

# Cache Satellite Distribution Systems: Modeling, Analysis, and Efficient Operation

Aner Armon and Hanoch Levy, *Member, IEEE*

**Abstract**—Web caches have become an integral component contributing to the improvement of the performance observed by Web clients. Cache satellite distribution systems (CSDSs) have emerged as a technology for feeding the caches with the information clients are expected to request, ahead of time. In such a system, the participating proxies periodically report to a central station about requests received from their clients. The central station selects a collection of Web documents, which are “pushed” via a satellite broadcast to the participating proxies, so that upon a future local request for the documents, they will already reside in the local cache, and will not need to be fetched from the terrestrial network.

In this paper, our aim is addressing the issues of how to operate the CSDS, how to design it, and how to estimate its effect. Questions of interest are: 1) what Web documents should be transmitted by the central station and 2) what is the benefit of adding a particular proxy into a CSDS? We offer a model for CSDS that accounts for the request streams addressed to the proxies and which captures the intricate interaction between the proxy caches. Unlike models that are based only on the access frequency of the various documents, this model captures both their frequency and their locality of reference. We provide an analysis that is based on the stochastic properties of the traffic streams that can be derived from HTTP logs, examine it on real traffic, and demonstrate its applicability in selecting a set of proxies into a CSDS.

**Index Terms**—Communication system performance, Internet, modeling, queueing analysis, satellite applications, satellite broadcasting, system analysis and design.

## I. INTRODUCTION

WEB CACHES have become an integral component contributing to the improvement of the performance observed by Web clients. Cache satellite distribution systems (CSDS) and content distribution networks (CDN) have emerged as technologies for feeding the caches with the information clients are expected to request, ahead of time. A typical CSDS, of which a schematic drawing is given in Fig. 1, consists of a set  $P$  of proxy caches and one central station. The proxies participating in the CSDS periodically report to the central station about the requests they are receiving from their local clients. The central station processes this information and uses it to predict what documents will be desired by other proxy caches in the system in the near future. The central station

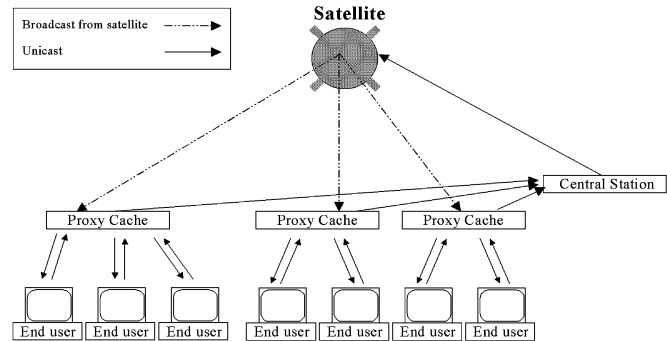


Fig. 1. Schematic diagram of the system.

then selects a collection of Web documents, which it retrieves, usually from the terrestrial network,<sup>1</sup> and then “pushes” these documents via a satellite broadcast link to all, or some, of the participating proxies. The result is that upon such request, the documents will reside in the local cache and will not need to be fetched using the local terrestrial network.

The advantage of broadcasting a document over CSDS is that once the document is requested at any participating proxy, the document is available in all caches. Thus, the user delay is reduced and the bandwidth cost is saved (assuming that the cost of broadcasting a document to  $K$  destinations is cheaper than to retrieve it on the terrestrial network  $K$  times). Nonetheless, such benefits do depend on whether the document will indeed be needed at the receiving proxy. In the case that the document is *not needed*, no benefit is gained. In fact, some damage may be caused, since the unwanted document “contaminates” the cache by pushing all documents residing in the cache, and potentially causing another document (that may be needed by the proxy) to be removed from the cache. Thus, broadcasting a document over CSDS is not always beneficial.

In a similar manner, the benefit for a proxy from participating in a CSDS may vary, depending on the mutual properties of the participating proxies. To demonstrate this, consider two proxies,  $A$  and  $B$ , which consider sharing a CSDS. If the proxies have interest in similar documents then it is likely that they will mutually benefit from a CSDS. For example, such mutual benefit is expected from two proxies of commercial Internet service providers (ISPs) in the same country. On the other hand, consider proxies  $A$  and  $B$  located in two different countries whose audience is interested mainly in their native language documents. In this case, the documents requested by the proxies are highly disjoint and the use of a CSDS will only cause the two proxies to “contaminate” each other.

<sup>1</sup>Although satellite retrieval does not affect the analysis.

Manuscript received December 15, 2002; revised July 1, 2003 and September 20, 2003. This work was supported in part by MAGNET, Chief Scientist Office, Ministry of Trade and Commerce, Israel. This paper was presented in part at the IEEE INFOCOM 2003 Conference, San Francisco, CA.

The authors are with the School of Computer Science, Tel-Aviv University, Tel-Aviv 44925, Israel (e-mail: aner@post.tau.ac.il; a\_armon@iname.com; hanoch@cs.tau.ac.il).

Digital Object Identifier 10.1109/JSAC.2003.819970

The aim of this paper is to devise an analytic tool that can be used in the design decisions involved in the operations of a CSDS. These include issues such as the selection of classes of Web documents for broadcast in the system and the selection of proxies to participate in the system. To this end, one must take into consideration the properties of the various proxies and the properties of the traffic streams<sup>2</sup> they are exposed to and to use this data to analyze the system. Two properties are inherent to streams directed at Web caches (as to other caches as well). These are the relative request frequencies of the various documents (some documents are accessed more frequently and some are less) and the locality of reference (a document that was requested recently is more likely to be requested again). Unfortunately, accounting for both properties is difficult and most common models do not account for both.

We propose a stochastic model that accounts for both properties of the request streams directed to a proxy. The statistics of these properties can be derived from the HTTP log of the proxy. The model uses these properties to predict the hit ratio, which is the fraction of the requests made to the cache and which result with the requested document found at the cache. Thus, using the analysis one can evaluate system performance under various conditions.

The model accounts for the intricate interaction of the different streams at the different caches and makes use of two document request distributions, associated with each stream: The stack depth distribution and the request count distribution. The stack depth distribution has been used successfully to model the behavior of cache systems (see, e.g., [8]), and in [12] and [13] it was used to model and analyze the interaction of database caches. The request count distribution helps us to estimate the potential gain of sharing the first miss of every document, among all participating proxies.

Our analysis yields a set of recursive equations whose computation complexity is  $O((K\alpha)^2)$ , where  $K$  is the cache size, and  $\alpha \leq 1$  is a modeling parameter that can be chosen to be quite small as to make the computation very efficient. The results of the analysis provide expressions for the improvement in the cache hit ratio in each of the participating proxies, as function of the document classes (or streams) broadcast to participating proxies. Thus, the analysis can be used for selecting the more effective streams to be broadcast, as well as for assessing the value of joining a proxy into a CSDS.

The analysis is supported by numerical results in which we examine the predictions of the analysis by comparing it to simulation. We observe good agreement between the simulation and the analysis especially in predicting the relative effect of different streams or different proxies on the performance. As such, the analysis fits well in operational rules and system design assistance.

We further demonstrate how the model can be used in designing a CSDS. In this context, we use the model to predict which proxies it is mostly beneficial (in terms of maximizing the hit rate gain) to add into a CSDS. We use simulation results conducted on real data to examine the quality of this prediction and find the prediction to be very good.

<sup>2</sup>Where a stream is a set of Web document requests sharing some commonality.

The structure of the rest of this paper is as follows: In Section II, we provide the model and the analysis approach. In Section III, we provide the analysis of the system. In Section IV, we discuss how the results can be used for devising operation rules and designing the system. Numerical results are provided in Section V. Section VI provides experimental results, based on real data, for the proxy assignment problem. Finally, Section VII summarizes the paper and its results.

### A. Related Work

Cache modeling and analysis received attention in the past in the context of operating systems and databases. For example, see Coffman and Denning [8]. In the context of databases, Levy and Morris [12], [13] proposed a model and analysis for evaluating the interaction of various traffic streams in a cache system. A similar cache modeling is addressed in Brehob and Enbody [5]. Recently, cache modeling on the World Wide Web received attention as well. For example, Breslau *et al.* [6] tested the sufficiency of using Zipf distribution to describe behavior of Web proxies. They provided a comparison to experimental observations, and proposed some cache replacement algorithms. Some attention was also given to properties of different request types on the Web, e.g., Almeida *et al.* [2], that provided classification and interpretation of specific proxy logs, based on geographical characterization.

The work on the modeling and analysis of CSDSs has been quite limited. Rodriguez and Biersack [14] provided an analysis of the performance of CSDSs, but without accounting for modeling aspects of the cache capacity, the interrelation between the document requests and their interreference, and the effects of the different streams on the cache performance. Hu *et al.* [10], provided an analytic model based on Poisson distributed request rate and a suggestion of client filtering policy in a CSDS based on Web servers visited in previous days. Chang [7] formalized the cache satellite distribution problem as an optimization problem, based on the assumption that documents are distributed according to the Zipf distribution; that paper does not account for the interreference of document requests as well.

Several references focused on the problem of the selection of documents to be broadcast, based on very detailed information dynamically sent to the central station from the proxies regarding their needs of specifically identified documents. Such scheduling problems apply for satellite broadcast (e.g., Cohen *et al.* [9] and Askoy *et al.* [4]), and for other applications (e.g., Acharya and Muthukrishnan [1], Sharaf and Chrysanthis [15]). These may require a significant amount of data sent from the proxies to the central server, and just like all other studies mentioned above, they do not deal with the system design issues that are being handled in this paper.

## II. GENERAL MODEL AND ANALYSIS APPROACH

### A. System Modeling and Analysis Approach

A CSDS can be modeled as follows. Each cache is subject to a stream of local document requests, originated by the end users of the proxy. Upon receiving a local request, the proxy will examine whether the requested document is in the cache,

in which case it will return the requested document to the user. If the requested document is not in the cache, the proxy will retrieve it from the terrestrial network and will announce the central station that the document was requested and was not found. The central station then will decide (typically by identifying the document as a part of some stream, defined herein) whether to retrieve the document from the terrestrial network and to broadcast it to all the participating proxies via the satellite link. Upon receiving a document from the satellite link, a proxy may select to either save it in its cache or to discard it.

*Remark 1:* The notion that not all participating proxies must accept all broadcast documents and put them in their cache, can actually split the set of recipient proxies of the satellite broadcast to several subsets, by common proxy interests. We will model this behavior by defining subsets of proxies, each sharing interest in similar document classes, i.e., only a portion of each document class, originating at each subset, is being broadcast. Most of our analysis focuses on a single subset, but is constructed in a way that can handle multiple proxy subsets.

Our aim in the analysis is to examine how the broadcast of a document or a class of documents (e.g., all the documents of xyz.com) will affect the performance of the individual caches. This analysis is to be carried as function of the parameters of the document class, as to assist us in the decision of which documents (or classes of documents) should be broadcast and which not. We assume that the system designer is equipped with the HTTP logs of the proxies (these are commonly available), and that the statistical data available in the logs can be used by the analysis to predict the relative benefit (or damage) of broadcasting a class. This benefit will be used to decide which classes of documents to broadcast and which proxies to join to the system.

Two models that were used to analyze cache systems are the independent reference model (IRM) and the stack reference model (SRM). Their advantage is in capturing both the locality of reference of the requests and their relative frequencies, which are both important properties for cache modeling. In contrast, other models, such as the common used Zipf model (see, e.g., [6]), capture only one of these properties. SRM has been used successfully in [12] and [13] to model the interaction between various caches. More specifically, it was used to examine the performance resulting from merging several disjoint streams into one cache. Due to the interaction between the caches in CSDS, it is appealing to attempt modeling CSDS by the SRM model. That approach however, seems to be hard to apply since the interaction between the streams in CSDS is much more complicated than in the problem addressed in [12] and [13]. This is true since this problem includes interaction of streams that are not mutually exclusive and it also includes cross-cache effects.

To overcome the complexity of the problem, we decompose the problem and its analysis to two parts, in each of which using suitable modeling tools to focus on capturing the major factors. The approach is as follows. Consider a document  $R$  that is broadcast by the central station to cache-proxy  $C_p$ , which stores it in the cache. The storing of  $R$  in  $C_p$  affects the performance of  $C_p$  in two aspects.

- 1) It affects the hit ratio of future requests to  $R$ . This effect is nonnegative (that is the hit ratio of these requests may only go up due to the storage of  $R$ ).

- 2) It affects the hit ratio of future requests to documents other than  $R$ . This effect is nonpositive since the presence of  $R$  in  $C_p$  may push the other documents out of the cache and decrease their hit-ratio.

We will focus on the *change in performance* of  $C_p$  due to the broadcast, that is, how the hit ratio changes due to the broadcast operation. As was previously hinted, we decompose our analysis to two cases.

Case 1) Here, we deal with the situation at which there is *at least one request* to  $R$  by  $C_p$ . This is either the request that triggered the broadcast of  $R$ , or a later request. Let  $E_B$  denote the broadcast-triggering event, and let  $E_p$  denote the event of the first request to  $R$  by  $C_p$  ( $E_B$  and  $E_p$  may coincide). At this case, the *change in performance* due to the broadcast of  $R$  (and its storage by  $C_p$ ) is twofold.

- a) A nonnegative effect on the hit ratio of requests to  $R$  (these are the requests to  $R$  happening after  $E_p$ ).
- b) A nonpositive effect on requests to documents other than  $R$ . A more precise examination shows that this effect is limited to the period between the events  $E_B$  and  $E_p$ . This limitation applies, since after the epoch at which  $E_p$  occurs,  $R$  is present in  $C_p$  regardless of the broadcast operation  $E_B$ , and thus the *change in performance*, due to  $E_B$ , past the epoch of  $E_p$ , is zero.

Case 2) Here, we deal with the situation at which there is *no request* to  $R$  by  $C_p$ . In this case, the broadcast *does not affect* the hit ratio for requests to  $R$ . However, the broadcast *does affect* (nonpositively) the hit ratio of documents other than  $R$ , since  $R$  now “contaminates”  $C_p$ .

In our analysis, we will neglect the effects on performance of Case 1b) above, due to its limited scope (which makes it negligible compared with the other effects). Our analysis will, therefore, focus on the above Case 1a) (Section III-A) and Case 2 (Section III-B). The reason the effect of Case 1b) is so minor, is because we assume that the probability of  $R$  to be accessed locally is much higher than the probability the document it replaced in the cache will be requested. The reason for this assumption to hold, is because of the inherent property of every stream (that a model such as SRM captures so naturally), that a page that was just requested has a much higher probability to be requested again, than a page that was not requested for a long time (so it is thrown out of the cache). The assumption of negligibility of Case 1b) tends to be less accurate, of course, if the cache of  $C_p$  is very small compared with its local stack depth distribution, or the rate of  $R$ 's class is much smaller than the rate of the class of the page that was thrown out.

## B. Document Request Modeling and Model Notation

We consider a set  $P$  of proxy caches  $C_1, \dots, C_{|P|}$ . Considering the entire incoming requests set for all proxies, we assume that we can classify this request set as a set  $S$  of disjoint streams, whereas each stream  $s \in S$  represents all requests for a class of

documents, i.e., a collection of documents that share some commonality. For example, a stream may represent all requests to all Web pages of xyz.com, or alternatively, only to the home page of abc.com. Intuitively speaking, one would tend to break large Web sites (e.g., Yahoo.com) into multiple streams (e.g., based on the site's directory structure) and join very small sites into a single stream.

We further assume that each stream  $s$  is associated with a random variable  $N_s$  denoting the request count, which is the number of requests that are made for a specific document  $R \in s$ , from the *whole* set  $P$ , until document  $R$  expires. Such expiration may represent document refresh, document removal from its hosting Web server or expiration of the document at all proxy caches. In all cases, the implication is that after the document is requested  $N_s$  times it may never be requested again or will be considered new and will have to be read again by at least one of the proxy caches. Let the *request count distribution*, which is the distribution of  $N_s$ , be denoted as  $c_s(n) = \Pr[N_s = n]$ . We further assume that the association of a request with a specific proxy  $C_p \in P$  is given by a Bernoulli process, that is the probability that an arbitrary request of stream  $s$  originates at proxy  $C_p$  is given by  $\lambda_{s,p}$ , where  $\sum_{p \in P} \lambda_{s,p} = 1$ . To clarify the last definition, note that  $\lambda_{s,p}$  determines only the distribution among different streams and proxies, while the actual document within the stream is selected according to the stream's distribution.

Since we want to construct our model in a way that can handle multiple proxy groups (see Remark 1), we will assume only a subset of proxies  $P' \subseteq P$  needs to participate in the CSDS. Let  $\Lambda_{s,P'} = \sum_{p \in P'} \lambda_{s,p}$  denote the probability that a request of stream  $s$  originates at  $P'$ .

We assume that the association of a request with a specific stream  $s \in S$ , is also given by a Bernoulli process. For the clarity of presentation, we can define the combined ratio  $\gamma_{s,p}$ , which is the probability that an arbitrary request originates at stream  $s$  and at proxy  $C_p$ . We can further widen the definition to stand for any subset  $S^*$  of streams, and/or for any subset  $P^*$  of proxies, yielding  $\gamma_{s,P^*} = \sum_{p \in P^*} \gamma_{s,p}$ ,  $\gamma_{S^*,p} = \sum_{s \in S^*} \gamma_{s,p}$ , and  $\gamma_{S^*,P^*} = \sum_{p \in P^*} \sum_{s \in S^*} \gamma_{s,p}$ .

*Remark 2:* Note that we will never use any  $\gamma$  as a stand alone value, but only as ratios of  $\gamma$ 's—so that  $\gamma$  can be represented in any units—requests per time unit, requests per log length, etc. To demonstrate, we can use  $\gamma$  for an alternative definition of  $\lambda_{s,p}$ , as  $\lambda_{s,p} = (\gamma_{s,p}/\gamma_{s,P})$ , and we can immediately see that the previous summation,  $\sum_{p \in P} \lambda_{s,p} = \sum_{p \in P} (\gamma_{s,p}/\gamma_{s,P}) = 1$ , still holds.

*Remark 3:* More detailed modeling of the SRM, which is specific to Section III-B, is given within that section.

### III. ANALYSIS

#### A. Effect of Broadcasting Document $R$ on the Performance of Requests to $R$

In this section, we aim at evaluating the potential performance improvement resulting from broadcasting stream  $s$  to a subset of proxies  $P' \subseteq P$ . We will consider a tagged document  $R \in s$ , whose number of requests is given by the random variable  $N_s$ , as noted in Section II-B. In addition to the modeling assumptions given in Section II, we further assume that all proxies has large enough caches to hold the document from the first request until

the last local request to the proxy is made. This assumption is reasonable for capturing the behavior of document  $R$ .

Consider first a system *without CSDS*. Under this system, for every document under every proxy, the first request for the document will result with a cache miss (since the document is not in the cache yet). All later requests from that proxy to that document will result with a hit since the document will be present in the cache. Focusing on cache  $C_p$  and conditioning on  $N = n$ , the probability that  $C_p$  will experience a (single) miss on the tagged document  $R$  is given by the probability that out of the  $n$  requests made to  $R$ , it will be requested at least once at  $C_p$

$$\Pr[\text{document miss at } C_p \mid N = n, \text{ no CSDS}] = 1 - \overline{\lambda_{s,p}}^n \quad (1)$$

where  $\overline{\lambda_{s,p}} := (1 - \lambda_{s,p})$ .

Let  $M_{s,p}$  be the total *request miss ratio* of stream  $s$  at  $C_p$  with no CSDS, defined to be the expected number of misses per request, and which equals to the ratio between the number of misses and the expected number of requests made at  $C_p$ . Using (1) and the definition of the stream's request count distribution  $c_s(n)$ , given in Section II-B, we can derive

$$\begin{aligned} M_{s,p} &= \frac{\sum_{n=1}^{\infty} c_s(n) (1 - \overline{\lambda_{s,p}}^n)}{\sum_{n=1}^{\infty} c_s(n) \lambda_{s,p} n} \\ &= \frac{1 - \sum_{n=1}^{\infty} c_s(n) \overline{\lambda_{s,p}}^n}{E[N] \lambda_{s,p}}. \end{aligned} \quad (2)$$

Now, consider the system with CSDS, in which when the document is requested by the first proxy to request it, the document is broadcast to all other proxies in its subset  $P'$ . In this case, the first proxy in  $P'$  will experience a single miss while none of the other proxies in  $P'$  will experience any miss.

Under this case,  $C_p$  experiences a miss on the tagged document, only if there is at least one request to the document that originates from a proxy in  $P'$ , and the first such request was issued by  $C_p$ . This event is given by

$$\begin{aligned} \Pr[\text{miss at } C_p \mid N = n, \text{ with CSDS}] \\ &= \frac{\lambda_{s,p}}{\Lambda_{s,P'}} (1 - \overline{\Lambda_{s,P'}}^n). \end{aligned} \quad (3)$$

Let  $M_{s,p}^{\text{CSDS}(P')}$  be the request miss ratio of stream  $s$  at  $C_p$  using a CSDS system connecting the subset of proxies  $P'$ , defined similarly to  $M_{s,p}$ . Its value is given by

$$\begin{aligned} M_{s,p}^{\text{CSDS}(P')} &= \frac{\sum_{n=1}^{\infty} (c_s(n) \frac{\lambda_{s,p}}{\Lambda_{s,P'}} (1 - \overline{\Lambda_{s,P'}}^n))}{\sum_{n=1}^{\infty} c_s(n) \lambda_{s,p} n} \\ &= \frac{1 - \sum_{n=1}^{\infty} (c_s(n) \overline{\Lambda_{s,P'}}^n)}{E[N] \Lambda_{s,P'}}. \end{aligned} \quad (4)$$

*Remark 4:* Note the similarity in form of (4) to that of (2). The only difference is that while (2) is dependent on  $\lambda_{s,p}$ , (4) is dependent on  $\Lambda_{s,P'}$ . The explanation is that from a hit/miss point of view, under CSDS the collection of participating proxies behaves as one large proxy, since they all share the same single miss per document. Note also that under CSDS, the miss ratio is independent of  $\lambda_{s,p}$  and of  $p$ , and is the same for every participating proxy.

Let us define  $G_{s,p}^{\text{CSDS}(P')}$ , the hit ratio gain of  $C_p$  due to using CSDS on stream  $s$ , which is simply

$$G_{s,p}^{\text{CSDS}(P')} = M_{s,p} - M_{s,p}^{\text{CSDS}(P')}. \quad (5)$$

Now, we can calculate  $G_p^{\text{CSDS}(P')}$ , the total hit ratio gain of  $C_p$  from using CSDS on all streams in  $S$ . This is done using the probabilities that a request of proxy  $C_p$  belongs to stream  $s$

$$G_p^{\text{CSDS}(P')} = \sum_{s \in S} \frac{\gamma_{s,p}}{\gamma_{S,p}} G_{s,p}^{\text{CSDS}(P')} = \frac{\sum_{s \in S} \gamma_{s,p} G_{s,p}^{\text{CSDS}(P')}}{\gamma_{S,p}}. \quad (6)$$

*Remark 5:* Note that for every stream  $s$  that is not requested by  $C_p$ ,  $\gamma_{s,p} = 0$ , and for all streams that are not broadcast, by definition  $G_{s,p}^{\text{CSDS}(P')} = 0$ . Practically, we can, therefore, define a subset of streams  $S' \subseteq S$  of broadcast streams, and to sum  $G_{s,p}^{\text{CSDS}(P')}$  only on  $S'$  ( $S'$  will be used in Section IV).

### B. Effect of Broadcasting Document $R$ to a Proxy That Will Not Request $R$

In the previous section, we assumed that every proxy cache is of sufficient size—so it can contain all local accessed documents, as well as additional CSDS broadcast documents, as many as we choose to broadcast. Under that assumption, and if broadcast costs or constraints are not accounted for, it is clearly optimal to broadcast *all streams* to *all proxies*, because there is no constraint on the proxy cache capacity, nor on the broadcast capacity.

To properly model the negative effects of document  $R$  on other documents, one needs to model a limited capacity proxy cache. In this more realistic situation, the broadcast of an unneeded document may have adverse effects, since it may occupy space in the cache and preclude more needed documents from the cache. Thus, the modeling of a finite size cache should allow one to account for the negative effects of a broadcast document. These effects can be termed as “cache contamination,” namely a document that is kept in the cache while it is not needed can be considered to contaminate the cache. Wrong broadcast policy (e.g., one that broadcasts “everything”) may contaminate the cache to a high degree, thus preventing it from serving the locally wanted documents.

Using the same definitions of stream  $s$  and proxy cache  $C_p$  as in Section II-B, a broadcast document from stream  $s$  is called a *waste document* in cache  $C_p$ , if it will not have any request directed to it at  $C_p$ . For a given proxy  $C_p \in P'$ , and a document whose number of requests is  $N$ , the document is a waste at  $C_p$  if it has at least one request by  $P'$  and no request by  $C_p$ . The probability of this event is given by

$$\begin{aligned} \Pr[\text{document of stream } s \text{ is waste at } C_p \mid N = n] \\ &= \overline{\lambda_{s,p}}^n \left[ 1 - \left( \frac{\overline{\Lambda_{s,P'}}}{\overline{\lambda_{s,p}}} \right)^n \right] \\ &= \overline{\lambda_{s,p}}^n - \overline{\Lambda_{s,P'}}^n. \end{aligned} \quad (7)$$

Now, using the same definition of Distribution  $c_s(n)$  as in Section II-B, we can calculate the fraction of documents in  $s$ , which end up being a waste at  $C_p$

$$\begin{aligned} \Pr[\text{document of stream } s \text{ is waste at } C_p] \\ &= \sum_{n=1}^{\infty} \left[ c_s(n) \left( \overline{\lambda_{s,p}}^n - \overline{\Lambda_{s,P'}}^n \right) \right]. \end{aligned} \quad (8)$$

In a similar manner to what we did in (1)–(2), let  $W_{s,p}$  be the total *request waste ratio* of stream  $s$  at proxy  $C_p$ , defined to be the ratio between the expected number of waste documents and the expected number of requests for document of stream  $s$ . Using (8), we can derive

$$W_{s,p} = \frac{\sum_{n=1}^{\infty} \left[ c_s(n) \left( \overline{\lambda_{s,p}}^n - \overline{\Lambda_{s,P'}}^n \right) \right]}{E[N]}. \quad (9)$$

Since the different broadcast streams are disjoint, and each broadcast document is different, we can simply total the waste for all incoming streams, after normalizing it by the relative stream rate ( $\gamma_{s,P}/\gamma_{S,P}$ ). Thus, we get

$$W_{S,p} = \sum_{s \in S} \frac{\gamma_{s,P}}{\gamma_{S,P}} W_{s,p}. \quad (10)$$

To end our definitions, we will further define the proxy’s total request ratio  $L_p = (\gamma_{S,p}/\gamma_{S,P})$ .

To model the interaction of the waste streams with the other streams on the cache assume that the streams are having *stack depth distributions*. When a stream is said to have a stack depth distribution it means that when the stream is applied to a LRU-managed cache, the probability that the current reference finds the element that it references at depth  $n$  ( $n = 1$  is the most recently used element in the cache) is  $d(n)$ . Given the above distribution, we can define the corresponding cumulative distribution  $D(k) = \sum_{n=1}^k d(n)$ , which is actually the stream hit ratio for cache size  $k$ , i.e., the probability that the element will be found at depth smaller or equal to  $k$ . Note that  $d(n)$  is a “defective” distribution in that it may not sum to unity: elements that have never previously been accessed will be assumed to be found at an infinite stack depth.

We will assume that each of the streams (the waste stream and the stream consisting of the other requests) obey the *stack reference model* (SRM), that is, they are stochastic processes which choose their next reference according to independent samplings of the stack depth distribution. Such processes are called *stack depth processes* and their hit ratio curves as a function of cache size coincide with their cumulative stack depth distributions. This is one of several simple models for reference streams (used, e.g., in [5]) and it has been reported, for example, in [8], [12], and [13], that this model tends to be quite successful in capturing temporal locality of references within a trace, and is superior to the so called *independent reference model* (IRM).

Given the stack depth distribution of proxy  $C_p$ ,  $d_p(n)$ , and the corresponding cumulative distribution  $D_p(n)$ , we will calculate the effect of the waste stream on the stack depth distribution. To this end, we will use a methodology similar to the one developed in [12] and [13], and we will track an arbitrary tagged locally requested document  $R$  through its journey through the local cache at which both the local documents and the waste

documents accumulate. This will allow us to derive the distribution of total depth (consisting of local and waste pages) at which  $R$  is requested, which will form the depth distribution of the combined cache. To clarify, note that under SRM the identities of individual documents within a stream are not modeled.  $R$  serves as an anonymous document that we track only along its journey over the cache.

Let  $n_L$  denote the number of local documents (nonwaste documents) residing “above” (that is, at a lower depth)  $R$  in the cache, and let  $n_W$  denote the number of waste documents residing above  $R$  in the cache. Then, at any time in the journey, the position of  $R$  in the cache is described by the state  $(n_L, n_W)$ . To track the behavior of this cache, we will focus only on the events where there is access to this cache, either by the waste documents or by the local requests. The probabilities that an event is a waste document arrival or a local access are given by  $x_p^W := W_{S,p}/(W_{S,p} + L_p)$  and  $x_p^L := L_p/(W_{S,p} + L_p)$ , respectively.

Since these probabilities are fixed and do not depend on the other events, the state  $(n_L, n_W)$  is sufficient to predict the future of  $R$ . To track the behavior of  $R$  in the cache, recall that the cache operates under the LRU policy. This means that when any document  $R'$  is requested, the document is placed at the top position (whose depth is 1) of the cache. This causes all the documents which have been prior to the operation above  $R'$  to be pushed one position downward (deeper). In the event that  $R'$  was not in the cache prior to the request, this should cause the deepest document in the cache to be pushed out.

Now, to track the journey of the tagged document  $R$  along the cache, assume that  $R$  is in state  $(n_L, n_W)$  just after the  $i$ th request (“time  $i$ ”) and examine the transition of  $R$  due to the  $i + 1$ st request.  $R$  will encounter one of the following events.

- 1)  $R$  remains at  $(n_L, n_W)$ . This occurs if a request is made for a local document of depth smaller than or equal to  $n_L$ . Thus, the probability of this event is given by  $x_p^L D_p(n_L)$ .
- 2)  $R$  moves to  $(n_L + 1, n_W)$ . This occurs if a request is made for a local document of depth greater than  $n_L + 1$ . Thus, the probability of this event is given by  $x_p^L D_p(n_L + 1)$ .
- 3)  $R$  moves to  $(n_L, n_W + 1)$ . This occurs if a request is made for any waste document. Thus, the probability of this event is given by  $x_p^W$ .
- 4)  $R$  finishes its journey. This occurs if a request is made for a local document of depth  $n_L + 1$ . Thus, the probability of this event is given by  $x_p^L d_p(n_L + 1)$ .

The list of the states to which  $R$  moves given that it leaves  $(n_L, n_W)$ , and the corresponding transition probabilities are the following. (Note that the following probabilities are independent of  $n_W$ .)

- 1)  $R$  moves to  $(n_L + 1, n_W)$ . The probability of this event is  $(x_p^L D_p(n_L + 1))/(x_p^L D_p(n_L) + x_p^W)$ .
- 2)  $R$  moves to  $(n_L, n_W + 1)$ . The probability of this event is  $(x_p^W)/(x_p^L D_p(n_L) + x_p^W)$ .
- 3)  $R$  finishes its journey. The probability of this event is  $(x_p^L d_p(n_L + 1))/(x_p^L D_p(n_L) + x_p^W)$ .

Now, let  $q(n_L, n_W)$  denote the probability that  $R$  will eventually reach the state  $(n_L, n_W)$  in its journey, assuming that the journey starts at the top of the stack, that is defined as state  $(0,$

$0)$ . Thus, we clearly have  $q(0, 0) = 1$ , and  $q(n_L, n_W)$  can be calculated recursively as follows:

$$q(n_L, n_W) = q(n_L - 1, n_W) \frac{x_p^L \overline{D_p(n_L)}}{x_p^L \overline{D_p(n_L - 1)} + x_p^W} + q(n_L, n_W - 1) \frac{x_p^W}{x_p^L \overline{D_p(n_L)} + x_p^W}. \quad (11)$$

Now, let  $e(n_L, n_W)$  denote the probability that  $R$  will finish its journey at  $(n_L, n_W)$ , i.e., that when  $R$  is retrieved it has  $n_L$  local documents and  $n_W$  waste documents above it. This value is given by the probability that  $R$  will reach  $(n_L, n_W)$  and then requested at  $C_p$ , thus finishing its journey

$$e(n_L, n_W) = q(n_L, n_W) \frac{x_p^L d_p(n_L + 1)}{x_p^L \overline{D_p(n_L)} + x_p^W}. \quad (12)$$

Now, we can compute the depth distribution of a local document in proxy  $C_p$  in the presence of all waste documents  $d_{p,L}(n)$

$$d_{p,L}(n) = \sum_{n_L=0}^{n-1} e(n_L, n - 1 - n_L). \quad (13)$$

Since all requests originating at proxy  $C_p$  are for local documents (i.e., there are no requests for waste documents), this depth distribution is the distribution of an arbitrary document in the merged cache in proxy  $C_p$ .

Finally, the hit rate in this cache, when its size is  $K$ , is given by

$$D_{p,L}(K) = \sum_{k=1}^K d_{p,L}(k). \quad (14)$$

### C. Net Hit Ratio Gain: Accounting for Gain and Loss

Following the analysis in Sections III-A and III-B, we can derive the *net* hit ratio gain on proxy  $C_p$ , with cache size  $K_p$ . This is done by subtracting the loss (Section III-B) from the gain (Section III-A), to get the net hit ratio gain  $G_p^{\text{CSDS\_NET}(P')}$ , using (6) and (14)

$$G_p^{\text{CSDS\_NET}(P')} = G_p^{\text{CSDS}(P')} - (D_p(K_p) - D_{p,L}(K_p)). \quad (15)$$

Finally, let us define  $G^{\text{CSDS}(P')}$ , as the total hit ratio gain over all proxies  $C_p \in P'$ . Just like in deriving (6), we need to take caution in selecting the normalizing factor, which is now  $(\gamma_{S,p}/\gamma_{S,P'})$ , i.e., the probability that an arbitrary request originated at proxy  $C_p$ , given it originated at one of the proxies participating in CSDS. Thus, we get

$$G^{\text{CSDS}(P')} = \sum_{p \in P'} \frac{\gamma_{S,p}}{\gamma_{S,P'}} G_p^{\text{CSDS\_NET}(P')}. \quad (16)$$

### D. Computational Complexity

Equations (11)–(14) can be computed in a recursive manner (starting from the low indices and going upwards). Thus, the computation complexity per proxy is  $O(J^2)$ , where  $J$  is the

number of entries in the stack depth distribution. In a straightforward approach, one would take  $J$  to be the cache size ( $K$ ), in which case the complexity is  $O(K^2)$ . However, if one is interested in reducing the complexity, one can represent the distribution by  $J = \alpha K$  values, where  $\alpha < 1$ . In this case, the complexity is  $O(\alpha^2 K^2)$ . Proper selection of  $\alpha$  (small value) can lead to a drastic reduction in the complexity without significantly affecting the accuracy.

#### IV. OPERATIONAL RULES AND EFFICIENT DESIGN

The model and analysis provided in this paper form a tool that can be used in the design and operations of a CSDS. An important question that needs to be addressed by the operator of the CSDS is the stream assignment problem, i.e., which classes (or streams) of documents should be broadcast in the system. Intuitively, it is expected that the operator should decide to broadcast classes of documents of “common interest” to the various proxies. The results provided in this analysis, namely in (15), represent the net hit gain on  $C_p$  due to broadcasting a set  $S'$  of streams (as introduced in Remark 5). Thus, to answer the question of which streams to broadcast the operator can compute (15) for various sets of streams, and derive the relative benefits to the systems as the result of broadcasting alternative sets of streams. For example, suppose that set  $S'$  is broadcast, and the operator would like to extend the set to  $S' \cup s$ , where  $s$  could be any of  $J$  alternative streams  $s_1, \dots, s_J$ . The operator can then evaluate (15) for  $S' \cup s_i, i = 1, \dots, J$ , and use it to rank the streams  $s_i, i = 1, \dots, J$ , and to decide which of them will be broadcast. Having selected the stream  $s_i$ , a new set is now formed,  $S' \cup s_i$ . The process may now repeat on how to extend  $S' \cup s_i$ .

Another important design question is the proxy assignment problem, i.e., which proxies should participate in a system. A simpler question can be whether to add a proxy  $C_p$  to an existing system where the participating proxies are the set  $P'$ . The value of adding  $C_p$  to the system can again be evaluated by applying (15) to the set  $P'$  and to the set  $P' \cup C_p$ .

These questions and other relevant questions, and a more detailed analysis of them, are a topic of current ongoing research. A greedy algorithm for the solution of the proxy assignment problem is presented and tested in Section VI below.

#### V. NUMERICAL RESULTS

In this section, we aim to examine the quality of the model and analysis. We will first describe the system we generated and then test our analytic models against simulations.

We consider a system composed of four proxy caches, receiving requests on five Zipf distributed SRM streams. Any other distribution might be used,<sup>3</sup> but the Zipf distribution might emulate real life the best. The relative request rate for each proxy and stream, used to generate the  $\gamma_{s,p}$  matrix, is given in Table I. Stream 1 is requested mainly (90%) by proxy 1 and the rest of the requests are divided uniformly across the rest of the proxies. Stream 2 is requested mainly (80%) by proxy 2 and the rest of

TABLE I  
MATRIX OF RELATIVE PROXY AND STREAM RATIOS

	Proxy1	Proxy2	Proxy3	Proxy4	Total/St.
Stream1	900	33	33	34	1000
Stream2	66	800	67	67	1000
Stream3	250	250	250	250	1000
Stream4	50	50	50	50	200
Stream5	333	333	333	1	1000
Total/Pr	1599	1466	733	402	4200

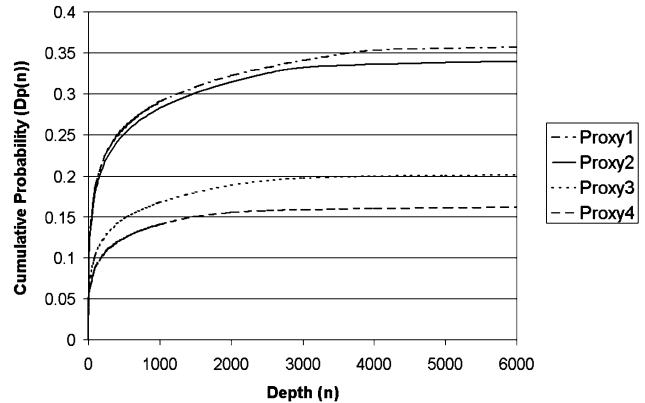


Fig. 2. Calculated local stack depth distribution.

the requests are divided uniformly across the rest of the proxies. Streams 3 and 4 are divided uniformly across all proxies. Stream 5 is divided uniformly across all proxies except proxy 4. All streams have the same total request share, except for Stream 4, whose request volume is five times lower than the volume of each of the other streams.

The results are generated as follows. We start by generating a simulative request log, composed of 5 000 000 (five million) requests, based on the relative request rate matrix and the global streams' SRM stack depth distributions, classifying each request to arrive from a stream and a proxy. Then, we can simulate any broadcast scenario. In addition, we generate from the log the request count distribution  $c_s(n)$  for every global stream, and the local stack depth distribution  $d_p(n)$  for every proxy. The generated cumulative local stack depth distribution (across all streams) for each proxy, which is equal to the local hit ratio of every proxy without using CSDS, is given in Fig. 2. Now, for every broadcast scenario we need to run the analysis of Section III-C above and the simulation, and compare the total net gain for the given scenario.

First, we examine the effect of broadcasting each of the various streams on every proxy. We consider a cache size of 4000 documents,<sup>4</sup> which is on the order of the saturation point of our distribution (the point at which increasing the cache size does not increase the local hit rate), and run separate CSDS, each broadcasting one stream. The results are reported in Fig. 3. First, we notice the good correlation between the analytic results and the simulation results, for the entire test scenario. We observe that the most beneficial stream to broadcast is, as expected, the high-volume, evenly distributed stream 3 [Fig. 3(c)]. We see a

<sup>3</sup>We tested the system on various distributions, including handpicked distributions and geometric distributions, and received similar accuracy.

<sup>4</sup>We repeated the numerical tests for cache sizes of 400 and 40 000 (with proper caches and distributions) as well, and received similar accuracy.

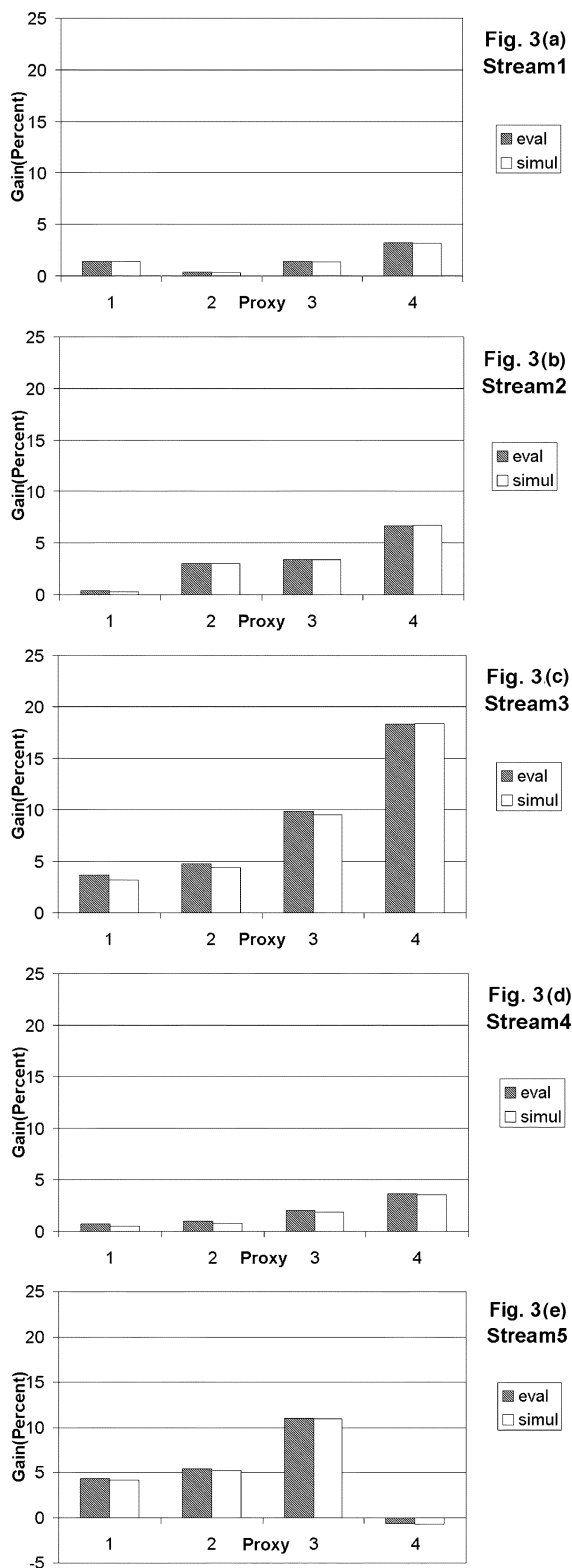


Fig. 3. Simulative and analytic gain for a single stream broadcast for cache size = 4000 documents.

difference of its contribution for different proxies, since its relative presence in proxy 3, is the highest. Observe [Fig. 3(e)] that stream 5 actually reduces the hit rate of proxy 4, because broadcasting it only contaminates proxy 4's cache, but does not contribute any hits. For the other proxies it contributes quite a

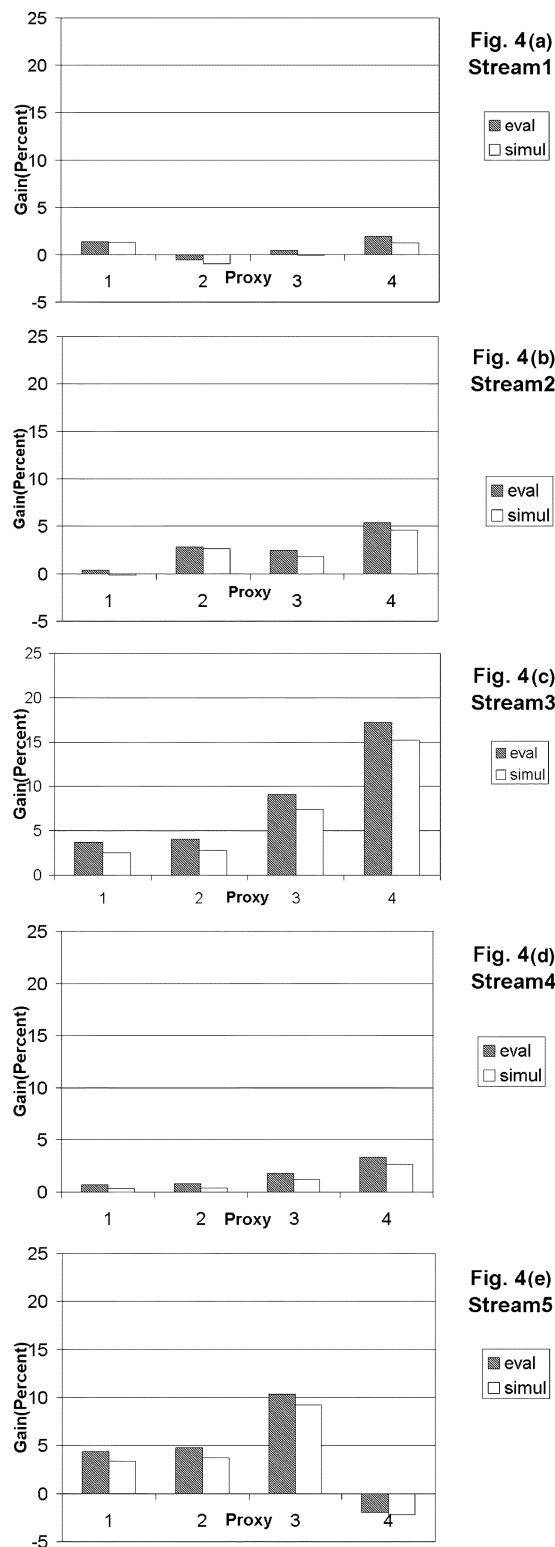


Fig. 4. Simulative and analytic gain for a single stream broadcast for cache size = 1000 documents.

lot, because, again, it is distributed evenly across those proxies. Stream 1 is not distributed uniformly [Fig. 3(a)]: Proxy 1 receives most requests and, thus, it will broadcast a high amount of documents to the other proxies. Now, if the receiving proxy has a low amount of other requests (like proxy 4), receiving stream 1 will not degrade the local hit rate by a substantial amount. For



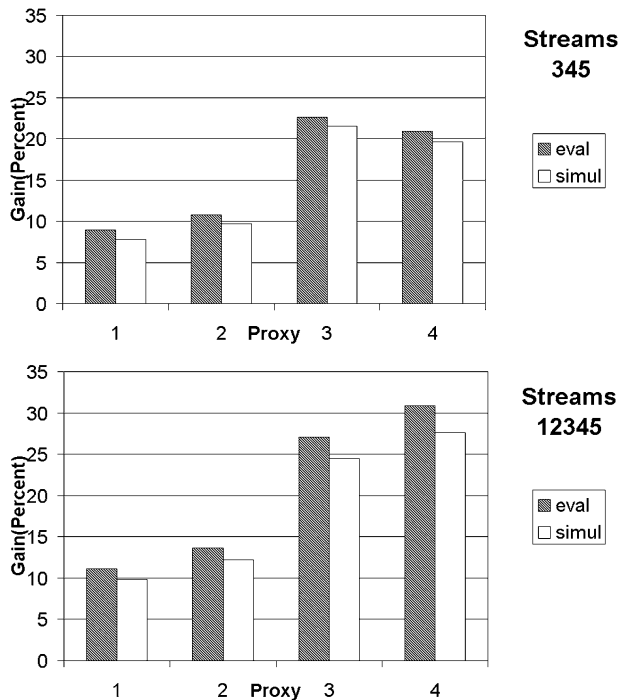


Fig. 5. Simulative and analytic gain for multiple streams broadcast, for cache size = 4000 documents.

“loaded” proxies, like proxy 2, the benefit we observe is small, because the arriving documents throw out other important documents. Stream 4 [Fig. 3(d)] is evenly distributed, but due to its low volume, it does not contribute much gain.

We then repeat the experiment for a much smaller sized cache of 1000 documents. The results are reported in Fig. 4. We notice our analytic results are still very well correlated to the experimental results, although the differences are more noticeable. Notice also that now the gains are smaller (since the contamination has a higher impact), especially on proxies with low correlation with the broadcast stream (e.g., stream 2 on proxy 1, or stream 1 on proxy 2).

For a comparison, we give in Fig. 5 a partial broadcast set scenario of streams 3–5 for a 4000 documents sized cache, and a full broadcast set scenario of streams 1–5 for the same cache size. We notice, again, the good correlation between the analytic results and the simulation results. We can also see the high gain of proxy 3 and proxy 4, which are relatively small proxies.

We can also check our main source of inaccuracy—misses that are not “first access” misses. Such misses appear when the cache sizes are small in a considerable measure than their stack depth distribution, so they can not hold even sufficient local requests. As already mentioned, this is not a common case in real life, and in our model we assumed it could be neglected. In Fig. 6, the effect of decreasing cache sizes on a CSDS that broadcasts only the uniformly distributed stream 3 can be seen. Since we mostly care about the relative error in the gain calculation, we plot the difference of the error in calculating each proxy’s gain, from the weighted average error of the entire system. In Fig. 7, we repeat the experiment for a CSDS that broadcasts only stream 5, which is not distributed uniformly. First, notice in both cases the error diverges for very small caches. Of course,

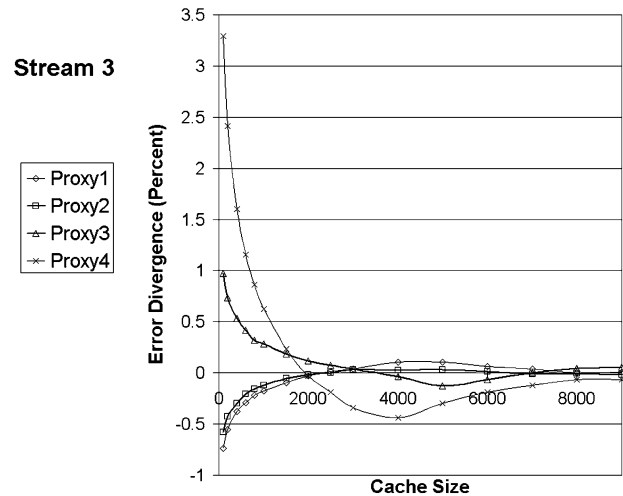


Fig. 6. Difference of stream 3 per-proxy error ratio from the total error expected value.

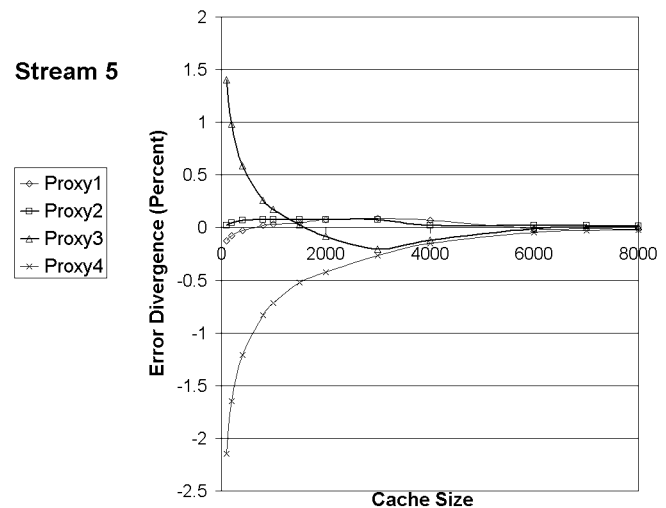


Fig. 7. Difference of stream 5 per-proxy error ratio from the total error expected value.

as the cache is smaller it diverges faster (so we notice cache 4 diverges the fastest). For all caches larger than 700 documents, we get in both cases a reasonable error divergence (less than 1 percent apart). Caches smaller than 700 documents are not likely to be practical (see their very low hit ratio in Fig. 2).

Thus, we observe that for all reasonable cache sizes the estimation error of our analysis is limited and the model will yield good results, especially for optimization and operational rules.

## VI. EXPERIMENTAL RESULTS: DESIGN OPTIMIZATION ON REAL DATA

In this section, we demonstrate the applicability of the model and analysis by addressing one of the operational and design problems presented in Section IV, the proxy assignment problem. We aim at selecting a set  $P'$  of proxies for a CSDS, such that the selected set will yield the maximal hit rate improvement. We will examine a simple greedy algorithm. First, take an arbitrary proxy and try to add to it one of a few other proxies. The proxy that will yield the best improvement

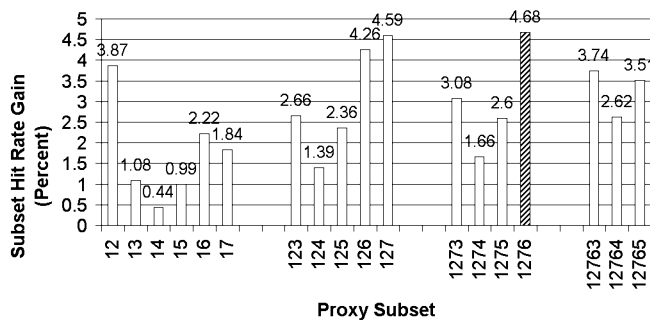


Fig. 8. Proxy allocation via the analytic model and the corresponding hit rate gains (real data).

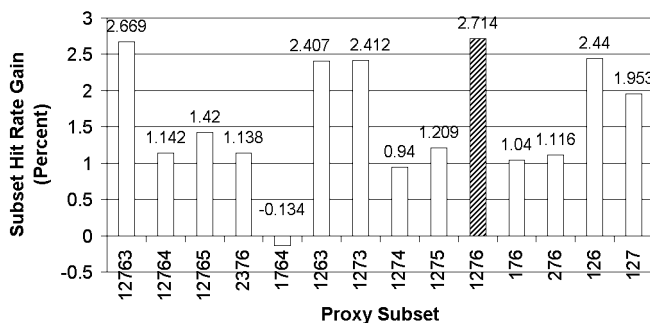


Fig. 9. Emulation results: Normalized hit rate gain of the selected set of Fig. 8 and of its neighboring sets.

will be added to the set, and then the process of adding one proxy at a time will repeat. The process will end when adding a proxy to the set will result in a lower improvement than without the proxy. At each stage of the algorithm, the value of the set examined is evaluated via the analytic model and is used as the criterion that guides the algorithm. Once the winning set is selected, we evaluate the quality of our analysis by comparing it to emulation results. We emulate the performance of the various proxy sets on real data, taken from HTTP logs.

To conduct this analysis we use logs taken from ircache.net [11], which included over 2 700 000 requests. We use a simplistic stream selection heuristics, where we associate a stream with each Website. Using the same method as for the synthetic log in Section V, we calculate the local SRM for each of the proxies and the request count distribution for each of the streams, and used it in our analytic model to predict the hit rate to be experienced by the various proxy sets.

Fig. 8 depicts the proxy sets examined by the algorithm, ordered from left to right, and their corresponding hit rate gain. The greedy algorithm selects the broadcast group {1, 2, 7, 6}. The analytic model (driving the algorithm) indicates that either adding or removing proxies from this set only degrade the gain.

Fig. 9 examines the quality of the algorithm by depicting the gain of the broadcast set {1, 2, 7, 6} and other alternative sets as derived via emulation on the actual logs. The gain is normalized to the volume of requests of the chosen broadcast set to make the comparison meaningful. The figure demonstrates that any modification to the selected set by either a single proxy addition, a single proxy removal, or a single proxy replacement end-up with lower gains. Thus, the set selected by the analytic model is indeed superior to other alternative sets.

The design of a generic efficient algorithm for proxy and stream selection into a CSDS is a topic of an ongoing research.

## VII. SUMMARY

In this paper, we dealt with the CSDS and aimed at providing a framework for the analysis and efficient operation of this system. We proposed a model that accounts for the intrinsic behavior of caches and request streams and which captures both the locality of reference experienced in these streams and their relative frequencies. Using the model, we provided an analysis that predicts the hit ratio gain for each stream and each proxy as function of the stream properties. We presented how these predictions, whose computational complexity is relatively low, can be used directly in solving the operational and design questions. Numerical examination versus simulation indicates that the analysis indeed captures the relative effect of various streams on various proxies and, thus, will predict well the relative merits of alternative operational rules.

## ACKNOWLEDGMENT

Real data logs were taken from <http://www.ircache.net> [11] (in reference to the National Science Foundation [Grants NCR-9616602 and NCR-9521745], and the National Laboratory for Applied Network Research).

## REFERENCES

- [1] S. Acharya and S. Muthukrishnan, "Scheduling on-demand broadcasts: New metrics and algorithms," in *Proc. 4th Annual ACM/IEEE Int. Conf. Mobile Computing and Networking (MobiCom'98)*, Oct. 1998, pp. 43–54.
- [2] V. F. Almeida, M. G. Cesario, R. C. Fonseca, W. Meira Jr., and C. D. Murta, "Analyzing the behavior of a proxy server in light of regional and cultural issues," presented at the 3rd Int. WWW Caching Workshop, June 1998, [Online]. Available: <http://hermes.wwwcache.ja.net/events/workshop/21>.
- [3] A. Armon and H. Levy, "Cache satellite distribution systems: Modeling and analysis," in *Proc. IEEE INFOCOM'03*, vol. 1, Apr. 2003, [Online]. Available: [http://www.ieee-infocom.org/2003/papers/06\\_04.PDF](http://www.ieee-infocom.org/2003/papers/06_04.PDF), pp. 240–250.
- [4] D. Askoy, M. Franklin, and S. Zodnik, "Data staging for on-demand broadcast," in *Proc. 27th VLDB Conf.*, Sept. 2001, pp. 571–580.
- [5] M. Brehob and R. J. Enbody, "An Analytic Model of Locality and Caching," Dept. Comput. Sci., Michigan State Univ., East Lansing, [Online]. Available: <http://www.cse.msu.edu/publications/tech/TR/MSU-CSE-99-31.pdf>, Tech. Rep. MSU-CSE-99-31, Aug. 1999.
- [6] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenkar, "On the implications of Zipf's law for Web caching," in *Proc. IEEE INFOCOM'99*, June 1999, pp. 126–134.
- [7] S.-G. Chang, "Caching strategy and service policy optimization in a cache-satellite distribution service," presented at the 5th INFORMS Telecommunications Conf., Mar. 2000, [Online]. Available: [http://www.digital.re.kr/changsg/\\_private/changsg06.pdf](http://www.digital.re.kr/changsg/_private/changsg06.pdf).
- [8] E. G. Coffman and P. J. Denning, *Operating System Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [9] R. Cohen, L. Kazir, and D. Raz, "Scheduling algorithms for a cache pre-filling content distribution network," presented at the IEEE INFOCOM'02, June 2002, [Online]. Available: <http://www.ieee-infocom.org/2002/papers/442.pdf>.
- [10] X.-Y. Hu, P. Rodriguez, and E. W. Biersack, "Performance study of satellite-linked Web caches and filtering policies," in *Networking 2000*, May 2000, pp. 580–595.
- [11] NLANR, National Science Foundation, Bethesda, MD, "Sanitized Cache Access Logs," [Online]. Available: <http://www.ircache.net>, Sept. 15, 2002.
- [12] H. Levy and R. J. T. Morris, "Exact analysis of Bernoulli superposition of streams into a least recently used cache," *IEEE Trans. Software Eng.*, vol. 21, pp. 682–688, Aug. 1995.

- [13] —, “Should caches be split or combined? Analysis using the superposition of bursty stack depth processes,” *Perform. Eval.*, vol. 27 & 28, pp. 175–188, 1996.
- [14] P. Rodriguez and E. W. Biersack, “Bringing the Web to the network edge: Large caches and satellite distribution,” *Mobile Networks Applicat.*, vol. 7, no. 1, pp. 67–78, Jan. 2002.
- [15] M. A. Sharaf and P. K. Chrysanthis, “On-demand broadcast: New challenges and algorithms,” presented at the 1st Hellenic Data Management Symp, July 2002, [Online]. Available: [http://www.dbnet.ece.ntua.gr/HDMS02-proceedings/HDMS02\\_files/chrysanthis.pdf](http://www.dbnet.ece.ntua.gr/HDMS02-proceedings/HDMS02_files/chrysanthis.pdf).



**Aner Armon** received the B.Sc. degree in mathematics, physics, and computer science from the Hebrew University, Jerusalem, Israel, in 1993, and is currently working toward the M.Sc. degree in computer science at Tel-Aviv University, Tel-Aviv, Israel.

Until 2001, he worked as a Vice President of Derivatives Technology, Laboratory for Computerized Trading (LCT), Knight Financial Products (formerly Arbitrade), a subsidiary of Knight Trading Group, New York. Since 2001, he has been with Cloverleaf Communications, Petach Tiqwa, Israel,

a storage networking start-up.



**Hanoch Levy** (S'83–M'87) received the B.A. degree in computer science with distinctions from The Technion—Israel Institute of Technology, Haifa, in 1980 and the M.Sc. and the Ph.D. degrees in computer science from University of California at Los Angeles, in 1982 and 1984, respectively.

From 1984 to 1987, he was a Member of Technical Staff in the Department of Teletraffic Theory, AT&T Bell Laboratories, Holmdel, NJ. Since 1987, he has been with the School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel. From 1992 to

1996, he has been at the School of Business and RUTCOR, Rutgers University, New Brunswick, NJ, on leave from Tel-Aviv University. His interests include computer communications networks, wireless networks, quality-of-service in high-speed networks, performance evaluation of computer systems, and queueing theory.