RC23775 (W0511-031) November 4, 2005
Mathematics

IBM Research Report

Near Optimality of the Discrete Persistent Access Caching Algorithm

Predrag Jelenkovic, Xiaozhu Kang
Columbia University
New York, NY

Ana Radovanovic
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Renewed interest in caching techniques stems from their application to improving the performance of the World Wide Web, where storing popular documents in proxy caches closer to end-users can significantly reduce the document download latency and overall network congestion. Rules used to update the collection of frequently accessed documents inside a cache are referred to as cache replacement algorithms. Due to many different factors that influence the Web performance, one of the key attributes of a cache replacement rule are low complexity and high adaptability to variability in Web access patterns. These properties are primarily the reason why most of the practical Web caching algorithms are based on the easily implemented Least-Recently-Used (LRU) cache replacement heuristic.

In our recent paper (8), we introduce a new algorithm, termed Persistent Access Caching (PAC), that, in addition to desirable low complexity and adaptability, somewhat surprisingly achieves nearly optimal performance for the independent reference model and generalized Zipf's law request probabilities. Two drawbacks of the PAC algorithm are its dependence on the request arrival times and variable storage requirements. In this paper, we resolve these problems by introducing a discrete version of the PAC policy (DPAC) that, after a cache miss, places the requested document in the cache only if it is requested at least k times among the last m, $m \geq k$, requests. However, from a mathematical perspective, due to the inherent coupling of the replacement decisions for different documents, the DPAC algorithm is considerably harder to analyze than the original PAC policy. In this regard, we develop a new analytical technique for estimating the performance of the DPAC rule. Using our analysis, we show that this algorithm is close to optimal even for small values of k and m, and, therefore, adds negligible additional storage and processing complexity in comparison to the ordinary LRU policy.

Keywords: persistent-access-caching, least-recently-used caching, least-frequently-used caching, move-to-front searching, generalized Zipf's law distributions, heavy-tailed distributions, Web caching, cache fault probability, average-case analysis

---

Contents

1 Introduction

Since the recent invention of the World Wide Web (WWW), there have been an explosive growth in distributed multimedia content and services that are now integral part of modern communication networks (e.g., the Internet). This massively distributed network information is repeatedly used by groups of users implying that bringing some of the more popular items closer to end-users can improve the network performance, e.g., reduce the download latency and network congestion. This type of information replication and redistribution system is often termed Web caching.

One of the key components of engineering efficient Web caching systems is designing document placement/replacement algorithms that are selecting and possibly dynamically updating a collection of frequently accessed documents. The design of these algorithms has to be done with special care since the latency and network congestion may actually increase if documents with low access frequency are cached. Thus, the main objective is to achieve high cache hit ratios, while maintaining ease of implementation and scalability. Furthermore, these algorithms need to be self-organizing and robust since the document access patterns exhibit a high degree of spatial as well as time fluctuations. The well-known heuristic named the Least-Recently-Used (LRU) cache replacement rule satisfies all of the previously mentioned attributes and, therefore, represents a basis for designing many practical replacement algorithms. However, as shown in (5) in the context of the stationary independent reference model with generalized Zipf's law requests, this rule is by a constant factor away from the optimal frequency algorithm that keeps in the cache most frequently used documents, i.e., replaces Least-Frequently-Used (LFU) items. On the other hand, the drawback of the LFU algorithm is that it needs to know (measure) the document access frequencies and employ aging schemes based on reference counters in order to cope with evolving access patterns, which results in high complexity. In the context of database disk buffering, (10) proposes a modification of the LRU policy, called LRU-K, that uses the information of the last K reference times for each document in order to make replacement decisions. It is shown in (10) that the fault probability of the LRU-K policy approaches, as K increases, the performance of the optimal LFU scheme. However, practical implementation of the LRU-K policy would still be of the same order of complexity as the LFU rule. Furthermore, for larger values of K, that might be required for nearly optimal performance, the adaptability of this algorithm to changes in traffic patterns will be significantly reduced.

In our recent paper (8), we designed a new cache replacement policy, termed the Persistent Access Caching (PAC) rule, that essentially preserves all the desirable features of LRU caching, while achieving arbitrarily close performance to the optimal LFU algorithm. Furthermore, the PAC algorithm has only negligible additional complexity in comparison to the widely used LRU policy. However, the drawback of the PAC policy is that its implementation and analysis depend on the Poisson assumption for the request arrival process. In this paper, we propose a discrete version of the PAC rule (DPAC), that, upon a miss for a document, stores the requested document in the cache only if there are at least k requests for it among m, $m \geq k$, previously requested documents; therefore, DPAC does not depend on request arrival times. Furthermore, the DPAC policy requires only a fixed amount of additional storage for m pointers and a small processing overhead that make it easier to implement than the original PAC rule. On the other hand, due to coupling of the request decisions, as pointed in the abstract, DPAC is significantly more difficult to analyze. To this end, we develop a new analytic technique that, in conjunction with the large deviation analysis and asymptotic results developed in (8; 6; 5), show that the DPAC policy performs near optimal. It is surprising that even for small values of k, m, the performance ratio between the DPAC and optimal LFU algorithm significantly improves when compared to the ordinary LRU; for example, this ratio drops from approximately 1.78 for LRU ($k = 1$) to $1.18, 1.08$ for $k = 2, 3$, respectively. In other words, with only negligible computational complexity relative to the LRU rule, the DPAC algorithm approaches the performance of the optimal LFU scheme without ever having to compute the document access frequencies. Furthermore, we show that the derived asymptotic result and simulation experiments match each other very well, even for relatively small cache sizes.

This paper is organized as follows. First, in Section 2, we formally describe the DPAC policy. Then, we develop a representation theorem for the stationary search cost of the related, Discrete Persistent Move-To-Front algorithm. This representation formula and lemmas of Section 3 provide necessary tools for proving our main theorem in Section 4. Informally, our main result shows that for large cache sizes and independent reference model with generalized Zipf's law request distributions, the fault probability of the DPAC algorithm approaches the optimal LFU policy while maintaining low implementation complexity. Furthermore, in Section 5, the numerical experiments show an excellent agreement between our analytical result and simulations. A brief discussion of our results and their possible extensions is presented in Section 6. In order to alleviate the reading process, we state the proof of a technical lemma in Section 7.

2 Model description and preliminary results

Consider a set $L = \{1, 2, \ldots, N\}$ of N (possibly infinite) documents, out of which x can be stored in an easily accessible location, called cache. The remaining $N - x$ documents (items) are placed outside of the cache in a slower access medium. Documents are requested at moments $\{\tau_n\}_{n\geq 1}$, with increments $\{\tau_{n+1} - \tau_n\}_{n\geq 0}, \tau_0 = 0$, being stationary and ergodic having $\mathbb{E}\tau_1 = 1/\lambda$ for some $\lambda > 0$, and $\tau_{n+1} - \tau_n > 0$ a.s. for $n \geq 0$. Furthermore, define a sequence of i.i.d. random variables $\{R_n\}_{n\geq 1}$, independent from $\{\tau_n\}_{n\geq 1}$, where $\{R_n = i\}$ represents a request for item i at time τ_n. We denote request probabilities as $\mathbb{P}[R_n = i] = q_i$ and, without loss of generality, we assume $q_1 \geq q_2 \geq \ldots$.

Now, we describe the cache replacement algorithm. First, we select fixed design parameters $m \geq k \geq 1$. Then, let $M_i(\tau_n)$ be the number of requests for item i among the m consecutive requests $\tau_n, \tau_{n+1}, \ldots, \tau_{n+m-1}$. Documents stored in the cache are ordered in a list, which is sequentially searched upon a request for a document and is updated as follows. If a requested document at time τ_n, say i, is found in the cache, we have a cache hit. In this case, if the number of requests for document i among the last m requests (including the current request) is at least k, i.e. $M_i(\tau_{n-m+1}) \geq k$, item i is moved to the front of the list while documents that were in front of item i are shifted one position down; otherwise, the list stays unchanged. Furthermore, if document i is not found in the cache, we call it a cache miss or fault. Then, similarly as before, if $M_i(\tau_{n-m+1}) \geq k$, document i is brought to the first position of the cache list and the least recently moved item, i.e., the one at the last position of the list, is evicted from the cache. We name the previously described cache replacement policy the Discrete Persistent Access Caching (DPAC (m,k)) algorithm. Note that in the special case of $m = k = 1$, DPAC reduces to the ordinary LRU heuristic. The performance measure of interest is the cache fault probability, i.e., the probability that a requested document is not found in the cache. We would like to mention that the probabilistic evaluation of an algorithm is typically referred to as the average-case analysis; the pointers to combinatorial (competitive, worse case) approaches can be found in (8).

Analyzing the DPAC(m,k) algorithm is equivalent to investigating the corresponding Move-To-Front (MTF) scheme that is defined as follows. Consider the same arrival model $\{R_n\}$, $\{\tau_n\}$ as in the first paragraph and assume that all documents are ordered in a list $L = \{1,2,\ldots,N\}, N \leq \infty$. When a request for a document arrives, say $R_n = i$, the list is searched and the requested item is moved to the front of the list only when $M_i(\tau_{n-m+1}) \geq k$; otherwise the list stays unchanged. We term the previously described searching algorithm the Discrete Persistent-MTF (DPMTF(m,k)). The performance measure of interest for this algorithm is the search cost $C_n^{(N)}$ that represents the position of the requested document at time τ_n.

Now, we claim that computing the cache fault probability of the DPAC(m,k) algorithm is equivalent to evaluating the tail of the searching cost $C_n^{(N)}$ of the DPMTF(m,k) searching scheme. Note that the fault probability of the DPAC(m,k) algorithm stays the same regardless of the ordering of documents in the slower access medium. In particular, these documents can be also ordered in an increasing order of the last times they are moved to the front of the cache list. Therefore, it is clear that the fault probability of the DPAC(m,k) policy for the cache of size x after the nth request is the same as the probability that the search cost of the DPMTF(m,k) algorithm is greater than x, i.e. $\mathbb{P}[C_n^{(N)} > x]$. Hence, even though DPAC(m,k) and DPMTF(m,k) belong to different application areas, their performance analysis is essentially equivalent. Thus, in the rest of the paper we investigate the tail of the stationary search cost distribution.

First, we prove the convergence of the search cost $C_n^{(N)}$ to stationarity. Suppose that the system starts at τ_0 with initial conditions given by an arbitrary initial permutation Π_0 of the list and a sequence of requests $R_0 = \{r_{-m+1}, r_{-m+2}, \ldots, r_{-1}\}$.

In order to prove the convergence of $C_n^{(N)}$ to stationarity, we construct a sequence of DPMTF searching schemes that start at negative time points and are observed at time $\tau_0 = 0$. To that end, let $\{R_{-n}\}_{n\geq 0}$ be a sequence of i.i.d. requests that are equal in distribution to R_1 that arrive at points on the negative axis $\{\tau_{-n}\}_{n\geq 1}$; these arrival points are equal in distribution to $\{\tau_n\}_{n\geq 1}$ and are independent from $\{R_{-n}\}_{n\geq 0}$. Then, for each $n > 0$ we construct a DPMTF(m,k) algorithm starting at τ_{-n}, with a sequence of requests $\{R_l : l = -n+1, \ldots, -1, 0\}$ and having the same initial condition as in the previous paragraph, given by Π_0 and R_0; let $C_{-n}^{(N)}$ be the search cost at time $\tau_0 = 0$. Note that in this construction we assume that for the DPMTF(m,k) algorithm starting at τ_{-n}, $n > 0$, there is no request at time τ_{-n}.

Now, if we consider the shift mapping $R_{n-k} \to R_{-k}, \tau_{n-k} \to \tau_{-k}$ for $k = 0, 1, \ldots n-1$, we conclude that, since the corresponding sequences are equal in distribution, the search costs $C_{-n}^{(N)}$ and $C_n^{(N)}$ are also equal in distribution, i.e. $C_n^{(N)} \stackrel{d}{=} C_{-n}^{(N)}$. Thus, instead of computing the tail of the search cost $C_n^{(N)}$, we continue with evaluating the tail of $C_{-n}^{(N)}$. In this regard, we define a sequence of random times $\{T_i^{(-n)}\}_{n\geq 1}$, where $-T_i^{(-n)}$ represents the last time before $t = 0$ that item i was moved to the front of the list in the case of the DPMTF(m,k) algorithm that started at τ_{-n}; if item i is not moved in $(\tau_{-n}, 0)$, we set $T_i^{(-n)} = -\tau_{-n}$. Next, we define random times T_i, $i \geq 1$, as

$$T_i \triangleq -\sup\{\tau_{-n} < 0 : R_{-n} = i, M_i(\tau_{-n-m+1}) \geq k\}. \tag{1}$$

From the definitions of T_i and $T_i^{(-n)}$, we conclude that the equality $T_i = T_i^{(-n)}$ a.s. holds on event $\{T_i^{(-n)} < -\tau_{-n+m-1}\}$. Therefore, the complementary sets of events are the same, i.e. $\{T_i \geq \tau_{n-m+1}\} = \{T_i^{(-n)} \geq -\tau_{-n+m-1}\}$.

Then, given the previous observations, we bound the tail of the search cost $C_{-n}^{(N)}$ as

$$\begin{aligned}\mathbb{P}[C_{-n}^{(N)} > x, R_0 = i, T_i^{(-n)} < -\tau_{-n+m-1}] \leq \mathbb{P}[C_{-n}^{(N)} > x, R_0 = i] \leq \\ \mathbb{P}[C_{-n}^{(N)} > x, R_0 = i, T_i^{(-n)} < -\tau_{-n+m-1}] + \mathbb{P}[C_{-n}^{(N)} > x, R_0 = i, T_i^{(-n)} \geq -\tau_{-n+m-1}].\end{aligned} \tag{2}$$

Next, observe that on event $\{R_0 = i, T_i^{(-n)} < -\tau_{-n+m-1}\}$ the search cost $C_{-n}^{(N)}$ is equal to the number of different documents (including i) that are moved to the front of the list from the last time that item i was brought to the first position. Thus, we derive

$$\begin{aligned}\mathbb{P}[C_{-n}^{(N)} > x, R_0 = i, T_i^{(-n)} < -\tau_{-n+m-1}] &= \mathbb{P}\left[R_0 = i, \sum_{j\neq i} 1[T_j^{(-n)} < T_i^{(-n)} < -\tau_{-n+m-1}] \geq x\right] \\ &= q_i \mathbb{P}\left[\sum_{j\neq i} 1[T_j < T_i < \tau_{n-m+1}] \geq x\right], \end{aligned} \tag{3}$$

where the last equality follows from the independence assumption on $\{\tau_{-n}\}_{n\geq 0}$, $\{R_{-n}\}_{n\geq 0}$ and the equality $T_i = T_i^{(-n)}$, $i \geq 1$, on $\{T_i < \tau_{n-m+1}\}$.

Hence, by monotone convergence theorem,

$$\lim_{n\to\infty} \sum_{i=1}^{N} \mathbb{P}[C_{-n}^{(N)} > x, R_0 = i, T_i < \tau_{n-m+1}] = \sum_{i=1}^{N} q_i \mathbb{P}\left[\sum_{j\neq i} 1[T_j < T_i] \geq x\right]. \tag{4}$$

Furthermore, due to stationarity and ergodicity of the arrival process $\{\tau_n\}$ and finiteness of $T_i < \infty$ a.s. we obtain

$$\lim_{n\to\infty} \mathbb{P}[T_i > \tau_n] = 0. \tag{5}$$

Finally, equality of events $\{T_i^{(-n)} \geq -\tau_{-n+m-1}\} = \{T_i \geq \tau_{n-m+1}\}$, independence of requests, limit in (5) and dominated convergence theorem imply

$$\lim_{n\to\infty} \sum_{i=1}^{N} \mathbb{P}[R_0 = i, T_i^{(-n)} \geq -\tau_{-n+m-1}] = \lim_{n\to\infty} \sum_{i=1}^{N} q_i \mathbb{P}[T_i \geq \tau_{n-m+1}] = 0.$$

The previous expression, in conjunction with (4) and (2), implies the following representation result:

Lemma 1 *For any $1 \leq N \leq \infty$, arbitrary initial conditions (Π_0, R_0) and any $x \geq 0$, the search cost $C_n^{(N)}$ converges in distribution to $C^{(N)}$ as $n \to \infty$, where*

$$\mathbb{P}[C^{(N)} > x] \triangleq \sum_{i=1}^{N} q_i \mathbb{P}[S_i(T_i) \geq x]; \tag{6}$$

we define $S_i(t) \triangleq \sum_{j\neq i} 1[T_j < t]$, $i \geq 1$.

Remarks: (i) Note that the expression in (6) is independent from the selection of the arrival process $\{\tau_n\}_{n\geq 1}$. To see this, assume two arrival processes $\{\tau_n\}_{n\geq 1}$ and $\{\tau'_n\}_{n\geq 1}$ that are independent from requests $\{R_n\}_{n\geq 1}$ and satisfy the stationarity, ergodicity and monotonicity conditions from the beginning of this section. Using definition (1), we define random times T_i, T'_i, $i \geq 1$, that correspond to processes $\{\tau_n\}_{n\geq 1}$, $\{\tau'_n\}_{n\geq 1}$, respectively. Then, it is easy to observe that $1[T_j < T_i] = 1[T'_j < T'_i]$ a.s. for any $j \neq i$, $i, j \geq 1$, i.e., the sequences of random times T_i, T'_i, $i \geq 1$ are ordered in exactly the same way. Thus, since $S_i(T_i)$ is completely determined by the ordering of these random times, it is clear that the distributions of the corresponding search costs are exactly the same. (ii) Using the preceding observation we can assume in the rest of the paper without loss of generality that $\{\tau_n\}_{n\geq 1}$ is a Poisson sequence of points with rate 1. This assumption will be helpful in Section 3 in decoupling the dependency among random times T_i, $i \geq 1$. The Poisson embedding technique for LRU policy with i.i.d. requests was first introduced in (3).

3 Preliminary results on Poisson processes

This section provides bounds on random times T_i and the sum $S_i(t)$, as defined in Lemma 1, that represent necessary building blocks for the asymptotic analysis of the stationary search cost from Lemma 1. Furthermore, it is worth noting that Lemma 4 and 5 develop a new technique that allows the decoupling of the dependency among random times T_i and, thus, enable us to estimate the sum in $S_i(t)$.

Recall the definition of T_i from (1). In order to avoid dealing with negative indices and signs we define here a sequence of random times on the positive sequence $\{\tau_n\}_{n\geq 1}$ that are equal in distribution to $T_i, i \geq 1$. Thus, with a small abuse of notation, we use the same name T_i for the following random times

$$T_i \triangleq \inf\{\tau_n > 0 : R_n = i, M_i(\tau_n) \geq k\}. \tag{7}$$

Next, as proposed in the remark after Lemma 1, we assume that $\{\tau_n\}_{n\geq 1}$ is a Poisson process of rate 1. Then, let $\{\tau_n^{(i)}\}_{n\geq 1}$ be a sequence of requests for document i. Given the i.i.d. assumption on $\{R_n\}_{n\geq 1}$ and its independence from the arrival points $\{\tau_n\}_{n\geq 1}$, the Poisson decomposition theorem implies that processes $\{\tau_n^{(i)}\}_{n\geq 1}$ are also Poisson and mutually independent for different i with rate q_i. This observation will be used in the proofs of the subsequent lemmas.

In order to ease the notation, throughout the paper we use H to denote a sufficiently large positive constant and h to denote a sufficiently small positive constant. The values of H and h are generally different in different places. For example, $H/2 = H$, $H^2 = H$, $H+1 = H$, etc. Next, we compute a bound on the tail of distribution of T_i for large i.

Lemma 2 *For any* $\varepsilon > 0$, *there exists* i_0, *such that for all* $i \geq i_0$,

$$\mathbb{P}[T_i > t] \leq e^{-\binom{m-1}{k-1}(1-\varepsilon)^2 q_i^k t} + m e^{-h\varepsilon q_i^{k-1} t}. \tag{8}$$

Proof: For $k = 1$ the bound trivially holds since $T_i \equiv \tau_1^{(i)}$ and, thus, we assume that $k \geq 2$.

First, we define a sequence of random times $\{\Theta_j\}$. We set $\Theta_1 = \tau_1^{(i)}$, and define $n(j)$, $j \geq 1$ to be the indices of points $\{\tau_j^{(i)}\}_{j\geq 1}$ in the original sequence $\{\tau_n\}_{n\geq 1}$, i.e. $\tau_j^{(i)} = \tau_{n(j)}$, $j \geq 1$. Then, if the first point from the sequence $\{\tau_j^{(i)}\}$ after time $\tau_{n(1)+m-1}$ is $\tau_{j_1}^{(i)}$, we define $\Theta_2 = \tau_{j_1}^{(i)}$. Similarly, Θ_3 is defined to be the first point from $\{\tau_j^{(i)}\}$ after time $\tau_{n(j_1)+m-1}$, etc. Observe that $\{\Theta_j\}$ is a renewal process with its increments for $j \geq 1$ equal to

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(i)} + \sum_{l=1}^{m-1} \xi_l, \tag{9}$$

where $\stackrel{d}{=}$ denotes equality in distribution and $\{\xi_j\}_{j\geq 1}$ are independent, exponentially distributed random variables with mean 1 that are independent of $\{\tau_n^{(i)}\}_{n\geq 1}$.

Next, we define

$$U_i \triangleq \inf\{\Theta_j : j \geq 1, M_i(\Theta_j) \geq k\}.$$

Note that this definition of U_i has identical form to the one for T_i in (7) since $\{\Theta_j\} \subset \{\tau_j^{(i)}\}$ and thus $R(\Theta_j) \equiv i$. Therefore, given $\{\Theta_j\} \subset \{\tau_n\}$, it is clear that

$$T_i \leq U_i. \tag{10}$$

Similarly, we define

$$X \triangleq \inf\{j \geq 1 : M_i(\Theta_j) \geq k\}.$$

Since $\{R_n\}$ is i.i.d and independent of $\{\tau_n\}$, X is independent of $\{\Theta_j\}_{j\geq 1}$ with a geometric distribution $\mathbb{P}[X = j] = (1-p)^{j-1}p$, $j \geq 1$, where p is equal to

$$p = \mathbb{P}[M_i(\tau_1^{(i)}) \geq k].$$

Then, from the definition of U_i and (9) we obtain

$$U_i = \Theta_X \stackrel{d}{=} \tau_X^{(i)} + \sum_{j=1}^{(m-1)(X-1)} \xi_j \stackrel{d}{\leq} \tau_X^{(i)} + \sum_{j=1}^{(m-1)X} \xi_j, \tag{11}$$

where $\stackrel{d}{\leq}$ represents inequality in distribution and X is independent of $\{\tau_n^{(i)}\}$ and $\{\xi_j\}$.

Next, since $\tau_X^{(i)}$ is a geometric sum of exponential random variables with X and $\{\tau_n^{(i)}\}_{n\geq 1}$ independent, it is easy to show (see Theorem 5.3, p. 89 of (2)) that $\tau_X^{(i)}$ is also exponential with parameter pq_i. Similarly, $\sum_{j=1}^{X} \xi_j$ is exponential with parameter p. Now, from monotonicity of q_i and $q_i \to 0$ as $i \to \infty$, it is also straightforward to derive that for any $\varepsilon > 0$, there exists i_0, such that for all $i \geq i_0$

$$p = \mathbb{P}[M_i(\tau_1^{(i)}) \geq k] = \sum_{l=k-1}^{m-1} \binom{m-1}{l} q_i^l (1-q_i)^{m-1-l} \geq (1-\varepsilon)\binom{m-1}{k-1} q_i^{k-1}. \tag{12}$$

At this point, using the observations from the previous paragraph, (10) and (11), we obtain, for all i large enough ($i \geq i_0$),

$$\begin{aligned}
\mathbb{P}[T_i > t] &\leq \mathbb{P}[U_i > t] \\
&\leq \mathbb{P}\left[\tau_X^{(i)} > (1-\varepsilon)t\right] + \mathbb{P}\left[\sum_{j=1}^{X(m-1)} \xi_j > \varepsilon t\right] \\
&\leq e^{-pq(1-\varepsilon)t} + (m-1)\mathbb{P}\left[\sum_{j=1}^{X} \xi_j > \frac{\varepsilon t}{m-1}\right] \\
&\leq e^{-\binom{m-1}{k-1} q_i^k (1-\varepsilon)^2 t} + (m-1) e^{-\frac{p\varepsilon t}{m-1}} \\
&\leq e^{-\binom{m-1}{k-1} q_i^k (1-\varepsilon)^2 t} + m e^{-h\varepsilon q_i^{k-1} t};
\end{aligned} \tag{13}$$

this completes the proof. $\diamond$

Next, we will prove the lower bound for the random time T_i defined in (7).

Lemma 3 *For any $\varepsilon > 0$, there exists i_0 such that for all $i \geq i_0$*

$$\mathbb{P}[T_i > t] \geq e^{-(1+\varepsilon)\binom{m-1}{k-1} q_i^k t}. \tag{14}$$

Proof: Since the bound is immediate for $k = 1$, we assume $k \geq 2$.

First, we group the points $\{\tau_n^{(i)}\}_{n\geq 1}$ into cycles using the following procedure. Let $\Theta_1 = \tau_1^{(i)}$ and define a random time

$$Z_1 = \inf\{j \geq 2 : M_i(\tau_{n(1)+m(j-1)}) = 0\},$$

where $n(1)$ is the index of the point $\Theta_1 \equiv \tau_1^{(i)}$ in the original sequence $\{\tau_n\}_{n\geq 1}$, i.e. $\tau_1^{(i)} = \tau_{n(1)}$. Then, the first cycle is the interval of time $C_1 = [\tau_{n(1)}, \tau_{n(1)+mZ_1-1}]$. Next, the first point of process $\{\tau_n^{(i)}\}$ after time $\tau_{n(1)+mZ_1-1}$, say $\tau_l^{(i)}$, we label as $\Theta_2 = \tau_l^{(i)}$ and, similarly as before, we define a random time

$$Z_2 = \inf\{j \geq 2 : M_i(\tau_{n(2)+m(j-1)}) = 0\},$$

where $n(2)$ is the index of the point Θ_2 in the original sequence $\{\tau_n\}_{n\geq 1}$, i.e. $\Theta_2 = \tau_{n(2)}$. We continue this procedure indefinitely. Note that the sequences $\{\Theta_j\}_{j\geq 1}, \{n(j)\}_{j\geq 1}$ as well as the other auxiliary variables (e.g., p, X) are different from the ones in the proof of Lemma 2. The same will apply for the proofs of the remaining lemmas in this section.

Now, due to the i.i.d structure of $\{R_n\}$ and its independence from $\{\tau_n\}$, the sequence of random times $\{Z_j\}$ is i.i.d. with geometric distribution

$$\mathbb{P}[Z_i = j] = (\mathbb{P}[M_i(\tau_1) > 0])^{j-2}\mathbb{P}[M_i(\tau_1) = 0], \quad j \geq 2. \tag{15}$$

Furthermore, $\{\Theta_j\}$ is a renewal process with renewal intervals equal to, for $j \geq 1$,

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(i)} + \sum_{i=1}^{mZ_1} \xi_i, \tag{16}$$

where $\{\xi_i\}$ is an i.i.d. sequence of exponential random variables with mean 1 that is independent of $\tau_1^{(i)}, Z_1$.

Next, we define for $j \geq 1$ sets

$$A_j \triangleq \{\omega : \exists \tau_n \in C_j, R(\tau_n) = i, M_i(\tau_n) \geq k\};$$

note that events A_j are independent since $M_i(\tau_{n(j)+m(Z_j-1)}) = 0$. Then, since the union of the arrival points in all cycles $\cup_j C_j$ contains all requests $\{\tau_n^{(i)}\}$,

$$\begin{aligned} T_i &= \inf\{\tau_n : R(\tau_n) = i, M_i(\tau_n) \geq k, \tau_n \in C_j, j \geq 1\} \\ &\geq L_i \triangleq \inf\{\Theta_j(\omega) : \omega \in A_j, j \geq 1\}, \end{aligned} \tag{17}$$

where the inequality is implied by $\tau_n \geq \Theta_j$ for any $\tau_n \in C_j$, $j \geq 1$.

Furthermore, we claim that

$$L_i = \Theta_X \overset{d}{\geq} \tau_X^{(i)}, \tag{18}$$

where X is independent of $\{\tau_n^{(i)}\}$, $\{\Theta_j\}$ and has geometric distribution $\mathbb{P}[X = j] = (1-p)^{j-1}p, j \geq 1$, with success probability

$$p = \mathbb{P}[A_n] \leq \mathbb{P}[\{M_i(\tau_1^{(i)}) \geq k\} \cup \{M_i(\tau_1^{(i)}, mZ_1) \geq k+1\}],$$

where $M_i(\tau_n, k)$ is defined as the number of references for document i among the requests occurring at $\tau_n, \tau_{n+1}, \ldots, \tau_{n+k-1}$. The equality in (18) follows from the renewal property of the Poisson process, the definition of Z_i, the independence of $\{R_n\}$ and $\{\tau_n\}$; and the inequality follows by neglecting the sum in (16). Furthermore, similarly as in the proof of Lemma 2, $\tau_X^{(i)}$ is an exponential random variable with distribution

$$\mathbb{P}[\tau_X^{(i)} > t] = e^{-pq_i t}. \tag{19}$$

Thus, in order to complete the proof, we need an upper bound on p. In this respect, using the union bound, we upper bound the success probability p as

$$\begin{aligned} p &\leq \mathbb{P}[\{M_i(\tau_1^{(i)}) \geq k\} \cup \{M_i(\tau_1^{(i)}, mZ_1) \geq k+1\}] \\ &\leq \mathbb{P}[M_i(\tau_1, m-1) \geq k-1] + \mathbb{P}[Z_1 > k+1] + \mathbb{P}[M_i(\tau_1, m(k+1)) \geq k] \\ &= \mathbb{P}[M_i(\tau_1, m-1) \geq k-1] + \mathbb{P}[M_i(\tau_1) \geq 1]^k + \mathbb{P}[M_i(\tau_1, m(k+1)) \geq k], \end{aligned} \tag{20}$$

where in the last equality we used the geometric distribution of Z_1 from (15). Finally, (17),(18), (19), (20) and the fact that uniformly for all $1 \leq l_2 \leq l_1 \leq m(k+1)$ and any fixed $\varepsilon > 0$,

$$\mathbb{P}[M_i(\tau_1, l_1) \geq l_2] = \sum_{s=l_2}^{l_1} \binom{l_1}{s} q_i^s (1-q_i)^{l_1 - s} \leq (1+\varepsilon)\binom{l_1}{l_2} q_i^{l_2}$$

for all i large enough ($i \geq i_0$), yield the stated bound in the lemma. ◇

In this paper we are using the following standard notation. For any two real functions $a(t)$ and $b(t)$ and fixed $t_0 \in \mathbb{R} \cup \{\infty\}$ we will use $a(t) \sim b(t)$ as $t \to t_0$ to denote $\lim_{t \to t_0}[a(t)/b(t)] = 1$. Similarly, we say that $a(t) \gtrsim b(t)$ as $t \to t_0$ if $\liminf_{t \to t_0} a(t)/b(t) \geq 1$; $a(t) \lesssim b(t)$ has a complementary definition.

In the following lemmas, we develop an analytic technique that allows us to decouple the dependency of random times T_i for i large, which in conjunction with the large deviation bound proved in Lemma 4 of (6), provides necessary estimates used in the proof of our main result in Section 4.

Lemma 4 *Let T_i, $i \geq 1$, be random variables defined in (7). Then, for $q_i \sim c/i^\alpha$ as $i \to \infty$, $\alpha > 1$ and*

$$t_0(x) = \frac{(1+\varepsilon)^{\alpha k} x^{\alpha k}}{(1-\varepsilon)^{k+2} c^k \binom{m-1}{k-1} [\Gamma(1 - \frac{1}{\alpha k})]^{\alpha k}},$$

we obtain

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < t_0(x)] \leq x\right] = o\left(\frac{1}{x^{\alpha-1}}\right) \quad \text{as } x \to \infty.$$

Proof: Note that for any $i_0 \geq 1$

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < t_0(x)] \leq x\right] \leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < t_0(x)] \leq x\right]. \tag{21}$$

Let $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1}$ be an ordered sequence of request times for documents $i \geq i_0$, i.e. $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1} = \cup_{i\geq i_0}\{\tau_n^{(i)}\}_{n\geq 1}$. We use $n(j),\ j \geq 1$, to denote the index of point $\tau_j^{(\overline{i_0})}$ in the original sequence $\{\tau_n\}_{n\geq 1}$, i.e. $\tau_j^{(\overline{i_0})} = \tau_{n(j)}$. Then, since process $\{\tau_n\}$ is Poisson and $\{R_n\}$ is i.i.d. sequence independent of $\{\tau_n\}$, by Poisson decomposition theorem, process $\{\tau_j^{(\overline{i_0})} \equiv \tau_{n(j)}\}_{j\geq 1}$ is also Poisson with rate $\sum_{i\geq i_0} q_i$. Next, in order to estimate an upper bound for random times T_i, $i \geq i_0$, we proceed as follows.

First, we define a sequence of random times $\{\Theta_j\}$. We set $\Theta_1 = \tau_{n(1)} \equiv \tau_1^{(\overline{i_0})}$; then, if the first point from the sequence $\{\tau_{n(j)}\}$ after time $\tau_{n(1)+m-1}$ is $\tau_{n(j_1)}$, we define $\Theta_2 = \tau_{n(j_1)}$. Similarly, Θ_3 is defined to be the first point from $\{\tau_{n(j)}\}$ after time $\tau_{n(j_1)+m-1}$, etc. Note that, due to the renewal structure of $\{\tau_n\}$, $\{\Theta_j\}$ is a renewal process whose increments for $j \geq 1$ satisfy

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(\overline{i_0})} + \sum_{l=1}^{m-1} \xi_l, \tag{22}$$

where $\tau_1^{(\overline{i_0})}, \{\xi_l\}_{l\geq 1}$ are independent exponential random variables with $\tau_1^{(\overline{i_0})}$ having parameter $\sum_{i\geq i_0} q_i$ and ξ_1 having parameter 1.

Next, for all $i \geq i_0$ define

$$U_i \triangleq \inf\{\Theta_j : j \geq 1, R(\Theta_j) = i, M_i(\Theta_j) \geq k\}.$$

Similarly as in the proof of Lemma 2, the definition of U_i, $i \geq i_0$, has identical form to the one for T_i in (7). The only difference is that points $\{\Theta_j\} \subset \{\tau_n\}$ and, therefore,

$$T_i \leq U_i. \tag{23}$$

Then, using (22), for $j \geq 1$ we have that

$$\Theta_j \stackrel{d}{=} \tau_j^{(\overline{i_0})} + \nu_{j-1}^{(1)} + \cdots + \nu_{j-1}^{(m-1)}, \tag{24}$$

where $\nu_0^{(l)} = 0, 1 \leq l \leq m-1$ and $\{\nu_j^{(l)}\}_{j\geq 1}, 1 \leq l \leq m-1$ are independent Poisson processes of rate 1 that are also independent of the Poisson process $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1}$ having rate $\sum_{i\geq i_0} q_i$. Using this observation and the fact that $\{R_n\}$ and $\{\tau_n\}$ are independent, we arrive at the following representation

$$U_i \stackrel{d}{=} \tau_{X_i}^{(\overline{i_0})} + \sum_{l=1}^{m-1} \nu_{X_i-1}^{(l)} \leq \tau_{X_i}^{(\overline{i_0})} + \sum_{l=1}^{m-1} \nu_{X_i}^{(l)}, \tag{25}$$

where X_i is a geometric random variable independent from $\{\nu_j^{(l)}\}_{j\geq 1}, 1 \leq l \leq m-1$ and $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1}$ with $\mathbb{P}[X_i = j] = (1-p_i)^{j-1} p_i$, where

$$p_i \triangleq \sum_{l\geq k-1} \frac{q_i}{\sum_{j\geq i_0} q_j} \binom{m-1}{l} q_i^l (1-q_i)^{m-1-l}.$$

Then, again, due to the Poisson decomposition theorem, variables $\{\tau_{X_i}^{(\overline{i_0})}\}_{i\geq i_0}$ are independent and exponentially distributed with $\tau_{X_i}^{(\overline{i_0})}$ having parameter $p_i \sum_{i\geq i_0} q_i$. Similarly, for each fixed $1 \leq l \leq m-1$, variables $\{\nu_{X_i}^{(l)}\}_{i\geq i_0}$ are also independent and exponential with $\nu_{X_i}^{(l)}$ having parameter p_i. (Note that for different l the sequences $\{\nu_{X_i}^{(l)}\}_{i\geq i_0}$ can be mutually dependent and also potentially dependent on $\{\tau_{X_i}^{(\overline{i_0})}\}_{i\geq i_0}$.) Furthermore, observe that for any $\varepsilon > 0$ and i_0 large enough

$$(1+\varepsilon)\frac{q_i}{\sum_{j\geq i_0} q_j} q_i^{k-1} \binom{m-1}{k-1} \geq p_i \geq (1-\varepsilon)\frac{q_i}{\sum_{j\geq i_0} q_j} q_i^{k-1} \binom{m-1}{k-1}. \tag{26}$$

Next, inequalities (23) and (25) imply, for any $i \geq i_0$ and $\varepsilon > 0$

$$
\begin{aligned}
1[T_i < t_0(x)] &\geq 1\left[\tau_{X_i}^{(\overline{i_0})} + \sum_{l=1}^{m-1} \nu_{X_i}^{(l)} < t_0(x)\right] \\
&\geq 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)] - m1\left[\nu_{X_i}^{(1)} > \frac{\varepsilon t_0(x)}{m}\right],
\end{aligned}
$$

and, therefore,

$$
\begin{aligned}
\mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < t_0(x)] \leq x\right] &\leq \mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)] - m\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1\left[\nu_{X_i}^{(1)} > \frac{\varepsilon t_0(x)}{m}\right] \leq x\right] \\
&\leq \mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)] \leq (1+\varepsilon/2)x\right] \\
&\quad + \mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1\left[\nu_{X_i}^{(1)} > \frac{\varepsilon t_0(x)}{m}\right] > \frac{x\varepsilon}{2m}\right]. \qquad (27)
\end{aligned}
$$

Now, using using Lemma 2 of (5), setting $i_0 = \lfloor\sqrt{x}\rfloor$, we derive

$$
\mathbb{E}\left[\sum_{i=\lfloor\sqrt{x}\rfloor}^{\infty} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)]\right] \gtrsim x(1+\varepsilon). \qquad (28)
$$

Using $q_i \sim c/i^{\alpha}$ and $1 - e^{-x} \leq x$, we arrive at

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i=\lfloor x\log x\rfloor+1}^{\infty} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)]\right] &= \sum_{i=\lfloor x\log x\rfloor+1}^{\infty} \mathbb{P}\left[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)\right] \\
&\leq \sum_{i=\lfloor x\log x\rfloor+1}^{\infty} 1 - e^{-Hq_i^k t_0(x)} \leq H \sum_{i=\lfloor x\log x\rfloor+1}^{\infty} \frac{x^{\alpha k}}{i^{\alpha k}} \\
&\leq H\frac{x^{\alpha k}}{(x\log x)^{(\alpha k-1)}} = \frac{Hx}{(\log x)^{\alpha k-1}} = o(x) \text{ as } x \to \infty.
\end{aligned}
$$

Thus, using the preceding estimate and (28), we obtain

$$
\mathbb{E}\left[\sum_{i=\lfloor\sqrt{x}\rfloor}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)]\right] \sim \mathbb{E}\left[\sum_{i=\lfloor\sqrt{x}\rfloor}^{\infty} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)]\right] \gtrsim x(1+\varepsilon) \text{ as } x \to \infty.
$$

Then, the previous expression, in conjunction with Lemma 4 of (6) implies that the first term of (27) satisfies, as $x \to \infty$,

$$
\mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\varepsilon)t_0(x)] \leq (1+\varepsilon/2)x\right] = o\left(\frac{1}{x^{\alpha-1}}\right). \qquad (29)
$$

Next, it is left to estimate the second term of (27). To this end, by using the monotonicity of q_i-s, assumption $q_i \sim c/i^{\alpha}$ as $i \to \infty$, inequality (26), $i_0 = \lfloor\sqrt{x}\rfloor$ and replacing $t_0(x)$, we obtain

$$
\begin{aligned}
\sum_{i=\lfloor\sqrt{x}\rfloor}^{\lfloor x\log x\rfloor} \mathbb{P}\left[\nu_{X_i}^{(1)} > \frac{\varepsilon t_0(x)}{m}\right] &\leq x\log x e^{-(1-\varepsilon)\frac{\varepsilon}{m} q_{\lfloor x\log x\rfloor}^k \binom{m-1}{k-1} \frac{1}{\sum_{i\geq\lfloor\sqrt{x}\rfloor} q_i} t_0(x)} \\
&= x\log x e^{-\frac{hx^{\alpha k} x^{\frac{\alpha-1}{2}}}{(x\log x)^{\alpha k}}} = x\log x e^{-\frac{hx^{\frac{\alpha-1}{2}}}{(\log x)^{\alpha k}}} = o(x) \text{ as } x \to \infty.
\end{aligned}
$$

Finally, applying Lemma 4 of (6), we derive

$$\mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1\left[\nu_{X_i}^{(1)} > \frac{\varepsilon t_0(x)}{m}\right] > \frac{x\varepsilon}{2m}\right] = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x\to\infty,$$

which, in conjunction with (21), (27) and (29), completes the proof of this lemma. ◇

Lemma 5 *Let T_i, $i \geq 1$, be random variables defined in (7). Then, for $q_i \sim c/i^{\alpha}$ as $i\to\infty$, $\alpha > 1$ and*

$$t_0(x) = \frac{x^{\alpha k}(1-2\varepsilon)^{\alpha k}}{(1+\varepsilon)^{k+1}c^k\binom{m-1}{k-1}[\Gamma(1-\frac{1}{\alpha k})]^{\alpha k}},$$

we obtain

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < t_0(x)] \geq x\right] = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x\to\infty.$$

Proof: The proof of this lemma uses the idea of cycles from the proof of Lemma 3 in order to lower bound the random times $\{T_i\}$ with a sequence of independent random variables, similarly as in Lemma 4. Thus, since many of the arguments are repetitive, we postpone this proof to Section 7. ◇

4 Near optimality of the DPAC algorithm

Consider the class of so-called online caching algorithms that make their replacement decisions using only the knowledge of the past requests and cache contents. Assume also that, at times of cache faults, the replacement decisions are only optional, i.e., the algorithm may keep the cache content constant (static). Within this context and the independent reference model, it is well known that the static LFU policy that stores the most popular documents in the cache is the optimal. For direct arguments that justify this intuitively apparent statement see the first paragraph of Subsection 4.1 in (7); this is also recently shown in (1) using the formalism of Markov Decision Theory. Therefore, $\mathbb{P}[R > x]$ is the fault probability of the optimal static policy and $\mathbb{P}[C > x]/\mathbb{P}[R > x]$ is an average-case competitive ratio between the stationary fault probabilities of the DPAC and optimal static algorithm. In the following theorems we show that, for the case of generalized Zipf's law distributions and large caches, this competitive ratio approaches 1 very rapidly as k grows.

First, we estimate the asymptotics of the tail of the stationary search cost $C^{(N)}$ for Zipf's distributions with $\alpha > 1$. In addition to the analytic techniques developed in Section 3, our method of proof uses probabilistic and sample path arguments introduced in (6) and (8) for the case of ordinary LRU and continuous time PAC algorithms, respectively. The starting point of our analysis is given by representation formula in (6) from Section 2. We assume that $N = \infty$ and denote $C \equiv C^{(\infty)}$.

Theorem 1 *Assume that $q_i \sim c/i^{\alpha}$ as $i\to\infty$ and $\alpha > 1$. Then, as $x\to\infty$*

$$\mathbb{P}[C > x] \sim K_k(\alpha)\mathbb{P}[R > x], \tag{30}$$

where

$$K_k(\alpha) \triangleq \left[\Gamma\left(1-\frac{1}{\alpha k}\right)\right]^{\alpha-1}\Gamma\left(1+\frac{1}{k}-\frac{1}{\alpha k}\right). \tag{31}$$

Furthermore, function $K_k(\alpha)$ is monotonically increasing in α, for fixed k, with

$$\lim_{\alpha\downarrow 1} K_k(\alpha) = 1, \quad \lim_{\alpha\uparrow\infty} K_k(\alpha) = K_k(\infty) \triangleq \frac{1}{k}\Gamma\left(\frac{1}{k}\right)e^{\gamma/k}, \tag{32}$$

where γ is the Euler constant, i.e. $\gamma \approx 0.57721\ldots$, and monotonically decreasing in k, for fixed α, with

$$\lim_{k\to\infty} K_k(\alpha) = 1. \tag{33}$$

Remarks: (i) The same asymptotic result holds for the case of the continuous time PAC policy that was recently derived in Theorem 1 of (8). (ii) After computing the second limit in (32) for $k = 1,2,3$, we notice a significant improvement in performance of the DPAC (m,k) algorithm when compared to the LRU policy ($k = 1$). Observe that already for $k = 3$, the DPAC policy performs approximately within 8% of the optimal static algorithm ($K_k(\infty) \approx 1.08$), which implies near optimality of the DPAC rule even for small values of k.

Proof: The proofs of monotonicity of $K_k(\alpha)$ and limits (32) and (33) can be found in Theorem 1 of (8).

Next, we prove the upper bound for the asymptotic relationship in (30). Define the sum of indicator functions $S(t) \triangleq \sum_{j=1}^{\infty} 1[T_j < t]$; note that $S(t)$ is a.s. non-decreasing in t, i.e. $S(t) \le S(t_0(x))$ a.s. for all $t \le t_0(x)$, where $t_0(x)$ is as defined in Lemma 5. Then, after conditioning on T_i being larger or smaller than $t_0(x)$, the expression in (6) can be upper bounded as

$$\begin{aligned}\mathbb{P}[C > x] &\le \mathbb{P}[S(t_0(x)) > x] + \sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge t_0(x)] \\ &\le \mathbb{P}[\sum_{i \ge i_0} 1[T_i < t_0(x)] \ge x - i_0] + \sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge t_0(x)], \end{aligned} \tag{34}$$

where in the previous expression we applied $\sum_{i=1}^{\infty} q_i = 1$ and i_0 can be an arbitrary, finite integer. Then, applying Lemma 5 we obtain that the tail of the search cost C is upper bounded by

$$\mathbb{P}[C > x] \le o\left(\frac{1}{x^{\alpha-1}}\right) + \sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge t_0(x)]. \tag{35}$$

Note that bound derived in Lemma 2 holds for i_0 sufficiently large. Then, due to the Poisson decomposition theorem, processes of requests for different documents i are Poisson as well and, therefore, since $q_i \ge q_{i_0}$ for $i \le i_0$, Poisson process of rate q_i can be constructed by superposition of Poisson processes with rates q_{i_0} and $q_i - q_{i_0}$. Thus, it is not hard to conclude that $T_i \le T_{i_0}$ (a.s.), implying

$$\mathbb{P}[T_i \ge t] \le \mathbb{P}[T_{i_0} \ge t]. \tag{36}$$

Therefore, we obtain

$$\begin{aligned}\sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge t_0(x)] &\le \sum_{i=1}^{i_0} q_i \mathbb{P}[T_{i_0} \ge t_0(x)] + \sum_{i=i_0}^{\infty} q_i e^{-q_i^k (1-\varepsilon)^2 \binom{m-1}{k-1} t_0(x)} + \sum_{i=i_0}^{\infty} m q_i e^{-h\varepsilon q_i^{k-1} t_0(x)} \\ &\triangleq I_1(x) + I_2(x) + I_3(x), \end{aligned} \tag{37}$$

where in the last two sums we used the result of Lemma 2.

After using the bound (8) and replacing $t_0(x)$, it immediately follows that

$$I_1(x) \le \sum_{i=1}^{i_0} q_i \left[e^{-q_{i_0}^k (1-\varepsilon)^2 t_0(x) \binom{m-1}{k-1}} + m e^{-h\varepsilon q_{i_0}^{k-1} t_0(x)} \right] = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \to \infty. \tag{38}$$

Now, by assumption of the theorem, for all i large enough ($i \ge i_0$, where i_0 is possibly larger than in (37))

$$(1-\varepsilon)c/i^{\alpha} < q_i < (1+\varepsilon)c/i^{\alpha}. \tag{39}$$

Furthermore, for i large enough ($i \ge i_0$) inequality $c/i^{\alpha} \le (1+\varepsilon)c/u^{\alpha}$ holds for any $u \in [i, i+1]$ and, therefore, using this bound, (39), the monotonicity of the exponential function and replacing $t_0(x)$ from Lemma 5, yields

$$\begin{aligned} I_2(x) &\le (1+\varepsilon) \sum_{i=i_0}^{\infty} \frac{c}{i^{\alpha}} e^{-\iota(\varepsilon) \frac{x^{\alpha k}}{[\Gamma(1-\frac{1}{\alpha k})]^{\alpha k} i^{\alpha k}}} \\ &\le (1+\varepsilon)^2 \int_1^{\infty} \frac{c}{u^{\alpha}} e^{-\iota(\varepsilon) \frac{x^{\alpha k}}{[\Gamma(1-\frac{1}{\alpha k})]^{\alpha k} u^{\alpha k}}} du, \end{aligned} \tag{40}$$

where $\iota(\varepsilon) \triangleq (1+\varepsilon)^{-(k+1)}(1-\varepsilon)^{2(k+1)}(1-2\varepsilon)^{\alpha k}$. Next, applying the change of variable method for evaluating the integral with $z = x^{\alpha k}\iota(\varepsilon)[\Gamma(1-\frac{1}{\alpha k})]^{-\alpha k}u^{-\alpha k}$, we obtain that the integral in (40) is equal to

$$\frac{c}{x^{\alpha-1}(\alpha-1)}\left[\Gamma\left(1-\frac{1}{\alpha k}\right)\right]^{\alpha-1}(\iota(\varepsilon))^{\frac{1}{\alpha k}-\frac{1}{k}}\frac{\alpha-1}{\alpha k}\int_0^{\frac{x^{\alpha k}\iota(\varepsilon)}{[\Gamma(1-\frac{1}{\alpha k})]^{\alpha k}}} e^{-z}z^{\frac{1}{k}-\frac{1}{\alpha k}-1}dz,$$

which, in conjunction with (40), implies

$$\limsup_{x\to\infty}\frac{I_2(x)}{\mathbb{P}[R>x]} \le K_k(\alpha)(\iota(\varepsilon))^{\frac{1}{\alpha k}-\frac{1}{k}} \to K_k(\alpha) \text{ as } \varepsilon\to 0, \tag{41}$$

where $K_k(\alpha)$ is defined in (31).

In order to estimate the asymptotics of $I_3(x)$, we use analogous steps to those we applied in evaluating $I_2(x)$. Thus, using inequalities from (39), $c/i^\alpha \le (1+\varepsilon)c/u^\alpha$ hold for all $u\in[i,i+1]$ and replacing $t_0(x)$, we obtain

$$\begin{aligned} I_3(x) &\le m(1+\varepsilon)\sum_{i=i_0}^{\infty}\frac{c}{i^\alpha}e^{-h\varepsilon\frac{x^{\alpha k}}{i^{\alpha(k-1)}}} \\ &\le m(1+\varepsilon)^2\int_1^\infty \frac{c}{u^\alpha}e^{-h\varepsilon\frac{x^{\alpha k}}{u^{\alpha(k-1)}}}du. \end{aligned} \tag{42}$$

Now, if $k=1$, it is straightforward to compute the integral in the preceding expression and obtain $I_3(x) \le m(1+\varepsilon)^2(c/(\alpha-1))e^{-h\varepsilon x^\alpha} = o(1/x^{\alpha-1})$ as $x\to\infty$. Otherwise, for $k\ge 2$, after using the change of variable method for solving the integral in (42) with $z = h\varepsilon x^{\alpha k}u^{-\alpha(k-1)}$, we obtain, as $x\to\infty$

$$I_3(x) \le m(1+\varepsilon)^3\frac{c}{(h\varepsilon)^{\frac{1}{k-1}(1-\frac{1}{\alpha})}}\frac{1}{\alpha(k-1)}\frac{1}{x^{\frac{k}{k-1}(\alpha-1)}}\Gamma\left(\frac{1}{k-1}-\frac{1}{\alpha(k-1)}\right) = o\left(\frac{1}{x^{\alpha-1}}\right). \tag{43}$$

Therefore (43), (41),(38), (37) and (35), yields, as $x\to\infty$

$$\mathbb{P}[C>x] \lesssim K_k(\alpha)\mathbb{P}[R>x]. \tag{44}$$

For the lower bound on $\mathbb{P}[C>x]$, starting from (6), we derive

$$\begin{aligned} \mathbb{P}[C>x] = \sum_{i=1}^{\infty}q_i\mathbb{P}[S_i(T_i)\ge x] &\ge \sum_{i=1}^{\infty}q_i\mathbb{P}[S(t_0(x))\ge x-1, T_i\ge t_0(x)] \\ &\ge \sum_{i=i_0}^{\infty}q_i\mathbb{P}[T_i\ge t_0(x)] - \mathbb{P}[S(t_0(x))\le x-1], \end{aligned}$$

where we choose $t_0(x)$ as in Lemma 4. Next, we apply Lemma 4 to estimate the second term in the preceding expression. Then, after applying Lemma 3 to lower bound the tail of random times T_i, $i\ge i_0$, in conjunction with the analogous reasoning as estimating $I_2(x)$ (see also expressions (39)-(40) of (8)), we complete the proof of this theorem. $\diamond$

In addition to the previous result, one can estimate the asymptotic performance of the DPAC algorithm when $0\le\alpha\le 1$. In this case our proofs would involve analogous arguments to those applied in the proof of preceding Theorem!1, in conjunction with techniques developed in Theorems 2 and 3 of (8). Since the results are the same as those in Theorems 2 and 3 of (8), we omit their proofs and just state them here.

Theorem 2 *Assume that $q_i = h_N/i$, $1\le i\le N$, where h_N is the normalization constant. Then, for any $0<\delta<1$, as $N\to\infty$*

$$(\log N)\mathbb{P}[C^{(N)}>\delta N] \sim F_k(\delta) \triangleq \frac{1}{k}\Gamma(0,\eta_\delta), \tag{45}$$

where η_δ uniquely solves the equation

$$1-\frac{1}{k}\eta^{\frac{1}{k}}\Gamma\left(-\frac{1}{k},\eta\right) = \delta;$$

note that, $\Gamma(x,y)$, $y > 0$, *is the incomplete Gamma function, i.e.* $\Gamma(x,y) = \int_y^\infty e^{-t}t^{x-1}dt$. *Furthermore, for any* $0 < \delta < 1$,

$$\lim_{k\to\infty} F_k(\delta) = \log\left(\frac{1}{\delta}\right). \tag{46}$$

$\diamond$

Theorem 3 *Assume that* $q_i = h_N/i^\alpha$, $1 \le i \le N$, *where* h_N *is the normalization constant and* $0 < \alpha < 1$. *Then, for any* $0 < \delta < 1$, *as* $N \to \infty$

$$\mathbb{P}[C^{(N)} > \delta N] \sim F_k(\delta) \triangleq \frac{1-\alpha}{\alpha k}(\eta_\delta)^{\frac{1}{\alpha k}-\frac{1}{k}}\Gamma\left(\frac{1}{k}-\frac{1}{\alpha k},\eta_\delta\right), \tag{47}$$

where η_δ *is the unique solution of the equation*

$$1-\frac{1}{\alpha k}\Gamma\left(-\frac{1}{\alpha k},\eta\right)\eta^{\frac{1}{\alpha k}} = \delta;$$

note that $\Gamma(x,y)$, $y > 0$, *is the incomplete Gamma function, i.e.* $\Gamma(x,y) = \int_y^\infty e^{-t}t^{x-1}dt$. *Furthermore,* $F_k(\delta)$, $\delta \in (0,1)$, *is a proper distribution, with* $\lim_{\delta\to 0} F_k(\delta) = 1$, $\lim_{\delta\to 1} F_k(\delta) = 1$ *and*

$$\lim_{k\to\infty} F_k(\delta) = 1-\delta^{1-\alpha}. \tag{48}$$

$\diamond$

5 Numerical experiments

In this section we illustrate our main results stated in Theorems 1, 2 and 3. Even though the results are valid for large cache sizes, our simulations show that the fault probability approximations, suggested by formula (30), (45) and (47), work very well for small caches as well.

5.1 Convergence to stationarity

In the presented experiments, we use a discrete time model without The Poisson embedding, i.e., $\tau_n = n$. In order to ensure that the simulated values of the fault probabilities do not deviate significantly from the stationary ones, we first estimate the difference between the distributions of $C^{(N)}$ and $C_n^{(N)}$, where $C_n^{(N)}$ is the search cost after n requests with arbitrary initial conditions.

Thus, using(2 – 3), we can upper bound the difference between the tails of these distributions as

$$\sup_x \left|\mathbb{P}[C_n^{(N)} > x] - \mathbb{P}[C^{(N)} > x]\right| \le e_n \triangleq \sum_{i=1}^N q_i \mathbb{P}\left[T_i \ge n-m+1\right].$$

Now, using similar argument as in Lemma 2, we obtain

$$\begin{aligned}\mathbb{P}[T_i > t] &\le \mathbb{P}(\tau_{X_i}^{(i)} + (m-1)X_i > t)\\ &\le \mathbb{P}[\tau_{X_i}^{(i)} > \frac{t}{2}] + \mathbb{P}[(m-1)X_i > \frac{t}{2}],\end{aligned} \tag{49}$$

where now $\{\tau_n^{(i)}\}$ denotes success times in a Bernoulli process with parameter q_i and X_i is independent of $\{\tau_n^{(i)}\}$ with geometric distribution having parameter

$$p_i = \mathbb{P}[M_i(\tau_1^{(i)} \ge k] = \sum_{l=k-1}^{m-1}\binom{m-1}{l} q_i^l (1-q_i)^{m-1-l}.$$

Next, it is well known that $\tau_{X_i}^{(i)}$ is geometric with parameter $q_i p_i$ and, therefore, from (49)

$$\mathbb{P}[T_i > t] \le (1-p_i q_i)^{\frac{t}{2}} + (1-p_i)^{\frac{t}{2(m-1)}}.$$

Thus, using the preceding bound, we obtain

$$e_n \leq \sum_{i=1}^{N} q_i \left[(1 - p_i q_i)^{\frac{n-m}{2}} + (1 - p_i)^{\frac{n-m}{2(m-1)}} \right]. \tag{50}$$

Note that, since p_i is increasing in m, the larger values of m speed up the convergence of the search cost process $\{C_n^{(N)}\}$ to stationarity. In other words, increase of m makes the algorithm more adaptable. On the other hand, the larger m implies the larger size of the additional storage needed to keep track of the past requests. Thus, although the stationary performance of the DPAC algorithm is invariant to m, this parameter provides an important design component whose choice has to balance between the algorithm complexity and adaptability.

5.2 Experiments

In the presented experiments we take the number of documents to be $N = 1300$ with popularities satisfying $q_i = h_N / i^{\alpha}$, $1 \leq i \leq 1300$, where $h_N = (\sum_{i=1}^{N} 1/i^{\alpha})^{-1}$. Also, we select $m = 20$ and α: **1)** $\alpha = 1.4$, **2)** $\alpha = 1$ and **3)** $\alpha = 0.8$. The initial permutation of the list is chosen uniformly at random and the set of initial m requests is taken to be empty. The fault probabilities are measured for cache sizes $x = 50j$, $1 \leq j \leq 15$. Simulation results are presented with "*" symbols on Figures 1,2, and 3, while the optimal static performance is presented with a thick solid line on the same figures.

In our first experiment, we illustrate Theorem 1. Since our asymptotic formula is obtained for infinite number of documents N, it can be expected that asymptotic expression gives reasonable approximation of the fault probability $\mathbb{P}[C^{(N)} > x]$ only if both N and x are large (with N much larger than x). However, our experiment shows that the obtained approximation works well for relatively small values of N and almost all cache sizes $x < N$.

Experiment 1: Here we select $\alpha = 1.4$. Before we conduct measurements, we allow the time of the first $n = 10^{10}$ requests to be a warm-up time for the system to reach its stationarity. Then, the actual measurement time is also set to be 10^{10} requests long. We measure the cache fault probabilities of the DPAC$(20, k)$ policy for values $k = 1, 2$. The experimental results for the cases when $k \geq 3$ are almost indistinguishable from the optimal performance, $\mathbb{P}[R > x]$, and, thus, we do not present them on Figure 1. After estimating e_n in (50) for a given warm-up time of 10^{10} requests, we obtain that $e_n < 10^{-12}$, which is negligible compared to the smallest measured probabilities ($> 10^{-2}$). Therefore, the measured fault probabilities are essentially the stationary ones. The accuracy of approximation $P^{(e)}(x)$ and the improvement in performance is apparent from Figure 1.

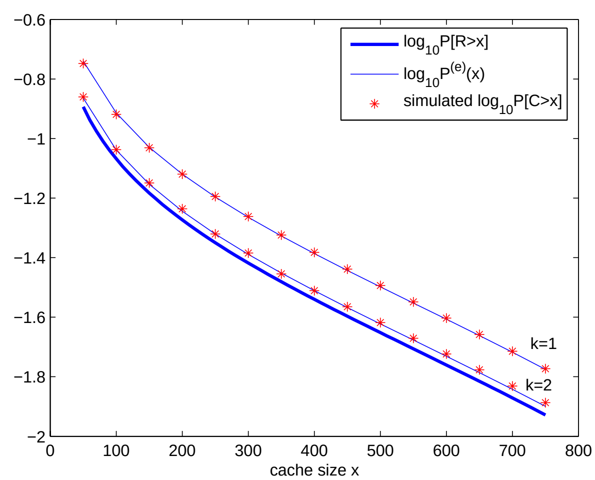

Fig. 1: Illustration for Experiment 1

Experiment 2: Here, we set $\alpha = 1$ and measure the cache fault probabilities for $k = 1, 2, 3$. Before we conduct measurements, we allow the time of the first n requests to be a warm-up time for the system to reach its stationarity; Then, the actual measurement time is also set to be n requests long. We set $n = 2\times 10^8$ for $k = 1,2$ and $n = 10^{11}$ for $k = 3$. Since the normalization constant $\frac{1}{h_N} = \log N + \gamma + o(1)$ as $N \to \infty$, where γ is the Euler's constant, the ratio $h_N \log N$ converges slowly to one and, therefore, instead of using the approximation $\mathbb{P}[C^{(N)} > x] \approx \frac{F_k(x/N)}{(\log N)}$, as suggested by Theorem 2 of (8), we define $\mathbb{P}^{(e)}(x) = h_N F_k(x/N)$. We obtain that for $k = 1,2$, $e_n < 3\times 10^{-11}$, while for $k = 3$, $e_n < 2\times 10^{-6}$, which is insignificant compared to the smallest measured probabilities. Thus, the process is basically stationary. The accuracy of approximation $P^{(e)}(x)$ and the improvement in performance is apparent from Figure 2.

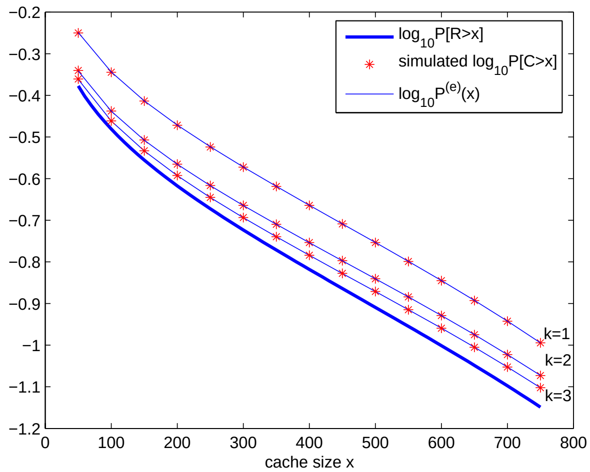

Fig. 2: Illustration for Experiment 2

Experiment 3: Finally, the third example assumes $\alpha = 0.8$ and considers cases $k = 1,2,3$. Before we conduct measurements, we allow the time of the first $n = 10^{10}$ requests to be a warm-up time for the system to reach its stationarity. Then, the actual measurement time is also set to be 10^{10} requests long. Similarly as in the case of $\alpha = 1$, due to the slow convergence of $h_N N^{1-\alpha}/(1-\alpha)$ to one as $N \to \infty$, we use an estimate $\mathbb{P}^{(e)}(x) = h_N(N^{1-\alpha}/(1-\alpha))F_k(x/N)$ instead of $F_k(x/N)$ that can be inferred from Theorems 3 of (8). We compute $e_n < 3\times 10^{-5}$, which is insignificant compared to the smallest measured probabilities. Thus, the process is basically stationary. Once again, the validity of approximation $P^{(e)}(x)$ and the benefit of the DPAC algorithm is evident from Figure 3.

6 Concluding remarks

In this paper we introduce a discrete version of the recently proposed continuous time PAC replacement rule (8) that possesses all desirable properties of the LRU policy, such as low complexity, ease of implementation and adaptability to variable Web access patterns. In addition to these attributes, the new DPAC policy eliminates drawbacks of the PAC rule, such as its dependence on the request arrival times and variable storage requirements. However, even in the case of the independent reference model, the DPAC policy is significantly harder to analyze than the continuous PAC rule. In this respect, we develop a new analytic technique that allows us to decouple replacement decisions of the requested documents and show that the performance (fault probability) of the DPAC algorithm, for large cache sizes, is very close to the optimal frequency algorithm even for small values of $k = 2,3$, which implies negligible additional complexity relative to the classical LRI policy. In addition, theoretical result is further validated using simulation that shows a significant improvement of the DPAC algorithm in comparison to the LRU scheme, even for small cache sizes x and the number of documents N. The excellent agreement between the analytical and experimental results and simplicity of the DPAC policy itself implies a very high potential of the proposed policy for practical purposes.

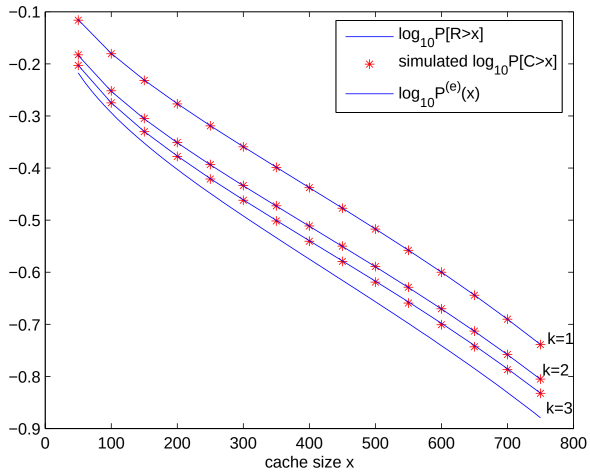

Fig. 3: Illustration for Experiment 3

Finally, we would like to note that earlier proposed "k-in-a-row" rule ((9; 4)) that was studied in the context of the expected list search cost, is a special case of the DPAC(m,k) algorithm when $m = k$.

7 Proof of Lemma 5

The case $k = 1$ that corresponds to the ordinary LRU algorithm is easy since the variables T_i are independent and exponentially distributed with parameters q_i. Thus, the result follows from Lemmas 3 and 4 of (6). Hence, in the rest of the proof we assume $m \geq k \geq 2$.

Note that for any $i_0 \geq 1$

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < t_0(x)] \geq x\right] \leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < t_0(x)] \geq x - i_0\right]; \tag{51}$$

a specific i_0 will be selected later in the proof. Let $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1}$ be an ordered sequence of request times for documents $i \geq i_0$, i.e. $\{\tau_j^{(\overline{i_0})}\}_{j\geq 1} = \cup_{i\geq i_0}\{\tau_n^{(i)}\}_{n\geq 1}$. We use $n(j),\ j \geq 1$ to denote the index of point $\tau_j^{(\overline{i_0})}$ in the original sequence $\{\tau_n\}_{n\geq 1}$, i.e. $\tau_j^{(\overline{i_0})} = \tau_{n(j)}$. Then, since process $\{\tau_n\}$ is Poisson and $\{R_n\}$ is i.i.d. sequence independent of $\{\tau_n\}$, by Poisson decomposition theorem, process $\{\tau_{n(j)} \equiv \tau_j^{(\overline{i_0})}\}_{j\geq 1}$ is also Poisson with rate $\sum_{i\geq i_0} q_i$.

Next, similarly as in Lemma 3, we group the points $\{\tau_n\}$ into cycles. The first cycle C_1, will be closed interval of time that starts with $\tau_{n(1)}$ and its length is determined by the following procedure. Let random variable Z_1 be defined as

$$Z_1 \triangleq \inf\{j > 0 : M_{\overline{i_0}}(\tau_{n(1)+(j-1)m+1}) = 0\},$$

where $M_{\overline{i_0}}(\tau_n) \triangleq \sum_{i\geq i_0} M_i(\tau_n)$. In other words, we observe groups of m consecutive requests until we come to a group of m requests where there are no requests for documents $i \geq i_0$. Then, the first cycle, C_1 will be an interval $[\tau_{n(1)}, \tau_{n(1)+mZ_1}]$. Next, starting from the first point of process $\{\tau_{n(j)}\}_{j\geq 1}$ after request $\tau_{n(1)+mZ_1}$, say $\tau_{n(l)}$, we define

$$Z_2 \triangleq \inf\{j > 0 : M_{\overline{i_0}}(\tau_{n(l)+(j-1)m+1}) = 0\},$$

and, therefore, the second cycle is interval $C_2 = [\tau_{n(l)}, \tau_{n(l)+mZ_2}]$. We continue this procedure indefinitely.

Then, denote the points of time that represent the beginnings of the previously defined cycles $C_j,\ j \geq 1$, by $\{\Theta_j\}_{j\geq 1}$. Clearly, from the independence assumptions on $\{\tau_n\}$ and $\{R_n\}$, $\{\Theta_j\}$ is a renewal process with renewal intervals, for $j \geq 1$, satisfying

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(\overline{i_0})} + \sum_{i=1}^{mZ_1} \xi_i,$$

where $\{\xi_i\}$ is an i.i.d. sequence of exponential random variables with mean 1 that is independent from $\tau_1^{(\overline{i_0})}$ and Z_1. Thus, by neglecting the sum in the preceding expression (i.e. the length of the cycles), the beginning of each cycle can be lower bounded with

$$\Theta_j \overset{d}{\geq} \tau_j^{(\overline{i_0})}. \tag{52}$$

Next, on each cycle $C_j,\ j \geq 1$, define an event that at least two distinct items are moved to the first position of the list during that cycle

$$A_j^{(0)} \triangleq \{\omega : \exists i_1, i_2 \geq i_0, i_1 \neq i_2, \exists \tau_{n_1}, \tau_{n_2} \in C_j, R(\tau_{n_l}) = i_l, M_{i_l}(\tau_{n_l}) \geq k \text{ for } l = 1,2\}.$$

Similarly, for each $i \geq i_0$ we define an event that exactly one document i (but no other documents) is moved to the first position of the list

$$A_j^{(i)} \triangleq \{\omega : \exists \tau_n \in C_j, R(\tau_n) = i, M_i(\tau_n) \geq k\} \cap \left(A_j^{(0)}\right)^c,$$

where A^c denotes a complement of event A. Then, for each fixed j these events are disjoint, and, due to the independence properties of our reference model, they are independent on different cycles and for fixed i equally distributed; let $p_i \triangleq \mathbb{P}[A_1^{(i)}], i \geq i_0$ or $i = 0$.

Now, using the bound in (52) and the Poisson decomposition theorem, it is easy to see that, for each fixed i, the beginning of the first cycle where event $A_j^{(i)}$ happens is lower bounded by L_i, where L_i are independent exponential random variables with parameters equal to $p_i \sum_{i \geq i_0} q_i$. Then, for $i \geq i_0$, the random times defined in (7) are lower bounded by the beginning of the first cycle where event $A_j^{(0)} \cup A_j^{(i)}$ occurs, which is further lower bounded by

$$T_i \geq L_i \wedge L_0, \tag{53}$$

where $x \wedge y = \min(x,y)$.

Next, we provide upper bounds on each of the probabilities p_i. Using the same arguments as in (20) of Lemma 3, we obtain that for any $\varepsilon > 0$, we can chose i_0 large enough, such that for all $i \geq i_0$

$$p_i \leq (1+\varepsilon)\frac{q_i}{\sum_{j=i_0}^{\infty} q_j}\binom{m-1}{k-1} q_i^{k-1}. \tag{54}$$

The probability p_0 can be bounded as

$$\begin{aligned}
p_0 &\leq \mathbb{P}(Z_j > l) + \mathbb{P}(A_j^{(0)}, Z_j \leq l) \\
&\leq H\left(\sum_{j=i_0}^{\infty} q_j\right)^l + \mathbb{P}[A_j^{(0)}, Z_j \leq l] \\
&\leq H\left(\sum_{j=i_0}^{\infty} q_j\right)^l + \sum_{j_1, j_2 \geq i_0, j_1 \neq j_2} \mathbb{P}[A_j^{(0)}(j_1, j_2), Z_j \leq l],
\end{aligned} \tag{55}$$

where l is a fixed constant that will be selected later. $A_j^{(0)}(j_1, j_2)$ is the event that during cycle C_j documents j_1 and j_2 are moved to the first position of the list.

$$\begin{aligned}
\mathbb{P}\left[A_1^{(0)}(j_1, j_2), Z_1 \leq l\right] &= \mathbb{P}\left[R(\Theta_1) = j_1, A_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\
&\quad + \mathbb{P}\left[R(\Theta_1) = j_2, A_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\
&\quad + \mathbb{P}\left[R(\Theta_1) \neq j_1, j_2, A_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\
&\triangleq p_{01}(j_1, j_2) + p_{02}(j_1, j_2) + p_{03}(j_1, j_2).
\end{aligned} \tag{56}$$

Now, we upper bound the first term of (56),

$$\begin{aligned}
p_{01}(j_1,j_2) &\le \frac{q_{j_1}}{\sum_{j=i_0}^{\infty} q_j}\mathbb{P}\left[M_{j_1}(\tau_{n(1)+1},ml)\ge k-1, M_{j_2}(\tau_{n(1)+1},ml)\ge k\right]\\
&\le \frac{q_{j_1}}{\sum_{j=i_0}^{\infty} q_j}\mathbb{P}[M_{j_1}(\tau_1,ml)\ge k-1, M_{j_2}(\tau_1,ml)\ge k]\\
&\le \frac{q_{j_1}}{\sum_{j=i_0}^{\infty} q_j}\left[\sum_{l_1=k-1}^{ml-k}\sum_{l_2=k}^{ml-l_1}\frac{(ml)!}{l_1!l_2!(ml-l_1-l_2)!}q_{j_1}^{l_1}q_{j_2}^{l_2}(1-q_{j_1}-q_{j_2})^{ml-l_1-l_2}\right]\\
&\le \frac{q_{j_1}}{\sum_{j=i_0}^{\infty} q_j}(Hq_{j_1}^{k-1}q_{j_2}^{k});
\end{aligned}$$

the last inequality holds because ml is fixed and finite, $j_1, j_2 \ge i_0$, and i_0 is large enough. Thus, we obtain

$$p_{01}(j_1,j_2) \le \frac{Hq_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}. \tag{57}$$

Similarly, we derive

$$p_{02}(j_1,j_2) \le \frac{Hq_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j} \tag{58}$$

and, by applying the same type of arguments, we bound

$$p_{03}(j_1,j_2) \le Hq_{j_1}^k q_{j_2}^k \le \frac{Hq_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}. \tag{59}$$

Therefore, (57),(58),(59) and (56) give that for any $j_1, j_2 \ge i_0, j_1 \ne j_2$,

$$\mathbb{P}[A_j^{(0)}(j_1,j_2), Z_j \le l] \le H\frac{q_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}, \tag{60}$$

where constant H is independent of j_1 and j_2. Now, by replacing the preceding bound in (55), we derive that for all i_0 large enough,

$$p_0 \le H\left(\sum_{j=i_0}^{\infty} q_j\right)^l + \frac{H}{(i_0-1)^{2\alpha k-2}\sum_{j=i_0}^{\infty} q_j}. \tag{61}$$

After setting the necessary ground for our analysis, we upper bound the left hand side of (51) as

$$\begin{aligned}
\mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < t_0(x)] \ge x-i_0\right] &\le \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i\wedge L_0 < t_0(x)] \ge x-i_0\right]\\
&\le \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i\wedge L_0 < t_0(x), L_0 > t_0(x)] \ge x-i_0\right] + \mathbb{P}[L_0 \le t_0(x)]\\
&\le \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i < t_0(x)] \ge x-i_0\right] + \mathbb{P}[L_0 \le t_0(x)].
\end{aligned} \tag{62}$$

Now, from (54), $\mathbb{P}[L_i < t_0(x)] \le 1 - e^{-(1+\varepsilon)\binom{m-1}{k-1}q_i^k t_0(x)}$ for $i \ge i_0$ and i_0 large enough. Furthermore, assigning $i_0 = \lceil \varepsilon x\rceil$ and applying Lemma 3 of (6), we derive as $x \to \infty$

$$\mathbb{E}\left[\sum_{i\ge\varepsilon x} 1[L_i < t_0(x)]\right] \lesssim \Gamma\left(1-\frac{1}{\alpha k}\right) c^{\frac{1}{\alpha}} \left(\binom{m-1}{k-1}\right)^{\frac{1}{\alpha k}} (1+\varepsilon)^{\frac{k+1}{\alpha k}} t_0(x)^{\frac{1}{\alpha k}}. \tag{63}$$

Then, if we replace $t_0(x)$ and use (63), it follows that $\mathbb{E}\left[\sum_{i\geq\varepsilon x} 1[L_i < t_0(x)]\right] \lesssim (1-2\varepsilon)x = (1-\frac{\varepsilon}{1-\varepsilon})(x-\varepsilon x)$ for all x large enough. Then, since L_i, $i \geq i_0$, are mutually independent, using large deviation result from Lemma 4 of (6), we show that the first term in (62) is bounded, for some $\theta > 0$, by

$$\mathbb{P}\left[\sum_{i\geq\varepsilon x} 1[L_i < t_0(x)] \geq x - \varepsilon x\right] \leq 2e^{-\theta x} = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \to \infty. \tag{64}$$

Next, we estimate the second term of (62). Using (61) with $i_0 = \lceil \varepsilon x \rceil$ and choosing $l = \lceil \frac{2\alpha k-2}{\alpha-1} \rceil - 1$, such that $(\alpha-1)(l+1) \geq 2\alpha k - 2$, we derive

$$\begin{aligned}\mathbb{P}[L_0 \leq t_0(x)] &\leq 1 - e^{-t_0(x)\, p_0 \sum_{i\geq\varepsilon x} q_i} \\ &\leq Hx^{\alpha k}\left[\frac{1}{x^{(\alpha-1)(l+1)}} + \frac{1}{x^{2\alpha k-2}}\right] \\ &\leq \frac{Hx^{\alpha k}}{x^{2\alpha k-2}} \leq \frac{H}{x^{\alpha k-2}} = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \to \infty. \end{aligned} \tag{65}$$

since $k \geq 2$ and $\alpha > 1$. Finally, replacing (65) and (64) in (62) imply the statement of the lemma. $\diamond$

References

[1] O. Bahat and A. M. Makowski. Optimal replacement policies for non-uniform cache objects with optional eviction. In *Proceedings of Infocom 2003*, San Francisco, California, USA, April 2003.

[2] E. Cinlar. *Introduction to Stochastic Processes*. Prentice–Hall, 1975.

[3] J. A. Fill and L. Holst. On the distribution of search cost for the move-to-front rule. *Random Structures and Algorithms*, 8(3):179, 1996.

[4] G. H. Gonnet, J. I. Munro, and H. Suwanda. Exegesis of self-organizing linear search. *SIAM J. Comput.*, 10(3):613–637, 1981.

[5] P. R. Jelenković. Asymptotic approximation of the move-to-front search cost distribution and least-recently-used caching fault probabilities. *Annals of Applied Probability*, 9(2):430–464, 1999.

[6] P. R. Jelenković and A. Radovanović. Least-recently-used caching with dependent requests. *Theoretical Computer Science*, 326:293–327, 2004.

[7] P. R. Jelenković and A. Radovanović. Optimizing LRU for variable document sizes. *Combinatorics, Probability & Computing*, 13:1–17, 2004.

[8] P. R. Jelenković and A. Radovanović. The Persistent-Access-Caching Algorithm. Technical Report EE2004-03-05, Department of Electrical Engineering, Columbia University, New York, April 2004.

[9] Y. C. Kan and S. M. Ross. Optimal list order under partial memory constraints. *Journal of Applied Probability*, 17:1004–1015, 1980.

[10] Elizabeth J. O'Neil, Patrick E. O'Neil, and Gerhard Weikum. An optimality proof of the LRU-K page replacement algorithm. *Journal of the ACM*, 46:92–112, 1999.