The Output of a Cache under the Independent Reference Model – Where did the Locality of Reference Go?

Sarut Vanichpun and Armand M. Makowski

ECE & ISR/CSHCN
University of Maryland at College Park
`{sarut,armand}@isr.umd.edu`

June 16, 2004

- S. Vanichpun and A. M. Makowski, “Comparing strength of locality of reference – Popularity, majorization, and some folk theorems,” Infocom’04, Hong Kong (PRC), April 2004.
- A.M. Makowski and S. Vanichpun, “Comparing locality of reference – Popularity, majorization and some folk theorems for the miss rates and output of caches,” in *Performance Evaluation and Planning Methods for the Next Generation Internet*, A. Girard, B. Sansò and F.J. Vázquez-Abad (Editors), Kluwer Academic Press.

Demand-driven caching

- N cacheable documents
- Cache of size M with $M < N$
- User requests $\boldsymbol{R} \equiv \{R_t,\ t = 0, 1, \ldots\}$ form a sequence of $\{1, \ldots, N\}$-valued rvs with known statistics
- Cache sets $\{S_t,\ t = 0, 1, \ldots\}$ evolve according to

$$S_{t+1} = \begin{cases} S_t & \text{if } R_t \in S_t \\ S_t + R_t & \text{if } R_t \notin S_t, |S_t| < M \\ S_t - U_t + R_t & \text{if } R_t \notin S_t, |S_t| = M \end{cases}$$

where $U_t \in S_t, t = 0, 1, \ldots$ is the document to be evicted according to the replacement policy π in force.

The output of a cache

- When studying hierarchical caching, misses at one cache become the requests submitted to cache(s) at next level
- Let $\{\nu_k,\ k = 0, 1, \ldots\}$ be the time indices when a miss occurs, i.e.,

$$\nu_0 = 0; \quad \nu_{k+1} := \nu_k + \mu_{k+1}, \quad k = 0, 1, \ldots$$

  with

$$\mu_{k+1} := \inf \{\ell = 1, 2, \ldots : \ R_{\nu_k + \ell} \notin S_{\nu_k + \ell}\}$$

- The `output` process $\boldsymbol{R}^\star \equiv \{R_k^\star,\ k = 0, 1, \ldots\}$ of the cache is simply the stream of misses given by

$$R_k^\star := R_{\nu_k}, \quad k = 1, 2, \ldots$$

- The statistics of the output stream $\boldsymbol{R}^\star$ are determined by
  - replacement policy π
  - statistics of $\boldsymbol{R}$

  with

$$\boldsymbol{R}^\star_\pi = \boldsymbol{R}^\star_\pi(\boldsymbol{R})$$

- `Key problem:` Understand which properties of the stream of requests $\boldsymbol{R}$ matter

Locality of reference

- A key property of the stream of requests $\boldsymbol{R}$
  - "Bursts of references are made in the near future to objects referenced in the recent past"
- Close cousin to the notion of `burstiness` in traffic engineering
- Elusive despite numerous efforts
  - "Categorical data"
  - Additional dimensions (e.g., temporal/spatial LR)
  - Operational significance with respect to simple `transformations` (e.g., output) not well understood

Two contributors to LR

- Popularity vector $\boldsymbol{p} = (p(1), \ldots, p(N))$ of $\boldsymbol{R}$ is the pmf on $\{1, \ldots, N\}$ given by

  $$p(i) := \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbf{1}\left[R_\tau = i\right] \quad a.s.$$

  for all $i = 1, \ldots, N$

- Temporal correlations
  - Correlations $\mathrm{Cov}[R_s, R_t]$ will not do
  - Inter-object correlations
  - Working set, stack distance, inter-request times

Comparing strength of LR

- We express the fact

  stream $\boldsymbol{R}_1$ exhibits less LR than stream $\boldsymbol{R}_2$

  by writing

  $$\boldsymbol{R}_1 \leq_{LR} \boldsymbol{R}_2$$

- `Objective:` How to formalize the notion of LR to
  - capture its intuitive meaning
  - validate operational expectations associated with LR

A Folk Theorem

- “Good cache replacement strategies produce an output stream of requests exhibiting `less` locality of reference than the input stream of requests,” i.e.,

$$\boldsymbol{R}^{\star}_{\pi}(\boldsymbol{R}) \leq_{LR} \boldsymbol{R}$$

- In the context of multi-level caching, this reduction property is often perceived as one of the main reasons for why caching looses its effectiveness after some point in the hierarchy of caches

IRM

- The independent reference model IRM $(\boldsymbol{p})$
  - Benchmark model
  - Finding good models is hard
- Under the IRM with popularity pmf

$$\boldsymbol{p} = (p(1), \ldots, p(N)),$$

the successive requests $\boldsymbol{R}$ are `i.i.d.` $\{1, \ldots, N\}$-valued rvs, each distributed according to the pmf $\boldsymbol{p}$, i.e.,

$$\mathbf{P}\left[R_t = i\right] = p(i), \quad i = 1, \ldots, N;\ t = 0, 1, \ldots$$

with

$$p(i) = \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbf{1}\left[R_\tau = i\right] \quad a.s.$$

Comparing LR under IRM

- Case 1 – `Skewed` pmf $\boldsymbol{p}$

  $$\boldsymbol{p} = (1-\delta, \varepsilon, \ldots, \varepsilon)$$

  with

  $$\delta = (N-1)\varepsilon \ll 1$$

- Case 2 – `Uniform` pmf $\boldsymbol{u}$

  $$u(1) = \ldots = u(N) = \frac{1}{N}$$

- Although `no` temporal correlations, it is reasonable to think

  $$\mathrm{IRM}(\boldsymbol{u}) \leq_{LR} \mathrm{IRM}(\boldsymbol{p})$$

  - $\boldsymbol{u}$ is “more balanced” than $\boldsymbol{p}$

Majorization

- For $\boldsymbol{x}$ and $\boldsymbol{y}$ elements of $\mathbb{R}^N$, let

$$x_{[1]} \geq x_{[2]} \geq \ldots \geq x_{[N]} \quad \text{and} \quad y_{[1]} \geq y_{[2]} \geq \ldots \geq y_{[N]}$$

denote the components of $\boldsymbol{x}$ and $\boldsymbol{y}$ arranged in `decreasing order`.

We write

$$\boldsymbol{x} \prec \boldsymbol{y}$$

whenever

$$\sum_{i=1}^{n} x_{[i]} \leq \sum_{i=1}^{n} y_{[i]}, \quad n = 1, 2, \ldots, N-1$$

and

$$\sum_{i=1}^{N} x_i = \sum_{i=1}^{N} y_i$$

Remarks

- Majorization formalizes the notion that $\boldsymbol{x}$ is more balanced than $\boldsymbol{y}$
- `Preordering` on $\mathbb{R}^N$ but not a total ordering
  - Permutation invariant
  - Constraint on the sum of components
- Extremal points: For any element $\boldsymbol{x}$ of $\mathbb{R}^N$,

  $$\xi \cdot \boldsymbol{u} \prec \boldsymbol{x} \prec \xi \cdot (1, 0, \ldots, 0)$$

  with

  $$\xi := x_1 + \ldots + x_N$$

- A.W. Marshall and I. Olkin, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, New York (NY), 1979.

Comparing LR via majorization

- Consider two streams of requests $\boldsymbol{R}_1$ and $\boldsymbol{R}_2$ with popularity pmfs $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$, respectively. We interpret

$$\boldsymbol{R}_1 \leq_{LR} \boldsymbol{R}_2$$

to mean

$$\boldsymbol{p}_1 \prec \boldsymbol{p}_2$$

- From now on assume IRM for the input to the cache

Back to the Folk Theorem

- For a given replacement policy π, need to
  - Find the popularity pmf of the output

    $$\boldsymbol{p}^\star_\pi = (p^\star_\pi(1), \ldots, p^\star_\pi(N))$$

    with

    $$p^\star_\pi(i) = \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \mathbf{1}\left[R^\star_k = i\right] \quad a.s.$$

    for all $i = 1, \ldots, N$
  - Show that

    $$\boldsymbol{p}^\star_\pi \prec \boldsymbol{p}$$

Finding $\boldsymbol{p}_\pi^\star$

- For all $i = 1, \ldots, N$, we have

$$\begin{aligned}
p_\pi^\star(i) &= \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \mathbf{1}\left[R_k^\star = i\right] \\
&= \lim_{t\to\infty} \frac{\sum_{\tau=1}^{t} \mathbf{1}\left[R_\tau = i, R_\tau \notin S_\tau\right]}{\sum_{\tau=1}^{t} \mathbf{1}\left[R_\tau \notin S_\tau\right]} \\
&= \frac{\lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}\left[R_\tau = i, R_\tau \notin S_\tau\right]}{\lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}\left[R_\tau \notin S_\tau\right]} \\
&= \frac{p(i) m_\pi(i; \boldsymbol{p})}{\sum_{j=1}^{N} p(j) m_\pi(j; \boldsymbol{p})}
\end{aligned}$$

with

$$m_\pi(i;\boldsymbol{p}) = \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \mathbf{1}\left[i \notin S_\tau\right] \quad a.s.$$

- Note that

$$p^\star_\pi(i) = \frac{p(i)m_\pi(i;\boldsymbol{p})}{M_\pi(\boldsymbol{p})}$$

Caching and compression?

- Whenever

  $$\boldsymbol{p}^\star_\pi \prec \boldsymbol{p},$$

  we also have the `entropy` comparison

  $$H(\boldsymbol{p}) \le H(\boldsymbol{p}^\star_\pi)$$

  by the Schur-concavity of the entropy mapping

  $$\boldsymbol{x} \to H(\boldsymbol{x}) = -\sum_{i=1}^{N} x(i) \log x(i)$$

- Cache as compressor
  - "Absorbing LR is a form of compression"

The random policy

- For each $t = 0, 1, \ldots$, select U_t randomly in S_t
- For all values of $M < N$, it holds that

$$\boldsymbol{p}^\star_{\text{Rand}} \prec \boldsymbol{p}$$

so that

$$\boldsymbol{R}^\star_{\text{Rand}} \leq_{LR} \boldsymbol{R}$$

The policy A_0

- For each $t = 0, 1, \ldots$, select U_t according to

$$U_t := \arg\min\{p(i) : \ i \in S_t\}$$

  - Minimizes the miss rate under IRM($\boldsymbol{p}$)

- `Partially (pre)loaded` policy with $M - 1$ most popular documents loaded in cache (in steady state), say

$$\Sigma := \{1, \ldots, M - 1\}$$

  if we assume

$$p(1) \geq p(2) \geq \ldots \geq p(N)$$

The output under A_0

- Note that $\quad p^\star_{A_0}(i) = 0, \quad i \in \Sigma$
- Introduce the pmfs $\boldsymbol{\pi}_0$ and $\boldsymbol{p}^\star_0$ on Σ^c by

$$\pi_0(i) = \frac{p(i)}{\sum_{j \notin \Sigma} p(j)}, \quad i \notin \Sigma$$

and

$$p^\star_0(i) = p^\star_{A_0}(i) = \frac{\pi_0(i)(1-\pi_0(i))}{\sum_{j \notin \Sigma} \pi_0(j)(1-\pi_0(j))}, \quad i \notin \Sigma$$

- For all values of $M < N$, it holds that

$$\boldsymbol{p}^\star_0 \prec \boldsymbol{\pi}_0 \quad \text{on } \Sigma^c$$

so that

$$\text{“}\boldsymbol{R}^\star_{A_0} \leq_{LR} \boldsymbol{R}\text{”}$$

The LRU policy

1	2	3						$M-1$	M
i_1	i_2	i_3						i_{M-1}	i_M

- The LRU (Least-Recently-Used) policy evicts the document which was requested the least recently at the time the replacement is required
- Very popular policy
  - Self-organizing (via stack implementation)
  - Good performance in terms of miss rate
  - Statistics of $\boldsymbol{R}$ are not required (universal)

- Expressions are available: For each $i = 1, \ldots, N$, we have

$$m_{\mathrm{LRU}}(i;\boldsymbol{p}) = \sum_{s \in \Lambda_i(M;\mathcal{N})} \frac{p(i_1)\cdots p(i_M)}{\prod_{k=1}^{M-1}(1 - \sum_{j=1}^{k} p(i_j))}$$

- The folk theorem is `not` always true!

Zipf-like pmfs

- One-parameter family of pmfs sometimes called `generalized` Zipf pmfs
- With $\alpha \geq 0$, the Zipf-like pmf $\boldsymbol{p}_\alpha$ on $\{1, \ldots, N\}$ is given by

$$p_\alpha(i) = \frac{i^{-\alpha}}{C_\alpha(N)}, \quad i = 1, \ldots, N$$

with

$$C_\alpha(N) := \sum_{i=1}^{N} i^{-\alpha}$$

Remarks

- Special cases
  - If $\alpha = 0$, $\boldsymbol{p}_\alpha = \boldsymbol{u}$
  - The case $\alpha = 1$ corresponds to the standard Zipf pmf
  - If $\alpha = \infty$, $\boldsymbol{p}_\alpha = (1, 0, \ldots, 0)$
- The larger α, the more skewed $\boldsymbol{p}_\alpha$, i.e.,

$$\boldsymbol{p}_\alpha \prec \boldsymbol{p}_\beta, \quad \alpha < \beta$$

A counterexample

- Assume the popularity pmf to be the Zipf-like distribution with $\alpha \geq 0$. Under the LRU policy, if

  $$N < M!,$$

  then there exists $\alpha^\star = \alpha^\star(M, N)$ such that the comparison

  $$\boldsymbol{p}^\star_{\mathrm{LRU},\alpha} \prec \boldsymbol{p}_\alpha$$

  `cannot` hold whenever $\alpha > \alpha^\star$

- Key idea: We show for large α that

  $$\min_{i=1,\ldots,N} p^\star_{\mathrm{LRU},\alpha}(i) < p_\alpha(N)$$

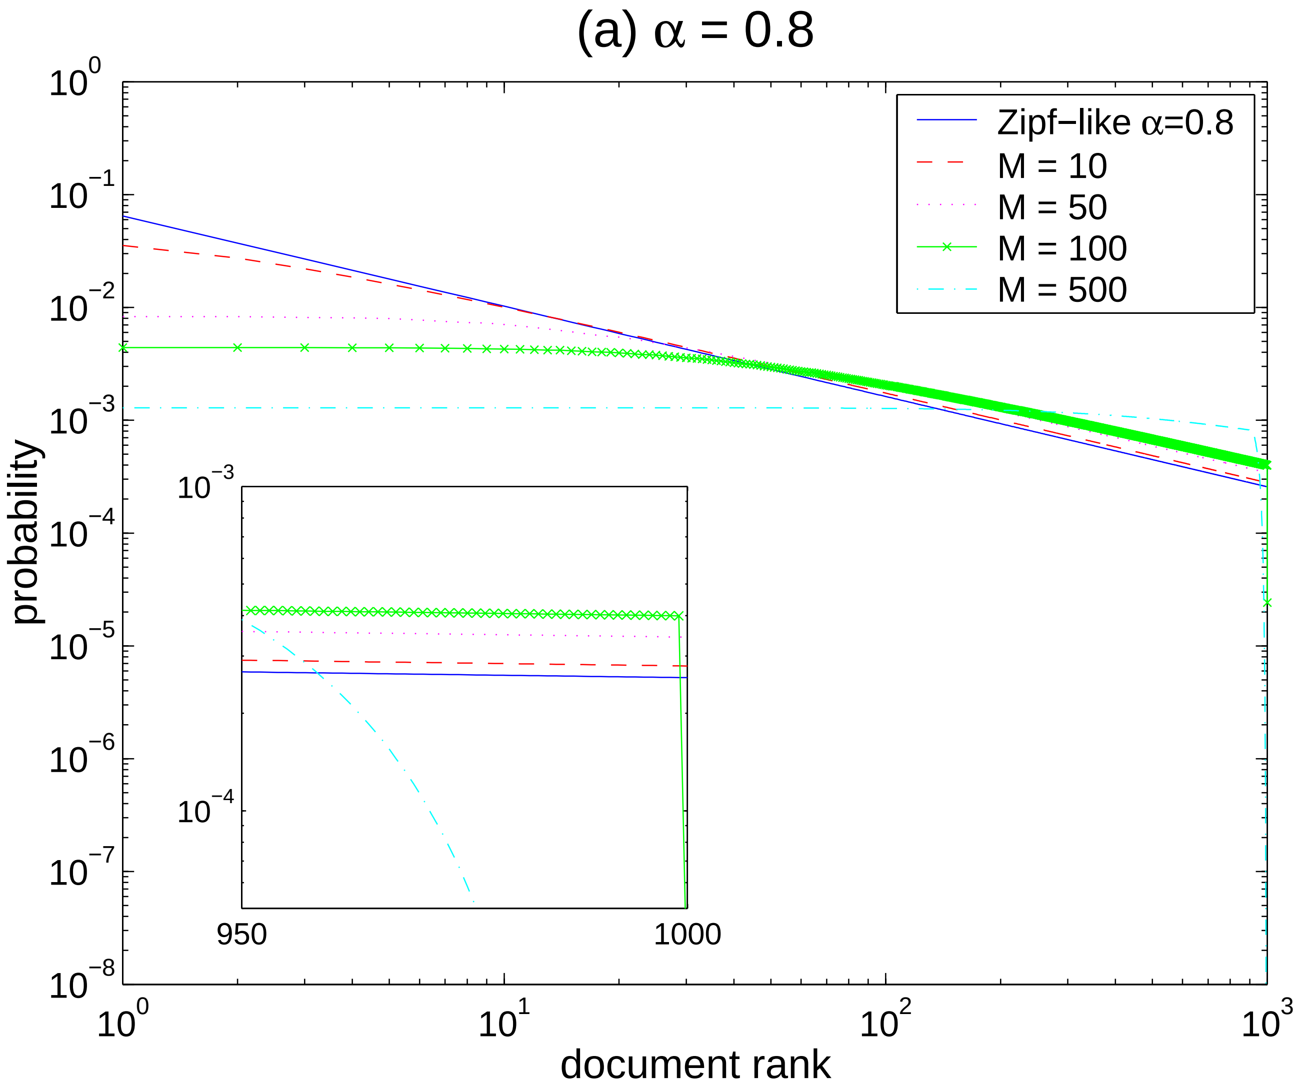
(a) α = 0.8
probability
document rank
10^0
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
10^{-8}
10^0
10^1
10^2
10^3
Zipf−like α=0.8
M = 10
M = 50
M = 100
M = 500
10^{-3}
10^{-4}
950
1000

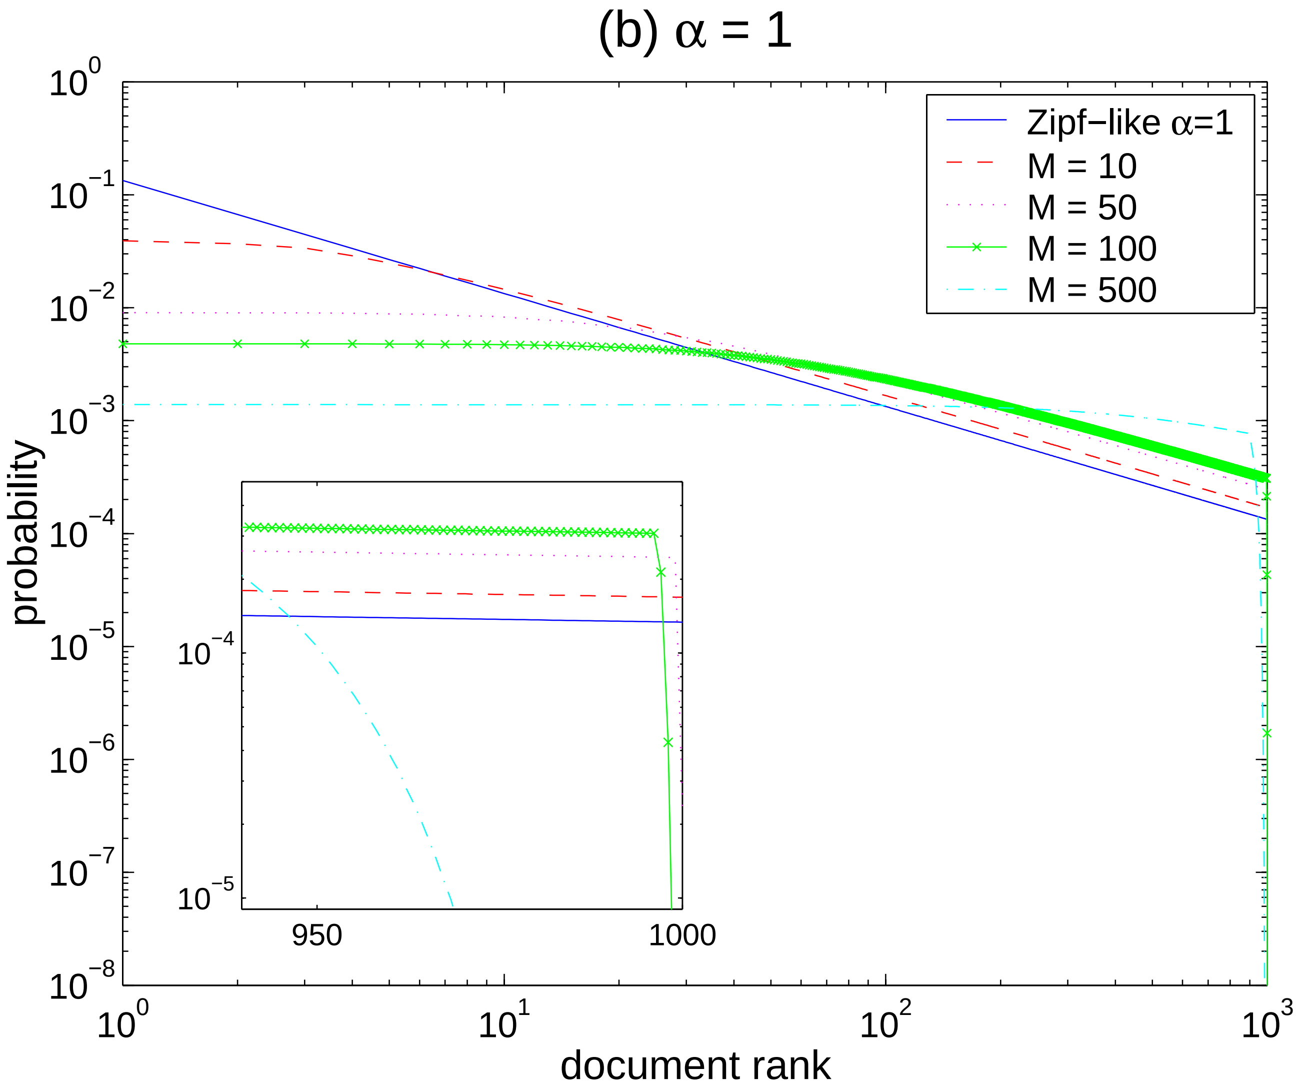
(b) α = 1
probability
document rank
10^0
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
10^{-8}
10^0
10^1
10^2
10^3
Zipf−like α=1
M = 10
M = 50
M = 100
M = 500
10^{-4}
10^{-5}
950
1000

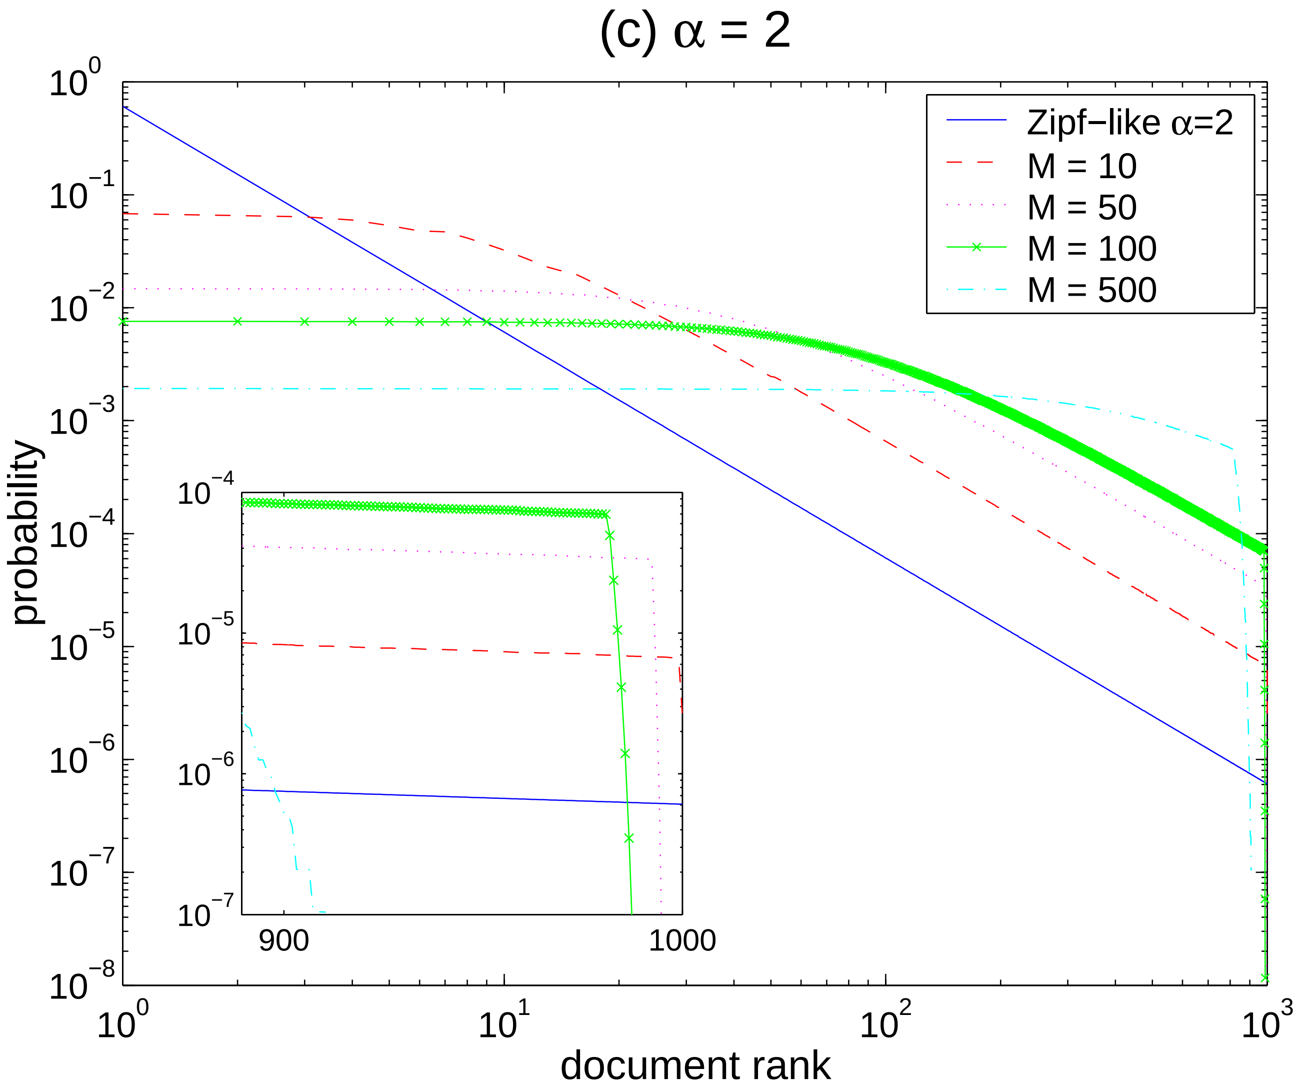
(c) α = 2
probability
document rank
10^0
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-6}
10^{-7}
10^{-8}
10^0
10^1
10^2
10^3
Zipf−like α=2
M = 10
M = 50
M = 100
M = 500
10^{-4}
10^{-5}
10^{-6}
10^{-7}
900
1000

A conjecture

- For each $N = 1, \ldots$, there exists an integer $M^\star = M^\star(\alpha; N)$ with $1 \leq M^\star < N$ such that

$$\boldsymbol{p}^\star_{\mathrm{LRU},\alpha} \prec \boldsymbol{p}_\alpha$$

whenever $M = 1, \ldots, M^\star$. Moreover, $M^\star(\alpha; N)$ decreases as α increases

- Cues
  - With $M = 1$, the comparison $\boldsymbol{p}^\star_\pi \prec \boldsymbol{p}$ valid for arbitrary N, $\boldsymbol{p}$ and π
  - With $\boldsymbol{p} = \boldsymbol{p}_0 = \boldsymbol{u}$, $\boldsymbol{p}^\star_{\mathrm{LRU},0} = \boldsymbol{p}_0$ for arbitrary N and M

Conclusions (I)

- Popularity not strong enough as a means to capture the operational meaning of LR
  - The folk theorem on the output of a cache is not necessarily true!

- Additional results available for
  - the CLIMB policy
  - the large class of Random On-demand Replacement (RORA) policies

Conclusions (II)

- Very partial results obtained only under the IRM
  - Based on known formulae
  - No probabilistic arguments
  - Zipf-like pmfs and asymptotics

- Extensions to more general models
  - Desirable
  - No algebraic structure describing the evolution of the caches (in contrast with Lindley recursion of Queueing Theory)