

Optimal Projective Algorithms for the List Update Problem*

Christoph Ambühl[†] Bernd Gärtner[‡] Bernhard von Stengel[§]

January 16, 2010

Abstract

The list update problem is a classical online problem, with an optimal competitive ratio that is still open, known to be somewhere between 1.5 and 1.6. An algorithm with competitive ratio 1.6, the smallest known to date, is COMB, a randomized combination of BIT and TIMESTAMP(0). This and many other known algorithms, like MTF, are *projective* in the sense that they can be defined by looking only at any pair of list items at a time. Projectivity simplifies both the description of the algorithm and its analysis, and so far seems to be the only way to define a good online algorithm for lists of arbitrary length. In this paper we characterize all projective list update algorithms and show that their competitive ratio is never smaller than 1.6 in the partial cost model. Therefore, COMB is a best possible projective algorithm in this model.

Keywords: linear lists, online algorithms, competitive analysis.

AMS subject classifications: 68W27, 68W40, 68P05, 68P10.

*A preliminary version of this paper appeared in [5].

[†]Department of Computer Science, University of Liverpool, Liverpool L69 3BX, United Kingdom.
Email: christoph.ambuhl@googlemail.com

[‡]Institute for Theoretical Computer Science, ETH Zürich, 8092 Zürich, Switzerland.
Email: gaertner@inf.ethz.ch

[§]Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom.
Email: stengel@maths.lse.ac.uk

1 Introduction

The *list update problem* is an online problem in the area of self-organizing data structures [3]. Requests to items in an unsorted linear list must be served by accessing the requested item. We assume the *partial cost model* where accessing the i th item in the list incurs a cost of $i - 1$ units. This is simpler to analyze than the original *full cost model* [10] where that cost is i . The goal is to keep access costs small by rearranging the items in the list. After an item has been requested, it may be moved free of charge closer to the front of the list. This is called a *free exchange*. Any other exchange of two consecutive items in the list incurs cost one and is called a *paid exchange*.

An *online* algorithm must serve the sequence σ of requests one item at a time, without knowledge of future requests. An optimum *offline* algorithm knows the entire sequence σ in advance and can serve it with minimum cost $\text{OPT}(\sigma)$. If the online algorithm serves σ with cost $A(\sigma)$, then it is called *c-competitive* if for a suitable constant b

$$A(\sigma) \leq c \cdot \text{OPT}(\sigma) + b \tag{1}$$

for all request sequences σ and all initial list states. The constant c is also called the *competitive ratio*. If the above inequality holds even for $b = 0$, the algorithm A is called *strictly c-competitive*.

The *move-to-front* rule MTF, for example, which moves each item to the front of the list after it has been requested, is strictly 2-competitive [10, 11]. This is also the best possible competitiveness for any deterministic online algorithm for the list update problem [10]. Another 2-competitive deterministic algorithm is TS, which is the simplest member of the TIMESTAMP class due to Albers [1]. TS moves the requested item x in front of all items which have been requested at most once since the last request to x .

As shown first by Irani [9], *randomized* algorithms can perform better on average. Such an algorithm is called *c-competitive* if

$$E[A(\sigma)] \leq c \cdot \text{OPT}(\sigma) + b,$$

for all σ and all initial list states, where the expectation is taken over the randomized choices of the online algorithm. The best randomized list update algorithm known to date is the 1.6-competitive algorithm COMB [2]. It serves the request sequence with probability $4/5$ using the algorithm BIT [10]. With probability $1/5$, COMB treats the request sequence using TS.

Lower bounds for the competitive ratio of randomized algorithms are harder to find; the first nontrivial bounds are due to Karp and Raghavan, see the remark in [10]. In the partial cost model, a lower bound of 1.5 is easy to find as only two items are needed. Teia [12] generalized this idea to prove the same bound in the full cost model, which requires long lists. The authors [6] showed a lower bound of 1.50084 (improved to 1.50115 in [4, p. 38]) for lists with five items in the partial cost model, using game trees and a modification of Teia's approach. The optimal competitive ratio for the list update problem (in the partial cost model) is therefore between 1.50115 and 1.6, but the true value is as yet unknown.

Our results. With the exception of Irani’s algorithm SPLIT [9], all the specific list update algorithms mentioned above are *projective*, meaning that the relative order of any two items x and y in the list after a request sequence σ only depends on the initial list state and the requests to x and y in σ . The simplest example of a projective algorithm is MTF. In order to determine whether x is in front of y after σ , all that matters is whether the last request to x was before the last request to y . The requests to other items are irrelevant.

A simple example of a non-projective algorithm is TRANSPOSE, which moves the requested item just one position further to the front.

The main result of this paper is a proof that 1.6 is the best possible competitive ratio attainable by a projective algorithm. As a tool, we develop an explicit characterization of deterministic projective algorithms.

These results are significant in two respects. First, they show that the successful approach of combining existing projective algorithms to obtain improved ones has reached its limit with the development of the COMB algorithm. New and better algorithms (if they exist) have to be non-projective, and must derive from new, yet to be discovered, design principles.

Second, the characterization of projective algorithms is a step forward in understanding the structural properties of list update algorithms. With this characterization, the largest and so far most significant class of algorithms appears in a new, unified way.

The complete characterization of projective algorithms turns out to be rather involved. However, there is a simple subclass of projective algorithms which already covers all reasonable projective algorithms. We call them *critical request algorithms*. A list update algorithm is completely described by the list state after a request sequence σ has been served, without specifying the particular rearrangement of items after each request; this can be done since we can assume that all changes in the list state are due to paid exchanges, as explained in further detail at the beginning of the next section. For critical request algorithms, the *unary projections* to individual items suffice to describe that list state. For a request sequence σ and list item x , deleting all requests to other items defines the unary projection σ_x , which is an i -fold repetition of requests to x , written as x^i , for some $i \geq 0$. In Section 5 it will be necessary to consider unary projections x^0 and y^0 of length zero as different if the items x and y are different; for the moment, this distinction does not matter. With L as the set of list items, let the set of these unary projections be

$$U = \{x^i \mid x \in L, i \geq 0\}. \quad (2)$$

Definition 1 (Critical request algorithm).

A *deterministic* critical request algorithm is defined by a function

$$F : U \rightarrow \{0, 1, 2, \dots\}, \quad \text{with } F(x^i) \leq i \text{ for any } x \in L.$$

We call the $F(\sigma_x)$ th request to x in σ the *critical request* to x . If $F(\sigma_x)$ is zero (for example if σ_x is the empty sequence \emptyset), then x has no critical request. In the list state after σ , all items with a critical request are grouped together in front of the items without critical request. The items with critical requests are ordered according to the time of the

$F(\sigma_x)$ th request to x in σ . The later a critical request took place in the sequence, the closer the item is to the front. The items without critical request are placed at the end of the list according to their order in the initial sequence. A *randomized critical request algorithm* is a probability distribution on the set of deterministic critical request algorithms.

As an example, consider the online algorithm for three items a , b , and c with the function F shown in the following table for requests up to four items.

i	0	1	2	3	4
$F(a^i)$	0	1	0	2	2
$F(b^i)$	0	0	2	2	4
$F(c^i)$	0	1	2	2	2

In the rest of this paper, list states are written as $[x_1x_2 \dots x_n]$ where x_1 is the item at the front of the list. Let the initial list state be $[abc]$. Consider the list state after $\sigma = abbcab$. We have $F(\sigma_a) = F(aa) = 0$, hence a does not have a critical request. For b we have $F(\sigma_b) = F(bbb) = 2$, therefore the second request to b in σ is its critical request. For c we have $F(\sigma_c) = F(c) = 1$. Thus after σ , the list state is $[cba]$. If we augment σ by another request to a , item a moves to the front, because its critical request is the second.

Algorithms based on critical request functions are clearly projective, since the relative order of any pair of items just depends on the relative order of the requests to x and y in σ and the relative order of x and y in the initial list state.

In good online algorithms, the critical requests are very recent, like in MTF which is described by the critical request function $F(x^i) = i$ for all items x . We define the critical request *relative* to the current position by

$$f(x^i) = i - F(x^i), \quad (3)$$

from which the critical request function is recovered as $F(x^i) = i - f(x^i)$. Then MTF is given by $f(x^i) = 0$. Algorithm TS is described by $f(x^i) = 1$ for all items and all $i > 0$ (and $f(\emptyset) = 0$). Because the BIT algorithm [10] is randomized, its critical requests are also randomized. For every item x , its relative critical request function can be written as $f(\emptyset) = 0$ and, for $i > 0$,

$$f(x^i) = (b_x + i) \pmod{2} \quad (4)$$

where $b_x \in \{0, 1\}$ is chosen once uniformly at random. Hence the critical request is the last or the second-to-last request with equal probability.

The structure of the paper is as follows. In the next section, we explain projective algorithms in more detail and how they can be analyzed. In Section 3, we give a characterization of M -regular projective algorithms, followed by the lower bound of 1.6 for this class of algorithms in Section 4. In Section 5, we characterize projective algorithms completely. We extend the lower bound to the full class in Section 6.

2 Projective Algorithms

In order to characterize list update algorithms, we first simplify their formal definition. The standard definition (of the partial cost model) considers a list state and a sequence of requests. For each request to one of the items to the list, the item can be accessed with access cost $i - 1$ if the item is in position i , and then moved free of charge closer to the front. In addition, paid exchanges are allowed which can be applied both before and after accessing the item, at a cost of one unit for exchanging any two consecutive items.

Contrary to the claim of [11, Theorem 3], paid exchanges may strictly improve costs. For example, let the initial list state be $[abc]$ and $\sigma = cbbc$. Then an optimal algorithm moves a behind b and c before the first request to c . This requires paid exchanges.

On the other hand, free exchanges can be mimicked by paid exchanges as follows: Instead of first paying k units in order to access item x and then moving it at no charge t positions closer to the front, one can first move the item t positions forward and then access the item. In both cases, one pays exactly k units.

Furthermore, one can restrict paid exchanges to take place just before the access to a requested item. This holds because after accessing an item, one can postpone paid exchanges until the next request is revealed.

The above considerations lead to a simplified but still equivalent model of list update algorithms. With the simplified model, we can specify any deterministic online algorithm A by a function

$$S^A : \Sigma \rightarrow \mathcal{L}.$$

Here, Σ denotes the set of finite request sequences, whereas \mathcal{L} denotes the set of the $n!$ states the list of n items can attain. By definition, $S^A(\sigma)$ denotes the list state after the last request of σ has been served by algorithm A .

Consider a request sequence σ and assume it is followed by a request to item x , the extended sequence denoted by σx . Then the cost of serving request x is defined by: the cost of re-arranging the list from state $S^A(\sigma)$ to $S^A(\sigma x)$ by paid exchanges, plus the cost of accessing x in state $S^A(\sigma x)$.

Using this notation, the initial list state can be denoted by $S^A(\emptyset)$. We will omit the superscript A in $S^A(\sigma)$ when the algorithm used is determined by the context.

In order to describe projective algorithms, we have to introduce the concept of *projections* of request sequences and list states. Let a request sequence σ be given and fix a pair of items x, y . The projection σ_{xy} of σ to x and y is the request sequence σ where all requests which are not to x or y are removed. Similarly, σ_x is σ with all requests other than x removed.

Given a list state L , the projection to x and y is obtained by removing all items except for x and y from the list. This is denoted by L_{xy} .

Definition 2. Let $S_{xy}(\sigma)$ be the projection of $S(\sigma)$ to x and y . A deterministic algorithm A is called *projective* if for all pairs of items x, y and all request sequences σ

$$S_{xy}(\sigma) = S_{xy}(\sigma_{xy}). \tag{5}$$

A randomized algorithm is projective if all deterministic algorithms that it chooses with positive probability are projective.

Thus, an algorithm is projective if the relative position of any pair of items depends only on the initial list state and the requests to x and y in the request sequence.

Projective algorithms have a natural generalization, where the relative order of any k -tuple of list items depends only on the requests to these k items. It turns out that for lists with more than k items, only projective algorithms satisfy this condition. This follows from the fact that, for example for $k = 3$, $S_{xyz}^A(\sigma) = S_{xyz}^A(\sigma_{xyz})$ (so the relative position of x and y does not depend on requests to w), and $S_{xyw}^A(\sigma) = S_{xyw}^A(\sigma_{xyw})$ (so the relative position of x and y does not depend on z), which implies that $S_{xy}(\sigma)$ depends only on σ_{xy} .

Already in [7], Bentley and McGeoch observed that MTF is projective: Item x is in front of y if and only if y has not been requested yet or if the last request to x took place after the last request to y .

With the exception of Irani's SPLIT algorithm [9], projective algorithms are the only family of algorithms that have been analyzed so far, typically using the following theorem, for example in [1, 2, 8].

Theorem 3. If a (strictly) projective algorithm is c -competitive on lists with two items, then it is also (strictly) c -competitive on lists of arbitrary length.

Proof. Consider first an arbitrary list update algorithm A . We define the *projected cost* $A_{xy}(\sigma)$ that A serves a request sequence σ , projected to the pair x, y , as follows: Let $\sigma'z$ be any prefix of σ . Then $A_{xy}(\sigma)$ is the number of times, for all requests z in σ , where $S_{xy}(\sigma')$ and $S_{xy}(\sigma'z)$ differ (which counts the necessary paid exchanges of x and y ; this may happen even if $z \notin \{x, y\}$ in case A is not projective), plus the number of times where $z = x$ and $S_{xy}(\sigma'z) = [yx]$ or $z = y$ and $S_{xy}(\sigma'z) = [xy]$. Let L be the set of list items. Then

$$A(\sigma) = \sum_{\{x,y\} \subseteq L} A_{xy}(\sigma), \quad (6)$$

because the costs $A(\sigma)$ are given by the update costs for changing $S(\sigma')$ to $S(\sigma'z)$, which is the sum of the costs of paid exchanges of pairs of items, plus the cost of accessing z in state $S(\sigma'z)$.

For a projective algorithm A the relative behavior of a pair of items is, according to (5), independent of the requests to other items. It is therefore easy to see that $A_{xy}(\sigma) = A_{xy}(\sigma_{xy})$ for projective algorithms: Because A is projective, $A_{xy}(\sigma_{xy})$ is also the cost of A for serving σ_{xy} on the two-item list containing x and y starting from $S_{xy}(\emptyset)$.

For the algorithm OPT, one can interpret $\text{OPT}_{xy}(\sigma_{xy})$ as the cost of optimally serving σ_{xy} on the two-item list $S_{xy}(\emptyset)$. To see this, note that in order to serve σ optimally, one can always start by moving the items not requested in σ to the tail of the list, without changing their relative order. Concerning OPT, we have $\text{OPT}_{xy}(\sigma) \geq \text{OPT}_{xy}(\sigma_{xy})$. Then

$$\text{OPT}(\sigma) = \sum_{\{x,y\} \subseteq L} \text{OPT}_{xy}(\sigma) \geq \sum_{\{x,y\} \subseteq L} \text{OPT}_{xy}(\sigma_{xy}) =: \overline{\text{OPT}}(\sigma). \quad (7)$$

Let A be a projective algorithm that is c -competitive on two items. Then for every pair of items x, y there is a constant b_{xy} such that for all σ

$$A_{xy}(\sigma_{xy}) \leq c \cdot \text{OPT}_{xy}(\sigma_{xy}) + b_{xy}.$$

Then

$$\begin{aligned} A(\sigma) &= \sum_{\{x,y\} \subseteq L} A_{xy}(\sigma_{xy}) \\ &\leq \sum_{\{x,y\} \subseteq L} (c \cdot \text{OPT}_{xy}(\sigma_{xy}) + b_{xy}) \\ &\leq c \cdot \overline{\text{OPT}}(\sigma) + \sum_{\{x,y\} \subseteq L} b_{xy} \\ &= c \cdot \overline{\text{OPT}}(\sigma) + b \\ &\leq c \cdot \text{OPT}(\sigma) + b. \end{aligned}$$

For the strict case, just set all $b_{xy} := 0$. □

Not all algorithms are projective. Let LMTF be the algorithm that moves the requested item x in front of all items which have not been requested since the previous request to x , if there has been such a request.

It is easy to prove that on lists with two items, combining LMTF and MTF with equal probability would lead to a 1.5-competitive randomized algorithm. Obviously, if LMTF was projective, this bound would hold for lists of arbitrary length.

However, LMTF is not projective. This can be seen from the request sequence $\sigma = baacbc$ with initial list $L_0 = [abc]$. It holds that $S^{\text{LMTF}}(\sigma) = cab$, whereas $S^{\text{LMTF}}(\sigma_{bc}) = S^{\text{LMTF}}(bcbc) = bca$. Hence $S_{bc}^{\text{LMTF}}(\sigma) \neq S_{bc}^{\text{LMTF}}(\sigma_{bc})$.

3 Critical Requests and M -regular Algorithms

In this section, we consider deterministic projective list update algorithms. In order to refer to the individual requests to an item x , we write unary projections as

$$x^i = x_{(1)}x_{(2)} \dots x_{(i)},$$

that is, $x_{(q)}$ is the q th request to x in σ if $\sigma_x = x^i$, for $1 \leq q \leq i$.

Let $\mathcal{P}(\sigma)$ be the set of all permutations of the sequence σ . In particular, $\mathcal{P}(x^i y^j)$ consists of all sequences with i requests to x and j requests to y .

Swapping two requests $x_{(q)}$ and $y_{(l)}$ in a request sequence σ means that $x_{(q)}$ and $y_{(l)}$, which are assumed to be adjacent, change their position in σ . If two requests are not adjacent, they cannot be swapped.

Definition 4. A pair of unary projections x^i, y^j is called *agile* if there exist two request sequences τ and τ' in $\mathcal{P}(x^i y^j)$ with $S_{xy}(\tau) = [xy]$ and $S_{xy}(\tau') = [yx]$.

Definition 5. We call a pair of requests $x_{(q)}, y_{(l)}$ an *agile pair* of σ if $x_{(q)}$ and $y_{(l)}$ are adjacent in σ and swapping $x_{(q)}$ and $y_{(l)}$ in σ changes $S_{xy}(\sigma)$.

Clearly, if x^i and y^j are agile, then there exists an agile pair in at least one sequence belonging to $\mathcal{P}(x^i y^j)$.

Lemma 6. If $x_{(q)}, y_{(l)}$ is an agile pair of σ , then x and y are adjacent in $S(\sigma)$.

Proof. Let σ' be σ with $x_{(q)}$ and $y_{(l)}$ swapped. Clearly, $S_{xy}(\sigma) \neq S_{xy}(\sigma')$ and $S_{st}(\sigma) = S_{st}(\sigma')$ holds for all $\{s, t\} \subseteq L$ except $\{x, y\}$. But this is possible only if x and y are adjacent in $S(\sigma)$. \square

Definition 7. For every $x^i \in U$ let $R(x^i)$ be the set defined as follows: $x_{(q)} \in R(x^i)$ if and only if there exists $y_{(l)}$ and σ such that $x_{(q)}, y_{(l)}$ is an agile pair in σ .

Lemma 8. If both pairs x^i, y^j and x^i, z^k are agile, then $|R(x^i)| = 1$.

Proof. Obviously, $R(x^i) > 0$. Suppose that $|R(x^i)| > 1$; we will show that this leads to a contradiction. If $|R(x^i)| > 1$ then there exists $\tau \in \mathcal{P}(x^i y^j)$ with an agile pair $x_{(q)}, y_{(l)}$. Similarly, there exists a sequence $\lambda \in \mathcal{P}(x^i z^k)$ with an agile pair $x_{(q')}, z_{(m)}$ and $q \neq q'$. By suitably inserting k requests to z into τ , we create a sequence σ with $\sigma_{xy} = \tau$ and $\sigma_{xz} = \lambda$ in which both $x_{(q)}, y_{(l)}$ and $x_{(q')}, z_{(m)}$ are adjacent pairs.

By (5), swapping the agile pair $x_{(q)}, y_{(l)}$ will change $S_{xy}(\sigma)$, but leave the relative order of the pairs of items unchanged. The same holds for the agile pair $x_{(q')}, z_{(m)}$ and the pair x, z .

Lemma 6 enforces x to be adjacent to both y and z in $S(\sigma)$. By swapping $x_{(q)}, y_{(l)}$, we obtain a sequence σ' in which $x_{(q')}, z_{(m)}$ is still an agile pair, but x and z are no longer adjacent in $S(\sigma')$, which contradicts Lemma 6. \square

Lemma 9. If $x_{(q)}, y_{(l)}$ is an agile pair in $\lambda \in \mathcal{P}(x^i y^j)$ and $|R(x^i)| = 1$ and $|R(y^j)| = 1$, then swapping $x_{(q)}$ and $y_{(l)}$ changes $S_{xy}(\sigma)$ in any sequence $\sigma \in \mathcal{P}(x^i y^j)$ where $x_{(q)}$ and $y_{(l)}$ are adjacent.

Proof. $|R(x^i)| = 1$ and $|R(y^j)| = 1$ implies that the only swap of requests that can change the relative order of x and y in a request sequence is swapping $x_{(q)}$ and $y_{(l)}$. If the lemma does not hold, there exists a sequence in σ in which we can swap $x_{(q)}$ and $y_{(l)}$ to obtain σ' with $S_{xy}(\sigma) = S_{xy}(\sigma')$. Then we can obtain any sequence in $\mathcal{P}(\sigma)$ by successively transposing adjacent requests, starting from either σ or σ' , without ever swapping $x_{(q)}$ and $y_{(l)}$. Thus, the relative order of x and y would be the same for all request sequences. But we know that swapping $x_{(q)}$ and $y_{(l)}$ changes $S_{xy}(\lambda)$. This is a contradiction. \square

In this and the next section, we consider online list update algorithms that move an item to the front of the list after sufficiently many consecutive requests to that item. This behavior is certainly expected for algorithms with a small competitive ratio. In this section, we show that such algorithms, which we call M -regular, can be characterized in terms of “critical requests”. In the next section, we use this characterization to show that such algorithms are at best 1.6-competitive.

Definition 10. For a given integer $M > 0$, a deterministic algorithm is called M -regular if for each item x and each request sequence σ , item x is in front of all other items after the sequence σx^M .

A randomized algorithm is called M -regular if it is a probability distribution over deterministic M -regular algorithms.

The algorithms discussed at the end of the introduction are all 1-regular or 2-regular. A projective algorithm that is not M -regular is FREQUENCY COUNT, which maintains the items sorted according to decreasing number of past requests; two items which have been requested equally often are ordered by recency of their last request, like in MTF. Hence, after serving the request sequence $x^{M+1}y^M$, item x is still in front of y , which shows that FREQUENCY COUNT is not M -regular for any M . Algorithms that are not M -regular are characterized in Section 5 below, but such “irregular” behavior must vanish in the long run for any algorithm with a good competitive ratio (see Section 6). Hence, the important projective algorithms are M -regular.

The following theorem asserts the existence of critical requests, essentially the unique element of $R(x^i)$ in Lemma 8, for those unary projections x^i where this lemma applies. For projectivity, the list items may also be maintained in reverse order, described as case (b) in the following theorem; competitive algorithms do not behave like this, as we will show later.

Theorem 11. Let A be a deterministic projective algorithm over a set L of list items. Then there exists a function

$$F : U \rightarrow \mathbb{N}, \quad F(x^i) \leq i \quad \text{for all } i$$

so that the following holds:

Let Q be a set of unary projections containing unary projections to at least three different items. Let all unary projections to different items in Q be pairwise agile. Then one of the following two cases (a) or (b) applies.

- (a) For all pairs of unary projections x^i, y^j from Q it holds that if $q = F(x^i)$ and $l = F(y^j)$, then

$$S_{xy}(\sigma) = \begin{cases} [xy] & \text{if } x_{(q)} \text{ is requested after } y_{(l)} \text{ in } \sigma \\ [yx] & \text{if } x_{(q)} \text{ is requested before } y_{(l)} \text{ in } \sigma \end{cases} \quad (8)$$

- (b) For all pairs of unary projections x^i, y^j from Q it holds that if $q = F(x^i)$ and $l = F(y^j)$, then

$$S_{xy}(\sigma) = \begin{cases} [xy] & \text{if } x_{(q)} \text{ is requested before } y_{(l)} \text{ in } \sigma \\ [yx] & \text{if } x_{(q)} \text{ is requested after } y_{(l)} \text{ in } \sigma \end{cases} \quad (9)$$

Proof. Since all pairs of unary projections in Q are pairwise agile, we can conclude $|R(x^i)| = 1$ for all $x^i \in Q$ by Lemma 8. This allows us to define $F(x^i) = q$ if $x_{(q)} \in R(x^i)$. From Lemma 9 we can conclude that for every pair x^i, y^j , either (8) or (9) holds.

It remains to prove that either all pairs are operated by (8) or by (9). If this was not the case, then it would be possible to construct a sequence σ which has a pair of critical requests adjacent to each other in σ without the corresponding items being adjacent in $S(\sigma)$, which contradicts Lemma 6. \square

The following theorem asserts that an M -regular algorithm operates, after M or more requests to a items in a list with at least three items, according to critical requests as in Definition 1. That is, case (b) of Theorem 11, where the list items are arranged backwards, does not apply.

Theorem 12. Let A be a deterministic projective M -regular algorithm over a set L of at least three list items. Then there exists a function

$$F : U \rightarrow \mathbb{N}, \quad F(x^i) \leq i \quad \text{for all } i$$

so that the following holds. Let $x, y \in L$. Let σ be any request sequence with $|\sigma_x| \geq M$ and $|\sigma_y| \geq M$. Then

$$S_{xy}(\sigma) = \begin{cases} [xy] & \text{if } x_{(q)} \text{ is requested after } y_{(l)} \text{ in } \sigma \\ [yx] & \text{if } x_{(q)} \text{ is requested before } y_{(l)} \text{ in } \sigma \end{cases}$$

Proof. Let Q be the set of all unary projections x^i with $i \geq M$. This set has the all the properties of the set Q in Theorem 11, where clearly case (a) applies because $S_{xy}(x^M y^M) = [yx]$. \square

4 The Lower Bound for M -regular Algorithms

In this section, we use Theorem 11 to prove the following result.

Theorem 13. No M -regular projective algorithm is better than 1.6-competitive.

We first give an outline of the proof. Given any $\varepsilon > 0$ and b , we will show that there is a probability distribution π on a finite set Λ of request sequences so that

$$\sum_{\lambda \in \Lambda} \pi(\lambda) \frac{\mathbf{A}(\lambda)}{\text{OPT}(\lambda) + b} \geq 1.6 - \varepsilon, \quad (10)$$

for any deterministic M -regular algorithm \mathbf{A} . Then *Yao's theorem* [13] asserts that also any randomized M -regular algorithm has competitive ratio $1.6 - \varepsilon$ or larger. This holds for any $\varepsilon > 0$, so the competitive ratio is at least 1.6. This ratio is achieved by COMB, and therefore 1.6 is a tight bound for the competitive ratio of M -regular algorithms.

All $\lambda \in \Lambda$ will consist only of requests to two items x and y . In what follows, let $\hat{M} > M$ and $\hat{M} \geq 3$ and let the request sequence ϕ be

$$\phi := x^{\hat{M}} y x^{\hat{M}} y^{\hat{M}} x y^{\hat{M}} x^{\hat{M}} y x y x^{\hat{M}} y^{\hat{M}} x y x y^{\hat{M}}. \quad (11)$$

Let K and T be positive integers and let H be the number of requests to x (and to y) in ϕ , that is,

$$H := |\phi|/2 = 4\hat{M} + 4. \quad (12)$$

Then the set Λ of sequences in (10) is given by

$$\Lambda = \Lambda(K, T) := \{x^{\hat{M}+t} y^{\hat{M}+h} \phi^K \mid 0 \leq h < H, 0 \leq t < HT\}, \quad (13)$$

where π chooses any λ in Λ with equal probability $\pi(\lambda) = 1/H^2T$.

OPT pays exactly ten units for each repetition of ϕ (which always starts in offline list state $[yx]$). Assuming that also the initial list state is $[yx]$, all sequences in Λ have offline cost $10K + 2$. This and the fact that $\pi(\lambda)$ for $\lambda \in \Lambda$ is constant allows us to show (10) once we can prove

$$\sum_{\lambda \in \Lambda} \mathbf{A}(\lambda) \geq 16KH^2T - o(KH^2T). \quad (14)$$

because then

$$\sum_{\lambda \in \Lambda} \pi(\lambda) \frac{\mathbf{A}(\lambda)}{\text{OPT}(\lambda) + b} = \frac{\sum_{\lambda \in \Lambda} \mathbf{A}(\lambda)}{\sum_{\lambda \in \Lambda} (\text{OPT}(\lambda) + b)} \geq \frac{16KH^2T - o(KH^2T)}{(10K + 2 + b)H^2T} \geq 1.6 - \varepsilon$$

for K and T large enough.

Definition 14. A request sequence σ ends at state (i, j) if $|\sigma_x| = i$ and $|\sigma_y| = j$. The request sequence λ passes state (i, j) if there is a proper prefix σ of λ , with $\lambda = \sigma\tau$ for non-empty τ , so that σ ends at (i, j) . The request in λ after (i, j) is the first request in τ .

Definition 15. Let $A_\lambda(i, j)$ denote the online cost of serving the request in λ after (i, j) . If λ does not pass (i, j) , let $A_\lambda(i, j) = 0$.

Definition 16. A state (i, j) is called *good* if for every proper prefix σ of ϕ (that is, $0 \leq |\sigma| < 2H$) there is exactly one $\lambda \in \Lambda$ so that a prefix $x^{\hat{M}+t}y^{\hat{M}+h}\phi^k\sigma$ of λ ends at (i, j) .

After proving that good states incur large costs in Lemma 17, we will only need to prove that almost all states are good to complete the proof of Theorem 13.

Lemma 17. Let (i, j) be a good state. Then

$$\sum_{\lambda \in \Lambda} A_\lambda(i, j) \geq 16.$$

Proof. Consider any $\lambda \in \Lambda$. We can assume that λ passes (i, j) (otherwise $A_\lambda(i, j) = 0$). The request in λ after (i, j) is some request in ϕ . The cost $A_\lambda(i, j)$ of serving that request depends on whether the requested item x or y is in front or not. This, in turn, is determined by the terms $f(x^i)$ and $f(y^j)$ as defined in (3), which determine the relative critical requests to x and y in λ (recall that the item with the more recent critical request is in front).

Because (i, j) is a good state, we obtain exactly all the requests in ϕ as the requests after (i, j) in λ when considering all λ in Λ that pass (i, j) . Therefore, the total cost $\sum_{\lambda \in \Lambda} A_\lambda(i, j)$ is the cost of serving exactly the requests in ϕ according to the critical requests as given by $f(x^i)$ and $f(y^j)$.

$f(x^i)$	$f(y^j)$	$x^{\hat{M}}$	$yx^{\hat{M}}$	$y^{\hat{M}}$	$xy^{\hat{M}}$	$x^{\hat{M}}$	$xyx^{\hat{M}}$	$y^{\hat{M}}$	$xyxy^{\hat{M}}$	$\sum_{\lambda \in \Lambda} A_\lambda(i, j)$
0	0	1..	11..	1..	11..	1..	1111..	1..	1111..	16
0	≥ 1	1..	1..	11..	111..	1..	101..	11..	11011..	≥ 16
1	1	11..	1..	11..	1..	11..	1011..	11..	1011..	16
1	≥ 2	11..	1..	111..	1..	11..	101..	111..	10111..	≥ 18
≥ 2	≥ 2	111..	1..	111..	1..	111..	101..	111..	101..	≥ 18

Table 1: Online costs $A_\lambda(i, j)$ for all λ that pass a good state (i, j) , which are the costs of serving the requests in ϕ . They depend on the relative critical requests $f(x^i)$ and $f(y^j)$.

The rows in Table 1 show the costs $A_\lambda(i, j)$ for the possible combinations of $f(x^i)$ and $f(y^j)$, up to symmetry in x and y . For example, consider the first case $f(x^i) = 0$ and $f(y^j) = 0$, where the critical request to an item is always the most recent request to that item. Suppose that the request after (i, j) is the first request, to x , in the subsequence $xy^{\hat{M}}$ of ϕ . The critical request to x is the last request to x earlier in $yx^{\hat{M}}$, and the critical request to y is the last request to y earlier (and more recent) in $y^{\hat{M}}$. The critical request to y is later than that to x , so y is in front of x , and serving x incurs cost 1, which is the

first 1 in the table entry 11.. in the column for $xy^{\hat{M}}$. The second 1 in 11.. is the cost of serving the first y . It is 1 because here the critical request to x is more recent than the critical request to y . The “.” in 11.. correspond to the costs of later requests to y in $y^{\hat{M}}$, which are zero for $f(x^i) = 0$ and $f(y^j) = 0$ (so for $\hat{M} = 4$ the complete cost sequence would be 11000). In a good state, each cost 0 or 1 in the table (in correspondence to the respective position in ϕ) is incurred by a sequence λ in Λ .

By construction of Λ , the requests before $x^{\hat{M}}$ in the first column of Table 1 are of the form $y^{\hat{M}}$, so y is in front of x , and the first request of $x^{\hat{M}}$ has always cost 1.

In the second row in Table 1, $f(x^i) = 0$ and $f(y^j) \geq 1$. As an illustration of a more complicated case, consider the subsequence $xyxy^{\hat{M}}$ of ϕ in the last column, with associated costs 11011.. The first 1 is the cost of serving the first request to x , because the preceding requests are $\hat{M} \geq M$ requests to y in $y^{\hat{M}}$ and because the algorithm is M -regular, which means $f(y^j) \leq M$, so y is in front of x . Because $f(x^i) = 0$, the cost of serving the first y in $xyxy^{\hat{M}}$ is also 1, because x is in front of y . The second request to x has cost 0 (the first 0 in 11011..) because y is not moved in front of x (the critical request to y is earlier than that to x because $f(y^j) \geq 1$). The next two costs 11 are for the second and third request to y in $xyxy^{\hat{M}}$, because the critical request to x is more recent.

The rows in Table 1 describe all possible cases for $f(x^i)$ and $f(y^j)$, because the costs for requests to x and y apply in the same manner when x and y are interchanged. The respective costs in Table 1 are easily verified. The right column shows that the total cost $\sum_{\lambda \in \Lambda} A_\lambda(i, j)$ is at least 16 in all these cases, which proves the claim. \square

Proof of Theorem 13. We only have to prove (14). Because of Lemma 17, it suffices to show that the number of good states is at least

$$KH^2T - o(KH^2T).$$

By Definition 16, state (i, j) is good if for every proper prefix σ of ϕ there exist unique k, h, t with $0 \leq k < K$, $0 \leq h < H$ and $0 \leq t < HT$ so that $x^{\hat{M}+t}y^{\hat{M}+h}\phi^k\sigma$ ends in (i, j) . This is equivalent to

$$\begin{aligned} i &= \hat{M} + t + kH + |\sigma_x|, \\ j &= \hat{M} + h + kH + |\sigma_y|, \end{aligned}$$

or equivalently

$$\begin{aligned} t &= i + h - j - (|\sigma_x| - |\sigma_y|), \\ h + kH &= j - \hat{M} - |\sigma_y|. \end{aligned} \tag{15}$$

For $0 \leq k < K$ and $0 \leq h < H$, the term $h + kH$ takes the values $0, \dots, KH - 1$. The second equation in (15) therefore has a unique solution in h, k , for any σ (where $0 \leq |\sigma_y| < H$) whenever $\hat{M} + H - 1 \leq j < \hat{M} + KH$. Because $0 \leq |\sigma_x| - |\sigma_y| < H$, the first equation in (15) has a unique solution t in $\{0, \dots, HT - 1\}$ if $j + H - 1 \leq i < j + HT - H$, for every fixed j . Hence the number of good states is at least

$$(KH - H + 1) \cdot (HT - 2H + 1) = KH^2T - o(KH^2T). \quad \square$$

5 The Full Characterization

In this section, we give the full characterization of deterministic projective algorithms. We consider the set U of unary projections of request sequences defined in (2) as the set of nodes of the directed graph $G = (U, E)$ with arcs (x^i, y^j) in E whenever there is a request sequence σ in $\mathcal{P}(x^i y^j)$ with $S(\sigma) = [xy]$.

For any two distinct items x and y and any $i, j \geq 0$, there is at least one arc between x^i and y^j . If the pair x^i, y^j is agile according to Definition 4, then there are arcs in both directions. Only pairs of nodes of the form x^i, x^j do not have arcs between them.

Let \mathcal{W} be the set of strongly connected components of G , and let $C(x^i)$ be the strongly connected component that x^i belongs to. We think of $C(x^i)$ as a “container” that contains x^i and all other unary projections y^j with $C(y^j) = C(x^i)$.

There exists a total order $<$ on these containers so that $C(x^i) < C(y^j)$ if $S_{xy}(\sigma) = [xy]$ after serving any $\sigma \in \mathcal{P}(x^i y^j)$. To see this, we define the following binary relation P on \mathcal{W} : Let $C(x^i) P C(y^j)$ if there is a path in G from x^i to y^j . Then P defines a partial order on \mathcal{W} . The only pairs of containers which are not ordered in P are those of the form $\{x^i\}, \{x^j\}$ for which there does not exist a container $C(y^k)$ with $C(x^i) < C(y^k) < C(x^j)$ or $C(x^j) < C(y^k) < C(x^i)$. By stipulating $\{x^i\} < \{x^j\}$ if and only if $i < j$ for such pairs, we can extend P to the desired total order $<$.

An easy case are the empty unary projections x^0 for items x : Note that x^0 and y^j are never in the same container because there is only a single sequence σ in $\mathcal{P}(x^0 y^j)$ which determines $S_{xy}(\sigma)$ uniquely, so there cannot be paths in both directions between x^0 and y^j in G . Hence $C(x^0) = \{x^0\}$, and $C(x^0) < C(y^0)$ if and only if x is in front of y in the initial list.

In summary, for a request sequence σ , the total order $<$ on \mathcal{W} determines the list order between two items x and y whose unary projections σ_x and σ_y belong to different containers in \mathcal{W} .

If σ_x and σ_y belong to the same container, then the list order between x and y can be described by essentially two possibilities. First, if the container contains only projections to at most two items x and y , nothing further can be said because the algorithm is trivially projective with respect to x and y ; the set of these containers will be denoted by \mathcal{W}_2 .

Second, if a container contains unary projections for three or more distinct items, then the algorithm’s behavior can be described by critical requests similar to Theorem 11; the set of such containers will be denoted by \mathcal{W}^+ . There is a symmetric set \mathcal{W}^- where the algorithm behaves in the same manner but with the list order reversed (which does not define competitive algorithms).

These assertions are summarized in the following theorem.

Theorem 18. Consider a deterministic projective list update algorithm. Then there are pairwise disjoint sets \mathcal{W}^+ , \mathcal{W}^- , \mathcal{W}_2 whose union is \mathcal{W} and a total order $<$ on \mathcal{W} and a function $C : U \rightarrow \mathcal{W}$ with

- (I) $C(x^0) = \{x^0\} \in \mathcal{W}_2$ for all $x \in L$.

(II) for any three items x, y, z , if $C(x^i) = C(y^j) = C(z^k) = w$, then $w \notin \mathcal{W}_2$.

Furthermore, if $C(x^i) \notin \mathcal{W}_2$, then there exists $F(x^i) \in \{1, \dots, i\}$ with the following properties: For all request sequences σ with $\sigma_x = x^i$ and $\sigma_y = y^j$,

(III) if $C(x^i) < C(y^j)$ then $S_{xy}(\sigma) = [xy]$;

(IVa) if $C(x^i) = C(y^j) \in \mathcal{W}^+$ then $S_{xy}(\sigma) = [xy]$ if and only if the $F(x^i)$ th request to x is *after* the $F(y^j)$ th request to y in σ ;

(IVb) if $C(x^i) = C(y^j) \in \mathcal{W}^-$ then $S_{xy}(\sigma) = [xy]$ if and only if the $F(x^i)$ th request to x is *before* the $F(y^j)$ th request to y in σ .

Proof. The set \mathcal{W} and the order $<$ have been defined above with the help of the graph G , which shows (III). We have also shown (I) above.

As before, let \mathcal{W}_2 be the set of containers with unary projections to at most two distinct items, which implies (II).

It remains to show (IVa) and (IVb). Consider a request sequence σ with $\sigma_x = x^i$ and $\sigma_y = y^j$. Let $C(x^i) = C(y^j) \notin \mathcal{W}_2$, so that there is a third item $z \notin \{x, y\}$ with $C(x^i) = C(y^j) = C(z^k)$. We want to apply Lemma 8. To this end, we first show:

$$(x^i, y^j) \in E \quad \text{and} \quad (y^j, z^k) \in E \quad \implies \quad (x^i, z^k) \in E. \quad (16)$$

Let $(x^i, y^j) \in E$, so that $S_{xy}(\sigma) = [xy]$ for some $\sigma \in \mathcal{P}(x^i y^j)$. If $(y^j, z^k) \in E$, then one can insert k requests to z into σ so that $S_{yz}(\sigma) = [yz]$. Adding the requests to z does not change $S_{xy}(\sigma)$, so $S(\sigma) = [xyz]$, which implies $(x^i, z^k) \in E$. This shows (16).

With the help of (16), we now show that if $C(x^i) = C(y^j)$, then the pair x^i, y^j is agile according to Definition 4. We will prove this by showing that

$$(x^i, y^j) \in E \quad \text{and} \quad (y^j, x^i) \in E. \quad (17)$$

To prove (17), recall that $C(x^i)$ is a strongly connected component of the graph G which also contains y^j and z^k . Therefore there exists a path in G from x^i to y^j via z^k . This path is a sequence of unary projections u_1, \dots, u_n with $u_0 = x^i$, $u_l = z^j$ for some $1 < l < n$, and $u_n = y^j$. Let s_i be the item of the corresponding unary projection u^i , in particular $s_0 = x$, $s_l = z$, $s_n = y$.

We call a path $u_1 \dots u_n$ between x^i and y^j *valid* if $|\{s_1, \dots, s_n\}| \geq 3$. We claim that if there exists a valid path between x^i and y^j of length $n > 3$, then there also exists a valid path of length $n - 1$.

To show this claim, consider the smallest q so that s_{q-1} , s_q , and s_{q+1} are three distinct items. Because of (16), clearly $(s_{q-1}, s_{q+1}) \in E$. If the path $u_1 \dots u_n$ remains valid after removing u_q , we are done.

Otherwise, clearly $|\{s_1, \dots, s_n\}| = 3$, and $q = 2$ because otherwise $s_q = s_2$ and one could remove u_2 . We consider the two cases $n = 4$ and $n > 4$. If $n = 4$, then s_1, \dots, s_4 must be of the form a, b, c, s_4 and it is easy to see that $s_4 \notin \{a, b, c\}$. Hence

this case cannot occur. If $n > 4$, then s_1, \dots, s_n is of the form a, b, c, a, c, \dots and one can apply (16) in order to remove u_3 and still get a valid path. This shows the claim.

It follows from the claim that there is a valid path of length two between x^i and y^j . A final application of (16) then gives $(x^i, y^j) \in E$. The same argument shows $(y^j, x^i) \in E$. This proves (17).

Because all pairs of unary projections are agile in $C(x^i)$, we can apply Theorem 11, whose cases (a) and (b) prove (IVa) and (IVb). This proves the theorem. \square

6 The Lower Bound for Irregular Algorithms

In Section 4, we showed that no deterministic M -regular projective list update algorithm can be better than 1.6-competitive. For this we gave, for any $\varepsilon > 0$, a suitable distribution on request sequences that bound the competitive ratio of the algorithm from below by $1.6 - \varepsilon$.

We extend this analysis to arbitrary randomized projective list update algorithms using the full characterization from the previous section.

In brief, the proof works as follows. Using the crucial notion of a good state (i, j) in Definition 16, we call a deterministic algorithm \hat{M} -regular *in state* (i, j) if it fulfills a certain condition, (18) below, where the algorithm only uses the containers from Theorem 18 in the normal way that one expects from competitive algorithms. The lower bound from Lemma 17 applies in expectation for algorithms that fulfill condition (18). If the condition fails, we can give simple request sequences that show that the randomized algorithm is not 1.6-competitive, so the old analysis and the lower bound $1.6 - \varepsilon$ do apply in expectation.

Theorem 19. Any randomized projective list update algorithm that accesses a list of at least three items is at best 1.6-competitive.

Proof. Assume the list has at least three items. Consider a randomized projective algorithm \mathcal{A} and assume that \mathcal{A} is c -competitive with $c < 1.6$. That is, there exists a constant b such that $\mathcal{A}(\sigma) \leq c \cdot \text{OPT}(\sigma) + b$ for all request sequences σ .

We adapt the proof for M -regular algorithms of Section 4. Let $\hat{M} \geq 3$, consider Λ in (13) and consider a good state (i, j) as defined in Definition 16.

Let \mathbf{A} be a deterministic projective algorithm. We say that algorithm \mathbf{A} is \hat{M} -regular *in state* (i, j) if, with \mathcal{W}^+ as in Theorem 18,

$$C(x^i) = C(y^j) \in \mathcal{W}^+, \quad f(x^i) < \hat{M}, \quad f(y^j) < \hat{M}. \quad (18)$$

It is easy to see that the proof of Lemma 17 applies if \mathbf{A} is \hat{M} -regular in (i, j) .

Recall that \mathcal{A} is just a probability distribution on the set of deterministic projective algorithms. Let r_{ij} be the event that \mathcal{A} is \hat{M} -regular in state (i, j) . Then in a good state (i, j) , then the expected cost of \mathcal{A} is bounded as follows:

$$E \left[\sum_{\lambda \in \Lambda} \mathcal{A}_\lambda(i, j) \right] \geq 16 \cdot \text{prob}(r_{ij}). \quad (19)$$

Analogous to the analysis for \hat{M} -regular deterministic algorithms, generalizing inequality (14), the right-hand side of (19) fulfills

$$\sum_{(i,j) \text{ good}} 16 \text{prob}(r_{ij}) \geq KH^2T - o(KH^2T) - \sum_{(i,j) \text{ good}} 16(1 - \text{prob}(r_{ij})). \quad (20)$$

Let $X := H(K + T) + \hat{M}$ and $Y := KH + \hat{M}$. For all good states (i, j) we have

$$1 \leq i \leq X \text{ and } 1 \leq j \leq Y. \quad (21)$$

If we can prove that, with growing \hat{M} , K , and T ,

$$\sum_{(i,j) \text{ good}} 16(1 - \text{prob}(r_{ij})) \leq \sum_{j=1}^Y \sum_{i=1}^X 16(1 - \text{prob}(r_{ij})) = o(KH^2T), \quad (22)$$

then we have proved (14) for irregular algorithms. This can be done by analyzing where (18) fails, that is, for each of the six cases according to

$$\sum_{j=1}^Y \sum_{i=1}^X 1 - \text{prob}(r_{ij}) \leq \sum_{j=1}^Y \sum_{i=1}^X \left(\begin{array}{l} \text{prob}(C(x^i) < C(y^j)) \\ + \text{prob}(C(x^i) > C(y^j)) \\ + \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^-) \\ + \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}_2) \\ + \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(x^i) \geq \hat{M}) \\ + \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(y^j) \geq \hat{M}) \end{array} \right).$$

We start by proving

$$\sum_{j=1}^Y \sum_{i=1}^X \text{prob}(C(x^i) < C(y^j)) \leq o(KH^2T). \quad (23)$$

To this aim, consider the sequence $x^i y^Y$ for $1 \leq i \leq X$. When serving this sequence, a request to y will be served in each of the states $(i, 1), (i, 2), \dots, (i, Y)$. Since every deterministic algorithm with $C(x^i) < C(y^j)$ pays one unit for accessing y in state (i, j) , the expected cost of \mathcal{A} for serving a request to y in a state (i, j) is at least $\text{prob}(C(x^i) < C(y^j))$. Therefore

$$\mathcal{A}(x^i y^Y) \geq \sum_{j=1}^Y \text{prob}(C(x^i) < C(y^j)). \quad (24)$$

On the other hand, it must hold that $\mathcal{A}(x^i y^Y) \leq c \cdot \text{OPT}(x^i y^Y) + b$ because \mathcal{A} is c -competitive. Since $\text{OPT}(x^i y^Y) = 1$ it follows that

$$\sum_{i=1}^X \sum_{j=1}^Y \text{prob}(C(x^i) < C(y^j)) \leq \sum_{i=1}^X \mathcal{A}(x^i y^Y) \leq X \cdot (c + b) = o(KH^2T). \quad (25)$$

The bound on $\text{prob}(C(x^i) > C(y^j))$ is very similar.

For $\text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^-)$, we also use request sequences of the form $\sigma = x^i y^Y$. Clearly, from the first request to y on, the critical requests to x is always earlier in σ than the critical request to y . Therefore $C(x^i) = C(y^j) \in \mathcal{W}^-$ implies that y is behind x in the list.

If $C(x^i) = C(y^j) \in \mathcal{W}_2$, the container $C(x^i)$ does not contain any unary projections to items other than x or y . The list has at least a third item z and either $C(x^i) < C(z^k)$ or $C(z^k) < C(x^i)$ for any k . We consider only the first case, where we can bound $\text{prob}(C(x^i) < C(z^k))$ similarly to (23). By considering the request sequence $x^i z^Y$ for $1 \leq i \leq X$, we obtain in the same way as with (24) and (25) that $\sum_{i=1}^X \mathcal{A}(x^i z^Y) = o(KH^2T)$.

We have

$$\text{prob}(C(x^i) = C(y^j) \in \mathcal{W}_2) \leq \text{prob}(C(x^i) < C(z^k)) + \text{prob}(C(x^i) > C(z^k))$$

for all z^k with $z \neq x, y$. Hence the left hand side can be bounded by the bound on the first two cases.

In a similar fashion, we bound $\text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(x^i) \geq \hat{M})$:

$$\begin{aligned} & \sum_{j=1}^Y \sum_{i=1}^X \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(x^i) \geq \hat{M}) \\ = & \sum_{j=1}^Y \sum_{i=1}^{\hat{M}-1} \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(x^i) \geq \hat{M}) + \\ & \sum_{j=1}^Y \sum_{i=\hat{M}}^X \text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(x^i) \geq \hat{M}) \\ \leq & 0 + \sum_{j=1}^Y \frac{1}{\hat{M}} \sum_{i=1}^X \sum_{\ell=0}^{\hat{M}-1} \text{prob}(C(x^{i+\ell}) = C(y^j) \in \mathcal{W}^+, f(x^{i+\ell}) \geq \hat{M}) \\ \leq & \sum_{j=1}^Y \frac{1}{\hat{M}} \sum_{i=1}^X \mathcal{A}(x^i y^j x^{\hat{M}}) \\ \leq & Y \frac{1}{\hat{M}} X (c \cdot \text{OPT}(x^i y^j x^{\hat{M}}) + b) = o(KH^2T). \end{aligned}$$

The bound on $\text{prob}(C(x^i) = C(y^j) \in \mathcal{W}^+, f(y^j) \geq \hat{M})$ is similar to the previous bound. \square

7 Conclusion

An open problem is to extend the lower bound to the full cost model, even though this model is not very natural in connection with projective algorithms. This would require request sequences over arbitrarily many items, and it is not clear whether an approach similar to the one given here can work.

Another ambitious goal is to further improve the lower bound in case of non-projective algorithms. Here, the techniques of the paper do not apply at all, and to get improvements that are substantially larger than the ones obtainable with the methods of [6] requires new insights.

Finally, the search for good non-projective algorithms has become an issue with our result. Irani's SPLIT algorithm [9] is the only one known of this kind with a competitive ratio below 2. A major obstacle for finding such algorithms is the difficulty of their analysis, because pairwise methods are not applicable, and other methods (e.g. the potential function method) have not been studied in depth. We hope that our result can stimulate further research in this direction.

References

- [1] S. Albers. Improved randomized on-line algorithms for the list update problem. *SIAM Journal on Computing*, 27(3):682–693, 1998.
- [2] S. Albers, B. von Stengel, and R. Werchner. A combined BIT and TIMESTAMP algorithm for the list update problem. *Information Processing Letters*, 56(3):135–139, 1995.
- [3] S. Albers and J. Westbrook. Self-organizing data structures. In A. Fiat and G. J. Woeginger, editors, *Online Algorithms*, volume 1442 of *Lecture Notes in Computer Science*, pages 13–51. Springer, 1998.
- [4] C. Ambühl. *On the List Update Problem*. PhD thesis, ETH Zürich, 2002.
- [5] C. Ambühl, B. Gärtner, and B. von Stengel. Optimal projective algorithms for the list update problem. In *Proceedings of 27th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 305–316, 2000.
- [6] C. Ambühl, B. Gärtner, and B. von Stengel. A new lower bound for the list update problem in the partial cost model. *Theoretical Computer Science*, 268(1):3–16, 2001.
- [7] J. L. Bentley and C. C. McGeoch. Amortized analyses of self-organizing sequential search heuristics. *Communications of the ACM*, 28(4):404–411, 1985.
- [8] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, New York, NY, USA, 1998.

- [9] S. Irani. Two results on the list update problem. *Information Processing Letters*, 38(6):301–306, 1991. Corrected version appeared as Technical Report 96-53, ICS Department, U.C. Irvine, CA, USA 1996.
- [10] N. Reingold, J. Westbrook, and D. D. Sleator. Randomized competitive algorithms for the list update problem. *Algorithmica*, 11(1):15–32, 1994.
- [11] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.
- [12] B. Teia. A lower bound for randomized list update algorithms. *Information Processing Letters*, 47(1):5–9, 1993.
- [13] A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *Proceedings of 19th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.