Near optimality of the discrete persistent access caching algorithm

Predrag R. Jelenković[1†] and Xiaozhu Kang[1‡] and Ana Radovanović[2]

[1]*Department of Electrical Engineering, Columbia University, NY 10025, USA. Email:* $\{$*predrag, xiaozhu*$\}$*@ee.columbia.edu*
[2]*Ana Radovanović, IBM Thomas J. Watson Research Center, Mathematical Sciences Department, NY 10598, USA. Email: aradovan@us.ibm.com*

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, the most desirable attributes of a cache replacement scheme 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 Jelenković and Radovanović (2004c), 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

1 Introduction

Since the recent invention of the World Wide Web, there has been an explosive growth in distributed multimedia content and services that are now an 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

[†]Supported by NSF Grant No. 0092113.
[‡]Supported by NSF Grant No. 0117738.

mentioned attributes and, therefore, represents a basis for designing many practical replacement algorithms. However, as shown in Jelenković (1999) 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 most frequently used documents in the cache, i.e., replaces Least-Frequently-Used (LFU) items. On the other hand, the drawbacks of the LFU algorithm are its need to know (measure) the document access frequencies and employ aging schemes based on reference counters in order to cope with evolving access patterns, which result in high complexity. In the context of database disk buffering, O'Neil et al. (1999) 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 O'Neil et al. (1999) 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 Jelenković and Radovanović (2004c), we designed a new cache replacement policy, termed the Persistent Access Caching (PAC) rule, that is essentially preserving 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 with the widely used LRU policy. However, the drawback of the PAC policy is that its implementation and analysis depend on the Poisson assumption on the request arrival times. 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 the coupling of the request decisions, as pointed out in the abstract, DPAC is significantly more difficult to analyze. To this end, we develop a new analytic technique, which, in conjunction with the large deviation analysis and asymptotic results developed in Jelenković and Radovanović (2004c,a); Jelenković (1999), shows that the performance of the DPAC policy is nearly 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 our asymptotic formulas 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 and develop a representation theorem for the stationary cache fault probability. This representation formula and lemmas of Section 3 provide necessary tools for proving our main result, stated in Theorem 1, in Section 4. Informally, this theorem shows that for large cache sizes, independent reference model and generalized Zipf's law request distributions with $\alpha > 1$, the fault probability of the DPAC algorithm approaches the optimal LFU policy while maintaining low implementation complexity. Furthermore, in Section 5, we provide an additional validation of our asymptotic approximations with simulation experiments. A brief discussion of our results is presented in Section 6. In order to alleviate the reading process, we present the proof of a technical lemma in Section 7.

2 Model description and preliminary results

Consider a set $L = \{1, 2, \ldots, N\}$ of $N \leq \infty$ documents (items) of unit size, out of which x documents can be stored in an easily accessible location, called cache. The remaining $N - x$ 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 $\tau_1 > 0$ a.s.. Furthermore, define i.i.d. random variables $R^{(N)}, \{R_n^{(N)}\}_{n\geq 1}$, independent from $\{\tau_n\}_{n\geq 1}$, where $\{R_n^{(N)} = i\}$ represents a request for item i at time τ_n. We denote the request probabilities as $q_i^{(N)} \triangleq \mathbb{P}[R^{(N)} = i]$ and, unless explicitly required for clarity, we omit the superscript N and simply write R, R_n, q_i; also, without loss of generality, we assume that $q_1 \geq q_2 \geq \ldots$.

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 $m \geq k = 1$, DPAC reduces to the ordinary LRU heuristic. Also, the earlier proposed "k-in-a-row" rule (Kan and Ross (1980); Gonnet et al. (1981)), that was studied in the context of the expected list search cost, is a special case of the DPAC(m,k) algorithm when $m = k$. 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 the literature on combinatorial (competitive, worst case) approach can be found in Jelenković and Radovanović (2004c).

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 in the list of the document requested 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 search 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 time $\tau_0 = 0$ with initial conditions given by an arbitrary initial permutation Π_0 of the list and a sequence of the preceding $m-1$ requests $\mathcal{R}_0 = \{r_{-m+2}, r_{-m+2}, \ldots, r_{-1}, r_0\}$.

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, equal in distribution to R, that arrive at points $\{\tau_{-n}\}_{n\geq 0}$, $\tau_{-n} \leq 0$. These arrival points are constructed such that sequence $\{\tau_{n+1} - \tau_n\}_{-\infty<n<\infty}$ is stationary and ergodic; in particular, for every $n \geq 1$, $\{-\tau_k\}_{k=-n}^{-1}$ is equal in distribution to $\{\tau_k\}_{k=1}^{n}$. Then, for each $n > 0$, we construct a DPMTF(m,k) algorithm starting at τ_{-n} with the same initial condition as in the previous paragraph, given by Π_0 and $\mathcal{R}_0$, and having a sequence of requests $\{R_l\}_{l=-n+1}^{0}$; note that in this construction we assume that there is no request at time τ_{-n}. Let $C_{-n}^{(N)}$ be the search cost at time $\tau_0 = 0$ for the DPMTF(m,k) algorithm starting at 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 $\{T_i^{(-n)}\}_{i=1}^{N}$, $n \geq 1$, where $-T_i^{(-n)}$ represents the last time before $t = 0$ that item i was moved to the front of the list under 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 $\{T_i\}_{i=1}^{N}$ 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 $\{T_i^{(-n)} <$

$-\tau_{-n+m-1}\}$, $n > m-1$. 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

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

Next, observe that on event $\{R_0 = i, T_i^{(-n)} < -\tau_{-n+m-1}\}$, $n > m-1$, 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

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

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

Hence, by the 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 the stationarity and ergodicity of the arrival process $\{\tau_n\}$, τ_n/n satisfies the strong law of large numbers that, in conjunction with the a.s. finiteness of $T_i < \infty$, results in

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

Finally, equality of events $\{T_i^{(-n)} \geq -\tau_{-n+m-1}\} = \{T_i \geq \tau_{n-m+1}\}$, independence of requests, the preceding limit and the 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 $(\Pi_0, \mathcal{R}_0)$ and any $x \geq 0$, the search cost $C_n^{(N)}$ converges in distribution to $C^{(N)}$ as $n \to \infty$, where*

$$\mathbb{P}[C^{(N)} > x] \triangleq \sum_{i=1}^{N} q_i \mathbb{P}\left[S_i(T_i) \geq x\right] \tag{5}$$

and $S_i(t) \triangleq \sum_{j\neq i} 1[T_j < t]$.

Remark 1 (i) Note that the expression in (5) 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 (1), we define random times $\{T_i\}_{i=1}^N$, $\{T'_i\}_{i=1}^N$, 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.e., the sequences of random times $\{T_i\}$, $\{T'_i\}$ 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 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 $\{T_i\}_{i=1}^N$. In the context of the LRU policy with i.i.d. requests, the Poisson embedding technique was first used in Fill and Holst (1996) to derive a representation for the fault probability, as stated in (5), with independent random times $\{T_i\}$. The work in Fill and Holst (1996) provides a probabilistic proof of an earlier result from Flajolet et al. (1992) that was obtained using formal languages.

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 Lemmas 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 $S_i(t)$. Throughout this section we assume that $N = \infty$.

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

$$T_i \triangleq \inf\{\tau_n > 0 : R_n = i, M_i(\tau_n) \geq k\}. \tag{6}$$

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 $\{\tau_n^{(i)}\}_{n\geq 1}$, $i \geq 1$, are also Poisson with rate q_i and mutually independent for different 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. Also, we use 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.

The following two lemmas compute upper and lower bounds on the tail of the distribution of T_i for large i.

Lemma 2 *For any $0 < \epsilon < 1$, there exists i_0 and a fixed constant $h > 0$, such that for all $i \geq i_0$,*

$$\mathbb{P}[T_i > t] \leq e^{-\binom{m-1}{k-1}(1-\epsilon)^2 q_i^k t} + m e^{-h\epsilon q_i^{k-1} t}. \tag{7}$$

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

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(i)} + \sum_{l=1}^{m-1} \xi_l, \tag{8}$$

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

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

Note that this definition of U_i has identical form to the one for T_i in (6) since $R(\Theta_j) \equiv i$ is implied by $\{\Theta_j\} \subset \{\tau_j^{(i)}\}$. Therefore, given $\{\Theta_j\} \subset \{\tau_n\}$, it is clear that

$$T_i \leq U_i. \tag{9}$$

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 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 (8) we obtain

$$U_i = \Theta_{X'} \overset{d}{=} \tau_X^{(i)} + \sum_{j=1}^{(m-1)(X-1)} \xi_j \leq \tau_X^{(i)} + \sum_{j=1}^{(m-1)X} \xi_j, \tag{10}$$

where X is constructed such that $X \overset{d}{=} X'$ and it is independent of $\{\tau_n^{(i)}\}$, $\{\xi_j\}$.

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 well known (see Theorem 5.3, p. 89 of Cinlar (1975)) 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 $\sum_i q_i = 1$, it follows that for any $\epsilon > 0$, there exists i_0, such that for all $i \geq i_0$, q_i is small enough to ensure

$$p = \mathbb{P}[M_i(\tau_1^{(i)}) \geq k] = \sum_{l=k-1}^{m-1} \binom{m-1}{l} q_i^l (1-q_i)^{m-1-l} \geq \binom{m-1}{k-1} q_i^{k-1}(1-q_i)^{m-k} \geq (1-\epsilon)\binom{m-1}{k-1} q_i^{k-1}. \tag{11}$$

At this point, using the observations from the previous paragraph, (9) and (10), we obtain, for all i large enough ($i \geq i_0$),

$$\begin{aligned}
\mathbb{P}[T_i > t] &\leq \mathbb{P}[U_i > t] \\
&\leq \mathbb{P}\left[\tau_X^{(i)} > (1-\epsilon)t\right] + \mathbb{P}\left[\sum_{j=1}^{X(m-1)} \xi_j > \epsilon t\right] \\
&\leq e^{-pq_i(1-\epsilon)t} + (m-1)\mathbb{P}\left[\sum_{j=1}^{X} \xi_j > \frac{\epsilon t}{m-1}\right] \\
&\leq e^{-\binom{m-1}{k-1} q_i^k (1-\epsilon)^2 t} + (m-1)e^{-\frac{p\epsilon t}{m-1}} \\
&\leq e^{-\binom{m-1}{k-1} q_i^k (1-\epsilon)^2 t} + m e^{-h\epsilon q_i^{k-1} t};
\end{aligned} \tag{12}$$

this completes the proof. ◇

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

$$\mathbb{P}[T_i > t] \geq e^{-(1+\epsilon)\binom{m-1}{k-1} q_i^k t}. \tag{13}$$

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 $\mathcal{C}_1 = [\tau_{n(1)}, \tau_{n(1)+mZ_1-1}]$. 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.

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

$$\mathbb{P}[Z_i = j] = (\mathbb{P}[M_i(\tau_1) > 0])^{j-2}\mathbb{P}[M_i(\tau_1) = 0], \quad j \geq 2. \tag{14}$$

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

$$\Theta_{j+1} - \Theta_j \overset{d}{=} \tau_1^{(i)} + \sum_{l=1}^{mZ_1-1} \xi_l, \tag{15}$$

where $\{\xi_l\}$ is an i.i.d. sequence of exponential random variables with mean 1 that is independent of $\tau_1^{(i)}$, Z_1. Note that the sequences $\{\Theta_j\}_{j\geq 1}, \{n(j)\}_{j\geq 1}$ as well as the other auxiliary variables (e.g., X', X, p) 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.

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

$$\mathcal{A}_j \triangleq \{\omega : \exists \tau_n \in \mathcal{C}_j, R(\tau_n) = i, M_i(\tau_n) \geq k\};$$

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

$$\begin{aligned} T_i &= \inf\{\tau_n : R(\tau_n) = i, M_i(\tau_n) \geq k, \tau_n \in \mathcal{C}_j, j \geq 1\} \\ &\geq L_i \triangleq \inf\{\Theta_j(\omega) : \omega \in \mathcal{A}_j, j \geq 1\}, \end{aligned} \tag{16}$$

where the inequality is implied by $\tau_n \geq \Theta_j$ for any $\tau_n \in \mathcal{C}_j, j \geq 1$.

Now, we define

$$X' = \inf\{j \geq 1 : \exists \tau_n \in \mathcal{C}_j, R(\tau_n) = i, M_i(\tau_n) \geq k\}.$$

The independence of events $\mathcal{A}_j$ implies that X' has a geometric distribution $\mathbb{P}[X' = j] = (1-p)^{j-1}p, j \geq 1$ with p satisfying

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

where $M_i(\tau_n, k)$ is defined as the number of references for document i among the requests occurring at $\tau_n, \tau_{n+1}, \ldots, \tau_{n+k-1}$. Furthermore, using the preceding definitions, we arrive at

$$L_i = \Theta_{X'} \overset{d}{\geq} \tau_X^{(i)}, \tag{17}$$

where X is selected to be independent of $\{\tau_n^{(i)}\}$ and $X \overset{d}{=} X'$; the inequality is obtained by neglecting the sum in (15). Furthermore, similarly as in the proof of Lemma 2, $\tau_X^{(i)}$ is an exponential random variable with distribution

$$\mathbb{P}[\tau_X^{(i)} > t] = e^{-pq_i t}. \tag{18}$$

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

$$\begin{aligned} p &\leq \mathbb{P}[\{M_i(\tau_1^{(i)}) \geq k\} \cup \{M_i(\tau_1^{(i)}, mZ_1) \geq k+1\}] \\ &\leq \mathbb{P}[M_i(\tau_1, m-1) \geq k-1] + \mathbb{P}[Z_1 > k+1] + \mathbb{P}[M_i(\tau_1, m(k+1)) \geq k] \\ &= \mathbb{P}[M_i(\tau_1, m-1) \geq k-1] + \mathbb{P}[M_i(\tau_1) \geq 1]^k + \mathbb{P}[M_i(\tau_1, m(k+1)) \geq k], \end{aligned} \tag{19}$$

where in the last equality we used the geometric distribution of Z_1 from (14). Finally, (16), (17), (18), (19) and the fact that uniformly for all $1 \leq l_2 \leq l_1 \leq m(k+1)$ and any fixed $\epsilon > 0$,

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

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

In the following lemmas, we develop an analytic technique that allows us to decouple the dependency of random times T_i for i large and to compute necessary estimates used in the proof of Theorem 1 in Section 4. The following lemmas use the asymptotic estimate from Lemma 2 of Jelenković (1999) and the large deviation bound from Lemma 4 of Jelenković and Radovanović (2004a); for reasons of completeness, we restate these lemmas in the Appendix.

Lemma 4 *Let $\{T_i\}_{i\geq 1}$ be random variables defined in (6). Then, for $q_i \sim c/i^\alpha$ as $i \to \infty$, $\alpha > 1$ and*

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

we obtain

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

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

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < \vartheta_x] \leq x\right] \leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < \vartheta_x] \leq x\right]. \tag{20}$$

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 the 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

$$\Theta_{j+1} - \Theta_j \stackrel{d}{=} \tau_1^{(\overline{i_0})} + \sum_{l=1}^{m-1} \xi_l, \tag{21}$$

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_l having parameter 1.

Next, for all $i \geq i_0$, define

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

Similarly as in the proof of Lemma 2, the definition of U_i has identical form to the one for T_i in (6). The only difference is that $\{\Theta_j\} \subset \{\tau_n\}$ and, therefore,

$$T_i \leq U_i. \tag{22}$$

Then, using (21), we have that

$$\Theta_j \stackrel{d}{=} \tau_j^{(\overline{i_0})} + \nu_{j-1}^{(1)} + \cdots + \nu_{j-1}^{(m-1)}, \quad j \geq 1, \tag{23}$$

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

$$U_i \stackrel{d}{=} \tau_{X_i}^{(\overline{i_0})} + \sum_{l=1}^{m-1} \nu_{X_i-1}^{(l)} \leq \tau_{X_i}^{(\overline{i_0})} + \sum_{l=1}^{m-1} \nu_{X_i}^{(l)}, \tag{24}$$

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 $\epsilon > 0$ and i_0 large enough

$$(1+\epsilon)\frac{q_i}{\sum_{j\geq i_0} q_j} q_i^{k-1}\binom{m-1}{k-1} \geq p_i \geq (1-\epsilon)\frac{q_i}{\sum_{j\geq i_0} q_j} q_i^{k-1}\binom{m-1}{k-1}. \tag{25}$$

Next, inequalities (22) and (24) imply, for any $i \geq i_0$ and $\epsilon > 0$,

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

and, therefore,

$$\begin{aligned}
\mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < \vartheta_x] \leq x\right] &\leq \mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\epsilon)\vartheta_x] - \sum_{l=1}^{m-1}\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1\left[\nu_{X_i}^{(l)} > \frac{\epsilon\vartheta_x}{m}\right] \leq x\right] \\
&\leq \mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\epsilon)\vartheta_x] \leq (1+\epsilon/2)x\right] \\
&\quad + m\mathbb{P}\left[\sum_{i=i_0}^{\lfloor x\log x\rfloor} 1\left[\nu_{X_i}^{(1)} > \frac{\epsilon\vartheta_x}{m}\right] > \frac{x\epsilon}{2m}\right].
\end{aligned} \tag{26}$$

Now, using $q_i \sim c/i^\alpha$ as $i \to \infty$, Lemma 6 of the Appendix and setting $i_0 = \lfloor\sqrt{x}\rfloor$, we derive, as $x \to \infty$,

$$\mathbb{E}\left[\sum_{i=\lfloor\sqrt{x}\rfloor}^{\infty} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\epsilon)\vartheta_x]\right] \gtrsim x(1+\epsilon). \tag{27}$$

Then, using $q_i \sim c/i^\alpha$ as $i \to \infty$ and $1 - e^{-x} \leq x$, we arrive at

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

Thus, applying the preceding estimate and (27), we obtain

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

The previous expression, in conjunction with Lemma 7 of the Appendix, implies that the first term of (26) satisfies, as $x \to \infty$,

$$\mathbb{P}\left[\sum_{i=\lfloor\sqrt{x}\rfloor}^{\lfloor x\log x\rfloor} 1[\tau_{X_i}^{(\overline{i_0})} \leq (1-\epsilon)\vartheta_x] \leq (1+\epsilon/2)x\right] = o\left(\frac{1}{x^{\alpha-1}}\right). \tag{28}$$

Next, it is left to estimate the second term of (26). To this end, by using the monotonicity of q_i-s, assumption $q_i \sim c/i^\alpha$ as $i \to \infty$, inequality (25), and replacing ϑ_x, we obtain

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

Finally, applying Lemma 7 of the Appendix, we derive

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

which, in conjunction with (20), (26) and (28), completes the proof of this lemma. ◇

Note that in the following lemma, with a small abuse of notation, we assign a different value to ϑ_x from the one in Lemma 4.

Lemma 5 *Let $\{T_i\}_{i\geq 1}$ be random variables defined in (6). Then, for $q_i \sim c/i^\alpha$ as $i \to \infty$, $\alpha > 1$ and*

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

we obtain

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

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. Thus, since many of the arguments are repetitive, we postpone this proof until Section 7. ◇

4 Near optimality of the DPAC algorithm

Consider the class of 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 optimal. For direct arguments that justify this intuitively apparent statement see the first paragraph of Subsection 4.1 in Jelenković and Radovanović (2004b); this is also recently shown in Bahat and Makowski (2003) 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 the average-case competitive ratio between the stationary fault probabilities of the DPAC and the optimal static algorithm.

In the following theorem we show that, for the case of generalized Zipf's law request distributions with $\alpha > 1$ and large caches, the competitive ratio $\mathbb{P}[C > x]/\mathbb{P}[R > x]$ approaches 1 very rapidly as k grows. The proof of the theorem is based on preliminary results from Section 3 as well as the probabilistic and sample path arguments introduced in Jelenković and Radovanović (2004a,c) for the case of ordinary LRU and continuous time PAC algorithm, respectively. The starting point of our analysis is the representation formula in (5) 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 $x \to \infty$,*

$$\mathbb{P}[C > x] \sim K_k(\alpha)\mathbb{P}[R > x], \tag{29}$$

where

$$K_k(\alpha) \triangleq \left[\Gamma\left(1 - \frac{1}{\alpha k}\right)\right]^{\alpha-1} \Gamma\left(1 + \frac{1}{k} - \frac{1}{\alpha k}\right). \tag{30}$$

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

$$\lim_{\alpha \downarrow 1} K_k(\alpha) = 1, \quad \lim_{\alpha \uparrow \infty} K_k(\alpha) = K_k(\infty) \triangleq \frac{1}{k}\Gamma\left(\frac{1}{k}\right) e^{\gamma/k}, \tag{31}$$

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

$$\lim_{k \to \infty} K_k(\alpha) = 1. \tag{32}$$

Remark 2 (i) The same asymptotic result holds for the case of the continuous time PAC policy that was recently derived in Theorem 1 of Jelenković and Radovanović (2004c). (ii) After computing the second limit in (31) 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 ($K_3(\infty) \approx 1.08$), which shows the near optimality of the DPAC rule even for small values of k.

Proof: The proofs of monotonicity of $K_k(\alpha)$ and limits (31) and (32) can be found in Jelenković and Radovanović (2004c).

Next, we prove the upper bound for the asymptotic relationship in (29). 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., $S(t) \le S(\vartheta_x)$ a.s. for all $t \le \vartheta_x$, where ϑ_x is defined in Lemma 5. Then, after conditioning on T_i being larger or smaller than ϑ_x, the expression in (5) can be upper bounded as

$$\mathbb{P}[C > x] \le \mathbb{P}[S(\vartheta_x) > x] + \sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge \vartheta_x],$$

where in the previous expression we applied $\sum_{i=1}^{\infty} q_i = 1$ and $\mathbb{P}[S(t) > x] \le 1$. Then, applying Lemma 5, we obtain that the tail of the search cost C is upper bounded by

$$\mathbb{P}[C > x] \le o\left(\frac{1}{x^{\alpha-1}}\right) + \sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \ge \vartheta_x] \quad \text{as} \quad x \to \infty. \tag{33}$$

Next, due to the Poisson decomposition theorem, times of requests for document i are Poisson of rate q_i and mutually independent for different i. 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 easy to conclude that, for $i \ge i_0$,

$$\mathbb{P}[T_i \ge t] \le \mathbb{P}[T_{i_0} \ge t]. \tag{34}$$

Therefore, using Lemma 2, we obtain, for i_0 large enough,

$$\sum_{i=1}^{\infty} q_i \mathbb{P}[T_i \geq \vartheta_x] \leq \mathbb{P}[T_{i_0} \geq \vartheta_x] \sum_{i=1}^{i_0} q_i + \sum_{i=i_0}^{\infty} q_i e^{-q_i^k (1-\epsilon)^2 \binom{m-1}{k-1} \vartheta_x} + \sum_{i=i_0}^{\infty} m q_i e^{-h\epsilon q_i^{k-1} \vartheta_x}$$
$$\triangleq I_1(x) + I_2(x) + I_3(x). \tag{35}$$

After using bound (7) and replacing ϑ_x, it immediately follows that

$$I_1(x) \leq e^{-q_{i_0}^k (1-\epsilon)^2 \vartheta_x \binom{m-1}{k-1}} + m e^{-h\epsilon q_{i_0}^{k-1} \vartheta_x} = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \to \infty. \tag{36}$$

Now, by assumption of the theorem, for all i large enough ($i \geq i_0$, where i_0 is possibly larger than in (35))

$$(1-\epsilon)c/i^{\alpha} < q_i < (1+\epsilon)c/i^{\alpha}. \tag{37}$$

Furthermore, for i large enough ($i \geq i_0$) inequality $c/i^{\alpha} \leq (1+\epsilon)c/u^{\alpha}$ holds for any $u \in [i, i+1]$ and, therefore, using this bound, (37), the monotonicity of the exponential function and replacing ϑ_x from Lemma 5, yields

$$I_2(x) \leq (1+\epsilon) \sum_{i=i_0}^{\infty} \frac{c}{i^{\alpha}} e^{-\iota(\epsilon)[\Gamma(1-\frac{1}{\alpha k})]^{-\alpha k} \frac{x^{\alpha k}}{i^{\alpha k}}}$$
$$\leq (1+\epsilon)^2 \int_1^{\infty} \frac{c}{u^{\alpha}} e^{-\iota(\epsilon)[\Gamma(1-\frac{1}{\alpha k})]^{-\alpha k} \frac{x^{\alpha k}}{u^{\alpha k}}} du, \tag{38}$$

where $\iota(\epsilon) \triangleq (1+\epsilon)^{-(k+1)}(1-\epsilon)^{k+2}(1-2\epsilon)^{\alpha k}$. Next, applying the change of variable method for evaluating the integral with $z = \iota(\epsilon)[\Gamma(1-\frac{1}{\alpha k})]^{-\alpha k} x^{\alpha k} u^{-\alpha k}$, we obtain that the integral in (38) is equal to

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

which, in conjunction with (38), implies

$$\limsup_{x\to\infty} \frac{I_2(x)}{\mathbb{P}[R > x]} \leq K_k(\alpha)(\iota(\epsilon))^{\frac{1}{\alpha k}-\frac{1}{k}}(1+\epsilon)^2 \to K_k(\alpha) \text{ as } \epsilon \to 0, \tag{39}$$

where $K_k(\alpha)$ is defined in (30).

In order to estimate the asymptotics of $I_3(x)$, we use analogous steps as we applied in evaluating $I_2(x)$. Thus, using inequalities from (37), $c/i^{\alpha} \leq (1+\epsilon)c/u^{\alpha}$ for $u \in [i, i+1]$ and replacing ϑ_x, we obtain

$$I_3(x) \leq m(1+\epsilon) \sum_{i=i_0}^{\infty} \frac{c}{i^{\alpha}} e^{-h\epsilon \frac{x^{\alpha k}}{i^{\alpha(k-1)}}}$$
$$\leq m(1+\epsilon)^2 \int_1^{\infty} \frac{c}{u^{\alpha}} e^{-h\epsilon \frac{x^{\alpha k}}{u^{\alpha(k-1)}}} du. \tag{40}$$

Now, if $k = 1$, it is straightforward to compute the integral in the preceding expression and obtain $I_3(x) \leq m(1+\epsilon)^2(c/(\alpha-1))e^{-h\epsilon x^{\alpha}} = o(1/x^{\alpha-1})$ as $x \to \infty$. Otherwise, for $k \geq 2$, after using the change of variable method for solving the integral in (40) with $z = h\epsilon x^{\alpha k} u^{-\alpha(k-1)}$, we obtain, as $x \to \infty$,

$$I_3(x) \leq m(1+\epsilon)^3 \frac{c}{(h\epsilon)^{\frac{1}{k-1}(1-\frac{1}{\alpha})}} \frac{1}{\alpha(k-1)} \frac{1}{x^{\frac{k}{k-1}(\alpha-1)}} \Gamma\left(\frac{1}{k-1} - \frac{1}{\alpha(k-1)}\right) = o\left(\frac{1}{x^{\alpha-1}}\right). \tag{41}$$

Therefore, (41), (39), (36), (35) and (33), yield, as $x \to \infty$,

$$\mathbb{P}[C > x] \lesssim K_k(\alpha)\mathbb{P}[R > x]. \tag{42}$$

For the lower bound on $\mathbb{P}[C > x]$, starting from (5), we derive

$$\mathbb{P}[C > x] = \sum_{i=1}^{\infty} q_i \mathbb{P}[S_i(T_i) \geq x] \geq \sum_{i=1}^{\infty} q_i \mathbb{P}[S(\vartheta_x) \geq x - 1, T_i \geq \vartheta_x]$$
$$\geq \sum_{i=i_0}^{\infty} q_i \mathbb{P}[T_i \geq \vartheta_x] - \mathbb{P}[S(\vartheta_x) \leq x - 1],$$

where we choose ϑ_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 $\{T_i\}_{i \geq i_0}$, in conjunction with the analogous reasoning as in estimating $I_2(x)$, we complete the proof of this theorem. $\diamond$

In Theorems 2 and 3 of Jelenković and Radovanović (2004c), we characterize the asymptotic performance of the continuous time PAC algorithm for Zipf's law distributions with $0 < \alpha \leq 1$. Careful examination of the proofs of these results and the lemmas from Section 3 strongly suggests that exactly the same results hold for the DPAC algorithm as well. Since the rigorous proofs of these results would involve lengthy calculations and repetitive arguments without basically adding any new insights, we refrain from proving these results. Instead, for reasons of convenience, we just restate them here and illustrate them with simulation experiments in the following section.

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

$$(\log N)\mathbb{P}[C^{(N)} > \delta N] \sim F_k(\delta) \triangleq \frac{1}{k}\Gamma(0, \eta_\delta), \tag{43}$$

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* $0 < \delta < 1$,

$$\lim_{k\to\infty} F_k(\delta) = \log\left(\frac{1}{\delta}\right). \tag{44}$$

Remark 3 (i) Note that $(\log N)\mathbb{P}[R^{(N)} > \delta N] \to \log(1/\delta)$ as $N \to \infty$. Thus, for large caches, the limit in (44) shows that the DPAC policy approaches the optimal static algorithm as k increases; (ii) A related result from Lemma 4.7 of Fill (1996), in the context of the ordinary MTF searching ($k = 1$), shows the convergence in distribution of the ratio $\log C^{(N)}/\log N$ to a uniform random variable on a unit interval; this result corresponds to a sub-linear scaling $\mathbb{P}[C^{(N)} > N^u] \to 1 - u$ as $N \to \infty$, $0 < u < 1$.

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

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

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,* $1 - F_k(\delta)$, $\delta \in (0, 1)$, *is a proper distribution, with* $\lim_{\delta\downarrow 0} F_k(\delta) = 1$, $\lim_{\delta\uparrow 1} F_k(\delta) = 0$ *and*

$$\lim_{k\to\infty} F_k(\delta) = 1 - \delta^{1-\alpha}. \tag{46}$$

Remark 4 (i) Similarly, (46) and $\mathbb{P}[R^{(N)} > \delta N] \to 1 - \delta^{1-\alpha}$ as $N \to \infty$ demonstrate the asymptotic near optimality of the DPAC policy; (ii) For the ordinary MTF searching ($k = 1$), the convergence of $C^{(N)}/N$ in distribution as $N \to \infty$ and the Laplace transform of the limiting distribution function were obtained in Lemma 4.5 of Fill (1996). The result in its presented form, also for the ordinary LRU, was derived in Jelenković (2002).

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 formulas (29), (43) and (45), work very well for small caches as well.

5.1 Convergence to stationarity

In the presented experiments, we use a discrete time model without Poisson embedding, i.e., $\tau_n = n$; recall from the remark after Lemma 1 that the fault probability does not depend on the selection of the request arrival times $\{\tau_n\}_{n\geq 1}$. 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), it is not hard to show that the difference between the tails of these distributions can be upper bounded as

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

Now, using similar arguments as in (12) of Lemma 2, we obtain

$$\begin{aligned} \mathbb{P}[T_i > t] &\leq \mathbb{P}[\tau_{X_i}^{(i)} + (m-1)X_i > t] \\ &\leq \mathbb{P}\left[\tau_{X_i}^{(i)} > \frac{t}{2}\right] + \mathbb{P}\left[(m-1)X_i > \frac{t}{2}\right], \end{aligned} \tag{47}$$

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

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

Next, (47) and the well known fact that $\tau_{X_i}^{(i)}$ is geometric with parameter $q_i p_i$ yield

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

Thus, using the preceding bound, we obtain

$$e_n \leq \sum_{i=1}^{N} q_i \left[(1 - p_i q_i)^{\frac{n-m}{2}} + (1 - p_i)^{\frac{n-m}{2(m-1)}} \right]. \tag{48}$$

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 the stationary value. In other words, increasing 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 algorithm complexity and adaptability.

5.2 Experiments

In the presented experiments we choose 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, for the respective three experiments, 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 conducting the measurements, we allow for a warm-up time of the first n requests to let the system reach its stationarity; then, the actual measurement time is also set to be n 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 (48) 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 are apparent from Figure 1.

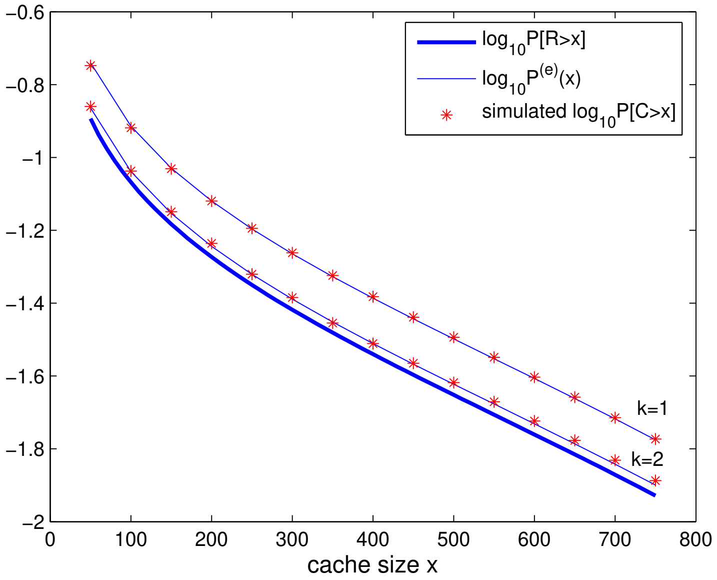

Fig. 1: Illustration for Experiment 1.

Experiment 2 Here, we set $\alpha = 1$ and measure the cache fault probabilities for $k = 1, 2, 3$. Again, both the warm-up and measurement time are chosen to be n requests long and equal to: $n = 2 \times 10^8$ for $k = 1, 2$ and $n = 10^{11}$ for $k = 3$. Since the normalization constant $(h_N)^{-1} = \log N + \gamma + o(1)$ as $N \to \infty$, where γ is the Euler's constant, the product $h_N \log N$ converges slowly to 1 and, therefore, instead of using the approximation $\mathbb{P}[C^{(N)} > x] \approx F_k(x/N)/\log N$, as suggested by Theorem 2, 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 are insignificant when compared to the smallest measured probabilities. Thus, the process is basically stationary. The accuracy of approximation $P^{(e)}(x)$ and the improvement in performance are apparent from Figure 2.

Experiment 3 Finally, the third example assumes $\alpha = 0.8$ and considers cases $k = 1, 2, 3$. Here, we select $n = 10^{10}$ to be the warm-up time as well as the actual measurement time. Similarly as in the case of $\alpha = 1$, due to the slow convergence of $h_N N^{1-\alpha}/(1-\alpha) \to 1$ 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. 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 are evident from Figure 3.

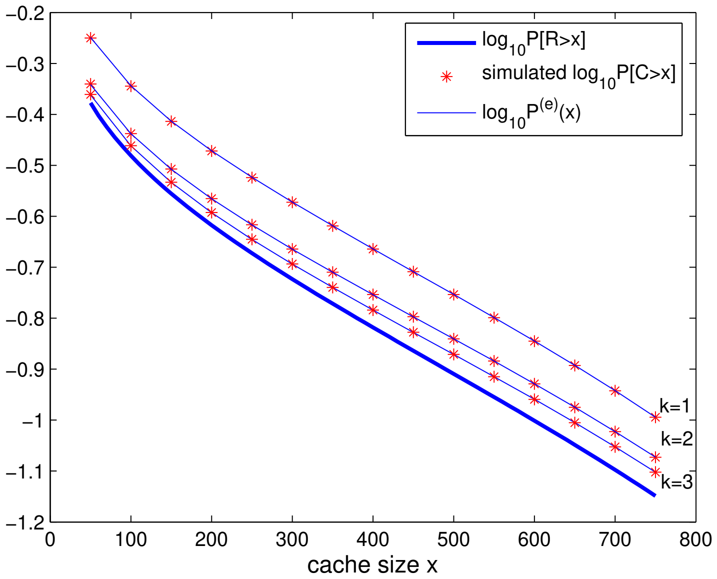

Fig. 2: Illustration for Experiment 2.

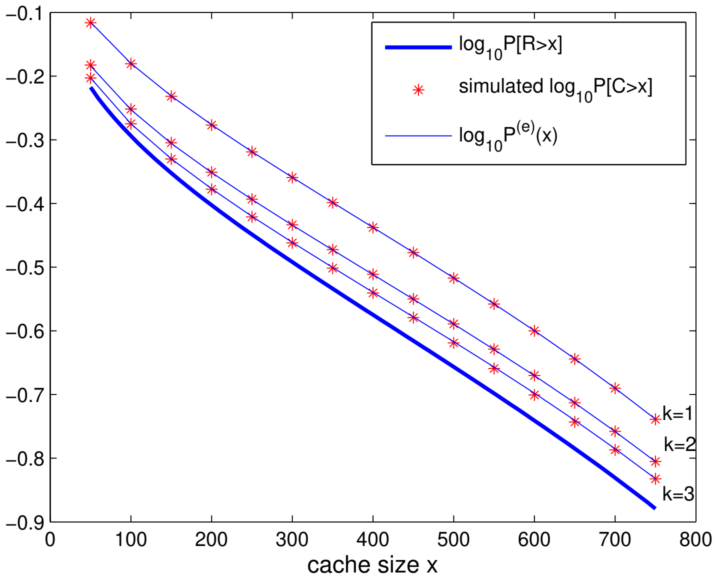

Fig. 3: Illustration for Experiment 3.

6 Concluding remarks

In this paper we introduce a discrete version of the recently proposed continuous time PAC replacement rule in Jelenković and Radovanović (2004c) that possesses all the 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, 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 fault probability of the DPAC algorithm, for large cache sizes, is very close to the optimal frequency algorithm even for small values of $k = 2, 3$; this implies negligible additional complexity relative to the classical LRU policy. In addition, the theoretical results are further validated using simulations that show significant improvement in performance of the DPAC algorithm in comparison with the LRU scheme, even for small cache sizes x and the number of documents N. Overall, given the considerable improvement in performance and low implementation complexity, we expect that the DPAC algorithm has a high potential for practical use.

Finally, in Web caching applications, documents typically have variable sizes. In this regard, by straightforward extension of the randomization procedure from Jelenković and Radovanović (2004b), one can easily design a randomized DPAC algorithm that is nearly optimal for variable document sizes as well.

7 Proof of Lemma 5

The case $k = 1$, corresponding 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 Jelenković and Radovanović (2004a). Hence, in the rest of the proof we assume $m \ge k \ge 2$.

Note that, for any $i_0 \ge 1$,

$$\mathbb{P}\left[\sum_{i=1}^{\infty} 1[T_i < \vartheta_x] \ge x\right] \le \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < \vartheta_x] \ge x - i_0\right]; \tag{49}$$

a specific i_0 will be selected later in the proof. Let $\{\tau_j^{(\overline{i_0})}\}_{j\ge1}$ be an ordered sequence of request times for documents $i \ge i_0$, i.e., $\{\tau_j^{(\overline{i_0})}\}_{j\ge1} = \cup_{i\ge i_0}\{\tau_n^{(i)}\}_{n\ge1}$. We use $n(j),\ j \ge 1$, to denote the index of point $\tau_j^{(\overline{i_0})}$ in the original sequence $\{\tau_n\}_{n\ge1}$, i.e., $\tau_j^{(\overline{i_0})} = \tau_{n(j)}$. Then, since process $\{\tau_n\}$ is Poisson and $\{R_n\}$ is an i.i.d. sequence independent of $\{\tau_n\}$, by Poisson decomposition theorem, process $\{\tau_{n(j)} \equiv \tau_j^{(\overline{i_0})}\}_{j\ge1}$ is also Poisson with rate $\sum_{i\ge i_0} q_i$.

Next, similarly as in Lemma 3, we group the points $\{\tau_j^{(\overline{i_0})}\}$ into cycles. The first cycle $\mathcal{C}_1$ 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 $M_{\overline{i_0}}(\tau_n) \triangleq \sum_{i\ge i_0} M_i(\tau_n)$. 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 $i \ge i_0$. Then, the first cycle, $\mathcal{C}_1$, will be the interval $[\tau_{n(1)}, \tau_{n(1)+mZ_1}]$. Next, starting from the first point of process $\{\tau_{n(j)}\}_{j\ge1}$ 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 $\mathcal{C}_2 = [\tau_{n(l)}, \tau_{n(l)+mZ_2}]$. We continue this procedure indefinitely.

Then, denote the points of time that represent the beginnings of the previously defined cycles $\mathcal{C}_j$, $j \ge 1$, by $\{\Theta_j\}_{j\ge1}$. Clearly, from the independence assumptions on $\{\tau_n\}$ and $\{R_n\}$, $\{\Theta_j\}$ is a renewal process with renewal intervals, for $j \ge 1$, 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 lengths of the cycles), the beginning of each cycle can be lower bounded with

$$\Theta_j \stackrel{d}{\ge} \tau_j^{(\overline{i_0})}. \tag{50}$$

Next, on each cycle $\mathcal{C}_j$, $j \ge 1$, define an event that at least two distinct items are moved to the first position of the list during that cycle

$$\mathcal{A}_j^{(0)} \triangleq \{\omega : \exists i_1, i_2 \ge i_0, i_1 \ne i_2, \exists \tau_{n_1}, \tau_{n_2} \in \mathcal{C}_j, R(\tau_{n_l}) = i_l, M_{i_l}(\tau_{n_l}) \ge k \text{ for } l = 1, 2\}.$$

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 during cycle $\mathcal{C}_j$

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

where $\mathcal{A}^c$ denotes a complement of event $\mathcal{A}$; for each fixed j, these events are disjoint. Then, due to the independence properties of our reference model and the fact that the last m points in each cycle do not contain requests for documents $i \geq i_0$, they are independent on different cycles and for fixed i equally distributed; let $p_i \triangleq \mathbb{P}[\mathcal{A}_1^{(i)}]$, $i \geq i_0$ or $i = 0$.

Now, using the bound in (50) and the Poisson decomposition theorem, it is easy to see that, for each fixed i, the beginning of the first cycle where event $\mathcal{A}_j^{(i)}$ 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 (6) are lower bounded by the beginning of the first cycle where event $\mathcal{A}_j^{(0)} \cup \mathcal{A}_j^{(i)}$ occurs, which is further lower bounded by

$$T_i \geq L_i \wedge L_0, \tag{51}$$

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 (19) of Lemma 3, we obtain that for any $\epsilon > 0$, we can chose i_0 large enough, such that for all $i \geq i_0$,

$$p_i \leq (1+\epsilon) \frac{q_i}{\sum_{j=i_0}^{\infty} q_j} \binom{m-1}{k-1} q_i^{k-1}. \tag{52}$$

The probability p_0 can be bounded as

$$\begin{aligned} p_0 &\leq \mathbb{P}[Z_j > l] + \mathbb{P}[\mathcal{A}_j^{(0)}, Z_j \leq l] \\ &\leq H \left(\sum_{j=i_0}^{\infty} q_j \right)^l + \mathbb{P}[\mathcal{A}_j^{(0)}, Z_j \leq l] \\ &\leq H \left(\sum_{j=i_0}^{\infty} q_j \right)^l + \sum_{j_1, j_2 \geq i_0, j_1 \neq j_2} \mathbb{P}[\mathcal{A}_j^{(0)}(j_1, j_2), Z_j \leq l], \end{aligned} \tag{53}$$

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

$$\begin{aligned} \mathbb{P}\left[\mathcal{A}_1^{(0)}(j_1, j_2), Z_1 \leq l\right] &= \mathbb{P}\left[R(\Theta_1) = j_1, \mathcal{A}_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\ &\quad + \mathbb{P}\left[R(\Theta_1) = j_2, \mathcal{A}_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\ &\quad + \mathbb{P}\left[R(\Theta_1) \neq j_1, j_2, \mathcal{A}_1^{(0)}(j_1, j_2), Z_1 \leq l\right] \\ &\triangleq p_{01}(j_1, j_2) + p_{02}(j_1, j_2) + p_{03}(j_1, j_2). \end{aligned} \tag{54}$$

Now, we upper bound the first term of (54),

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

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

$$p_{01}(j_1, j_2) \leq \frac{H q_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}. \tag{55}$$

Similarly, we derive

$$p_{02}(j_1, j_2) \leq \frac{H q_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j} \tag{56}$$

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

$$p_{03}(j_1, j_2) \leq H q_{j_1}^k q_{j_2}^k \leq \frac{H q_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}. \tag{57}$$

Therefore, (55), (56), (57) and (54) imply that for any $j_1, j_2 \geq i_0, j_1 \neq j_2$,

$$\mathbb{P}[\mathcal{A}_j^{(0)}(j_1, j_2), Z_j \leq l] \leq H \frac{q_{j_1}^k q_{j_2}^k}{\sum_{j=i_0}^{\infty} q_j}, \tag{58}$$

where constant H is independent of j_1 and j_2. Now, by replacing the preceding bound in (53), we derive that for all i_0 large enough

$$p_0 \leq H \left(\sum_{j=i_0}^{\infty} q_j \right)^l + \frac{H}{(i_0 - 1)^{2\alpha k - 2} \sum_{j=i_0}^{\infty} q_j}. \tag{59}$$

After setting the necessary ground for our analysis, we upper bound the left hand side of (49) as

$$\begin{aligned}
\mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[T_i < \vartheta_x] \geq x - i_0 \right] &\leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i \wedge L_0 < \vartheta_x] \geq x - i_0 \right] \\
&\leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i \wedge L_0 < \vartheta_x, L_0 > \vartheta_x] \geq x - i_0 \right] + \mathbb{P}[L_0 \leq \vartheta_x] \\
&\leq \mathbb{P}\left[\sum_{i=i_0}^{\infty} 1[L_i < \vartheta_x] \geq x - i_0 \right] + \mathbb{P}[L_0 \leq \vartheta_x].
\end{aligned} \tag{60}$$

Now, from (52), $\mathbb{P}[L_i < \vartheta_x] \leq 1 - e^{-(1+\epsilon)\binom{m-1}{k-1} q_i^k \vartheta_x}$ for $i \geq i_0$ and i_0 large enough. Furthermore, assigning $i_0 = \lceil \epsilon x \rceil$ and applying Lemma 6 of the Appendix, we derive as $x \to \infty$,

$$\mathbb{E}\left[\sum_{i \geq \epsilon x} 1[L_i < \vartheta_x] \right] \lesssim \Gamma\left(1 - \frac{1}{\alpha k}\right) c^{\frac{1}{\alpha}} \left(\binom{m-1}{k-1} \right)^{\frac{1}{\alpha k}} (1+\epsilon)^{\frac{k+1}{\alpha k}} \vartheta_x^{\frac{1}{\alpha k}}. \tag{61}$$

Then, if we replace ϑ_x and use (61), it follows, as $x \to \infty$, that $\mathbb{E}\left[\sum_{i \geq \epsilon x} 1[L_i < \vartheta_x]\right] \lesssim (1-2\epsilon)x \leq (1-\epsilon)(x - \epsilon x)$. Thus, since L_i, $i \geq i_0$, are mutually independent, using large deviation result from Lemma 7, we show that the first term in (60) is bounded, for some $\theta > 0$, by

$$\mathbb{P}\left[\sum_{i \geq \epsilon x} 1[L_i < \vartheta_x] \geq x - \epsilon x \right] \leq 2e^{-\theta x} = o\left(\frac{1}{x^{\alpha - 1}} \right) \quad \text{as} \quad x \to \infty. \tag{62}$$

Next, we estimate the second term of (60). Using (59) with $i_0 = \lceil \epsilon x \rceil$ and choosing $l = \lceil (2\alpha k - 2)/(\alpha - 1) \rceil - 1$, such that $(\alpha - 1)(l + 1) \geq 2\alpha k - 2$, we derive

$$\begin{aligned}\mathbb{P}[L_0 \leq \vartheta_x] &\leq 1 - e^{-\vartheta_x p_0 \sum_{i \geq \epsilon x} q_i} \\ &\leq H x^{\alpha k} \left[\frac{1}{x^{(\alpha-1)(l+1)}} + \frac{1}{x^{2\alpha k - 2}} \right] \\ &\leq \frac{H x^{\alpha k}}{x^{2\alpha k - 2}} \leq \frac{H}{x^{\alpha k - 2}} = o\left(\frac{1}{x^{\alpha - 1}} \right) \quad \text{as } x \to \infty, \end{aligned} \tag{63}$$

since $k \geq 2$ and $\alpha > 1$. Finally, replacing (63) and (62) in (60), implies the statement of the lemma. $\diamond$

Appendix

The following lemmas correspond to Lemma 2 of Jelenković (1999) and Lemma 4 of Jelenković and Radovanović (2004a), respectively.

Lemma 6 *Let $S(t) = \sum_{i=1}^{\infty} B_i(t)$ and assume $q_i \sim c/i^{\alpha}$ as $i \to \infty$, with $\alpha > 1$ and $c > 0$. Then, as $t \to \infty$,*

$$m(t) \triangleq \mathbb{E}S(t) \sim \Gamma\left(1 - \frac{1}{\alpha}\right) c^{\frac{1}{\alpha}} t^{\frac{1}{\alpha}}.$$

Lemma 7 *Let $\{B_i, 1 \leq i \leq N\}$, $N \leq \infty$, be a sequence of independent Bernoulli random variables, $S = \sum_{i=1}^{N} B_i$ and $m = \mathbb{E}[S]$. Then for any $\epsilon > 0$, there exists $\theta_\epsilon > 0$, such that*

$$\mathbb{P}[|S - m| > m\epsilon] \leq 2e^{-\theta_\epsilon m}.$$

References

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.

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

J. A. Fill. Limits and rate of convergence for the distribution of search cost under the move-to-front rule. *Theoretical Computer Science*, 164:185–206, 1996.

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.

P. Flajolet, D. Gardy, and L. Thimonier. Birthday paradox, coupon collector, caching algorithms and self-organizing search. *Discrete Applied Mathematics*, 39:207–229, 1992.

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

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.

P. R. Jelenković. Least-Recently-Used Caching with Zipf's Law Requests. In *The Sixth INFORMS Telecommunications Conference*, Boca Raton, Florida, 2002.

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

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

P. R. Jelenković and A. Radovanović. The Persistent-Access-Caching Algorithm. Technical Report EE2004-03-05, Department of Electrical Engineering, Columbia University, New York, March 2004c. Available at `http://www.comet.columbia.edu/~predrag/mypub/PAC04.ps`.

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

E. J. O'Neil, P. E. O'Neil, and G. Weikum. An optimality proof of the LRU-K page replacement algorithm. *Journal of the ACM*, 46:92–112, 1999.