

The Content and Access Dynamics of a Busy Web Server

Abstract

In this paper, we present a detailed study of the dynamics of a busy Web server, at present one of the largest in the Internet. We analyze the dynamics of both the server content and the client accesses made to the server. The former includes a characterization of the content creation and modification process while the latter considers page popularity, and the temporal stability and spatial locality in client accesses. Some of our key results are: (a) file modification, though more frequent than file creation, often results in little change in the file being modified, (b) a small set of files tends to get modified repeatedly, (c) file popularity follows a Zipf-like distribution with a parameter α that is much larger than reported in previous, proxy-based studies, and (d) there is significant temporal stability in file popularity but not much stability in the domains from which clients access the popular content. We point out several broad implications of these findings.

1 Introduction

The rapid growth of the World Wide Web has spawned several research efforts aimed at characterizing the Web workload and traffic. Such a characterization is vital in many ways. It enables identification of existing or potential bottlenecks, and the design and the evaluation of new algorithms to improve the Web.

The Web basically consists of servers, clients, and proxies. The servers are typically the originators of content while the clients are the consumers of content. Proxies, when present, mediate the communication between a set of clients and a subset/all of the servers. As such, each of these three components provides a unique perspective on the functioning

of the Web. However, by far the majority of Web characterization studies have focused on data gathered at a proxy host (or by a network packet sniffer placed at a location where a proxy host might have been). Proxy-based studies are useful for many purposes: the design and evaluation of caching and prefetching algorithms, the characterization of server popularity, etc. However, proxy-based studies have their limitations because they offer only a limited perspective on the goings-on at clients and at servers, i.e., a proxy is not in a position to observe *all* of the events occurring either at clients (e.g., scrolling up and down in a browser window) or at servers (e.g., the servers' communication with clients that do not connect via the proxy).

A significant difficulty that researchers face in doing a server-based study is the very limited availability of server traces. The few pioneering studies in this area [2] [3] have had to make do with data from relatively small departmental servers, typically at universities. While they have certainly been valuable, the main limitation of these studies is that the bulk of Web traffic is served out by large commercial servers. It is unclear how well inferences drawn from the study of a small departmental server would scale to the large commercial sites in the real world.

We have been fortunate to have obtained access to detailed traces from a large commercial server site, which, to preserve anonymity, we will refer to simply as *FooBar*. *FooBar* is a large commercial site in the same category as the likes of CNN [7], MSNBC [21], and ABCNews [1], and is consistently ranked among the top sites in the Web [20]. The trace data we obtained was of two kinds: (a) *content logs*, which record file creation and modification events, and also include copies of successive versions of (a subset of the) files, and

(b) *access logs*, which record client accesses to the HTML content (but not to the inline images). Thus, we are in a position to study (certain aspects of) both the *back-end* (i.e., content dynamics) and the *front-end* (i.e., access dynamics) of the FooBar site.

A detailed discussion of the results appears later in the paper, but here are some of our more interesting findings: (a) file modification, although more frequent than file creation, often tends to change little in the file being modified, (b) a small subset of the files tends to get modified repeatedly, (c) file popularity follows a Zipf-like distribution with a parameter α that is much larger than reported in previous, proxy-based studies, and (d) there is significant temporal stability in file popularity but not much stability in the domains from which clients request the popular content.

A limitation of our study is that it is difficult to determine how well the results derived from the FooBar site generalize to the other large sites in the Internet. Nevertheless, we believe our study is a valuable step towards characterizing the workload of large, commercial servers.

The rest of this paper is organized as follows. In Section 2, we survey previous work. We present a discussion of the architecture of the FooBar site, our trace collection methodology, and the traces themselves in Section 3. Then in Sections 4 and 5, we present a detailed analysis of the content logs and the access logs, respectively. In Section 6, we summarize our key results and discuss the broader implications of our findings. Finally, in Section 7, we touch upon ongoing and future work.

2 Previous Work

As discussed above, much of the work thus far in Web workload characterization has focussed on proxies, often with a view to evaluating the effectiveness of proxy caching. The hit rate of proxy caches have been found to be quite low, often not much higher than 50% [9] [11]. A substantial fraction of the misses arise from *first-time* accesses to files (i.e., compulsory misses). Proxy logs have also been used to study the

effectiveness of cooperative caching. In this context, a recent study [26] [27] reports that the organizational membership of clients is significant in that clients belonging to the same organization are more likely to request the same documents than clients picked at random. Our analysis of spatial locality (Section 5.5) shows this significance can be diminished, for instance, by the occurrence of a "hot" news event that is popular globally, across organizational boundaries.

The relative popularity of Web pages accessed via a proxy has also been studied extensively. The almost universal consensus is that page popularity follows a Zipf-like distribution where the popularity of the i th most popular file is proportional to $1/i^\alpha$. The value of α is typically less than 1, which makes popularity distribution as seen by the proxy rather flat (e.g., [5] reports that it takes 25-40% of pages to draw 70% of the client accesses). Our results (Section 5.2) show that while the Zipf-like distribution holds for the FooBar server site as well, α tends to be much larger, typically 1.4-1.6.

Proxy logs have also been used to study the rate of change and age distribution of files (e.g., [8]). These have been deduced indirectly using the last-modified timestamp in the HTTP response header, which opens up the possibility of missed updates. In contrast, we use file modification logs obtained directly from the FooBar site back-end in our study, so the possibility of missed updates is diminished/eliminated.

Server-based studies are far fewer in number. A few of these have focused primarily on the network dynamics of busy Web servers [18] [4]. These are interesting but orthogonal to the focus of this paper.

A few of the server-based studies have been along lines similar to this paper. [2] studied access logs from a set of Web servers, the busiest of which saw under 50000 accesses in a day. They showed that file popularity followed Zipf's distribution (i.e., Zipf-like with $\alpha = 1$). They also demonstrated the presence of temporal locality in file accesses. [3] studied various aspects of server behavior using data from a set of servers. They reported that 10% of the files accessed accounted for 90% of the server requests, and that 10% of the (client) domains accounted for over 75% of the server re-

quests. However, the busiest of the servers they studied only saw a *total* of around 350,000 accesses in a day. In contrast, the FooBar server cluster sees, on average, over 25 million accesses each day to its HTML content alone (image accesses, which are not included in our traces, would increase this number significantly).

In summary, we see our study of the FooBar site as complementing the existing body of literature on Web workload and traffic characterization. We believe both the size of the dataset we have analyzed and the use of back-end logs to characterize the content dynamics make our study valuable.

3 Experimental Setup and Methodology

In this section, we briefly describe the essential aspects of the FooBar server site, and discuss the trace data that we gathered and processed.

3.1 Server Site Architecture

The FooBar server site comprises a cluster of over 40 server nodes, each running the Microsoft Internet Information Server (IIS) [17]. These server nodes are grouped together into sub-clusters containing approximately 6 nodes each. Load balancing is done at two levels. First, each sub-cluster is assigned a *virtual IP* (VIP) addresses and DNS round-robin cycles through the 6 VIP addresses for the entire FooBar site. Second, within each sub-cluster that shares a VIP address, the Windows Load Balancing Service (WLBS) [17] is used to spread load evenly across the nodes. The main point to take away is that at different times, a particular client’s request may be served by any of the 40 nodes.

3.2 Server Access Logs

Each server node maintains a standard HTTP access log that records several pieces of information for each client access: a timestamp, the client’s IP address, the URL accessed, the response size, the server status code, etc. For administrative

reasons, the server site operator chose to turn off logging for image accesses. So the logs only record accesses to HTML content. While the absence of image access logs is a limitation of our data set, we do not believe it interferes with our study in any significant way since our analysis of access dynamics focuses on *Web pages* (as defined by the HTML content) rather than on the individual files.

Table 1 summarizes the overall statistics of the server access logs we used in our study. For our analysis, we picked traces from several different periods, each spanning a few consecutive days. In some periods, we had only an hour’s worth of traces per day, while in others we had traces from the entire day. The traces on 12/17/98 and 12/18/98¹ were especially interesting, because they correspond to a “hot” news event, namely the launching of *Operation Desert Fox* by the US military against Iraq.

The FooBar server site saw, on average, over 25 million client accesses for its HTML content alone (image hits were over and above this). Due to disk and memory limitations, we only used logs from 9 or 12 (i.e., 22.5–30%) of the server nodes out of the cluster of 40 when analyzing the logs from a whole day period. Since requests are randomly dispatched to the server nodes, considering only a subset of the server nodes is unlikely to significantly bias or otherwise impact the results of our analysis.

In some of our analyses, we clustered clients together into *domains*, which we defined to be all but the hostname part of the clients’ DNS names (e.g., the domain for foo.bar.com is bar.com). We determined the DNS name of a host via reverse DNS lookup on its IP address. We had a fairly high success rate — typically over 70% — as reported in Table 1. For the analyses that involved domain information, we ignored clients for which the reverse DNS lookup failed. We realize that our definition of a domain is simplistic, and are currently looking at more sophisticated alternatives that also consider network topology.

¹Throughout this paper, dates appear in the format *month/day/year*.

	12/17 - 12/18 (1998)	8/1 - 8/5	8/3 - 8/5	9/27 - 10/1	10/7 - 10/11	10/14 - 10/18
Period	9 AM - 12 AM	10 - 11 AM	9 AM - 12 AM	all	all day	all day
% total server logs used	100	100	100	30	22.50	22.50
HTTP Requests	10413866	7061572	14183820	38191388	28102751	30000981
Objects	34443	30199	35359	57053	60323	52429
Clients	253484	440151	656449	1948609	1831027	1938437
% Domain discovered ²	58.587	-	76.227	78.967	78.344	-
Domains	41025	-	75569	117615	396528	-
% GET	99.818	99.050	99.065	99.065	99.008	99.059
% POST	0.088	0.448	0.464	0.474	0.512	0.475
% HEAD	0.082	0.406	0.392	0.327	0.350	0.336
% Other methods	0.012	0.096	0.079	0.134	0.130	0.130
% status=200	58.084	55.913	57.104	56.195	55.088	54.744
% status=302	4.554	15.017	15.267	17.647	16.047	18.443
% status=304	36.946	27.529	26.231	23.812	26.517	24.501
% status=400	0.010	0.023	0.025	0.031	0.029	0.026
% status=403	0.003	0.024	0.018	0.013	0.015	0.018
% status=404	0.327	1.347	1.241	1.738	1.654	1.661
% status=500	0.012	0.089	0.070	0.131	0.125	0.126
% Other status	0.064	0.058	0.044	0.433	0.525	0.481

Table 1: Overall trace statistics

3.3 Content Creation and Modification Logs

The back-end of the FooBar site uses the Microsoft Content Replication System (CRS) [17] to replicate content from a staging server to each of the 40 server nodes. Each replication event is logged, specifying the time of replication and the file being replicated together with its size. However, all that a CRS log entry says is that a file was replicated, so by itself it does not enable us to determine whether a new file was created or an existing one was updated. We disambiguate between file creation and modification by using several days' worth of CRS logs to prime our list of files that already exist, and thereafter (once the spike in the number of file "creation" events has subsided) treating CRS replication events for files not seen before as file creations³. The CRS system did not log file deletions, although, in general, it could have. The CRS logs we analyzed corresponded to the 4-week period from 10/1/99 through 10/28/99.

3.4 Content Logs

For a small fraction of the content hosted by FooBar, we were able to obtain a log of HTML content itself, i.e., successive versions of the files as and when they were modified. A new version of the file was logged on the hour if the file had been modified (at least once) in the past hour. The subset of files logged in this manner was determined by the server site operator⁴. The content log, although limited in scope, enables us to get some insight into the evolution of files as they undergo modification.

3.5 Proxy Logs

In some of our analyses (Section 5.2), we compare characteristics of the FooBar server access log with that of a busy caching proxy. We obtained logs from a proxy cache that serves a large campus population with over 50000 client

³As a validation of this heuristic, we confirmed that the number of file creation events beyond the priming period does not diminish with time, as would have happened had the file creation events been "bogus".

⁴One of the reasons they generated the content log was to feed it into various search engines for re-indexing the corresponding pages.

hosts. The logs were gathered over a 2-day period — 10/6/99 and 10/7/99.

4 Server Content Dynamics

In this section, we analyze the dynamics of file creation and modification.

4.1 File Creation and Modification Processes

We studied the dynamics of file creation and modification using information derived from the CRS logs. Figure 1(a) shows the number of file creation and modification events (computed hourly) over a one-week period (midnight Saturday, 10/9/99 through midnight Friday, 10/15/99 local time). Not surprisingly there is a clear diurnal cycle in the file creation and modification process. There is a trough at nighttime, and several peaks during the daytime. We believe the reason for these peaks is that even if the actual content generation/modification process is spread out uniformly, the CRS replication process is only scheduled to run from time to time (i.e., it replicates a bunch of files at a time rather than individual files; the only exception is a "hot" file that needs to be updated immediately). The time of replication is what really matters because only beyond it does the new content become available to clients.

A weekly cycle is also evident from Figure 1(a). There tend to be fewer events on the weekend (the first two days shown in the figure) than during the week.

In Table 2, we tabulate a more detailed breakdown of the event counts. We note there are nearly four times as many file modification events as creation events during the course of the week. A closer examination of the modification events reveals that they tend to be concentrated on a small number of files. On average, there were around 10 modification events per file (only considering files that were modified at least once during the week).

The large numbers of creation and modification events have broad implications. The creation of new files poses

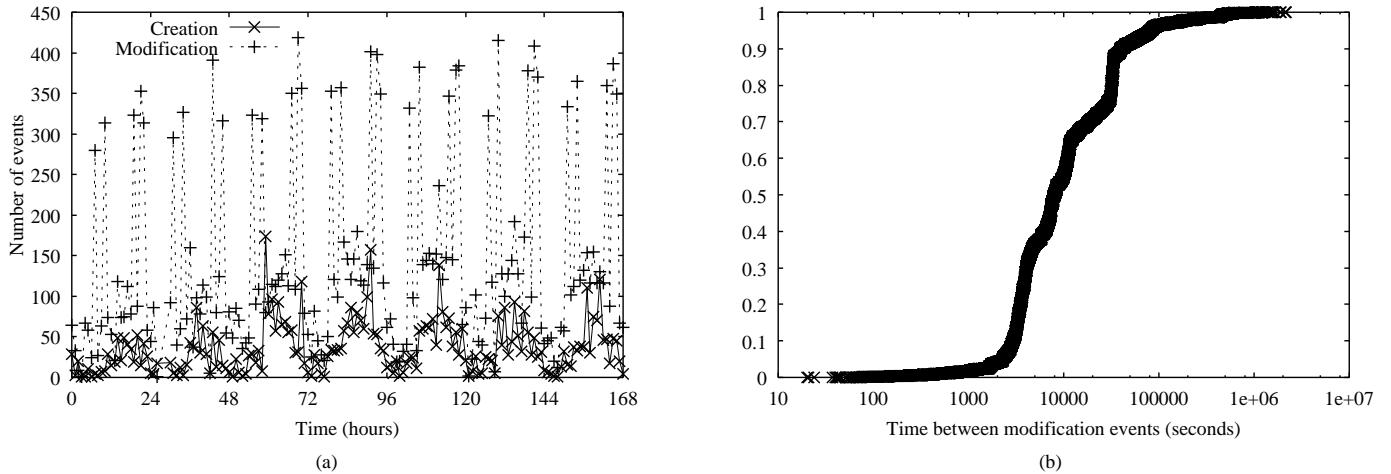


Figure 1: (a) An hourly count of file creation and modification events over a one-week period (Saturday through Friday). (b) CDF of the time interval between successive modifications of a file (conditioned on the file being modified).

a challenge to latency/bandwidth saving schemes such as prefetching [22] [10] or server-initiated “push” [24]. These schemes depend on the past history of accesses to the file. But for a newly created file, there exists no such history. The large number of modification events suggests that it may be worthwhile to deploy efficient protocol mechanisms for the validation/invalidation of files cached at proxies and/or clients (e.g., citeCKR98 [23]).

Table 2 also reveals that around 1% of the modification events corresponded to GIF/JPEG image files. Intuitively, we would not expect images to get modified very much; instead we would expect new images to be assigned new names. On closer examination, we discovered that the images being modified were almost exclusively maps, primarily weather maps but also maps of other kinds (e.g., a weekly “health” map indicating the current status of disease outbreak in the country).

4.2 Distribution of Modification Intervals

Next, we turn to Figure 1(b), which shows the cumulative distribution function (CDF) of the time duration between two successive modifications (i.e., the *modification interval*) of a file. (Note that we only considered files that were modified during the 4-week period of our CRS logs. In other words, the

CDF is conditioned on the file being modified during the 4-week period.) The CDF exhibits two distinct knees. The first is around the 5% level and occurs at a modification interval of around 3000 seconds (about 1 hour). The second is around the 95% level and occurs at a modification interval of around 80000 seconds (approximately 1 day). Both of these observations are in agreement with our intuition. By default, the CRS replication process is scheduled to run approximately once an hour. More frequent updates happen only when there are “hot” news events, which is typically an infrequent occurrence. Hence there are only a small number of instances of modification happening in within an hour of the previous one. At the other end of the spectrum, a day seems like a natural “upper bound” for the update period. Of course, it is certainly possible for files to be modified on longer timescales such as weekly, monthly, etc. (or even aperiodically), but this is likely to be an infrequent occurrence.

4.3 Implications of Modification History

We now turn to examining the relationship between successive modification intervals of a file. The motivation is to determine whether the modification dynamics of a file in the past is a good predictor of the future. This finding would have significant implications for Web cache consistency mecha-

Creation	Modification	Unique Files Modified	GIF/JPEG Modification
6007	23343	2453	287

Table 2: Count of various events during the week from 10/9/99 through 10/15/99.

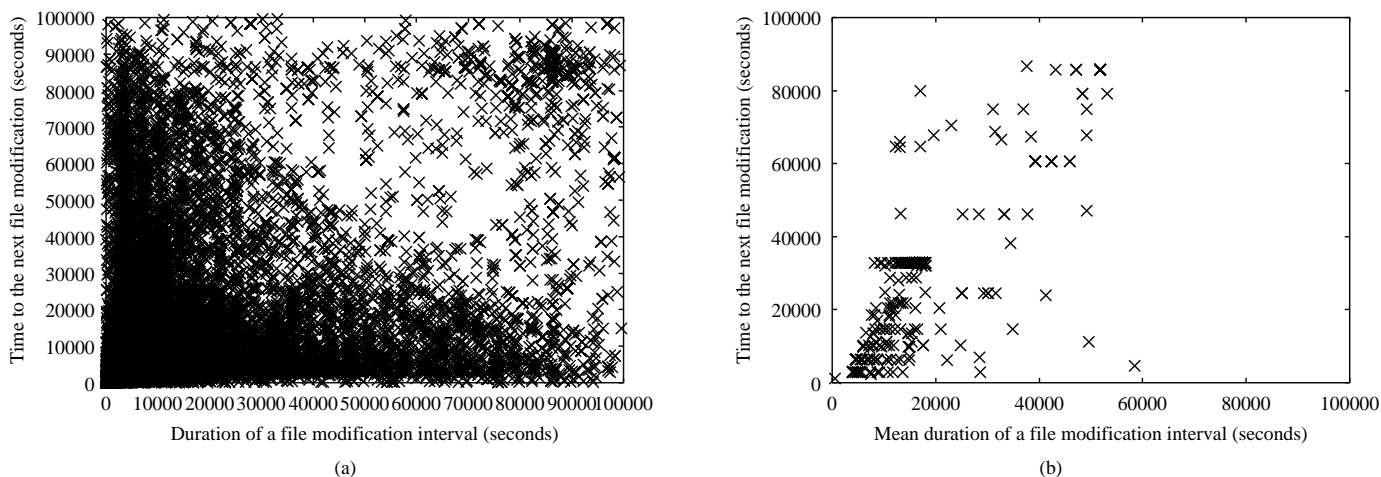


Figure 2: (a) A scatter plot where a point (x, y) represents two successive modification intervals for a file, the first of duration x and the second of duration y . (b) A scatter plot of points (x, y) where x is the mean duration of the modification interval for a file (computed using past history), and y is the duration of a new modification interval for that file.

nisms such as adaptive TTL [13].

Figure 2(a) shows a scatter plot of pairs of successive modification intervals for a file. If duration of the previous modification interval were indicative of the next one, we would have expected the points to be clustered along a positively-sloped line. It is clear from the figure that this is not the case. There are several instances where a short modification interval for a file is followed immediately by a long one, and vice versa. Indeed, the coefficient of correlation [14] is only mildly positive — 0.23.

We then explored the possibility of using a longer history than just the duration of the previous modification interval. We divided the 4-week duration of our CRS logs into roughly two halves. Using the data from the first half, we determined the mean modification interval for each file. In the end, we only retained the set of files for which the mean was computed over 10 or more samples (i.e., there was “sufficient” modification history). Then in using the second half of the log data, we determined the length of the first modification inter-

val for each of the files in the set. The question we wanted to answer was how good a predictor the mean was of the duration of the new modification interval. Figure 2(b) shows the scatter plot of these two quantities. We observe that there is a fairly strong positive correlation between the mean duration of the modification interval from the past and the duration of the new interval. The value of the coefficient of correlation — 0.79 — confirms this.

Thus, we believe it may be appropriate to use the modification interval from the past as a predictor of the future so long as a sufficient number of samples from the past are averaged.

4.4 Extent of Change upon File Modification

We now examine more closely just how much a file changes when it is modified. First, we compute how much the file size changes when it is modified. Figure 3(a) shows the CDF for the change in file size (unsigned magnitude) in terms of bytes. We observe that there is little change in file size despite

the modification. Over 70% of the modifications cause less than a 1% change in the file size (due to space limitations, we have not shown the graph of file size change expressed as a percentage). While it is certainly possible for the file size to remain virtually unchanged but for the content to change significantly, intuitively it seems too coincidental to be likely. The analysis we present next sheds more light on this.

We used the content logs (which, as described in Section 3.4, were available for a subset of the HTML content) to explore more carefully how much a file changes upon modification. Since we were examining just HTML content, we decided to focus on the *visible* textual content, i.e., text that would be displayed by a client browser. To this end, we extracted the visible text by stripping out the HTML tags from each version of a file. We then quantified how similar two successive versions of a file were using the *cosine similarity metric* [12]. For each document, a term (i.e., word) frequency vector is constructed. The cosine similarity metric is then just the inner-product of the two vectors divided by the product of their lengths, i.e., the cosine of the angle between the vectors. The more similar the files are, the closer to 1 the metric is. The similarity metric is simple in that it only considers word frequency without caring about the order in which the words appear. To prevent "insignificant" words such as "and", "the", etc. from overwhelming the more significant ones, we created a list of common words to disregard in the similarity computation.

Figure 3(b) shows the CDF of the text similarity metric computed over successive versions of files. We see that the metric tends to be close to 1, which indicates that little changes in terms of the visible textual content. On closer examination, we found (at least) a couple of common modes of (minor) textual change: (a) a date/time string contained within the TITLE HTML tags is updated from time to time, and (b) the links to related pages are updated while leaving the bulk of the page unchanged.

In conclusion, our analysis suggests that successive versions of files are very similar, both in terms of size and in terms of content. This implies that techniques such as delta-

encoding [19] would be very useful.

5 Server Access Dynamics

In this section, we discuss the dynamics of the client accesses made to the FooBar server site.

5.1 Type of Access

As shown in Table 1, over 99% requests employ the GET method. Moreover, among all the requests, the ones with HTTP response status code 200 (action successful) account for 55%. Around 15% to 18.5% of the requests have response status code 302 (moved temporarily), and around 23% to 37% have response status code 304 (not modified). The latter implies that the potential benefit of an efficient Web objects consistency protocol could be large. Unless otherwise specified, in our analysis of the access dynamics, we do not distinguish among different HTTP methods (i.e., GET, HEAD, and POST) and different status code. Instead we treat them equally as Web accesses.

5.2 The Applicability of Zipf's Law to Web Requests

Several prior studies have investigated the application of Zipf's law to web accesses and arrived at different conclusions. [5] gives a comprehensive summary of previous work on this issue. Their studies shows the distribution of web requests from a fixed group of users follows a Zipf-like distribution, C/i^α , very well. The value of α varies from trace to trace, ranging from 0.64 to 0.83.

We investigate this issue further by studying access logs from both the server and the proxy. Due to space limitations, we show the graph of document accesses versus document ranking only for the server traces (Figure 4).

We make the following observations:

- The curves for both the server traces and the proxy traces (not shown) fit a straight line reasonably well. The straight line on the log-log scale implies that the




Figure 3: (a) The CDF of change in file size (in bytes) between successive versions. (b) The CDF of the text similarity metric computed over successive versions of files.

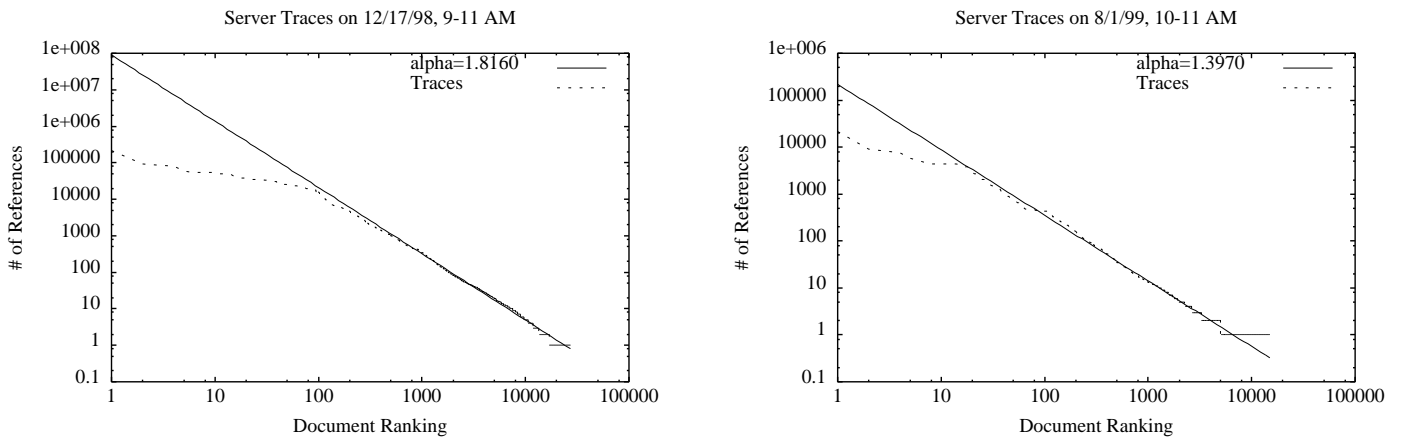


Figure 4: Frequency of document accesses versus document ranking (server traces)

request frequency is proportional to $1/i^\alpha$. The values of α are obtained using least square fitting, excluding the top 100 documents (as in [5]), and also excluding the flat tail.

- The value of α varies from trace to trace. The values of α in the server traces are consistently and significantly higher than those in the proxy traces.

More specifically, The values of α in the server traces are mostly around 1.4 - 1.6, with lowest being 1.3970 and highest being 1.816. In comparison, the α in the proxy traces are much lower, around 1.0 for both the days we had proxy logs for.

We think that it is not coincidental that the traces collected at a popular web site has higher α than proxy trace, but is likely to be a common phenomenon. This can be explained as follows: A number of previous studies [5, 26] show popular web pages are spread almost evenly across hot web servers. This implies that the proxy aggregating requests to a number of popular servers should have a slower decay in popularity than any individual server. This can be illustrated by the following simplified example:

Suppose a proxy accesses s web servers. The web server i has n_i documents, with decay coefficient α_i . In addition, it receives A_i accesses to its most popular web page. Then according to the Zipf-like distribution, we have

$$\log(A_i) = \alpha_i * \log(n_i)$$

The decay coefficient at the proxy can be computed as follows:

$$\alpha_{proxy} = \frac{\log(\max(A_1, A_2, \dots, A_s))}{\log(n_1 + n_2 + \dots + n_s)}$$

For a simple case when $A_1 = A_2 = \dots = A_s = A$, $n_1 = n_2 = \dots = n_s = n$, and $\alpha_1 = \alpha_2 = \dots = \alpha_s = \alpha_{server}$, we have

$$\alpha_{proxy} = \frac{\log(A)}{\log(s * n)} = \alpha_{server} * \frac{\log(n)}{\log(s * n)} < \alpha_{server}$$

The exact difference between α_{proxy} and α_{server} depends on s (the total number of popular servers). Of

course, things are much more complicated in practice: (i) the Web accesses does not strictly follow Zipf like distribution especially for the most popular documents; (ii) Web servers are very heterogeneous. It is very hard to compute α exactly. On the other hand, as long as it is true that the popular web pages are spread almost evenly across hot web servers, the α at the proxy is lower than at each individual popular web server.

- The value of α is highest on 12/17/98, when there was an unusual global event interesting to people all over the world. This is what we would expect. Since during such period users all over the world are interested in a small set of pages related to the event, making hot documents extremely hot and cold documents colder than usual, the difference in the hits count of hot pages and cold pages are thus enlarged, which contributes to a larger α value.

To further study the impact of different α values, we also plot the cumulative distribution of requests to popular documents. Figure 5 shows the cumulative probability of access for the top $r\%$ of documents for the server traces on 12/17/98 (the highest α) and the proxy traces on 10/6/99. As we can see, the top 2% documents account for 90% accesses in the server traces. In contrast, it takes 36% to 39% documents to account for 90% accesses in the proxy traces. So, as [5] discovered, 10/90 rule (i.e. 90% accesses go to 10% documents) does not apply for the web accesses seen at the proxies. On the other hand, according to our server traces, the web accesses observed at the server side are sometimes even more concentrated than the 10/90 rule.

In summary, we observe the access patterns at both the server and the proxy exhibit Zipf-like distribution. The value of α is much higher for the server accesses than for the proxy. In some cases, the top 2% documents account for 90% accesses. In contrast, the accesses seen at the proxy traces is more heterogeneous, which takes up to 40% to account for 90% accesses.

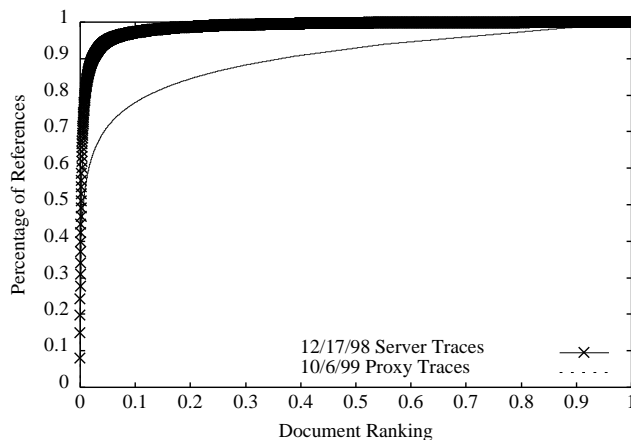


Figure 5: Cumulative distribution of requests to documents

5.3 Access Pattern at Lower-level Proxies vs. Higher-level Proxies

In the previous section, we have shown the access pattern seen at the Web server has Zipf like distribution, with $\alpha > 1$ for server traces and $\alpha \approx 1$ for proxy traces. Given the increasing widespread deployment of Web caching both at the edges and in the core of network, an interesting question arises how such deployment affects the web access pattern. In particular, we want to answer the following questions: (i) given the access pattern of the lower level proxies (or end-users), what is the access pattern at the higher- level proxies (HP) look like? (ii) If the access pattern at the lower level proxies exhibits a Zipf-like distribution, will the access pattern still have a Zipf-like distribution at the higher level? If so, what will the value of α at the higher level proxies be? Answers to these questions depend on how lower-level proxies are assigned to higher-level proxies (By assigning a lower-level proxy A to a higher-level proxy B , we mean whenever there is a cache miss at the lower-level proxy A , the request will be forwarded to the higher-level proxy B for service.) To make our analysis simple, here we consider random assignments from lower-level proxies to higher-level proxies.

5.3.1 Analysis

We formulate the problem as follows:

Suppose without hierarchical caching, a page receives m number of accesses. When we employ a two-level caching hierarchy with x number of higher level proxies, shown in Figure 6, what is the average number of higher level proxies, z , that access the page?

To simplify our analysis, we ignore document invalidation, and also assume infinite cache size at all the proxies (both higher-level and lower-level proxies). Therefore multiple requests for a single document will be forwarded up towards the root only once by a higher-level proxy cache (basically, upon receipt of the request for the time from any of its children).

Based on these assumptions, we can derive

$$z = \frac{x * m}{x + m - 1}$$

The detailed analysis is shown in Appendix A.

Using the above result, we know if the i th popular page has C/i^α accesses from the lower level proxies, then it will have $C/(i^\alpha + \frac{C}{x})$ accesses from the higher level proxies (ignoring the -1 in the denominator since $x + m$ is likely to be much larger). When $\frac{C}{x}$ is much smaller than i^α , equivalently when i is large enough, then the access pattern at the higher level proxies looks very similar to the access pattern at the




Figure 6: Two-level caching hierarchy

lower level. So the α values in both cases are close to each other for large i .

5.3.2 Validation using Traces

We use our server traces to validate the above analysis. Specifically, we consider the requests recorded in the server traces are from the lower-level proxies. (Note that in the context of the server traces, the "lower-level proxies" are just the client hosts identified in the traces. Some of these may have been real proxies while the remaining were just end-hosts; we were not in a position to tell one from the other.) We then insert a number of higher-level proxies between the lower-level proxies and the server, where lower-level proxies are assigned randomly to some higher-level proxy. Again, as in our analysis, we ignore the page invalidation, and assume infinite cache size at all the proxies (both higher-level and lower-level).

The graph in the left of Figure 7 shows the access pattern seen at the Web server when 2-level hierarchical caching is employed. As we can see, the curves with different numbers of higher-level proxies only differ in the initial portions. The fewer higher-level proxies, the longer the flat region is, where the flat region occurs when almost all the higher-level proxies access the document. This is intuitive, because given random assignment from lower-level proxies to higher-level proxies, as long as the popularity of a document is above a certain threshold, almost every higher-level proxy will request it. Furthermore, the threshold depends on the total number of higher-level proxies. The fewer higher-level proxies we have, the smaller the threshold is. As the documents' popularity decreases, the difference in the access pattern with different

number of higher-level proxies diminishes. As predicted in our previous analysis, the decay coefficients α with different number of higher-level proxies are very close to each other.

To further explore the accesses pattern at higher-level proxies, we also consider a more natural mapping from lower-level proxies to higher-level proxies wherein all lower-level proxies (i.e., clients) within a domain share a domain-wide higher-level proxy. In particular, we study the number of requesting domains for each web object. As shown in right of Figure 7, the curve on the log-log scale matches well with a straight line, which implies a Zipf-like distribution. However the value of α becomes a little lower: 1.5006, as opposed to 1.6511 when considering the number of references.

To summarize, we examine the access pattern at the higher-level proxies through both mathematical analysis and traces studies. Our results show that the Zipf-like distribution also holds at the higher-level proxies, and the value of α is likely to be a little lower than at the lower-level proxies.

5.4 Temporal Stability

In this section, we present our analysis of temporal stability of Web accesses. In particular, we are interested in answering the following key questions:

- How much does the web pages ranking vary with time? That is, do popular web pages on one day remain popular on the following days?
- For each web page, how much does the set of domains interested in it vary from one day to the next?

Answers to these questions are very critical for designing sensible prefetching or server-initiated push algorithms. Any




Figure 7: Frequency of document accesses versus document ranking when two-level caching hierarchy is employed

prefetching algorithms based on past history rely on some degree of stability (both in ranking and in the interest group (ie. the set of domains requesting the pages)). Our trace analysis helps to shed light on how well such reactive prefetching algorithms can perform in practice.

5.4.1 Stability of Web Pages Ranking

We study the stability of web pages ranking as follows: We use the web access logs of several consecutive days. For each day, we pick n most popular documents. We then compare the overlap in the most popular pages selected from one day to the next. Our results are illustrated in Figure 8. We make the following observations: First, the overlap is mostly over 60%, which essentially means many documents are hot on both days. Second, for the several consecutive days period, the overlap is mostly the same. For example, the amount of overlap between 8/1/99 vs 8/2/99 is quite close to that between 8/1/99 vs 8/5/99. This implies the web accesses in the same short time-frame (within 1 week period) are equally good for predicting which Web pages will be most popular in the near future. That is, last week trace is almost as useful as yesterday trace for predicting web pages ranking. The reason for this is that many of the very popular pages are *index* pages (such as the default front page for the FooBar site) that contain pointers to other pages. Third, the ranking stability tends

to decrease as the number of documents selected increases. This indicates that very hot documents are more likely to remain hot than moderate hot documents. Similar results are observed in other periods of traces.

We also study the ranking stability in a larger time window. Namely, we consider the overlap between the two days that are more widely separated. Our results are shown in Figure 9. As we would expect, the overlap decreases as the time interval gets longer. For example, the overlap between 12/17/98 and 10/18/99, which are 10 months apart, is considerably smaller, mostly below 20%. On the other hand, even for the two months separation (8/1/99 and 10/18/99), though the overlap is lower than one day separation, it is still quite significant. For the top 100 documents, the overlap is above 60%. However compared to the one day separation, the overlap in the two month separation decreases much faster as the number of documents selected increases.

We further explore the issue by breaking down the overlap and disjoint regions into four groups (assuming Day 1 is prior to Day 2):

- Common & unmodified: documents that are popular on both days and have not been modified since Day 1
- Common & modified: documents that are popular on both days, and have been modified since Day 1
- Different & old: documents that that are popular on

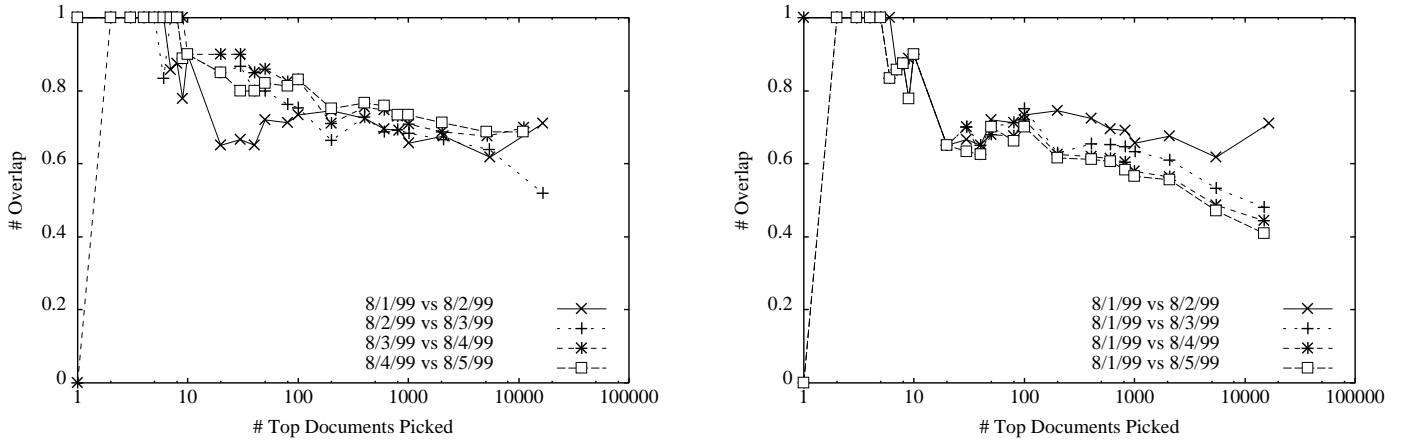


Figure 8: Stability of document popularity in several days period

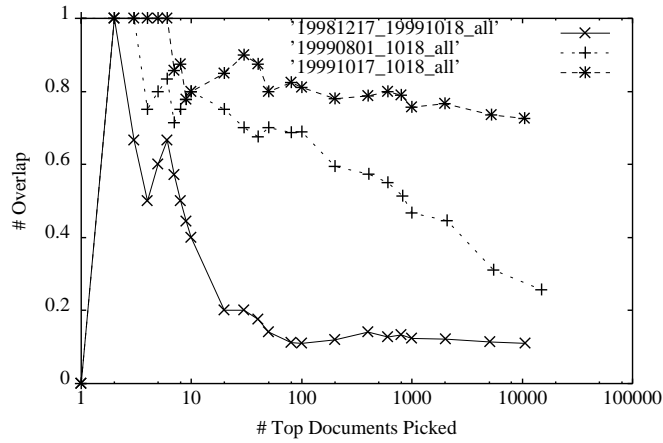


Figure 9: Stability of document popularity in large time window

only one of the days, but were in existence on both days

- Different & new: documents that were created after Day 1, and are popular only on Day 2

Our results are shown in Figure 10. We observe that many of the popular files in common between the two days are remain unmodified through both days. Moreover the disjoint region in the set of popular files on two days are mostly not due to the creation of new files, since most of the disjoint region consists of the accesses to Web pages that were in existence on both days. Even if the number of modified files is small, designing a good invalidation techniques for web documents is still very desirable because of the frequent updates made to the small set of files (as observed in Section 4.1).

To summarize, in this section, we study the ranking stability of web pages, and find the stability is reasonably high on the scale of days. The ranking tends to change only gradually over time. We also examine what contributes to the overlap and disjoint regions of popular documents on two days. Our results show the disjoint region mostly consists of old pages, not newly created pages. In the overlap region, the number of modified documents exceeds the number of unmodified documents, although both counts are of the same order of magnitude.

5.4.2 Stability of Interest Group

We now consider how the interest group for each web page changes over time. Our approach is as follows: For every web page that receives over 100 accesses on the earlier of the two days under our study, we find the set of domains which access the page on each day, and compare the overlap region. As mentioned in Section 3.2, we ignore requests from clients for which the reverse DNS lookup fails. Since the percentage of failure (Table 1) is reasonably low, it should not make a significant difference in our results.

Figure 11 shows the percentage of overlap region for several pairs of days we studied. As we can see, the overlap is not very large. Only a few documents that have over half of the domains making requests on both days. We observe similar results during other periods of traces as well. One explana-

tion for this may be that the interest groups for the documents do not stabilize within a day, possibly because our definition of domains is too fine-grained. Another explanation could be that employing proxy cache can reduce the likelihood of requesting the web pages multiple times. As part of our future work, we will further investigate this issue, and find out the reason that accounts for the large variation in the set of domains that request the web page.

5.5 Spatial Locality

In this section, we present our analysis of web traces with respect to spatial locality, i.e., the extent to which there is sharing of requests (i.e., accesses to the same object) within a domain (*intra-domain/local*), across domains (*inter-domain/global*), or both. Our analysis is similar in flavor to that in [26], although they analyzed proxy traces while we analyze server logs.

[26] examines the sharing of web documents from an organizational point of view. Their analysis is based on a one-week proxy trace at University of Washington (UW), USA taken in mid-May 1999. Their results show that organization membership appears to be significant. However the vast majority of the requests made are to objects that are shared among multiple organizations.

We explore the organization-based sharing using the web server traces as before. The server traces differ from proxy traces in the following ways: (i) The clients in the server traces are more heterogeneous. They are from all over the world. (ii) Requests seen in the server traces are for the objects at that server only, while in the proxy traces record requests made to different servers.

Figure 12 shows the distribution of clients, objects, and requests in different domains. As we can see, there is large variation in the domain size: the largest domain has over 10,000 clients making requests to the web server, which generate over 100,000 requests for around 10,000 objects. In contrast, some domains have only one client making only one request to the server. In the cumulative distribution plot, we see

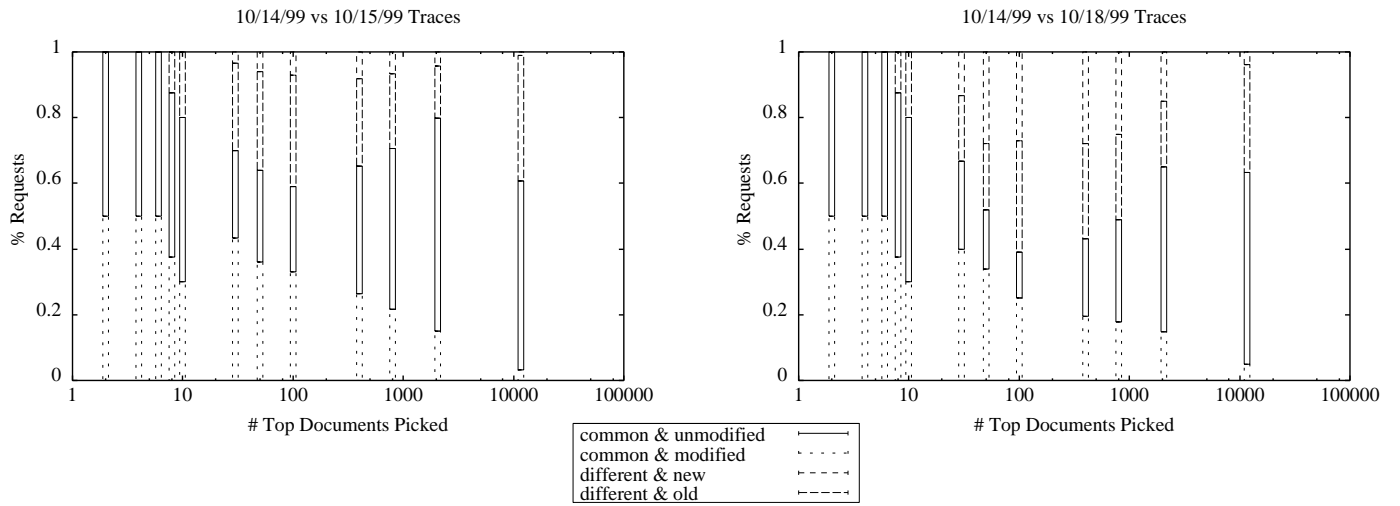


Figure 10: Stability of document popularity

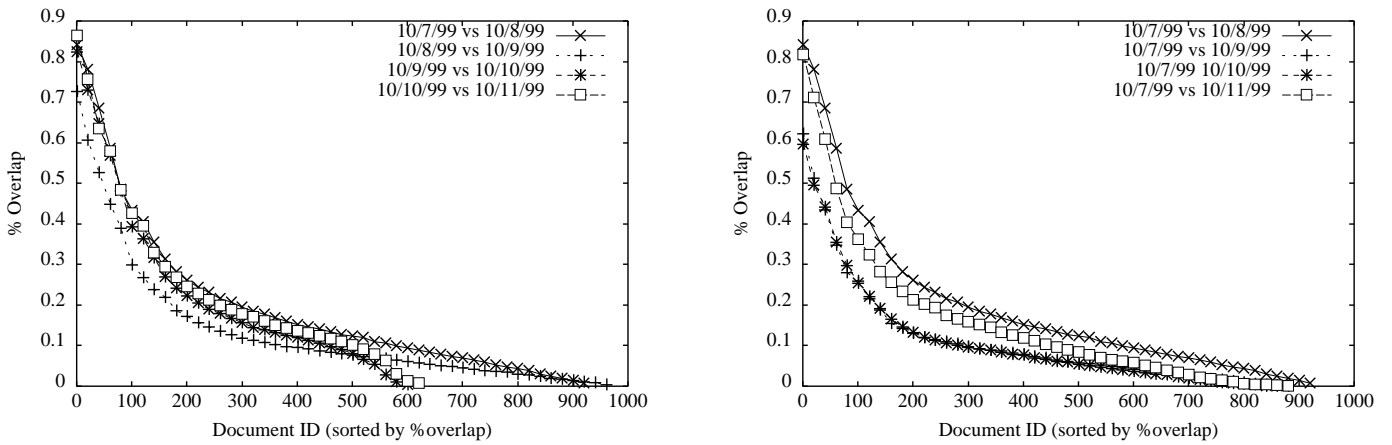


Figure 11: Stability of interest group

in the 12/17/98 trace, the top 10 domains account for 7.83% requests, the top 100 domains account for 18.63% requests, and the top 1000 domains account for 43.63% requests. In the 10/17/99 trace, the top 10 domains generate 11.86% requests; the top 100 domains generate 23.00% requests, and the top 1000 domains generate 44.43% requests.

Figure 13 shows the extent intra-domain (local) and inter-domain (global) sharing. For each domain (on the x-axis), the graphs show the percentage of all requests, and separately the percentage of all objects requested, by clients in that domain, that were shared. The curve marked "first shared requests" counts just the first request to a shared object as a percentage of all requests. As we can see, the inter-domain sharing is significantly higher than intra-domain, which means most objects are globally popular. Moreover compared to the proxy study in [26], the sharing in the server traces here is consistently higher. We believe this is not surprising. Since the proxy aggregates requests to different servers, the requests in the proxy traces are more heterogeneous, which possibly leads to a lower degree of sharing.

As in [26], we also categorize the sharing objects and requests based on the way in which they are shared: shared globally only, shared locally only, shared both globally and locally, and not shared. Table 3 shows the mean and variance of the percentage of objects and requests (computed over all the domains) that fall into each category. We observe that most objects are shared either globally only or both globally and locally. There are no objects that are only shared locally. This is not surprising, because it is very unlikely to have objects repeatedly accessed by a single domain but have no accesses from any other domain (i.e., locally shared only). The percentage of objects that are not shared is also very low. These are objects that received only one access in all and there are usually not very many of them.

Figure 14 plots the number of objects that were shared by exactly x organizations. As it is shown, a large number of objects (over 10,000) are accessed by a single domain. On the other hand, there is a significant portion of objects that are accessed by many domains. Some objects receive requests

from over 10,000 domains.

Finally, we compute the degree of sharing (both intra-domain and inter-domain) when clients are assigned randomly to domains rather than based on their DNS names. (The number of clients assigned to each domain in this manner is the same as before.) We compare the degree of sharing with random assignment with that when clients are assigned to their true domains. The goal is to determine if domain membership is significant.

The top two graphs in Figure 15 show the intra-domain and inter-domain sharing on 12/17/98, and the lower two graphs show the results for 10/7/99. In both cases, the inter-domain sharing is with random assignment is comparable to that with the true assignment, as we would expect. In contrast, intra-domain sharing with true assignment is noticeably higher than with random assignment for 10/7/99. This is also observed in the traces of many other periods. On the other hand, the intra-domain sharing on 12/17/98 (the day of Operation Desert Fox) is comparable with both true and random assignment. From these results, we conclude the following. In most cases domain membership is significant, i.e., clients belonging to the same domain are more likely to share requests than clients picked at random. However, when there is a "hot" event, the global interest can become so dominant that even clients picked at random tend to share many requests, thus diminishing the significance of domain membership.

5.6 Other Results

5.6.1 Correlation between document age and popularity

We examined the correlation between document age (i.e., the time elapse since it's creation) and popularity. Our results are shown in Figure 16, where x-axis denotes the time elapse since the document creation, and y-axis denotes the document ID sorted in increasing order by the total number of accesses. It is evident from the graphs that most documents receive a lot more accesses soon after their creation than afterwards. On the other hand, there are a number of documents (the ones denoted with a large document ID) that remain hot for the

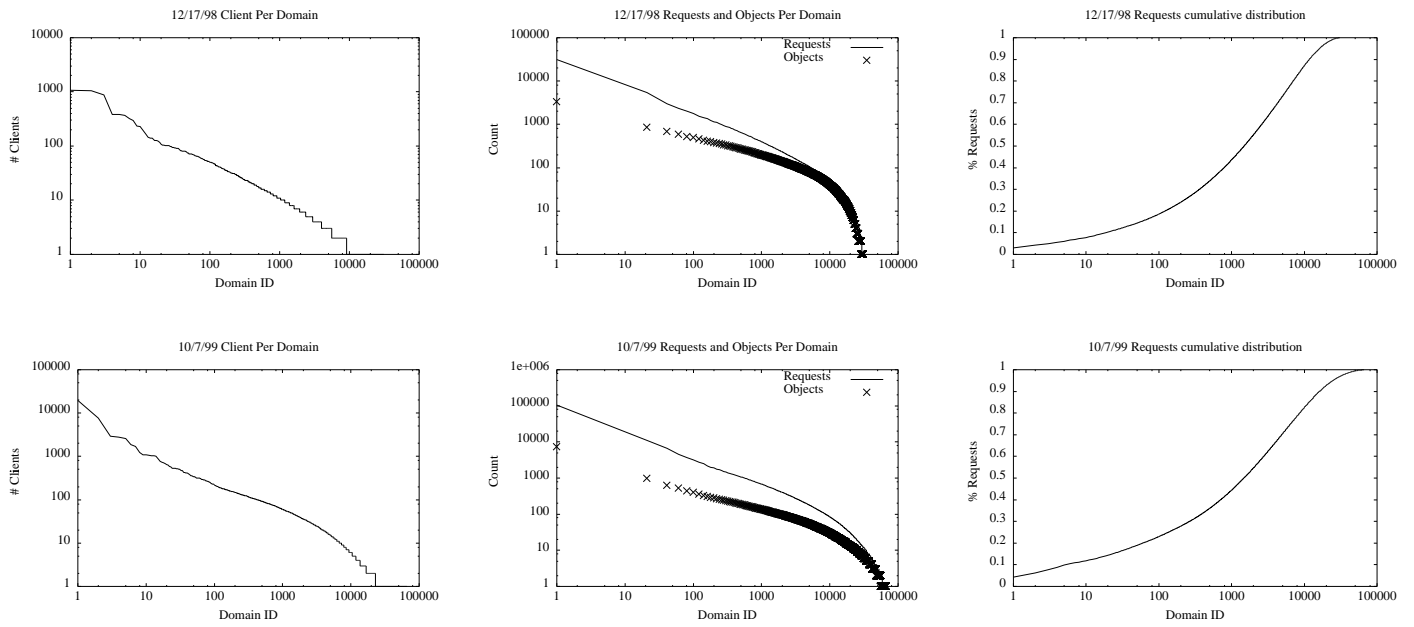


Figure 12: Distribution of clients, objects, and requests in domains. The object and request graph is sorted by the number of objects in the domain.

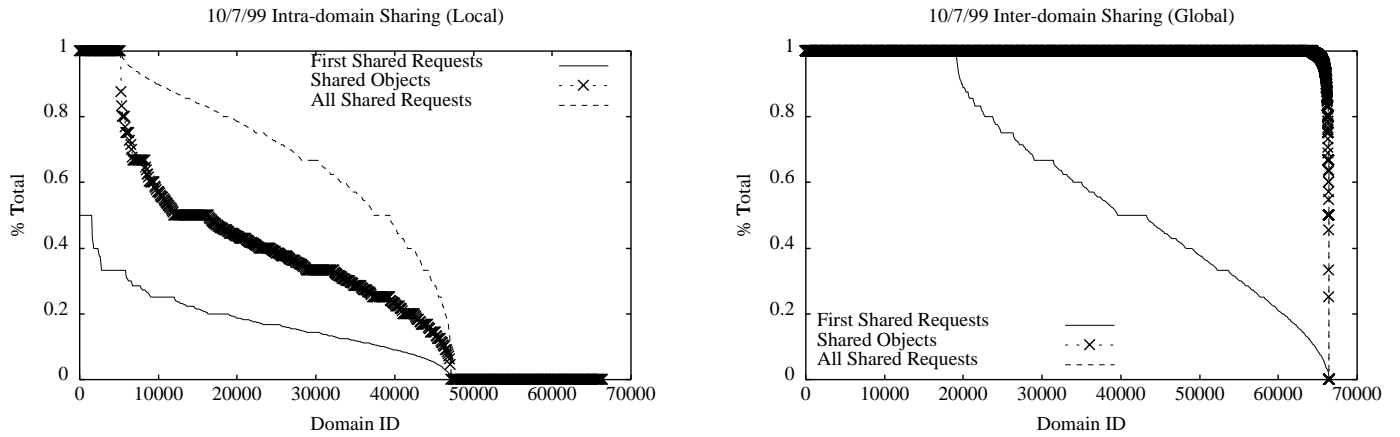


Figure 13: Intra-domain and inter-domain sharing

	Globally shared only		Globally and locally shared		Locally shared only		Not shared	
	mean	variance	mean	variance	mean	variance	mean	variance
Objects	0.6774	0.0846	0.3202	0.0844	0	0	0.0023	0.0009
Requests	0.4965	0.1362	0.5018	0.1364	0	0	0.0017	0.0008

Table 3: Breakdown of objects and requests according to the way in which they are shared for 10/7/99 traces

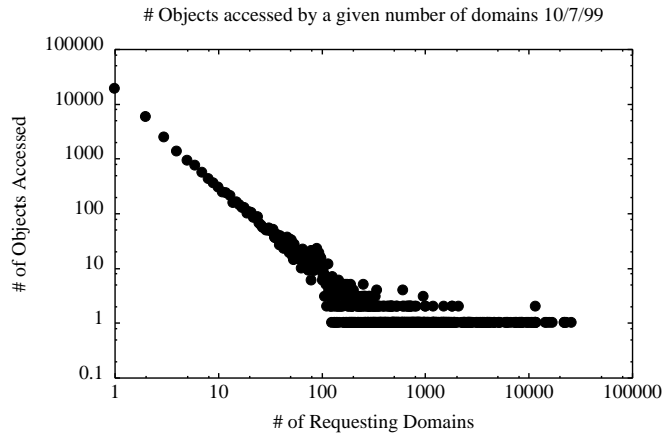


Figure 14: The number of domains that access the objects

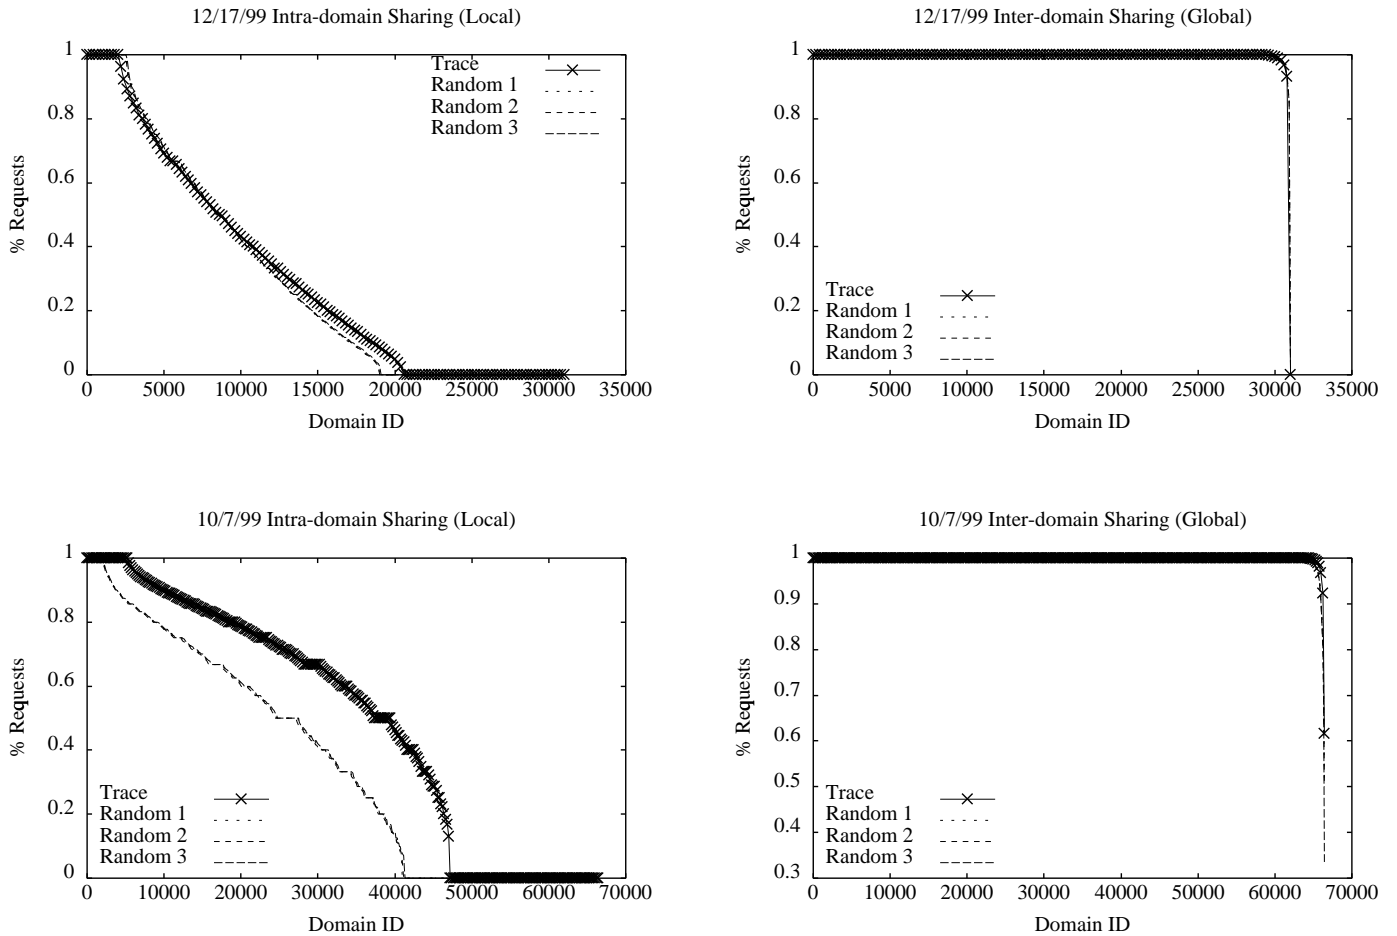


Figure 15: Compare the intra-domain and inter-domain sharing in the real traces with four random client to domain assignments

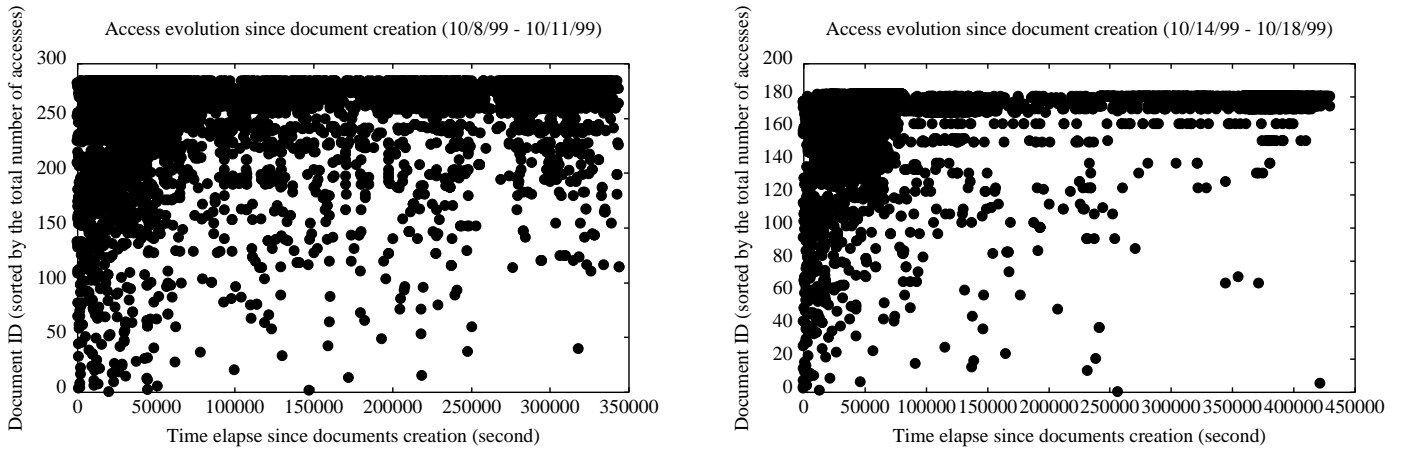


Figure 16: Web Accesses Time Series

entire five-day period under study.

5.6.2 Classify accesses according to modification/creation

Previous research [25] shows that up to 40% of accesses could to the objects that have not been accessed before by clients in a domain. In a caching context, these lead to first-time (i.e., compulsory) misses. We are very interested in understanding what compose these first-time accesses/misses. Using both access logs and modification logs, we have found that most first-time accesses (which we determine on a per-domain basis) are to old objects that were created at least a day ago. This is shown in Table 4. We believe such accesses are very hard to predict. This is because objects that haven't been accessed by a certain domain several days following its creation are considered unpopular. Accesses to unpopular documents are relatively hard to predict. Moreover because these documents are unpopular, the benefit of prefetching them in advance is not cost-effective.

We also study the repeated accesses to the objects according to its modification history. More specifically, we divide the accesses into two groups: accesses to objects that are modified on the day of access, and those that not. As shown in Table 4, over half the repeated accesses are to the modified objects. This implies the potential benefit of good object

invalidation algorithms is large.

6 Conclusions

In this paper, we have studied the dynamics of a large commercial Web site, which is consistently ranked among the busiest in the Web. We have analyzed the dynamics of both the content and the client accesses at this site.

6.1 Summary of Key Results

Our main findings are:

1. The server content tends to be highly dynamic with around 6000 files created and 24000 files modified over a one-week period. For the subset of files that are modified, the time gap between successive modifications tends to lie between an hour and 24 hours (i.e., a day).
2. Past modification behavior of a file, if averaged over a sufficient number of samples, tends to be a reasonably good predictor of future modification behavior.
3. Most (HTML) file modifications tend to be minor both in terms of the change in the file size and in the visible textual content.
4. File popularity tends to be distributed according to a Zipf-like distribution. However, the parameter α tends

Date	First & New	First & Old	Repeated & Mod	Repeated & Unmod
10/8/99	72362	240119	1317246	1014821
10/9/99	14652	96167	697338	576113
10/10/99	15156	99248	758626	610435
10/11/99	47619	206743	1163424	919085

Table 4: Breakdown the web accesses according to the modification/creation

to be in the range 1.4–1.6, which is much larger than has been reported in the literature (based on the analysis of proxy logs). The large value of α implies, for instance, that just the top 2% of documents account for 90% of the accesses.

5. The popularity of files tends to be stable over a timescale of days. Of the top 100 documents in terms of popularity on a given day, 60–100% tend to remain among the top 100 for up to 5 days. However, the set of domains from which accesses to the popular documents are made tends to change significantly from day to day. For example, there is only a 40% overlap when considering the top 100 documents. This may be a consequence of our (fine-grained) definition of a domain, and we are currently exploring this further.
6. Organizational (i.e., domain) membership of clients tends to have a significant (positive) impact on the degree of local sharing, unless there is a globally-interesting event (such as Operation Desert Fox in December 1998) that cuts across organizational boundaries.
7. In the case of most documents, their popularity tends to drop off with age. However, some documents tend to maintain their popularity for a significant length of time.
8. First-time accesses (i.e., the first access to a document by any client in a domain) generally tend to be to documents that are at least a day or more old and are unpopular.

6.2 Broad Implications

We believe our findings have broad implications for the future evolution of the Web:

1. The frequency of file modification underscores the importance of having efficient cache consistency mechanisms (e.g., [6] [23]).
2. File creation also tends to be a frequent event. When a popular news story is updated, it may be assigned a new file name. The absence of past access history for this new file may make the task of prefetching or preemptively pushing out such content challenging.
3. Since the degree to which files are modified tends to be small, techniques such as delta encoding [19] appear promising.
4. The large value of α in the Zipf-like distribution of file popularity together with the stability of the popular set of files over time suggests that caching just a small fraction of the files at proxies may help eliminate a large fraction of the requests.
5. Since first-time accesses tend to be to old and unpopular documents, it appears that it will be hard to cut down on the (significant number of) first-time misses experienced by Web caches using techniques such as prefetching.

7 Ongoing and Future Work

We are currently analyzing a larger content log set, investigating better heuristics for identifying client domains, and examining why the stability in interest group is low. For the longer term, we would like to study data sets from other large server

sites to confirm or refute the findings reported in this paper. We also plan to develop prefetching and/or preemptive push algorithms that are optimized based on the insights we have gained from this paper.

References

- [1] <http://www.abcnews.com>
- [2] V. Almeida, A. Bestavros, M. Crovella, and A. Oliveira. Characterizing Reference Locality in the WWW. *Technical Report TR-96-11, Boston University*, 1996.
- [3] M. F. Arlitt and C. L. Williamson. Web Server Workload Characterization: The Search for Invariants. In *Proc. of SIGMETRICS'96*, May 1996.
- [4] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, M. Stemm, and R. H. Katz. TCP Behavior of a Busy Web Server: Analysis and Improvements. In *Proc. Infocom 1998*, March 1998.
- [5] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web Caching and Zipf-like Distributions: Evidence and Implications. In *Proc. of INFOCOMM'99*, March 1999.
- [6] E. Cohen, B. Krishnamurthy, and J. Rexford. Improving End-to-End Performance of the Web Using Server Volumes and Proxy Filters. In *Proc. SIGCOMM'98*, September 1998.
- [7] <http://www.cnn.com>
- [8] F. Douglis, A. Feldman, B. Krishnamurthy, and J. C. Mogul. Rate of Change and Other Metrics: A Live Study of the World Wide Web. In *Proc. USITS '97*, December 1997.
- [9] B. M. Duska, D. Marwood, and M. J. Feeley. The Measured Access of World Wide Web Proxy Caches. In *Proc. USITS '97*, December 1997.
- [10] L. Fan, Q. Jacobson, P. Cao, and W. Lin. Web Prefetching Between Low-Bandwidth Clients and Proxies: Potential and Performance. In *Proc. of SIGMETRICS'99*, May 1999.
- [11] S. D. Gribble and E. A. Brewer. System Design Issues for Internet Middleware Services: Deductions from a Large Client Trace. In *Proc. USITS '97*, December 1997.
- [12] D. A. Grossman, O. Frieder. Information Retrieval — Algorithms and Heuristics. *Kluwer International Series in Engineering and Computer Science*, September 1998.
- [13] J. Gwertzman, M. Seltzer. World-Wide Web Cache Consistency. In *Proc. USENIX '96*, January 1996.
- [14] R. Jain. The Art of Computer Systems Performance Analysis. *John Wiley and Sons*, 1991.
- [15] D. Li and D. R. Cheriton. OTERS (On-Tree Efficient Recovery using Subcasting): A Reliable Multicast Protocol, In *Proc. of 6th IEEE International Conference on Network Protocols (ICNP'98)*. October 1998, pp. 237-245.
- [16] D. Li, and D. R. Cheriton. Scalable Web Caching of Frequently Updated Objects Using Reliable Multicast. In *Proc. USITS '99*, October 1999.
- [17] Microsoft Corporation. <http://www.microsoft.com>
- [18] J. C. Mogul. Network Behavior of a Busy Web Server and its Clients. *Research Report 95/5, Compaq Western Research Lab*, October 1995.
- [19] J. C. Mogul, F. Douglis, A. Feldman, and B. Krishnamurthy. Potential Benefits of Delta Encoding and Data Compression for HTTP. In *Proc. SIGCOMM '97*, September 1997.
- [20] <http://www.mediametrix.com>
- [21] <http://www.msnbc.com>
- [22] V. N. Padmanabhan and J. C. Mogul. Using Predictive Prefetching to Improve World Wide Web Latency. *ACM SIGCOMM Computer Communication Review*, July 1996.

- [23] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of Multicast Trees: Comments on the Chuang-Sirbu scaling law. In *Proc. SIGCOMM '99*, September 1999.
- [24] J. Touch. The LSAM Proxy Cache - a Multicast Distributed Virtual Cache.
- [25] A. Vahdat, M. Dahlin, T. Anderson, and A. Aggarwal. Active Names: Flexible Location and Transport of Wide-Area Resources. In *Proc. USITS '99*, October 1999.
- [26] A. Wolman, G. Voelker, N. Sharma, N. Cardwell, M. Brown, T. Landray, D. Pinnel, A. Karlin, H. Levy. Organization-based Analysis of Web-Object Sharing and Caching. In *Proc. USITS '99*, October 1999.
- [27] A. Wolman, G. Voelker, N. Sharma, N. Cardwell, M. Brown, T. Landray, D. Pinnel, A. Karlin, H. Levy. On the Scale and Performance of Cooperative Web Proxy Caching. In *Proc. SOSP '99*, December 1999
- [28] H. Yu, L. Breslau, and S. Shenker. A Scalable Web Cache Consistency Architecture. In *Proc. of SIGCOMM'99*, September 1999.

A Appendix

Given there are m lower-level proxies requesting the object, and x higher-level proxies, derive the average number of higher-level proxies that access the page.

Let n_j be the number of different ways a page is accessed by j Higher-level proxies, and z be the average number of accesses the page receives from the higher level proxies. Then we have

$$z = \frac{\sum_{j=1}^x j * n_j}{\sum_{j=1}^x n_j}$$

Now let's compute n_j . As we know, n_j denotes the total number of ways of assigning m accesses from lower level proxies to j (out of m) higher level proxies. This means

$$n_j = \binom{x}{j} * N$$

where N is the number of ways of putting m balls into j bins. By definition, N is the total number of ways of assigning a_i such that $a_1 + a_2 + \dots + a_j = m$, where $a_i > 0$. This is equivalent to the number of ways of assigning s_i such that $s_1 = a_1, s_2 = a_1 + a_2, \dots, s_j = a_1 + a_2 + \dots + a_j = m$. Since $a_i > 0$, s_i is an increasing serie. Therefore $N = \binom{m-1}{j-1}$. So we have

$$n_j = \binom{x}{j} * \binom{m-1}{j-1}$$

So

$$z = \frac{\sum_{j=1}^x j * \binom{x}{j} * \binom{m-1}{j-1}}{\sum_{j=1}^x \binom{x}{j} * \binom{m-1}{j-1}}$$

Note that

$$\begin{aligned} & \sum_{j=1}^x \binom{x}{j} * \binom{m-1}{j-1} \\ &= \sum_{j=1}^x \binom{x}{x-j} * \binom{m-1}{j-1} \\ &= \binom{x+m-1}{x-1} \end{aligned}$$

$$\begin{aligned} & \sum_{j=1}^x j * \binom{x}{j} * \binom{m-1}{j-1} \\ &= x * \sum_{j=1}^x \binom{x-1}{j-1} * \binom{m-1}{j-1} \\ &= x * \sum_{j=1}^x \binom{x-1}{x-j} * \binom{m-1}{j-1} \\ &= x * \binom{x+m-2}{x-1} \end{aligned}$$

With simple algebraic manipulations, we immediately get

$$z = \frac{x * m}{x + m - 1}$$