

Reordering Buffer Management with Advice

Anna Adamaszek¹, Marc P. Renault², Adi Rosen³, and Rob van Stee¹

¹ Max-Planck-Institut für Informatik, Saarbrücken, Germany,
`{anna, vanstee}@mpi-inf.mpg.de`

² Univerité Paris Diderot - Paris 7 and UPMC, France.
`mrenault@liafa.univ-paris-diderot.fr`

³ CNRS and Univerité Paris Diderot - Paris 7, France.
`adiro@liafa.univ-paris-diderot.fr`

Abstract. In the reordering buffer management problem, a sequence of colored items arrives at a service station to be processed. Each color change between two consecutively processed items generates a cost. A reordering buffer of capacity k items can be used to preprocess the input sequence in order to decrease the number of color changes. The goal is to find a scheduling strategy that, using the reordering buffer, minimizes the number of color changes in the given sequence of items.

We consider the problem in the setting of online computation with advice. In this model, the color of an item becomes known only at the time when the item enters the reordering buffer. Additionally, together with each item entering the buffer, we get a fixed number of advice bits, which can be seen as information about the future or as information about an optimal solution (or an approximation thereof) for the whole input sequence. We show that for any $\varepsilon > 0$ there is a $(1 + \varepsilon)$ -competitive algorithm for the problem which uses only a constant (depending on ε) number of advice bits per input item.

We complement the above result by presenting a lower bound of $\Omega(\log k)$ bits of advice per request for an optimal deterministic algorithm.

1 Introduction

In the *reordering buffer management problem*, a sequence of colored items arrives at a service station for processing. At any time, the service station is configured to process items of a certain color c . Changing the configuration of the service station, i.e., preparing the station to service items of a different color, generates a cost. The total cost of processing a sequence of items is equal to the number of color changes in the sequence. In order to reduce the processing cost, the service station is equipped with a *reordering buffer* that has a capacity to hold k items. At each point in time, the buffer contains the first k items of the input sequence that have not yet been processed, and one arbitrary item has to be chosen to be extracted from the buffer and processed by the service station. If the color of the chosen item differs from the color of the previously processed item, the configuration of the service station has to be changed, generating a cost. The goal is to find a scheduling strategy that, using the reordering buffer, preprocesses the input sequence to minimize the number of color changes. In the *online setting*, which is considered in this paper, the color of an item becomes known only at the time when the item enters the reordering buffer, i.e., we do not know the whole input sequence in advance.

This framework has many applications in various areas, such as production engineering, storage systems, network optimization and computer graphics (e.g., [?, ?, ?, ?, ?]). One simple example is a paint shop of a car manufacturing plant, where switching colors between two consecutively painted cars generates costs and delays due to the necessary cleaning and set-up. Therefore, some paint shops are equipped with a reordering buffer to preprocess the sequence of incoming cars (see [?]).

In this paper, we study the problem in the setting of *online computation with advice* [EFKR11]. This setting has gained increasing attention in recent years in that it relaxes the traditional online setting of no information about the future, and allows the online algorithm to get some partial information about future requests or about the structure of an optimal (or near-optimal) solution. It also allows one to quantify the amount of information about the future available to the online algorithm and to study the interplay between the amount of information and the attainable competitive ratio. Informally (See Section ?? for a formal definition), in this setting, the online algorithm receives with each request some b bits of *advice* which are the value of a function, defined by the algorithm, of the whole input sequence (including the future). In this manner the online algorithm receives some information about the future. Note that we view the advice as given by an oracle, i.e., we do not consider the issue if the advice can be computed efficiently when knowing the whole input sequence.

Related work. The reordering buffer management problem has been introduced by Racke et al. [?], and has been extensively studied. In the online setting, the best known results are a deterministic $O(\sqrt{\log k})$ -competitive algorithm by Adamaszek et al. [?], and a randomized $O(\log \log k)$ -competitive algorithm by

Avigdor-Elgrabli and Rabani [?]. To complement this, there are (nearly) matching lower bounds of $\Omega\left(\sqrt{\frac{\log k}{\log \log k}}\right)$ and $\Omega(\log \log k)$ on the competitive ratio of any online deterministic and randomized algorithms, respectively [?]. In the offline setting, i.e., when the whole input sequence is known in advance, the best known result is a constant factor approximation by Avigdor-Elgrabli and Rabani [?], while the problem is known to be NP-hard [?,?].

More general versions of the problem have been studied, where the context switching cost for switching from an item of color c to an item of color c' depends on c' (e.g., [?,?]), or on both c and c' [?].

The model of online computation with advice considered in the present paper was introduced by Emek et al. [EFKR11]. In that paper, the authors gave tight bounds of $\Theta(\log n/b)$ on the competitive ratio of deterministic and randomized online algorithms with advice for Metrical Task Systems, where n is the number of states of the system and b is the number of advice bits per request. They also gave a deterministic online algorithm with advice for the k -server problem which is $k^{O(\frac{1}{b})}$ -competitive, where $\Theta(1) \leq b < \log k$. This has been first improved by Böckenhauer et al. [BKKK11], and later improved to $\left\lceil \frac{\lceil \log k \rceil}{b-2} \right\rceil$ by Renault and Rosén [?].

Böckenhauer et al. [BKK⁺09] introduced a somewhat similar model for online algorithms with advice, where the advice is a single tape of bits instead of being given separately for each request. This allows an algorithm to read a different number of bits of advice per request, but it requires that the online algorithm knows how many bits of advice to read with each request. Thus, the two models are, in general, incomparable. Several results were given in this related model [DKP08,BKK⁺09,HKK10,KK11,BKKK11]. For example, in [DKP08,BKK⁺09], the authors explore the number of bits of advice required for deterministic and randomized paging algorithms, algorithms for the DiffServe problem, algorithms for a special case of the Job Shop Scheduling problem, and algorithms for the Disjoint Path Allocation problem to be 1-competitive.

Our contribution. We give an online algorithm with advice that, for any $\varepsilon > 0$, achieves a competitive ratio of $1 + \varepsilon$, using only $O(\log(1/\varepsilon))$ advice bits per request. We show how to construct advice for any input sequence, based on an optimal solution for the given sequence, which allows us to obtain a good competitive ratio. The overview of the construction is as follows. The advice bits for each element of a color c encode how the algorithm should handle all “adjacent” elements of color c , i.e., if the algorithm should keep all the elements in the buffer until more elements of the color arrive, output these elements at once, or output them, but only after a certain waiting period. The idea is that the order of the colors with the latter property is not contained in the advice of single elements (as that would require too many bits of advice), but it is encoded together among all the the advice bits of the given color. The key obstacle is that with a small number of advice bits per item we cannot encode the exact order in which the colors should be output. Therefore, we have to modify the optimal solution by removing some elements earlier from the buffer and keeping some

other elements until later, so that we can encode the order of all colors which have to be output “soon, but not yet”. We show how to perform this operation with only a slight increase of the cost of the generated output sequence.

We complement the above result by giving a lower bound that shows that, in order for an online algorithm to have optimal performance, the number of bits of advice must depend on k . More precisely, $\Omega(\log k)$ bits are required. This lower bound applies even if all the advice bits are given to the algorithms before the sequence starts, and it matches (up to constant factors) the trivial upper bound, where the advice indicates which color switch has to be performed at each step.

2 The Model

We use the definition of *deterministic online algorithms with advice* as presented in [EFKR11]. An online algorithm is defined as a request-answer game that consists of a request set, R ; a sequence of finite nonempty answer sets, A_1, A_2, \dots ; and a sequence of cost functions, $\text{cost}_n : R^n \times A_1 \times A_2 \times \dots \times A_n \rightarrow \mathcal{R}^+ \cup \{\infty\}$ for $n = 1, 2, \dots$. In addition, online algorithms with advice have access via a query to an advice space, U , which is a finite set. The advice space has a size of 2^b , where $b \geq 0$ is the number of bits of advice provided to the algorithm with each request. With each request, the online algorithm receives some advice that is defined by a function, $u_i : R^* \rightarrow U$, where i is the request index, that is applied to the whole request sequence, including future requests. A deterministic online algorithm with advice can, thus, be represented as a sequence of pairs (g_i, u_i) , where g_i is the function defining the action of the online algorithm at step i and is defined $g_i : R^i \times U^i \rightarrow A_i$ for $i = 1, 2, \dots$. The action taken by the online algorithm after receiving request r_i is therefore a function of the first i requests, r_1, \dots, r_i , and the advice received so far, $u_1(\sigma), \dots, u_i(\sigma)$.

We use the standard definitions of competitive analysis. We say that an algorithm is α -*competitive* if, for every finite request sequence σ , $|\text{ALG}(\sigma)| \leq \alpha \cdot |\text{OPT}(\sigma)| + \zeta$, where ζ is a constant which does not depend on the request sequence σ , $|\text{ALG}(\sigma)|$ denotes the cost of the solution generated by ALG for σ , and $|\text{OPT}(\sigma)|$ is the cost of an optimal solution for σ .

3 Structure of Advice and the Algorithm

The advice for each input element e consists of two parts: the *type* t_e and the *value* v_e . The type of an element e describes how e , or rather a whole collection of elements of the same color as e , should be handled by the algorithm ALG, i.e., if ALG should keep them in the buffer until more elements of the same color arrive, if ALG should output them at once, or output them, but only after a certain waiting period. There are four possible types of an element, i.e., $t_e \in \{\text{WAIT}, \text{LIST}, \text{READY}, \text{COMPLETE}\}$. The value v_e of an element e is used to encode an order, according to which some input elements will be output. For each element e we have $v_e \in \{0, \dots, D - 1\}$, where D is a constant depending on ε and will be fixed later.

The advice sequence will be constructed based on an optimal solution OPT for the input sequence.

Time In this paper, we use the following notion of time. In each time step $1, \dots, n$, one element arrives and is stored in the buffer. In each time step $k, \dots, k + n - 1$, one element is removed from the buffer, after the element which arrives at that time (if any) has been stored.

Color blocks The elements which have the same color are partitioned into *color blocks*. A *complete color block* is a maximal set of elements of one color ending with an element of type COMPLETE, and containing no other elements of type COMPLETE. The advice sequence will be constructed in such a way that the types of the consecutive elements from one complete color block always form an expression

$$(\text{WAIT})^* ((\text{LIST})^* \cup (\text{READY})^*) \text{COMPLETE},$$

i.e., first there is an arbitrary number (possibly zero) of elements of type WAIT, then an arbitrary number of elements of type LIST or an arbitrary number of type READY, and at the end exactly one element of type COMPLETE.

At time t , a *color block* is that part of a complete color block which is contained in the buffer (some elements may have been already served, and some may not have been read yet). The notion of a color block here differs from the standard definition of a color block for the reordering buffer management problem, where it denoted the set of all elements of one color contained in the buffer, which were then output consecutively from the buffer. We now can have more than one color block of the same color in the buffer. The advice sequence will be constructed in such a way, that the algorithm ALG will always output all elements of each complete color block consecutively.

We define the *type* t_B of a color block B in the following way. If there is an element of type COMPLETE in B , then the type of B is COMPLETE. If all the elements of B have type WAIT, B has type WAIT. Otherwise, the type of B is LIST or READY, depending on whether B contains elements of type LIST or of type READY.

The values v_e of all elements of a color block B together encode a *value of a color block* $v_B \in \{0, \dots, D^{|B|} - 1\}$. The value of a color block will be encoded only for color blocks containing elements of type LIST, and it will define the order in which these color blocks will be output by the algorithm ALG.

Waiting list The algorithm ALG maintains a *waiting list* of color blocks of type LIST. This list contains only (a subset of) those color blocks that have the property that some elements of their color are read into the buffer of OPT while OPT is serving it. At the beginning of the algorithm, the waiting list is empty. Whenever a color block B of type LIST appears (i.e. the first element of type LIST for some color block is read from the input), it is inserted into the waiting list. The *initial* position of the block on the waiting list, i.e., the position where the block is inserted, is defined by the value v_B of the block. Value 0 denotes the

head of the list. Note that when a new color block is inserted into the waiting list, the position of other blocks on the list can change.

The value v_B of each color block is read from the advice of its elements. However, some short color blocks may not contain sufficiently many elements to encode their required position v_B . We will check for this case and ensure that such blocks are never stored in the waiting list. Instead, we serve parts of them immediately and keep only part in the buffer.

The algorithm Starting at time k , the algorithm ALG chooses a color block to be output in the following way.

1. If ALG has a complete color block in the buffer, it chooses the oldest of them (i.e., the block which first became complete is chosen).
2. Otherwise, if ALG has a color block of type READY in the buffer, it chooses the oldest of them (i.e., the block which first obtained type READY is chosen).
3. Otherwise ALG chooses a color block of type LIST which is at the head of the waiting list. The color block is then removed from the waiting list.

The algorithm then outputs consecutively all elements of the chosen color block, one element at each time step. Notice that when ALG starts outputting a color block, possibly some elements of the block have not yet been read from the input. When ALG gets such elements from the input sequence, it appends them to the color block and outputs them without making any color change.

The complete construction of the advice sequence is given in Section 4. In Section 5, we show that the construction of the advice sequence guarantees that ALG can always choose a color block to be output (i.e., there is always a block of type COMPLETE, READY or LIST in the buffer of ALG). We also show that when ALG starts outputting a color block, it can output *all* elements from that color block without making any color change. That means in particular that the cost of the algorithm is upper bounded by the number of color blocks in the advice sequence.

4 Constructing the Advice Sequence

Overview The advice sequence $A_\varepsilon(\sigma)$ is constructed *offline*, based on an *optimal* solution OPT for the instance σ of the reordering buffer management problem. The idea of the construction is as follows. We initially assign each input element e type t_e based on the way OPT handles e . Let c be the color of e . If OPT makes a color change right after outputting e , e is assigned type COMPLETE. Otherwise, if the whole color block containing e is kept in the buffer of OPT until the next element element of color c is read from the input sequence, e is gets type WAIT. Otherwise, e gets type LIST. This ensures the following invariant.

Invariant 1 *For each color block, its elements are output by OPT in a single block (i.e., without making a color change).*




Fig. 1. The procedures $\text{Split}(B)$ (on the left) and $\text{Postpone}(B)$ (on the right). None of the resulting blocks remain on the waiting list. The letters w, l, r, c are used to denote elements of types WAIT, LIST, READY and COMPLETE, respectively. The frames represent color blocks.

For each color block B which contains elements of type LIST, we want to assign a value v_B , which is the initial position of the block on the waiting list. The blocks in the waiting list are ordered according to the order in which they are output by OPT. If for each color block the advice included the exact position of the block in the waiting list, based on the output sequence of OPT, the algorithm ALG would output an optimal solution. However, for some short color blocks we cannot encode their position on the waiting list. To deal with this problem, we will modify the advice data, to decrease the number of blocks which are in the waiting list at any time. The modification increases the number of color blocks (i.e., the cost of ALG becomes larger than the cost of OPT), and introduces elements of type READY. Blocks containing elements of type READY are not inserted into the waiting list. We now give the overview of the needed modifications of the advice data.

Procedure Remove(B) Remove the block B from the waiting list. For each block B' that was inserted into the waiting list later than B , and behind B , decrease $v_{B'}$ by one.

Procedure Split(B) Run $\text{Remove}(B)$. Reassign type COMPLETE to the first element of block B which had type LIST, and reassign type READY to the remaining elements of B which had type LIST assigned (see Figure 1, left). Note that possibly some of these elements have not been read from the input yet.

The block B has been split into two blocks, both ending with an element of type COMPLETE. The elements of the first block (called the *early* block) are removed from the buffer. The second block is called a *late* block. This block is kept in the buffer (but not on the waiting list).

Procedure Postpone(B) Run $\text{Remove}(B)$. Reassign type WAIT to the first element of B which originally had type LIST assigned, and reassign type READY to the remaining elements of B which had type LIST assigned (see Figure 1, right). Possibly some of these elements have not been read from the input yet. Blocks processed in this way are called *postponed blocks*.

The procedure $\text{Split}(B)$ makes ALG evict the elements from the first block of B (i.e., the early block) earlier than they were evicted by OPT, generating free space in the buffer of ALG. This allows some blocks to stay in the buffer until the next element of the block is read from the input, and only then to be output by ALG. These are the blocks for which we run procedure $\text{Postpone}(B)$. We will specify later which blocks will be treated in this way. Whenever a block B is removed from the waiting list using these procedures, the effect is that ALG will never insert it into the waiting list. This is the reason we update the values $v_{B'}$ of other blocks on the waiting list when applying procedure $\text{Remove}(B)$. We have the following observation.

Observation 1 *The procedures $\text{Split}(B)$ and $\text{Postpone}(B)$ maintain Invariant 1.*

We are now ready to describe the details of the construction of the advice sequence $A_\varepsilon(\sigma)$. The input elements from the sequence σ are processed one by one, using a buffer of size k (which we call an *advice buffer*) and a waiting list. Notice that at this point we are only creating an advice sequence, that is assigning type t_e and value v_e for each input element e , and not creating an output sequence. In particular, "removing elements from the buffer" does not mean "appending elements to the output sequence".

Processing an input element Whenever there is empty space in the buffer, we read the next element e from the input sequence σ . We proceed as follows. First, we assign the initial type of the element e as described above, based on the way OPT handles e . Then, if e is the first element in its block which is of type LIST, i.e., its color block changes type to LIST, we insert the color block of e at the appropriate position, based on the optimal output sequence OPT, into the waiting list. If after this operation the waiting list has too many blocks of length similar to the one added, we remove all such blocks from the waiting list, by applying procedure $\text{Split}(B)$ to some of the blocks, and procedure $\text{Postpone}(B)$ to the remaining blocks of the given length. We use the following definition.

Definition 1. *The class of a block B on the waiting list is $\lfloor \log |B| \rfloor$, where $|B|$ is the number of items of this block up to and including the first element of type LIST.*

We now give the detailed description of actions performed for a newly read element e of color c . Here C is another constant which will depend on ε .

1. (*Type assignment*) We assign the element e type t_e based on the way OPT handles e (see Figure 2).
 - (a) If OPT makes a color change right after outputting e , e is assigned type COMPLETE.
 - (b) Else, if e has already been assigned type READY due to some previous Split or Postpone operation (i.e., before e has been read into the buffer), it keeps the type READY.




Fig. 2. Possible initial types of elements (excluding the special case of READY). The type LIST indicates that OPT starts serving this color block before all elements have been read into the buffer; the remaining elements enter the buffer while the color block is being served (and while other elements of the color block are being removed).

- (c) Else, if the whole color block containing e is kept in the buffer of OPT until the next element element of color c is read from the input sequence, e is gets type WAIT.
- (d) Else, e gets type LIST.

Note that the type of the element e can be modified later, by applying procedure $\text{Split}(B)$ or $\text{Postpone}(B)$ to the color block B of e .

2. (*Insertion into the waiting list*) If e is the *first* element of its color block of type LIST, insert the color block B of e into the waiting list at position v_B (derived from the solution OPT). Let i be the class of block B . Count the number n_i of color blocks of class i which are currently in the waiting list.

If $n_i = C$, let B_1 and B_2 be the two last color blocks of class i in the waiting list. Run $\text{Split}(B_1)$ and $\text{Split}(B_2)$. For each other block B_0 of class i on the waiting list, run $\text{Postpone}(B_0)$. By performing these operations we ensure that ALG will never insert these C blocks into the waiting list.

Note that in step 2 some blocks may be split, with the resulting early blocks being removed from the buffer.

Whenever the advice buffer becomes full, and there are no elements of the current color block in the buffer, we choose a new color block B to be processed (i.e., removed from the buffer), according to the rules of ALG. The following lemma shows that the rules of ALG can always be successfully applied.

Lemma 1. *While creating the advice sequence, if blocks are removed from the buffer according to the rules of ALG and Step 2 above, then starting from time k there is always an element of type different than WAIT in the buffer.*

Let B be the color block chosen from the full advice buffer according to the rules of ALG. If B is not on the waiting list, we simply remove it from the buffer. If B is on the waiting list, we have to consider two cases.

1. The value v_B is smaller than $D^{|B|}$, i.e., it can be encoded using $|B|$ values v_e of the elements of the block. We set appropriately the values v_e , and we

remove B from the buffer and from the waiting list. (Note that this time we do not apply the procedure $\text{Remove}(B)$, i.e., we do not modify any values $v_{B'}$ of blocks B' .)

Only blocks processed in this way will ever be inserted by ALG into the waiting list.

2. The value v_B is at least $D^{|B|}$. Run $\text{Split}(B)$. In this case, only the early block obtained from B is removed from the buffer.

Notice that, as above, in determining the length of B in Step 1 we count only the elements up to and including the first element of type LIST. The reason is that these are the only elements of B that have been inserted into the buffer before B has to be inserted into the waiting list, i.e., before the value of B has to be computed.

Note also that the initial value of v_B might decrease (due to calls to the procedure Remove when B is already in the waiting list) before B is processed by ALG. Specifically, it does not matter if v_B was initially too high, as long as it drops below $D^{|B|}$ before the block is chosen to be processed.

5 Analysis of ALG

Lemma 1 leaves open the possibility that ALG may be able to serve only *part* of a block. This could happen if not all elements of the (complete) block enter the buffer while ALG is serving its color. In such a case, ALG would have to return to this color several times, and its cost would be higher than the number of color blocks given by the advice sequence $A_\varepsilon(\sigma)$. At the beginning of this Section we will show that this does not happen, and ALG always outputs complete color blocks. Then we will bound the competitive ratio of ALG.

First, we present two technical lemmas.

Lemma 2. *In the advice sequence $A_\varepsilon(\sigma)$, types of the consecutive elements from one complete color block always form an expression*

$$(\text{WAIT})^* ((\text{LIST})^* \cup (\text{READY})^*) \text{COMPLETE}.$$

Lemma 3. *Let c_1, \dots, c_ℓ be a collection of consecutive elements of color c , output by OPT with no color change in between, and let t be the time when OPT outputs the first element c_1 of the collection. If ALG outputs c_1 at some time $t' \geq t$, and it has not output any c_i before time t' , it can output all elements c_1, \dots, c_ℓ with no color change in between.*

We can now consider all types of blocks which ALG starts outputting before they become complete, and show that OPT cannot keep the elements of such blocks longer in the buffer. Applying Lemma 3 gives us the following results.

Lemma 4. *When ALG starts outputting a color block, it can finish it with no color changes.*

Theorem 1. *The algorithm ALG always outputs complete color blocks.*

In the remaining part of this Section we will bound the competitive ratio of ALG, if ALG uses the advice sequence $A_\varepsilon(\sigma)$. For this, we will need an upper bound on the cost of ALG (denoted by $|\text{ALG}|$), as well as a lower bound on the cost of OPT (denoted by $|\text{OPT}|$). From the construction of the advice sequence, we know the following.

Observation 2 *The cost of ALG is upper bounded by the number of color blocks given by the advice sequence $A_\varepsilon(\sigma)$.*

Observation 3 *The cost of OPT is lower bounded by the number of color blocks other than the late blocks given by the advice sequence $A_\varepsilon(\sigma)$.*

To bound the competitive ratio, it is enough to bound the number of late color blocks in $A_\varepsilon(\sigma)$, compared to the number of color blocks of other types. Late blocks can be generated in two ways: while inserting new elements into the waiting list, and when processing a color block at the head of the waiting list. We call the sets of late blocks generated during these two operations by LATE_I and LATE_P , respectively.

Lemma 5. *We have $|\text{LATE}_I| \leq \frac{2}{C}(|\text{POSTPONED}| + |\text{EARLY}|)$, where POSTPONED and EARLY denote the sets of postponed and early blocks for $A_\varepsilon(\sigma)$.*

Lemma 6. *We have $|\text{LATE}_P| \leq \frac{C}{D-1}|\text{LISTED}|$, where LISTED denotes the set of color blocks which contain elements of type LIST .*

Combining the two Lemmas above and setting $C = \lceil 2/\varepsilon \rceil$ and $D = \lceil C/\varepsilon \rceil + 1$ gives us the main theorem of the paper.

Theorem 2. *For any $\varepsilon > 0$ there is a $(1 + \varepsilon)$ -competitive algorithm for the reordering buffer management problem which uses only $O(\log(1/\varepsilon))$ advice bits per input item.*

6 Optimality Lower Bound

In this Section we show that a deterministic algorithm requires $\Omega(\log k)$ bits of advice per request in order to be optimal. Throughout this Section, we assume without loss of generality that an algorithm will only perform a color switch when there are no more items of the current color in the buffer, and the buffer is either full or contains the last request. Let k be the size of the buffer, and let

$$\sigma = \langle c_1, c_2, \dots, c_k, \pi_1(1), c_{k+1}, \pi_1(2), c_{k+2}, \dots, \pi_1(k), c_{2k}, \pi_2(1), \pi_2(2), \dots, \pi_2(k) \rangle,$$

where π_1 is a permutation of the colors c_1, \dots, c_k , and π_2 is a permutation of the colors c_{k+1}, \dots, c_{2k} . Note that $|\sigma| = 4k$.

As long as the buffer of an optimal algorithm remains full, the algorithm must switch to the color c_i when next element waiting to enter the buffer has color c_i . In particular, the first k consecutive colors output by the algorithm must be $\pi_1(1), \dots, \pi_1(k)$. More formally,

Lemma 7. *Given σ , let S be the sequence of colors $\pi_1(1), \dots, \pi_1(k)$. Let \mathcal{C} be the set of all the possible optimal sequences of color switches for σ . The k -prefix of all the sequences in \mathcal{C} is S .*

Using Lemma 5, we can prove the main theorem of this section.

Theorem 3. *At least $\frac{\log k}{8}$ bits of advice per request are required for a deterministic algorithm with advice to be optimal.*

References

- [BKK⁺09] Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Královic, Richard Královic, and Tobias Mömke. On the advice complexity of online problems. In Yingfei Dong, Ding-Zhu Du, and Oscar H. Ibarra, editors, *ISAAC*, volume 5878 of *Lecture Notes in Computer Science*, pages 331–340. Springer, 2009.
- [BKKK11] Hans-Joachim Böckenhauer, Dennis Komm, Rastislav Královic, and Richard Královic. On the advice complexity of the k -server problem. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP (1)*, volume 6755 of *Lecture Notes in Computer Science*, pages 207–218. Springer, 2011. Also as technical report at <ftp://ftp.inf.ethz.ch/pub/publications/tech-reports/7xx/703.pdf>.
- [DKP08] Stefan Dobrev, Rastislav Královic, and Dana Pardubská. How much information about the future is needed? In *SOFSEM'08: Proceedings of the 34th conference on Current trends in theory and practice of computer science*, pages 247–258, Berlin, Heidelberg, 2008. Springer-Verlag.
- [EFKR11] Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. *Theor. Comput. Sci.*, 412(24):2642–2656, 2011.
- [HKK10] Juraj Hromkovic, Rastislav Královic, and Richard Královic. Information complexity of online problems. In Petr Hlinený and Antonín Kucera, editors, *MFCs*, volume 6281 of *Lecture Notes in Computer Science*, pages 24–36. Springer, 2010.
- [KK11] Dennis Komm and Richard Královic. Advice complexity and barely random algorithms. In Ivana Cerná, Tibor Gyimóthy, Juraj Hromkovic, Keith G. Jeffery, Rastislav Královic, Marko Vukolic, and Stefan Wolf, editors, *SOFSEM*, volume 6543 of *Lecture Notes in Computer Science*, pages 332–343. Springer, 2011.

A Proofs from Section 4

Proof (of Lemma 1). Assume for contradiction that at some time t , i.e., after reading exactly t items from the input, all items have type WAIT. If all these elements originally had type WAIT assigned to them, they are kept in the buffer of OPT at time t , and OPT does not remove any of them before it reads an element with a matching color from the input sequence. Contradiction, as OPT has to remove an element at time t , and its buffer is completely filled with elements which cannot be removed at that time.

Therefore, we must have run $\text{Postpone}(B)$ for a block B at some point, and reassigned a type from LIST to WAIT. If at time t , OPT has already finished outputting any postponed block B from the advice buffer, the element of B of type COMPLETE would be already read from the input and it would be in the advice buffer — contradiction. So no postponed blocks have been output by OPT yet. That means that at time t , OPT is in the process (or is just beginning) of outputting one of the postponed blocks B , and all but $k - |B|$ other elements from the advice buffer are still in the buffer of OPT (and have type WAIT). However, as B is a postponed block, there are two blocks of size greater than $|B|/2$ each, which are kept in the buffer of OPT, and which have been removed from the advice buffer. Contradiction, as OPT does not have enough buffer space. \square

B Proofs from Section 5

Proof (of Lemma 2). The type COMPLETE is only assigned to final elements of blocks (in Step 1, and also in the procedure $\text{Split}(B)$). In Step 1, the type WAIT is only ever assigned to elements that form a *prefix* of a color block, and $\text{Postpone}(B)$ only extends such a prefix by one element, while $\text{Split}(B)$ keeps the types of such a prefix unchanged. Finally, the types LIST and READY are never combined in one block, neither in Step 1 nor in Step 2 (see also Figure 1). \square

Proof (of Lemma 3). At time t , OPT has the elements c_1, \dots, c_m in its buffer for some $m \in [1, \ell]$ (along with $k - m$ other items). It switches to the color c of c_1 and serves this color for exactly the next ℓ time steps. Afterwards, OPT has read the input until item exactly $t + \ell$. The buffer content of OPT is the same as at time t , except that the m elements of color c have been replaced by a set S of m new elements with different colors which arrive in the interval $[t + 1, t + \ell]$.

At time t' , based on the assumptions of the lemma, ALG has the elements $c_1, \dots, c_{m'}$ in its buffer for some $m' \in [m, \ell]$ (and $k - m'$ other elements). ALG switches to color c and starts serving it. It continues serving this color and reading the input until it has read exactly $m' \geq m$ elements of different colors or until it has served all ℓ items of the block, whichever comes first.

To find out how far ALG can read the input, we must count the number of elements that are not of color c in the input starting from time $t' + 1 \geq t + 1$, and continue until we find $m' \geq m$ of them. But then we certainly count all elements of S (that ALG had not yet read at time t'). Hence ALG reads the input at least as far as OPT, and can indeed serve the entire color block. \square

Proof (of Lemma 4). The only color blocks which ALG can start outputting before reading all the elements from the input are READY blocks (postponed and late blocks) and blocks of type LIST. The other blocks are COMPLETE and can trivially be output at once.

First consider postponed blocks. From the construction of the advice sequence, we know that ALG starts outputting a postponed block later than OPT: OPT outputs it strictly before reading the block element f that follows the first element e of type LIST, but the advice for e gets modified to WAIT in the Postpone operation and therefore ALG only outputs e *after* reading f (of type READY or COMPLETE). From Lemma 3, ALG can finish outputting all the elements of the block without a color change.

Now consider non-complete blocks of type LIST. Let t' be the time when ALG starts outputting a block B of type LIST. By Invariant 1 and Observation 1, OPT outputs all the elements of B in a single block. We want to show that OPT starts outputting the block B at some time $t \leq t'$, as then by Lemma 3 we are done.

Assume that OPT starts outputting B at a time $t > t'$. As B has type LIST at time t' , we get that at time t' there are no blocks of type READY or COMPLETE in the buffer of ALG. All the blocks in the buffer of ALG at time t' are

1. original blocks of type WAIT, which OPT has still in the buffer at time t' ;
2. postponed blocks, which still have type WAIT because of the reassignment of the type, where OPT can be in the process of outputting the first of them at time t' . In this case, OPT still has two early blocks in its buffer, of total size greater than $|B|$, which ALG has already output;
3. blocks of type LIST, which are in the waiting list behind B , i.e. OPT still has them in its buffer at time t , and therefore also at time t' .

OPT also keeps the whole block B in its buffer — contradiction, as OPT does not have enough buffer capacity to have an item which it would remove at time t' .

Finally, consider late READY blocks. Let t' be the time when ALG starts outputting a late READY block B of color c . Say the first element f of this block is loaded into the buffer at time $t \leq t'$ (note that t does not depend on the algorithm used). Originally, f was part of a longer block B' and was directly preceded in this block (not necessarily directly preceded in the input) by e , the first element of B' of type LIST. Since e had type LIST, by definition OPT starts to serve B' strictly before f arrives, i.e., before time $t \leq t'$. Since the late block B is not complete at time t' , but it is processed together with the corresponding early block by OPT, OPT must still be serving color c at time t' (either the early or the late block that was produced by the splitting of B').

Since there are no complete blocks in the buffer of ALG at time t' , all blocks that OPT served completely before time t' have also been served by ALG by time t' (possibly using multiple color changes per block). Hence, all items that are in the buffer of ALG at time t' and that are not of the current color of OPT (which is c) are also in the buffer of OPT. Therefore, in the time period which starts at time t' in which *both* ALG and OPT are serving color c , OPT does not have more room than ALG to store new items of different colors, so it cannot read more of the input than ALG, and ALG manages to serve the entire late block B . \square

Proof (of Theorem 1). From Lemma 4, whenever ALG starts outputting a color block, it can finish it with no color changes. From Lemma 1, one of the rules of ALG can always be applied. \square

Proof (of Lemma ??). Color blocks from the set LATE_I are obtained while inserting new elements into the waiting list, when there are C blocks of the same class in the waiting list. The C blocks are then modified to obtain two late blocks, two early blocks and $C - 2$ postponed blocks. The inequality follows. \square

Proof (of Lemma ??). Let $B \in \text{LATE}_P$, and let $\text{orig}(B)$ be the original block which later has been split into an early block and the late block B . Let i be the class of $\text{orig}(B)$. As $\text{orig}(B)$ has been split, at the time when $\text{orig}(B)$ has been processed we had $v_{\text{orig}(B)} \geq D^{|B|} \geq D^{2^i}$.

Let $\mathcal{B}(B)$ be the set of color blocks which, during the construction of the advice sequence, have been in the waiting list at the time $\text{orig}(B)$ has been inserted there, and which have been successfully removed from the waiting list (i.e., they did not become postponed block, and did not split into early and late blocks) before $\text{orig}(B)$ has been split. From the construction of the advice sequence we get that $|\mathcal{B}(B)| = v_{\text{orig}(B)}$. Each block from the set $\mathcal{B}(B)$ belongs to the set LISTED.

Let $B_0 \in \text{LISTED}$, and let t be the time when B_0 has been (successfully) removed from the waiting list. For any block $B \in \text{LATE}_P$ such that $B_0 \in \mathcal{B}(B)$, the block $\text{orig}(B)$ must have been in the waiting list at the time t . As for any class i the number of blocks of class i which are in the waiting list at time t is smaller than C , B_0 belongs to less than C sets $\mathcal{B}(B)$, such that $\text{orig}(B)$ is in class i .

Now, for each block $B \in \text{LATE}_P$, we charge a cost of $\frac{1}{v_{\text{orig}(B)}} \leq \frac{1}{D^{2^i}}$ to each of the $v_{\text{orig}(B)}$ blocks from the set $\mathcal{B}(B)$. The total cost charged to all blocks in LISTED equals $|\text{LATE}_P|$. From the reasoning above we know that each block from the set LISTED can belong to at most C sets $\mathcal{B}(B)$ per each class i , and therefore is charged a total cost less than

$$\sum_{i=0}^{\lceil \log k \rceil} \frac{C}{D^{2^i}} \leq \frac{C}{D-1} .$$

That gives us that $|\text{LISTED}| \cdot \frac{C}{D-1} \geq |\text{LATE}_P|$. \square

Proof (of Theorem ??). We set $C = \lceil 2/\varepsilon \rceil$ and $D = \lceil C/\varepsilon \rceil + 1$. From Lemma ?? and Lemma ?? we get that

$$|\text{LATE}_I| + |\text{LATE}_P| \leq \left(\max \left\{ \frac{2}{C}, \frac{C}{D-1} \right\} \right) |\text{OPT}| \leq \varepsilon |\text{OPT}| .$$

That gives us $|\text{ALG}| \leq (1 + \varepsilon)|\text{OPT}|$. The number of advice bits per input item is $2 + \lceil \log D \rceil = O(\log(1/\varepsilon))$. \square

C Proofs from Section ??

Proof (of Lemma 5). Let OFF be an offline algorithm for σ that performs $2k$ color switches, denoted by C^{OFF} , such that the first k color switches are those of S and the last k color switches are $\pi_2(1), \pi_2(2), \dots, \pi_2(k)$. The correctness proof for OFF is by induction on j over the indexes of C^{OFF} . Initially, OFF fills its buffer with the first k requests which are k distinct colors. By definition, the request waiting to enter the buffer is $\pi_1(1)$ which is the color switched to by OFF. OFF serves both requests of color $\pi_1(1)$ in σ with one color switch and a new color enters the buffer. Assume that for every color switch for $j = 1$ to $u - 1$, the 2 requests of $C^{\text{OFF}}(j)$ are served. For the inductive step from $u - 1$ to u , let σ_u be σ with the pairs of requests with colors of $C^{\text{OFF}}(1), \dots, C^{\text{OFF}}(u - 1)$ removed. The 2 requests of color $C^{\text{OFF}}(u)$ occur in the first $k + 1$ requests of σ_u . Therefore, both requests will be served when OFF switches to $C^{\text{OFF}}(u)$. Since OFF correctly serves σ and the number of color switches is equal to the number of distinct colors in the request sequence, OFF is optimal.

Finally, we show by induction on j over the indexes of S , that S is the only sequence of color switches that an optimal algorithm can perform. Note that no algorithm can perform better than the number of distinct color switches. Therefore, by induction, we will show that any deviation from S results in an additional color switch which implies more color switches than OFF. At the point of the first color switch, the buffer of an algorithm will contain the first k distinct colors and $S(1)$ is waiting to enter the buffer. If any other color besides $S(1)$ is switched to at this point, the algorithm will only serve the first request of that color resulting in an additional color switch to serve the second request of the same color. Therefore, an optimal algorithm must switch to $S(1)$ initially. Assume that this is the case for j from 1 to $u - 1$. For the inductive step from $u - 1$ to u , let σ_u be σ with the pairs of requests with colors of $S(1), \dots, S(u - 1)$ removed. The first k requests of σ_u will be the first request to k distinct colors including $S(u)$ and the $k + 1$ -th request will have the color $S(u)$. As in the base case, an optimal algorithm must switch to $S(u)$ at this point. \square

Proof (of Theorem ??). There are $(k!)^2$ different input sequences of the form

$$\sigma = \langle c_1, c_2, \dots, c_k, \pi_1(1), c_{k+1}, \pi_1(2), c_{k+2}, \dots, \pi_1(k), c_{2k}, \pi_2(1), \pi_2(2), \dots, \pi_2(k) \rangle,$$

with $k! \geq k^{\frac{k}{2}}$ different permutations π_1 . By Lemma 5, the sequence of the first k colors output by an optimal algorithm must be exactly the sequence π_1 .

Assume, for contradiction, that there is a deterministic algorithm using less than $\frac{\log k}{8}$ bits of advice per input item and which is optimal. As the algorithm uses less than $\frac{\log k}{8}$ bits of advice per input item, i.e., less than $\frac{k \log k}{2}$ bits of advice in total, there must be two sequences σ_1 and σ_2 which have different permutations π_1 , and for which the whole advice sequence $A = \{a_1, \dots, a_{4k}\}$ is the same.

The first color switch, for both sequences σ_1 and σ_2 , is to the same color $C(1) = g(c_1, \dots, c_k, a_1, \dots, a_k)$. As the algorithm is optimal both for σ_1 and σ_2 ,

the color $\pi_1(1)$ must be the same for both σ_1 and σ_2 . Assume that $\pi_1(j)$, for some $j \in \{2, \dots, k\}$, is the first position at which σ_1 and σ_2 differ. Then, as the algorithm is optimal, it performs the same color switches for both σ_1 and σ_2 over the first $j - 1$ color switches. Now, let

$$C(j) = g(c_1, \dots, c_k, \pi_1(1), c_{k+1}, \pi_1(2), c_{k+2}, \dots, \pi_1(j-1), c_{k+j-1}, a_1, \dots, a_{k+2(j-1)}).$$

The j -th color switch of the algorithm both for σ_1 and σ_2 is the same, as it is to the color $C(j)$. However, $C(j)$ must be different from $\pi_1(j)$ for at least one of the sequences σ_1, σ_2 . We obtain contradiction with the optimality of the algorithm.

A deterministic algorithm using less than $\frac{\log k}{8}$ bits of advice per input item cannot be optimal.

□