

CRITICAL SIZING OF LRU CACHES WITH DEPENDENT REQUESTS

PREDRAG R. JELENKOVIĆ,* *Columbia University*

ANA RADOVANOVIĆ** *** AND

MARK S. SQUILLANTE,** **** *IBM T. J. Watson Research Center*

Abstract

It was recently proved by Jelenković and Radovanović (2004) that the least-recently-used (LRU) caching policy, in the presence of semi-Markov-modulated requests that have a generalized Zipf's law popularity distribution, is asymptotically insensitive to the correlation in the request process. However, since the previous result is asymptotic, it remains unclear how small the cache size can become while still retaining the preceding insensitivity property. In this paper, assuming that requests are generated by a nearly completely decomposable Markov-modulated process, we characterize the critical cache size below which the dependency of requests dominates the cache performance. This critical cache size is small relative to the dynamics of the modulating process, and in fact is sublinear with respect to the sojourn times of the modulated chain that determines the dependence structure.

Keywords: Network caching; least-recently-used caching; move-to-front searching; nearly completely decomposable Markov process; Markov-modulated process; Zipf's law

2000 Mathematics Subject Classification: Primary 60J25; 60F10

Secondary 68P10; 60G50

1. Introduction

The basic idea of caching is to maintain high-speed access to a subset of x popular items out of a larger collection of N documents that are otherwise accessed at a slower rate. In the context of Internet applications and services, such as Web access and content delivery, caching has been widely recognized as an effective way to reduce the latency for downloading Internet documents. This is achieved by keeping the popular documents in high-speed caches that are located close to the users requesting these documents. Naturally, the problem of selecting and possibly dynamically updating the content of a cache is central to the efficient operation of any caching system. The broad popularity of the least-recently-used (LRU) policy stems from its many desirable characteristics, including a high hit ratio, a low complexity, and the flexibility to adapt dynamically to possible changes in the request patterns.

Received 13 February 2006; revision received 22 August 2006.

* Postal address: Department of Electrical Engineering, Columbia University, New York, NY 10027, USA.

Email address: predrag@ee.columbia.edu

Supported by the NSF, grant no. 0092113.

** Postal address: Department of Mathematical Sciences, IBM T. J. Watson Research Center, PO Box 218, Yorktown Heights, NY 10598, USA.

*** Email address: aradovan@us.ibm.com

**** Email address: mss@watson.ibm.com

Owing to its importance in practice, LRU caching has received significant attention in the research literature, in the context of both combinatorial (worst-case) [3], [4], [15], [16] and probabilistic (average-case) analyses; the latter is the focus of this paper. In particular, we consider the LRU algorithm in the presence of strong statistical correlation that often characterizes the access patterns for Internet documents; see, e.g. [1], [5], [12], and the references therein. However, most of the existing work on the average-case analysis of LRU caches is either performed under the assumption of independent and identically distributed (i.i.d.) requests or is computationally intractable. To alleviate this problem, in the recent work [13] two of the present authors developed a novel, analytically explicit asymptotic method for analyzing LRU caches in the presence of semi-Markov-modulated requests. This way of modeling dependency in the request process provides the desired flexibility for capturing possibly strong statistical correlation, including the widely reported long-range dependence of the access patterns for Web documents. In fact, Markov modulation techniques are widely used to model dependencies in the arrival processes in the context of queueing and insurance risk theories; see, e.g. [2], [11], and the references therein. The main results of [12] and [13] imply that, asymptotically, for large cache sizes, the cache fault probability in the presence of semi-Markov-modulated requests behaves the same way as in the corresponding LRU system with i.i.d. requests [10]. This surprising insensitivity result was further validated experimentally in [12], where excellent agreement with the asymptotic results was found, even in the cases of actual trace-driven simulations and relatively small cache sizes, further supporting our way of modeling dependency in the request process.

Since the results of [13] are asymptotic, they do not provide information on how small the cache sizes can become while still retaining the discovered insensitivity property. The present work attempts to answer this question by studying the cache performance through a joint scaling of the dependence structure of the requests and the cache size. In this paper, the request sequence is modeled as a nearly completely decomposable Markov-modulated process with the modulating Markov process having transition rates linearly proportional to a scaling parameter δ . The jumps in this modulating process occur on a time scale of the order $1/\delta$, which implies that the dependency in the request process increases as $\delta \downarrow 0$. We scale the cache size as an increasing function of $1/\delta$ and identify a critical cache size below which the dependency dominates the cache performance. Our main results show that this critical cache size is sublinear in comparison with the time scale of transitions in the modulating process, i.e. $1/\delta$. Thus, informally, our results show that the discovered insensitivity property is indeed robust.

The remainder of the paper is organized as follows. In Section 2 we define the model used in our study, while in Section 3 we present a summary of the results that are used in our main theorems. The main results are provided in Theorems 1 and 2 of Section 4, together with a discussion of their implications. In Section 5 we conclude the paper.

2. Model description

A LRU cache of size x can be described as follows. Consider a universe of N documents (items), of which x can be placed in an efficiently accessible location called the cache. Each time a request for a document is made, the cache is searched first. If the document is not found in the cache (cache fault), additional delay is incurred to access the item from the outside universe and the item is added to the cache by replacing the least-recently-accessed document in the cache. The performance measure of interest for this algorithm is the LRU cache fault probability, i.e. the probability that the requested document is not found in the cache.

Analyzing the LRU policy is equivalent to investigating the move-to-front (MTF) searching algorithm. In order to justify this claim, we assume that the x documents in the cache, under the LRU rule, are arranged in increasing order of their last access times. Every time there is a request for a document that is not in the cache, the document is moved to the first position of the cache and the last document in the cache is moved to the outside universe. Clearly, the fault probability stays the same if the remaining $N - x$ documents in the outside universe are arranged in any specific order. In particular, they can be arranged in increasing order of their last access times. The obtained searching scheme performed on the ordered list of all documents is called the MTF algorithm. Furthermore, it is clear from the previous arguments that the LRU fault probability is equal to the tail of the MTF search cost, i.e. the position of the requested document evaluated at the cache size. Additional arguments that justify the connection between the MTF search cost distribution and the LRU cache fault probability can be found in [7], [9], and [10]. We therefore proceed with a description of the MTF algorithm.

More formally, consider a finite set of documents $L = \{1, \dots, N\}$ and a sequence of document requests that arrive at time points $\{\tau_n, -\infty < n < \infty\}$, which represent a Poisson process of unit rate. At each point τ_n , we use R_n to denote the document that has been requested, i.e. the event $\{R_n = i\}$ represents a request for document i at time τ_n . The sequence $\{R_n\}$ is assumed to be independent of the Poisson arrival points $\{\tau_n\}$. The dynamics of the MTF algorithm are defined as follows. Suppose that the system starts at the arrival instant, τ_0 , of the 0th request with an initial permutation Π_0 of the MTF list. Then, every time $\tau_n, n \geq 0$, that a document is requested, its position in the list is determined and this value is taken to represent the searching cost, $C_n^{(N)}$, at time τ_n . The list is then updated by moving the requested document to the first position of the list and shifting those documents that were in front of the requested item one position down. Note that, according to the discussion in the preceding paragraph, $P[C_n^{(N)} > x]$ represents the fault probability of a cache of size x at time τ_n .

Next, we characterize the dependence structure of the request process. Let $N_\delta = \{T_n, -\infty < n < \infty\}, T_0 \leq 0 < T_1$, be a Poisson point process of rate $\delta > 0$. Furthermore, let $\{\mathcal{J}_n, -\infty < n < \infty\}$ be a finite-state, irreducible, aperiodic Markov chain, independent of N_δ , taking values in $\{1, \dots, M\}$, where M is a finite, positive integer. This process is assumed to be stationary with marginal distribution $\pi_k = P[\mathcal{J}_n = k]$. By embedding this Markov chain into the Poisson process N_δ , we construct a piecewise-constant, right-continuous *modulating process* J , where J is defined by $J_t = \mathcal{J}_n$ for $T_n \leq t < T_{n+1}$. Note that the transition rates in J are linearly proportional to δ and, therefore, this is a nearly completely decomposable process for small δ .

For each $k, 1 \leq k \leq M$, let $q_i^{(k)}$ be a probability mass function used to denote the probability of requesting document $i, 1 \leq i \leq N$, when the underlying process J is in state k . The dynamics of R_n are then uniquely determined by the modulating process J according to the equation

$$P[R_l = i_l, 1 \leq l \leq n \mid J_t, t \leq \tau_n] = \prod_{l=1}^n q_{i_l}^{(J_{\tau_l})},$$

where $n \geq 1$; that is, the sequence of requests $\{R_n\}$ is conditionally independent given the modulating process J . We use $q_i = P[R = i] = \sum_{k=1}^M \pi_k q_i^{(k)}$ to express the marginal request distribution and assume that $q_i > 0$ for all $i, 1 \leq i \leq N$.

3. Preliminary results

The model described in the previous section, for a fixed δ , is a special case of the more general one introduced in [13], and some of the results of [13] are therefore used in this paper. In

particular, Lemma 1 of [13] shows that the search cost $C_n^{(N)}$, $N < \infty$, converges in distribution to the stationary value $C^{(N)}$ when the request process $\{R_n\}$ is stationary and ergodic. In the following subsection we outline this convergence and provide a characterization of the tail of the limiting search cost distribution when the number of documents tends to ∞ . Subsection 3.2 contains results on MTF searching with i.i.d. requests that were stated and proved in [10] and [13] and will be used in proving our main theorems.

3.1. Representation results

Section 3.1 of [13] contains a general characterization of the stationary distribution of $C^{(N)}$, $N < \infty$. Assume that the probability mass functions $q_i^{(k)}$ are defined for every $i \geq 1$ and k , $1 \leq k \leq M$. Then let σ_t be the σ -algebra $\sigma(J_u, -t \leq u \leq 0)$ containing the history of the process J_t in the interval $[-t, 0]$. We denote the conditional probability $P[\cdot | \sigma_t]$ by $P_{\sigma_t}[\cdot]$. Furthermore, let $N_j(u; J)$ be the number of requests for document j in $[-u, 0)$, $0 < u \leq t$, and define an indicator function $B_j(t; J) = \mathbf{1}_{\{N_j(t; J) > 0\}}$, $j \geq 1$, which equals 1 if item j was requested in $[-t, 0)$. Then $S_i(t; J)$, the number of distinct documents different from i that were requested in $[-t, 0)$, can be expressed as

$$S_i(t; J) := \sum_{j \neq i} B_j(t; J),$$

and

$$P_{\sigma_t}[B_j(t; J) = 1] = 1 - e^{-\hat{q}_j t}, \tag{1}$$

where the $\hat{q}_j \equiv \hat{q}_j(t)$, $j \geq 1$, denote the empirical probabilities of requesting document j in the interval $[-t, 0)$, i.e.

$$\hat{q}_j := \sum_{k=1}^M q_j^{(k)} \hat{\pi}_k \quad \text{with} \quad \hat{\pi}_k := \frac{1}{t} \int_{-t}^0 \mathbf{1}_{\{J_u=k\}} du.$$

Now we can construct a sequence of finite lists and show that their search costs $C^{(N)}$ converge in distribution, as $N \rightarrow \infty$, to

$$P[C > x] = E \int_0^\infty \hat{f}(t) P_{\sigma_t}[S_i(t; J) > x - 1] dt, \tag{2}$$

where $\hat{f}(t)$ is defined by

$$\hat{f}(t) := \sum_{i=1}^\infty q_i^{(J_0)} q_i^{(J_{-t})} e^{-\hat{q}_i t}. \tag{3}$$

The reader is referred to the proof of Proposition 1 of [13] for details, where the above results were established under more general model assumptions on the request process $\{R_n\}$ than those introduced in Section 2. The preceding representation of the distribution of C is the starting point of our analysis.

Remark 1. (i) Throughout the paper we will exploit the facts that the variables $S_j(t; J)$ and $B_j(t; J)$, $j \geq 1$, are monotonically increasing in t and that the variables $B_j(t; J)$, $j \geq 1$, are conditionally independent given σ_t . This conditional independence arises from the Poisson arrival structure, as is apparent from the derivation in [13]. In general, when the request times are not Poisson, e.g. integer time arrivals, these variables may not be conditionally independent. For i.i.d. requests, the Poisson embedding technique was first introduced in [8].

(ii) It is clear that the derivation of the above results does not rely on the fact that the requests arrive at a constant rate [13]. Thus, our results can be generalized to the case where the arrival rate depends on the state of the modulating process J ; i.e. the rate can be set to λ_k when $J_t = k$. We do not consider this extension, since it further complicates the notation without providing any significant new insight.

3.2. Results for i.i.d. requests

We next provide several lemmas that consider the LRU caching scheme under independent requests and which will be used in proving our main theorems. The MTF model with i.i.d. requests is equivalent to our general problem formulation when the modulating process is assumed to be a constant, i.e. $J_t \equiv \text{const}$. In this case, the Bernoulli variables $B_j(t)$, $j \geq 1$, indicating that a document j was requested in $[-t, 0)$ are independent and have success probabilities $\mathbb{P}[B_j(t) = 1] = 1 - e^{-q_j t}$. Using the notation $S_i(t) := \sum_{j \neq i} B_j(t)$, it is thus easy to see that the distribution of the limiting stationary search cost C from (2) reduces to

$$\mathbb{P}[C > x] = \int_0^\infty \sum_{i=1}^\infty q_i^2 e^{-q_i t} \mathbb{P}[S_i(t) > x - 1] dt. \tag{4}$$

Throughout the paper we shall use some standard notation. For any two real functions $a(t)$ and $b(t)$ and a fixed $t_0 \in \mathbb{R} \cup \{\infty\}$, we will write $a(t) \sim b(t)$ as $t \rightarrow t_0$ to mean that $\lim_{t \rightarrow t_0} a(t)/b(t) = 1$. Similarly, we will write $a(t) \gtrsim b(t)$ as $t \rightarrow t_0$ to mean that $\liminf_{t \rightarrow t_0} a(t)/b(t) \geq 1$; $a(t) \lesssim b(t)$ has a complementary definition. The following two results, originally proved in Lemmas 1 and 2 of [10], are restated here for convenience.

Lemma 1. *Assume that $q_i \sim c/i^\alpha$ as $i \rightarrow \infty$, with $\alpha > 1$ and $c > 0$. Then, as $t \rightarrow \infty$,*

$$\sum_{i=1}^\infty q_i^2 e^{-q_i t} \sim \frac{c^{1/\alpha}}{\alpha} \Gamma\left(2 - \frac{1}{\alpha}\right) t^{-2+1/\alpha},$$

where Γ is the gamma function.

Lemma 2. *Let $S(t) = \sum_{i=1}^\infty B_i(t)$ and assume that $q_i \sim c/i^\alpha$ as $i \rightarrow \infty$, with $\alpha > 1$ and $c > 0$. Then, as $t \rightarrow \infty$,*

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

Throughout the paper we shall use H to denote a sufficiently large, positive constant, whereas h will be used to denote a sufficiently small, positive constant (where the notion of sufficiency will depend on the context). The values of H and h are generally different in different places. For example, $H/2 = H$, $H^2 = H$, $H + 1 = H$, and similarly for h . The next two lemmas, which are repeatedly used in establishing our main results, were originally proved in [13].

Lemma 3. *Let $\{B_i, i \geq 1\}$ be a sequence of independent Bernoulli random variables, let $S = \sum_{i=1}^\infty B_i$, and let $m = \mathbb{E} S$. Then, for any $\varepsilon > 0$, there exists a $\theta_\varepsilon > 0$ such that*

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

Lemma 4. *If $0 \leq q_i \leq H/i^\alpha$ for some fixed $\alpha > 1$, then, for any $x \geq 1$,*

$$P[C > x] \leq \frac{H}{x^{\alpha-1}}.$$

The result established in the following lemma is repeatedly used in the proof of Theorem 1. By $\lfloor x \rfloor$ we denote the largest integer less than or equal to x .

Lemma 5. *Let $c/i^\alpha \leq q_i \leq c^{-1}/i^\alpha$, $\alpha > 1$, for some positive constant c . Then, for any $x > 0$,*

$$P[C > x] \geq \frac{h}{x^{\alpha-1}}.$$

Proof. Note that, for any $\varepsilon > 0$ and a sufficiently large x , by using (4) the tail of the search cost C can be bounded from below as follows:

$$P[C > x] \geq P[S(Hx^\alpha) > x - 1] \int_{Hx^\alpha}^\infty \sum_{i=1}^\infty q_i^2 e^{-q_i t} dt \geq (1 - \varepsilon) \sum_{i=1}^\infty q_i e^{-Hx^\alpha q_i}. \quad (5)$$

The first inequality follows from the monotonicity of $S(t)$ and the second inequality is obtained by applying Lemmas 2 and 3 and integrating. Next, from the assumptions of the lemma, we have

$$\sum_{i=1}^\infty q_i e^{-Hq_i x^\alpha} \geq \sum_{\lfloor x \rfloor + 1}^\infty \frac{c}{i^\alpha} e^{-Hx^\alpha/i^\alpha} \geq ce^{-H} \int_{x+1}^\infty \frac{1}{u^\alpha} du \geq \frac{h}{x^{\alpha-1}},$$

which, in conjunction with (5), proves the result.

4. Main results

In this section we state and prove our main results. We show that the cache fault probability exhibits different performance characteristics depending on the scaling between the cache size, x , and the parameter δ .

In preparation for these proofs, we denote the epochs of reversed jump points by $\mathcal{T}_n := -T_{-n}$, $n \geq 0$; this notation is convenient since C depends on J_t for $t \leq 0$. Furthermore, we define

$$S^{(k)}(t) := B_i^{(k)}(t) + S_i^{(k)}(t) = B_i^{(k)}(t) + \sum_{j \neq i} B_j^{(k)}(t), \quad 1 \leq k \leq M,$$

where $B_i^{(k)}(t)$, $i \geq 1$, are Bernoulli random variables with $P[B_i^{(k)}(t) = 1] = 1 - e^{-q_i^{(k)} t}$. In addition, let $C^{(k)}$ correspond to the stationary search cost with i.i.d. requests when $J_t \equiv k$.

4.1. Asymptotic decomposability

The following theorem establishes the critical cache size scaling as a function of the parameter δ , below which the dependency in the request process dominates cache performance, i.e. the insensitivity result does not hold.

Theorem 1. *Let $q_i \leq c_1/i^\alpha$, $c_1 > 0$, $\alpha > 1$, and suppose that there exists a k , $1 \leq k \leq M$, such that $q_i^{(k)} \geq c_2/i^\alpha$, $c_2 > 0$. If x_δ satisfies $x_\delta \delta^{1/\alpha} \rightarrow 0$ as $\delta \rightarrow 0$, then*

$$P[C > x_\delta] \sim \sum_{k=1}^M \pi_k P[C^{(k)} > x_\delta] \quad \text{as } x_\delta \rightarrow \infty.$$

Proof. To simplify notation we write $x \equiv x_\delta$. We first prove the lower bound. Since $S(t; J) = S^{(k)}(t)$ almost surely on $\{J_0 = k\}$ for all t , $-\mathcal{T}_0 \leq t \leq 0$, the representation formula given in (2) implies that

$$\begin{aligned}
 P[C > x] &= E \int_0^\infty \hat{f}(t) P_{\sigma_t}[S_i(t; J) > x - 1] dt \\
 &\geq E \int_0^{\mathcal{T}_0} \sum_{i=1}^\infty (q_i^{(J_0)})^2 e^{-q_i^{(J_0)} t} P[S_i^{(J_0)}(t) > x - 1 \mid J_0] dt \\
 &\geq \sum_{k=1}^M P[J_0 = k, \mathcal{T}_0 > Hx^\alpha] \int_0^\infty \sum_{i=1}^\infty (q_i^{(k)})^2 e^{-q_i^{(k)} t} P[S_i^{(k)}(t) > x - 1] dt \\
 &\quad - \sum_{k=1}^M \pi_k \int_{Hx^\alpha}^\infty \sum_{i=1}^\infty (q_i^{(k)})^2 e^{-q_i^{(k)} t} dt. \tag{6}
 \end{aligned}$$

Now, since

$$q_i^{(k)} \leq \bar{q}_i := \frac{q_i}{\min_k \pi_k}, \quad 1 \leq k \leq M,$$

$q_i \leq c_1/i^\alpha$, and $xe^{-x} \leq e^{-1}$ (for $x \geq 0$), the second summand in (6) can be bounded as follows:

$$\begin{aligned}
 \sum_{k=1}^M \pi_k \int_{Hx^\alpha}^\infty \sum_{i=1}^\infty (q_i^{(k)})^2 e^{-q_i^{(k)} t} dt &\leq \sum_{k=1}^M \pi_k \frac{1}{Hx^\alpha} \sum_{i=1}^{\lfloor H^{1/\alpha} x \rfloor} q_i^{(k)} Hx^\alpha e^{-q_i^{(k)} Hx^\alpha} \\
 &\quad + \frac{1}{(\min_k \pi_k)} \int_{H^{1/\alpha} x}^\infty \frac{c_1}{y^\alpha} dy \\
 &\leq \frac{1}{H^{1-1/\alpha}} \frac{1}{x^{\alpha-1}} \left(e^{-1} + \frac{c_1}{\min_k \pi_k (\alpha - 1)} \right). \tag{7}
 \end{aligned}$$

Then, by the assumption of the theorem, $P[J_0 = k, \mathcal{T}_0 > Hx^\alpha] = \pi_k e^{-H\delta x^\alpha} \rightarrow \pi_k$ as $\delta \rightarrow 0$ ($x \rightarrow \infty$) and, therefore, from (6) and (7) we obtain

$$P[C > x] \gtrsim \sum_{k=1}^M \pi_k P[C^{(k)} > x] - \frac{1}{H^{1-1/\alpha}} \left(e^{-1} + \frac{c_1}{\min_k \pi_k (\alpha - 1)} \right) \frac{1}{x^{\alpha-1}} \quad \text{as } x \rightarrow \infty.$$

To simplify the notation in the remainder of the paper we will use the relations ‘ \gtrsim ’, ‘ \sim ’, and ‘ \lesssim ’ without explicit reference to $x \rightarrow \infty$.

Next, by applying Lemma 5 and letting $H \rightarrow \infty$, we conclude that

$$P[C > x] \gtrsim \sum_{k=1}^M \pi_k P[C^{(k)} > x]. \tag{8}$$

Let us now prove the upper bound. After splitting the integral in (2), we (schematically) define

$$P[C > x] = E \int_0^{\mathcal{T}_0} + E \int_{\mathcal{T}_0}^\infty =: I_1(x) + I_2(x). \tag{9}$$

First, we provide an upper bound for $I_1(x)$. Since $S(t; J) = S^{(k)}(t)$ almost surely on $\{J_0 = k\}$, we obtain

$$\begin{aligned}
 I_1(x) &= \mathbb{E} \sum_{k=1}^M \mathbf{1}_{\{J_0=k\}} \int_0^{\mathcal{T}_0} \sum_{i=1}^{\infty} (q_i^{(k)})^2 e^{-q_i^{(k)}t} \mathbb{P}[S_i^{(k)}(t) > x - 1] dt \\
 &\leq \sum_{k=1}^M \pi_k \int_0^{\infty} \sum_{i=1}^{\infty} (q_i^{(k)})^2 e^{-q_i^{(k)}t} \mathbb{P}[S_i^{(k)}(t) > x - 1] dt \\
 &= \sum_{k=1}^M \pi_k \mathbb{P}[C^{(k)} > x],
 \end{aligned} \tag{10}$$

where the inequality is obtained after replacing \mathcal{T}_0 with ∞ .

Next, in deriving an asymptotic estimate of $I_2(x)$, we use $q_i^{(J-t)} e^{-\hat{q}_i t} dt = -d(e^{-\hat{q}_i t})$ as follows:

$$\begin{aligned}
 I_2(x) &\leq \mathbb{E} \sum_{i=1}^{\infty} q_i^{(J_0)} \int_{\mathcal{T}_0}^{\infty} q_i^{(J-t)} e^{-\hat{q}_i t} dt \\
 &= \mathbb{E} \sum_{i=1}^{\infty} q_i^{(J_0)} \int_{\mathcal{T}_0}^{\infty} -d(e^{-\hat{q}_i t}) \\
 &= \mathbb{E} \sum_{i=1}^{\infty} q_i^{(J_0)} e^{-q_i^{(J_0)} \mathcal{T}_0} \\
 &= \sum_{k=1}^M \pi_k \sum_{i=1}^{\infty} \frac{q_i^{(k)} \delta}{q_i^{(k)} + \delta}.
 \end{aligned}$$

Since the first assumption of the theorem implies that $q_i^{(k)} \leq H/i^\alpha$, $1 \leq k \leq M$, use of the inequalities

$$\sum_{i=1}^{\infty} \frac{q_i^{(k)}}{q_i^{(k)} + \delta} \leq \sum_{i=1}^{\infty} \frac{1}{1 + h\delta i^\alpha} \leq \int_0^{\infty} \frac{1}{1 + h\delta z^\alpha} dz \leq \frac{1}{(h\delta)^{1/\alpha}} \int_0^{\infty} \frac{1}{1 + y^\alpha} dy$$

yields

$$I_2(x) \leq H\delta^{1-1/\alpha} = o\left(\frac{1}{x^{\alpha-1}}\right),$$

where the equality is implied by the assumption of the theorem, since $x\delta^{1/\alpha} \rightarrow 0$ as $\delta \rightarrow 0$ yields $\delta^{1-1/\alpha} = o(1/x^{\alpha-1})$. Finally, this last observation, together with (10) and Lemma 5, implies that $I_2(x) = o(I_1(x))$, which, in conjunction with (9) and (8), concludes the proof.

4.2. Asymptotic insensitivity

The following theorem establishes the scaling of the cache size as a function of the parameter, δ , for which the insensitivity result holds.

Theorem 2. *Let $q_i \sim c/i^\alpha$ as $i \rightarrow \infty$, with $\alpha > 1$. If x_δ satisfies $x_\delta \delta^{1/\alpha} / \log x_\delta \rightarrow \infty$ as $\delta \rightarrow 0$, then*

$$\mathbb{P}[C > x_\delta] \sim K(\alpha) \mathbb{P}[R > x_\delta] \quad \text{as } x_\delta \rightarrow \infty,$$

where

$$K(\alpha) := \left(1 - \frac{1}{\alpha}\right) \left[\Gamma\left(1 - \frac{1}{\alpha}\right)\right]^\alpha$$

and Γ is the gamma function.

Remark 2. The following proof is based on Theorem 3 of [13]. The main technical novelty of the present paper is to demonstrate that the estimates in this proof hold *uniformly* for all small δ . In this regard, the proof below draws these parallels and emphasizes the derivation of the uniform bounds.

Proof of Theorem 2. To simplify notation we again set $x \equiv x_\delta$. We first prove the upper bound. After splitting the integral in (2), we (schematically) define

$$P[C > x] = E \int_0^{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor}} + E \int_{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor}}^\infty =: I_1(x) + I_2(x). \tag{11}$$

Next, we show that $I_1(x)$ is negligible for large x , i.e.

$$I_1(x) = o\left(\frac{1}{x^{\alpha-1}}\right). \tag{12}$$

To this end, after conditioning on the value of $\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor}$, we obtain

$$I_1(x) \leq E \left[\mathbf{1}_{\{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor} > 2hx^\alpha\}} \int_0^{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor}} \hat{f}(t) P_{\sigma_t}[S(t; J) \geq x] dt \right] + E \left[\mathbf{1}_{\{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor} \leq 2hx^\alpha\}} \int_0^{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor}} \hat{f}(t) P_{\sigma_t}[S(t; J) \geq x] dt \right],$$

where $\hat{f}(t)$ is as defined in (3). Note that $\hat{f}(t) \leq \sum_{i=1}^\infty q_i^{(J_0)} = 1$ and

$$\int_0^\infty \hat{f}(t) dt = 1, \tag{13}$$

since

$$-d(e^{-\hat{q}_i t}) = e^{-\hat{q}_i t} d\left(\sum_{k=1}^M q_i^{(k)} \int_{-t}^0 \mathbf{1}_{\{J_u=k\}} du\right) = e^{-\hat{q}_i t} q_i^{(J-t)} dt.$$

Then, using (13),

$$I_1(x) \leq E \left[\mathbf{1}_{\{\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor} > 2hx^\alpha\}} \int_0^\infty \hat{f}(t) dt \right] + E \int_0^{2hx^\alpha} P_{\sigma_t}[S(t; J) \geq x] dt \leq P[\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor} > 2hx^\alpha] + 2hx^\alpha P[\bar{S}(2hx^\alpha) \geq x], \tag{14}$$

where $\bar{S}(t) := \sum_{i \geq 1} \bar{B}_i$ and the \bar{B}_i , $i \geq 1$, are independent Bernoulli random variables with $P[\bar{B}_i = 1] = 1 - e^{-\hat{q}_i t}$, in analogy with Equation (28) of [13]. Now, by Lemmas 2 and 3, the last term of (14) is $o(1/x^{\alpha-1})$.

Next, after using the Chernoff bound for the sum of exponential i.i.d. random variables, we obtain, for any $\theta, \delta > \theta > 0$,

$$\begin{aligned} \mathbb{P}[\mathcal{T}_{\lfloor hx^\alpha \delta \rfloor} > 2hx^\alpha] &\leq e^{-\theta 2hx^\alpha} \exp\left(-\lfloor hx^\alpha \delta \rfloor \log\left(1 - \frac{\theta}{\delta}\right)\right) \\ &\leq \exp\left(-\frac{1}{2}h \frac{x^\alpha \delta}{\log x} \log x\right) \\ &= o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \rightarrow \infty, \end{aligned}$$

since $x^\alpha \delta / \log x \rightarrow \infty$ as $x \rightarrow \infty$, which follows directly from the assumption that $x \delta^{1/\alpha} / \log x \rightarrow \infty$ as $\delta \rightarrow 0$. This, along with the estimates in (14), proves (12).

In order to estimate $I_2(x)$, we define the set $\mathcal{A}(n)$ as

$$\mathcal{A}(n) := \bigcap_{1 \leq k \leq M} \left\{ \left| \tau_k(\mathcal{T}_n) - \frac{\pi_k(n+1)}{\delta} \right| \leq 2\varepsilon \frac{\pi_k(n+1)}{\delta} \right\},$$

where $\tau_k(\mathcal{T}_n)$ represents the total time that process J spends in state k in the interval $(-\mathcal{T}_n, 0)$. Next, note that, owing to the memoryless property of the exponential distribution, we have $\tau_k(\mathcal{T}_n) \stackrel{D}{=} \sum_{i=0}^{N_n(k)} \varepsilon_i$, where, for $0 \leq i \leq n$, $N_n(k)$ represents the number of times that the Markov chain $\{J_{-\mathcal{T}_i}\}$ visits state k and ε_i are exponential i.i.d. random variables with mean $1/\delta$. Then

$$\begin{aligned} \mathbb{P}\left[\tau_k(\mathcal{T}_n) > (1 + \varepsilon) \frac{\pi_k(n+1)}{\delta}\right] &\leq \mathbb{P}[N_n(k) \geq (1 + \varepsilon)\pi_k(n+1)] \\ &\quad + \mathbb{P}\left[\sum_{i=0}^{\lceil (1+\varepsilon)\pi_k(n+1) \rceil} \varepsilon_i > (1 + 2\varepsilon) \frac{\pi_k(n+1)}{\delta}\right], \end{aligned} \tag{15}$$

where $\lceil x \rceil$ denotes the smallest integer greater than or equal to x .

Next note that, for any $\theta, 0 < \theta < \delta$, and any positive integer n ,

$$\begin{aligned} \mathbb{P}\left[\sum_{i=1}^n \varepsilon_i > (1 + \varepsilon) \frac{n}{\delta}\right] &= \mathbb{P}\left[\exp\left(\theta \sum_{i=1}^n \varepsilon_i\right) > e^{\theta(1+\varepsilon)n/\delta}\right] \\ &\leq \exp\left(-n \left[\frac{\theta}{\delta}(1 + \varepsilon) + \log\left(1 - \frac{\theta}{\delta}\right)\right]\right), \end{aligned}$$

where in the last expression we have used the Markov inequality. Therefore, by setting $u = \theta/\delta$ in the preceding expression, we obtain

$$\mathbb{P}\left[\sum_{i=1}^n \varepsilon_i > (1 + \varepsilon) \frac{n}{\delta}\right] \leq \inf_{0 < u < 1} \exp(-n[u(1 + \varepsilon) + \log(1 - u)]) = \exp(-n[\varepsilon + \log(1 + \varepsilon)]), \tag{16}$$

where the infimum is achieved for $u = \varepsilon/(1 + \varepsilon)$. Then, after applying a well-known large deviation result on finite-state ergodic Markov chains (see, e.g. Section 3.1.2 of [6]) to bound the first term of (15), and using (16), we conclude that there exists a constant $\theta_k(\varepsilon) > 0$, independent of δ , which satisfies

$$\mathbb{P}\left[\tau_k(\mathcal{T}_n) > (1 + \varepsilon) \frac{\pi_k(n+1)}{\delta}\right] \leq e^{-\theta_k(\varepsilon)n}. \tag{17}$$

Using arguments analogous to those in (15), (16), and (17) to estimate the exponential upper bound for $P[\tau_k(\mathcal{T}_n) < (1 - \varepsilon)\pi_k(n + 1)/\delta]$, from these results in conjunction with the union bound we conclude that

$$P[\mathcal{A}^c(n)] \leq \max_k P\left[\left|\tau_k(\mathcal{T}_n) - \frac{\pi_k(n + 1)}{\delta}\right| > 2\varepsilon \frac{\pi_k(n + 1)}{\delta}\right] \leq He^{-\theta_\varepsilon n} \tag{18}$$

for some positive constant $\theta_\varepsilon > 0$ independent of δ .

At this point, we are ready to proceed with the estimation of the integral $I_2(x)$. After multiplying $I_2(x)$ by $\mathbf{1}_{\{\mathcal{A}(n)\}}$ and $\mathbf{1}_{\{\mathcal{A}^c(n)\}}$, we define

$$\begin{aligned} I_2(x) &\leq E \sum_{n=\lfloor hx^\alpha \delta \rfloor}^\infty \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_t}[S(t; J) \geq x] dt \\ &= E \sum_{n=\lfloor hx^\alpha \delta \rfloor}^\infty \mathbf{1}_{\{\mathcal{A}^c(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_t}[S(t; J) \geq x] dt \\ &\quad + E \sum_{n=\lfloor hx^\alpha \delta \rfloor}^\infty \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_t}[S(t; J) \geq x] dt \\ &=: I_{21}(x) + I_{22}(x). \end{aligned} \tag{19}$$

Then, using (18), we obtain

$$I_{21}(x) \leq \sum_{n=\lfloor hx^\alpha \delta \rfloor}^\infty P[\mathcal{A}^c(n)] \leq He^{-\theta_\varepsilon hx^\alpha \delta} = o\left(\frac{1}{x^{\alpha-1}}\right) \text{ as } x \rightarrow \infty. \tag{20}$$

Next we estimate $I_{22}(x)$. Since $S(t; J)$ is almost surely nonincreasing in t , after splitting the sum we obtain

$$\begin{aligned} I_{22}(x) &\leq E \sum_{n=\lfloor hx^\alpha \delta \rfloor}^{\lfloor g_\varepsilon x^\alpha \delta \rfloor} \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_{\mathcal{T}_{n+1}}}[S(\mathcal{T}_{n+1}; J) \geq x] dt \\ &\quad + E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor + 1}^\infty \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) dt, \end{aligned} \tag{21}$$

where g_ε is to be defined later. Note that, for every $\omega \in \mathcal{A}(n)$ and $k, 1 \leq k \leq M$,

$$(1 - 2\varepsilon)\pi_k \frac{n + 1}{\delta} \leq \tau_k(\mathcal{T}_n) \leq (1 + 2\varepsilon)\pi_k \frac{n + 1}{\delta}. \tag{22}$$

Therefore, by definition (1),

$$P_{\sigma_{\mathcal{T}_n}}[B_i(\mathcal{T}_n; J) = 1] = 1 - \exp\left(-\sum_{k=1}^M q_i^{(k)} \tau_k(\mathcal{T}_n)\right) \leq P[B_i^*(n) = 1]$$

and

$$P_{\sigma_{\mathcal{T}_n}}[S(\mathcal{T}_n; J) \geq x] \leq P[S^*(n) \geq x], \tag{23}$$

where we define $S^*(n) := \sum_{i=1}^{\infty} B_i^*(n)$ with $\{B_i^*(n), i \geq 1\}$ representing a sequence of independent Bernoulli random variables such that

$$P[B_i^*(n) = 1] = 1 - \exp\left(-\frac{(1 + 2\varepsilon)q_i(n + 1)}{\delta}\right);$$

$S^*(n)$ is constructed to be nondecreasing in n . Then, if we define g_ε to be

$$g_\varepsilon := \frac{(1 - 2\varepsilon)^\alpha}{\Gamma(1 - 1/\alpha)^\alpha c(1 + 2\varepsilon)},$$

using arguments analogous to those leading to (66)–(67) of [13], we conclude that, for any large x , $E S^*(g_\varepsilon x^\alpha \delta) < (1 - \varepsilon)x$ and, therefore, by Lemma 3,

$$\begin{aligned} E \sum_{n=\lfloor hx^\alpha \delta \rfloor}^{\lfloor g_\varepsilon x^\alpha \delta \rfloor} \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_{\mathcal{T}_{n+1}}} [S(\mathcal{T}_{n+1}; J) \geq x] dt &\leq g_\varepsilon x^\alpha \delta P[S^*(g_\varepsilon x^\alpha \delta) \geq x] \\ &= o\left(\frac{1}{x^{\alpha-1}}\right). \end{aligned} \tag{24}$$

Next we derive the asymptotics of the second term in (21). Note that, for every $\omega \in \mathcal{A}(n)$ and $t \in (\mathcal{T}_n, \mathcal{T}_{n+1}]$, the bound from (22) results in

$$\hat{f}(t) \leq \sum_{i=1}^{\infty} q_i^{(J_0)} q_i^{(\mathcal{J}_{-n-1})} \exp\left(- (1 - 2\varepsilon)q_i \frac{n}{\delta}\right) \exp(-q_i^{(\mathcal{J}_{-n-1})}(t - \mathcal{T}_n)); \tag{25}$$

recall that $\mathcal{T}_n = -T_{-n}$ from above and that $J_{-\mathcal{T}_{n+1}} = J_{T_{-n-1}} = \mathcal{J}_{-n-1}$ from Section 2. Using the preceding bound in the second term of (21), computing the integration with respect to t and applying $1 - e^{-x} \leq x$ ($x \geq 0$) yields

$$\begin{aligned} E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor + 1}^{\infty} \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) dt &\leq E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor + 1}^{\infty} \mathbf{1}_{\{\mathcal{A}(n)\}} \sum_{i=1}^{\infty} q_i^{(J_0)} q_i^{(\mathcal{J}_{-n-1})} (\mathcal{T}_{n+1} - \mathcal{T}_n) \exp\left(- (1 - 2\varepsilon)q_i \frac{n}{\delta}\right) \\ &\leq \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor + 1}^{\infty} \sum_{i=1}^{\infty} \frac{1}{\delta} E[q_i^{(J_0)} E[q_i^{(\mathcal{J}_{-n-1})} \mid J_0]] \exp\left(- (1 - 2\varepsilon)q_i \frac{n}{\delta}\right). \end{aligned} \tag{26}$$

In this expression we now employ the asymptotic independence of the Markov chain J_n (as in Equations (19) and (20) of [13]) and the independence between $\{\mathcal{T}_n\}$ and $\{\mathcal{J}_{-n}\}$, and bound the resulting sum by an integral using the change of variable $t = n/\delta$. In conjunction with (24), this yields

$$\begin{aligned} I_{22}(x) &\leq (1 + \varepsilon) \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor + 1}^{\infty} \sum_{i=1}^{\infty} q_i^2 \frac{1}{\delta} \exp\left(- (1 - 2\varepsilon)q_i \frac{n}{\delta}\right) + o\left(\frac{1}{x^{\alpha-1}}\right) \\ &\leq (1 + \varepsilon) \int_{g_\varepsilon x^\alpha}^{\infty} \sum_{i=1}^{\infty} q_i^2 \exp(- (1 - 2\varepsilon)q_i t) dt + o\left(\frac{1}{x^{\alpha-1}}\right). \end{aligned}$$

Finally, by applying Lemma 1, we find that

$$\limsup_{x \rightarrow \infty} I_{22}(x)x^{\alpha-1} \leq K(\alpha) \frac{(1 + \varepsilon)^2(1 + 2\varepsilon)^{1-1/\alpha}}{(1 - 2\varepsilon)^{1+\alpha-1/\alpha}}.$$

In conjunction with (21), (20), (19), and (11), this, as we let $\varepsilon \rightarrow 0$, proves the upper bound.

The estimation of the lower bound of (2) starts from

$$P[C > x] \geq E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor}^\infty \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_{\mathcal{T}_n}}[S(\mathcal{T}_n; J) \geq x] dt, \tag{27}$$

where $g_\varepsilon := (1 + 2\varepsilon)^\alpha \Gamma(1 - 1/\alpha)^{-\alpha} c^{-1} (1 - \varepsilon)^{-1}$. Using arguments analogous to those leading to (23), and redefining

$$P[B_i^*(n) = 1] = 1 - \exp\left(-\frac{(1 - 2\varepsilon)q_i(n + 1)}{\delta}\right), \quad i \geq 1,$$

we obtain

$$\begin{aligned} E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor}^\infty \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) P_{\sigma_{\mathcal{T}_n}}[S(\mathcal{T}_n; J) \geq x] dt \\ \geq P[S^*(g_\varepsilon x^\alpha \delta) > x] E \sum_{n=\lfloor g_\varepsilon x^\alpha \delta \rfloor}^\infty \mathbf{1}_{\{\mathcal{A}(n)\}} \int_{\mathcal{T}_n}^{\mathcal{T}_{n+1}} \hat{f}(t) dt. \end{aligned}$$

We complete the proof of the theorem by applying arguments analogous to those used in (25) and (26), to bound $\hat{f}(t)$ and, therefore, the integral on the right-hand side of (27) from below for all $\omega \in \mathcal{A}(n)$ and $t \in [\mathcal{T}_n, \mathcal{T}_{n+1})$. Then, in conjunction with asymptotic independence and bounding arguments analogous to those used in Equations (19)–(21) of [13], we can easily complete the proof of the lower bound for all sufficiently small δ . The details of this are omitted because the arguments are repetitive.

4.3. Discussion

Note that when $x < 1/\delta^p$ for some $p > 0$, the condition $x\delta^{1/\alpha} / \log x \rightarrow \infty$ of Theorem 2 is implied by $x\delta^{1/\alpha} / \log(1/\delta) \rightarrow \infty$. Thus, for sufficiently large H and all $x > H \log(1/\delta) / \delta^{1/\alpha}$, the cache behaves like the corresponding i.i.d. system with marginal distribution $\{q_i\}$. Hence, under this asymptotic scaling, the correlation structure plays no role. On the other hand, Theorem 1 states that for very small caches, $x \leq 1/(H\delta^{1/\alpha})$, the cache performance is distinctly different from that of the corresponding i.i.d. system; in fact, the fault probability is decomposed into a mixture of i.i.d. systems. Informally, we see that this qualitative transition in the cache performance occurs around cache sizes on the order of $1/\delta^{1/\alpha}$. As previously noted, this value is sublinear (relatively negligible) in comparison with the time scale of jumps ($1/\delta$) in the modulating process J .

In order to gain additional insight into the qualitative behavior underlying Theorems 1 and 2, consider the expected time between two successive requests for a document during which the underlying Markov chain J_t is in a fixed state. The expected length of this time interval is inversely proportional to the document’s conditional access frequency, and the LRU algorithm thus has a tendency, in stationarity, to arrange the documents in the cache list in (approximately) descending order of their access probabilities. Therefore, it can be intuitively expected that the

access probabilities of documents at the end of the cache list are on the order of $x^{-\alpha}$, which, from the above arguments, implies that the time period during which every document in the cache is accessed at least once is of the order x^α . Hence, if the expected sojourn time that the modulated process spends in a particular state, $1/\delta$, is much greater than x^α , i.e. $1/\delta \gg x^\alpha$, then the cache content basically goes through many replacement cycles and the cache essentially reaches stationarity while the underlying modulating process remains in the same state, resulting in the decomposition result presented in Theorem 1. On the other hand, if $1/\delta \ll x^\alpha$ then the Markov chain J_t undergoes significant mixing between the successive requests for documents that are at the end of the cache list. Therefore, successive requests for documents with probabilities smaller than $1/x^\alpha$, which essentially determine the cache fault probability, appear to be nearly independent, implying that the cache fault probability is the same as it would be if the requests were i.i.d., as we have rigorously shown in Theorem 2.

In the critical regime, when $1/\delta \approx x^\alpha$, the above arguments suggest that the time scale of jumps in J_t and the access frequencies of documents at the end of the cache list are comparable. Therefore, the cache fault probability is a result of an intricate and complex interplay between the modulating chain dependence structure and the conditional access frequencies. While deeper understanding of this critical regime is important for a complete mathematical understanding of the problem, we believe that the potential asymptotic results will likely be more difficult to prove, and, unlike those in Theorems 1 and 2, will not be explicit. In addition, even asymptotically, the cache fault probability in the critical regime will be a complex functional of the transition probabilities of J_t and the conditional access frequencies $q_i^{(k)}$, whose further understanding will require numerical studies. Therefore, we do not pursue this direction further.

5. Concluding remarks

In this paper we have investigated the performance, namely the cache fault probability, of LRU caches in the presence of correlated requests. It has been recently discovered, in [12] and [13], that, for semi-Markov-modulated requests and marginal access frequencies obeying a generalized Zipf's law, the caching performance does not depend on the correlation in the request traffic for large cache sizes. Specifically, LRU cache performance is asymptotically identical to the case of i.i.d. requests that have the same access frequencies. However, for small caches this is clearly not the case. Hence, in our present study we have investigated the smallest (critical) cache size above which the discovered asymptotic insensitivity property still holds. Our approach is based on the use of a joint scaling between the request process dependence structure and the cache size. More precisely, we consider requests that are modulated by nearly completely decomposable Markov processes with small transition rate, δ , such that the cache size, x_δ , grows to infinity as $\delta \downarrow 0$. Then, extending the analytical techniques of [12] and [13], we showed in Theorems 1 and 2 that, maybe somewhat surprisingly, the critical scaling of the cache size is very small in relation to the time scale of the request process dependence structure; basically, it is sublinear in the average time, $1/\delta$, between the jumps in the modulating process. Hence, from a practical perspective, it may not be necessary to model in great detail the request process dependence structure found in Web environments. In addition, it is worth noting that our results can be extended to the case with variable document sizes by exploiting the recent analysis in [14].

Acknowledgement

We are very grateful to the anonymous reviewer for valuable suggestions.

References

- [1] ALMEIDA, V., BESTAVROS, A., CROVELLA, M. AND DE OLIVIERA, A. (1996). Characterizing reference locality in the WWW. In *Proc. 4th Internat. Conf. Parallel Distributed Inf. Systems*, IEEE Computer Society, Washington, DC, pp. 92–107.
- [2] ASMUSSEN, S., HENRIKSEN, L. F. AND KLÜPPELBERG, C. (1994). Large claims approximations for risk processes in a Markovian environment. *Stoch. Process. Appl.* **54**, 29–43.
- [3] BENTLEY, J. L. AND MCGEOCH, C. C. (1985). Amortized analyses of self-organizing sequential search heuristics. *Commun. Assoc. Comput. Mach.* **28**, 404–411.
- [4] BORODIN, A. AND EL-YANIV, R. (1998). *Online Computation and Competitive Analysis*. Cambridge University Press.
- [5] BRESLAU, L., CAO, P., FAN, L., PHILLIPS, G. AND SHENKER, S. (1999). Web caching and Zipf-like distributions: evidence and implications. In *Proc. INFOCOM '99*, Vol. 1 (New York, March 1999), pp. 126–134.
- [6] DEMBO, A. AND ZEITOUNI, O. (1993). *Large Deviations Techniques and Applications*. Jones and Bartlett, Boston, MA.
- [7] FILL, J. A. (1996). An exact formula for the move-to-front rule for self-organizing lists. *J. Theoret. Prob.* **9**, 113–160.
- [8] FILL, J. A. AND HOLST, L. (1996). On the distribution of search cost for the move-to-front rule. *Random Structures Algorithms* **8**, 179–186.
- [9] FLAJOLET, P., GARDY, D. AND THIMONIER, L. (1992). Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Appl. Math.* **39**, 207–229.
- [10] JELENKOVIĆ, P. R. (1999). Asymptotic approximation of the move-to-front search cost distribution and least-recently used caching fault probabilities. *Ann. Appl. Prob.* **9**, 430–464.
- [11] JELENKOVIĆ, P. R. AND LAZAR, A. A. (1998). Subexponential asymptotics of a Markov-modulated random walk with queueing applications. *J. Appl. Prob.* **35**, 325–347.
- [12] JELENKOVIĆ, P. R. AND RADOVANOVIĆ, A. (2003.). Asymptotic insensitivity of least-recently-used caching to statistical dependency. In *Proc. INFOCOM 2003*, Vol. 1 (San Francisco, CA, March–April 2003), pp. 438–447.
- [13] JELENKOVIĆ, P. R. AND RADOVANOVIĆ, A. (2004). Least-recently-used caching with dependent requests. *Theoret. Comput. Sci.* **326**, 293–327.
- [14] JELENKOVIĆ, P. R. AND RADOVANOVIĆ, A. (2004). Optimizing LRU caching for variable document sizes. *Combinatorics Prob. Comput.* **13**, 627–643.
- [15] SLEATOR, D. D. AND TARJAN, R. E. (1985). Self-adjusting binary search trees. *J. Assoc. Comput. Mach.* **32**, 652–686.
- [16] YOUNG, N. E. (2000). On-line paging against adversarially biased random inputs. *J. Algorithms* **37**, 218–235.