes all elements belonging to the selected slots from the buffer.

LEMMA 2.5. *The expected cost of RND in round r is $O(\sqrt{\log\log k/\log k}) \cdot \mathrm{cost}_{\mathsf{ONL}}^r - 1$.*

PROOF. Since RND never chooses to evict the first slot of a class, this slot is never cleared by RND. The probability that a specific other slot of a class is cleared is $(\alpha N_i + 1)/(N_i - 1) < 2\alpha$. Therefore, the expected cost of RND in round r is at most $2\alpha\,\mathrm{cost}_{\mathsf{ONL}}^r$. This is because, due to the way the slots are chosen for promotion, at most $\mathrm{cost}_{\mathsf{ONL}}^r$ slots are promoted that are not the first slot of a class.

From Lemma 2.4b we know that $\mathrm{cost}_{\mathsf{ONL}}^r \geq \alpha d/4$. Now since $\alpha = 1/\sqrt{2d} = \sqrt{\log\log k/\log k}$, we get that the expected cost of RND is at most $2\alpha \cdot \mathrm{cost}_{\mathsf{ONL}}^r \leq 10\alpha \cdot \mathrm{cost}_{\mathsf{ONL}}^r - 1$, where the inequality holds since $8\alpha\,\mathrm{cost}_{\mathsf{ONL}}^r \geq 2\alpha^2 d = 1$. $\square$

The lemma implies the following theorem.

THEOREM 2.6. *Any deterministic algorithm for the reordering buffer management problem has a competitive ratio of $\Omega(\sqrt{\log k/\log\log k})$.*

PROOF. Clearly, the cost of OPT is at most the expected cost of RND which is, due to Lemma 2.2 and Lemma 2.5, at most $\sum_{i=1}^{d} N_i + O(\sqrt{\log\log k/\log k})\sum_r \mathrm{cost}_{\mathsf{ONL}}^r$. Due to Lemma 2.1, the cost of ONL is at least $\sum_r \mathrm{cost}_{\mathsf{ONL}}^r$. Therefore, the competitive ratio tends to $\Omega(\sqrt{\log k/\log\log k})$ as the number of rounds tends to infinity. $\square$

2.3 Lower bound for randomized algorithms

In this section we provide a lower bound of $\Omega(\log\log k)$ on the competitive ratio of any randomized online algorithm for the reordering buffer problem. For this section, we define α to be $(\log\log k)^2/\log k$.

For the analysis in this section, for every class C_i, we choose a slot for promotion in the following way:

> For class C_i choose a slot z in the class uniformly at random. Promote z.

We start by giving a bound on the expected cost of any online algorithm on the resulting input sequence.

LEMMA 2.7. *For a randomized online algorithm ONL, for an input sequence consisting of R rounds, and for sufficiently large k it holds that $\mathbf{E}[\sum_r \mathrm{cost}_{\mathsf{ONL}}^r] \geq R \cdot \log\log k/8$.*

PROOF. Let ONL be an arbitrary online algorithm using a buffer of size k. We fix a round r and analyze $\mathbf{E}[\mathrm{cost}_{\mathsf{ONL}}^r]$. For class C_i, let ℓ_i denote the number of slots cleared by ONL before round r. Note that the ℓ_i's are (dependent) random variables but the following holds for any valid fixed choice of values.

According to Lemma 2.4a we have $\sum_{i=1}^{d}\ell_i \cdot \log^{i-1} k \geq \alpha k/2$. If during the round we promote one of the ℓ_i cleared slots in C_i, $\mathrm{cost}_{\mathsf{ONL}}^r$ increases by one. This happens with probability $\ell_i/N_i = \ell_i \cdot \log^{i-1} k \cdot d/(k+k\alpha) \geq \ell_i \cdot \log^{i-1} k \cdot d/(2k)$. Summing this over all classes we obtain $\mathbf{E}[\mathrm{cost}_{\mathsf{ONL}}^r] \geq \alpha d/4 = \log\log k/8$. Taking the sum over all R rounds completes the proof. $\square$

Next we need to show that the expected optimal cost on the input sequence is significantly smaller.

LEMMA 2.8. *There is an offline algorithm OFF using a buffer of size k such that for an input sequence consisting of R rounds, and for sufficiently large k, $\mathbf{E}[\sum_r \mathrm{cost}_{\mathsf{OFF}}^r] \leq 2\alpha\sum_{i=1}^{d} N_i + O(R)$.*

PROOF. We present an offline algorithm OFF using a buffer of size k that has the desired upper bound on the expected cost. The algorithm ensures that in the beginning of a round at least $\alpha N_i + 1$ slots are cleared from every class C_i. This means that the algorithm has more than $\alpha k + \log^{d-1} k$ elements removed that are held by OPT'.

For a class C_i, one slot is promoted in each round. Consider the sequence of slots that are promoted over all rounds. We partition this sequence into *phases* such that each phase contains $N_i - \alpha N_i - 1$ pairwise different slots. Then, in the beginning of each phase, i.e., in the beginning of the round in

which the first promotion of the new phase takes place, OFF clears all slots that are not contained in the phase. Note that due to the definition of a phase, exactly $\alpha N_i + 1$ slots in class C_i are cleared. Also note that these slots remain cleared during the whole phase, since none of them is promoted.

Let $\mathrm{cost}^r_{\mathsf{OFF}}(i)$ denote the contribution of class C_i to $\mathrm{cost}^r_{\mathsf{OFF}}$, i.e., $\mathrm{cost}^r_{\mathsf{OFF}}(i)$ is 1 if a cleared slot in class C_i is promoted in round r, and 0 otherwise. Clearly $\sum_{i=1}^d \sum_r \mathrm{cost}^r_{\mathsf{OFF}}(i) = \sum_r \mathrm{cost}^r_{\mathsf{OFF}}$. In the following we analyze $\mathbf{E}[\sum_r \mathrm{cost}^r_{\mathsf{OFF}}(i)]$ for $i \in \{1, \ldots, d\}$.

Observe that the contribution of one phase to $\sum_r \mathrm{cost}^r_{\mathsf{OFF}}(i)$ is at most $\alpha N_i + 1$. Let p_i be the total number of phases of class C_i, then we get $\mathbf{E}[\sum_r \mathrm{cost}^r_{\mathsf{OFF}}(i)] \leq \mathbf{E}[p_i](\alpha N_i + 1)$. Let X be the length of a single phase (except the last phase which may be incomplete and, therefore, shorter). Clearly,

$$\mathbf{E}[p_i] \leq 1 + \frac{1}{\mathbf{Pr}[X \geq N_i \ln(1/\alpha)/4]} \cdot \frac{4R}{N_i \ln(1/\alpha)} \ ,$$

where R denotes the total number of rounds. If we can show that $\mathbf{Pr}[X \geq N_i \ln(1/\alpha)/4] \geq 1/2$, the lemma follows since

$$\mathbf{E}[\textstyle\sum_r \mathrm{cost}^r_{\mathsf{OFF}}] = \sum_{i=1}^d \mathbf{E}[\textstyle\sum_r \mathrm{cost}^r_{\mathsf{OFF}}(i)] \leq \sum_{i=1}^d \mathbf{E}[p_i](\alpha N_i + 1)$$

$$\leq \sum_{i=1}^d 2\,\mathbf{E}[p_i]\alpha N_i \leq \sum_{i=1}^d \left(2\alpha N_i + \frac{16R\alpha}{\ln(1/\alpha)} \right)$$

$$\leq 2\alpha \sum_{i=1}^d N_i + \frac{16 d R\alpha}{\ln(1/\alpha)} = 2\alpha \sum_{i=1}^d N_i + O(R) \ .$$

Here, the second inequality uses the fact that $\alpha N_i \geq 1$, which follows from the integrality of αN_i. The proof of the remaining fact that $\mathbf{Pr}[X \geq N_i \ln(1/\alpha)/4] \geq 1/2$ is deferred to Section B in the appendix. $\qquad\square$

THEOREM 2.9. *Any online algorithm for the reordering buffer management problem has competitive ratio at least* $\Omega(\log \log k)$.

PROOF. Combining Lemma 2.7 with Lemma 2.1 shows that the expected cost of an online algorithm ONL on the input sequence consisting of R rounds is at least $\sum_{i=1}^d N_i + R \cdot \log \log k / 8$. Combining Lemma 2.8 with Lemma 2.2 shows that the expected cost of OPT on the input sequence consisting of R rounds is at most $(2\alpha + 1)\sum_{i=1}^d N_i + O(R)$. Therefore, the competitive ratio is at least

$$\frac{\sum_{i=1}^d N_i + R \cdot \log \log k / 8}{(2\alpha + 1)\sum_{i=1}^d N_i + O(R)} \ .$$

Letting R tend to infinity gives the theorem. $\qquad\square$

3. THE DETERMINISTIC UPPER BOUND

In this section we present a deterministic, $O(\sqrt{\log k})$-competitive online algorithm for the reordering buffer management problem. The cost for switching to a color c can be described by a weight w_c for this color. We assume that for all colors c it holds that $1 \leq w_c \leq W$, where W is polynomially bounded in k.

Without loss of generality we can assume that an algorithm for the reordering buffer management problem works according to the following general scheme. In each step the algorithm has an *active output color*, which is equal to the color of the last element that was appended to the output sequence. If there is at least one element with this active color in the buffer, the earliest among these elements is removed, appended to the output sequence, and the next element from the input sequence takes its place in the buffer. Otherwise, if there are no more elements of the active output color in the buffer, the algorithm performs a color-change and chooses a new color (among the colors present in the buffer) to output next.

Note that the algorithm only has to make a decision if a color-change is performed. Therefore, we describe our algorithm LCC (Largest Color Class) by specifying how the new output color is chosen when a color-change is required. But first, we introduce some further notation. The *i-th step* of an algorithm is the step in which the algorithm appends the i-th element to the output sequence. The *buffer content at step i* for an algorithm ALG is the set of elements in ALG's buffer right before the i-th element is moved to the output. For the analysis we will assume that the buffer always contains k elements. This may not be true at the end of the input sequence as the algorithm runs out of elements to fill the buffer. However, this part of the sequence does not influence our asymptotic bounds.

Let for a given color c at a given time t, $\phi_c^t = w_c/n_c^t$ denote the *cost-effectiveness* of color c at time t, where n_c^t denotes the number of elements of color c that are in LCC's buffer at time t. In the following we drop the superscript, as the time step t will be clear from the context.

For each time step, we partition the colors into classes according to cost effectiveness. For $i \in \{-\lceil \log k \rceil, \ldots, \lceil \log W \rceil\}$, the class C_i consists of colors with cost-effectiveness between 2^i and 2^{i+1}. Let $d = O(\log k)$ denote the number of different classes.

The general idea behind the algorithm is that it aims to remove colors from classes that occupy a large fraction of the space in the buffer. To this end the algorithm selects the class that currently occupies the largest space in the buffer and *marks* all colors in this class for eviction (Line 12 in Algorithm 1). Whenever a color-change is required, one of these marked colors is chosen as the new output color (and unmarked). If there are no marked colors left, the new class that occupies the largest space is selected and the process is repeated.

This algorithmic idea is combined with a mechanism that penalizes colors for using up space in the buffer at the time a color-change occurs. This is similar to techniques used e.g. in [14, 8, 3], and ensures that colors (in particular colors with a low weight) do not stay in the buffer for too long, thereby blocking valuable resources.

To realize all this, our algorithm LCC maintains a counter P and additional penalty counters P_c for every color c. LCC also maintains a flag for every color that indicates if the color is marked. Whenever a color is not in the buffer its penalty counter is zero. In particular, in the beginning of the algorithm all penalty counters (including the counter P) are zero. The formal description of our algorithm for selecting a new output color is given as Algorithm 1.

Before the algorithm selects a marked color c_m as the new output color, it assigns a value of w_{c_m} to a penalty counter P (Line 15). In a *post-processing phase* (after outputting all elements of color c_m) this penalty is distributed to penalty counters of individual colors, as follows. The penalty counter P is continuously decreased at rate 1, while the penalty-

Algorithm 1 Largest Color Class (LCC)

1: **Output:** a new output color
2: // let n_c denote the number of elements
3: // with color c in the buffer
4: $\forall$ colors $c: t_c \leftarrow \frac{w_c - P_c}{n_c/k}; \quad t \leftarrow \min(\{t_c \mid \text{color } c\} \cup \{P\});$
5: $P \leftarrow P - t; \quad \forall$ colors $c: P_c \leftarrow P_c + \frac{n_c}{k} \cdot t$
6: // the above ensures that t is small enough such that
7: // $P \geq 0$ and $P_c \leq w_c$ for all c
8: **if** $P = 0$ **then**
9:   **if** no marked color exists **then**
10:     // let $C_{\max}$ denote the class that occupies
11:     // the largest space in the buffer
12:     mark all colors in $C_{\max}$
13:   **end if**
14:   // let c_m denote an arbitrary marked color
15:   $P \leftarrow w_{c_m}$
16:   $P_{c_m} \leftarrow 0$
17:   unmark color c_m
18:   **return** color c_m as the new output color
19: **else**
20:   $c_a \leftarrow \arg\min_c t_c$ // pick color c_a such that $P_{c_a} = w_{c_a}$
21:   $P_{c_a} \leftarrow 0$
22:   unmark color c_a if it was marked
23:   **return** color c_a as the new output color
24: **end if**

counters of colors in the buffer are increased at rate n_c/k where n_c denotes the number of elements of color c that are in the buffer (Lines 4 and 5). Note, that we assume that the buffer is full, hence, the rate of decrease of the P-counter equals the total rate of increase of P_c-counters.

When a counter P_c reaches w_c the penalty distribution is interrupted; the P_c-counter is reset to 0; and the corresponding color c returned as the new output color (Lines 20–23). The penalty distribution resumes when all elements with color c have been removed and the next color-change takes place. The penalty distribution and the post-processing phase ends once the P-counter reaches 0.

We note that the algorithm can be significantly simplified if all colors c have weight $w_c = 1$.

3.1 The analysis

Let for a reordering algorithm ALG and an input sequence σ, ALG(σ) denote the output sequence generated by ALG on input σ. A *color-block* of an output sequence is a maximal subsequence of consecutive elements with the same color. The *cost* of a color-block of color c is equal to the weight w_c of c. The cost $|\text{ALG}(\sigma)|$ of algorithm ALG on input σ is defined as the sum of the costs taken over all color-blocks in the output sequence ALG(σ).

For a color-block b we use $s_{\text{start}}(b)$ and $s_{\text{end}}(b)$ to denote the start index of b and end index of b, respectively, in the output sequence. This is the same as the time-step when the first and the last element of b is appended to the output sequence.

3.1.1 A few simple cases

In this section we first identify different types of color-blocks for which we can fairly easily derive a bound on their respective contribution to the cost $|\text{LCC}(\sigma)|$ of our online algorithm. In Section 3.1.2 we will then introduce a technique that enables us to handle the remaining color-blocks, as well.

We call a color-block of LCC that is not generated in a post-processing phase a *normal* color-block (these are the color-blocks produced when the algorithm switches to the respective color in Line 18). Other color-blocks are called *forced color-blocks* (the ones caused by Line 23). The following lemma shows that we can focus our analysis on normal color-blocks.

LEMMA 3.1. *The sum of the cost of forced color-blocks is at most the sum of the cost of normal color-blocks.*

PROOF. The total cost for forced color-blocks does not exceed the total penalty that is distributed to colors during the post-processing phase of normal color-blocks. The penalty that is distributed during the post-processing phase of a (normal) color-block b with color c is equal to w_c, i.e., the cost for b. Summing over all normal color-blocks gives the lemma. □

In the following we use OPT to denote an optimum offline algorithm. We say that an element is *online-exclusive in step i*, if in this step the element is in LCC's buffer but has already been removed from OPT's buffer. Similarly we call an element *opt-exclusive in step i*, if it is in OPT's buffer but not in LCC's buffer at this time. Note that by this definition in every step the number of online-exclusive elements equals the number of opt-exclusive elements, since the size of LCC's and OPT's buffer is the same.

We extend the first definition to colors: We say that a color is online-exclusive in step i, if there exists an element of this color that is online-exclusive. Finally, we say that a color-block b is online-exclusive if its color is online-exclusive in step $s_{\text{start}}(b)$. The following lemma derives a bound on the cost of online-exclusive color-blocks.

LEMMA 3.2. *The cost of LCC for online-exclusive color-blocks is at most $|\text{OPT}(\sigma)|$.*

PROOF. Let b denote an online-exclusive color-block in LCC(σ), let e denote its first element, and let c be the color of b. Let b' denote the color-block of color c that precedes b (in case b is the first color-block of color c we define $s_{\text{end}}(b') = -1$ for the following argument). Note that element e is not yet in the buffer at step $s_{\text{end}}(b') + 1$ as in this case it would be appended to the output sequence in step $s_{\text{end}}(b') + 1$.

Let b_{opt} denote the color-block of OPT that contains element e. Clearly, this block ends after step $s_{\text{end}}(b') + 1$ as e only arrives after this step. Since b is online-exclusive, its first element (i.e. e) is removed from the OPT-buffer before step $s_{\text{start}}(b)$.

Altogether, we have shown that there exists a color-block b_{opt} in OPT(σ) (with color c) that ends in the interval $]s_{\text{end}}(b') + 1, s_{\text{end}}(b)[$. We match the online-exclusive block b to b_{opt}. In this way we can match every online-exclusive block to a unique block in OPT(σ) with the same color. This gives the claim. □

Another class of color-blocks for which we can easily derive a bound on the contribution to the cost $|\text{LCC}(\sigma)|$ is given by so-called *opt-far color-blocks* defined as follows. A normal color-block b from the sequence LCC(σ) is called *opt-far*, if during its post-processing phase the number of online-exclusive elements never drops below $k/\sqrt{\log k}$. This means that throughout the whole post-processing phase for b the buffers of LCC and OPT are fairly different. The following lemma derives an upper bound on the cost of opt-far blocks in an output sequence generated by LCC.

LEMMA 3.3. *The cost of* LCC *for opt-far color-blocks is at most* $O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$.

PROOF. Fix an opt-far color-block b, and let c denote the color of b. During the post-processing phase for b the number of online-exclusive elements is always above $k/\sqrt{\log k}$. Therefore, at least a $1/\sqrt{\log k}$-fraction of the penalty distributed during the post-processing phase goes to online-exclusive colors. The total cost for online-exclusive color-blocks is at least as large as the penalty that these colors receive, since the penalty of a color c cannot increase beyond its cost w_c.

Hence, the total penalty distributed during the post processing phases of opt-far color-blocks is at most $\sqrt{\log k}$ times the cost for online-exclusive color-changes. This in turn is at most as large as $|\mathsf{OPT}(\sigma)|$ due to Lemma 3.2. The lemma follows by observing that the total penalty distributed during post-processing phases of opt-far color-blocks is equal to the cost of these blocks. $\square$

3.1.2 The potential

A crucial ingredient for the proof of the upper bound in Section 3.1.3 is the way how we handle *normal* color-blocks that are neither *online-exclusive* nor *opt-far*. For this we introduce the notion of *potential*. The idea is that, on the one hand, the total potential is bounded by some function in terms of the optimum cost $|\mathsf{OPT}(\sigma)|$ (see Claim 3.4a). On the other hand, we will show that *normal* color-blocks that are neither *opt-exclusive* nor *opt-far* generate a large potential. This allows us to derive a bound on the contribution of these color-blocks to the cost $|\mathsf{LCC}(\sigma)|$.

The definition of potential is based on the differences in the buffer between LCC and OPT. In the following we use τ_j to denote the start index of the j-th color-block of OPT.

For an element e_τ that is appended to the output sequence $\mathsf{LCC}(\sigma)$ at time τ we define for $\tau_j > \tau$

$$\varphi(\tau, \tau_j) = \begin{cases} 0 & \text{if OPT processed } e_\tau \text{ before step } \tau_j, \\ 1 & \text{otherwise.} \end{cases}$$

$\varphi(\tau, \tau_j)$ simply measures whether the element e_τ occupies a slot in OPT's buffer at time τ_j. We say that element e_τ generates *potential* $w_{c_j} \cdot \varphi(\tau, \tau_j)$ for time step τ_j, where c_j denotes the color of the j-th color-block in $\mathsf{OPT}(\sigma)$.

For technical reasons we also introduce a *capped potential* as follows. We define

$$\hat{\varphi}(\tau, \tau_j) = \begin{cases} 0 & \begin{array}{l} \text{if OPT processed } e_\tau \text{ before step } \tau_j \text{ or at} \\ \text{least } k/\sqrt{\log k} \text{ elements } e_{\tau'} \text{ with } \tau' < \tau \\ \text{have } \varphi(\tau', \tau_j) = 1, \end{array} \\ 1 & \text{otherwise.} \end{cases}$$

$\hat{\varphi}(\tau, \tau_j)$ measures whether the element e_τ is one of the first $k/\sqrt{\log k}$ elements to occupy a slot in OPT's buffer at time τ_j, where elements are ordered according to their appearance in $\mathsf{LCC}(\sigma)$. We say that element e_τ generates *capped potential* $w_{c_j} \cdot \hat{\varphi}(\tau, \tau_j)$ for time step τ_j, where c_j denotes the color that is processed by OPT at time τ_j.

We use $\hat{\varphi}(\tau) := \sum_{j:\tau_j > \tau} w_{c_j} \hat{\varphi}(\tau, \tau_j)$ to denote the total capped potential generated by the element at position τ in $\mathsf{LCC}(\sigma)$. We define the *total capped potential* $\hat{\varphi}$ by $\hat{\varphi} := \sum_\tau \hat{\varphi}(\tau)$.

CLAIM 3.4. *The capped potential fulfills the following properties:*

(a) $\hat{\varphi} \leq k/\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|$.

(b) *Let* $\tau < t < \tau_j$, *and assume that the number of online-exclusive items in step t is at most* $k/\sqrt{\log k}$. *Then* $\hat{\varphi}(\tau, \tau_j) = \varphi(\tau, \tau_j)$, *and, hence, the* capped potential $w_{c_j} \cdot \hat{\varphi}(\tau, \tau_j)$ *generated by e_τ for position τ_j is equal to the potential. In other words the contribution of e_τ is not capped.*

PROOF. The first statement follows from the fact that the capped potential generated for a time-step τ_j cannot exceed $k/\sqrt{\log k} \cdot w_{c_j}$, where w_{c_j} is the cost of OPT in the step. This holds because of the cap. Since the potential is generated for time-steps τ_j that correspond to color-changes by OPT the statement follows.

Now, assume for contradiction that the second statement does not hold. Since, obviously $\hat{\varphi}(\tau, \tau_j) \leq \varphi(\tau, \tau_j)$ it must hold that $\hat{\varphi}(\tau, \tau_j) = 0$ and $\varphi(\tau, \tau_j) = 1$. This means that element e_τ occupies a place in OPT's buffer at time τ_j but there are at least $k/\sqrt{\log k}$ elements $e_{\tau'}$, $\tau' < \tau < t$ that also occupy a place in OPT's buffer at time τ_j, and therefore e_τ's contribution is "capped". But all these elements are opt-exclusive at time t. Since, at any time step the number of opt-exclusive elements must be equal to the number of online-exclusive elements, we can conclude that in step t there are more than $k/\sqrt{\log k}$ online-exclusive elements. This is a contradiction. $\square$

3.1.3 The main theorem

THEOREM 3.5. LCC *is a deterministic online algorithm with competitive ratio* $O(\sqrt{\log k})$.

PROOF. The algorithm LCC marks all colors in a class, and then selects an arbitrary marked color whenever it has to do a normal color-change. When no marked colors are left, it again marks all colors in some class and continues.

We call the time between two marking operations (or after the last marking operation) a *phase*. Fix some phase P and let C denote the set of colors that get marked in the beginning of the phase. Let for $c \in C$, s_c denote the number of elements of color c in LCC's buffer at the time of the marking operation that starts P. Further, let ϕ denote the lower bound on the cost-effectiveness of colors in C, i.e., $\phi \leq w_c/s_c \leq 2\phi$ holds for all colors $c \in C$. We call a color-change normal (forced) if it starts a normal (forced) color-block in the output sequence $\mathsf{LCC}(\sigma)$. In $\mathsf{LCC}(\sigma)$ the phase consists of a consecutive subsequence of elements, starting with an element of a color in C and ending with the last element of a color-block from the post-processing phase of the last normal color-change of the phase.

Let $\mathrm{cost}(P)$ denote the cost incurred by LCC during the phase. This cost consists of color-changes to colors in C (either *normal* or *forced*), and of color-changes to colors not in C (these are *forced*). Let $\mathrm{ncost}(P)$ and $\mathrm{fcost}(P)$ denote the cost incurred by LCC during the phase for *normal* and *forced* color-changes, respectively. Further, let $\mathrm{ncost} := \sum_{\text{phases } P} \mathrm{ncost}(P)$ denote the total normal cost summed over all phases. In light of Lemma 3.1 it is sufficient to relate ncost to the optimum cost $|\mathsf{OPT}(\sigma)|$. In order to do this we distinguish the following cases:

Case 1 The normal cost $\mathrm{ncost}(P)$ is at most $9/10 \cdot \mathrm{fcost}(P)$. Let $\mathrm{ncost}_{\mathrm{small}}$ denote the normal cost summed over all phases P that fulfill this condition and let $\mathrm{ncost}_{\mathrm{large}}$ denote the normal cost summed over other phases (i.e., $\mathrm{ncost}_{\mathrm{large}} =$

ncost − ncost$_{\text{small}}$). Then

$$\text{ncost}_{\text{small}} \leq \frac{9}{10} \sum_P \text{fcost}(P) \leq \frac{9}{10} \text{ncost}$$

$$= \frac{9}{10}(\text{ncost}_{\text{small}} + \text{ncost}_{\text{large}}) \ ,$$

where the second inequality follows from Lemma 3.1. This gives ncost$_{\text{small}} \leq 10\,\text{ncost}_{\text{large}}$. In the following cases we show that ncost$_{\text{large}} \leq O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$. With this fact we have that the normal cost ncost$_{\text{small}}$ generated by phases that fulfill the condition for Case 1 is at most $O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$.

Case 2 The cost of OPT during the phase is at least ncost$(P)/4$. The total normal cost generated by phases that fulfill this condition is at most $O(|\mathsf{OPT}(\sigma)|)$.

Case 3 The cost of online-exclusive color-blocks generated during the phase is at least ncost$(P)/4$. Then, we can amortize the normal cost of the phase against the cost of online-exclusive color-blocks, which in turn can be amortized against the cost of OPT by Lemma 3.2. This gives that the total normal cost generated by phases that fulfill the conditions for this case is at most $O(|\mathsf{OPT}(\sigma)|)$.

Case 4 The cost of opt-far color-blocks generated during the phase is at least ncost$(P)/4$. Then we can amortize the cost of the phase against the cost of opt-far color-blocks, which in turn can be amortized against the cost of OPT by Lemma 3.3.

Hence, the total cost for phases that fulfill the conditions for this case is at most $O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$.

Case 5 In the following we assume that none of the above cases occurs. This means there must exist a subset $C' \subset C$ of colors such that for each color $c \in C'$ its first color block in the phase is

(a) not online-exclusive,

(b) not opt-far, and

(c) not forced.

Further, we have that

(d) elements of colors in C' are not appended to the output sequence by OPT during the phase;

(e) and cost$(C') \geq$ cost$(C)/10$,

where we use cost$(S) := \sum_{c \in S} w_c$ for a set S of colors.

To see this we generate C' as follows. First take all colors from C (colors initially marked in the phase) and remove colors among them for which the first color-change is forced (this ensures Property c). The cost of the remaining set of colors is exactly ncost(P). Then remove colors for which the first block of the phase is onl-exclusive or opt-far, and colors that are requested by OPT during the phase. Since, we are not in Case 2, Case 3 or Case 4 this step can only remove colors with a total cost of $3/4 \cdot$ ncost(P). After this properties a, b and d hold. This gives the set C'.

Property e can be seen as follows. By the construction cost$(C') \geq$ ncost$(P)/4$. From the fact that Case 1 does not hold we get that ncost$(P) \geq 9/10 \cdot$ fcost(P), and, hence, $2 \cdot$ ncost$(P) \geq 9/10 \cdot$ cost$(P) \geq 9/10 \cdot$ cost(C). This gives cost$(C') \geq 9/80 \cdot$ cost$(C) \geq 1/10 \cdot$ cost(C).

Let S denote the set of elements with colors in C' that are in LCC's buffer in the beginning of the phase. We will show that these elements generate a large potential after the end of the phase. From this it follows that we can amortize the cost of the phase against $|\mathsf{OPT}(\sigma)|$ because of the following argument.

Assume that for some value Z the elements in S generate a potential of at least $Z \cdot$ cost(C') after time t, where t is the index of the last time-step of the phase. Observe that, according to Property b above, the (first) color-blocks of colors in C' that are generated during the phase are not opt-far. This means that during the post-processing phase of each of these blocks, the number of online-exclusive items falls below $k/\sqrt{\log k}$ at some point. This means that we can apply Claim 3.4b to all elements in S, which gives that elements in S also generate $Z \cdot$ cost(C') *capped potential* after time t, as their contribution to the potential is not capped.

Claim 3.4a tells us that the total capped potential is at most $O(k/\sqrt{\log k}) \cdot |\mathsf{OPT}(\sigma)|$. Therefore, the total normal cost generated by phases that fulfill the conditions for Case 5 is at most

$$\frac{O(k/\sqrt{\log k}) \cdot |\mathsf{OPT}(\sigma)|}{Z \cdot \text{cost}(C')} \text{cost}(C) = O\Big(\frac{k}{Z\sqrt{\log k}}|\mathsf{OPT}(\sigma)|\Big).$$

By showing that elements in S generate at least potential $Z \cdot$ cost$(C') = \Omega(k/\log k) \cdot$ cost(C') we get that this generated cost is at most $O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$.

For completing the analysis of Case 5 it remains to show the above bound on the potential generated by elements of S. For this, we first show that the cardinality of the set S is large. We have

$$|S| = \sum_{c \in C'} s_c \geq \sum_{c \in C'} w_c/2\phi \geq \frac{1}{20\phi} \sum_{c \in C} w_c \geq \frac{1}{20} \sum_{c \in C} s_c$$

$$\geq \frac{k}{20d} = \Omega(k/\log k) \ ,$$

where the first and third inequality follows since $\phi \leq w_c/s_c \leq 2\phi$; the second inequality holds since $\sum_{c \in C'} w_c =$ cost$(C') \geq$ cost$(C)/10 \geq \sum_{c \in C} w_c/10$. The last inequality follows since the algorithm LCC selects a class that occupies the largest space in the buffer, and, hence, occupies at least space k/d, where d denotes the number of classes.

With this lower bound on $|S|$ the statement follows directly from the following claim that is proved in Section A in the appendix.

CLAIM 3.6. *Let S denote a set of elements that are opt-exclusive at time t, and let s_c denote the number of elements of color c in S. Let* cost$(S) = \sum_{c:s_c>0} w_c$, *and assume that there is a value ϕ such that $\phi \leq w_c/s_c \leq 2\phi$ holds for all colors with elements in S. Then the contribution to the potential by elements from S generated after time t is at least* $|S| \cdot$ cost$(S)/8$.

Applying the claim with t being the last step of the phase, gives that the elements from S generate potential at least $\Omega(k/\log k) \cdot$ cost(C') after the end of the phase. This finishes the analysis of Case 5.

The above cases show that the contribution of all phases to the cost of LCC is at most $O(\sqrt{\log k} \cdot |\mathsf{OPT}(\sigma)|)$. This gives the theorem. □

4. REFERENCES

[1] Amjad Aboud. Correlation clustering with penalties and approximating the reordering buffer management problem. *Master's thesis, Computer Science Department, The Technion — Israel Institute of Technology*, 2008.

[2] Yuichi Asahiro, Kenichi Kawahara, and Eiji Miyano. NP-hardness of the sorting buffer problem on the uniform metric. Unpublished manuscript, 2010.

[3] Noa Avigdor-Elgrabli and Yuval Rabani. An improved competitive algorithm for reordering buffer management. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 13–21, 2010.

[4] Reuven Bar-Yehuda and Jonathan Laserson. Exploiting locality: Approximating sorting buffers. In *Proceedings of the 3rd Workshop on Approximation and Online Algorithms (WAOA)*, pages 69–81, 2005.

[5] Dan Blandford and Guy Blelloch. Index compression through document reordering. In *Proceedings of the Data Compression Conference (DCC)*, pages 342–351, 2002.

[6] Ho-Leung Chan, Nicole Megow, Rob van Stee, and René Sitters. The sorting buffer problem is NP-hard. *CoRR*, abs/1009.4355, 2010.

[7] Matthias Englert, Harald Räcke, and Matthias Westermann. Reordering buffers for general metric spaces. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC)*, pages 556–564, 2007.

[8] Matthias Englert and Matthias Westermann. Reordering buffer management for non-uniform cost models. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 627–638, 2005.

[9] Iftah Gamzu and Danny Segev. Improved online algorithms for the sorting buffer problem. In *Proceedings of the 24th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 658–669, 2007.

[10] Kai Gutenschwager, Sven Spiekermann, and Stefan Voß. A sequential ordering problem in automotive paint shops. *International Journal of Production Research*, 42(9):1865–1878, 2004.

[11] Rohit Khandekar and Vinayaka Pandit. Online and offline algorithms for the sorting buffers problem on the line metric. *Journal of Discrete Algorithms*, 8(1):24 –35, 2010.

[12] Jens S. Kohrt and Kirk Pruhs. A constant factor approximation algorithm for sorting buffers. In *Proceedings of the 6th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 193–202, 2004.

[13] Jens Krokowski, Harald Räcke, Christian Sohler, and Matthias Westermann. Reducing state changes with a pipeline buffer. In *Proceedings of the 9th International Fall Workshop Vision, Modeling, and Visualization (VMV)*, pages 217–224, 2004.

[14] Harald Räcke, Christian Sohler, and Matthias Westermann. Online scheduling for sorting buffers. In *Proceedings of the 10th European Symposium on Algorithms (ESA)*, pages 820–832, 2002.

APPENDIX

A. PROOF OF CLAIM 3.6

P R O O F. Let $c_1, \ldots, c_\ell$ denote the colors in S, ordered according to the times $\tau_1 < \cdots < \tau_\ell$ at which the first element of a color is evicted by OPT. Let i denote the smallest number such that $\sum_{j=1}^{i} w_{c_j} \geq \text{cost}(S)/2$.

We show that the number of elements with colors $c_i, \ldots, c_\ell$ is large. For any j we have $\phi \leq w_{c_j}/s_{c_j} \leq 2\phi$, and, hence, also $\text{cost}(S) \geq \phi|S|$. Therefore $\sum_{j=i}^{\ell} s_{c_j} \geq \frac{1}{2\phi} \sum_{j=i}^{\ell} w_{c_j} \geq \frac{1}{4\phi}\text{cost}(S) \geq \frac{1}{4}|S|$.

Let e_τ denote an element from S that is appended to $\text{LCC}(\sigma)$ at time τ, and let c be its color. For any color c_j that is removed by OPT before color c, e_τ generates potential (not necessarily capped potential) $w_{c_j} \cdot \varphi(\tau, \tau_j) = w_{c_j}$ at time-step τ_j. If c is evicted by OPT after c_i we get that e_τ generates potential $\sum_{j=1}^{i} w_{c_j} \geq \text{cost}(S)/2$. As the number of elements with colors evicted after c_i is at least $|S|/4$, the potential generated by them is at least $|S| \cdot \text{cost}(S)/8$. □

B. PROOF OF CLAIM B.1

The analysis of the random variable X is based on a straightforward coupon collector type argument and included for completeness.

C L A I M B.1. $\mathbf{Pr}[X \geq N_i \ln(1/\alpha)/4] \geq 1/2$.

P R O O F. Consider a phase and let X_j be the number of rounds between the promotion of the $(j-1)$-th distinct slot and the j-th distinct slot. Then $X = X_1 + X_2 + \cdots + X_{N_i - \alpha N_i} - 1$. We have $\mathbf{E}[X_j] = \frac{N_i}{N_i - j + 1}$ and we get

$$
\begin{aligned}
\mathbf{E}[X] &= \frac{N_i}{N_i} + \frac{N_i}{N_i - 1} + \cdots + \frac{N_i}{\alpha N_i + 1} - 1 \\
&= N_i(H_{N_i} - H_{\alpha N_i}) - 1 \\
&\geq N_i \cdot \ln N_i - N_i \cdot (\ln(\alpha N_i) + 1) - 1 \\
&= N_i \ln(1/\alpha) - (N_i + 1) \geq N_i \ln(1/\alpha)/2 \;,
\end{aligned}
$$

where the last step follows for sufficiently small α (i.e., for sufficiently large k).

From Chebyshev's inequality we get

$$
\begin{aligned}
\mathbf{Pr}&\left[X \leq N_i \ln(1/\alpha)/4\right] \\
&\leq \mathbf{Pr}\left[|X - \mathbf{E}[X]| \geq N_i \ln(1/\alpha)/4\right] \\
&\leq \frac{16}{N_i^2 \ln^2(1/\alpha)} \cdot \mathbf{Var}[X] \\
&= \frac{16}{N_i^2 \ln^2(1/\alpha)} \cdot \sum_{j=1}^{(1-\alpha)N_i} \mathbf{Var}[X_j] \\
&\leq \frac{16}{N_i^2 \ln^2(1/\alpha)} \cdot \sum_{j=1}^{(1-\alpha)N_i} \mathbf{E}[X_j]^2 \\
&\leq \frac{16}{\ln^2(1/\alpha)} \cdot \sum_{j=1}^{\infty} \frac{1}{j^2} \leq \frac{32}{\ln^2(1/\alpha)} \leq \frac{1}{2} \;,
\end{aligned}
$$

where the third step follows since the X_j's are independent and the last step follows for sufficiently small α. □