LRU Caching with Dependent Competing Requests

Guocong Quan, Kaiyi Ji, Jian Tan
Department of Electrical and Computer Engineering
The Ohio State University, Columbus, 43210
{quan.72, ji.367, tan.252}@osu.edu

***Abstract*—Least-recently-used (LRU) caching systems have been widely used, and are increasingly deployed driven by emerging trends for big data. In a typical scenario, these systems are used to serve multiple flows of dependent data item requests that are also correlated over time. These flows compete for the limited cache space. Characterizing the miss ratios of these competing flows can facilitate the design and improve the system performance. The existing asymptotic analyses for correlated requests give explicit results for Zipf's distributions with the index greater than a critical value (one). Consequently, the asymptotic result is inaccurate around this critical point, which notably is also the typical parameter region reported by many empirical measurements. In contrast, we derive the asymptotic miss ratios of multiple flows for a large class of truncated heavy-tailed data item popularity distributions with time dependency. Importantly, it significantly improves the accuracy in numerical computations when the index of a Zipf's distribution is close to one. Moreover, the result generalizes beyond Zipf's distributions, e.g., to Weibull, for multiple flows of correlated data item requests. Our asymptotic result directly exploits the critical properties of the distribution and the truncated support region. As our versatile expression is explicit, it avoids the numerical computations required by the characteristic time approximation. Interestingly, it also validates the characteristic time approximation with new forms for multiple flows of competing requests that are correlated over time under certain conditions.**

I. Introduction

Caching systems are a core component of Internet data infrastructures. They greatly improve the performance of various Web services, e.g., for information retrieval, data analytics, social networks and e-commence, by enabling low-cost access to a fast, but limited cache space. From a large collection of data records that are stored in slow but persistent media, a selective subset of popular data items can be temporarily put in the cache to accelerate data processing. Driven by emerging trends for big data, caching systems are widely deployed [1], [2], [3], [4]. They typically serve multiple flows of data requests that are also correlated over time [5], [6], which compete for the limited cache space. These flows of requests could have different data popularities, time correlations, varying item sizes, different request rates and even overlapped data items shared across different flows, which jointly impact system performances.

In current engineering, the least-recently-used (LRU) algorithm and its extensions [7], [8], [9], [10] have been widely used in data caching systems [1], [11], [5], [2], [3], [4]. The LRU algorithm is appealing since it is self-organizing with low cost in tracking history data. When a request arrives, if the requested data can be found in the cache, we call it a "hit"; otherwise, we call it a "miss". When a miss occurs and the cache is full, LRU moves the item or items that have not been requested for the longest time (least recently used) out of the cache to make room for the newly requested one.

This work was supported by the National Science Foundation under Grant No. 1717060.

Due to the importance of LRU caching, we derive new, easy-to-compute and accurate asymptotic results for the miss ratios (a.k.a. miss probabilities) of multiple flows correlated over time under a large class of truncated heavy-tailed popularity distributions. For example, we can address a generalized Zipf's distribution with requests modulated by a finite-state Markov chain. The introduced correlations capture the dynamics when the requests have time-varying popularity distributions. Although there have been extensive efforts to characterize the miss ratios of LRU caching systems [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], most of them either focus on independent requests or a single flow with equal data item sizes. We develop a framework for multiple competing flows to jointly address different factors in the same model, including, e.g., time correlations, a large class of (truncated) heavy-tailed popularity distributions, and variable data item sizes.

We resort to asymptotic analysis, since it gives explicit results and provides direct insights based on the popularity distributions. Existing works often assume a Zipf's distribution (i.e. $p_i = c/i^{\alpha}$) and treat the index $\alpha > 1$ and $0 < \alpha < 1$ separately. Consequently, the asymptotic result is not accurate around the critical point $\alpha \approx 1$. In contrast, we derive a uniform result that coherently covers both cases with a smooth transition. More importantly, it significantly improves the accuracy in numerical computations when $\alpha \approx 1$, and generalizes beyond Zipf's distributions, e.g., to regularly varying and heavy-tailed Weibull distributions. Notably, numerous empirical measurements have shown that the data popularity distributions often follow a Zipf's distribution, e.g., $\alpha \in (0.6, 0.8)$ in web caching systems [22] and $\alpha \in (0.2, 1.2)$ at the block I/O level [23]; other reports also exist [5]. These observations strongly motivate our analysis for $\alpha \approx 1$.

In addition to asymptotic analysis, another commonly used technique in analyzing miss ratios is based on the characteristic time approximation [16], [24]. Compared to this approach, our asymptotic result is explicit, since it directly exploits the critical properties of the distribution and the truncated support region. Thus, it avoids the numerical computations required by the characteristic time approximation. Interestingly, our result implies a new form of the characteristic time approximation for time-dependent competing requests, which can be proved asymptotically for large cache sizes under certain conditions.

Related work

Under independent reference model (IRM), the miss probability of LRU caching has been extensively studied. When cache sizes are relatively small, the miss ratio can be explicitly calculated [12], [13], [14], [15]. For large cache sizes, several approaches have been adopted, such as the commonly used characteristic time approximation [16], [24] and mean field approximations [17], [18], [19]. For some specific popularity distributions (e.g., Zipf's and Weibull distributions), the asymptotic miss ratio can be explicitly expressed. The studies on a single flow of independent requests following a Zipf's distribution with $\alpha > 1$ and $0 < \alpha < 1$ are investigated in [20] and [21], respectively. For these two cases, a unified form for Zipf's distributions with $\alpha > 0, \alpha \neq 1$ is proposed in [25]. The asymptotic miss ratios for Weibull distributions are studied in [26]. For multiple flows, whether caching systems should use resource pooling or separation is investigated in [27].

For correlated requests, a few approximations for miss ratios have been derived [28], [29], [30], which, however, incur a high computation cost. For Zipf's law with $\alpha > 1$, the asymptotic miss ratios can be explicitly expressed [31], [32], but the estimation is not accurate for $\alpha \approx 1$ and does not address multiple competing flows. A fluid limit for general request processes is derived [33], whose asymptotic result for Zipf's law agrees with [31], [32]. Due to the lack of a general analytic framework, we propose a unified approach to study the asymptotic miss ratio for multiple dependent competing flows for a large class of (truncated) heavy tailed distributions.

Summary of contributions

(1) We derive new, easy-to-compute and accurate asymptotic results for the miss probabilities of multiple flows with dependent requests served on a common LRU cache. Compared with existing asymptotic analyses [32], [31], our result is able to analyze generalized Zipf's popularity distributions regularly varying with $\alpha > 1$ and $0 < \alpha < 1$ in a unified model and is more accurate for $\alpha \approx 1$ with small support regions.

(2) We extend the existing works on Zipf's popularity distributions [31], [32] to a broad class of heavy-tailed distributions, including regularly varying and heavy-tailed Weibull distributions with time correlations. In addition, we prove a new form of the characteristic time approximation for dependent competing requests under certain conditions.

(3) We conduct extensive numerical simulations to verify the theoretical results. All results show a good match.

II. Model Description

Consider K flows of dependent requests sharing a common LRU cache. To model correlations of the requests, define $\{\Pi_t\}_{t\in\mathbb{R}}$ as a stationary and ergodic modulating process with finitely many states $\{1, 2, \cdots, M\}$ and stationary distribution $\pi_m = \mathbb{P}[\Pi_t = m]$. Data popularities and request rates vary in different states. Assume that the arrivals of flow k in state m follow a Poisson process with rate $\lambda_{k,m}$, $1 \leq k \leq K, 1 \leq m \leq M$. The requests of the aggregated flow arrive at time points $\{\tau_n, -\infty < n < \infty\}$. For simplicity, we use $\Pi_n \equiv \Pi_{\tau_n}$ to denote the state at time τ_n. Let I_n denote the index of the flow for the request arrived at τ_n. The event $\{I_n = k\}$ represents that the request at τ_n is from flow k. We have $\mathbb{P}[I_n = k|\Pi_n = m] = \lambda_{k,m}/\sum_{i=1}^{K}\lambda_{i,m}$. Let data

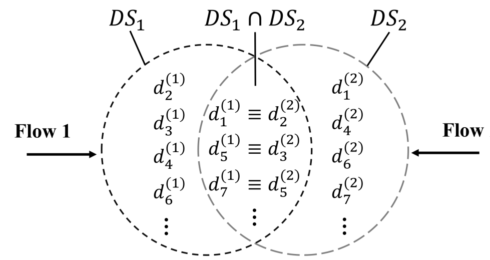

Fig. 1. Data items overlap between two flows

set $DS_k = \{d_1^{(k)}, d_2^{(k)}, \cdots, d_{N_k}^{(k)}\}$ denote the set of data items requested by flow k, where $d_i^{(k)}$ is the i'th data item of the set DS_k, $1 \leq i \leq N_k$, $1 \leq k \leq K$. Let $s_i^{(k)}$ denote the item size of $d_i^{(k)}$. The data items in the same data set are distinct. However, different data sets can have overlap. We use $d_i^{(k)} \equiv d_j^{(g)}$ to represent that $d_i^{(k)}$ and $d_j^{(g)}$ are the same data item. For example, in Fig. 1, we illustrate two overlapped data sets with common data items $d_1^{(1)} \equiv d_2^{(2)}$, $d_5^{(1)} \equiv d_3^{(2)}$ and $d_7^{(1)} \equiv d_5^{(2)}$. Let R_n be the data item requested at τ_n. The event $\{I_0 = k, R_0 = d_i^{(k)}\}$ represents that the request at τ_n is from flow k to fetch data item $d_i^{(k)}$. Note that $\{R_n = d_i^{(k)}\}$ and $\{R_n = d_j^{(g)}\}$ are the same event if we have $d_i^{(k)} \equiv d_j^{(g)}$. Define the popularities

$$\mathbb{P}\left[R_0 = d_i^{(k)} \middle| \Pi_0 = m\right] = p_i^{(k,m)}, \quad (1)$$

$$\mathbb{P}\left[R_0 = d_i^{(k)} \middle| I_0 = k, \Pi_0 = m\right] = q_i^{(k,m)}. \quad (2)$$

Let $\nu_{k,m} \triangleq \mathbb{P}[I_0 = k|\Pi_0 = m]$. We have

$$p_i^{(k,m)} = \sum_{j=1}^{K} \nu_{j,m}\mathbb{P}\left[R_0 = d_i^{(k)} \middle| I_0 = j, \Pi_0 = m\right].$$

In general, $p_i^{(k,m)}$ can be very different from $q_i^{(k,m)}$, since some data items can be requested by multiple flows, as shown in Fig. 1.

Let $\pi_{k,m} = \mathbb{P}[\Pi_0 = m|I_0 = k]$, $q_i^{(k)} = \sum_{m=1}^{M}\pi_{k,m}q_i^{(k,m)}$, $p_i^{(k)} = \sum_{m=1}^{M}\pi_m p_i^{(k,m)}$. In DS_k, assume the data items $d_i^{(k)}$ are sorted such that the sequence $p_i^{(k)}$ is non-increasing with respect to i. Note that $q_i^{(k)}$ is not necessarily non-increasing by this ordering. Now, for flow k, we introduce two functions $(\Psi_k(\cdot), \Theta_k(\cdot))$, both defined in a neighborhood of infinity. For any $\zeta > 1$ and $(\Psi_k(\cdot), \Theta_k(\cdot))$ independent of ζ, we investigate how the following functional relationship impacts the miss ratio, by keeping $N_k = \zeta y$ and letting $y \to \infty$,

$$\sum_{i=y}^{N_k} q_i^{(k)} \sim \Psi_k\left(\left(p_y^{(k)}\right)^{-1}\right) + \Theta_k(N_k). \quad (3)$$

Note that (3) contains two additively separable terms, which holds for all $\zeta > 1$. The notation $g(x) \sim h(x)$ means $\lim_{x\to\infty} g(x)/h(x) = 1$.

Specifically, we consider a class of distributions that satisfy

$$\lim_{n\to\infty} q_n^{(k)}/q_{n+1}^{(k)} = 1, \tag{4}$$

and

$$\Psi_k(x) \sim x^{-\beta_k} l_k(x), \beta_k \neq 0, \tag{5}$$

where $l_k(x)$ is slowly varying [34]. A function $l_k(x) : \mathbb{R}^+ \to \mathbb{R}^+$ is slowly varying if for any $\lambda > 0, l_k(\lambda x)/l_k(x) \to 1$ as $x \to \infty$; and $\Psi_k(x) = x^{-\beta_k} l_k(x)$ is called regularly varying of index $-\beta_k$, where β_k can be even negative. A large class of heavy-tailed distributions satisfy (4) and (5), e.g., heavy-tailed Weibull distributions $q_i^{(k)} \sim d\exp\left(-ci^{\xi}\right)$ with $c, d > 0, 0 < \xi < 1$ and regularly varying distributions $q_i^{(k)} \sim l(i)/i^{\alpha}$, $\alpha > 0, \alpha \neq 1$ with $l(i)$ being slowly varying. Additionally, for $\epsilon > 0, 1 \leq m \leq M$ and $\beta^* = \max_{1\leq k\leq K}\{|\beta_k|\}$, the modulating process is assumed to satisfy, as $n \to \infty$,

$$\mathbb{P}\left[\left|\sum_{i=1}^{n} \mathbf{1}\{\Pi_i = m\} - \pi_m n\right| > \epsilon n\right] = o(n^{-\beta^*}), \tag{6}$$

where $g(n) = o(h(n))$ means $\lim_{n\to\infty} g(n)/h(n) = 0$.

LRU is equivalent to the move-to-front (MTF) policy [35], [36], [20], [31], which sorts the data items in the increasing order of their last access times. Each time a request is made, the requested data item is moved to the first position of the list and all the items that were in front of it increase their positions by one. The move-to-front (MTF) policy has the same miss probability as a LRU cache. Thus, we can define the miss ratio of a LRU by the searching cost of MTF.

Definition 1. *Define C_n to be the summation of the sizes for all the data items in the sorted list under MTF that are in front of the position of the data item corresponding to the request R_n made at time τ_n.*

If the cache capacity is x, a cache miss under LRU policy, which is equivalent for MTF, can be denoted by $\{C_n > x\}$. If the data items are unit-sized, i.e. $s_i^{(k)} \equiv 1$ for all k, i, the event $\{C_n > x\}$ means the position of the requested data item in the list is larger than x under MTF. When the system reaches its stationary, we only need to consider the miss probability at time τ_0 [31].

III. Main results

For the K flows sharing a cache, denote by $DS = \bigcup_{k=1}^{K} DS_k = \{d_i^{\circ}, i = 1, 2, \cdots, N\}$ the set of data items requested by the entirety of these flows, with $\mathbb{P}[R_0 = d_i^{\circ} | \Pi_0 = m] = p_i^{(\circ,m)}$. Let s_i be the size of data item d_i° and assume $\bar{s} \triangleq \sup_i s_i < \infty$. In general, s_i can take different values when the items have various sizes. Define

$$m(x) = \sum_{i=1}^{N} s_i \left(1 - \prod_{m=1}^{M}\left(1 - p_i^{(\circ,m)}\right)^{\pi_m x}\right), \tag{7}$$

which is an increasing function with an inverse $m^{\leftarrow}(x)$.

Assume that, for $\epsilon \in (0, 1)$ and a function $0 < \delta(x) \leq 1$,

$$\lim_{x\to\infty} \frac{m^{\leftarrow}((1+\epsilon\delta(x))x)}{m^{\leftarrow}(x)} = f(\epsilon), \qquad \lim_{\epsilon\to 0} f(\epsilon) = 1, \tag{8}$$

$$\lim_{x\to\infty} \frac{m((1+\epsilon)x)}{m(x)} = g(\epsilon), \qquad \lim_{\epsilon\to 0} g(\epsilon) = 1. \tag{9}$$

In addition, there exist $h_2 > h_1 > 0, h_4 > h_3 > 0$ and x_0, for $x > x_0$,

$$h_1 < \frac{\delta(x)}{\delta(x+\epsilon\delta(x))} < h_2, h_3 < \frac{\delta(x-\epsilon\delta(x))}{\delta(x)} < h_4. \tag{10}$$

A. Miss ratios of competing flows

For LRU caching with competing flows, the miss ratio of flow k with dependent requests can be characterized by the following theorem. Let $\Gamma(\beta, s) = \int_s^{\infty} x^{\beta-1}e^{-x}dx$ be the incomplete gamma function.

Theorem 1. *Under the conditions (4) - (10), for $N_k = \zeta m^{\leftarrow}(x), \zeta > 0$, $1 \leq k \leq K$, as $x \to \infty$, we have*

$$\mathbb{P}\left[C_0 > x | I_0 = k\right] \sim \beta_k \Gamma\left(\beta_k, m^{\leftarrow}(x) p_{N_k}^{(k)}\right) \Psi_k(m^{\leftarrow}(x)).$$

The proof is presented in Section VI.

A large class of heavy-tailed distributions including regularly varying, Weibull and Zipf's distributions satisfy the conditions in Theorem 1. For a special case with a single flow ($K = 1$) of correlated requests following a truncated Zipf's law ($p_i^{(1)} \sim c/i^{\alpha}, 1 \leq i \leq N_1$), the asymptotic miss ratios have been derived for $\alpha > 1$ [31], [32]. Applying Theorem 1, we can compute the asymptotic result for $\alpha > 1$ and $0 < \alpha < 1$ in a unified form in Corollary 1. More importantly, it gives more accurate numerical results compared to [31], [32]. The results for a single flow of independent requests following Zipf's law, including [20] ($\alpha > 1$), [21] ($0 < \alpha < 1$) and [25] ($\alpha > 0, \alpha \neq 1$) can be also covered by Thoerem 1; see Corollary 1. In addition, Theorem 1 also characterize the miss ratios for multiple competing flows of correlated requests with overlapped data items.

In order to explicitly compute $m(x)$, we establish the following lemma. Let $p_i^{\circ} = \sum_{m=1}^{M} \pi_m p_i^{(\circ,m)}$, $1 \leq i \leq N$. For $1 \leq k \leq K, 1 \leq m \leq M$, define

$$\bar{m}(x) = \sum_{i=1}^{N} s_i \left(1 - e^{-p_i^{\circ} x}\right), \tag{11}$$

$$\bar{m}^{(k)}(x) = \sum_{i=1}^{N_k} s_i^{(k)} \left(1 - \exp\left(-\sum_{m=1}^{M} \pi_m \nu_{k,m} q_i^{(k,m)} x\right)\right).$$

Lemma 1. *For K flows sharing a cache without overlapped data items, under (9), we have, as $x \to \infty$,*

$$m(x) \sim \bar{m}(x) = \sum_{k=1}^{K} \bar{m}^{(k)}(x). \tag{12}$$

The proof is presented in Section VI.

Using Theorem 1 and Lemma 1, we study some special cases that allow time correlated requests.

Corollary 1. *Under condition (6), consider a single flow with unit-sized items and* $q_i^{(1)} \sim c/i^{\alpha}, i = 1, 2, \cdots N, \alpha > 0, \alpha \neq 1$*. For any* $\zeta > 0$ *and* $N = \zeta m^{\leftarrow}(x)$*, we have, as* $x \to \infty$*,*

$$\mathbb{P}[C_0 > x] \sim \Gamma\left(1 - \frac{1}{\alpha}, \frac{cm^{\leftarrow}(x)}{N^{\alpha}}\right) \frac{c^{1/\alpha}}{\alpha m^{\leftarrow}(x)^{1-1/\alpha}}, \quad (13)$$

where $m^{\leftarrow}(x)$ *is the unique solution of the equation*

$$\begin{aligned} x =& \Gamma\left(1 - \frac{1}{\alpha}, \frac{cm^{\leftarrow}(x)}{N^{\alpha}}\right) (cm^{\leftarrow}(x))^{1/\alpha} \\ &+ N(1 - e^{-cm^{\leftarrow}(x)/N^{\alpha}}). \end{aligned} \quad (14)$$

In particular, for $\alpha > 1$, since $\lim_{x\to\infty} m^{\leftarrow}(x)/N^{\alpha} = 0$, equation (13) can be further simplified, at the expense of accuracy, to

$$\mathbb{P}[C_0 > x | I_0 = 1] \sim \Gamma\left(1 - \frac{1}{\alpha}\right)^{\alpha} \frac{c}{\alpha x^{\alpha-1}}, \quad (15)$$

which is the result of Theorem 2 in [31]. In Section IV, simulations verify that equation (13) is more accurate than [31], when N is relatively small or α is around 1. By setting $M = 1$, Corollary 1 also covers the results for a single flow of independent requests following Zipf's law [20], [21], [25].

Proof. Since $q_i^{(1)} = p_i^{(1)} = c/i^{\alpha}$, we obtain

$$\sum_{i=x}^{N} q_i^{(1)} \sim \int_x^N \frac{c}{t^{\alpha}} dt = \frac{c}{(\alpha-1)x^{\alpha-1}} - \frac{c}{(\alpha-1)N^{\alpha-1}},$$

which implies, $\Psi_1(x) \sim c^{1/\alpha} x^{1/\alpha - 1}/(\alpha - 1)$, as $x \to \infty$. Using Lemma 1, we have

$$\begin{aligned} m(x) &\sim \sum_{i=1}^{N} \left(1 - e^{-cx/i^{\alpha}}\right) \sim \int_1^N \left(1 - e^{-cx/t^{\alpha}}\right) dt \\ &\sim \Gamma\left(1 - \frac{1}{\alpha}, q_N^{(1)} x\right) (cx)^{1/\alpha} + N\left(1 - e^{-q_N^{(1)} x}\right). \end{aligned}$$

Then, applying Theorem 1, we complete the proof. ☐

Corollary 2. *Under condition (6), consider* K *flows of dependent requests without overlapped data items. For* $q_i^{(k,m)} = c_{k,m}/i^{\alpha_{k,m}}$*,* $\alpha_{k,m} > 0$*,* $\alpha_{k,m} \neq 1$*,* $1 \leq i \leq N_k, 1 \leq k \leq K, 1 \leq m \leq M$*, let* $\tilde{\alpha}_k = \min_{1\leq m\leq M} \alpha_{k,m}$*,* $S_k = \{m : \alpha_{m,k} = \tilde{\alpha}_k\}$*,* $\tilde{c}_k = \sum_{m\in S_k} \pi_{k,m} c_{k,m}$ *and* $\tilde{\nu}_k = \sum_{m\in S_k} \pi_m \nu_{k,m} c_{k,m}/\tilde{c}_k$ $1 \leq k \leq K$*. Assume that, for each* $1 \leq k \leq K$*, the data items of flow* k *have an identical size* $s^{(k)}$*, i.e.* $s_i^{(k)} = s^{(k)}$*,* $1 \leq i \leq N_k$*. Then, for any* $\zeta_k > 0$ *and* $N_k = \zeta_k m^{\leftarrow}(x)$*, we have, as* $x \to \infty$*,*

$$\begin{aligned} &\mathbb{P}[C_0 > x | I_0 = k] \\ &\sim \Gamma\left(1 - \frac{1}{\tilde{\alpha}_k}, \frac{\tilde{c}_k \tilde{\nu}_k}{N_k^{\tilde{\alpha}_k}} m^{\leftarrow}(x)\right) \frac{\tilde{c}_k^{1/\tilde{\alpha}_k}}{\tilde{\alpha}_k (\tilde{\nu}_k m^{\leftarrow}(x))^{1-1/\tilde{\alpha}_k}}, \end{aligned} \quad (16)$$

where $m^{\leftarrow}(x)$ *is the unique solution of*

$$\begin{aligned} x =& \sum_{k=1}^{K} \Gamma\left(1 - \frac{1}{\tilde{\alpha}_k}, \frac{\tilde{c}_k \tilde{\nu}_k}{N_k^{\tilde{\alpha}_k}} m^{\leftarrow}(x)\right) s^{(k)} \left(\tilde{c}_k \tilde{\nu}_k m^{\leftarrow}(x)\right)^{1/\tilde{\alpha}_k} \\ &+ \sum_{k=1}^{K} s^{(k)} N_k \left(1 - \exp\left(-\frac{\tilde{c}_k \tilde{\nu}_k}{N_k^{\tilde{\alpha}_k}} m^{\leftarrow}(x)\right)\right). \end{aligned}$$

Proof. Since $q_i^{(k,m)} = c_{k,m}/i^{\alpha_{k,m}}$, we have

$$q_i^{(k)} = \sum_{m=1}^{M} \frac{\pi_{k,m} c_{k,m}}{i^{\alpha_{k,m}}} \sim \sum_{m\in S_k} \frac{\pi_{k,m} c_{k,m}}{i^{\tilde{\alpha}_k}} = \frac{\tilde{c}_k}{i^{\tilde{\alpha}_k}}$$

and $p_i^{(k)} = \sum_{m=1}^{M} \pi_m \nu_{k,m} q_i^{(k,m)} \sim \tilde{\nu}_k q_i^{(k)}$. Using the same approach as in the proof of Corollary 1, we obtain, as $x \to \infty$,

$$\Psi_k(x) \sim \frac{\tilde{c}_k^{1/\tilde{\alpha}_k} (\tilde{\nu}_k x)^{1/\tilde{\alpha}_k - 1}}{\tilde{\alpha}_k - 1} \quad (17)$$

and

$$\begin{aligned} m(x) \sim& \sum_{k=1}^{K} \Gamma\left(1 - \frac{1}{\tilde{\alpha}_k}, q_{N_k}^{(k)} \tilde{\nu}_k x\right) s^{(k)} \left(\tilde{c}_k \tilde{\nu}_k x\right)^{1/\tilde{\alpha}_k} \\ &+ \sum_{k=1}^{K} s^{(k)} N_k \left(1 - e^{-q_{N_k}^{(k)} \tilde{\nu}_k x}\right). \end{aligned} \quad (18)$$

Combining (17), (18) and Theorem 1 completes the proof. ☐

For flows with overlapped data items, we can decompose them into sub-flows that have no overlapped data. Now, we use an example to illustrate this method. Let Π_n be a finite state Markov chain with the stationary distribution $(\pi_1, \pi_2, \cdots \pi_M)$. Let $A_j = \{d_1^{(A_j)}, d_2^{(A_j)}, \cdots d_N^{(A_j)}\}$, $1 \leq j \leq 3$, be three disjoint sets of data items. Let $DS_1 = A_1 \cup A_3$, $DS_2 = A_2 \cup A_3$, i.e., A_3 is the set of overlapped data items of flow 1 and flow 2. Let $q_i^{(A_j,m)} \triangleq \mathbb{P}[R_0 = d_i^{(A_j)} | R_0 \in A_j, \Pi_0 = m] = c_{j,m}/i^{\alpha_{j,m}}$. Define $r_1^{(A_3)} = \mathbb{P}[R_0 \in A_3 | I_0 = 1]$. Similarly, we can define $r_1^{(A_1)}, r_2^{(A_2)}, r_2^{(A_3)}$ with $r_1^{(A_1)} + r_1^{(A_3)} = 1$, $r_2^{(A_2)} + r_2^{(A_3)} = 1$. The miss ratio of flow k, $k = 1, 2$, is

$$\begin{aligned} \mathbb{P}[C_0 > x | I_0 = k] =& r_k^{(A_k)} \mathbb{P}[C_0 > x | R_0 \in A_k] \\ &+ r_k^{(A_3)} \mathbb{P}[C_0 > x | R_0 \in A_3], \end{aligned} \quad (19)$$

where $\mathbb{P}[C_0 > x | R_0 \in A_j]$, $1 \leq j \leq 3$, can be directly obtained from Corollary 2 if we view A_1, A_2, A_3 as the data sets of three flows without overlapped data items.

B. The Characteristic Time Approximation

In this section, we extend the characteristic time approximation from a single flow of independent requests to multiple competing flows with dependent requests. Define

$$\mathbb{P}_{CT}\left[C_0 > x | I_0 = k\right] = \sum_{i=1}^{N_k} q_i^{(k)} e^{-p_i^{(k)} \bar{m}^{\leftarrow}(x)}, \quad (20)$$

where $\bar{m}^{\leftarrow}(x)$ is the inverse of $\bar{m}(x)$ defined in (11). Note that, by letting $M = 1$, $q_i^{(k)} = q_i^{(1,k)}$, $p_i^{(k)} = p_i^{(1,k)}$ for all i, equation (20) degrades to the characteristic time approximation for independent requests presented in [37], [16]. Moreover, we verify the characteristic time approximation (20) by the following theorem.

Theorem 2. *Under the assumptions in Theorem 1, we have, as* $x \to \infty$*,*

$$\mathbb{P}\left[C_0 > x | I_0 = k\right] \sim \mathbb{P}_{CT}\left[C_0 > x | I_0 = k\right]. \quad (21)$$

The proof is presented in Section VI.

IV. Experiments

In this section, we conduct simulation experiments using C++ to verify the theoretical results. Experiment 1 shows that our asymptotic results are more accurate than the existing ones for a flow of correlated requests following a Zipf's distribution when α is close to 1 or the support region is relatively small. Experiment 2 validates Theorem 1 using multiple flows of requests with regularly varying popularity distributions that are modulated by a Markov chain. Experiment 3 investigates 2 flows of dependent requests with overlapped data items.

Experiment 1. We conduct two simulations for a single flow of correlated requests to show that Corollay 1 (labeled as "theoretical 1") is more accurate than Theorem 2 of [31] (labeled as "theoretical 2"). Assume all data items are unit-sized. Let Π_n be a two state ($\{1,2\}$) Markov chain with the transition matrix $(0.8, 0.2; 0.6, 0.4)$. Let $q_i^{(1,1)} = c_1/i^{\alpha}$, $q_i^{(1,2)} = c_2 \exp(-\xi i)$, $1 \le i \le N$. Since exponential distributions decay much faster than Zipf's distributions, we have $q_i^{(1)} \sim \pi_{1,1} q_i^{(1,1)}$ as $i \to \infty$. The first simulation studies Zipf's distributions with $\alpha \approx 1$. First, we set $\alpha = 1.1, \xi = 0.6, N = 10^4$, and plot the empirical miss ratios as well as the theoretical miss ratios by Corollary 1 and [31] in Fig. 2. Then, we choose $\alpha = 0.8 < 1$, which is not covered by [31], and repeat the simulation. The empirical miss ratios are compared with theoretical ones calculated by Corollary 1. It can be observed that theoretical results by Corollary 1 match very well with empirical ones for both $\alpha > 1$ and $0 < \alpha < 1$. In contrast, [31] cannot provide accurate estimations when $\alpha = 1.1$. In the second simulation, we study a relatively small support region by setting $\alpha = 2.0, \xi = 0.6, N = 10^3$. Corollary 1 also gives more accurate approximations than [31]. For independent requests, [25] proposes an accurate approximation even when $\alpha \approx 1$, but it does not cover correlated requests.

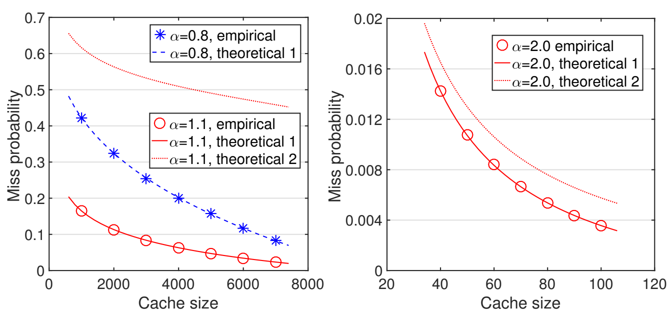

Fig. 2. A single flow following Zipf's law

Experiment 2. To validate Theorem 1, we consider 4 flows of dependent requests beyond Zipf's distribution. Assume these flows have unit-sized data without overlap. Let Π_n be a two state Markov chain with the transition matrix $(0.8, 0.2; 0.6, 0.4)$. Let $q_i^{(k,1)} = c_{k,1} \log(i)/i^{\alpha_{k,1}}$, $q_i^{(k,2)} = c_{k,2} \exp(-\xi_{k,2} i)$, $1 \le i \le 10^6$, $1 \le k \le 4$. Since exponential distributions decay much faster, we have $q_i^{(k)} \sim \pi_{k,1} q_i^{(k,1)}$ as $i \to \infty$. Let $(\alpha_{1,1}, \alpha_{2,1} \alpha_{3,1} \alpha_{4,1}) = (1.6, 1.8, 2.0, 2.4)$, $(\nu_{1,1}, \nu_{2,1}, \nu_{3,1}, \nu_{4,1}) = (0.3, 0.3, 0.2, 0.2)$, $(\nu_{1,2}, \nu_{2,2}, \nu_{3,2}, \nu_{4,1}) = (0.4, 0.3, 0.2, 0.1)$ and $\xi_{k,2} = 0.6$, $1 \le k \le 4$. For each flow, we compare the empirical miss ratios with theoretical ones calculated by Theorem 1. Note that $\Psi_k(m^{\leftarrow}(x))$ can be computed by numerical methods (e.g., binary search) based on (3) and (7). We plot the theoretical miss ratios and the empirical ones in Fig. 3. The perfect match validates Theorem 1.

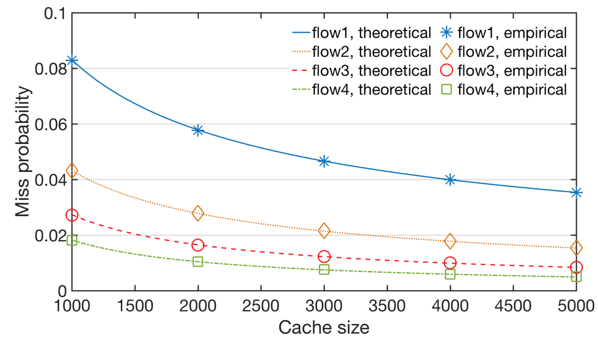

Fig. 3. Multiple dependent flows beyond Zipf's law

Experiment 3. Consider 2 flows of dependent requests with overlapped data that have been described in section III-A. We also use the notations introduced therein. Let Π_n be a two state Markov chain with transition matrix $(0.5, 0.5; 0.25, 0.75)$. Let $N_j = 10^4$, $\alpha_{j,1} = 0.8$, $\alpha_{j,2} = 2$, $1 \le j \le 3$, and $\nu_{1,1} = 0.2$, $\nu_{2,1} = 0.8$, $\nu_{1,2} = \nu_{2,2} = 0.5$, $r_1^{(A_3)} = r_2^{(A_3)} = 0.2$ and $r_1^{(A_1)} = r_2^{(A_2)} = 0.8$. In Fig. 4, we compare the theoretical miss ratios computed by (19) and the empirical ones obtained from simulations. It can be observed that the theoretical results match very well with the empirical ones.

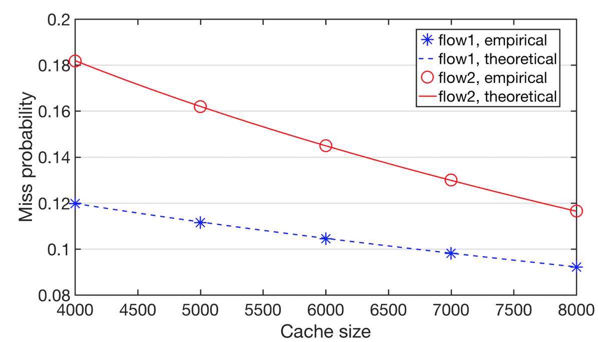

Fig. 4. Two dependent flows with overlapped data

V. Conclusion

We develop a unified framework to analyze the asymptotic miss ratios for multiple flows of dependent data item requests that share a common LRU cache. The analysis jointly considers time correlations, data popularities, varying item sizes, different request rates and even overlapped data items. The result is derived for a broad class of (truncated) heavy-tailed popularity distributions, e.g., regularly varying and heavy-tailed Weibull distributions. Existing asymptotic results often treat a Zipf's distribution of $\alpha > 1$ and $0 < \alpha < 1$ separately; we provide a uniform result. As a result, it is more accurate when the index α is around 1 or the support region is relatively small. Extensive simulation experiments validate our theoretical results. The high accuracy, low computation cost and applicability for a large class of popularity distributions make our analytical results useful in understanding LRU caching systems.

VI. Proofs

In order to prove Theorem 1, we need to establish a lemma. Let $M(n)$ denote the total size of all the distinct data items that have been requested on time points $\{\tau_{-1}, \tau_{-2}, \cdots, \tau_{-n}\}$. Define $M^{\leftarrow}(x) = \min\{n : M(n) \geq x\}$ to be the inverse function of $M(n)$.

Lemma 2. *For $\epsilon(x) = \epsilon\delta(x)$ as in (8) and $\bar{s} = \sup_i s_i < \infty$, we obtain*

$$\begin{aligned}\mathbb{P}\left[M(m^{\leftarrow}(x)) \geq (1+\epsilon(x))x\right] \\ \leq e^{-(1-\epsilon)(\epsilon(x))^2 x/4\bar{s}} + o(m^{\leftarrow}(x)^{-\beta^*}).\end{aligned} \tag{22}$$

Proof. Define a Bernoulli random variable $X_i^{(n)}$ with $X_i^{(n)} = 1$ if item d_i° has been requested in $\mathcal{R}_n = \{R_{-1}, R_{-2}, \cdots, R_{-n}\}$ and $X_i^{(n)} = 0$ otherwise. Then, we have $M(n) = \sum_{i=1}^{\infty} s_i X_i^{(n)}$.

Now, conditional on the modulating process $\mathcal{J}_n = (\Pi_{-1}, \Pi_{-2}, \cdots, \Pi_{-n})$, we can count the number of requests in $\mathcal{R}_n$ that observe the modulating process in state $m, 1 \leq m \leq M$. For each state m of the modulating process, define a random variable $n_m = |\{j : \Pi_{-j} = m, 1 \leq j \leq n\}|$. Let $p_i^{\mathcal{J}_n}(n) = \mathbb{P}\left[X_i^{(n)} = 1 | \mathcal{J}_n\right]$ and $p_i(n) = \mathbb{P}[X_i^{(n)} = 1] = \mathbb{E}\left[p_i^{\mathcal{J}_n}(n)\right]$. Then, we have

$$p_i^{\mathcal{J}_n}(n) = 1 - \left(\prod_{m=1}^{M}\left(1 - p_i^{(\circ,m)}\right)^{n_m}\right). \tag{23}$$

Define $\mathcal{E}_\lambda^+ = \bigcap_{1\leq m\leq M}\{n_m < (1+\lambda)\pi_m n\}$. By Markov's inequality, for $\theta > 0$, we obtain, using independence of $X_i^{(n)}$'s conditional on $\mathcal{J}_n$, and $\mathbb{E}\left[e^{\theta s_i X_i^{(n)}} | \mathcal{J}_n\right] = p_i^{\mathcal{J}_n}(n)e^{\theta s_i} + 1 - p_i^{\mathcal{J}_n}(n) = p_i^{\mathcal{J}_n}(n)\left(e^{\theta s_i} - 1\right) + 1 \leq e^{p_i^{\mathcal{J}_n}(n)\left(e^{\theta s_i}-1\right)}$,

$$\begin{aligned}&\mathbb{P}\left[M(n) \geq (1+\epsilon(m(n)))m(n)\right] \\ &\leq \mathbb{P}\left[M(n) \geq (1+\epsilon(m(n)))m(n)|\mathcal{E}_\lambda^+\right] + \mathbb{P}\left[\left(\mathcal{E}_\lambda^+\right)^C\right] \\ &\leq \mathbb{E}\left[\Pi_{i=1}^N \mathbb{E}\left[e^{\theta s_i X_i^{(n)}}\middle|\mathcal{J}_n\right]\middle|\mathcal{E}_\lambda^+\right] \Big/ e^{(1+\epsilon(m(n)))\theta m(n)} \\ &\quad + \mathbb{P}\left[\left(\mathcal{E}_\lambda^+\right)^C\right] \\ &\leq \mathbb{E}\left[e^{\sum_{i=1}^N p_i^{\mathcal{J}_n}(n)\left(e^{\theta s_i}-1\right)}\middle|\mathcal{E}_\lambda^+\right] \Big/ e^{\theta(1+\epsilon(m(n)))m(n)} \\ &\quad + \mathbb{P}\left[\left(\mathcal{E}_\lambda^+\right)^C\right].\end{aligned}$$

Using $e^x - 1 \leq (1+\xi)x, 0 < x < 2\xi/e^{\xi}, \xi > 0$, we obtain $e^{\theta s_i} - 1 \leq (1+\epsilon(m(n))/2)\theta s_i$ for $\theta = \epsilon(m(n))/(2\bar{s})$.

As a result, we have,

$$\begin{aligned}&\mathbb{P}\left[M(n) \geq (1+\epsilon(m(n)))m(n)\right] \\ &\leq \mathbb{E}\left[e^{(1+\epsilon(m(n))/2)\theta\sum_{i=1}^N p_i^{\mathcal{J}_n}(n)s_i}\middle|\mathcal{E}_\lambda^+\right] \Big/ e^{\theta(1+\epsilon(m(n)))m(n)} \\ &\quad + \mathbb{P}\left[\left(\mathcal{E}_\lambda^+\right)^C\right] \\ &\leq \exp\Bigg(-\frac{\epsilon(m(n))^2}{4\bar{s}}m(n) \\ &\qquad + \left(\frac{\epsilon(m(n))}{2\bar{s}} + \frac{\epsilon(m(n))^2}{4\bar{s}}\right)(m((1+\lambda)n) - m(n))\Bigg) \\ &\quad + \mathbb{P}\left[\left(\mathcal{E}_\lambda^+\right)^C\right].\end{aligned} \tag{24}$$

Recalling (9), we can choose λ small enough such that, for large n,

$$\frac{m((1+\lambda)n)}{m(n)} \leq 1 + \frac{\epsilon^2\delta(m(n))}{3}.$$

Therefore, for sufficiently large n, equation (24) can be further upper bounded by

$$\begin{aligned}&\mathbb{P}\left[M(n) \geq (1+\epsilon(m(n)))m(n)\right] \\ &\qquad \leq \exp\left(-(1-\epsilon)\frac{\epsilon(m(n))^2}{4\bar{s}}m(n)\right) + o(n^{-\beta^*}),\end{aligned}$$

implying (22) by replacing n with $m^{\leftarrow}(x)$. □

Proof of Theorem 1. First, we derive a representation for the miss probability $\mathbb{P}[C_0 > x | I_0 = k]$. We find the last request before R_0 such that the requested date item is the same with R_0, and denote it by $R_{-\sigma}$. Let $M(n)$ denote the total size of all the distinct data items that have been requested on the time interval $[\tau_{-n}, \tau_{-1}]$. Define the inverse function of $M(n)$ to be $M^{\leftarrow}(x) = \min\{n : M(n) \geq x\}$. We claim that

$$\{C_0 > x\} = \{\sigma > M^{\leftarrow}(x)\}. \tag{25}$$

If the event $\{\sigma > M^{\leftarrow}(x)\}$ occurs, then we have $M(\sigma) > x$, which means, under MTF, the total size of distinct items listed in front of the requested item R_0 is larger than the cache size x. This implies $\{\sigma > M^{\leftarrow}(x)\} \subseteq \{C_0 > x\}$. If $\{C_0 > x\}$ occurs, then there must be enough distinct data items that have been requested on $(\tau_{-\sigma}, 0)$ so that the data item $R_{-\sigma}$ is moved out of the cache. Therefore, we obtain $\{C_0 > x\} \subseteq \{\sigma > M^{\leftarrow}(x)\}$, which proves (25) and implies

$$\mathbb{P}[C_0 > x | I_0 = k] = \mathbb{P}[\sigma > M^{\leftarrow}(x) | I_0 = k]. \tag{26}$$

Given the cache size x, $M^{\leftarrow}(x)$ is a random variable. We estimate $M^{\leftarrow}(x)$ by a deterministic function $m^{\leftarrow}(x)$, to derive the explicit form of the asymptotic miss ratio.

From (26), we take two steps to calculate the miss ratio. The first step is to show, for $N_k = \zeta n$, $\zeta > 0$, as $n \to \infty$,

$$\mathbb{P}[\sigma > n | I_0 = k] \sim \beta_k \Gamma\left(\beta_k, n p_{N_k}^{(k)}\right)\Psi_k(n). \tag{27}$$

The second step is to relate $M^{\leftarrow}(x)$ to $m^{\leftarrow}(x)$ as $x \to \infty$.

Assume that $\Psi_k(x)$ is eventually absolutely continuous and strictly decreasing, by Proposition 1.5.8 and Proposition 1.5.10 of [34], we can construct such a function, for $x > x_0$,

$$\Psi_k^*(x) = \begin{cases} -\beta_k \int_{x_0}^x s^{-\beta_k - 1} l(s) ds & \text{if } \beta_k < 0, \\ \beta_k \int_x^\infty s^{-\beta_k - 1} l(s) ds & \text{if } \beta_k > 0. \end{cases} \tag{28}$$

For x_0 large enough, we have, as $y \to \infty$,

$$\begin{aligned} \sum_{i=y}^{N_k} q_i^{(k)} &\sim \Psi_k\left(\left(p_y^{(k)}\right)^{-1}\right) + \Theta_k(N_k) \\ &\sim \Psi_k^*\left(\left(p_y^{(k)}\right)^{-1}\right) + \Theta_k(N_k). \end{aligned} \tag{29}$$

Therefore, there exists x_0 such that, for $x > x_0$, $\Psi_k(x)$ is decreasing and has an inverse function. According to (3), for $\forall \epsilon \in (0,1)$, there exists i_ϵ such that, for $i \geq i_\epsilon$,

$$\begin{aligned} (1-\epsilon)\left(\sum_{j=i}^{N_k} q_j^{(k)}\right) &\leq \Psi_k\left(\left(p_i^{(k)}\right)^{-1}\right) + \Theta_k(N_k) \\ &\leq (1+\epsilon)\left(\sum_{j=i}^{N_k} q_j^{(k)}\right). \end{aligned} \tag{30}$$

For i_ϵ large enough such that $1/p_{i_\epsilon}^{(k)} > x_0$, we have, for $i \geq i_\epsilon$,

$$\begin{aligned} \Psi_k^{\leftarrow}\left((1-\epsilon)\left(\sum_{j=i}^{\infty} q_j^{(k)}\right) - \Theta_k(N_k)\right) &\geq \left(p_i^{(k)}\right)^{-1} \\ \geq \Psi_k^{\leftarrow}\left((1+\epsilon)\left(\sum_{j=i}^{\infty} q_j^{(k)}\right) - \Theta_k(N_k)\right). \end{aligned} \tag{31}$$

Step 1: Conditional on the event $\{R_0 = d_i^{(k)}, I_0 = k\}$ and the modulating process $\mathcal{J}_n = (\Pi_{-1}, \Pi_{-2}, \cdots, \Pi_{-n})$, the requests $R_{-1}, R_{-2}, \cdots R_{-n}$ are independent, which implies

$$\begin{aligned} &\mathbb{P}\left[\sigma > n \middle| R_0 = d_i^{(k)}, I_0 = k, \mathcal{J}_n\right] \\ &= \prod_{j=1}^n \mathbb{P}\left[R_{-j} \neq d_i^{(k)} \middle| R_0 = d_i^{(k)}, I_0 = k, \mathcal{J}_n\right] \\ &= \prod_{j=1}^n \left(1 - p_i^{(k,\Pi_{-j})}\right) = \prod_{j=1}^M \left(1 - p_i^{(k,j)}\right)^{n_j}, \end{aligned}$$

where $n_j = \sum_{i=1}^n \mathbf{1}\{\Pi_{-i} = j\}$, $1 \leq j \leq M$. Thus, unconditional on R_0, we obtain, recalling (1) and (2),

$$\begin{aligned} &\mathbb{P}[\sigma > n | I_0 = k] \\ &= \mathbb{E}\left[\sum_{m=1}^M \pi_{k,m} \mathbb{P}[\sigma > n | I_0 = k, \Pi_0 = m, \mathcal{J}_n]\right] \\ &= \mathbb{E}\left[\sum_{m=1}^M \pi_{k,m} \sum_{i=1}^{N_k} q_i^{(k,m)} \prod_{j=1}^M \left(1 - p_i^{(k,j)}\right)^{n_j}\right] \\ &= \mathbb{E}\left[\sum_{i=1}^{N_k} q_i^{(k)} \prod_{j=1}^M \left(1 - p_i^{(k,j)}\right)^{n_j}\right]. \end{aligned} \tag{32}$$

First, we derive the upper bound of (32),

$$\begin{aligned} &\mathbb{P}[\sigma > n | I_0 = k] \\ &\leq \mathbb{E}\left[\sum_{i=1}^{N_k} q_i^{(k)} \prod_{j=1}^M \left(1 - p_i^{(k,j)}\right)^{n_j} \mathbf{1}\{n_j \geq (1-\epsilon)\pi_j n\}\right] \\ &\quad + \mathbb{P}\left[\min_{1 \leq j \leq M} n_j / \pi_j < 1 - \epsilon\right] \\ &\leq (1+\epsilon)^M \sum_{i=1}^{N_k} q_i^{(k)} \exp\left(-\sum_{j=1}^M (1-\epsilon) n \pi_j p_i^{(k,j)}\right) \\ &\quad + \mathbb{P}\left[\min_{1 \leq j \leq M} n_j / \pi_j < 1 - \epsilon\right] \\ &\leq (1+\epsilon)^M \exp\left(-(1-\epsilon) n p_{i_\epsilon}^{(k)}\right) \\ &\quad + (1+\epsilon)^M \sum_{i=i_\epsilon+1}^{N_k} q_i^{(k)} \exp\left(-(1-\epsilon) n p_i^{(k)}\right) \\ &\quad + \mathbb{P}\left[\min_{1 \leq j \leq M} n_j / \pi_j < 1 - \epsilon\right] \triangleq I_1 + I_2 + I_3. \end{aligned} \tag{33}$$

For $\epsilon_2 \in (p_{N_k}^{(k)}, p_{i_{\epsilon_1}}^{(k)})$, large integer n and any integer l with $\lfloor \log n p_{N_k}^{(k)} \rfloor \triangleq n(p_{N_k}^{(k)}) \leq l \leq \lfloor \log n \epsilon_2 \rfloor \triangleq n(\epsilon_2)$, we can find i_l such that $p_{i_l+1}^{(k)} \leq e^l / n \leq p_{i_l}^{(k)} \leq \epsilon_2$. For an integer m with $n(p_{N_k}^{(k)}) < m < n(\epsilon_2)$, we have $i_m > i_{n(\epsilon_2)} > i_{\epsilon_1}$, and

$$\begin{aligned} &I_2 / (1+\epsilon)^M \\ &= \left(\sum_{i=i_{\epsilon_1}+1}^{i_{n(\epsilon_2)}-1} + \sum_{i=i_{n(\epsilon_2)}}^{i_m} + \sum_{i=i_m+1}^{N_k}\right) q_i^{(k)} e^{-(1-\epsilon) n p_i^{(k)}} \\ &\leq e^{-(1-\epsilon) n \epsilon_2} + \sum_{l=m}^{n(\epsilon_2)} \sum_{j=i_{l+1}+1}^{i_l} q_j^{(k)} e^{-(1-\epsilon) e^l} \\ &\qquad + \sum_{j=i_m+1}^{N_k} q_j^{(k)} e^{-(1-\epsilon) n p_j^{(k)}} \\ &\triangleq I_{21} + I_{22} + I_{23}. \end{aligned}$$

Using (31), we have, for $Q_j = \sum_{i=j}^\infty q_i^{(k)}$, $\Delta Q_j = Q_j - Q_{j+1}$,

$$I_{23} = \sum_{j=i_m+1}^{N_k} \exp\left(-\frac{(1-\epsilon) n}{\Psi_k^{\leftarrow}\left((1-\epsilon) Q_j - \Theta_k(N_k)\right)}\right) \Delta Q_j.$$

Since $e^{-1/\Psi_k^{\leftarrow}(u)}$ is decreasing with u, we have, for $\forall u \in (Q_{j+1}, Q_j)$, $\exp(-(1-\epsilon)n/(\Psi_k^{\leftarrow}((1-\epsilon)u - \Theta_k(N_k)))) \geq \exp(-(1-\epsilon)n/(\Psi_k^{\leftarrow}((1-\epsilon)Q_j - \Theta_k(N_k))))$, which implies

$$\begin{aligned} I_{23} &\leq \int_0^{Q_{i_m}} \exp\left(-\frac{(1-\epsilon) n}{\Psi_k^{\leftarrow}\left((1-\epsilon) u - \Theta_k(N_k)\right)}\right) du \\ &\leq \int_{z(\epsilon)}^{e^m} e^{-z} d\left(\Psi_k\left(\frac{(1-\epsilon) n}{z}\right)\right), \end{aligned}$$

where $z(\epsilon) = (1-\epsilon)n / \Psi_k^{\leftarrow}((1-\epsilon) Q_{N_k} - \Theta_k(N_k))$. By Theorem 1.2.1 of [34], (28) and (29), we obtain,

$$I_{23} / \Psi_k(n) \lesssim (1-\epsilon)^{-\beta_k} \int_{z(\epsilon)}^{e^m} \beta_k e^{-z} z^{\beta_k - 1} dz. \tag{34}$$

Using the same approach, I_{22} can be bounded by

$$I_{22}/\Psi_k(n) \lesssim \sum_{k=m}^{\infty}(1+\epsilon)e^{-e^k}\left(e^{k+1}\right)^{\beta_k} < \infty. \tag{35}$$

Recalling the assumption $\mathbb{P}\left[|\sum_{i=1}^n \mathbf{1}\{\Pi_i = m\} - \pi_m n| > \epsilon n\right] = o(n^{-\beta^*})$, $\epsilon \in (0,1)$, $1 \le m \le M$ and $\Psi_k(n) \gtrsim n^{-\beta^*} l(n)$, we have $I_1 = o(\Psi_k(n))$, $I_3 = o(\Psi_k(n))$ in (33). Combining (34) and (35), passing $\epsilon \to 0$, $n \to \infty$, and $m \to \infty$, we obtain

$$\begin{aligned}&\mathbb{P}[\sigma > n | I_0 = k]/\Psi_k(n) \\ &\quad\lesssim \int_{np^{(k)}_{N_k}}^{\infty} \beta_k e^{-z} z^{\beta_k - 1} dz = \beta_k \Gamma(\beta_k, np^{(k)}_{N_k}).\end{aligned} \tag{36}$$

Next, we derive the lower bound. According to the weak law of large number, for $\forall\ \epsilon \in (0,1)$ and $1 \le j \le M$, we can choose n large enough such that $\mathbb{P}[n_j \le (1+\epsilon)\pi_j n] \ge 1-\epsilon$. $\mathbb{P}[\sigma > n | I_0 = k]$ can be bounded as

$$\begin{aligned}&\mathbb{P}[\sigma > n | I_0 = k] \\ &\ge \mathbb{E}\left[\sum_{i=1}^{N_k} q_i^{(k)} \prod_{j=1}^{M}\left(1 - p_i^{(k,j)}\right)^{n_j} \mathbf{1}\{n_j \le (1+\epsilon)\pi_j n\}\right] \\ &\ge (1-\epsilon)^M \sum_{i=1}^{N_k} q_i^{(k)} \prod_{j=1}^{M}\left(1 - p_i^{(k,j)}\right)^{(1+\epsilon)\pi_j n}.\end{aligned}$$

Since $1-x \ge e^{-(1+\epsilon)x}$ for sufficiently small x, we can choose i_ϵ large enough such that $1 - p_i^{(k,j)} \ge e^{-(1+\epsilon)p_i^{(k,j)}}$ for $i \ge i_\epsilon$, which yields

$$\begin{aligned}&\mathbb{P}[\sigma > n | I_0 = k] \\ &\ge (1-\epsilon)^M \sum_{i=i_\epsilon}^{N_k} q_i^{(k)} \prod_{j=1}^{M} \exp\left(-(1+\epsilon) n \pi_j p_i^{(k,j)}\right) \\ &= (1-\epsilon)^M \sum_{i=i_\epsilon}^{N_k} q_i^{(k)} \exp\left(-(1+\epsilon) n p_i^{(k)}\right).\end{aligned}$$

Using (4), (30), we obtain, for $Q_i = \sum_{j=i}^{N_k} q_j^{(k)}$,

$$\begin{aligned}&\mathbb{P}[\sigma > n | I_0 = k]/(1-\epsilon)^{M+1} \\ &\ge \sum_{i=i_\epsilon}^{N_k} q_{i-1}^{(k)} \exp\left(-(1+\epsilon) n p_i^{(k)}\right) \\ &\ge \sum_{i=i_\epsilon}^{N_k} (Q_{i-1} - Q_i) \exp\left(-\frac{(1+\epsilon)n}{\Psi_k^{\leftarrow}((1+\epsilon)Q_i - \Theta_k(N_k))}\right) \\ &\ge \int_{Q_{N_k}}^{Q_{i_\epsilon - 1}} \exp\left(-\frac{(1+\epsilon)n}{\Psi_k^{\leftarrow}((1+\epsilon)u - \Theta_k(N_k))}\right) du,\end{aligned}$$

where the last inequality holds since $\exp(-1/\Psi_k^{\leftarrow}(x))$ is decreasing with x. Choose $i_n > i_\epsilon$ such that $\Psi_k^{\leftarrow}((1+\epsilon)Q_{i_n} - \Theta_k(N_k)) = n/W$. Using a similar approach as in deriving (34), we obtain, for $z(\epsilon) = (1+\epsilon)n/\Psi_k^{\leftarrow}((1+\epsilon)Q_{N_k} - \Theta_k(N_k))$,

$$\mathbb{P}[\sigma > n | I_0 = k]/\Psi_k(n) \gtrsim (1-\epsilon)^M \int_{z(\epsilon)}^{W} \beta_k e^{-z} z^{\beta_k - 1} dz,$$

implying, by passing $\epsilon \to 0$, $n \to \infty$ and $W \to \infty$,

$$\begin{aligned}&\mathbb{P}[\sigma > n | I_0 = k]/\Psi_k(n) \\ &\quad\gtrsim \int_{np^{(k)}_{N_k}}^{\infty} \beta_k e^{-z} z^{\beta_k - 1} dz = \beta_k \Gamma\left(\beta_k, np^{(k)}_{N_k}\right).\end{aligned} \tag{37}$$

Combining (36) and (37) yields (27).

Step 2: Using (22), we have, for $x_1 = m^{\leftarrow}(x/(1+\epsilon(x)))$

$$\begin{aligned}&\mathbb{P}[M^{\leftarrow}(x) < x_1] \le \mathbb{P}\left[M(m^{\leftarrow}(x/(1+\epsilon(x)))) \ge x\right] \\ &\quad= \mathbb{P}\left[M\left(m^{\leftarrow}\left(\frac{x}{1+\epsilon(x)}\right)\right) \ge \left(\frac{(1+\epsilon(x))x}{1+\epsilon(x)}\right)\right].\end{aligned} \tag{38}$$

Recalling (10) and $\delta(x) \le 1$, we have, for $x > x_0$, $\epsilon(x) \ge h_1 \epsilon\left(x/(1+\epsilon(x))\right)$, which implies, using (22) and (38),

$$\begin{aligned}&\mathbb{P}\left[M(x_1) \ge \left(1 + h_1 \epsilon\left(\frac{x}{1+\epsilon(x)}\right)\right)\left(\frac{x}{1+\epsilon(x)}\right)\right] \\ &\quad\le \exp\left(-(1-\epsilon)(h_1\epsilon(x/(1+\epsilon(x))))^2 x/\left(4\bar{s}(1+\epsilon(x))\right)\right) \\ &\qquad + o((m^{\leftarrow}(x/(1+\epsilon(x))))^{-\beta^*}) \\ &\quad\le \exp\left(-\frac{{h_1}^2}{{h_2}^2}\frac{(1-\epsilon)\epsilon(x)^2 x}{4(1+\epsilon)\bar{s}}\right) \\ &\qquad + o((m^{\leftarrow}(x/(1+\epsilon(x))))^{-\beta^*}).\end{aligned} \tag{39}$$

Combining (22), (26), (27) and (39) yields

$$\begin{aligned}&\mathbb{P}[C_0 > x | I_0 = k] \le \mathbb{P}[\sigma > M^{\leftarrow}(x), M^{\leftarrow}(x) \ge x_1 | I_0 = k] \\ &\qquad + \mathbb{P}[M^{\leftarrow}(x) < x_1 | I_0 = k] \\ &\le \mathbb{P}[\sigma > m^{\leftarrow}(x/(1+\epsilon(x)))] + \mathbb{P}[M^{\leftarrow}(x) < x_1] \\ &\lesssim \beta_k \Gamma\left(\beta_k, m^{\leftarrow}(x) p^{(k)}_{N_k}\right) \Psi_k\left(m^{\leftarrow}(x/(1+\epsilon(x)))\right) \\ &\qquad + \exp\left(-\frac{{h_1}^2}{{h_2}^2}\frac{(1-\epsilon)\epsilon(x)^2 x}{4\bar{s}(1+\epsilon)}\right) \\ &\qquad + o((m^{\leftarrow}(x/(1+\epsilon(x))))^{-\beta^*}).\end{aligned} \tag{40}$$

Using $\overline{\lim}_{x\to\infty} \log\left(m^{\leftarrow}(x)\right)/\left(\delta^2(x)x\right) = 0$, (5), (8), and passing $\epsilon \to 0$, we obtain,

$$\begin{aligned}\mathbb{P}[C_0 > x | I_0 = k] &\lesssim \beta_k \Gamma\left(\beta_k, m^{\leftarrow}(x) p^{(k)}_{N_k}\right) \Psi_k\left(m^{\leftarrow}(x)\right) \\ &\quad + o\left(\Psi_k\left(m^{\leftarrow}(x)\right)\right).\end{aligned} \tag{41}$$

Similarly, for $x_2 = m^{\leftarrow}(x/(1-\epsilon(x)))$, we have

$$\begin{aligned}\mathbb{P}[C_0 > x | I_0 = k] &\ge \mathbb{P}[\sigma > M^{\leftarrow}(x), M^{\leftarrow}(x) \le x_2 | I_0 = k] \\ &\quad - \mathbb{P}[M^{\leftarrow}(x) > x_2 | I_0 = k],\end{aligned}$$

which implies, by a similar approach as in deriving (41),

$$\begin{aligned}\mathbb{P}[C_0 > x | I_0 = k] &\gtrsim \beta_k \Gamma\left(\beta_k, m^{\leftarrow}(x) p^{(k)}_{N_k}\right) \Psi_k\left(m^{\leftarrow}(x)\right) \\ &\quad - o\left(\Psi_k\left(m^{\leftarrow}(x)\right)\right).\end{aligned} \tag{42}$$

Moreover, for $\forall\ \zeta > 0$, $N_k = \zeta n$, we have $0 < np^{(k)}_{N_k} < 2/\zeta$, which indicates

$$\liminf_{n\to\infty} \beta_k \Gamma(\beta_k, np^{(k)}_{N_k}) > \beta_k \Gamma\left(\beta_k, 2/\zeta\right) > 0.$$

Therefore, combining (41) and (42) finishes the proof. □

Proof of Lemma 1. Using $1-x \le e^{-x}$, we have

$$\bar{m}(x) = \sum_{i=1}^{N} s_i \left(1 - \prod_{m=1}^{M} \exp\left(-\pi_m p_i^{(\circ,m)} x\right)\right) \le m(x).$$

Moreover, for $\epsilon \in (0,1)$, there exists x_ϵ such that, for $0 \le x \le x_\epsilon$, $1-x \ge e^{-(1+\epsilon)x}$. Choosing i_ϵ such that $p_i^{(\circ,m)} \le x_0$ for any $i > i_\epsilon$, $1 \le m \le M$, we have

$$\begin{aligned}\bar{m}(x) &= \sum_{i=1}^{i_\epsilon} s_i \left(1 - \exp\left(-\sum_{m=1}^{M} \pi_m p_i^{(\circ,m)} x\right)\right) \\ &\quad + \sum_{i=i_\epsilon+1}^{N} s_i \left(1 - \exp\left(-\sum_{m=1}^{M} \pi_m p_i^{(\circ,m)} x\right)\right) \\ &\ge i_\epsilon \bar{s} + \sum_{i=i_\epsilon+1}^{N} s_i \left(1 - \prod_{m=1}^{M} \left(1 - p_i^{(\circ,m)}\right)^{\pi_m x}\right) \\ &\ge i_\epsilon \bar{s} + m((1+\epsilon)x),\end{aligned}$$

which implies, by (9), $\bar{m}(x) \gtrsim m(x)$ as $x \to \infty$. Combining the upper bound and the lower bound, we obtain, as $x \to \infty$, $m(x) \sim \bar{m}(x)$. Moreover, since there are no overlapped data items, we have $m(x) \sim \bar{m}(x) = \sum_{k=1}^{K} \bar{m}^{(k)}(x)$. □

Proof of Theorem 2. Using a similar approach as in the proof of Theorem 1 that derives the lower bound and upper bound of $\mathbb{P}[\sigma > n | I_0 = k]$, we have, as $x \to \infty$,

$$\mathbb{P}_{CT}\left[C_0 > x | I_0 = k\right] \sim \beta_k \Gamma\left(\beta_k, \bar{m}^{\leftarrow}(x) p_{N_k}^{(k)}\right) \Psi_k(\bar{m}^{\leftarrow}(x)).$$

Combining this with (5) and applying Theorem 1, Lemma 1, we can prove Theorem 2. □

References

[1] R. Nishtala, H. Fugal, S. Grimm, M. Kwiatkowski, H. Lee, H. C. Li, R. McElroy, M. Paleczny, D. Peek, P. Saab, D. Stafford, T. Tung, and V. Venkataramani, "Scaling memcache at Facebook," in *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. Lombard, IL: USENIX, 2013, pp. 385–398.

[2] "Voldemort," http://www.project-voldemort.com/voldemort/.

[3] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: Amazon's highly available key-value store," in *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles*, ser. SOSP '07. ACM, 2007, pp. 205–220.

[4] "Redis," http://redis.io/.

[5] B. Atikoglu, Y. Xu, E. Frachtenberg, S. Jiang, and M. Paleczny, "Workload analysis of a large-scale key-value store," *SIGMETRICS Performance Evaluation Review*, vol. 40, no. 1, pp. 53–64, Jun. 2012.

[6] Y. Xua, E. Frachtenbergb, and S. Jiang, "Building a high-performance key-value cache as an energy-efficient appliance," *Performance Evaluation*, vol. 79, pp. 24–37, September 2014.

[7] A. S. Tanenbaum, *Modern Operating Systems*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall Press, 2001.

[8] D. Starobinski and D. Tse, "Probabilistic methods for web caching," *Performance evaluation*, vol. 46, no. 2, pp. 125–137, 2001.

[9] M. Garetto, E. Leonardi, and V. Martina, "A unified approach to the performance analysis of caching systems," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 1, no. 3, pp. 12:1–12:28, May 2016.

[10] E. G. Coffman and P. J. Denning, *Operating systems theory*. Prentice-Hall Englewood Cliffs, NJ, 1973, vol. 973.

[11] "Spark," http://spark.apache.org/.

[12] O. I. Aven, L. B. Boguslavsky, and Y. A. Kogan, "Some results on distribution-free analysis of paging algorithms," *IEEE Trans. Computers*, vol. 25, no. 7, pp. 737–745, 1976.

[13] O. I. Aven, E. G. Coffman, and Y. A. Kogan, *Stochastic analysis of computer storage*. Springer Science & Business Media, 1987, vol. 38.

[14] O. Babaoglu and D. Ferrari, "Two-level replacement decisions in paging," *IEEE Transactions on Computers*, vol. 100, no. 32, 1983.

[15] E. Gelenbe, "A unified approach to the evaluation of a class of replacement algorithms," *IEEE Transactions on Computers*, vol. 100, no. 6, pp. 611–618, 1973.

[16] H. Che, Y. Tung, and Z. Wang, "Hierarchical web caching systems: Modeling, design and experimental results," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 7, pp. 1305–1314, Sep 2002.

[17] R. Hirade and T. Osogami, "Analysis of page replacement policies in the fluid limit," *Operations research*, vol. 58, no. 4-part-1, pp. 971–984, 2010.

[18] N. Tsukada, R. Hirade, and N. Miyoshi, "Fluid limit analysis of FIFO and RR caching for independent reference models," *Performance Evaluation*, vol. 69, no. 9, pp. 403–412, 2012.

[19] N. Gast and B. Van Houdt, "Transient and steady-state regime of a family of list-based cache replacement algorithms," *ACM SIGMETRICS Performance Evaluation Review*, vol. 43, no. 1, pp. 123–136, 2015.

[20] P. R. Jelenković, "Asymptotic approximation of the move-to-front search cost distribution and least-recently-used caching fault probabilities," *The Annals of Applied Probability*, no. 2, pp. 430–464, 1999.

[21] P. R. Jelenković, "Least-recently-used caching with Zipf's law requests," in *The Sixth INFORMS Telecommunications Conference*. Boca Raton, Florida, 2002.

[22] J. Roberts and N. Sbihi, "Exploring the memory-bandwidth tradeoff in an information-centric network," in *Teletraffic Congress (ITC), 2013 25th International*. IEEE, 2013, pp. 1–9.

[23] Y. Yang and J. Zhu, "Write skew and Zipf distribution: Evidence and implications," *ACM Transactions on Storage (TOS)*, vol. 12, no. 4, p. 21, 2016.

[24] C. Fricker, P. Robert, and J. Roberts, "A versatile and accurate approximation for LRU cache performance," in *Proceedings of the 24th International Teletraffic Congress*, ser. ITC '12, 2012, pp. 8:1–8:8.

[25] C. Berthet, "Approximation of LRU caches miss rate: Application to power-law popularities," *arXiv:1705.10738*, 2017.

[26] P. R. Jelenković and X. Kang, "LRU caching with moderately heavy request distributions," in *2007 Proceedings of the Fourth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*. SIAM, 2007, pp. 212–222.

[27] J. Tan, G. Quan, K. Ji, and N. Shroff, "On resource pooling and separation for LRU caching," in *Proceedings of the 2018 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*. ACM, 2018.

[28] E. G. Coffman Jr and P. Jelenković, "Performance of the move-to-front algorithm with Markov-modulated request sequences," *Operations Research Letters*, vol. 25, no. 3, pp. 109–118, 1999.

[29] R. P. Dobrow and J. A. Fill, "The move-to-front rule for self-organizing lists with Markov dependent requests," in *Discrete Probability and Algorithms*. New York, NY: Springer New York, 1995, pp. 57–80.

[30] E. R. Rodrigues, "The performance of the move-to-front scheme under some particular forms of Markov requests," *Journal of applied probability*, vol. 32, no. 4, pp. 1089–1102, 1995.

[31] P. R. Jelenković and A. Radovanović, "Least-recently-used caching with dependent requests," *Theoretical Computer Science*, vol. 326, no. 1-3, pp. 293–327, Oct. 2004.

[32] T. Sugimoto and N. Miyoshi, "On the asymptotics of fault probability in least-recently-used caching with Zipf-type request distribution," *Random Structures & Algorithms*, vol. 29, no. 3, pp. 296–323, 2006.

[33] T. Osogami, "A fluid limit for a cache algorithm with general request processes," *Advances in Applied Probability*, vol. 42, no. 3, pp. 816–833, 2010.

[34] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation*. Cambridge University Press, 1987.

[35] J. Fill, "An exact formula for the move-to-front rule for self-organizing lists," *Journal of Theoretical Probability*, vol. 9, no. 1, pp. 113–160, 1996.

[36] P. Flajolet, D. Gardy, and L. Thimonier, "Birthday paradox, coupon collectors, caching algorithms and self-organizing search," *Discrete Applied Mathematics*, vol. 39, no. 3, pp. 207–229, November 1992.

[37] R. Fagin, "Asymptotic miss ratios over independent references," *Journal of Computer and System Sciences*, vol. 14, no. 2, pp. 222–250, 1977.