

Towards Content Distribution Networks with Latency Guarantees

Chengdu Huang and Tarek Abdelzaher
Department of Computer Science
University of Virginia
Charlottesville, VA 22904
e-mail: {ch4pp, zaher}@cs.virginia.edu

Abstract

This paper investigates the performance of a content distribution network designed to provide bounded content access latency. Content can be divided into multiple classes with different configurable per-class delay bounds. The network uses a simple distributed algorithm to dynamically select a subset of its proxy servers for different classes such that a global per-class delay bound is achieved on content access. The content distribution algorithm is implemented and tested on PlanetLab [1], a world-wide distributed Internet testbed. Evaluation results demonstrate that despite Internet delay variability, subsecond delay bounds (of 200-500ms) can be guaranteed with a very high probability at only a moderate content replication cost. The distribution algorithm achieves a 4 to 5 fold reduction in the number of response-time violations compared to prior content distribution approaches that attempt to minimize average latency. To the authors' knowledge, this paper presents the first wide-area performance evaluation of an algorithm designed to bound maximum content access latency, as opposed to optimizing an average performance metric.

I. Introduction

This paper presents the first wide-area evaluation of a content distribution network (CDN) designed to provide delay bounds on content access. The predominant use of the Internet for Web content delivery has recently spurred much research on CDN performance. Existing research on CDNs includes techniques for efficient client request redirection [2], [3], [4], [5], server placement strategies to improve average response time or bandwidth consumption [6], [7], [8], [9], [10], logical overlay topologies for large-scale content distribution [11], [12], [13], consistency maintenance mechanisms [14], [15], and empirical CDN performance measurement studies [16], [17]. Little attention has been given to QoS guarantee issues in CDNs. In this paper, we address the specific problem of providing

subsecond guarantees on access delay. By “guarantee”, we do not mean deterministic delay bound guarantees. Instead, we use the word “guarantee” in a loose sense, to mean trying to achieve latency bounds with a “high probability”. Our work is motivated by empirical measurements, which show that Internet delays can exceed the human perception threshold by as much as an order of magnitude with a non-trivial probability [18]. At the same time, underutilization of ISP backbones creates delay patterns that are reasonably stable, suggesting that it is possible to choose stable locations for content replicas using a delay-sensitive algorithm that bounds access delay. This conjecture is evaluated in this paper.

Our evaluation results are drawn from an actual implementation and deployment of a service prototype in a world-wide WAN testbed. Operationally, our network is composed of a large set of content distribution proxies. Content providers negotiate with our distribution network the intended service access latency for their content. The content distribution network dynamically determines the number and locations of content distribution proxies that will need to host the content in order to globally guarantee the negotiated bound on access delay. Content access latency, defined as the time elapsed from the arrival of a client's request to some CDN server until the requested content has been sent to the client, includes the delay resulting from forwarding the request within the CDN if the initially accessed server does not have a local copy of the content and processing delays in CDN servers. In this work, we focus on the former part of the total latency.

The remainder of this paper is organized as follows. Section II elaborates on the service model. The distributed algorithm for proxy selection is described in Section III. Implementation is described in Section IV. An extensive performance evaluation is presented in Section V. Section VI presents a survey of previous work. The paper concludes with Section VII.

II. System Model

In our model, when a content provider authorizes a CDN provider to host and replicate its content objects, it also specifies the latency bound on access to its content in a QoS contract. Different content providers can have different latency bound requirements with the CDN provider. Content objects of a same provider but of different sizes can have different latencies. There are many ways to specify latency bound values for objects of different sizes. One way is to provide a uniform latency bound for objects of sizes up to a certain upper limit. Another way is to make the latency bound a function of object size, making larger objects have larger latency bounds. In our current implementation, the latency bound is a step function of object size: $L(s) = L_{base} + L_{incr} \times \lfloor s/S_{step} \rfloor$ where s and $L(s)$ are the size and latency bound of the object and L_{base} , L_{incr} and S_{step} are negotiable parameters. Each class has a QoS contract in this form. The discretization of achievable bounds allows the system to classify content into large aggregate classes by latency bounds, as opposed to having to maintain a different bound for each object individually depending on its size.

We assume client requests can be redirected to the nearest CDN server using various request routing mechanisms [2], [3]. When a CDN server receives a request, if the requested content object is available at the server's local storage, reply is sent back to the client directly. Otherwise, the CDN server acts as a proxy and forwards the request to some replica (or to the origin server). When a reply is received from the selected replica, the reply is sent back to the ISP proxy from which the request originated.

The latency perceived by a client when requesting a content object hosted by our CDN service actually consists of three parts, latency between the client and the ISP entry point/CDN server, which is known as the last-mile problem, latency caused by requesting the object from some peer CDN server or the origin server across the backbone when the requested object is not available at the CDN server that received the client request, and processing delay in the CDN server. The second part above is the focus of our service because the last-mile problem is usually out of a CDN provider's control. Also, as broadband services of ISPs become more and more popular, we envision that the last-mile problem will become less severe in the future. For the third part, we acknowledge that it is a very important factor and can even dominate the total latency in some cases. However, currently, we address only the case where servers are not overloaded. This assumption will be relaxed in future work.

III. Algorithm

In this section, we formulate our QoS guarantee problem as a graph domination problem. A distributed algorithm is

then presented to solve this problem. Our replica placement has to ensure that requests received by the CDN servers are served within the latency bound specified in the QoS contract. In the extreme case, if the content is replicated to *all* CDN servers, network latencies of requests for the content will be minimized. However, the cost of replicating the content everywhere is large. First, content update traffic is not free. We envision future content that is volatile and requires frequent updates (e.g., live coverage of sports events with photographic snapshots, commentary, and video clips). The fewer the number of replicas the lower the update cost. Second, having too many replicas for each content object can significantly reduce overall throughput of a CDN system. The reason lies in the fact that partitioning client requests to too many different replicas lessens temporal locality of content access requests. As a result, many client requests have to be served from servers' disks instead of memory. This phenomenon has been confirmed in previous research [4]. Finally, depending on content size, storage cost may or may not be an issue.

Hence, we state the content distribution problem as that of finding a replica placement that uses the minimum number of replicas and ensures that requests received by any CDN server of the content provider can be served within the latency bound.

Our CDN servers form an overlay network. Although every CDN server can reach every other server via the underlying IP routing, only links whose delay is less than the bound are valid in the overlay topology. Observe that, under these conditions, different content with different QoS requirements could have a different view of the overlay network. Due to the asymmetric nature of Internet routing, all the edges are directional, meaning an edge from server S_i to S_j does not imply the existence of an edge from server S_j to S_i . With this in mind, we can formally state our QoS guarantee problem as follows. Given a set of CDN servers $\mathcal{S} = \{S_i | i \in [1, N]\}$, a content object C , and its latency bound L , we construct an overlay network whose nodes are the CDN servers \mathcal{S} . Edges are added following the rule that $\overrightarrow{S_i S_j}$ exists if and only if S_i can download C from S_j within a time less than L . Trivially, $\overrightarrow{S_i S_i}$ exists for all $i \in [1, N]$. Our objective is to find a minimum size subset $\mathcal{D} \subset \mathcal{S}$ to hold replicas of C , such that for every node S in $\mathcal{S} - \mathcal{D}$ there exists at least one edge $\overrightarrow{S D}$ in the overlay network such that $D \in \mathcal{D}$. Observe that because of the way edges are defined, the aforementioned condition implies that the delay between any node S and at least one replica D is less than the bound.

The algorithm to solve the above replica placement problem classifies files into discretized categories by latency bound. The algorithm needs to run for each content class (with a common latency bound) instead of for each

content object. To further reduce the scale of the problem, observe that network edges which satisfy a shorter bound also satisfy a longer one. Hence, approximate solutions for a longer latency bound can be derived incrementally from solutions for shorter latency bounds. In this paper, we separate probing network delay from running the placement algorithm. Network probing is done in a common module. The placement algorithm runs concurrently (separately) for each content class.

A. Centralized Algorithm

Consider the graph domination problem formulated above for a single content class. A set of vertices of a graph is called dominating if every other vertex is adjacent to at least one vertex of the set. The graph domination problem asks for a dominating set of the minimum possible size in a given graph. Graph domination is known to be NP-Hard [19]. A common approximation algorithm for the graph domination problem is the greedy algorithm [20]. Vertex S_i is said to be able to “cover” vertex S_j if and only if $\vec{S}_j S_i$ is in the graph. In accordance with previous papers on dominating set algorithms (e.g. [21]), we call the number of uncovered vertices that a vertex can cover (including itself) the *span* of the vertex. The greedy algorithm repeatedly selects the vertex that has the highest remaining span to join the dominating set until the graph is dominated. The intuition behind the greedy algorithm is that the vertices with high span values are more likely to be selected in a dominating set.

It is known that the above greedy algorithm achieves an H_Δ -approximation where Δ is the maximum degree of the nodes and H_i is the i th harmonic number (i.e. $H_i = \sum_{k=1}^i \frac{1}{k}$) [20], [22]. This is actually the best approximation known so far. The greedy algorithm is very simple but has its drawbacks in practice. Most importantly, it requires a centralized coordinator to perform the replica assignment. It assumes that network topology is known to the centralized coordinator, which in this case means that the coordinator knows the latency from *any* node to *any* other node in the network. Such centralized algorithms generally have scalability and robustness limitations and lack the flexibility of doing online adjustments to accommodate changes in network conditions. We use the aforementioned centralized greedy algorithm only as a baseline to compare against. Our goal is to find a distributed algorithm that would perform close to the centralized one.

B. Distributed Algorithm

A few distributed algorithms for constructing dominating sets have been proposed. Liang and Haas used a distributed algorithm called DDCH [23] to generate virtual backbones for Ad Hoc networks. When synchronously executed, DDCH achieves the same approximation ratio as

the centralized greedy algorithm (Section III-A) but there exist networks for which DDCH takes a very long time ($\Omega(n)$ rounds, n is the number of nodes) to complete. LRG introduced by Jia et al.[21] is a refinement of DDCH, in which randomization is used to break symmetries when multiple nodes attempt to add to the dominating set. It’s proven that LRG terminates in $O(\log n \log \Delta)$ rounds with high probability, where Δ is the maximum span of the graph and yields an approximation ratio with expectation of $O(\log \Delta)$. Kuhn and Wattenhofer designed a more sophisticated distributed algorithm [24] based on LP relaxation techniques. Given an arbitrary constant k , the algorithm computes a dominating set of expected size $O(k\Delta^{2/k} \log \Delta)$ times optimal in $O(k^2)$ rounds. Though theoretically appealing, the performance of these algorithms when used in asynchronous systems was not studied. Moreover, these algorithms share a common disadvantage: they need multiple rounds to finish. Kuhn’s algorithm [24] can be made to finish in one round by setting $k = 1$ but then the performance will be far from optimal. In practice, more rounds translates into more overhead and a longer termination time, which may be problematic if network conditions change before the algorithm terminates.

In this paper, we seek an algorithm that is both very fast (ideally, constant-time) and yields a small dominating set in a highly asynchronous execution environment. We designed a simple distributed algorithm to meet these requirements. The algorithm is fully distributed and asynchronous. Figure 1 shows the pseudocode that is executed at each node independently. We require that every node know only its own span value and the number of nodes it can reach within the delay bound including itself. Each node sends its span (SPAN), i.e. the number of nodes that can download the object from it within the latency bound, to all the nodes it can cover and collects other nodes’ spans. The message specifies the span for each content class. If a node can only be covered by itself, the node has to be in the dominating set and sends a DOMINATOR message to nodes it can cover claiming that it is in the dominating set. Otherwise, after it gets SPAN messages from all the nodes that can cover it, or timeout, it chooses the node with highest span among them (which could be itself) and makes that node a member of the dominating set by sending it a NOMINATION message. When a node receives a NOMINATION message, it joins the dominating set. It then sends a DOMINATOR message to nodes it can cover and will not nominate another node if it has not sent out any NOMINATION message yet. Assuming no message loss, since every node selects a node that covers itself, when all nodes exit the algorithm the graph is dominated.

The mechanism that each node independently chooses

a nominee makes the algorithm be able to terminate in one round. This mechanism only uses information about direct neighbors but is very effective in choosing the right nodes to be added to the dominating set. The rule that receiving a DOMINATOR message from some adjacent node makes a node refrain from nominating another is very helpful for reducing the size of the final dominating set because senders of the DOMINATOR messages are already in the dominating set and nomination can only possibly cause a new node to be added to the dominating set. Note that the NOMINATION and DOMINATOR messages are class-specific. However, if the same server is nominated for multiple classes, the corresponding messages are coalesced.

```

// For node  $S_i$ 
send SPAN (span of  $S_i$ ) to all nodes;
if (I am the only node can cover me)
  join the dominating set;
  send DOMINATOR to nodes I can cover;
  return;
while (true)
  if (recv DOMINATOR from  $S_j$  &&  $S_j$  covers me)
    send COVERED to nodes I can cover;
    return; //  $S_i$  is covered by  $S_j$ 
  if (receive SPAN from  $S_j$ )
    updates span of  $S_j$ ;
  if (receive COVERED from  $S_j$ )
    decrease span  $S_j$  by 1;
  if (receive NOMINATION from  $S_j$ )
    join the dominating set;
    send DOMINATOR to nodes I can cover;
    return;
  if (got SPAN of all nodes that cover me)
    select the highest span node  $S_k$  (with
      random tie-breaking) from nodes that
      can cover me;
    send NOMINATION to  $S_k$ ; // covered by  $S_k$ 

```

Fig. 1. Distributed Graph Domination Algorithm

In Section V-A we give a detailed performance evaluation and comparison of the centralized greedy algorithm, previous distributed algorithms mentioned above, and our distributed algorithm.

IV. Implementation

We implemented our content distribution service by instrumenting a Squid Proxy Cache [25] and deployed it on PlanetLab. Although the distributed replica selection algorithm is simple in appearance, there are many technical issues that need great care when implemented on a real system and deployed in a realistic WAN environment. These issues are described below.

A. Probing

In Section III-B, we assumed that nodes know their span values. This knowledge does not come for free in the real

world. When our system is to host a content class, we may need to construct the overlay network for the class before invoking the distributed replica selection algorithm. This is done based on probing results. Probing is an independent operation that is performed for different files sizes and the results are interpolated to provide delay information for an arbitrary file size.

A CDN server probes each peer a few times by sending requests for content of different sizes. The measured delays are fitted to a linear equation describing delay as a function of content size. Curve fitting is done using a least squares estimator. The parameters (slope and intersect) of the delay equation for a given peer are then sent to that peer in a REACHABILITY message.

Since our probing actually downloads files from the remote servers, it has a higher overhead compared to other approaches such as sending ping packets. However, the overall probing cost remains acceptable because (1) probing happens between servers of the CDN provider at a very low frequency, (2) probing does not have to be active all the time. When the system is running, client requests forwarded by one CDN server to another can serve as measurements, and (3) traffic volume of probing messages is typically orders of magnitude smaller than traffic generated by clients' requests. Probing is a relatively independent module of our system and is not the focus of this work. Exploring more sophisticated probing techniques could help construct more accurate overlay topologies but is beyond the scope of this paper.

B. Asynchronous Execution

Imposing a synchronous execution mode on a large scale CDN system is not acceptable. The performance of a distributed algorithm in an asynchronous mode is the meaningful metric for our system. The distributed algorithm implemented in our system is able to run asynchronously. In the asynchronous mode nodes constantly collect control packets (SPAN, DOMINATOR, COVERED, and NOMINATION) from other nodes. After a certain predefined period, each node makes nomination decisions, if it has not sent any NOMINATION messages yet, based on the information it has at the instant. This guarantees that a dominating set will be generated but the cardinality of the dominating set could be larger than that generated in a synchronous execution. In [23] the authors discussed the asynchronous operation of DDCH, which uses a periodic mode too, except that their algorithm usually takes multiple rounds to finish. The mode can also be applied to LRG [21] since it is a refinement of DDCH. We followed that mode in our implementations of DDCH and LRG.

C. Replica Selection

A server can receive multiple DOMINATOR messages from different reachable servers for the same content class. Since all these servers are replicas and can be reached within the latency bound, any of them can be the server’s replica. Observe that when latencies between two servers are low on average, the variances of the latencies tend to be low too. Based on this rationale, we make the server choose the replica that has the lowest latency as its primary replica. Other reachable replicas are marked as backup replicas. Note that the origin server always has valid copies of content. It is automatically a backup replica for all the CDN servers that can reach it within the latency bound.

D. Periodical Invocation

The assignment of replicas is not fixed for the lifetime of the system. The distributed replica selection algorithm is actually invoked periodically using the most recent probing data. This enables the system to self-adjust to accommodate network traffic and server workload condition changes. Probing enables the servers to have a fairly up-to-date view of latencies to other servers all the time. When a server is about to run the replica selection algorithm, it first sends out REACHABILITY messages to other servers according to its current latency statistics. After these messages are exchanged among the servers, they can broadcast SPAN messages to report their span values and run the distributed algorithm.

Adjustments to replica selection can be done in between periodic algorithm invocations as well. We further discuss and evaluate this mechanism in Section V-D.

V. Evaluation

In this section, we evaluate the performance of our system and the effects of various design considerations on the efficiency of our distributed replica selection algorithm. We first evaluate the efficiency of our algorithm in terms of generating minimal number of replicas. Then we study the system’s performance of achieving delay bound, comparing to several baseline algorithms. We also explore techniques to handle different file sizes and achieve load-balancing and scalability. Finally, we explore the performance of the algorithm in the absence of accurate system knowledge. The evaluation demonstrates that our system is able to achieve the desired delay guarantees in the absence of global knowledge and at a very moderate cost.

A. Evaluation of the Number of Replicas

As discussed in Section III, the number of replicas (NOR) needed for a certain content object directly relates to the cost of hosting the object and is one of the primary performance metrics of our system. Our goal is to minimize the

NOR. We have formulated the replica selection for maintaining a latency bound as a graph domination problem and presented a decentralized approximation algorithm in Section III.

To understand the performance in terms of minimizing NOR, we compare our distributed algorithm with the centralized greedy heuristic described in Section III-A and two distributed algorithms mentioned in Section III, namely DDCH [23] and LRG [21]. We did not include Kuhn’s algorithm [24] in our comparison study because firstly the algorithm uses a purely synchronous model for communication and secondly when the number of rounds k is small, performance of the algorithm is much worse than the other three distributed algorithms we studied.

We ran all the three distributed algorithms and the centralized greedy algorithm on two CDNs with different scales. One consists of 30 servers of PlanetLab, the other consists of 80 servers. For each scale, algorithms were run at three different times and for three different latency bound values. We averaged the NOR of 10 invocations of our distributed algorithm. The results are summarized in Table I where our distributed algorithm is labeled “DG”. In this experiment, the centralized algorithm was implemented by a coordinator that we setup to collect probing results from the individual servers. Based on these probing results and the specified latency bound value, the coordinator would construct the network topology and run the centralized greedy algorithm.

Observe that although we used the same set of servers in all experiments of same scale, network conditions were slightly different, which resulted in a different NOR at different experiments for the same latency bound. We can see that compared with the centralized algorithm, the average NOR of our distributed replica selection algorithm is only slightly higher than that of the greedy algorithm when network scale is relatively small (30 nodes). When network scale is large (80 nodes), in some settings our distributed algorithm even yielded smaller NORs than the centralized algorithm. Also, the number of replicas as a fraction of the total network size did not increase with the increase in network scale. This makes us believe our distributed algorithm is simple yet has very good scalability. Moreover, it justifies using only 30 nodes in the remaining experiments instead of 80, as the results remain representative. Comparing the three distributed algorithms, DDCH, LRG and DG, we can see that no one clearly outperforms the other in terms of minimizing NOR. However, our algorithm is the simplest of the three and hence has an advantage in terms of implementation convenience.

Another metric we looked at is the termination time of the distributed algorithms which is the time elapsed until all nodes are covered. As mentioned in Section IV-B, when

30 Nodes	Centralized			DDCH			LRG			DG		
	200ms	300ms	400ms	200ms	300ms	400ms	200ms	300ms	400ms	200ms	300ms	400ms
set 1	8	7	6	8.8	7.7	6.3	9	7.8	6.5	8.6	7.5	6.4
set 2	6	5	5	6.5	5.5	5.5	6.4	5.5	5.8	6.9	5.6	5.2
set 3	4	4	4	4.4	4.3	4	5	4.8	4.1	5	4.7	4.6
80 Nodes	200ms	300ms	400ms	200ms	300ms	400ms	200ms	300ms	400ms	200ms	300ms	400ms
set 4	19	18	17	21.6	18.8	16.3	21.2	18.9	16.7	20.6	18.3	16.5
set 5	17	16	14	17.5	16.7	14.2	17.3	16.6	14.5	17.6	16.4	14.8
set 6	13	12	12	13.6	12.2	11.1	13.3	12	11.3	14	12.9	11.4

TABLE I. Number of Replicas

running in asynchronous mode, the algorithms will be executed periodically. For the fairness of comparison, we set the same period for all the three distributed algorithms. Note that our distributed algorithm can finish in constant time regardless of the network scale and topology. We define *normalized termination times* as the ratio of an algorithm’s termination time over our distributed algorithm’s termination time. Figure 2 plots the normalized termination times of DDCH and LRG in the experiment sets of Table I. From the figure, we can see that both

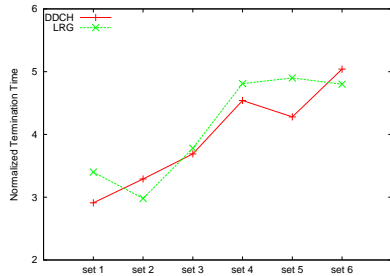


Fig. 2. Normalized Termination Time

DDCH and LRG take much longer to terminate. Another phenomenon observed about the two algorithms that makes them less attractive is that the normalized termination times for the 80-node network (set 4, 5, 6) are generally higher than those for the 30-node network. We conclude that our algorithm fits the goal of our system, which is finding a small dominating set very fast, and that it scales better than the rest. This algorithm will therefore be used in the rest of the evaluation section.

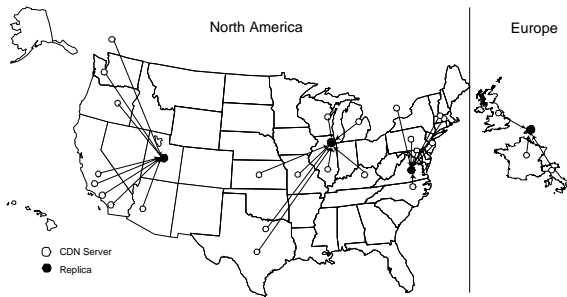


Fig. 3. One Instance of Experiments on PlanetLab

Figure 3 shows the actual replica locations in a representative instance of the above experiment. In this instance,

3 replicas are assigned in the United States and 1 in Europe. The figure also shows other CDN servers and the replicas they are connected to by our algorithm. Observe the natural clustering of CDN servers around their nearest replica, which provides an intuitive sanity check on the quality of our overlay structure formation.

B. Evaluation of the Latency Bound Guarantee

In this section, we evaluate our system’s ability to meet latency bounds with high confidence. In the previous section, we have chosen the best algorithm to use. The next question to ask is how well the algorithm performs in terms of achieving delay bounds. Note that meeting delay bounds is a different problem from the one typically addressed in previous efforts; namely, one of minimizing average latency. We confirm below that algorithms for solving the latter do not work well for the former.

To put ourselves at the maximum disadvantage, we compare our algorithm only against latency minimization algorithms that are centralized. A centralized algorithm has the advantage of global knowledge and hence the potential to produce better quality solutions. Nevertheless, we show that when it comes to meeting delay bounds, our algorithm outperforms its centralized rivals. Notice that, unlike our distributed replica selection algorithm which determines the NOR in a decentralized fashion on its own, all the centralized algorithms compared here take the NOR as an input parameter. For the fairness of our comparison, in each experiment, we first ran our distributed algorithm. After it fastly settled down, we collected information from all the CDN servers and determined the actual NOR yielded by the distributed algorithm. We then fed this value to all the centralized algorithms.

For the centralized algorithms we setup a centralized coordinator to synthesize information collected from each individual server. When enough information is obtained for the coordinator to construct the network topology, the centralized algorithms were invoked. The servers picked by the coordinator were then assigned as replicas. The other servers probed all the assigned replicas and, same as in our distributed algorithm, chose the one with the lowest latency. We use the following three centralized replica selection algorithms as baselines:

Single Server: In this baseline, every client request goes

to the origin server. No replicas are created. This algorithm is expected to have poor performance. We include it for completeness to offer a “worst-case” point of reference.

Random: This algorithm randomly chooses a specified number of replicas among the entire set of CDN nodes. This number is the same as that used in our distributed algorithm. This algorithm is included here to separate the performance enhancement attributed to the use of multiple replicas from the enhancement that results from their non-random proper placement by our algorithm.

Average Latency Greedy: A greedy network performance optimization algorithm has been studied in [7], [8] whose objective is to minimize overall (average) client perceived latency. We call this algorithm Average Latency Greedy (ALG). To minimize latency, ALG first computes the total cost associated with each server assuming that all the client requests received by all the N servers are forwarded to that server. It then picks the server that yields the lowest cost as a replica. This computation is repeated by looking for the next replica that, in conjunction with the previously selected ones, yields the lowest cost. The algorithm keeps selecting additional replicas until the NOR reaches a specified number. When computing cost, ALG assumes that the client requests received by a CDN non-replica server can always be directed to the nearest replica and takes the corresponding latency value as the cost. According to simulation results of [7], ALG performs very close to optimal in minimizing the average client perceived latency in a CDN.

In addition to the centralized algorithms mentioned above, we also introduce a variant of our original distributed algorithm. In this variant, each server nominates the nearest (as opposed to the highest-span) neighbor among all the neighbors that are within the latency bound. If no neighbor is reachable within latency bound, the server itself becomes a replica. We call the original algorithm and this variant “span-first” and “latency-first” distributed algorithms respectively. The latency-first algorithm explicitly attempts to minimize latency and guarantee the latency bound at the same time. We show that latency first can reduce the average latency but may increase NOR.

We installed the above baseline algorithms as well as our distributed algorithm on all servers. All algorithms were then run at the same time in different processes. The purpose is to make our system and baselines witness the same network conditions. For each CDN system process, clients were created that request files at a specified rate. Client processes run on servers in the same LAN as the servers running CDN system processes. We first evaluate the performance of the algorithms above for each of three different latency bound values. We collect client-perceived latencies of a total of 150,000 requests over the whole network and plot the cumulative distribution functions

(CDF) of latency values. Due to space limit, we only show the CDF of latencies downloading 16KB files here (Figure 4). Similar trends were observed for other file sizes.

From the graph we can tell that our distributed algorithms, span-first and latency-first, consistently perform better than the baselines. In Figure 4 the latency bound conformity ratios (i.e., percentage of requests that met their delay bounds) are 98%, 99% and 99% for latency bounds of 200ms, 300ms and 400ms respectively while those of ALG are 94%, 94% and 95%. The random algorithm has much poorer performance in terms of both latency bound conformity ratio and average latency. Having complete global information makes the ALG algorithm able to achieve comparable or even lower average latency but it is inferior to our distributed algorithms in providing the latency bound guarantee because it does not take the latency bound into consideration when selecting replicas. Figure 4 also shows that both random and ALG algorithm’s performance degrade as the latency bound value increases, because NOR decreases with the latency bound.

Compared to the span-first algorithm, latency-first can achieve the same or even higher conformity ratio and a lower average latency. However, it creates more replicas. The conformity ratio is better because latency-first nominates the lowest latency node among those that satisfy the delay bound. Hence, the system has a larger cushion against delay perturbations. The NOR is higher with latency-first because the metric used to nominate neighbors to become replicas (latency between two nodes) is a function of both the source and destination nodes. Hence, it may vary depending on the nominating node. In contrast, the span of a nominated node is a function of only that node. Hence, different nominating nodes are more likely to nominate the same neighbor. Consequently, the total number of nominated nodes decreases.

C. Effect of File Size

In our CDN system, each CDN node determines its set of neighbors using latency information. However, different file sizes have different latencies and web objects can essentially be of any size. Hence, we need techniques to estimate the latency of downloading an object as a function of file size using only a limited number of probes. Fortunately, our measurements show that the average network latency of downloading a file is roughly proportional to its size when the file size is between 1KB and 100KB. We chose this range as what our system targets for because on one hand this range encompasses the majority of web objects [26], [27] and on the other hand for large files (of multi-megabytes), like multimedia clips, maintaining a relatively constant downloading rate is the most important performance metric but is not the

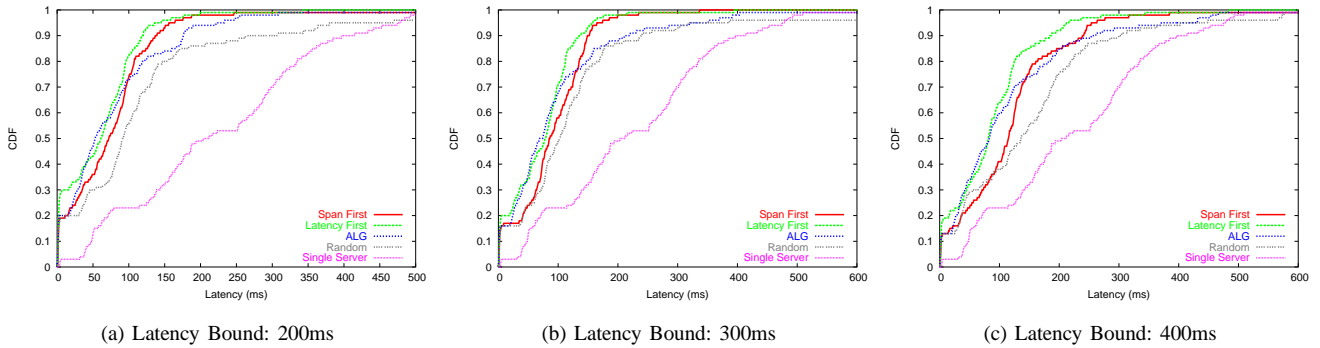


Fig. 4. Latency Bound Guarantee with Imperfect Visibility

focus of our system. Hence: $L(s) = As + B$. We leverage this property to quickly estimate the latency function. The estimation module first probes objects of several different sizes, then uses a recursive least square (RLS) estimator to estimate the parameters A and B of the linear equation above for the given client-server pair. We remove some outliers when doing the online estimation by setting a *noise factor* f . Noise factor simply means that if the result of a measurement is out of range $[E/f, f \times E]$ where E is the estimator’s prediction, the result of this measurement will not be fed into the estimator to avoid confusing it. Dynamic adjustment of estimator’s parameters helps our system adapt to network and system condition changes.

We validate our estimation technique by testing it on realistic Internet servers. We setup one node of PlanetLab as the server and chose four representative nodes as clients. The clients kept requesting files of sizes that range from 4KB to 100KB and ran the estimation algorithm. The experiment lasted for over 24 hours. Figure 5 plots the average latencies of files with different sizes and the estimation values of our algorithm. We can see that when the client is relatively close to the server (server #1, #2), our estimation is very accurate. When the client is far away (server #3, #4) but the file size is under 64KB, our estimation is still accurate. When client is far away and the file size is relatively big, the estimation error of our algorithm is no more than 10%.

While we believe that the above simple delay estimation is sufficient for our purposes, other more sophisticated network measurement/estimation techniques in existing research such as bprobe [28], Pathload [29], PTR/IGI [30], and Spruce [31] can be plugged into our system, if needed, without any impact on other modules.

D. Overload Avoidance

Our system includes explicit provisions for avoiding server overload. Firstly, when a server broadcasts its span to its neighbors, it also attaches its workload information. CDN servers will avoid nominating servers with very high

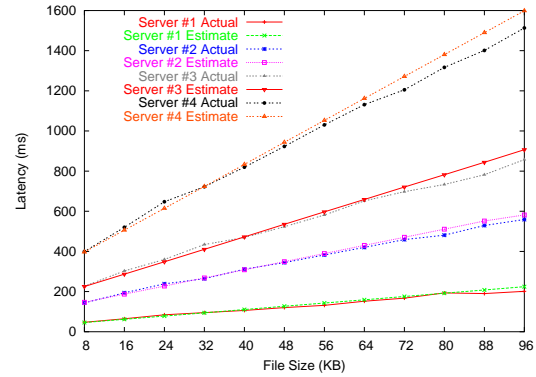


Fig. 5. Latency Estimation

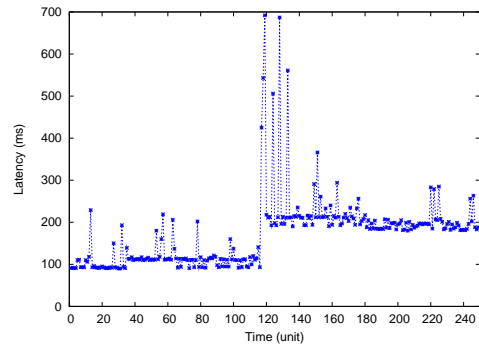


Fig. 6. Overloading Test

workload. Secondly, when a CDN server detects that it is overloaded, it can avoid being selected as replica by refraining from sending out its span to its neighbors. Lastly, when a CDN server is overloaded, its peers will notice an increase in the latencies seen when probing the server. Hence, the peers will exclude this server from being nominated as replica. When an overloaded CDN server is deselected as replica, some other replica(s) will be created. So when the overloaded server’s workload gets reduced and it is selected as a replica again later, it will share the total traffic with the newly created replica and therefore avoid becoming overloaded again.

When a server detects a high percentage of client

requests forwarded to its primary replica have recently exceeded the latency bound, it will pick one of its backup replicas, if any, or otherwise become a replica itself. In the following, we present an experiment that illustrates overload avoidance in our system. Since PlanetLab is a shared platform, we were explicitly discouraged from running extensive overload experiments on its servers. Instead, we used a server S on our LAN as the overload victim. A few PlanetLab servers then chose this server as their replica. Figure 6 plots the latencies of client requests and probings of one of the servers C that chose S as its replica. At time 115, we started multiple httpperf [32] clients in our LAN, sending requests at very high rates to S to overload it. C detected that latency to S became very high. Client requests directed to S continuously exceeded the 300ms latency bound. Server C then chose another replica as a temporary replica. Although the temporary replica was not as close to C as S , it was still within the latency bound. While using the temporary replica, C still periodically tried to fall back to its original replica S by probing the original replica, which are seen as the spikes from time 120 to 135 in the figure, because S used to have lower latency than its current temporary replica. After a few tries, C gave up and stayed with its temporary replica until the next invocation of distributed replica selection which happened from time 160 to 180. After that C acquired a new replica. From this experiment we can see that our system has the ability to avoid overloaded servers.

E. Visibility

An important consideration in the design of our algorithm was that it should perform well in the absence of any global knowledge shared among the CDN servers. This includes the knowledge of the identity of servers that are members of the CDN network. In practice, we believe that for moderately sized CDN networks that belong to the same administrative entity, it is not unreasonable to assume global knowledge of membership. Such knowledge can be preconfigured into the system and will not change frequently. However, larger systems with a decentralized administration can benefit from protocols that work well in the absence of any global information.

While many peer-to-peer frameworks adopt the extreme requirement that servers can only have a very local view of the network such as only a few immediate neighbors, our model lies in the middle in that we allow each server to know of some appreciable fraction of other servers. However, we do not require that this fraction be close to 1. We believe that our assumption is more appropriate for the current size of CDNs. Moreover, it avoids the negative performance implications of having to traverse application-layer multi-hop paths in a logical overlay topology, which may have a severe impact on the network’s ability to

efficiently meet subsecond delay bounds.

To evaluate the performance of our algorithm in the absence of global knowledge, we let each node know of only a certain percentage of other nodes of the whole system when it starts up. We call this parameter *visibility*. Imperfect visibility causes more replicas to be selected because some nodes with high spans are not visible to all other nodes. We evaluate two simple heuristics to reduce the NOR when visibility is low:

Reciprocal Mode: If servers know only an arbitrary subset of the network, cases may arise when server A knows server B but B does not know A . In this case, average visibility can be increased if each server simply introduced itself to all other servers it is aware of. This is called the reciprocal mode. Assuming there are N servers in the CDN system and every server’s initial visibility is x ($0 < x \leq 1$), then normally, the number of nodes each node knows is Nx . In the reciprocal mode, after each node introduces itself to the servers it initially knows, the average number of nodes each node knows increases by $Nx(1 - x)$, reaching a total of $N(2x - x^2) > Nx$. Thus, this mode increases the average visibility.

Highest Span Node Exchange: The reason why reducing visibility incurs higher NOR lies in that some high span nodes are not visible to all servers. To overcome this problem, we use another heuristic, namely, nodes exchange information of the node with the highest span among their known neighbors. When a node sends out its own span values, it attaches the highest span node it knows. When a node receives a span report and finds that the attached high span node is not known to itself, the node is added to its known server list. This technique does not increase the number of nodes each node knows by much, but gives nodes very important information for their replica selection process.

Figure 7 investigates the impact of reducing visibility on the NOR and the effectiveness of our heuristics in keeping the NOR small despite low visibility. For each latency bound and visibility value, we run our algorithms 10 times and take the average NOR. As shown in Figure 7, for all visibility values, the NOR of the system decreases with the increase in the latency bound. For all latency bound values, NOR increases as visibility decreases. Using the reciprocal mode greatly reduces NOR when visibility is above 0.5. An interesting observation from Figure 7 is that the system in the reciprocal mode with x as the initial visibility performs much better, in terms of minimizing NOR, than a system that simply drops packets from unknown servers, but has $2x - x^2$ as its initial visibility. This is true even though the average numbers of nodes each node ends up knowing about are the same in both cases. For example, the NOR at an initial visibility of 0.5 in the reciprocal mode is less than that at a visibility of 0.8 in a system that drops

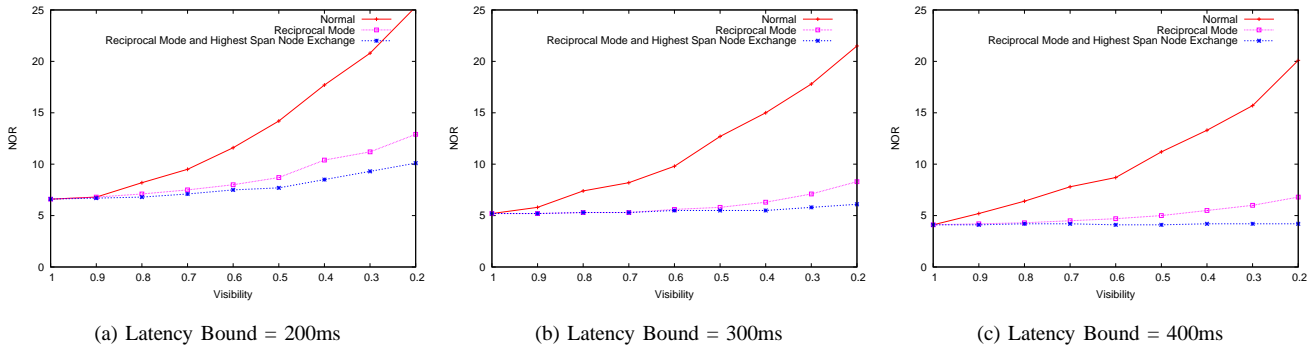


Fig. 7. NOR, Changing Visibility

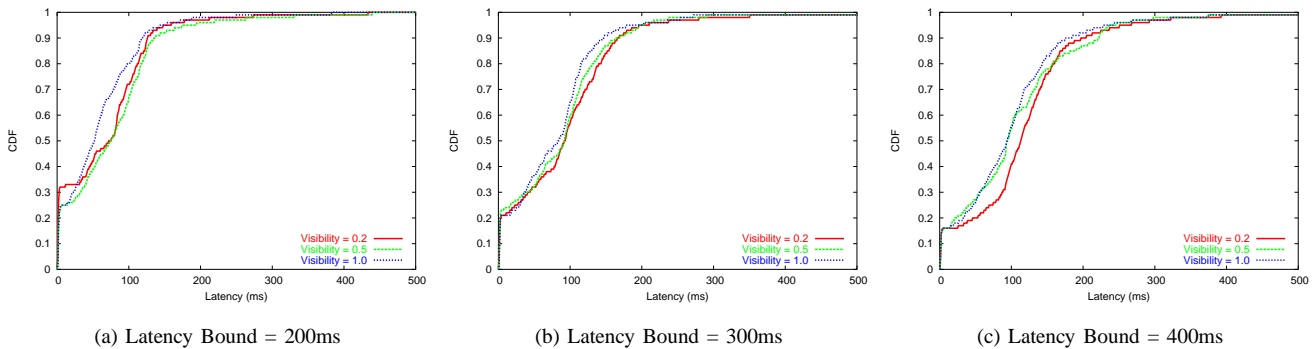


Fig. 8. Client Perceived Latencies with Imperfect Visibility

unknown server packets. The reason lies in the fact that in the reciprocal mode the peer knowledge relationships are bi-directional after the initial introduction, while that is not true in the second case.

When visibility is below 0.4, the reciprocal mode alone is not sufficient. However, using both this mode and the highest span node exchange heuristics can further reduce NOR. From Figure 7, we can tell that when both heuristics are used, the increase in NOR is negligible even when visibility is as low as 0.2, when latency bounds of 300ms and 400ms are requested. When the latency bound is very tight (200ms), the NOR increases somewhat more appreciably with decreased visibility. This is because when the latency bound is loose, more nodes have a high span and there is a better chance of finding a small number of replicas that dominate the network. When the latency bound is rigid, there are fewer choices for the dominating set.

In principle, imperfect visibility may also affect the latency bound conformity ratio because having low visibility makes it harder to find replicas that have low latencies. To show the effect of imperfect visibility on meeting latency bound guarantees, we ran three concurrent experiments configured with visibility values of 0.2, 0.5 and 1.0 respectively. Three processes were created on each

server, one on behalf of each experiment. These processes ran our distributed replica selection algorithm with both the reciprocal mode and the highest span node exchange heuristics enabled. The CDF of client perceived latencies of the whole network is plotted in Figure 8 in each case. Observe that in all cases the success ratio in meeting the bound remains very high, differing in less than 1% across visibility values. This result is very encouraging.

When latency bounds are more relaxed (Figures 8b and 8c), the system yields the same NOR's for all three visibility values. Average latencies increase slightly with reduced visibility. In the case where the bound is tight (200 ms), shown in Figure 8a, the NOR increases slightly with decreased visibility. An interesting consequence is that the average latency actually improves over that at a higher visibility as evident by comparing the latency curves at visibility values 0.2 and 0.5.

Based on the analysis and experiments above, we conclude that perfect visibility is not needed in our system. Our system's performance in terms of minimizing NOR and achieving a high confidence in meeting the latency bound are insensitive to the visibility parameter.

VI. Related Work

The efficacy of content distribution depends primarily

on the placement of replicas. While the replica placement problem in CDNs has been extensively studied in previous literature [6], [7], [8], [9], [10], the placement objectives have mostly been to optimize some *average* performance metric such as client-perceived latency, number of ASes traversed, or some notion of cost of link traversal.

Li et al. [6] simplified Internet topology as a tree and developed an optimal algorithm based on that topology. However, the authors did not evaluate the performance of their dynamic programming algorithm for a realistic Internet topology. Besides, the algorithm has a relatively high computational complexity $O(N^3M^2)$.

Qiu et al. [7] formulated the replica placement problem as that of choosing a fixed number of replicas among a given set of locations to minimize the overall client request latencies. This problem is mapped to a K-median problem in graph theory. The authors assumed that each individual CDN server has unlimited disk capacity and therefore each client uses a *single* replica. Based on simulations, a greedy algorithm was shown to perform very well and be relatively insensitive to imperfect input data. In the work by Jamin et al. [8], an algorithm called *Transit Node* has a similar flavor to our replica selection algorithm. The transit node algorithm assumes that the node with the highest degree can reach more nodes within a smaller latency. The algorithm therefore places replicas on candidate hosts in descending order of degree. In [8] the degree of a node is roughly the outdegree of the AS to which the node belongs. This is different from the degree concept in our distributed replica selection algorithm which is tightly related to network and system load conditions. A later work [10] further evaluated a router fanout-based replica placement algorithm and found that with careful design, the router-level fanout-based placement algorithm is almost as good as a greedy algorithm in [7] in most cases.

Kangasharju et al. explicitly took storage capacity of each individual CDN server into account and considered each AS as a node in a graph representing one CDN server [9]. The problem of optimizing average number of ASes traversed for client requests was formulated as a combinatorial optimization problem.

Venkataramani et al. [33] studied the replica placement problem of minimizing overall client access time under servers' bandwidth constraints. They used a simple hierarchical model: all cache/CDN servers are leaf nodes of the hierarchy tree. Intermediate nodes are dedicated for maintaining book-keeping information.

The above algorithms all share a same limitation, namely that they are all offline centralized algorithms. In a recent study [34] on general replica placement algorithms, a unified framework was designed to classify them. According to their findings, most existing algorithms for CDNs such as algorithms in [7], [8], [9] do not

scale for systems with more than 10^4 nodes because of the large number of content objects a CDN could host. They identified decentralized algorithms as an important direction for replica placement in CDNs. Our algorithm, being a purely distributed scheme, falls into this category, and can thus address scalability issues of existing replica placement algorithms.

SCAN [35] also studies a problem of minimizing the number of replicas while meeting client latency and server capacity constraints, but has a totally different context and solution in that it utilizes Tapestry [36] as the underlying distributed object routing and location system.

Rabinovich et al. [37] proposed a protocol suite for dynamic replication and migration of Internet objects. It has algorithms for deciding on the number and placement of replicas and an algorithm for distributing requests among available replicas. However their work does not address QoS guarantee issues.

Ko and Rubenstein [38] investigated an abstracted problem of placing replicas of different types of resources in large-scale network systems such as P2P networks and wireless sensor networks. They considered a different network model from ours in which each node must hold some resource. The placement is to achieve that any resource is reachable over a short path from any point in the network. However, their concept of distance between nodes is different from network latency of Internet. Hence, it is not clear how to easily extend their work to provide latency guarantee in the scenario described in this paper.

VII. Conclusions

In this paper, we presented a content distribution network (CDN) designed to provide latency bounds on content access and gave an extensive performance evaluation. We approached the delay guarantee problem by dynamically selecting replicas using a distributed algorithm that attempts to meet global latency bounds. The algorithm formulates the replica selection problem as a graph domination problem. Performance of our distributed replica selection algorithm in minimizing the overhead, the number of replicas, and meeting the delay bounds is evaluated and compared with other distributed graph domination algorithms as well as algorithms for latency minimization on an actual implementation on PlanetLab. Our results show that our system can achieve the needed latency with a very high confidence and at a very limited content replication cost. Being completely decentralized, our system enjoys the ability to adapt to various network and server condition changes and has very good scalability. For future work, we are interested in extending the performance evaluation with real traces of client traffic as opposed to synthetic workload. We are also interested in providing more mechanisms to tune the trade-offs between meeting

latency bounds and minimizing the number of replicas, as well as to allow the confidence in the latency guarantee to be a tunable parameter.

VIII. Acknowledgement

The work reported in this paper was supported in part by the National Science Foundation under grants CCR-0093144, ANI-0105873, EHS-0208769, and MURI N00014-01-1-0576. We are grateful to Ronghua Zhang and Ying Lu for their help and valuable discussions on this work.

References

- [1] PlanetLab, "<http://www.planet-lab.org>."
- [2] Akamai, "<http://www.akamai.com>."
- [3] A. Barbir, B. Cain, R. Nair, and O. Spatscheck, "Known content network (CN) request-routing mechanisms," July 2003.
- [4] L. Wang, V. Pai, and L. Peterson, "The effectiveness of request redirection on CDN robustness," in *Proceedings of The Fifth Symposium on Operating Systems Design and Implementation (OSDI'02)*, Boston, MA, December 2002.
- [5] V. Cardellini, M. Colajanni, and P. S. Yu, "Request redirection algorithms for distributed web systems," *IEEE Transactions on Parallel and Distributed Systems*, April 2003.
- [6] B. Li, M. Golin, G. Italiano, , and K. Sohraby, "On the optimal placement of web proxies in the Internet," in *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [7] L. Qiu, V. N. Padmanabhan, and G. M. Voelker, "On the placement of web server replicas," in *Proceedings of IEEE INFOCOM'01*, April 2001.
- [8] S. Jamin, C. Jin, A. R. Kurc, D. Raz, and Y. Shavitt, "Constrained mirror placement on the Internet," in *Proceedings of IEEE INFOCOM'01*, April 2001.
- [9] J. Kangasharju, J. Roberts, and K. W. Ross, "Object replication strategies in content distribution networks," in *Proceedings of The 6th International Web Caching Workshop and Content Delivery Workshop (WCW'01)*, Boston, MA, June 2001.
- [10] P. Radoslavov, R. Govindan, and D. Estrin, "Topology-informed Internet replica placement," in *Proceedings of The 6th International Web Caching Workshop and Content Delivery Workshop (WCW'01)*, Boston, MA, June 2001.
- [11] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for Internet applications," in *Proceedings of ACM SIGCOMM'01*, Aug 2001.
- [12] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A scalable content addressable network," in *Proceedings of ACM SIGCOMM'01*, 2001.
- [13] A. Rowstron and P. Druschel, "Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems," in *Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, Heidelberg, Germany, November 2001.
- [14] J. Yin, L. Alvisi, M. Dahlin, and A. Iyengar, "Engineering server-driven consistency for large scale dynamic web services," in *Proceedings of International World Wide Web Conference (WWW'01)*, 2001.
- [15] A. Ninan, P. Kulkarni, P. Shenoy, K. Ramamritham, and R. Tewari, "Cooperative leases: scalable consistency maintenance in content distribution networks," in *Proceedings of International World Wide Web Conference (WWW'02)*, 2002.
- [16] K. L. Johnson, J. F. Carr, M. S. Day, and M. F. Kaashoek, "The measured performance of content distribution networks," in *Proceedings of the 5th International Web Caching Workshop and Content Delivery Workshop (WCW'00)*, May 2000.
- [17] B. Krishnamurthy, C. Wills, and Y. Zhang, "On the use and performance of content distribution networks," in *Proceedings ACM SIGCOMM Internet Measurement Workshop 2001*, November 2001.
- [18] C. Huang and T. Abdelzaher, "Towards content distribution networks with latency guarantees," University of Virginia, Tech. Rep. UVA CS-2004-15, 2004.
- [19] M. R. Garey and D. S. Johnson, *Computers and Intractability*. W. H. Freeman, 1979.
- [20] V. Chvátal, "A greedy heuristic for the set-covering problem," in *Mathematics of Operations Research*, 1979.
- [21] L. Jia, R. Rajaraman, and T. Suel, "An efficient distributed algorithm for constructing small dominating sets," in *Proceedings of the 20th ACM Symposium on Principles of Distributed Computing (PODC'01)*, August 2001.
- [22] D. S. Johnson, "Approximation algorithms for combinatorial problems," in *Journal of Computer and System Sciences*, 1974.
- [23] B. Liang and Z. J. Hass, "Virtual backbone generation and maintenance in ad hoc network mobility management," in *Proceedings of IEEE INFOCOM'00*, March 2000.
- [24] F. Kuhn and R. Wattenhofer, "Constant-time distributed dominating set approximation," in *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC'03)*, July 2003.
- [25] Squid Web Cache, "<http://www.squid-cache.org>."
- [26] M. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," in *Proceedings of SIGMETRICS'96*, May 1996.
- [27] S. Saroiu, K. P. Gummadi, R. Dunn, S. D. Gribble, and H. M. Levy, "An analysis of Internet content delivery systems," in *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI'02)*, Boston, MA, December 2002.
- [28] R. L. Carter and M. E. Crovella, "Server selection using bandwidth probing in wide-area networks," in *Proceedings of IEEE INFOCOM'97*, 1997.
- [29] M. Jain and C. Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with tcp throughput," in *Proceedings of ACM SIGCOMM'02*, August 2002.
- [30] N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *IEEE JSAC Special Issue in Internet and WWW Measurement, Mapping, and Modeling*, August 2003.
- [31] J. Strauss, D. Katabi, and M. F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *Proceedings of Internet Measurement Conference (IMC) 2003*, October 2003.
- [32] D. Mosberger and T. Jin, "httpperf: A tool for measuring web server performance," in *Proceedings of The First Workshop on Internet Server Performance*, June 1998.
- [33] A. Venkataramani, P. Weidmann, and M. Dahlin, "Bandwidth constrained placement in a WAN," in *Proceedings of ACM Principles of Distributed Computing (PODC'01)*, August 2001.
- [34] M. Karlsson, C. Karamanolis, and M. Mahalingam, "A framework for evaluating replica placement algorithms," HP Labs, Tech. Rep. HPL-2002-219, 2002.
- [35] Y. Chen, R. H. Katz, and J. D. Kubiawicz, "Scan: A dynamic, scalable, and efficient content distribution network," in *Proceedings of The International Conference on Pervasive Computing (Pervasive'02)*, August 2002.
- [36] B. Y. Zhao, J. Kubiawicz, and A. D. Joseph, "Tapestry: An infrastructure for fault-tolerant wide-area location and routing," UC Berkeley, Tech. Rep. UCB/CSD-01-1141, 2001.
- [37] M. Rabinovich, I. Rabinovich, R. Rajaraman, and A. Aggarwal, "A dynamic object replication and migration protocol for an Internet hosting service," in *Proceedings of International Conference on Distributed Computing Systems (ICDCS'99)*, 1999.
- [38] B.-J. Ko and D. Rubenstein, "Distributed, self-stabilizing placement of replicated resources in emerging networks," in *Proceedings of the 11th International Conference on Network Protocols (ICNP'03)*, November 2003.