Competitive analysis of randomized paging algorithms

Dimitris Achlioptas [a,1], Marek Chrobak [b,2], John Noga [c,*,2]

[a] *Department of Computer Science, University of Toronto, Toronto, Ont., Canada M5S 3G4*
[b] *Department of Computer Science, University of California, Riverside, CA 92521, USA*
[c] *Department of Mathematics, University of California, Riverside, CA 92521, USA*

Abstract

The paging problem is defined as follows: we are given a two-level memory system, in which one level is a fast memory, called *cache*, capable of holding k items, and the second level is an unbounded but slow memory. At each given time step, a request to an item is issued. Given a request to an item p, a *miss* occurs if p is not present in the fast memory. In response to a miss, we need to choose an item q in the cache and replace it by p. The choice of q needs to be made *on-line*, without the knowledge of future requests. The objective is to design a replacement strategy with a small number of misses. In this paper we use competitive analysis to study the performance of randomized on-line paging algorithms. Our goal is to show how the concept of work functions, used previously mostly for the analysis of deterministic algorithms, can also be applied, in a systematic fashion, to the randomized case. We present two results: we first show that the competitive ratio of the marking algorithm is exactly $2H_k - 1$. Previously, it was known to be between H_k and $2H_k$. Then we provide a new, H_k-competitive algorithm for paging. Our algorithm, as well as its analysis, is simpler than the known algorithm by McGeoch and Sleator. Another advantage of our algorithm is that it can be implemented with complexity bounds independent of the number of past requests: $O(k^2 \log k)$ memory and $O(k^2)$ time per request. © 2000 Elsevier Science B.V. All rights reserved

Keywords: On-line algorithms; Analysis of algorithms; Competitive analysis; Paging; Randomized algorithms.

1. Introduction

The paging problem is defined as follows: we are given a two-level memory system, in which one level is a fast memory (that we refer to as *cache*) capable of holding k items, and the second level is an unbounded but slow memory. At each given time

---

* Corresponding author.
E-mail address: jnoga@cs.ucr.edu (J. Noga)

step, a request to an item is issued. Given a request to an item p, a *miss* occurs if p is not present in the cache. In response to a miss, we need to move p from the slow memory into the cache. In order to make room for p, one of the items residing currently in the cache, say q, needs to be evicted. The choice of q is made *on-line*, i.e., before the next request is issued, and a strategy for making such choices will be referred to as an *on-line algorithm*.

The cost function associated with the paging problem is the number of misses. In general, no on-line algorithm can achieve a minimum cost on all request sequences. Therefore, in order to evaluate various on-line algorithms, one needs to design a performance measure that takes into account the on-line nature of the problem. In this paper we use the competitive analysis approach: a given on-line algorithm is said to be *c-competitive*, if on every request sequence its cost is bounded (asymptotically) by c times the optimal cost for this sequence.

Paging is a classical on-line problem and has been extensively studied in the literature on competitive on-line algorithms. It can be viewed as a special case of the k-server problem (see, for example, [12, 13, 17, 3]), in which all distances are equal to one. For the deterministic case, it has been established that the well-known LRU (least recently used) strategy is k-competitive, and that no better competitiveness is possible (see [20]).

In this paper we concentrate on the randomized version of the paging problem. It is relatively easy to show (see [7]) that no randomized on-line algorithm can be better than H_k-competitive, where $H_k = \sum_{i=1}^{k} 1/i$ is the kth harmonic number. Two algorithms have been proposed for this problem in the past. Fiat et al. [7] gave a simple marking algorithm, called MARK, and proved that it is $2H_k$-competitive. Subsequently, McGeoch and Sleator [18] presented another algorithm, called PARTITION, and proved that it is H_k-competitive, and thus optimal.

Work functions have played an important role in the analysis of on-line problems. However, in the past, work functions have been used mostly in the analysis of deterministic algorithms, and only recently they have been applied to the randomized case (see [6, 15, 16, 8]). For example, in [6], some optimal randomized algorithms for the page migration problem were developed based on this technique. One of our goals is to show how work functions can also be applied in the competitive analysis of randomized algorithms for the paging problem. A work function approach usually starts with some characterization of work functions for a given problem. Then, the properties of work functions are used to design and analyze an on-line algorithm. Koutsoupias and Papadimitriou [11] presented a simple, elegant characterization of work functions for paging. In an earlier work, McGeoch and Sleator in [18] gave an equivalent characterization of the behavior of an optimal algorithm, although their formulation did not explicitly involve work functions.

1.1. Analysis of MARK

Algorithm MARK of [7], even though not optimal, is of independent interest. It is simpler, faster and more space-efficient than PARTITION. MARK can be implemented using

O(k) memory and O(1) time per request, while PARTITION may need as much as $\Omega(n)$ memory, where n is the number of past requests. Fiat et al. [7] provided an upper bound of $2H_k$ on the competitive ratio of MARK. The general lower bound is H_k. Thus, it would be interesting to know the exact competitiveness of this algorithm. Our first result is a proof that the competitive ratio of MARK is $2H_k - 1$.

1.2. A new optimal algorithm

The result of McGeoch and Sleator [18] gives a tight bound on the optimal competitive ratio in the randomized case. However, the algorithm PARTITION from [18] is somewhat counter-intuitive, and both the correctness and competitiveness proofs are rather difficult. In the second part of the paper, we present an alternative H_k-competitive algorithm for paging, called EQUITABLE, that we believe is simpler and more natural. In fact, our algorithm follows naturally from the concepts of work functions, stable algorithms, and some basic game-theoretic principles. We believe that these ideas can be extended to other on-line problems. We also show that our algorithm, can be implemented in space O($k^2 \log k$) and time O(k^2) per request. Unlike those for PARTITION, these bounds are independent of the length of the request sequence.

2. Preliminaries

Throughout the paper, by k we denote the cache size. By a *cache configuration*, or simply *configuration*, we will mean a k-tuple of items representing the cache content. We will assume that the initial configuration is fixed, and call it X_0. As explained in the introduction, an on-line paging algorithm $\mathscr{A}$ needs to respond to every request p before the next request is issued. If a miss occurs, that is, if p is not in the cache, $\mathscr{A}$ must decide which item q should be evicted from the cache to make room for p. Each miss has unit cost.

Mathematically, it is convenient to define an on-line algorithm as a function $\mathscr{A}(\varrho)$ that to a given request sequence ϱ assigns the configuration after serving ϱ. In order to ensure that the requests are satisfied, if r is the last request in ϱ then we require that $r \in \mathscr{A}(\varrho)$. Note that this definition allows $\mathscr{A}$ to swap an arbitrary number of items in the cache at each step, whether a miss occurred or not. However, we will charge $\mathscr{A}$ a cost of 1 for each such swap. It is easy to see that in this case, without loss of generality, $\mathscr{A}$ will never bring an item into the cache unless it is the current request.

Denote by $cost_{\mathscr{A}}(\varrho)$ the cost of $\mathscr{A}$ on request sequence ϱ, and by $opt(\varrho)$ the optimal cost on ϱ. We will say that $\mathscr{A}$ is *c-competitive* if there is a constant a such that on every request sequence ϱ

$$cost_{\mathscr{A}}(\varrho) \leqslant c\, opt(\varrho) + a. \tag{1}$$

In our algorithms the additive constant a will be zero. The *competitive ratio* of $\mathscr{A}$ is the minimum c for which $\mathscr{A}$ is c-competitive. (In our applications this minimum is well-defined.)

There are several ways to define a randomized algorithm. One can define a randomized algorithm as a probability distribution on the set of all deterministic algorithms for a given problem. Another way is to view a randomized algorithm as an algorithm that at every step chooses its move from a probability distribution on the set of possible moves. In the theory of multi-stage games these two approaches are sometimes called, respectively, *mixed strategies* and *behavior strategies* (see, for example, [14]). The definitions of cost and competitiveness extend naturally to randomized algorithms, independently of which of the two above definitions is being used. If $\mathscr{A}$ is a randomized algorithm then $cost_{\mathscr{A}}(\varrho)$ denotes the expected cost of $\mathscr{A}$ on ϱ, and inequality (1) remains unchanged. It is quite easy to show that these approaches are equivalent, in the sense that an algorithm of each type can be transformed into one of the other type without increasing its (expected) cost on any request sequence.

Yet another way is to consider an algorithm that at each request chooses, deterministically, its probability distribution on the configuration set. In that case $\mathscr{A}(\varrho)$ is a probability distribution on the set of all possible configurations. We call this a *distribution-based* algorithm. The cost of a move can be defined by a so-called transport distance between the distributions. It can be shown that this approach is equivalent to the other two (see [4, 6]).

Algorithm MARK is defined as a behavior algorithm, while our H_k-competitive algorithm EQUITABLE is easier to define using the distribution-based approach. However, for the sake of completeness and for cost estimation, we also show how to "implement" EQUITABLE as a behavior algorithm.

In analyzing the competitiveness of on-line algorithms it is often useful to know the optimal solution for each request sequence. A *work function* is a function from the set of possible configurations to real numbers which gives the optimal cost to serve a sequence of requests and end in a particular configuration. Specifically, the work function ω associated with a sequence of requests ϱ is defined as follows: $\omega(X)$ is the minimum cost of servicing ϱ, starting from the initial configuration X_0 and ending in X. Note that we do not insist that the last request r belongs to X. In such a case, define $\omega(X) = 1 + \min_{x \in X} \omega(X + r - x)$. In other words, we allow an optimal algorithm to swap r out of the cache after the request has been satisfied and before the next request was issued. The optimal cost of servicing a request sequence with associated work function ω is simply $\min(\omega)$.

Suppose that the current work function is ω and r is the new request. What is the new work function, after serving r? Denote this new, updated, work function by $\omega \wedge r$. It is straightforward to see that this function is

$$\omega \wedge r(X) = \begin{cases} \omega(X) & \text{if } r \in X, \\ 1 + \min_{x \in X} \omega(X + r - x) & \text{if } r \notin X. \end{cases}$$

With each work function ω we can associate an *offset function* $\omega - \min(\omega)$. Instead of keeping track of work functions, it is more convenient to deal with offset functions. If ω is the current work function and μ is the current offset function then by μ^r we

will denote the offset function after request r, that is $\mu^r = \omega \wedge r - \min(\omega \wedge r)$. The value $\min(\omega \wedge r) - \min(\omega) = \min(\mu \wedge r)$ is the optimal cost associated with serving the request for r.

Let K be a set of configurations and ω a work function. We say that ω is *coned-up* from K, if for every configuration X there is a $Y \in K$ such that $\omega(X) = \omega(Y) + |X - Y|$. In other words, the value of ω on all configurations is uniquely determined by its values in K. If K is a singleton, $K = \{X\}$, then we call ω a *cone on* X.

2.1. Potential argument

In order to prove competitiveness of a given algorithm $\mathscr{A}$, we will use amortized analysis. A potential function Φ assigns a real number to a given offset function and a configuration of $\mathscr{A}$. (To simplify notation, throughout the paper we will omit the function arguments when defining potentials.) We consider each move separately. Each move is described by a current offset function ω, a configuration of $\mathscr{A}$, and a request r. Denote by $\Delta cost_{\mathscr{A}}$, Δopt and $\Delta \Phi$, the cost of $\mathscr{A}$, the optimal cost, and the change of the potential in this move. Our goal is to show that for every move

$$\Delta cost_{\mathscr{A}} + \Delta \Phi \leqslant c \, \Delta opt. \tag{2}$$

Inequality (2) implies c-competitiveness of $\mathscr{A}$ by simple summation over the whole request sequence. This method works for randomized algorithms as well, the only difference being that the potential now depends on the current offset function and the distribution of $\mathscr{A}$, while $\Delta cost_{\mathscr{A}}$ is the expected cost of $\mathscr{A}$ in the given move.

2.2. Characterization of work functions

Koutsoupias and Papadimitriou [11] gave the following elegant characterization of the work functions for the paging problem with the cache of size k.

Lemma 1. *Every offset function is coned up from the set of configurations for which its value is zero. Moreover, if ω is the current offset function and r is the last request then there is a sequence of sets $L_1, L_2, \ldots, L_k$, with $L_1 = \{r\}$, such that $\omega(X) = 0$ iff $|X \cap \bigcup_{i \leqslant j} L_i| \geqslant j$ for all $1 \leqslant j \leqslant k$.*

The sets L_i are called the *layers* of ω. Note that the above representation is not always unique. Configurations for which the value of ω is zero will be called *valid* and the collection of all valid configurations will be denoted by $V(\omega)$.

The set $S(\omega) = \bigcup_{i \leqslant k} L_i$ will be called the *support* of ω. By Lemma 1, we can identify ω with its sequence of layers, and write $\omega = (L_1 | L_2 | \ldots | L_k)$. We always have $|L_1| = 1$. Let i be the largest number for which $L_1, \ldots, L_i$ are singletons. All items in $L_1 \cup \cdots \cup L_i$ are called *revealed*. We can assume that an optimal algorithm has all revealed items in the cache, since any optimal algorithm can be modified to one with this property without raising its cost. By $N(\omega)$ we denote the set of non-revealed items in $S(\omega)$.

Koutsoupias and Papadimitriou [11] also give a method for updating the layers after a request. Let $\omega = (L_1 | \ldots | L_k)$. Suppose that r is a new request. Then

$$
\omega^r = \begin{cases}
(r|L_1|\ldots|L_{j-1}|L_j \cup L_{j+1} - r|L_{j+2}|\ldots|L_k) & \text{if } r \in L_j \text{ and } j < k, \\[2mm]
(r|L_1|\ldots|\ldots|L_{k-1}) & \text{if } r \in L_k, \\[2mm]
(r|L_1 \cup L_2|L_3|\ldots|L_k). & \text{if } r \notin S(\omega).
\end{cases}
\tag{3}
$$

Furthermore, the optimal cost associated with this request is 0 if $r \in S(\omega)$ and 1 if $r \notin S(\omega)$.

To simplify notation, in the equations above and throughout the paper, we will often omit braces in the notation for sets. We will also write $X + x$ (or $X - x$) when adding an item x to a set X (respectively, removing x from X). If $Z \subseteq X$ for some $X \in V(\omega)$, then we will also use notation ω^Z for $\omega^{z_1 \ldots z_p}$, where $z_1 \ldots z_p$ is an arbitrary permutation of Z. (It is easy to see that this is well-defined.)

Example. Let $k = 3$, and suppose that originally items a, b, and c are in the cache. So the initial configuration is $X_0 = \{a, b, c\}$ and the initial work function is the cone on X_0, represented by $(a|b|c)$ (the order of a, b, c in this representation is arbitrary). Consider the request sequence d, e, b. The corresponding sequence of offset functions is

$$
(a|b|c) \;\to\; (d|a,b|c) \;\to\; (e|a,b,d|c) \;\to\; (b|e|a,c,d).
$$

The optimal cost of this sequence is 2.

3. Analysis of MARK

Algorithm MARK maintains up to k marks on the items that are in the cache. At the beginning, MARK initializes all items in the cache as marked. Suppose that an item p is requested. If p is in the cache, we mark it, unless it is already marked. If p is not in the cache, we first check if all items in the cache are already marked, and if so, we unmark them all. Then we mark p and swap it with a random, uniformly chosen, unmarked item in the cache.

We divide the execution into *phases*, where each phase starts when k items are marked and a non-marked item is about to be requested. Without loss of generality, we assume that the last phase is complete, since we can always complete it by requesting k items in the support, without increasing the optimal cost. Thus, each phase consists of requests to exactly k distinct items. Items that were marked when the phase was about to start and are now unmarked will be called *active*. The definition of MARK and the rules for updating the offset functions (3) imply:

Fact 1. *The cache contains only marked items and active items. All marked items are in the cache. If there are m marked items and v active items then each active item is in the cache with probability $(k - m)/v$.*

Fact 2. *Let $\omega = (L_1 | \ldots | L_k)$. If L_i contains a marked item then all items in $\bigcup_{j<i} L_j$ are marked.*

Theorem 1. *The competitive ratio of algorithm* MARK *is* $2H_k - 1$.

Proof. (*Lower bound*) We will present a cycle of requests on $k + 2$ items where the optimal cost is 1, while the cost of MARK is $2H_k - 1$.

Start from a state where the offset function is $\omega = (x_1 | x_2 | \ldots | x_k)$, a cone on $X = \{x_1, \ldots, x_k\}$, there is one marked item x_1 and k active items $x_2, \ldots, x_k, y$, where $y \notin X$. We first request an item $x_0 \notin X + y$. Now $\omega^{x_0} = (x_0 | x_1, x_2 | x_3 | \ldots | x_k)$ and x_0, x_1 are marked. Then we request $x_2, \ldots, x_{k-1}$, reaching an offset function $(x_{k-1} | \ldots | x_2 | x_0 | x_1, x_k)$, where $x_0, \ldots, x_{k-1}$ are marked. Finally, we request x_k, going back to a state equivalent to the one we started from. The optimal cost is 1 (x_0 replaces x_1), while the cost of MARK is

$$1 + \frac{2}{k} + \frac{2}{k-1} + \cdots \frac{2}{3} + \frac{2}{2} = 2H_k - 1.$$

To complete the proof it remains to show that the state from which we started the cycle can be reached from the initial state. The initial state is a cone on the items originally in the cache. Let these points be $\{x_1, x_2, \ldots, x_{k-1}, z\}$. By requesting $y, x_k, x_{k-1}, \ldots, x_1$, we reach the desired state. $\square$

(*Upper bound*). We will prove that MARK is $(2H_k - 1)$-competitive, using a potential function argument. If the current offset function has s layers that contain unmarked items then the potential is

$$\Phi = s(H_k - H_s + 1),$$

where we will take $s(H_k - H_s + 1) = 0$, for $s = 0$.

For an arbitrary phase let $\Delta cost$, Δopt and $\Delta \Phi$ denote MARK's cost, the optimal cost and the potential change, respectively, during this phase. Our goal is to prove that $\Delta cost + \Delta \Phi \leqslant (2H_k - 1)\Delta opt$. This will imply the $(2H_k - 1)$-competitiveness of MARK, as explained in the previous section. We will classify each request in one of three types depending on whether the requested item was (a) outside the support, (b) non-active and in the support, (c) active and in the support. Assume that there are t requests of type (a), l requests of type (b) and that initially there are s layers in the offset function containing unmarked items.

Each request of type (b) must come from one of the layers with nonmarked items and, because of Fact 2, it reduces the number of such layers by one. Therefore $l \leqslant s$.

After the first request the first layer consists of one marked item and the other $k - 1$ layers contain unmarked items. In the rest of the phase, depending on whether the first request is of type (a) or (b), there will be either $k - t$ or $k - 1 - t$ requests for unmarked items in the support. For each such request the number of layers containing unmarked items decreases by one. Hence, the offset function at the beginning of the next phase

will have s' layers that contain unmarked items, where $s' \leqslant k - 1 - (k - 1 - t) = t$. For $l = 0$ the first request must be of type (a), implying $s' \leqslant t - 1$.

Now we estimate $\Delta cost$. Since MARK faults on any item at most once in a phase, items of type (a) and (b) contribute at most $l + t$ to $\Delta cost$. Consider now the ith request of type (c). At the time of this request there are at most $l + t + i - 1$ marked items and exactly $k - i + 1$ active items. From Fact 1, the probability that MARK faults on this request is at most $(l+t)/(k-i+1)$. Since the range of i is from 1 to $k - l - t$, items of type (c) contribute at most $(l + t)(H_k - H_{l+t})$ to $\Delta cost$.

Clearly, $\Delta opt = t$. Now, we proceed as follows:

$$\Delta cost + \Delta \Phi \leqslant (l + t)(H_k - H_{l+t} + 1) + s'(H_k - H_{s'} + 1) - s(H_k - H_s + 1) \quad (4)$$

$$\leqslant (l + t)(H_k - H_{l+t} + 1) + s'H_k - s(H_k - H_s + 1)$$

$$\leqslant (l + t)(H_k - H_{l+t} + 1) + s'H_k - l(H_k - H_l + 1) \quad (5)$$

$$= (2H_k - 1)t + 2t - (l + t)H_{l+t} + lH_l - (t - s')H_k$$

$$\leqslant (2H_k - 1)t \quad (6)$$

$$= (2H_k - 1)\Delta opt.$$

Inequality (4) comes from our bound for $\Delta cost$ and the definition of Φ, s, s'; (5) holds because $l \geqslant s$ and Φ is increasing in s. For $l > 0$, inequality (6) holds because then $2t - (l + t)H_{l+t} \leqslant -lH_l$ and $s' \leqslant t$, and, for $l = 0$, (6) holds because in this case $s' \leqslant t - 1$ implying $2t - tH_t - (t - s')H_k \leqslant 0$ $(t, k \geqslant 1)$. $\square$

Since the initial and final potentials are 0, we can assume that the additive constant from (1) is 0.

4. A New optimal algorithm

A randomized algorithm is said to be *stable* if its probability distribution at each step does not depend upon anything other than the current offset function. Thus, a stable algorithm can be described as a function $\mathscr{A}(\omega)$ giving the distribution of $\mathscr{A}$ if the current offset function is ω. Naturally, for stable randomized algorithms the potential function will depend only on the current offset function.

PARTITION, the algorithm proposed by McGeoch and Sleator [18], can be classified as distribution-based and stable. Its probability distribution is defined by the following k-round tournament: Initially, each $x \in L_i$ is given rank $i + 1$. At step $i = k, k - 1, \ldots, 1$, pick uniformly $k - i + 1$ items from among those that have rank $i + 1$, and decrease their rank to i. At the end, the k winners of rank 1 are chosen as the configuration.

Call an arbitrary ranking valid if (a) each item of rank $i \geqslant 2$ belongs to $L_{i-1} \cup \cdots \cup L_k$, and (b) there are exactly $k - i + 1$ items of rank $\leqslant i$ in $L_i \cup \cdots \cup L_k$. Mc-Geoch and Sleator prove that the above tournament generates the uniform probability distribution on all valid rankings, and use this fact to prove that PARTITION is H_k-competitive.

Our approach is different. Assume that we are looking for a stable algorithm. Thus, we need to specify a probability distribution $P(\omega)$ on the set of configurations used by this algorithm when the current offset function is ω. How can we derive this distribution? Suppose that the requests are generated by a so-called *lazy adversary*, that is, one that starting at ω only makes requests that do not increase the optimal cost. In other words, the requests are in the support of ω. We can assume that the adversary never requests the revealed items, and thus each lazy sequence consists of at most $k-1$ requests, and the final offset function is a cone. We refer to such request sequences as adversary *lazy strategies*.

What is the probability distribution that would best protect us against such an adversary? We would expect that for the ideal distribution all such lazy strategies have the same on-line cost. (We leave it to the reader to verify that PARTITION does not satisfy this property for $k \geqslant 3$.) Now, view the problem as a two-person zero-sum game. Let our player determine his configuration. Let the adversary determine his lazy strategy, item by item. The value of a game is the same as the outcome of an optimal mixed strategy against a pure (deterministic) strategy which appears with nonzero probability in some optimal mixed strategy. Using this fact, it is not hard to see that the adversary can pick the first request uniformly from the support. It seems reasonable that the optimal distribution for our player would match the optimal adversary strategy.

Algorithm EQUITABLE. The algorithm is defined by a probability distribution $P_X(\omega)$ of being in a configuration X when the current offset function is ω. Select X as follows: Let $X = \emptyset$. While $|X| < k$ select an item x uniformly from $S(\omega^X) - X$ and add this item to X. $P_X(\omega)$ is the probability of selecting X using the above random process.

In other words, $P_X(\omega)$ is the probability of reaching the cone on X if, starting at ω, at each step a request is chosen uniformly from the support. For simplicity, we will use the same notation $P(\omega)$ for the distribution of EQUITABLE and for the random process that generates this distribution. It is easy to see that only valid configurations have positive probability of being used by EQUITABLE.

Denote the probability that EQUITABLE has x in the cache when the current offset function is ω by $P_x(\omega)$. Obviously, $P_x(\omega) = 0$ for $x \notin S(\omega)$ and $P_x(\omega) = 1$ for $x \in S(\omega) - N(\omega)$. Let M be any set such that $N(\omega) \subseteq M \subseteq S(\omega)$. We can as well assume $P(\omega)$ initializes X to $S(\omega) - M$ instead of $\emptyset$. Therefore

$$P_x(\omega) = \frac{1}{|M|} \sum_{z \in M} P_x(\omega^z) . \tag{7}$$

Lemma 2. *If the current offset function is ω, and the request is r, then the cost of* EQUITABLE *on this request is $1 - P_r(\omega)$, the probability that r is not in the cache.*

Proof. It is sufficient to show how EQUITABLE can be "implemented" as a behavior algorithm $\mathscr{B}$ that has the following properties: (i) if a requested item is in the cache, then $\mathscr{B}$ does not move, and (ii) if a request is not in the cache, then $\mathscr{B}$ only swaps one item. By "implement" we mean that at each step $\mathscr{B}$ induces the same probability distribution as EQUITABLE.

Given an offset function ω, a configuration X of $\mathscr{B}$, and a request r, we want to define how $\mathscr{B}$ will serve r. $\mathscr{B}$ orders X randomly as follows: a permutation $x_1 x_2 \ldots x_k$ of X is chosen with the same probability that this would be the order the items of X were chosen by $P(\omega)$, given that the set X was chosen. Let j be the largest integer such that $X + r - x_j \in V(\omega^r)$. Then $\mathscr{B}$ replaces x_j in the cache by r.

Note that if $r \in X$ then $x_j = r$ and no actual swap is needed. If $r \notin S(\omega)$ then we have $j = k$. In the last case, if $r \in S(\omega) - X$, let $X_i = \{x_1, \ldots, x_i\}$ for all i, and notice that there is a unique p such that $r \in S(\omega^{X_p-1}) - S(\omega^{X_p})$. Equivalently, both x_p and r are in the last layer of ω^{X_p-1}. By the choice of p, we have $X - x_p + r \in V(\omega^r)$ and $X - x_i + r \notin V(\omega^r)$ for $i > p$. Therefore $j = p$, and we conclude that j is the index for which x_j and r are in the last layer of ω^{X_j-1}. Additionally, we have $x_i \in S(\omega^{r+X_{i-1}}) - X_{i-1} - r$ for all $i < j$.

$\mathscr{B}$ clearly satisfies conditions (i) and (ii) above, and it generates only configurations in $V(\omega^r)$, so it remains to show that it induces the same distribution on $V(\omega^r)$ as EQUITABLE. The proof is by induction on the number of requests. It is certainly true in the initial state, so assume that it holds for some offset function ω. We will show that it also holds for ω^r. We can assume that when $P(\omega^r)$ chooses its configuration Y, it initializes $Y = \{r\}$ and then it selects the remaining $k - 1$ items $y_1, \ldots, y_{k-1}$ of Y. Denote $Y_i = \{y_1, \ldots, y_i\}$ for each i. Recall that each y_{i+1} is uniformly distributed in $S(\omega^{Y_i+r}) - Y_i - r$.

We can also think of $\mathscr{B}$ as generating a sequence of $k - 1$ items other than r, namely the $x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_k$. For convenience, rename this sequence $z_1, \ldots, z_{k-1}$. Let $Z_i = \{z_1, \ldots, z_i\}$ for each $i \leqslant k-1$, and $Z = Z_{k-1} + r$. We want to show that random variables Y and Z have the same distribution. In order to do so, it is sufficient to prove the following claim: for each i, z_{i+1} is uniformly distributed in $S(\omega^{Z_i+r}) - Z_i - r$. We break the proof into two cases.

Case A: $r \notin S(\omega)$. Then $j = k$, and $z_i = x_i$ for each $i \leqslant k-1$. Then, since $S(\omega^{r+X_i}) - X_i - r = S(\omega^{X_i}) - X_i$ for $i \leqslant k - 2$, both processes make the same random choices at each step, and the claim follows.

Case B: $r \in S(\omega)$. The proof is by induction on i. It holds vacuously for $i = 0$. If $z_{i+1} = x_{i+1}$ then $i + 1 < j$, $Z_i = X_i$ and x_{i+1} was chosen uniformly from $S(\omega^{X_i}) - X_i$. Since, as we explained before, $x_{i+1} \in S(\omega^{r+X_i}) - X_i - r \subseteq S(\omega^{X_i}) - X_i$, we obtain that x_{i+1} is distributed uniformly in $S(\omega^{X_i+r}) - X_i - r$. In the other case, if $z_{i+1} = x_{i+2}$ then x_{i+2} was chosen uniformly from $S(\omega^{X_{i+1}}) - X_{i+1} = S(\omega^{r+Z_i}) - Z_i - r$, completing the proof. $\square$

Example. Let $\omega = (a|b, c|d|e, f)$. The following chart summarizes the distribution of EQUITABLE:

X	$\{a,b,c,d\}$	$\{a,b,c,e\}$	$\{a,b,c,f\}$	$\{a,b,d,e\}$	$\{a,b,d,f\}$	$\{a,c,d,e\}$	$\{a,c,d,f\}$
$P_X(\omega)$	$\frac{1}{10}$	$\frac{3}{20}$	$\frac{3}{20}$	$\frac{3}{20}$	$\frac{3}{20}$	$\frac{3}{20}$	$\frac{3}{20}$

To illustrate how to obtain the above numbers, consider the probability of $X = \{a,b,c,e\}$. In the process used to define EQUITABLE, X can be generated in 24 different orders. The permutations (a,b,c,e), (a,c,b,e), (b,a,c,e), (b,c,a,e), (c,a,b,e), and (c,b,a,e) have probabilities $\frac{1}{6} \cdot \frac{1}{5} \cdot \frac{1}{4} \cdot \frac{1}{3} = \frac{1}{360}$. The permutations (a,b,e,c), (a,c,e,b), (b,a,e,c), and (c,a,e,b) have probabilities $\frac{1}{6} \cdot \frac{1}{5} \cdot \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{240}$. The permutations (a,e,b,c), (a,e,c,b), (b,e,a,c), and (c,e,a,b) have probabilities $\frac{1}{6} \cdot \frac{1}{5} \cdot \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{180}$. The permutations (b,c,e,a) and (c,b,e,a) have probabilities $\frac{1}{6} \cdot \frac{1}{5} \cdot \frac{1}{4} \cdot \frac{1}{1} = \frac{1}{120}$. Permutations (b,e,c,a) and (c,e,b,a) have probabilities $\frac{1}{6} \cdot \frac{1}{5} \cdot \frac{1}{3} \cdot \frac{1}{1} = \frac{1}{90}$. Permutations (e,a,b,c), (e,a,c,b), (e,b,a,c), and (e,c,a,b) have probabilities $\frac{1}{6} \cdot \frac{1}{4} \cdot \frac{1}{3} \cdot \frac{1}{2} = \frac{1}{144}$. Permutations (e,b,c,a) and (e,c,b,a) have probabilities $\frac{1}{6} \cdot \frac{1}{4} \cdot \frac{1}{3} \cdot \frac{1}{1} = \frac{1}{72}$. So the probability of X is $\frac{3}{20}$.

What is the probability $P_c(\omega)$? Item c can be generated in a number of ways. With probability $\frac{1}{6}$, it can be generated in the first step. With probability $\frac{1}{6} \cdot \frac{1}{5}$ it can be generated after b. By considering all possible sequences ending at c, we obtain $P_c(\omega) = \frac{7}{10}$.

Lemma 3 (The Commutativity Lemma). *For each offset function ω and any two items $x, y \in S(\omega)$*

$$P_x(\omega) + P_y(\omega^x) = P_y(\omega) + P_x(\omega^y).$$

Proof. The proof is by induction on $n = |N(\omega)|$. If $n = 0$, then the lemma holds vacuously. In the inductive step, suppose that $\omega = (L_1|\ldots|L_k)$, and that the lemma holds for every offset function with fewer non-revealed items, in particular for each ω^z where $z \in N(\omega)$.

The lemma is obvious if x, y are in the same layer. So we can assume that $x \in L_i$, $y \in L_j$, for some $1 \leqslant i < j \leqslant k$. If x is revealed then the lemma is clearly true. So assume $x, y \in N(\omega)$. Note that $\omega^{pq} = \omega^{qp}$ for all items $p \in S(\omega) - L_k$ and $q \in S(\omega)$. Then, if $j < k$, using (7) and the inductive assumption, we have

$$P_x(\omega) + P_y(\omega^x) = \frac{1}{n} \sum_{z \in N(\omega)} \left[P_x\left(\omega^z\right) + P_y\left(\omega^{zx}\right) \right]$$

$$= \frac{1}{n} \sum_{z \in N(\omega)} \left[P_y\left(\omega^z\right) + P_x\left(\omega^{zy}\right) \right] = P_y(\omega) + P_x(\omega^y).$$

Suppose now $j = k$, and let $\ell = |L_k|$. If $z \in L_k$ then $P_y(\omega^{zx}) = 0$ and $P_x(\omega^z) = P_x(\omega^y)$. Therefore, after using (7) and the inductive assumption for ω^z when $z \in N(\omega) - L_k$, we have

$$P_x(\omega) + P_y(\omega^x) = \frac{1}{n} \sum_{z \in N(\omega) - L_k} \left[P_y(\omega^z) + P_x(\omega^{zy}) \right] + \frac{1}{n} \sum_{z \in L_k} P_x(\omega^z)$$

$$= P_y(\omega) + \frac{1}{n} \sum_{z \in N(\omega) - L_k} P_x(\omega^{yz}) + \frac{1}{n} \sum_{z \in L_k} P_x(\omega^y)$$

$$= P_y(\omega) + \frac{n - \ell}{n} P_x(\omega^y) + \frac{\ell}{n} P_x(\omega^y) \; = \; P_y(\omega) + P_x(\omega^y).$$

In the last step we applied (7) to $P_x(\omega^y)$ with $M = N(\omega^y) = N(\omega) - L_k$. $\square$

Lemma 4. *For any offset function ω, the cost of* EQUITABLE *is the same against all lazy adversary strategies for ω.*

Proof. The proof is by induction on $n = |N(\omega)|$. The base case, $n = 0$, is trivial. For the inductive step, let $\omega = (L_1 | \ldots | L_k)$, and suppose we are given two lazy strategies $x\alpha$ and $y\beta$. Let L_i be the first non-revealed layer. Each lazy strategy must contain an item from L_i, and thus, by Lemmas 2 and 3, we can assume that $x, y \in L_i$. Then EQUITABLE has the same cost on x and y, and ω^x and ω^y are identical, up to a symmetry. Since, by the inductive assumption, the cost of EQUITABLE on α and β is the same, we are done. $\square$

Theorem 2. *Algorithm* EQUITABLE *is H_k-competitive.*

Proof. Let $\Phi(\omega)$ be cost of EQUITABLE in a lazy adversary strategy for $\omega = (L_1 | \ldots | L_k)$. This is well defined because of Lemma 4. By its very definition, Φ satisfies inequality (2) when $r \in S(\omega)$.

For $r \notin S(\omega)$, we proceed by induction on k. The base case, $k = 1$, is trivial: $\Delta\Phi = 0 = H_1 - 1$. For $k > 1$, since $\sum_{x \in S(\omega)} P_x(\omega) = \sum_{x \in S(\omega^r)} P_x(\omega^r) = k$, $S(\omega^r) = S(\omega) + r$, and $P_r(\omega^r) = 1$, we have $\sum_{x \in S(\omega)} [P_x(\omega) - P_x(\omega^r)] = 1$. We conclude that there exists an $x \in S(\omega)$ such that

$$P_x(\omega) - P_x(\omega^r) \leqslant \frac{1}{|S(\omega)|} \; \leqslant \; \frac{1}{k}. \tag{8}$$

We can assume that $x \in L_i$ with $i \neq 1$, for if $x \in L_1$ then $P_x(\omega) = 1$ and $P_x(\omega^r) = P_y(\omega^r)$ for any $y \in L_2$. Let μ be the same as ω^x but with the first layer $\{x\}$ removed.

Then μ is an offset function for the cache with $k - 1$ items. We have $\Phi(\omega^x) = \Phi(\mu)$ and, similarly, $\Phi(\omega^{rx}) = \Phi(\mu^r)$. Therefore

$$\Delta\Phi = \Phi(\omega^r) - \Phi(\omega) = [\Phi(\omega^{rx}) + 1 - P_x(\omega^r)] - [\Phi(\omega^x) + 1 - P_x(\omega)]$$

$$\leqslant \Phi(\mu^r) - \Phi(\mu) + \frac{1}{|S(\omega)|} \leqslant H_{k-1} - 1 + \frac{1}{k} = H_k - 1,$$

completing the proof.   □

4.1. Time and space complexity

In this paper we measure the space complexity by the maximum number of item identifiers stored in the memory at any given time. Although this is not the most general model (which is to count the number of used bits), it is sufficient for comparing memory requirements of realistic paging algorithms.

Let the behavior version of EQUITABLE, as defined in the proof of Lemma 2, be denoted by $\mathcal{B}$. As defined, the space required by $\mathcal{B}$ is O(n), where n is the number of past requests. We now show how to improve the space complexity to O($k^2 \log k$). Let $M = \lceil 5k^2 H_k \rceil$. We only make one simple modification of the algorithm: whenever the current offset function ω satisfies $|S(\omega)| = M$, $\mathcal{B}$ replaces ω in its memory by χ_X, where X is $\mathcal{B}$'s current configuration, and then it proceeds as if χ_X were the current offset function.

$\mathcal{B}$'s cost for this step is zero. Since we have raised values of the current offset function on some configurations, in order to make sure that $\mathcal{B}$ does not benefit from this action, we need to assign to it appropriate (negative) optimal cost Δopt. The offset function is raised by at most $\max(\chi_X - \omega) \leqslant k - 1$, so it is sufficient to set $\Delta opt = -k + 1$. We refer to it as a *forgiveness* step, since on some configurations the overall optimal cost can actually decrease. (See [6, 5] for other examples of the forgiveness method.)

The proof that $\mathcal{B}$ remains H_k-competitive uses a more subtle potential argument. Intuitively speaking, for every request that increases the size of the support (and hence the need for memory) the amortized cost is a little less than $H_k \cdot \Delta opt$. So, when the support gets "too big" the algorithm can use these savings to offset the cost of substituting its current configuration as its estimate for the support. Define a *stage* to be the sequence of steps in-between two forgiveness moves. Let Φ be the potential introduced in Theorem 2. Initially and after each forgiveness move, let $\Psi = 0$. For each move (other than forgiveness) define the quantity $\Delta\Psi = H_k \Delta opt - \Delta cost - \Delta\Phi$. We will refer to each $\Delta\Psi$ as the *savings* at a given step, and let Ψ stand for the total savings up to a given step of a given stage. Since Φ and EQUITABLE satisfy (2), $\Delta\Psi \geqslant 0$ for all moves in the stage. Therefore, $\Psi \geqslant 0$ at all times. By summing over all requests, if Φ and Ψ satisfy

$$\Delta cost_{\mathcal{B}} + \Delta\Phi + \Delta\Psi \leqslant H_k \cdot \Delta opt$$

at each step, then $\mathcal{B}$ is H_k competitive. By the very definition of Ψ, the above inequality is satisfied for each step other than forgiveness. In the forgiveness step we have $\Delta cost = 0$, $\Delta\Phi \leqslant 0$, $\Delta opt = -k + 1$ and $\Delta\Psi = -\Psi$, so it suffices to show the following claim: When $|S(\omega)| = M$ then $\Psi \geqslant (k-1)H_k$. To prove the claim, consider a stage ending at ω. For each $m = k, \ldots, M-1$ there was a step in this stage when the offset function μ satisfied $|S(\mu)| = m$ and the new request was $r \notin S(\mu)$. At that time, from the last inequality in the proof of Theorem 2, we know that $\Delta\Psi$ satisfied $\Delta\Psi \geqslant 1/k - (1/m)$. So at the end of the stage we have

$$
\begin{aligned}
\Psi &\geqslant \sum_{m=k}^{M-1} \left(\frac{1}{k} - \frac{1}{m} \right) = \frac{M}{k} - H_{M-1} + H_{k-1} - 1 \\
&\geqslant kH_k - H_k + 4kH_k - H_{5k^2 H_k} \geqslant (k-1)H_k,
\end{aligned}
$$

where the last inequality uses the fact that $\ln k \leqslant H_k \leqslant \ln k + 1$.

We now show how $\mathcal{B}$ can be implemented in $O(k^2)$ time per request. Let ω be the current offset function, X be the current configuration of $\mathcal{B}$, and r be the request. Updating ω can be accomplished in time $O(k \log k)$ using appropriate data structures to represent the layers of the offset function. Given an order for the items of X, finding the index j described in Lemma 2 can be easily accomplished in $O(k)$ time once we note that we only need to keep track of which layer contains each $x_i \in X$. So, it suffices to show an $O(k^2)$ method for finding an ordering of X which is consistent with Lemma 2. This will be done by induction on the number of requests since the last time ω was a cone. When ω is a cone, each permutation appears with equal probability. This can be accomplished in time $O(k)$ by uniformly selecting k items from X one at a time. Assume that X is ordered, and that Z is defined as in the proof of Lemma 2. Arrange the elements of Z in the order $(z_1, \ldots, z_{i-1}, r, z_i, \ldots z_{k-1})$ with probability proportional to the probability that Z would be chosen by $P(\omega^r)$ in this order. Finding the probabilities for each of these k orders takes time $O(k)$, giving the overall time $O(k^2)$ for ordering Z. An argument analogous to the proof of Lemma 2 shows that using this method for updating the ordering yields the same distribution at each step as that of EQUITABLE. We conclude this section with the following theorem.

Theorem 3. *Algorithm* EQUITABLE *can be implemented in* $O(k^2 \log k)$ *memory and* $O(k^2)$ *time per request.*

5. Final comments

Both PARTITION from [18], and the naive implementation of EQUITABLE are very time and space consuming. They both use $O(n)$ space, where n is the number of past requests. We have shown that EQUITABLE can be implemented in space $O(k^2 \log k)$ and time $O(k^2)$ per step, independent of n. Algorithm MARK, although not optimally competitive, uses only $O(k)$ memory and $O(1)$ time for each request. One problem that we leave open is whether there is a simple H_k-competitive algorithm for paging

which requires only $O(k)$ memory. Raghavan and Snir in [19] investigated a memory versus randomization trade-off in on-line algorithms. Is there also a trade-off between memory and competitiveness?

It would also be interesting to extend the applications of work functions to other approaches to competitive analysis. The competitive analysis has been criticized for being overly pessimistic. Addressing this problem, the most recent work on the competitive analysis of paging moves towards relaxing the definition of competitiveness. Koutsoupias and Papadimitriou [11] used work functions to analyze paging under two refined notions of competitiveness. The first one restricts the adversary to using a distribution chosen from a given set of distributions. The second compares two classes of algorithms by allowing the server to choose an algorithm from the first set, the adversary to choose an algorithm from the second set, and considering the worst case ratio of costs of the two algorithms.

There is also some recent research in this direction that does not involve work functions. In [10] the paging problem is considered in the case where the requests come from a Markov process. In [21, 22] both the paging and weighted cache problems are considered as linear programming problems and the resulting dual problem is investigated. Several results about the so-called loose competitiveness of these problems are proven. In [1, 2, 9] the paging problem is addressed with the assumption that after any particular request the next request is likely to be from a small set of nearby items.

We believe that the work function approach can be used to simplify and possibly refine the results from these papers, and ultimately lead to a more systematic and uniform treatment of different approaches to competitive analysis of paging.

Acknowledgements

References

[1] A. Borodin, S. Irani, P. Raghavan, B. Schieber, Competitive paging with locality of reference, in: Proc. 23rd ACM Symp. on Theory of Computing, 1991, pp. 249–259.

[2] A. Borodin, S. Irani, P. Raghavan, B. Schieber, Competitive paging with locality of reference, J. Comput. System Sci. 50 (2) (1995) 244–258.

[3] M. Chrobak, L.L. Larmore, An optimal online algorithm for k servers on trees, SIAM J. Comput. 20 (1991) 144–148.

[4] M. Chrobak, L.L. Larmore, Metrical service systems: randomized strategies. manuscript, 1992a.

[5] M. Chrobak, L.L. Larmore, Generosity helps or an 11-competitive algorithm for three servers, J. Algorithms 16 (1994) 234–263; also in Proc. ACM-SIAM Symp. on Discrete Algorithms, 1992b, pp. 196–202.

[6] M. Chrobak, L.L. Larmore, N. Reingold, J. Westbrook, Page migration algorithms using work functions, Tech. Report YALE/DCS/RR-910, Department of Computer Science, Yale University, 1992, J. Algorithms, to appear.

 [7] A. Fiat, R. Karp, M. Luby, L.A. McGeoch, D. Sleator, N.E. Young, Competitive paging algorithms, J. Algorithms 12 (1991) 685–699.
 [8] S. Irani, S. Seiden, Randomized algorithms for metrical task systems, in: Proc. 4th Workshop on Algorithms and Data Structures, 1995, pp. 159–170.
 [9] S. Irani, A. Karlin, S. Phillips, Strongly competitive algorithms for paging with locality of reference, in: 3rd Annu. ACM-SIAM Symp. on Discrete Algorithms, 1992, pp. 228–236.
[10] A. Karlin, S. Phillips, P. Raghavan, Markov paging, in: Proc. 33rd IEEE Symp. on Foundations of Computer Science, 1992, pp. 208–217.
[11] E. Koutsoupias, C. Papadimitriou, Beyond competitive analysis, in: Proc. 35th Symp. on Foundations of Computer Science, 1994a, pp. 394–400.
[12] E. Koutsoupias, C. Papadimitriou, On the k-server conjecture, in: Proc. 26th Symp. on Theory of Computing, 1994, pp. 507–511.
[13] E. Koutsoupias, C. Papadimitriou, On the k-server conjecture, J. Assoc. Comput. Mach. 42 (5) (1995) 971–983.
[14] H. Kuhn, Extensive games and the problem of information, in: Kuhn, H., Tucker, A. (Eds.), Contributions to the Theory of Games, Princeton University Press, 1953, pp. 193–216.
[15] C. Lund, N. Reingold, Linear programs for randomized on-line algorithms, in: Proc. 5th ACM-SIAM Symp. on Discrete Algorithms, 1994a, pp. 382–391.
[16] C. Lund, N. Reingold, J. Westbrook, D. Yan, On-line distributed data management, in: Proc. European Symp. on Algorithms, 1994b, pp. 202–214.
[17] M. Manasse, L.A. McGeoch, D. Sleator, Competitive algorithms for server problems, J. Algorithms 11 (1990) 208–230.
[18] L. McGeoch, D. Sleator, A strongly competitive randomized paging algorithm, J. Algorithms 6 (1991) 816–825.
[19] P. Raghavan, M. Snir, Memory versus randomization in online algorithms, in: 16th International Colloquium on Automata, Languages, and Programming, Lecture Notes in Computer Science, vol. 372. Springer, Berlin, 1989, pp. 687–703.
[20] D. Sleator, R.E. Tarjan, Amortized efficiency of list update and paging rules, Communications of the ACM, 28 (1985) 202–208.
[21] N. Young, On-line caching as cache size varies, in: Proc. 2nd Ann. ACM-SIAM Symp. on Discrete Algorithms, 1991, pp. 241–250.
[22] N. Young, The k-server dual and loose competitiveness for paging, Algorithmica 11 (6) (1994) 525–541.