Web Caching and Zipf-like Distributions: Evidence and Implications

Lee Breslau Pei Cao Li Fan Graham Phillips Scott Shenker*

Abstract

This paper addresses two unresolved issues about Web caching. The first issue is whether Web requests from a fixed user community are distributed according to Zipf's law [Zip29]. Several early studies have supported this claim [Gla94, CBC95, ABCdO96] while other recent studies have suggested otherwise [NHo+98, ACC+98]. The second issue relates to a number of recent studies on the characteristics of Web proxy traces, which have shown that the hit-ratios and temporal locality of the traces exhibit certain asymptotic properties that are uniform across the different sets of the traces [CI97, RV98, DMF97, GB97a, KLM97]. In particular, the second issue is whether these properties are inherent to Web accesses or whether they are simply an artifact of the traces. An answer to these unresolved issues will facilitate both Web cache resource planning and cache hierarchy design.

We show that the answers to the two questions are related. We first investigate the page request distribution seen by Web proxy caches using traces from a variety of sources. We find that the distribution does not follow Zipf's law precisely, but instead follows a Zipf-like distribution with the exponent varying from trace to trace. Furthermore, we find that there is (i) a weak correlation between the access frequency of a Web page and its size and (ii) a weak correlation between access frequency and its rate of change. We then consider a simple model where the Web accesses are independent and the reference

*Lee Breslau and Scott Shenker are with the Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304. Pei Cao and Li Fan are with the Computer Science Department, University of Wisconsin, Madison, WI 53706. Graham Phillips is with the Computer Science Department, University of Southern California, Los Angeles, CA 90089.

probability of the documents follows a Zipf-like distribution. We find that the model yields asymptotic behaviors that are consistent with the experimental observations, suggesting that the various observed properties of hit ratios and temporal locality are indeed inherent to Web accesses observed by proxies.

Finally, we revisit Web cache replacement algorithms and show that the algorithm that is suggested by this simple model performs best on real trace data. The results indicate that while page requests do indeed reveal short-term correlations and other structures, a simple model for an independent request stream following a Zipf-like distribution is sufficient to capture certain asymptotic properties observed at Web proxies.

1 Introduction

Due to the explosive growth of the Web, Web proxy caching has recently received considerable attention. It is considered one of the most important techniques at reducing Web traffic, which accounts for a large percentage of Internet traffic today. Several researchers have observed that the relative frequency with which Web pages are requested follows Zipf's law [Zip29]. Zipf's law states that the relative probability of a request for the i'th most popular page is inversely proportional to $1/i$. Glassman [Gla94] was perhaps the first use of Zipf's law to model the distribution of Web page requests and several other authors have also applied Zipf's law to the distribution of web requests [CBC95, ABCdO96]. However, recently several studies have questioned whether the requests do indeed follow Zipf's law and concluded otherwise [NHo$^+$98, ACC$^+$98]
.

One of our goals in this paper is to investigate the applicability of Zipf's law. Using six traces from proxies at academic institutions, corporations and ISPs, we find that *the distribution of page requests generally follow a Zipf-like distribution where the relative probability of a request for the i'th most popular page is inversely proportional to $1/i^\alpha$, with α typically taking on some value less than unity*. The observed values of the exponent α vary from trace to trace. That is, the request distribution does not follow the strict Zipf's law (for which $\alpha = 1$), but instead follows a more general Zipf-like distribution with varying α. Moreover, we find that there is little correlation between the access frequency of a document and its size; that is, popular documents are not necessarily smaller than unpopular ones. The correlation between the access frequency of a document and its rate of modification varies from very low to none, depending on the traces. These results raise the possibility that, for some purposes, one might be able to sufficiently model web accesses by a simple model that assumes independent references following a Zipf-like distribution and no correlation between request frequency and response size or rate of change.

In looking at Web proxy traces, researchers have also investigated how the hit-ratio depends, asymptotically, on the cache size and the number of requests, and have examined the temporal locality of these request streams [CI97, RV98, DMF97, GB97a, KLM97]. Although various studies have used different set of traces, the following three qualitative asymptotic properties have been identified:

- *For an infinite sized cache, the hit-ratio for a Web proxy grows in a log-like fashion*[1] *as a function of the client population of the proxy and of the number of requests seen by the proxy.* Cao et al. observed this property in Digital Equipment Corporation's proxy traces [CI97, KMM96], Gribble et al. observed this property in University of California at Berkeley's proxy traces [GB97a, GB97b] and Duska et al. observed this property in a number of traces from university proxies and ISP proxies [DMF97].

- *The hit-ratio of a Web cache grows in a log-like fashion as a function of the cache size.* Many Web caching studies reach this conclusion [ABCdO96, Gla94, CI97, WAS$^+$96, GB97a, RV98, CBC95, DMF97].

- *The probability that a document will be referenced k requests after it was last referenced is roughly proportional to $1/k$.* That is, Web traces exhibit excellent temporal locality. Of the two studies that

[1] *That is, the hit-ratio grows as either the log, or as a small power, of the argument. Given the coarseness of the data, it is impossible to distinguish between a logarithmic and a small power growth.*

investigated temporal locality, Rizzo et al. observed this property in the Web proxy traces collected at University of Pisa, Italy [RV98], while Cao et al. observed this property in the Digital Equipment Corporation's proxy traces [CI97, KMM96].

These observations were offered without an underlying explanation. It is not clear whether these properties follow certain inherent characteristics of Web accesses or are simply an artifact of the collection of traces studied. The properties are useful for designing caching algorithms and configuring proxy caches, and therefore it is important to understand how general they are.

In the second portion of this paper, we show that the asymptotic behavior of the cache-hit rate and temporal locality are related to the Zipf-like nature of the request stream. More specifically, we show that *if one assumes that Web page requests are independent and the probability that a page is accessed follows a Zipf-like distribution then the three asymptotic properties listed above all follow*. Although the assumption that the requests are independent is obviously an over-simplification, our results show that the model is powerful enough to explain the three asymptotic properties as observed in real proxy traces.

The rest of the paper is organized as follows. In Section 2 we describe the trace data and the resulting frequency distribution of requests. Section 3 discusses the implications of the Zipf-like behavior, with Section 3.1 describing the simple model for page requests and Sections 3.2, 3.3 and 3.4 deriving expressions for cache hit-ratio and page request interarrival times. We analyze Web cache replacement algorithms in Section 4 where we compare the LRU and LFU page replacement policies. We conclude in Section 5 with a short discussion of our results.

2 The applicability of Zipf's Law to Web requests

A number of prior studies have investigated the application of Zipf's law to Web accesses, and arrived at different conclusions.

- In 1994, Glassman et al [Gla94] traced Web accesses from Digital Equipment Corporation's Palo Alto facilities and gathered about 100,000 HTTP and Gopher requests from about 300 users; they found that the request distribution fit the Zipf's law, Ω/i, very well.
- In 1995, Cunha et al [CBC95] gathered over 500,000 Web accesses from the users at the Computer Science Department at Boston University for a period of four months and observed that the requests follow an Ω/i^{α} distribution where $\alpha = 0.986$, which is quite close to the true Zipf's law.
- In 1996, Almeida et al [ABCdO96] showed that Web accesses seen by a Web server follow Zipf's law. However, the Web accesses seen at a Web server and those seen at a Web proxy (which are what we are interested in) may well be different, since the former involve all the users on the Internet and one Web server, and the latter involves all the servers on the Internet and a fixed group of users.
- In 1998, Nishikawa et al [NHo$^+$98] from Hitachi, Ltd, Japan studied a 16-day access log involving over 2,000,000 requests that went through the company firewall and found that the request distribution followed the Zipf-like distribution Ω/i^{α} quite well, but with $\alpha = 0.75$ which is rather far from true Zipf behavior.
- In 1998, Almeida et al [ACC$^+$98] examined requests at a major proxy cache in Brazil and concluded that the request stream did not fit Zipf's law.

There might be several reasons for the conflicting results, including the fact that some of the traces are old and some are quite short. Adding further to the confusion is the phrase "Zipf's law," which in some papers refer to the Ω/i distribution, and in some other papers refer to the Ω/i^{α} distribution with $\alpha \leq 1$.

To address the issue, we study six traces from academic, corporate and ISP environments, and examine whether the request distribution follows Ω/i or Ω/i^{α} with $0 < \alpha < 1$. We call Ω/i^{α} with $0 < \alpha < 1$ a *Zipf-like behavior*, and reserve the phrase Zipf's law for the true Ω/i distribution.

The six traces are:

- **DEC** traces [KMM96]: Digital Equipment Corporation Web Proxy traces, servicing about 17,000 workstations. The trace is for a period of 25 days (Aug. 29 to Sep. 21, 1996). Here we present the results on the trace of the week of Aug. 29 to Sep. 4, 1996, covering 3,543,968 requests. Results on other traces are similar.
- **UCB** traces [GB97b]: traces of HTTP requests gathered from the Home IP service offered by UC Berkeley to its students, faculty, and staff. The total trace is for a period of 18 days from Nov. 1 till Nov. 19, 1996. Here, we present the results on the traces from Nov. 14 till Nov. 19, covering 1,907,762 requests. Again, results on other parts of the traces are similar.
- **UPisa** traces [Riz97]: traces of HTTP requests made by the users in Computer Science Department in Universita di Pisa, Italy, for a period of three months from January to March, 1997. We only examine GET requests, and only those whose URLs do not include query strings, since most proxies do not cache query requests. The trace includes 2,833,624 requests.
- **Questnet** traces [Wee98]: 7-days worth of logs of HTTP requests seen by the parent proxies at Questnet, which is a regional ISP in Australia, from Jan. 15 to Jan. 21, 1998. The proxies are parent proxies serving about 12 child proxies in the regional network. We extract successful GET requests seen by the parent proxies. The trace includes 2,885,285 requests.
- **NLANR** traces [oANR97]: one-day log (Dec. 22, 1997) of HTTP requests to four major parent proxy caches in the National Caching hierarchy by the National Lab for Applied Networking Research, USA. We extract successful GET requests, totaling 1,766,409.
- **FuNet** traces [Jar98]: 10-day log (around mid June, 1998) of HTTP requests to proxies at FUNET, a regional network serving the academic and research communities in Finland. It contains 4,815,551 requests.

Our main observations are as follows:

Traces	DEC	UPisa	FuNet	UCB	Questnet	NLANR
α	0.77	0.78	0.83	0.69	0.73	0.64

Table 1: α value for each trace.

- The distribution of web requests from a fixed group of users follows a Zipf-like distribution, Ω/i^{α}, very well. The value of α varies from trace to trace, ranging from 0.64 to 0.83. Traces from a homogeneous community (for example, DEC, Pisa and FuNet) have a larger α (centering around 0.8), and traces from a more diversified environment (for example, UCB, QuestNet and NLANR) have a smaller α (centering around 0.7).

- As follows immediately from the Zipf-like distribution, the α value affects the concentration of Web accesses to hot documents in each trace. For example, 25% of all documents accounts for 70% of Web accesses in DEC, Pisa and FuNet traces, while it takes 40% of all document to draw 70% of Web accesses in UCB, QuestNet and NLANR. The "10/90" rule (i.e. 90% of accesses go to 10% of items), evident in program execution, does not apply to Web accesses seen by a proxy.

- There is low statistical correlation between the frequency that a document is accessed and its size, though the average size of cold documents (for example, those accessed less than 10 times) is larger than that of hot documents (for example, those accessed more than 10 times).

- The statistical correlation between a document's access frequency and its average modifications per request varies from trace to trace, but is generally quite low.

The results are detailed in the following sections.

2.1 Is the distribution of Web requests Zipf-like?

Figure 1 shows, for each of the six traces, the number of times that a URL has been accessed versus the ranking of the URL in the collection of URLs in a trace (the 1st URL is the most popular one). The y-axis

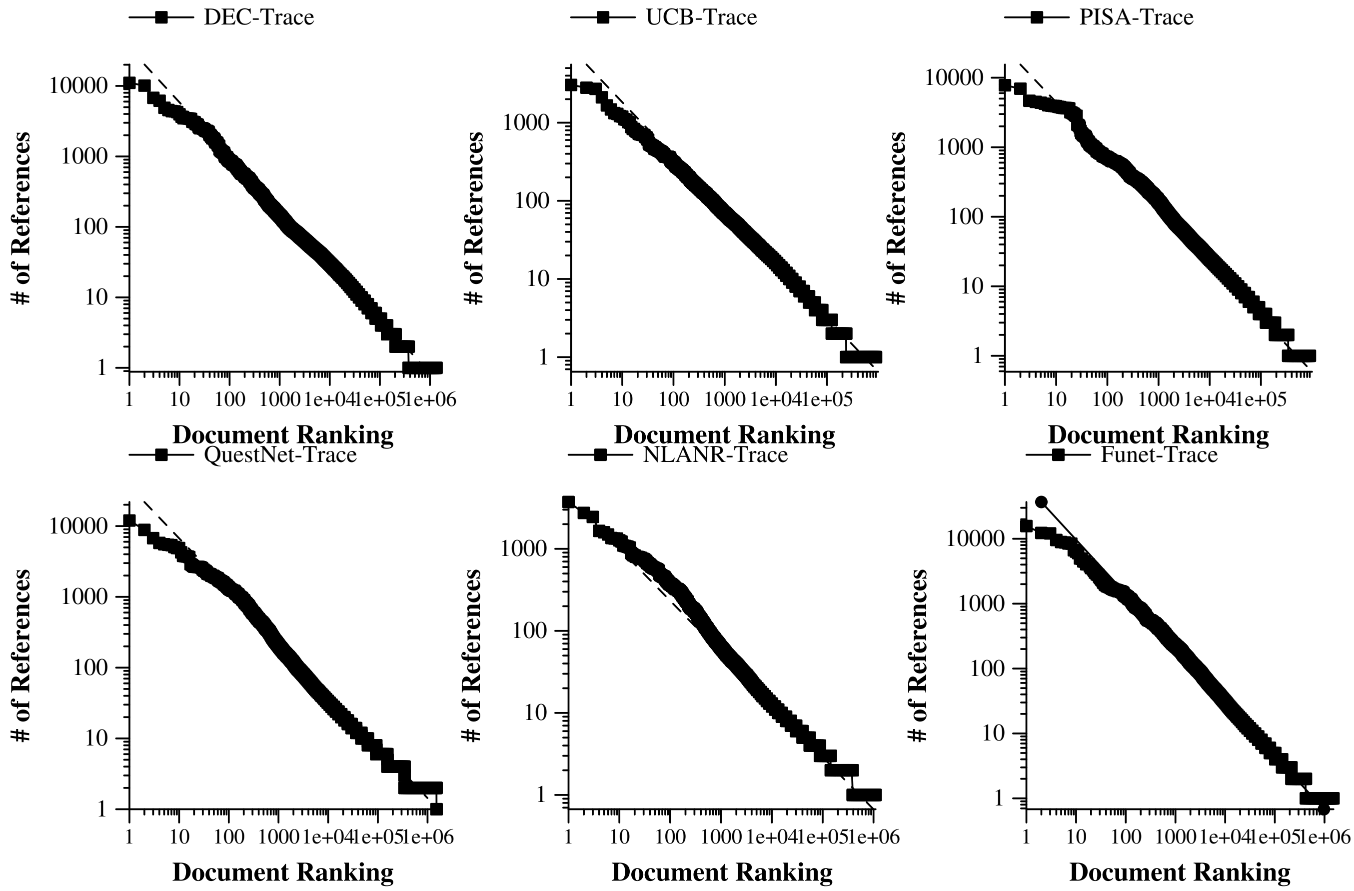

Figure 1: Frequency of document accesses versus document ranking.

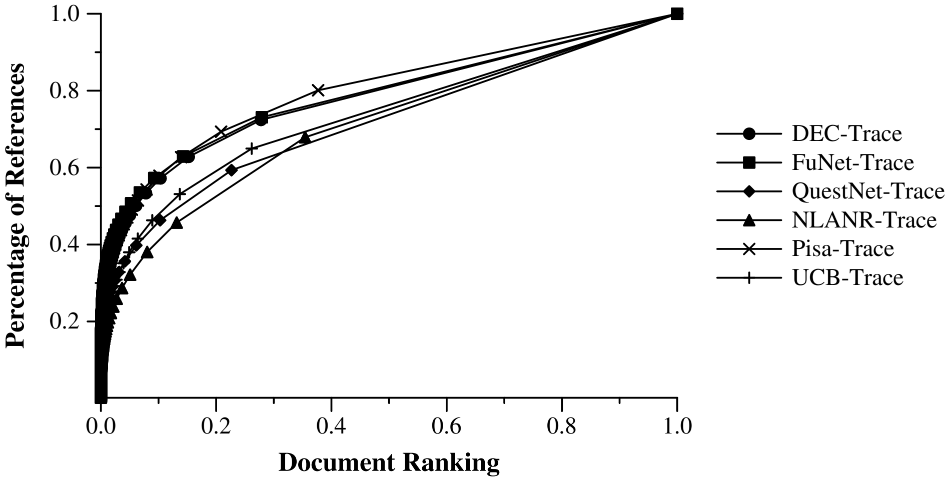

Figure 2: Cumulative distribution of requests to documents.

and x-axis in the figure are in log scale.

From the plots, one can observe that except for the beginning and the end of the curve, most of the curve fits a straight line well. Since this is a log-log curve, a straight line means that the request frequency is proportional to $1/i^{\alpha}$. The value of α for the six traces are shown in Table 1. Thus, the clear indication of this data is that a Zipf-like distribution fits the data fairly well, but that the value of α varies from trace to trace.

For traces from a homogeneous environment, as in a corporation (DEC) or a departmental (Pisa) setting, α appears to be larger, centering around 0.8. For traces from a more diversified user population (for example, UCB) or those that are filtered by first-level proxies (for example, QuestNet and NLANR), α appears to be smaller, centering around 0.7.

Perhaps the biggest impact of α lies in the concentration of Web requests to hot documents. For Zipf-like distributions, the cumulative probability that one of the top k page is accessed is given asymptotically by $\phi(k) = \sum_{i=1}^{k} \cdot \Omega / i^{\alpha} \approx \Omega k^{1-\alpha}/(1-\alpha)$. Since $\Omega \approx (1-\alpha)/N^{1-\alpha}$, where N is the total number of Web documents, $\phi(k) \approx (k/N)^{1-\alpha}$. Because $k/N < 1$ for all meaningful k, a larger α increases $\phi(k)$, meaning more requests are concentrated on a few hot documents.

Since the request concentration has broad implications on multicast/broadcast schemes for the Web and

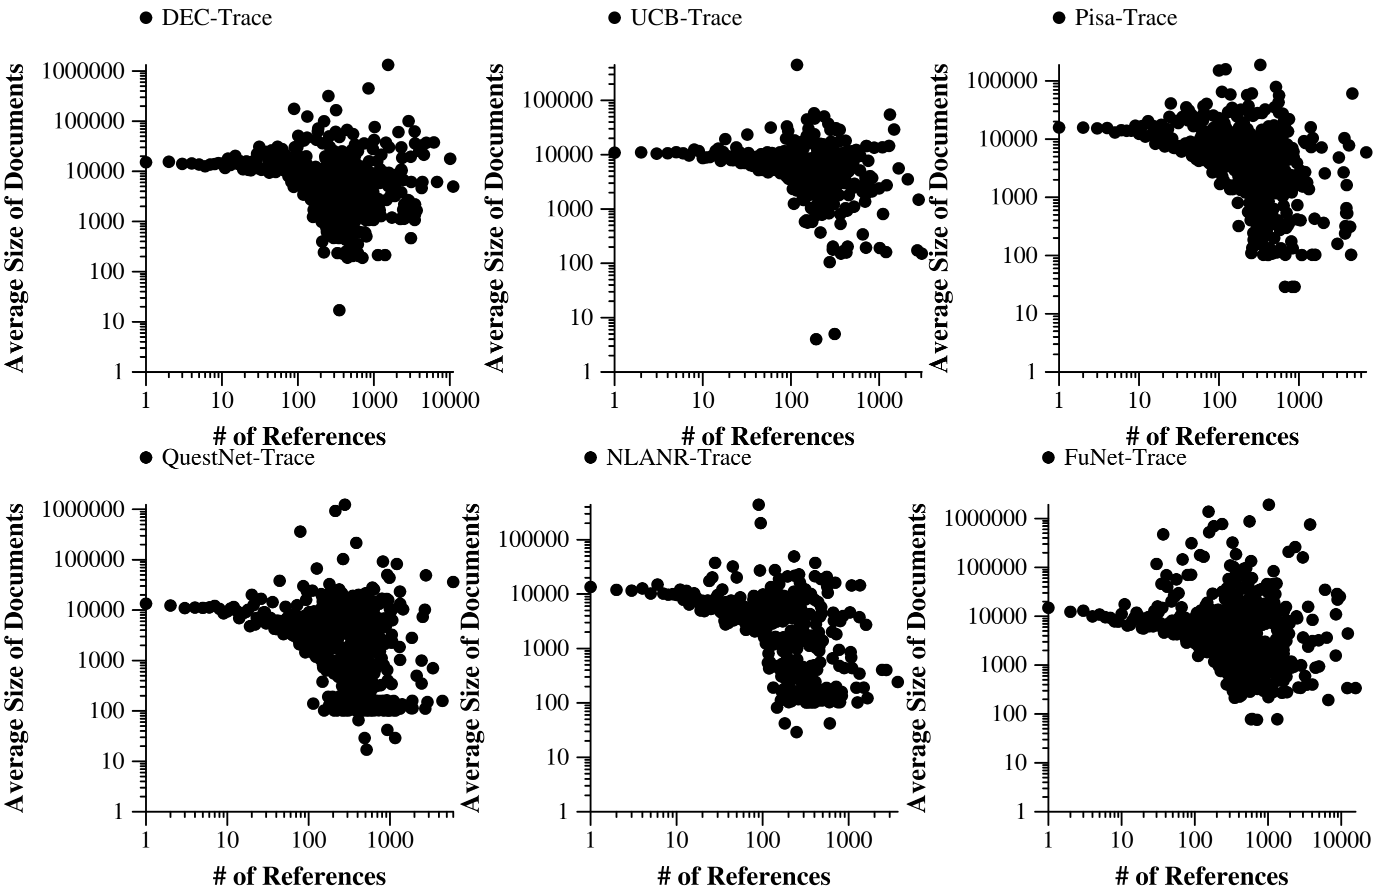

Figure 3: Average size of a document versus access frequency.

Web cache design, we examine the the cumulative distribution of requests to popular documents in each trace. Figure 2 shows the cumulative probability of access for top r% of documents in each trace. As shown clearly, the top 1% of the documents account for about 20% to 35% of all requests seen by the proxy, and the top 10% of the documents account for about 45% to 55% of all requests. However, it takes 25% to 40% of all documents to account for 70% of the requests, and it takes 70% to 80% of documents to account for 90% of requests. In other words, the "10/90" rule, evident in program execution, unfortunately does not apply to Web accesses. The results point to both the potential and the limitation of satellite-based broadcast schemes for the Web.

Traces	DEC	UPisa	FuNet	UCB	Questnet	NLANR
Cor. Coef.	0.04	-0.08	0.003	-0.04	-0.02	-0.09

Table 2: Correlation Coefficient between the frequency of accesses to a document and its size.

Traces	DEC	UPisa	FuNet	UCB	Questnet	NLANR
Cor. Coef.	-0.19	-0.27	0.005	0.002	-0.03	-0.08

Table 3: Correlation Coefficient between document change rate and access frequency.

2.2 Correlation between Access Frequency and Document Size

A natural question is whether there is any correlation between the frequency of access to a document and its size. The answer affects caching algorithm design and design of Web proxy benchmarks. Figure 3 shows the average size of the documents that are accessed exactly n times in the trace, with the x-axis showing n. Again, both x-axis and y-axis are in log scale. The figure shows that there are many popular documents that are larger than 15KB, which is the average size of documents that are accessed only once. Furthermore, across n, there does not seem to be strong correlation between document size and n, though the average size of popular documents is smaller than that of unpopular documents.

We also calculated the correlation coefficient between the access frequency and document size for the traces, shown in Table 2. The numbers show that the correlation, if any, is weak and can be ignored.

2.3 Correlation between Access Frequency and Change Rate

Another question of interest is whether popular documents change less often than unpopular ones. The answer affects designs of Web cache consistency mechanisms [GS96].

From the traces, we observe the changes in a document by detecting changes in the last-modified-time of each response or changes in the document size of each response. Only two of our traces, DEC and UCB, come with the last-modified-time information. For each document, we then measure the "change rate" r as

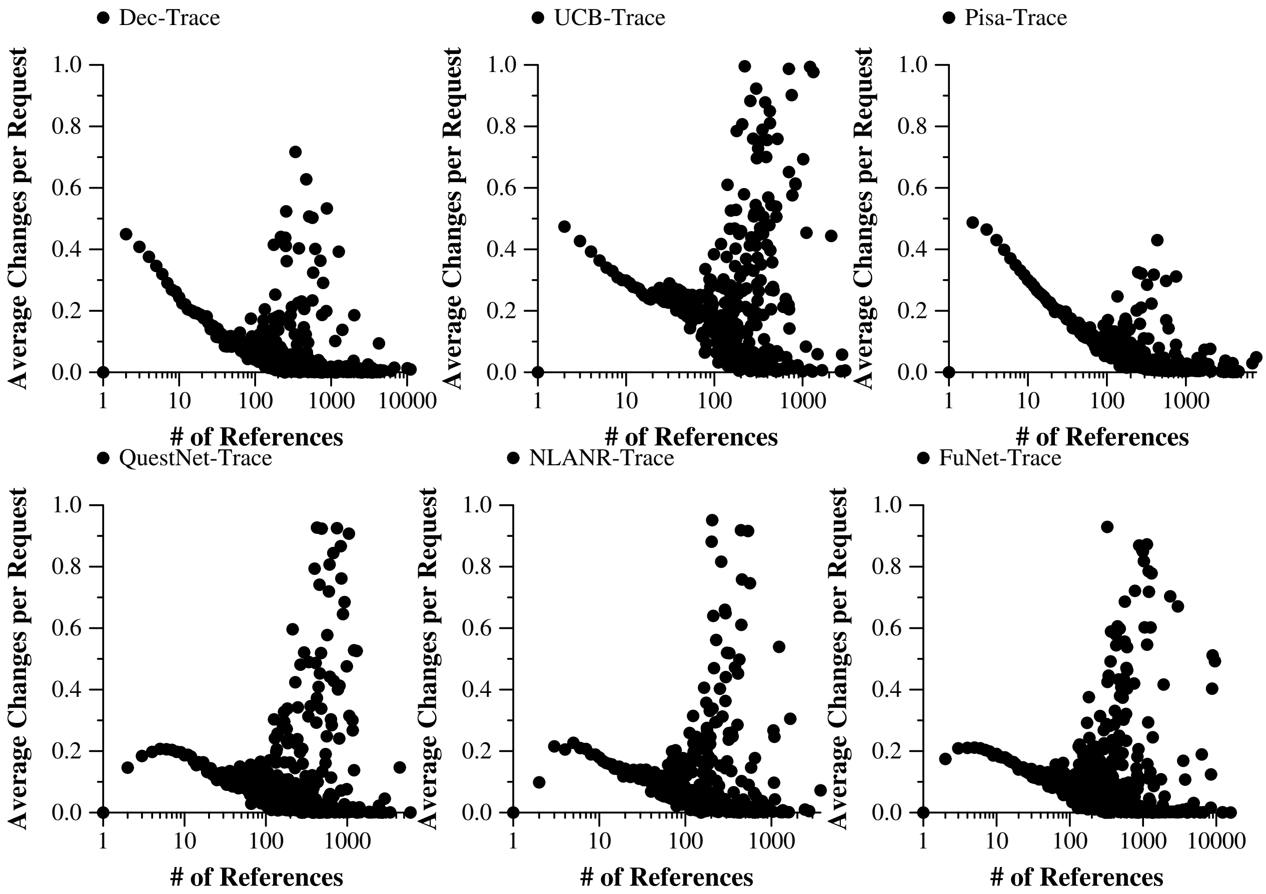

Figure 4: Change rate of a document versus access frequency.

the ratio between the number of observed changes and the number of accesses. Note that the change rate may be an under-estimation of the actual change rate, particularly for documents that are accessed a few times, because multiple changes between requests for a document are not picked up by the trace.

Figure 4 plots the average change rate of the documents as a function of the number of times they are accessed. Only the x-axis is in log scale. The graph show that while many popular documents change infrequently (for example, less than 10 user requests per change), some popular documents change very frequently. We again calculate the correlation coefficient between the document change rate and its access frequency. The results are shown in Table 3. Though the correlation coefficients vary depending on the traces, they are small in general. Thus, there is no strong evidence that popular documents change less often. In the design of any Web cache consistency scheme or any Web access model, it is perhaps best to assume that there is no correlation between document popularity and its rate of change.

2.4 Request Distribution to Web Servers

We have also investigated the distribution of requests to Web servers (i.e. how many requests are forwarded to a particular Web server). We found that, in general, a Zipf-like distribution does not characterize the request frequency well. This is particularly true in the case of the UCB trace. For other traces, Zipf-like distribution is at best a rough approximation. We also found that no single Web server contributes to most of popular pages, and popular Web pages are spread almost evenly across hot Web servers. Due to space limitation we cannot include the relevant figures here; interested readers are encouraged to contact the authors for more detail.

3 Implications of Zipf-like Behavior

We now proceed to investigate the implications of the Zipf-like behavior observed in the traces. First, we define a simple model for web requests where the requests are independent and distributed according to a

Zipf-like behavior. We then derive expressions for the hit-ratios and request interarrival times.

3.1 The Model

Consider a cache that receives a stream of requests for Web pages. Let N be the total number of Web pages in the universe. Let $P_N(i)$ be the conditional probability that, given the arrival of a page request, the arriving request is made for page i. Let all the pages be ranked in order of their popularity where page i is the i'th most popular page. We assume that $P_N(i)$, defined for $i = 1, 2, \ldots, N$, has a "cut-off" Zipf-like distribution given by

$$P_N(i) = \frac{\Omega}{i^\alpha},$$

where

$$\Omega = \left(\sum_{i=1}^{N} \frac{1}{i^\alpha} \right)^{-1}$$

The true Zipf's law [Zip29] has $\alpha = 1$ but in this paper we consider a broader class of distribution functions with exponents in the range $0 < \alpha \leq 1$. Each page request is drawn independently from the Zipf distribution, so we are neglecting any other source of correlations in the request stream. We additionally assume that, over the course of time, no pages are invalidated by the cache.

This model, having complete independence between cache requests, is completely unrealistic. However, it is a quite tractable model, and our question here is whether or not the model is sufficient to derive understanding about the asymptotic properties of hit-ratios and request interarrival times.

We calculate the hit-ratio for the cache in two extreme limiting cases and also the interrival times in a limiting case; we consider these in the following three subsections.

3.2 Infinite Cache, Finite Request Stream

We first consider the case where the cache has unlimited storage, so that all previously requested pages remain in the cache. In this case, we consider a finite request stream of R requests, and wish to determine the probability that the next request, the $R+1$'st request, is a request for a page that already resides in the cache. The hit-ratio $H(R)$ can be calculated as follows. If the $R+1$'st request is for page i then the probability that this page is in the cache is given by $(1-(1-P_N(i))^R)$. Thus, we have:

$$H(R) = \sum_{i=1}^{N} P_N(i)\left(1-(1-P_N(i))^R\right) \tag{1}$$

For $\alpha = 1$ the asymptotic behavior of the hit-ratio is $H(R) \approx \ln R$, while for other values of α the asymptotic behavior of the hit-ratio is $H(R) \approx R^{(\frac{1}{\alpha}-1)}$. This asymptotic behavior can be seen more directly by approximating the sum for $H(R)$ by assuming the function $f(i) = (1-(1-\frac{\Omega}{i})^R)$ is given, for $1 \ll R \ll N$, by:

$$f(i) = \begin{cases} 1, & \text{if } i^\alpha \leq R\Omega \\ 0, & \text{otherwise} \end{cases}$$

Then we have

$$H(R) = \sum_{i=1}^{N} \frac{\Omega}{i^\alpha} f(i) \approx \sum_{i=1}^{(R\Omega)^{\frac{1}{\alpha}}} \frac{\Omega}{i^\alpha}$$

For $\alpha = 1$, $H(R) \approx \Omega \ln(R\Omega)$, while for $0 < \alpha < 1$, $H(R) \approx \frac{\Omega}{1-\alpha}(R\Omega)^{(\frac{1}{\alpha}-1)}$. Note that when $R > N$ this approximation is no longer valid because $H(R)$ must approach unity while these expressions diverge for large R. This result that the hit-ratio grows logarithmically or like a small power is qualitatively consistent with previously observed behavior [GB97a, CI97, DMF97].

Figure 5 shows the hit-ratio $H(R)$ empirically observed in the cache traces as a function of R (on a log-log plot). Document modifications are not considered. The data for each individual trace is, on this log-log plot, roughly linear, thereby suggesting that indeed $H(R) \approx R^\beta$ for some β (where the power varies from

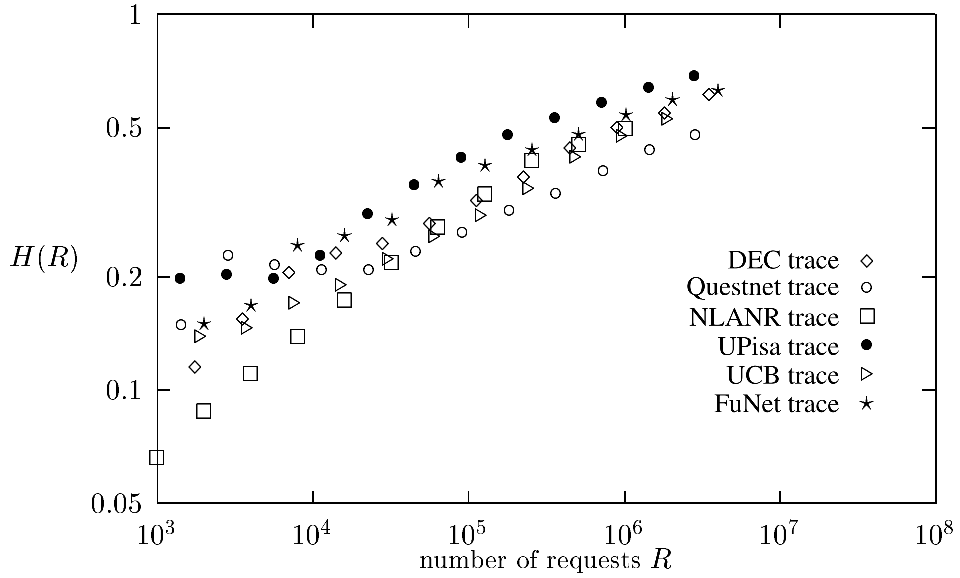

Figure 5: Hit-ratio as a function of the number of requests.

Traces	DEC	UPisa	FuNet	UCB	Questnet	NLANR
α	0.83	0.84	0.84	0.83	0.88	0.90

Table 4: α value computed from linear fit of $\log H(R)$ versus $\log R$ for each trace.

trace to trace). We have fit each of these log-log data sets to a straight line (yielding the value of β), and then used the fact that theory predicts that $H(R) \approx R^{(\frac{1}{\alpha}-1)}$ to compute an effective α from β: that is, $\alpha = \frac{1}{1+\beta}$. Table 4 displays these α values. The agreement with the α values computed directly from the distribution (in Table 1 is good, although the Questnet and NLANR results are more inaccurate than the others.

3.3 Finite Cache, Infinite Request Stream

In this section we consider a finite cache with a capacity of C Web pages subject to an infinitely long request stream. We assume that the cache can hold C Web pages regardless of the size of each Web page. Furthermore, we assume that the cache holds the C most popular pages as indicated by the ordering i. The ordering can be determined by measuring the request frequency of each page, which is equivalent to assuming that the cache has a Perfect-LFU page removal policy (see Section 4 for a discussion of the

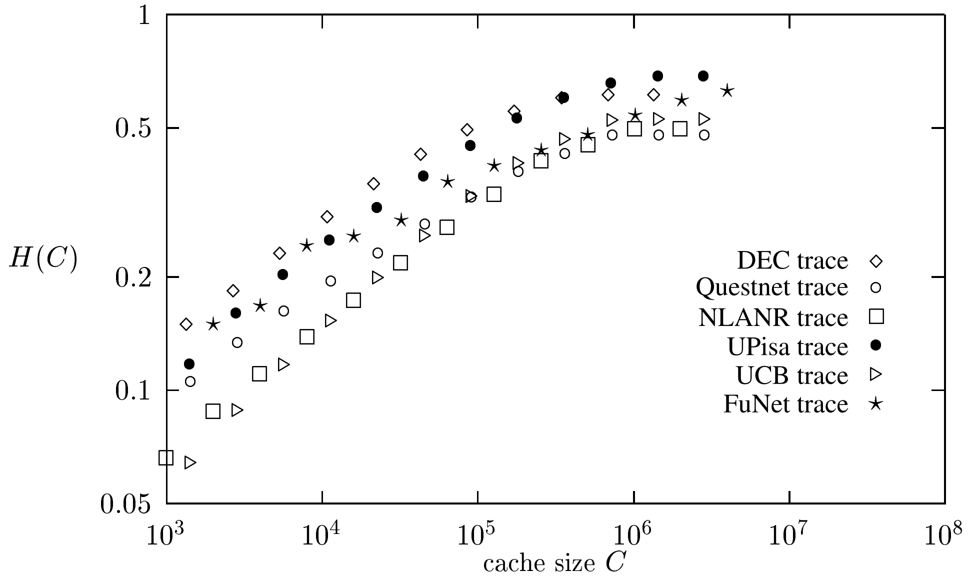

Figure 6: Hit-ratio as a function of cache size.

Perfect-LFU policy).

We are interested in the asymptotic hit-ratio $H(C)$, which is given by:

$$H(C) = \sum_{i=1}^{C} P_N(i)$$

In the case $\alpha = 1$, then asymptotically $H(C) \approx \ln C$, and for $0 < \alpha < 1$, then asymptotically $H(C) \approx C^{1-\alpha}$. This result is qualitatively consistent with previously observed behavior that the hit-ratio increases logarithmically or as a small power as a function of cache size [ABCdO96, Gla94, CI97, WAS$^+$96, GB97a, RV98, CBC95, DMF97]. These references do not claim any particular form for the asymptotic behavior, but their data appears to grow in a log-like (or small-power-like) fashion in the asymptotic regime.

Figure 6 displays the hit-ratio for the various Web cache traces on a log-log plot, ignoring document modifications. Note that some of the graphs for the traces flatten out for large R. This behavior is expected because, for finite sized traces, the hit-ratio will reach a maximum and then remain constant when the cache size becomes larger than the trace file; thus, this flattening out behavior is a limitation of the sizes of the trace files rather than an indication of the true asymptotic behavior.

Traces	DEC	UPisa	FuNet	UCB	Questnet	NLANR
α	0.77	0.76	0.81	0.70	0.76	0.67

Table 5: α value computed from linear fit of $\log H(C)$ versus $\log C$ for each trace.

The results for each individual trace is, on this log-log plot, roughly linear. As with the $H(R)$ data, this suggests that indeed $H(C) \approx C^{\beta}$ for some β (where the power varies from trace to trace). Since the theory predicts that $H(C) \approx C^{1-\alpha}$ we have computed the best linear fit to each data set, using the slope as the value of β and then computing $\alpha = 1 - \beta$. Table 4 displays these computed α values. The agreement with the α values computed directly from the distribution (shown in Table 1) is remarkably good. Thus, the simple analytical model has been very successful in predicting the general asymptotic behavior of $H(R)$ and $H(C)$.

3.4 Page Request Interarrival Times

We now investigate the distribution of times between requests for a given page.

Let us consider an infinite arrival stream and consider a request for a given page i. The quantity of interest is the probability distribution $d(k)$ that the next request for that page is k requests later (*i.e.*, that the request for page i is followed by $k-1$ requests for pages other than page i, followed by another request for page i). Assuming that page requests are independent, we find that

$$d(k) = \sum_{i=1}^{N} \left(P_N(i)\right)^2 \left(1 - P_N(i)\right)^{k-1} \tag{2}$$

For $\alpha = 1$ we have $\Omega \approx \frac{1}{\ln N}$ and therefore[2]

$$d(k) \approx \frac{1}{k \ln N} \left(\left(1 - \frac{1}{N \ln N}\right)^{k} - \left(1 - \frac{1}{\ln N}\right)^{k} \right)$$

[2]The approximation for $d(k)$ is made by replacing the sum with an integral in equation (2), substituting $P_N(i) = \Omega/i$, then substituting $u = 1 - \Omega/i$ and finally noting that $du = \Omega/i^2 di$.

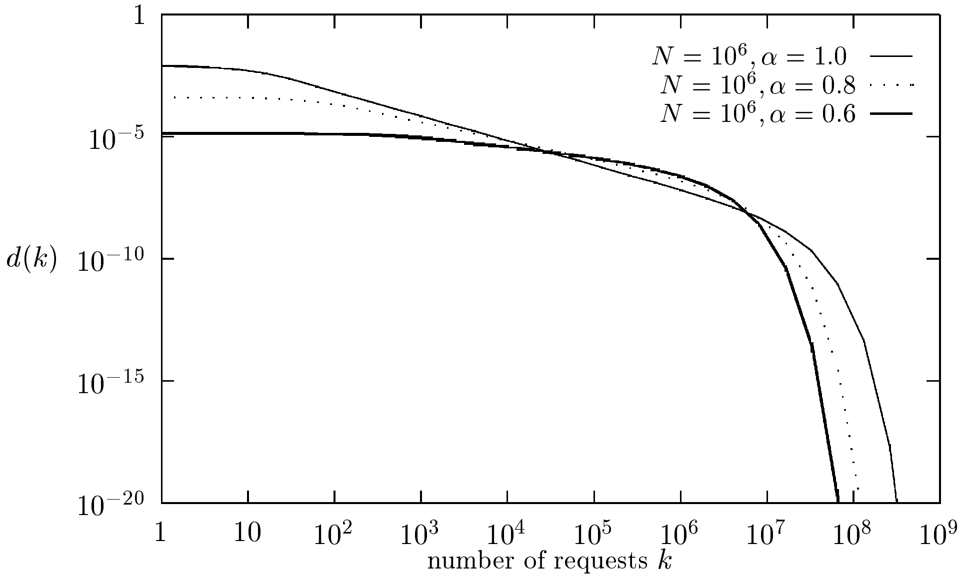

Figure 7: Probability distribution for request interarrival times.

Note that for $N \ln N \gg k \gg \ln N$, $d(k) \approx 1/(k \ln N)$.

For $0 < \alpha < 1$ the asymptotic behavior of this expression is less tractable, but as shown in Figure 7 the behavior in the intermediate regime $N \ln N \gg k \gg \ln N$ is qualitatively similar to the $\alpha = 1$ case, but with different slopes in the linear regime. The model's predictions, in the intermediate regime, are qualitatively consistent with data observed at operational web servers [RV98, CI97].

We have also examined the request interarrival behavior in our traces. Figure 8 shows a plot of the probability distribution for page request interarrival times $d(k)$ produced by our model and the distribution for three cache traces (we did not show the remaining three traces to keep the graph readable). The agreement between the simple model and the trace data is quite good in the intermediate regime.

4 Cache Replacement Algorithms

We next consider whether the simple model introduced in the previous section can be used to inform cache replacement strategies. The model that we have discussed so far is called an *independent reference model* [GCD73] in the early operating system paging studies [Den80]. It is well known in the operating

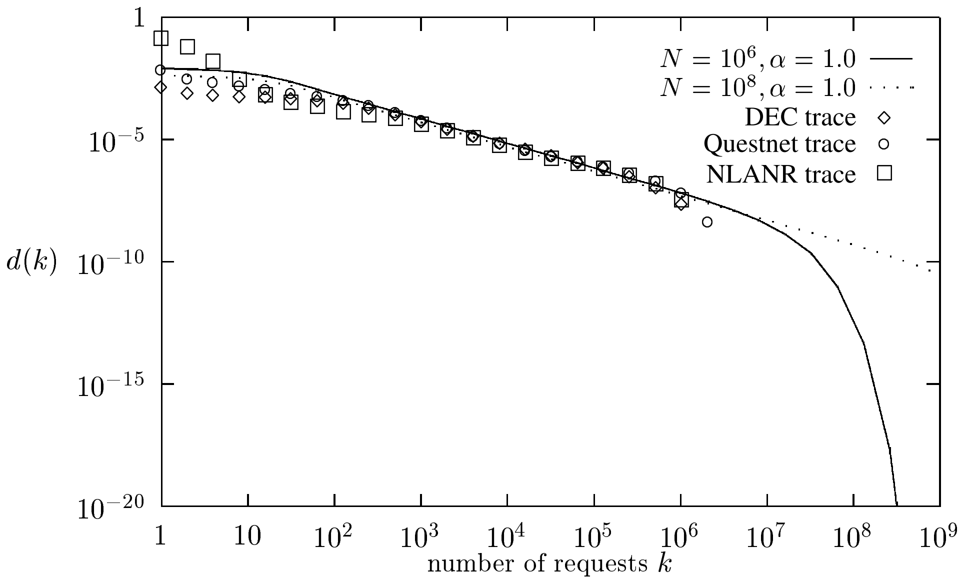

Figure 8: Probability distribution for page request interarrival times.

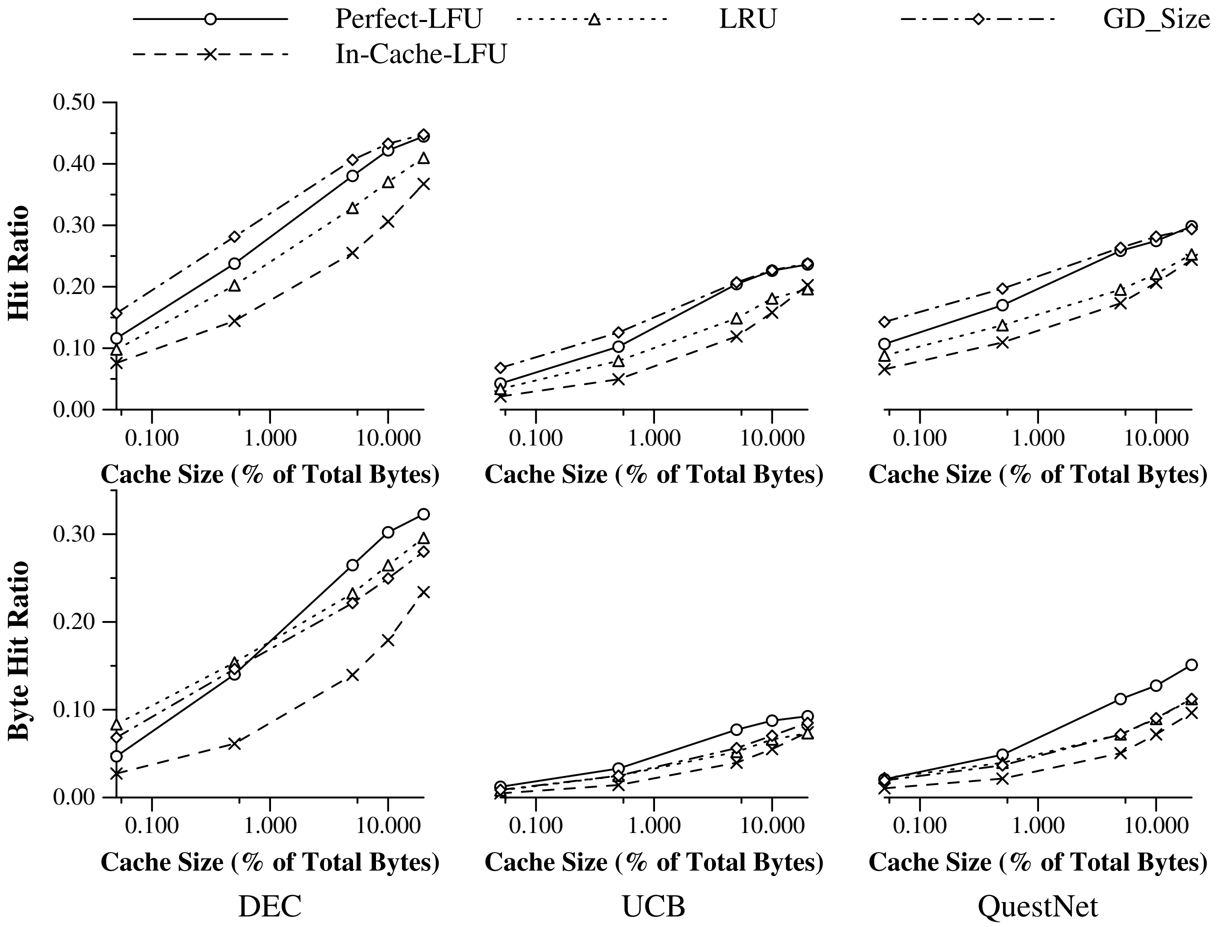

Figure 9: Hit-ratio and byte hit-ratio for four algorithms for a DEC, UCB and QuestNet traces. Results on other traces are similar.

system caching community that if (i) the requests are independent and have a fixed probability and (ii) the pages have the same size, then the optimal replacement algorithm is to keep those pages with the highest probabilities in the cache [GCD73]. In practice, an online algorithm detects those documents by keeping track of the number of references to each document and keeping those documents with the highest reference count in the cache. In other words, the best online algorithm under the independent reference model is the Least-Frequently-Used (LFU) algorithm.

However, there are two versions of LFU that are often confused in the literature: In-Cache-LFU, and Perfect-LFU. To make a clear distinction between the two policies, Perfect-LFU remembers page counters even when a page is evicted from the cache, while In-Cache-LFU removes the page counter together with the evicted page. Clearly, Perfect-LFU incurs more overhead and is less practical than In-Cache-LFU. However, Perfect-LFU is a better realization of the optimal cache replacement algorithm under the model.

The question we wish to answer is: which of the four replacement algorithms—Perfect-LFU, In-Cache-LFU, LRU and GD-Size—performs the best in terms of hit-ratio? Note that LRU is the most widely-used Web cache replacement algorithm, while GD-Size is a new algorithm that takes both document size and locality into account and was shown to outperform existing algorithms in terms of hit-ratio [CI97]. We answer the above question using trace-driven simulations.

Figure 9 shows the hit-ratios and byte hit ratios for the four replacement algorithms for the DEC, UCB and QuestNet traces. Results from other traces are consistent with those in the figure and are omitted due to space limitation.

Figure 9 shows that, as predicted by the independent reference model, Perfect-LFU performs best in terms of byte hit-ratios in most cases. GD-Size still performs best in terms of hit-ratios for small cache sizes because GD-Size takes document size into account (that is, small documents are favored over large ones) whereas the Perfect-LFU does not. Note that when cache sizes are large, Perfect-LFU performs comparably to GD-Size in hit-ratio and much better in byte hit-ratio. The figure also shows that In-Cache-LFU performs the worst and consequently is a poor choice for Web cache replacement algorithms.

The main drawback of Perfect-LFU is that it requires additional space to maintain the counts for documents that are evicted from the cache. In addition, it fails to take document size into account. We are currently designing an algorithm that has a bounded space requirement and considers both reference counts and document size.

However, our goal in this paper is not the design of specific replacement policies. Rather, we wish to observe that if temporal locality, due to strong correlations between references, was the dominant effect then LRU would have performed the best. However, because LRU did not perform as well as Perfect-LRU, it may suggest that the temporal locality effects may not be important when analyzing qualitative aspects of web caching performance.

However, for quantitative predictions of web caching performance, it may be important to consider correlations in the request stream. In particular, Almeida et al. compared plots of miss-ratio for a synthetic workload having a Zipf distribution against plots of miss-ratio for real workloads and concluded that the Zipf model was inaccurate because it did not capture the locality of reference in the real request stream [ABCdO96]. In this paper we looked at the asymptotic properties of the hit ratio and request inter-arrival times, and perhaps in the asymptotic regime the correlations focused on in [ABCdO96] are not so relevant.

5 Discussion

In this paper we first showed evidence that web requests follow a Zipf-like distribution. We then introduced a simple model for web requests, where requests are independent and distributed according to a Zipf-like distribution, and showed that this simple model can explain the asymptotic behavior for three properties that are observed in real Web cache traces. Our results suggest that these properties found in many studies are perhaps inherent to Web access streams and not an artifact of the traces studied.

Our results also suggest that a simple model for web requests may be sufficient to understand certain asymptotic properties of cache performance. Our model has many limitations. For instance, the model does not consider document modifications, nor does it consider the cache's replacement policy, both of which no doubt play an important role in a cache's performance. Nonetheless, we think the simple model presented here may worthwhile for studying the asymptotic behavior of various quantities. We are currently trying to improve the model and trying to understand its applicability to other proxy traces as well as its implications for Web cache design and configuration.

References

[ABCdO96] Virgilio Almeida, Azer Bestavros, Mark Crovella, and Adriana de Oliveira. Characterizing reference locality in the WWW. In *IEEE International Conference in Parallel and Distributed Information Systems*, Miami Beach, Florida, USA, December 1996. http://www.cs.bu.edu/groups/oceans/papers/Home.html.

[ACC+98] Virgilio Augusto F. Almeida, Marcio Anthony G. Cesirio, Rodrigo Fonseca Canado, Wagner Meira Junior, and Cristina Duarte Murta. Analyzing the behavior of a proxy server in the light of regional and cultural issues, 1998. http://www.anades.dcc.ufmg.br/paperSubmetidos/webcache/cultural/.

[CBC95] Carlos Cunha, Azer Bestavros, and Mark Crovella. Characteristics of WWW client-based traces. Technical Report TR-95-010, Boston University, Computer Science Dept., Boston, MA 02215, USA, April 1995. http://www.cs.bu.edu/groups/oceans/papers/Home.html.

[CI97] Pei Cao and Sandy Irani. Cost-aware www proxy caching algorithms. In *Proceedings of the 1997 USENIX Symposium on Internet Technology and Systems*, pages 193–206, December 1997. http://www.cs.wisc.edu/~cao/publications.html.

[Den80] Peter J. Denning. Working sets past and present. *IEEE Transaction on Software Engineering*, SE-6(1):64–84, January 1980.

[DMF97] Bradley M. Duska, David Marwood, and Michael J. Feeley. The measured access characteristics of world-wide-web client pr oxy caches. In *Proceedings of USENIX Symposium on Internet Technology and Systems*, December 1997.

[GB97a] Steven Gribble and Eric Brewer. System design issues for internet middleware services: Deductions from a large client trace. In *Proceedings of the 1997 Usenix Symposium on Internet Technologies and Systems*, Monterey, California, USA, December 1997. http://HTTP.CS.Berkeley.EDU/˜gribble/.

[GB97b] Steven Gribble and Eric Brewer. UCB home IP HTTP traces, June 1997. http://www.cs.berkeley.edu/ gribble/traces/index.html.

[GCD73] Edward G. Coffman, Jr. and Peter J. Denning. *Operating Systems Theory*. Prentice-Hall, Inc., 1973.

[Gla94] Steven Glassman. A caching relay for the world wide web. In *First International Conference on the World-Wide Web*, CERN, Geneva, Switzerland, May 1994. http://www1.cern.ch/WWW94/PrelimProcs.html.

[GS96] James Gwertzman and Margo Seltzer. World-wide web cache consistency. In *Proceedings of the 1996 USENIX Technical Conference, San Diego, CA*, January 1996.

[Jar98] Pekka Jarvelainen, 1998. mailto:pj@csc.fi.

[KLM97] Thomas M. Kroeger, Darrell D. E. Long, and Jeffrey C. Mogul. Exploring the bounds of web latency reduction from caching and prefetching. In *Proceedings of USENIX Symposium on Internet Technology and Systems*, December 1997.

[KMM96] T. M. Kroeger, J. Mogul, and C. Maltzahn. Digital's web proxy traces, August 1996. ftp://ftp.digital.com/pub/DEC/traces/proxy/webtraces.html.

[NHo+98] Norifumi Nishikawa, Takafumi Hosokawa, Yasuhide M ori, Kenichi Yoshidab, and Hiroshi Tsujia. Memory-based architecture for distributed www caching p roxy. In *The 7th WWW Conference*, April 1998. http://www7.conf.au/programme/fullpapers/1928/com1928.h tm.

[oANR97] National Lab of Applied Network Research. Sanitized access log, July 1997. ftp://ircache.nlanr.net/Traces/.

[Riz97] Luigi Rizzo. Web proxy traces, May 1997. http://info.iet.unipi.it/ luigi/proxy-traces/.

[RV98] Luigi Rizzo and Lorenzo Vicisano. Replacement policies for a proxy cache. Technical Report RN/98/13, University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK, 1998. http://www.iet.unipi.it/˜luigi/caching.ps.gz.

[WAS+96] Stephen Williams, Marc Abrams, Charles Standridge, Ghaleb Abdulla, and Edward Fox. Removal policies in network caches for WWW documents. In *ACM SIGCOMM '96*, August 1996. http://www.acm.org/sigcomm/sigcomm96/program.html.

[Wee98] Julianne Weekers, January 1998. mailto:julianne@cc.uq.edu.au.

[Zip29] George Kingsley Zipf. *Relative frequency as a determinant of phonetic change*. Reprinted from the Harvard Studies in Classical Philiology, Volume XL, 1929.