

Modeling object characteristics of dynamic Web content

W. Shi,^{a,*} E. Collins,^b and V. Karamcheti^b

^aDepartment of Computer Science, Wayne State University, 420 State Hall, 5143 Cass Ave, Detroit, MI 48202, USA

^bDepartment of Computer Science, New York University, Room 704, 715 Broadway, New York, NY 10003, USA

Received 14 May 2003

Abstract

Requests for dynamic and personalized content have increasingly become a significant part of Internet traffic, driven both by a growth in dynamic Web services and a “trickle-down” effect stemming from the effectiveness of caches and content-distribution networks at serving static content. To efficiently serve this trend, several server-side and cache-side techniques have recently been proposed. Although such techniques, which exploit different forms of reuse at the sub-document level, appear promising, a significant impediment to their widespread deployment is (1) the absence of good models describing characteristics of dynamic Web content and (2) the lack of effective synthetic content generators, which reduce the effort involved in verifying the effectiveness of a proposed solution.

This paper addresses both of these shortcomings. Its primary contribution is a set of models that capture the characteristics of dynamic content both in terms of independent parameters such as the distributions of object sizes and their freshness times, as well as derived parameters such as content reusability across time and linked documents. These models are derived from an analysis of the content from six representative news and e-commerce sites, using both size-based and level-based splitting techniques to infer document objects. A secondary contribution is a Java-based dynamic content emulator, which uses these models to generate edge-side include-based dynamic content and serve requests for whole documents as well as separate objects.

© 2003 Elsevier Inc. All rights reserved.

Keywords: Object characteristics; Dynamic Web content; Dynamic content emulator

1. Introduction

Dynamic and personalized content delivery has attracted a great deal of attention recently from both commercial and research communities. One reason is, of course, the growing popularity of dynamic Web services, exemplified by news sites and personalized sites (e.g., my.yahoo.com) both of which require dynamic generation of content. The other reason is the trickle-down effect [18,20] of widely deployed proxy caches and content delivery networks (CDNs), which effectively filter incoming requests for static Web content and presumably shift the traffic seen on the Internet from popular static objects to less popular objects and dynamic Web content.

To efficiently serve and deliver such dynamic and personalized content, researchers have recently pro-

posed several server-side and cache-side mechanisms. Server-side techniques, exemplified by techniques such as delta encoding [28], data update propagation [10], fragment-based page generation [11,17], and the Cachuma [42] class-based page classification scheme reduce the burden on the server by allowing reuse of previously generated content to serve new requests. Cache-side techniques, exemplified by systems such as Active Cache [9], Gemini [29], CONCA [34], and the content assembly technique proposed by Wills et al. [39] attempt to reduce the latency of dynamic content delivery by moving some functionality to the edge of the network. Similar trends are also visible in commercial caching and edge server products, most notably IBM’s WebSphere [22] and Akamai’s Edgesuite [1]. Despite their focus on different aspects of the problem, these approaches share the same rationale, specifically that it is possible to view the document in terms of a quasi-static *template* (expressed using formatting languages such as XSL-FO [37] or edge-side include (ESI) [36]), which is filled out with multiple individually cacheable and/or uncacheable

*Corresponding author.

E-mail addresses: weisong@wayne.edu (W. Shi), edc206@cs.nyu.edu (E. Collins), vijayk@cs.nyu.edu (V. Karamcheti).

objects. This object composition assumption enables surrogates and downstream proxy caches to reuse templates and cacheable objects to efficiently serve subsequent requests and additionally reduces server load, bandwidth requirements and user-perceived latencies by allowing only the unavailable objects to be fetched.

Although the above techniques appear promising, it is difficult to predict to what extent their stated benefits apply to a workload different than the one they were evaluated on. As an example of the challenges, consider the following questions that we were faced with when trying to come up with general policies for our CONCA architecture [34]: At what granularity must objects be cached to achieve sufficient reuse? Is there a sharp threshold for choosing this granularity, or is it the case that the benefits are continuously varying? Can we estimate likely freshness times of objects from the duration they have been present in the cache? Is their a correlation between object size and their reuse?

While trying to answer these questions we encountered the problem that unlike the wealth of models on the nature of static Web content, ranging from the Zipf-like request distribution [7] to the Pareto distribution of document sizes [13,26], there is an absence of corresponding models for dynamic Web content. Though some researchers have studied the rate of change for dynamic Web content [8,16], most previous work focuses on the document level granularity, which is not sufficient to answer the above questions. Even if such models were present, an additional problem one encounters while trying to evaluate a new technique is the absence of template-based content. At this time, Web servers do not explicitly supply the template and component objects making up the document (even if to a human eye, it is often clear what these are). The lack of such content rules out simulation-based studies of the kind that have resulted in tremendous advances in the state of the art for delivery and caching of static content.

This paper addresses both of these shortcomings. Its primary contribution is a set of models that capture the characteristics of dynamic Web content both in terms of independent parameters such as the distributions of component object sizes and their freshness times, as well as derived parameters such as whole-document content reusability across time and linked documents. To derive these models, we analyzed dynamic content collected every 10 min over a 2-week period from six representative news and e-commerce sites: www.cnn.com, dailynews.yahoo.com, www.nytimes.com, www.amazon.com, www.barnesnoble.com, and www.windowsmedia.com. To cope with the absence of an explicit template in the documents, we inferred both the template and the component objects using parameterized level-based and size-based splitting techniques. The main analysis results are summarized below:

- The sizes of component objects making up a dynamic Web page can be captured very well using an *Exponential* or *Weibull* distribution. This is in contrast to how static documents are modeled, usually with a heavy-tailed Pareto distribution.
- Except for a considerable fraction of objects that change very infrequently, the freshness times of component objects can be captured using a *Weibull* distribution. For the two e-commerce sites above (www.amazon.com and www.bn.com), this distribution degenerates into a sharp bimodal one: all of the objects either change on almost every access or change very infrequently.
- Content from all of the six sites demonstrated significant opportunity for reuse across both the temporal and spatial (in linked documents) dimensions. More interesting is the relationship between content reuse and the object granularity: for four of the sites (www.cnn.com, dailynews.yahoo.com, www.nytimes.com, www.windowsmedia.com), there was a graceful degradation of reusability with increasing object granularity, while the degradation was much sharper for the other two.

A secondary contribution of the paper is a dynamic content emulator (DYCE), which uses these models to generate parameterizable dynamic content that adheres to the ESI specification. DYCE builds on top of the Tomcat Web server from the Jakarta open source initiative [3], and can service requests for both whole documents as well as individual components. DYCE is easy to configure, extend, and use and should prove a useful tool for other researchers in this area. To validate both our models and their use in DYCE, we wrote an idealized cache simulator that can work off both the real trace data as well as the output of DYCE. Comparing the outputs of this simulator for the two cases verifies that DYCE effectively models real content, and at a significant simulation cost advantage.

The rest of this paper is organized as follows. Section 2 describes metrics of interest and the methodology used in this study, specifically the extraction of information about document templates and component objects. The analysis results for content from the six sites and derived models are presented in Section 3. Section 4 describes the design, implementation, and validation of the stand-alone dynamic content emulator. Section 5 discusses related work and we conclude in Section 6.

2. Characterizing dynamic content

We first discuss the metrics of interest in dynamic content and then describe our approach for characterizing them.

2.1. Metrics of interest

The consensus view for thinking about dynamic Web content appears to be based upon the notion of object composition. A document is viewed as consisting of a quasi-static document *template* and individual *objects* exhibiting different characteristics, which fill out this template. These characteristics can include attributes such as whether the object is cacheable or not, what its update frequency is, whether it can be shared across multiple requests/users or if it is personalized, etc. Fig. 1 shows the snapshot of a personalized my.yahoo.com page, and what the corresponding document template and component objects might look like. S* and P* represent objects that are shared and private, respectively, and TTL captures the length of time this object remains valid.

From the perspective of surrogates and proxy caches that are attempting to improve delivery of dynamic content, there are several metrics of interest. Clearly, both the *number of objects* and their *size distribution* are important: the first indicates the opportunity for reuse, while the second determines whether or not this opportunity can be traded off against object management overheads. Another metric of interest is the document-level *content reusability*, both across repeated requests for the same document (the temporal dimension) as well as across requests for other documents that are linked from this one (the spatial dimension). Content reusability, which we define to be the fraction of the document that can be reused is in some sense a derived metric that depends on several independent parameters: the sizes of individual objects, their *freshness time* (the length of time the object remains valid), and the correlation between the two. The larger the content reusability, the more the likely benefits from techniques that either reuse or cache individual component objects.

Section 3 characterizes dynamic content from a representative set of Web sites in terms of these independent and derived metrics.

2.2. Methodology

If one has access to the Web server and/or application server of content providers, it is a relatively straightforward exercise to instrument these to obtain the above metrics of interest. Unfortunately, many commercial Web sites which create dynamic content are proprietary. Therefore, we adopt an alternative approach based on the analysis of dynamic content (an HTML file) downloaded from various Web sites.

The biggest difficulty that must be overcome in our chosen approach is the lack of an explicit template associated with the document, which can help identify its component objects. To work around this difficulty, we make the assumption that *the document template can be expressed as a nested table*. Fig. 2(a) shows an example: here, the document includes head and body tags, with the body including one table and two rows. We have been unable to find a Web site whose pages are fully dynamically generated, where this assumption is not satisfied.

Given this assumption, our approach first extracts the template associated with the document, then identifies the (logical) objects that fill out this template (grouping neighboring objects that exhibit the same freshness time characteristics), and finally aggregates logical objects into physical objects that serve as the granularity at which document characteristics are modeled in Section 3. These steps are described in additional detail below:

- (1) *Data collection*: In this first step, a dynamic document is downloaded every *T* time units from its Web site to set up a series. Since the time interval

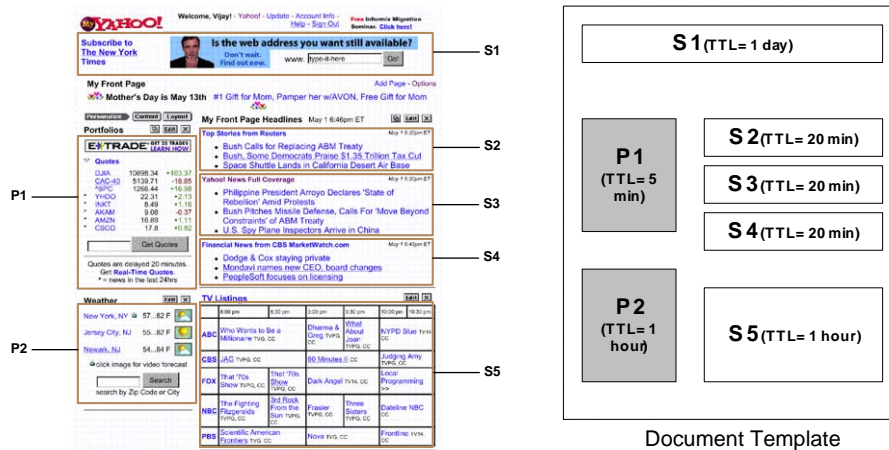


Fig. 1. Dynamic content can be viewed in terms of a quasi-static document template and individual objects, which exhibit different sharing, cacheability, and freshness time characteristics.

used in the data collection determines the minimum value of the freshness interval, this time interval should be set as small as possible. The documents in this series are denoted as (d_1, d_2, \dots, d_n) .

- (2) *Tree building*: Our assumption about the document structure allows the document to be represented as a tree of objects. In principle, this should be an easy affair: since HTML has an associated XML DTD associated with it, one could just build the object tree from the corresponding XML tree. However, we found that the HTML produced by many servers does not fully comply with the specification. Therefore, we first preprocess the document using tidy [33] to make it comply with the W3C XHTML 4.01 standard.

Each leaf node in the tree is a possible object. To assist with the comparison step described next, each node in the tree is annotated with a *content digest*. The latter is computed using the collision-resistant SHA-1 hash function [19]. Internal nodes are associated with a hash of the content digests of their child nodes. Continuing with our example, the corresponding object tree of the document in Fig. 2(a) is shown in Fig. 2(b). Such object trees are computed for all document instances producing (t_1, t_2, \dots, t_n) .

- (3) *Tree comparison*: Based on the object tree constructed in Step 2, the freshness time of each object is computed by comparing the rest of the series with the first document tree t_1 . The comparison is computed in a bottom-up fashion by comparing the digest of each node; the freshness time of the object is set to the time interval corresponding to the earliest document where the digests differ. For

example, if tree t_k is the first tree where the object has changed, then the freshness time of the object is tagged as being the time interval between t_k and t_1 . Fig. 2(c) shows the results of this step, which is also represented as an XML tree. To eliminate the effect of starting time, we computed the freshness time of each object of a series of document trees: the freshness time of the corresponding object is the maximum of the freshness times computed for that object over the series.

- (4) *Object grouping*: This step clusters contiguous objects (in the object tree) with similar properties into a single large logical object. The algorithm again proceeds in a bottom-up fashion from the leaf nodes: sibling nodes with the same freshness times are combined with each other, and a singleton (grouped) node is combined with its parent under the same conditions.

Fig. 2(d) shows the result of object grouping: the dark circles represent objects remaining as the leaf nodes. In this specific example, the first two td nodes 1,3 are combined with their parent node tr, while the third td node exhibits a different freshness time than its parent node, and hence cannot be combined. These remaining leaf nodes define the component objects of the document, with the rest of the tree representing the template.

- (5) *Tree splitting*: The last step of the object identification process groups one or more of the logical objects into physical objects; the latter represent the granularity at which surrogates and proxy caches can work with the document objects. A smaller number of objects would presumably minimize management overheads at these surrogates and

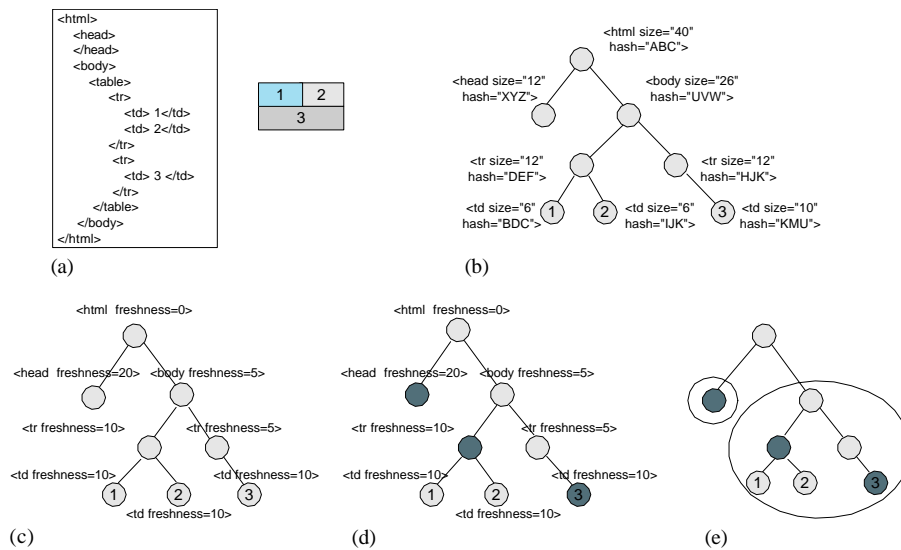


Fig. 2. Identifying the objects making up a document involves four steps, going from (a) a simple document d , to (b) the corresponding object tree of document d , to (c) the object tree annotated with freshness times, to (d) logical groupings of objects, and to (e) the final objects chosen to represent this document.

caches, while a larger number of objects would improve potential reuse. Ideally, this decision ought to be driven by the document semantics, taking into consideration factors such as freshness times, relationships between objects, etc.

In the absence of such information, we have to use heuristic means. In this paper, we use two techniques to split the object tree generated in Step 4 into component objects: *size-based* and *level-based*. Size-based splitting chooses nodes whose size exceeds a certain threshold (and whose child nodes have size smaller than the threshold). Level-based splitting follows the logical structure of the document: all nodes below a certain depth are grouped together. Returning to our example, if only two levels are allowed, there will be only two nodes in the document, the `<head>` and `<body>`. Note that although we employ these two heuristics independently, they can be combined to choose document objects.

3. Analysis of dynamic content characteristics

To model the characteristics of dynamic content, both in terms of independent parameters such as the number, sizes, and freshness times of objects, as well as derived parameters such as content reusability, we analyzed traces collected from six Web sites using the methodology described in the previous section.

The main pages at these sites, which included three news sites (cnn, yahoo, new york times), two e-commerce sites (www.amazon.com, www.barnesnoble.com), and an entertainment site (www.windowsmedia.com), were downloaded every 10 min over a 2-week interval from January 31 to February 14, 2002. A point that needs to be made here is that we intentionally chose collection periods, where the Web site was functioning in a “stable” mode. The primary reason for this choice is our perception that such information is more useful for longer-term surrogate and proxy cache planning as opposed to the unpredictable behavior of a Web site during “unstable” periods (consider, for example, the nature of dynamic content that news sites generated in the period 09/10–09/12, 2001).

Each document was split into physical objects according to four different size limits (0, 500, 1000, and 2000 bytes), and three different level limits. A size limit of 0 means that any object can be chosen as a physical object in our analysis. An additional point needs to be made here. In the size-based splitting scheme, it is possible to have some objects that are smaller in size than the specified limit. This comes about because we assume a static template structure across the document series, which has the consequence that

Table 1

The mean number of objects corresponding to different size and level limits for the six Web sites

Site	Level	Size limit			
		0	500	1000	2000
cnn	18	142	71	49	33
	12	118	71	49	33
	6	35	32	30	26
yahoo	26	113	64	38	35
	20	113	64	38	35
	14	85	58	37	35
nytimes	23	69	34	29	23
	17	67	34	29	23
	11	63	32	28	23
amazon	21	178	28	22	16
	15	106	27	22	16
	9	54	22	20	16
bn	21	114	51	36	19
	15	80	38	30	19
	9	55	27	22	11
wmedia	19	73	30	17	15
	13	47	30	17	15
	7	17	14	13	12

whenever we mark a logical object as being a physical object we also need to apply the same marking for its siblings (even if their size is smaller than the limit).

We only report on a subset of these experiments in this section, primarily to improve the readability of the figures and tables.

3.1. Number of objects

The number of objects in a dynamic document helps determine what opportunity, if any, is available for exploiting reuse. Unlike well-studied corresponding questions in the static content case, for example, the number of embedded images in a document [6], the number of physical objects in a document is influenced by the size and level parameters used for choosing. The average number of objects for each document is shown in Table 1. Each row in the table shows, for a fixed level limit, the number of chosen objects for different size limit settings. These counts verify intuition: the larger the size limit, the fewer the number of objects. Similarly, the deeper the chosen level, the more objects there are.

Unless explicitly stated to the contrary, the measurements in the rest of the section refer to the largest level limit for each of the documents. We have experimented with lower limit settings and have found that the results overall have the same flavor. Note also that although the number of objects in a given document is sometimes

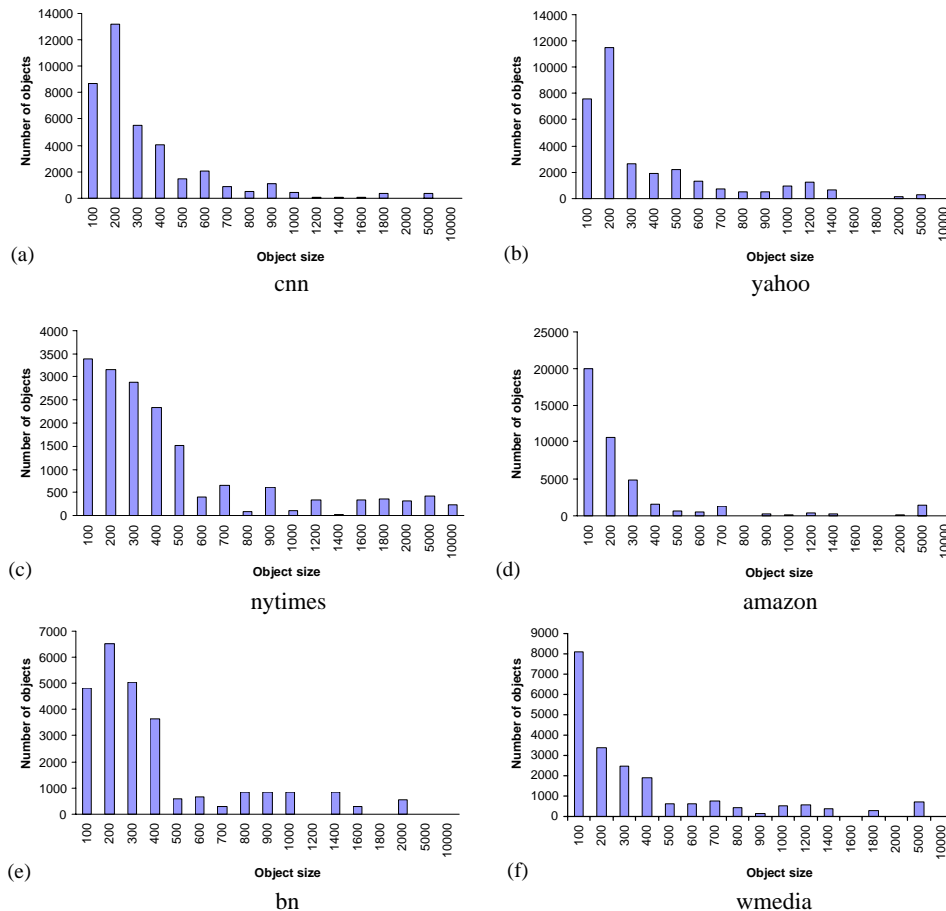


Fig. 3. The distribution of object sizes at different Web sites assuming a size limit of 0 for the tree splitting procedure. Each bar represents the number of objects contained in the size range bounded above by the value on the X -axis and below by the value of the bar to its immediate left.

small, considering these objects over the entire 2-week series provides a statistically significant sample.

3.2. Distribution of object sizes

Similar to the size distribution of static documents, we need to understand the size distribution of objects embedded in a dynamic Web document. Several previous studies have verified that the size of a static resource, characterized by a heavy tail, is captured well by a Pareto distribution for the tail (with $\alpha = 1.0 \sim 1.5$) [12].¹ However, intuition suggests that this characterization may not apply to objects within a dynamic document for two reasons: first, a preliminary characterization of several sites serving dynamic content seems to suggest that dynamic content itself tends not to have overly large sizes, and second, the size of an embedded object is restricted to be smaller than that of the overall document. Fig. 3 shows the distribution of object sizes for our six document series, which verifies this intuition. We find that a large number (about 90%)

of objects in our documents are of relatively small size (less than 500 bytes) when the size limit for tree splitting is set to 0. This fact is better highlighted in the cumulative distribution function graphs shown in Fig. 4.

To develop a model for the object size distributions seen for different size limits, we use standard statistical methods similar to those used by Paxson et al. [32]. We used both chi-square and Anderson-Darling (A^2) empirical distribution functions (EDF) for estimation of goodness-of-fit [14]. By comparing with several distributions, such as lognormal, exponential, Weibull and Pareto, we found that the best value (the smallest value) of chi-square was for the *exponential* distribution, whose cumulative distribution function is $F(x) = 1 - e^{-\lambda x}$. *amazon* and *wmedia* are two exceptions, exhibiting a slightly better fit to a *Weibull* distribution. We note that the Weibull distribution (with a CDF $F(x) = 1 - e^{-(\lambda x)^b}$) is a variant of the exponential distribution, so these anomalies are likely not very crucial. Table 2 shows the parameters of the exponential or Weibull distributions that best fit our trace data for various size limits. The results of the chi-square tests were in the range 0.1–2.5, showing a very close fit. The

¹The body of the resource size has lognormal distribution [6].

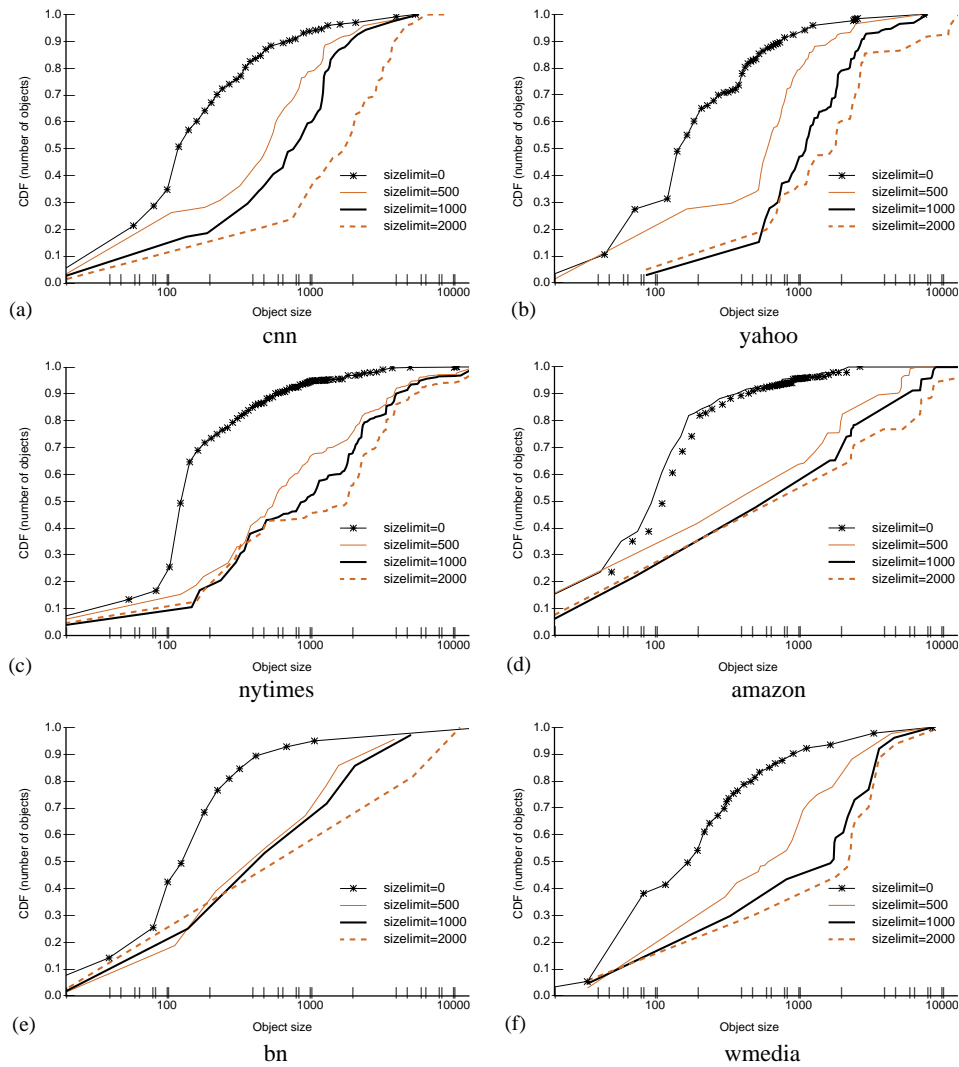


Fig. 4. The cumulative distribution function (of the number of objects) versus object size for different size limit settings.

Table 2

Parameters of exponential and Weibull distributions best fitting the measured object size distributions for the six Web sites. Items with λ and b values denote Weibull distributions

Site	Distribution	Size limit			
		0	500	1000	2000
cnn	Exponential	0.005	0.0015	0.001	4.59×10^{-4}
yahoo	Exponential	0.0038	0.0014	6.85×10^{-4}	4.72×10^{-4}
nytimes	Exponential	0.005	0.0011	$\lambda = 6.7 \times 10^{-4}$	$\lambda = 5.0 \times 10^{-4}$
	Weibull			$b = 0.7$	$b = 0.7$
amazon	Weibull	$\lambda = 0.0059$ $b = 0.9$	$\lambda = 0.0011$ $b = 0.5$	$\lambda = 7.09 \times 10^{-4}$ $b = 0.7$	$\lambda = 5.92 \times 10^{-4}$ $b = 0.5$
	Exponential	0.0055	0.0014	$\lambda = 0.0012$	$\lambda = 7.52 \times 10^{-4}$
bn	Weibull			$b = 0.7$	$b = 0.5$
	Weibull	$\lambda = 0.004$ $b = 0.7$	$\lambda = 0.001$ $b = 0.9$	$\lambda = 5.08 \times 10^{-4}$ $b = 0.9$	$\lambda = 4.02 \times 10^{-4}$ $b = 1.9$

results of the A^2 test showed no significance in terms of goodness of fit for most of the documents; however, this failure, which is a common problem with EDF [6,32],

was primarily due to the fact that a fairly large data set was used. The fit improved significantly when we sub-sampled the data set. Although not shown, we observed

similar distribution fits for different settings of the level limit.

3.3. Distribution of object freshness times

As described in Section 2, object freshness times are computed by comparing objects in a particular document against each subsequent document in the series. Each object is associated with the shortest time interval where the corresponding object in the template differs from the current one. Fig. 5(a)–(f) shows the measured cumulative distribution function, for different settings of the size limit, of the number of objects versus average object freshness times in the range 0 up to 10 080 min (1 week). Fig. 6(a)–(f) zooms into the results at 1-day granularity. A point (x, y) on these curves represents the number of objects (y) in the document that have a

freshness time below x . The reason the CDF curve does not reach 1 at the extreme right of the plots is because the freshness time of the remainder of the objects is significantly higher than the right limit of 10 080 min for the 1-week case, and 1440 min for the 1-day case. These latter objects essentially stay fixed throughout a document series.

The following observations can be made from the figures:

- As noted above, a large fraction of objects (in *cnn*'s case, this number is about 40%) in all six of the documents do not change within the timescale of a week. In fact, the objects do not change even over timescales of a month, indicating that significant reuse can be exploited even by just focusing on objects that exhibit this relatively simple behavior.

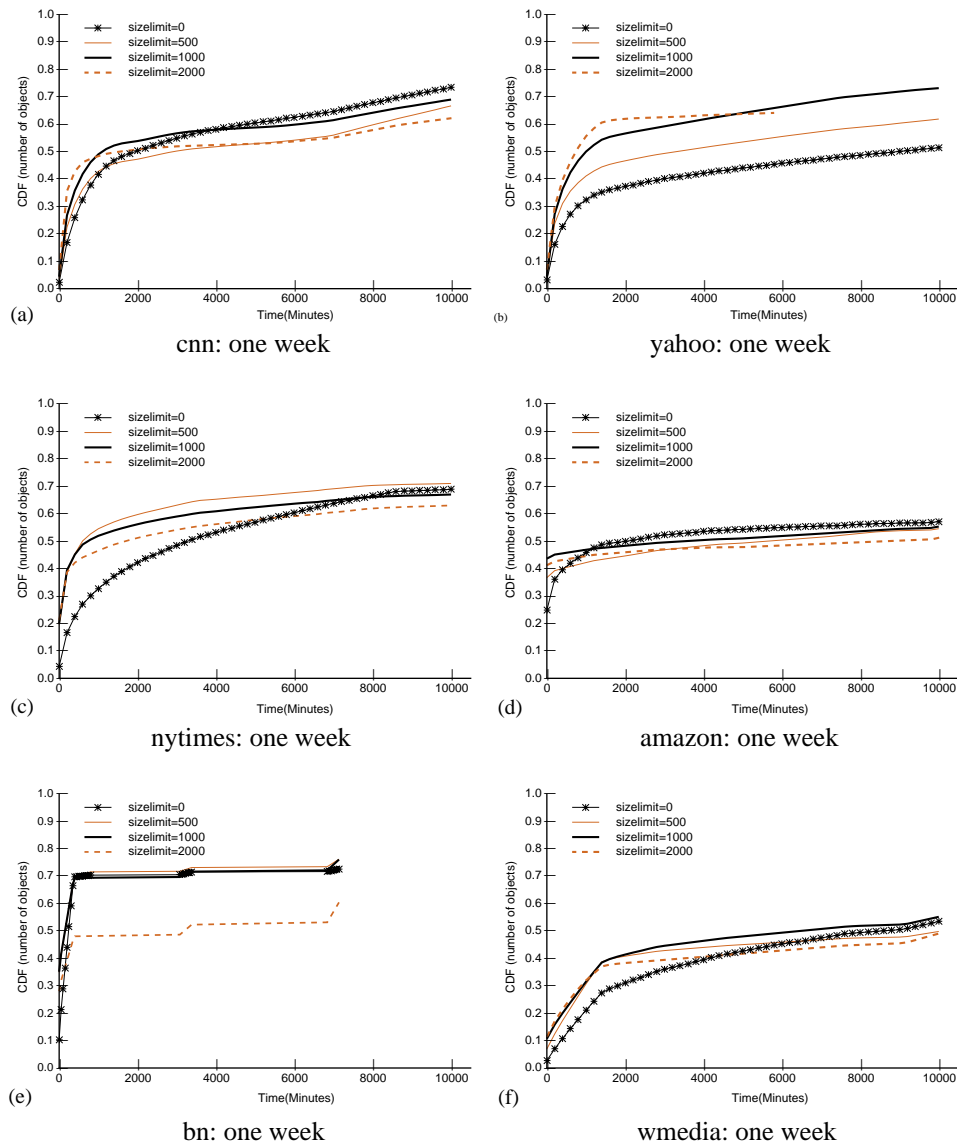


Fig. 5. The distribution of freshness times for different size limits, shown over the timescale of a week.

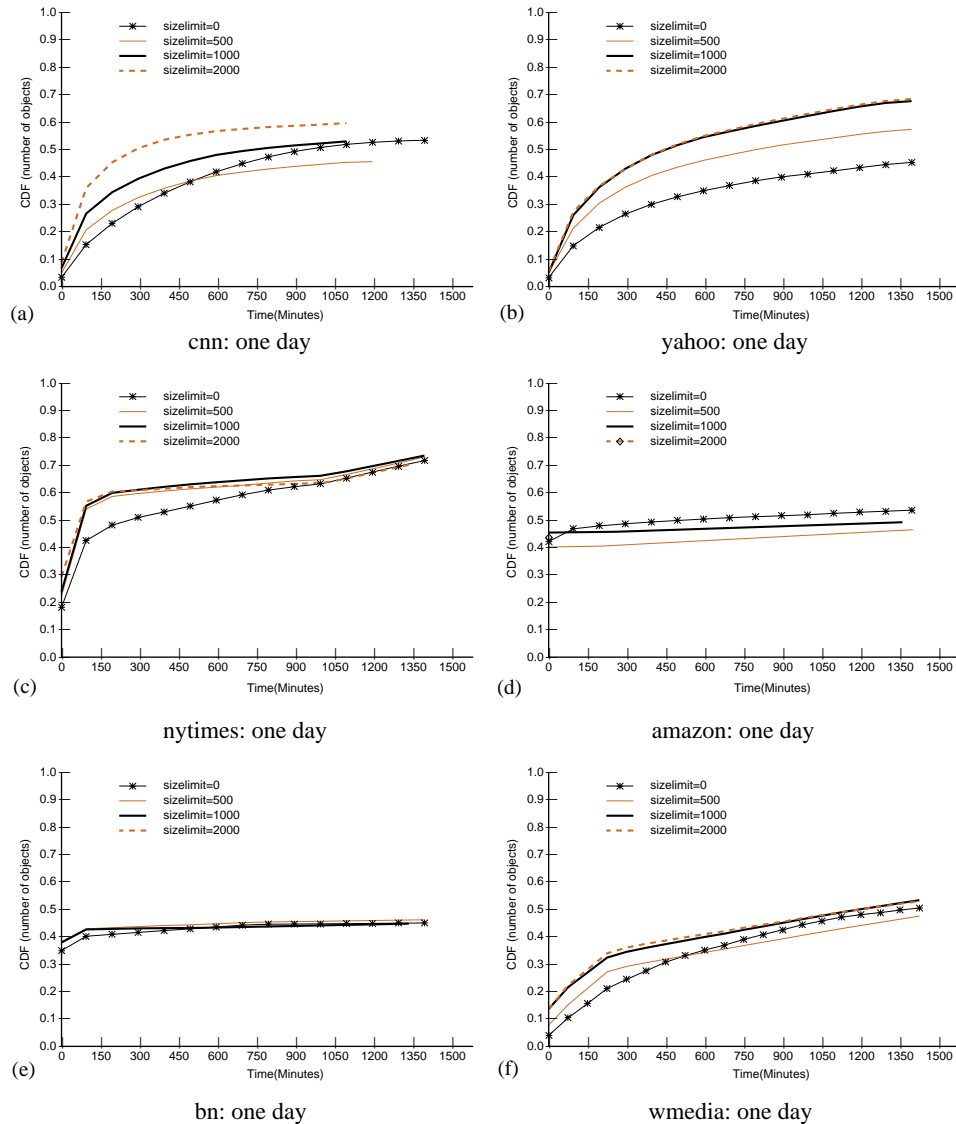


Fig. 6. The distribution of freshness times for different size limits, shown over the timescale of a day.

- For those objects that do change within the timescale of a day, we found that a large fraction of objects have small freshness times, and that this fraction drops off as the freshness time increases. It would seem that such behavior can be captured by the *Pareto* distribution $(F(x) = 1 - (k/x)^{-\alpha}, x > k)$. However, similar to the observation made by other researchers [6], we found that although Pareto is effective at capturing the heavy tail nature of the freshness time distribution, it does not fit well with the body, particularly for low freshness times. We then looked at the Weibull distribution, which has been previously used to model lifetimes of components in computer systems, and found that it characterizes our measured behavior much better.

Table 3 shows the parameters of the Weibull distribution that best fit the data: with the exception

of the commercial sites discussed below, our results show that the object freshness times can be effectively modeled using a Weibull distribution.

- The freshness time distributions for the two commercial Web sites (*amazon* and *bn*) markedly differ from those of the other sites in that all of the objects appear to fall into one of two categories corresponding to either a near-zero freshness time (accounts for about 40% of the objects) or those that have a very large freshness time exceeding 1 day. This behavior, which can be attributed to frequently changing dynamic advertisements included in the page, can be thought of as a degenerate case of the more general behavior that is captured using the Weibull distribution. Table 3 shows the fraction of the objects that fall into each category for the two sites corresponding to different size limits.

Table 3
Parameters of distributions best fitting the freshness time measurements for the six Web sites

Site	Distribution	Size limit			
		0	500	1000	2000
cnn	Weibull	$\lambda = 1.77 \times 10^{-4}$ $b = 0.4$	$\lambda = 1.03 \times 10^{-4}$ $b = 0.3$	$\lambda = 1.56 \times 10^{-4}$ $b = 0.3$	$\lambda = 1.03 \times 10^{-4}$ $b = 0.2$
yahoo	Weibull	$\lambda = 3.77 \times 10^{-4}$ $b = 0.3$	$\lambda = 9.57 \times 10^{-4}$ $b = 0.3$	$\lambda = 2.47 \times 10^{-4}$ $b = 0.3$	$\lambda = 3.23 \times 10^{-4}$ $b = 0.3$
nytimes	Weibull	$\lambda = 1.51 \times 10^{-4}$ $b = 0.5$	$\lambda = 3.0 \times 10^{-4}$ $b = 0.3$	$\lambda = 2.11 \times 10^{-4}$ $b = 0.3$	$\lambda = 1.39 \times 10^{-4}$ $b = 0.3$
amazon	Bimodal	0 : 0.36, > 1435 : 0.43	0 : 0.39, > 1435 : 0.45	0 : 0.43, > 1435 : 0.45	0 : 0.41, > 1435 : 0.42
bn	Bimodal	0 : 0.20, > 1435 : 0.28	0 : 0.35, > 1435 : 0.24	0 : 0.35, > 1435 : 0.24	0 : 0.28, > 1435 : 0.40
wmedia	Weibull	$\lambda = 6.42 \times 10^{-4}$ $b = 0.5$	$\lambda = 3.82 \times 10^{-4}$ $b = 0.3$	$\lambda = 5.16 \times 10^{-4}$ $b = 0.3$	$\lambda = 3.03 \times 10^{-4}$ $b = 0.3$

The two commercial sites show a degenerate case of the Weibull distribution: almost all of the objects have either near-zero freshness time or very large freshness times. The fraction of objects falling into each category is denoted using the form $a : b$, a indicates the freshness time (in min) and b reflects the fraction.

Table 4
Correlation between object size and freshness time for the six sites, expressed in terms of the correlation coefficient

Site	Size limit			
	0	500	1000	2000
cnn	-0.199	-0.280	-0.305	-0.477
yahoo	0.048	0.075	-0.013	-0.053
nytimes	-0.134	-0.196	-0.451	-0.549
amazon	-0.153	-0.900	-0.936	-0.999
bn	-0.064	-0.365	-0.329	-0.641
wmedia	-0.204	-0.254	-0.329	-0.316

3.4. Correlation between object size and freshness time

To ascertain if there was a relationship between the sizes and freshness times of objects in the documents, we computed the correlation coefficient for the two series (see Table 4). While we had hoped to see a strong positive correlation between object size and freshness time (i.e., smaller objects exhibit smaller freshness times and vice-versa), we found that for our traces there was no conclusive indication of any correlation between the two metrics.

3.5. Extent of content reusability

As noted in the beginning of this paper, one of our main motivations for undertaking this study was to better understand the nature of content reusability in dynamic Web content. Content reusability is a derived parameter, which is affected both by the distribution of object sizes and freshness times; however, this dependence is unpredictable, particularly in the absence of any correlation between the two.

There are two dimensions across which surrogates and proxy caches can benefit from content reuse. The first dimension, *time*, benefits repeated downloads of the same document: reuse in this case refers to the fraction of objects that remain valid across multiple downloads. The second dimension, *document links*, improves performance for downloading a group of linked documents. Reuse in this latter case refers to the fraction of objects that are shared across two (different) documents. We examine these two dimensions in turn.

3.5.1. Reusability across time

There are two ways to determine content reusability across time in the same document, based upon one's assumptions about the existence of a document template. The first, referred to hereafter as *content reusability*, assumes that each document is associated with a static template, and defines an object to be reusable *only* if it occupies the same place within the template. The second, which we call *object reusability*, takes a more general view in defining an object to be reusable if it appears *anywhere* in a new document.

Fig. 7 shows the content reusability (in terms of the bytes in the document that can be reused) over a 1-week interval for the homepages of our six Web sites. A few points need to be made here. First, each curve is an average of 1008 contiguous documents (corresponding to a 1-week trace spaced at 10-min intervals), each of which has been compared with 1008 subsequent documents in the series. Second, for the measurements reported here, the level limit was set to the highest possible in each case to ensure that we obtained the maximum potential content reusability. Finally, the curves include reuse of the template, which occupies between 10 and 20% of the whole document.

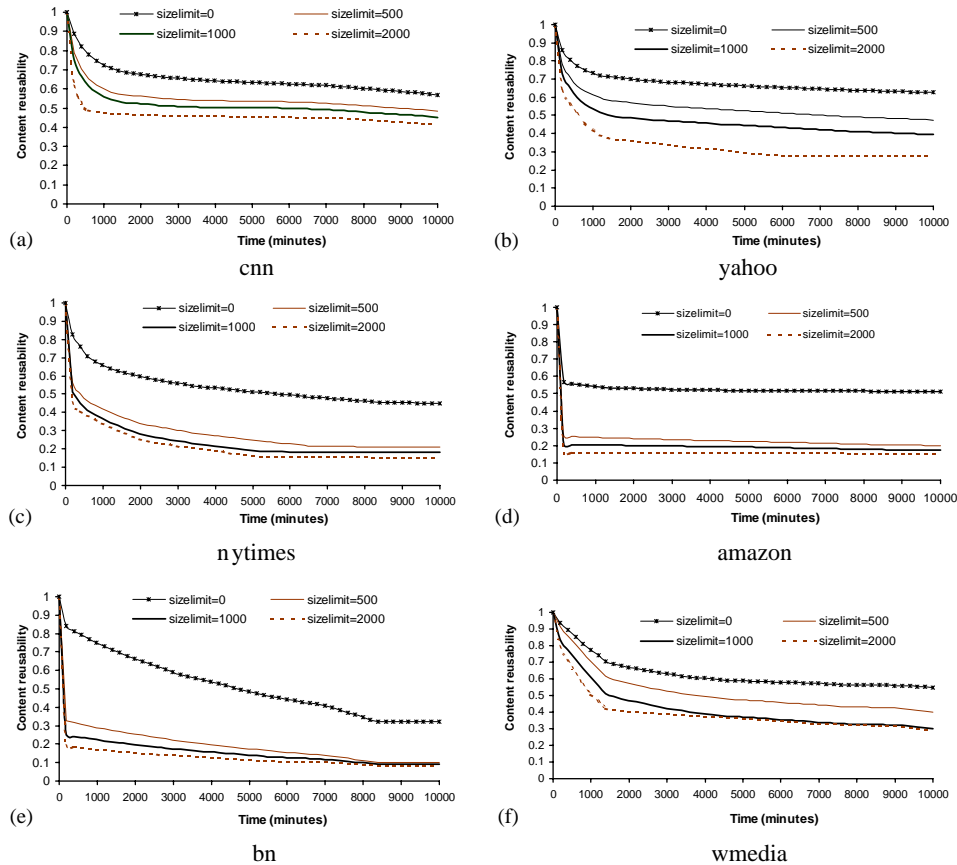


Fig. 7. Content reusability of different Web sites averaged over a 1-week period, assuming a static document template.

Several observations can be made from these figures:

- (1) The smaller the granularity of objects, the more the reusability, which complies with our intuition. For example, 75% of the document bytes are reusable in the case of *cnn* after 12 h have elapsed when the size limit for defining objects is set to 0 (i.e., there are no restrictions on how small an object can be), while the corresponding reusability is 50% when the size limit is set to 2000 bytes.
- (2) Content reusability stabilizes after a time period, and thereafter stays flat for a long time. We redid these measurements using 2-week traces but the results stayed the same, indicating that a certain fraction of the document remains usable over extended periods of time.
- (3) The content reusability curves for the two commercial Web sites are flat, corresponding to the bimodal nature of their freshness time distributions. It is only possible to reuse the unchanged portions of the document, since the rest changes too frequently for there to be any reuse.
- (4) For the remaining four sites, acceptable levels of reusability (larger than 40%) require that the size limit of objects be set less or equal to 1000 bytes.

However, what is encouraging (from the perspective of someone trying to design a proxy cache system) is that the amount of content reusability degrades gracefully as the size limit is increased. Thus, there is a range of size limits, based on the efficiency and overheads associated with caching at the object granularity, where a proxy cache can see benefits from reuse.

Fig. 8 shows the corresponding figures for object reusability. As noted earlier, this represents the fraction of the document bytes that can be reused after a certain time period, assuming that an object in the first document appears somewhere in the second.² Intuitively, object reusability should be greater than content reusability. However, while comparing the set of objects, we do not consider the document template and objects that are smaller than a certain threshold. Table 5 shows the fraction of document bytes accounted for by the objects that are considered. As we can see, the unaccounted for objects end up contributing a sizeable fraction of the overall document bytes and because they

²This is a very compute-intensive operation, therefore, Fig. 8 shows only the object reusability of the traces over a 1-day period.

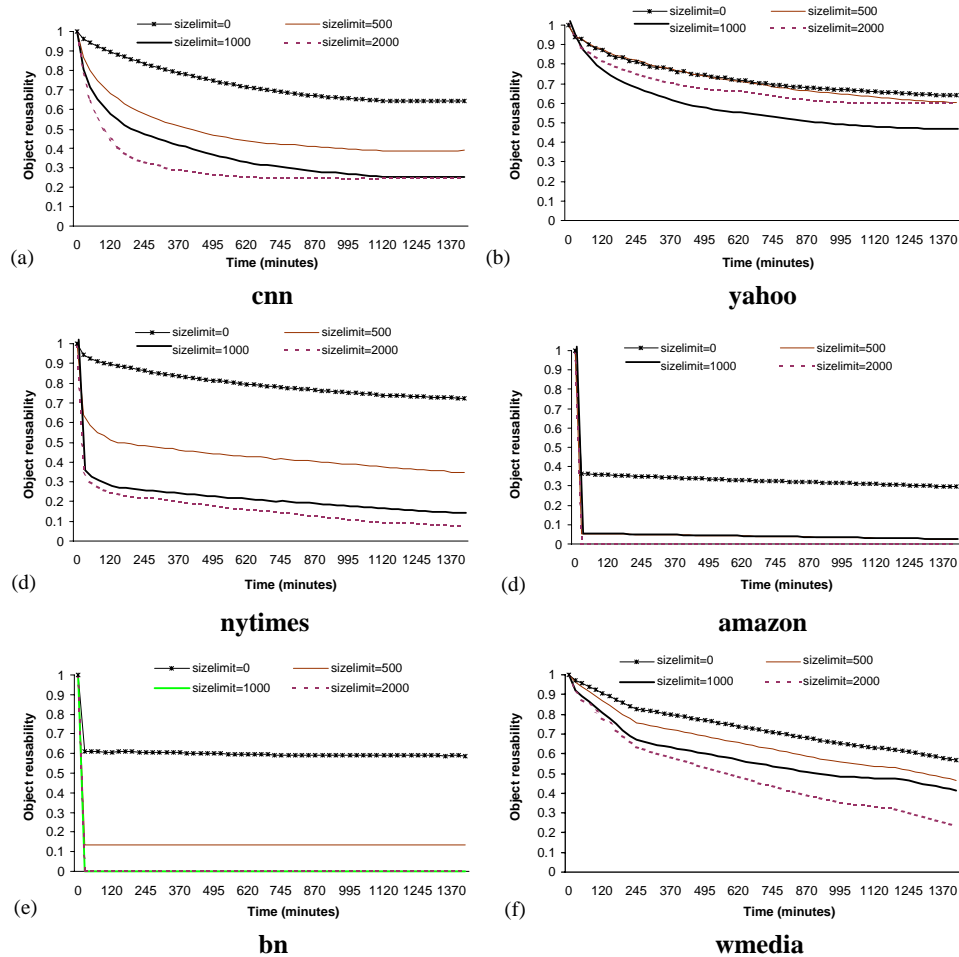


Fig. 8. Object reusability of different Web sites averaged over a 1-day period, assuming the document does not have a template.

Table 5

The fraction of the document bytes accounted for by objects selected in the object reusability computation

Site	Size limit			
	0	500	1000	2000
cnn	89.4%	73.5%	60.3%	64.2%
yahoo	93.5%	87.6%	73.0%	50.8%
nytimes	96.4%	88.0%	88.0%	79.3%
amazon	82.4%	93.3%	90.9%	93.4%
bn	98.2%	96.4%	96.2%	96.2%
wmedia	91.0%	91.8%	88.8%	89.7%

are deemed not reusable, object reusability ends up being slightly lower than the corresponding content reusability. Other than this difference, object reusability exhibits the same characteristics as content reusability and the same observations apply.

3.5.2. Reusability across linked documents

Content reuse is also possible across dynamically generated documents from the same site that end up, for

various reasons, to share the same set of objects. If these documents are linked to each other, then one can exploit this reuse across traversals of the links. As an example, consider the *cnn* Web page, which contains links to pages that contain additional detail about headlines, as well as news in various categories. Given that all of these pages are generated on the same server, it appears reasonable that they would share some objects.

To verify this intuition, as part of our trace collection process, we also collected a set of documents that were linked in from the home page of each of our sites. To take *cnn* as an example, we collected 15 other documents in addition to the home page: local (*loc*), travel (*tra*), weather (*wea*), career (*car*), health (*hea*), special (*spe*), technology (*tec*), US (*usa*), world (*wor*), and the top two headlines, all of which are hosted on the CNN server, as well as two links *finance* (*fn*) and *sports* (*spo*), which are hosted on different servers.

Although these 14 documents do not share the same document template, our analysis results show that large parts of them are in fact shared. Fig. 9 shows the object

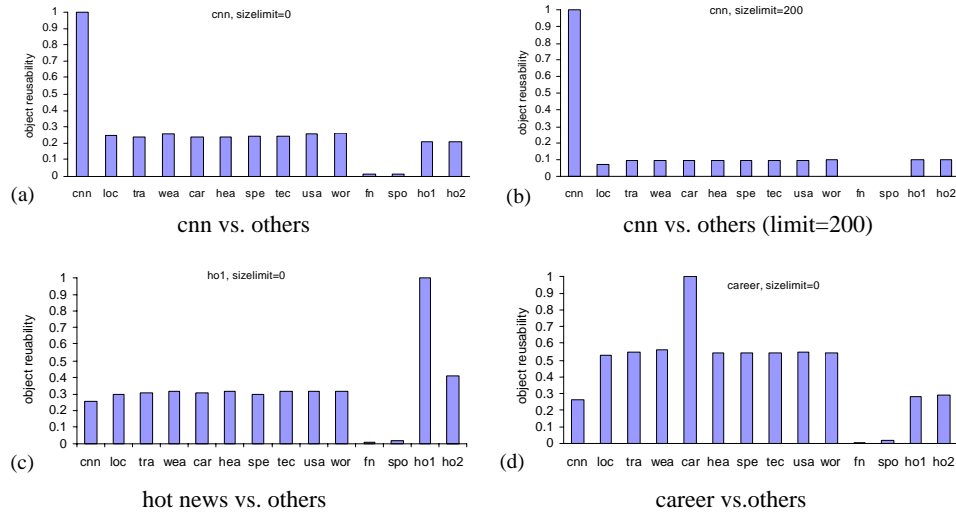


Fig. 9. Object reusability across documents from the www.cnn.com Web site with respect to (a) the *cnn* default page (with size limit 0); (b) the *cnn* default page (with size limit 200); (c) a headline page; and (d) the career news page.

reusability across the different documents, using as reference the document from a particular page. For instance, Fig. 9 (a) and (b) show the reuse with respect to the main *cnn* page when the size limit is 0 and 200 bytes, respectively. For the smaller size limit, we see that 25% of the document bytes can be shared across other documents that are also hosted on the same server, while the sharing across documents that are hosted elsewhere is not so good. This verifies our intuition about server activity. The measurements for the larger size limit shows that achieving these gains requires the ability to reuse objects at a fine granularity.

Fig. 9(c) and (d) show the corresponding measurements using one of the headline pages and the career page as a reference. We find that the amount of reusability is higher (between 40 and 50%), primarily attributable to the smaller size of the overall document in these cases. Thus, our results verify that there is substantial opportunity for reusing objects across dynamic documents that are linked from each other.

3.6. Implications of the content characteristics

Our study of object characteristics of dynamic Web content has several implications for future research in optimizing the generation and delivery of dynamic Web content:

- An immediate implication, which we explore in the rest of this paper, is that such models enable the development of synthetic content generators, a prerequisite for simulation-based studies of the kind that have previously proved very successful in the static content case.
- Our results also answer most of the questions we posed in Section 1. Specifically, our results show that

there is a significant opportunity for content reuse, although one needs to distinguish between content that exhibits a graceful degradation of reusability with object size (e.g., *cnn*, *yahoo*, *nytimes*, and *wmedia*) and the kind that exhibits a sharper degradation (e.g., *amazon* and *bn*). That some dynamic content exhibits the first kind of behavior is encouraging: it frees up a surrogate or proxy cache designer to work with a range of object sizes in order to achieve benefits from reusability. The second kind of behavior lends itself to simple prediction schemes: each object falls in one or the other of the categories (not both).

- The Weibull nature of the freshness times distribution suggests that because a large number of objects get updated very frequently, client-initiated cache consistency protocols as opposed to server-initiated ones such as volume leases [41], may be more appropriate for these dynamic objects. Note that server-initiated protocols can still be effective for sites such as *amazon* and *wmedia*, which contain objects with relatively large freshness times.
- Knowledge of object size and freshness time distributions can be exploited by a cache designer who can incorporate them into heuristics to decide which objects to cache and which to prefetch. For example, based on the caching history, an object can be tagged as being more likely to provide an increased reuse opportunity if it has already been reused a certain number of times, and then be prefetched in advance.

3.7. Limitations of our approach

It is worth revisiting some of the assumptions inherent in our methodology, to ensure that the results of our study are interpreted correctly.

First, since we use an object-splitting technique based on synthetically chosen size and level limit parameters, it is quite possible that the objects in our analysis may differ from the real objects in the content (as identified at the server). In particular, two objects in the real situation may be treated as one by our analysis or vice versa. Thus, our characterization of object sizes, freshness times, and reusability is predicated on how closely our approach approximates the real situation. We should note that in the absence of additional server-side information (such as ESI [36]), it is conceivable that surrogate and cache components would employ an object identification technique very similar to ours.

Second, our conclusions are drawn from a relatively small set of dynamic content sites and may or may not be reflective of a broader selection. Our choice of the sites was driven by an effort to pick a representative sample of dynamic content sites, which additionally received high traffic. We also believe that unlike sites with static content, where there is much more diversity, most dynamic content sites rely on the same low-level tools for content generation and hence may exhibit more similarity with each other.

4. DYCE: dynamic content emulator

Although several researchers have proposed server-side and cache-side techniques for improving generation and delivery of dynamic content, these techniques have not as yet been widely adopted by either content providers or content deliverers. We attribute this situation to the following: content providers need to be first convinced of the benefits of such techniques before they modify their content to support composition and reuse at the granularity of sub-document objects. Unfortunately, it is difficult to verify the efficacy of a proposed technique in the absence of either real or synthetic content that adheres to the consensus view of dynamic Web content as consisting of a document template and individual objects.

In this section, we describe the design of a dynamic content emulator (DYCE), which alleviates the above

situation. DYCE uses the models of object size and freshness time obtained in the previous section to emulate a Web server that serves dynamic Web content, but additionally also incorporates support for document templates and identification of individual or multiple component objects.

4.1. Design

DYCE separates out the functions of content generation and content representation into two modules: the *dynamic content generator* (DCG) and the *dynamic content presenter* (DCP). DCP interfaces with requesters and appears as a traditional Web server augmented with additional functionality described later in the section (see Fig. 10).

To simplify its use, DCP supports a fairly simple interface, servicing two kinds of requests, either for the *document template* or for a set of *selected objects*. In response to requests of the first kind, the DCP checks a local cache to see if the document template exists, and if it is still valid. If both conditions are satisfied, it replies with the template as described below. If not, it requests the DCG to generate this document. The DCG takes three parameters as input to create the corresponding document template and individual objects: the *dynamic content type*, the *size limit* and the *level limit*. The latter two parameters have the same meanings as discussed earlier, while the first one is used by the DCG to index into a database of canned distributions. For example, the request `<cnn, 200, -1>` results in the DCG generating a `cnn`-like document (i.e., with the characteristics reflected in the `cnn` trace) with objects larger than 200 bytes and a level limit set to infinity.

Requests for a single object are handled similarly. The DCP first checks the cache to verify validity, and if the object freshness time has expired, then it requests the DCG to get a new version. The DCG services this request by looking up the distribution corresponding to the document template. The DCP extends this basic scheme to support requests for a set of objects.

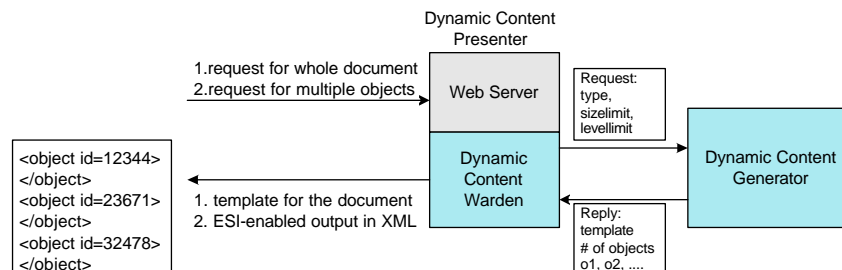


Fig. 10. The architecture of DYCE, which includes two components: a DCG and a Web server-like DCP. The extra functionality of DCP is provided by a module called the dynamic content warden.

4.2. Implementation

Our implementation builds the DCP on top of the Jakarta Tomcat 4.0 Web server infrastructure. All functionality for dealing with dynamic content (understanding templates, verifying validity, etc.) is encapsulated in a *dynamic content warden*, which is implemented as a servlet. The DCG is implemented as a standalone Java application, and interfaces with a database that contains parameters for modeling various kinds of dynamic content. Configuration parameters associated with the dynamic content warden allow simulation of multiple Web sites, and selection of different object granularities for a particular site. Both the DCG and the dynamic content warden are easy to extend to incorporate different representations of the document template and different distributions governing object size and freshness time characteristics.

The content presented to requesters by the DCP is currently packaged as an HTML document augmented with *ESI* directives. We considered using the *XSL-FO* formatting language from W3C, however felt that because of its complicated nature it would be a while before content providers and caching vendors would support this standard. *ESI*, on the other hand, benefits from the fact that it is a natural extension to server-side include (*SSI*), which is widely adopted.

Our implementation had to extend the *ESI* specification in two ways. These extensions, which concern requests and responses that involve multiple objects within the same document, might make sense to be included in the *ESI* specification as it evolves. Logically, to support *ESI*, downstream surrogates and/or proxy caches will send requests for all of the objects that are invalid in the local cache. To reduce the overhead for creating a different connection for each request, we have adopted an alternative, which combines the requests into a single message. Our suggestion for the combined requests is the comma separated values (*CSV*) format. For example, the request for three objects in our example might look like `http://webserver/pageServer?object=1,2,3` where the last three numbers refer

to unique IDs of objects. We also suggest using *XML* to describe the reply for multiple objects (see Fig. 10). The second extension associates freshness times with individual objects (instead of the entire document) by adding an attribute to the `<esi:include>` tag, which now looks like `<esi:include freshness='8' src='URL'>`. The source code of *DYCE* is available at <http://www.cs.wayne.edu/~weisong/dyce.htm>.

4.3. Validation of DYCE

To validate both *DYCE*, and as a consequence our analysis models, we developed an idealized proxy cache simulator that can be driven either from our collected traces or from emulated content generated by the *DYCE*. The output of the simulator is the number of response bytes transmitted between the server and the cache, assuming that requests for the Web page of interest arrive at the cache according to a poisson arrival model with mean $\lambda = 1$ min.

When driven by real trace data, the simulator computes the transferred bytes by comparing the currently cached document against the document corresponding to the request time, inferring the objects that have changed (using a method similar to that described in Section 2), and replaces the cache contents with the new document. When driven by *DYCE*, the simulator maintains a cache of the document objects and requests new versions upon detecting expiration: these requests are handled by the DCG, which looks up the database of distribution parameters (corresponding to the trace data) to send an object with size and freshness properties randomly drawn from the corresponding distributions. Fig. 11(a) and (b) show the results of the two cases while simulating client activity over a 1-day period: *RealSim* refers to the simulation using the trace data, while *DYCESim* refers to the simulation using *DYCE*.

As Fig. 11 shows, the *DYCE*-driven simulation is able to capture the behavior of the real trace data: the mean bandwidth requirements in the two cases are virtually indistinguishable, as is the spread between the lower and

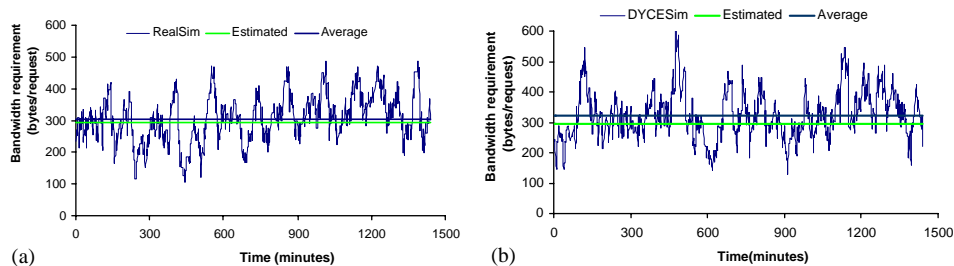


Fig. 11. The bandwidth requirements (bytes/request) in *RealSim* and *DYCESim*. The plots also show the predicted value of bandwidth, obtained as discussed in the text.

upper bounds of required bandwidth. More importantly, and this is something that the figure does not show, DYCE offers a significant advantage in simulation time attesting to its utility as an efficient tool for evaluating new techniques. In particular, to simulate the 1-day interval on a dual Pentium III-1 GHz machine with 512 MB of memory, the simulator took 0.5 s when driven by DYCE, in contrast to the 2.5 min that it took to process the real trace.

Fig. 11 also highlights an additional advantage of modeling the characteristics of dynamic content. In some situations, as here, these models can be used to obtain a quick analytical estimate of the values of interest. For example, we can use the following simple prediction model to estimate the average bandwidth requirement between the proxy cache and the server. Assuming that the document consists of N objects (o_1, o_2, \dots, o_N) , with sizes and freshness times (s_1, s_2, \dots, s_N) and (f_1, f_2, \dots, f_N) respectively, the average bandwidth requirement seen by an idealized proxy cache which receives requests according to a poisson-arrival model with mean λ is given by the following expression:

$$BW = \sum_{i=1}^M s_{z_i} \approx N \cdot \bar{S} \cdot F(\lambda), \quad (1)$$

where M is the number of objects that have expired when the request comes in, and (z_1, z_2, \dots, z_M) are the sequence numbers of those objects. F is the cumulative distribution function of the freshness times (f_1, f_2, \dots, f_N) and \bar{S} is the mean of the object sizes (s_1, s_2, \dots, s_N) . The approximation comes about because an average request, that arrives every λ time units, finds $N \cdot F(\lambda)$ objects whose freshness times have expired ($F(t)$ represents the fraction of objects with freshness time less than t). One can assume that each of these objects has an average size of (\bar{S}) , resulting in the expression above.

For the trace data simulated in this experiment, which consisted of $N = 137$ objects with an object size distribution given by exponential ($\lambda = 0.004$) and a freshness time distribution given by Weibull ($F(\lambda) = 0.0086$), the estimated bandwidth works out to be 294 bytes/request. As shown in Fig. 11, this estimate closely approximates the bandwidth requirements seen by both RealSim and DYCESim.

5. Related work

Web workload characterization has been extensively studied in the past 5 years from the perspective of proxies [7,16,38,40], client browsers [5,13,24], and servers [4,30]. However, many of these studies focus on the characteristics of Web resources at the granular-

ity of the whole document, examples of which include content type, resource size, response size, resource popularity, modification frequency, temporal locality, client access pattern, and the number of embedded resources. These previous research results accurately capture the characteristics of static Web content. However, for dynamic Web content which introduces the notion of object composition, several of these characteristics, such as request and response sizes need to be revisited. Moreover, dynamic content necessitates understanding of new characteristics, such as the number and sizes of objects making up a document, and the freshness times of these objects. To the best of our knowledge, the work described in this paper is one of the first efforts trying to model these latter characteristics, and complements our recent work on the characterization of a channel-based personalized Web site [35].

Although closest in spirit, the work by Brewington and Cybenko [8] and Douglass et al. [16] on understanding the dynamics of the Web differs from our efforts in two aspects. First, both these works focus on characteristics at the document level, while we are interested in those at the sub-document. Second, only dynamic Web content (which changes very frequently) is used in our analysis, while the Web content analyzed in their work has broader characteristics, with a large fraction actually corresponding to static content. Interestingly, the Weibull distribution of freshness time of objects found in our work is similar to that of the expected change time across pages found in [8]. We believe this behavior originates because of self-similarity characteristics between changes at the sub-document and document levels; however, additional work is needed to verify this hypothesis.

Wills and Mikhailov [39] have quantitatively analyzed the content reusability present in traces collected from several Web sites over a 1-day interval. Two notable differences of our work include the characterization of content reusability at finer granularity (making explicit the time dependence), and the relationship between object size and content reusability. Our results indicate the object granularity that must be supported in order to successfully take advantage of potential reusability.

Challenger et al. [11] analyzed object size distributions based on server traces from the 2000 Olympics site. Although related to our objectives, their work employs a different definition of what an object is: they include not only the individual objects embedded within a document, but also the entire generated document itself. Because of this reason, while a Pareto distribution fits their findings, we conclude that the sizes of individual objects in the document follow an exponential distribution.

Our work on the dynamic content emulator resembles work done by Barford et al. on the SURGE static

workload generator [6]. However, unlike SURGE, which was designed to model client access patterns, DYCE focuses on the complementary goal of emulating server behavior, both in terms of its load properties as well as the nature of the content itself. To simulate an appropriate delay to model overheads of dynamic content generation, our emulator uses delay models from previous research [2,21,23,25,31].

The work in this paper was motivated in part by our inability to extend, to our specific setting, the results previously obtained by researchers working on various aspects of dynamic and personalized content delivery. Such work, which has focused on both server-side [10,11,15,42] and cache-side [9,17,27,29,39] has typically been validated with specific, proprietary workloads. For example, Challenger et al. used the 1998 Olympic winter games workload in [10], and the 2000 Olympic games workload in [11], and Douglis et al. used a modified internal AT&T Web-based “recruiting database” to evaluate their idea of HPP [17]. The work in this paper addresses this shortcoming, providing both models and tools that allow researchers to efficiently evaluate their techniques on a variety of (synthetic) workloads. We must note that our intention is not to replace experimentation with real workloads, but instead to complement it.

The work reported in this paper represents a particular view of dynamic content, which currently dominates traffic on the Internet. With proposed techniques such as the class-based page classification scheme in [42], the Active Cache system proposed by Cao et al. [9], and more recently the Gemini system proposed by Myers et al. [29], which require Web servers to provide both content and specialized code (applets in Active Cache, Java classes in Gemini) to proxy caches, it is clear that the nature of dynamic content and therefore its characteristics will evolve in the future. While this may require that we revisit some of our results, we believe that the methodology we have adopted is sound and can be used to characterize these new forms of content as and when they become available.

6. Summary

This paper has proposed a methodology for evaluating characteristics of dynamic Web content, and used this methodology to obtain models for various independent and derived metrics of interest such as object sizes, freshness times, and content reusability. To summarize our main findings, we have found that (1) object sizes in dynamic documents can be modeled using an exponential or Weibull distribution; (2) for document objects that change over time, their freshness times can be modeled in terms of a Weibull distribution, which sometimes degenerates to a case where all of the objects

change on every access; (3) there is significant opportunity for object reuse across both multiple accesses of the same document as well as accesses of related documents. Our results have also made explicit the dependence between the size at which document objects are managed and the corresponding reuse potential.

These models have also served as the foundation for the design of a tool, DYCE, which emulates a Web server serving dynamic content, both in terms of the load properties as well as the nature of the generated content. DYCE incorporates emerging standards such as ESI, permitting its use for evaluating object composition-based optimizations for improving dynamic content delivery at the level of surrogates or caches.

References

- [1] Akamai Technologies Inc., Edgesuite services, http://www.akamai.com/html/en/sv/edgesuite_over.html.
- [2] M.A. Mendes, V.A. Almeida, Analyzing the impact of dynamic pages on the performance of web servers, Proceedings of the Computer Measurement Group Conference, Anaheim, CA, December 1998.
- [3] Apache Jakarta Project, <http://jakarta.apache.org>.
- [4] M. Arlitt, C. Williamson, Internet web servers: workload characterization and performance implications, IEEE/ACM Trans. Networking 5 (5) (1997) 631–645.
- [5] P. Barford, A. Bestavros, A. Bradley, M.E. Crovella, Changes in web client access patterns: characteristics and caching implications, World Wide Web, Characterization and Performance Evaluation 2 (Special issue) (1999) 15–28 <http://www.cs.bu.edu/techreports/1998-023-web-client-trace-changes-and-%implications.ps.Z>.
- [6] P. Barford, M.E. Crovella, Generating representative web workloads for network and server performance evaluation, Proceedings of the Performance '98/ACM SIGMETRICS '98, Madison, WI, July 1998.
- [7] L. Breslau, P. Cao, L. Fan, G. Phillips, S. Shenker, Web caching and zipf-like distributions: evidence and implications, Proceedings of the IEEE conference on Computer Communications (INFOCOM'99), March 1999, pp. 126–134, <http://www.research.att.com/~breslau/papers/zipf.ps.gz>.
- [8] B.E. Brewington, G. Cybenko, How dynamic is the web? Proceedings of the 9th International World Wide Web Conference (2000), Amsterdam, Netherlands, May 2000.
- [9] P. Cao, J. Zhang, K. Beach, Active cache: caching dynamic contents on the web, Proceedings of the IFIP International Conference on Dist. Sys. Platforms and Open Dist. Processing, Lake District, England, 1998, pp. 373–388, <http://www.cs.wisc.edu/~cao/papers/active-cache.ps>.
- [10] J. Challenger, A. Iyengar, P. Dantzig, A scalable system for consistently caching dynamic web data, Proceedings of the IEEE Conference on Computer Communications (INFOCOM'99), New York, NY, March 1999.
- [11] J. Challenger, A. Iyengar, K. Witting, C. Ferstat, P. Reed, A publishing system for efficiently creating dynamic web content, Proceedings of the IEEE Conference on Computer Communications (INFOCOM'00), Tel-Aviv, Israel, March 2000.
- [12] M.E. Crovella, A. Bestavros, Self-similarity in world wide web traffic: evidence and possible causes, Proceedings of the Performance '96/ACM SIGMETRICS '96, Philadelphia, PA, May 1996.

- [13] M.E. Crovella, A. Bestavros, Self-similarity in world wide web traffic: evidence and possible causes, *IEEE/ACM Trans. Networking* 5 (6) (1997) 835–846 <http://www.cs.bu.edu/fac/best/res/papers/ton97.ps>.
- [14] R.B. D'Agostino, M.A. Stephens (Eds.), *Goodness-of-Fit Techniques*, Marcel Dekker, Inc., New York, 1986.
- [15] A. Datta, K. Dutta, H. Thomas, D. VanderMeer, K. Ramamritham, Accelerating dynamic web content generation, *IEEE Internet Comput.* 6 (5) (2002) 26–35.
- [16] F. Douglis, A. Feldmann, B. Krishnamurthy, J. Mogul, Rate of change and other metrics: a live study of the world wide web, *Proceedings of the First USENIX Symposium on Internet Technologies and Systems (USITS'97)*, Monterey, CA, December 1997, pp. 147–158, <http://www.douglis.org/fred/work/papers/roc.pdf>.
- [17] F. Douglis, A. Haro, M. Rabinovich, HPP:HTML macro-preprocessing to support dynamic document caching, *Proceedings of the First USENIX Symposium on Internet Technologies and Systems (USITS'97)*, Monterey, CA, December 1997, pp. 83–94, <http://www.douglis.org/fred/work/papers/hpp.pdf>.
- [18] R. Doyle, J. Chase, S. Gadde, A. Vahdat, The trickle-down effect: web caching and server request distribution, *Proceedings of the Sixth International Workshop on Web Caching and Content Distribution (WCW'01)*, June 2001.
- [19] F. 180-1, Secure Hash Standard. U.S. Department of Commerce/N.I.S.T., National Technical Information Service, Springfield, VA, April 1995.
- [20] J. Gecsei, Determining hit ratios for multilevel hierarchies, *IBM J. Res. Dev.* 18 (1974) 316–327.
- [21] Y. Hu, A. Nanda, Q. Yang, Measurement, analysis and performance improvement of apache web server, *Proceedings of the IEEE International Performance, Computing, and Communications Conference*, Boston, MA, February 1999.
- [22] IBM Corp, Websphere platform, Phoenix, AZ, <http://www.ibm.com/websphere>.
- [23] A. Iyengar, J. Challenger, D. Dias, P. Dantzig, High-performance web site design techniques, *IEEE Internet Comput.* 4 (2) (March/April 2000) 17–26.
- [24] T. Kelly, Thin-client web access patterns: measurements for a cache busting proxy, *Proceedings of the Sixth International Workshop on Web Caching and Content Distribution (WCW'01)*, Boston, MA, June 2001, <http://www.cs.bu.edu/techreports/2001-017-wcw01-proceedings/129.beck.pdf>.
- [25] B. Kothari, M. Claypool, Performance analysis of dynamic web page generation technologies, *Proceedings of the International Network Conference (INC)*, Plymouth, UK, July 2000.
- [26] B. Krishnamurthy, J. Rexford, *Web protocols and practice: HTTP/1.1, Networking Protocols, Caching and Traffic Measurement*, Addison-Wesley, Inc., Reading, MA, 2001.
- [27] M. Mikhailov, C.E. Wills, Change and relationship-driven content caching, distribution and assembly, Technical Report WPI-CS-TR-01-03, Computer Science Department, WPI, March 2001, <http://www.cs.wpi.edu/~cew/papers/tr01-03.pdf>.
- [28] J.C. Mogul, F. Douglis, A. Feldmann, B. Krishnamurthy, Potential Benefits of Delta-Encoding and Data Compression for HTTP, *Proceedings of the 13th ACM SIGCOMM'97*, Cannes, France, September 1997, pp. 181–194, <http://www.douglis.org/fred/work/papers/sigcomm97.pdf>.
- [29] A. Myers, J. Chuang, U. Hengartner, Y. Xie, W. Zhang, H. Zhang, A secure and publisher-centric web caching infrastructure, *Proceedings of the IEEE Conference on Computer Communications (INFOCOM'01)*, Anchorage, AK, April 2001.
- [30] V.N. Padmanabhan, L. Qiu, The content and access dynamics of a busy web site: findings and implications, *Proceedings of the ACM SIGCOMM'00*, Stockholm, Sweden, August 2000, <http://research.microsoft.com/~liliq/papers/pub/sigcomm2000.ps>.
- [31] V. Pai, P. Druschel, W. Zwaenepoel, Flash: an efficient and portable web server, *Proceedings of the 1999 USENIX Annual Technical Conference*, Monterey, CA, June 1999.
- [32] V. Paxson, Empirically-derived analytic models of wide area TCP connections, *IEEE/ACM Trans. Networking* 2 (4) (1994) 316–336.
- [33] D. Raggett, HTML tidy tool, 2000, <http://www.w3c.org/People/Raggett/tidy>.
- [34] W. Shi, V. Karamcheti, CONCA: an architecture for consistent nomadic content access, *Workshop on Cache, Coherence, and Consistency (WC3'01)*, Sorrento, Italy, June 2001.
- [35] W. Shi, R. Wright, E. Collins, V. Karamcheti, Workload characterization of a personalized web site and its implication on dynamic content caching, *Proceedings of the Seventh International Workshop on Web Caching and Content Distribution (WCW'02)*, August 2002, pp. 1–16, <http://www.cs.wayne.edu/~weisong/papers/wcw02.pdf>.
- [36] M. Tsimelzon, B. Weihl, L. Jacobs, ESI language specification 1.0, 2000, <http://www.esi.org>.
- [37] W3C XSL Working Group, <http://www.w3.org/Style/XSL/>.
- [38] D. Wessels, *Web Caching*, O'Reilly Inc., Sebastopol, CA, 2001.
- [39] C.E. Wills, M. Mikhailov, Studying the impact of more complete server information on web caching, *Proceedings of the Fifth International Workshop on Web Caching and Content Distribution (WCW'00)*, 2000.
- [40] A. Wolman, G.M. Voelker, N. Sharma, N. Cardwell, A. Karlin, H.M. Levy, On the scale and performance of cooperative web proxy caching, *Proceedings of the 17th ACM Symposium on Operating Systems Principles (SOSP)*, December 1999, pp. 16–31.
- [41] J. Yin, L. Alvisi, M. Dahlin, C. Lin, Hierarchical cache consistency in a WAN, *Proceedings of the Second USENIX Symposium on Internet Technologies and Systems (USITS'99)*, October 1999.
- [42] H. Zhu, T. Yang, Class-based cache management for dynamic web content, *Proceedings of the IEEE, Conference on Computer Communications (INFOCOM'01)*, April 2001, <http://www.cs.ucsb.edu/projects/swala/cache2001.ps>.