

WPI-CS-TR-03-35

Dec. 2003

Using Related Domain Names to Improve DNS
Performance

by

Hao Shang
Craig E. Wills

Computer Science
Technical Report
Series

WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

Abstract

In this paper, we present a novel approach to exploit the relationships among domain names to improve the cache hit rate for a local DNS server. Using these relationships, an authoritative DNS server (ADNS) can piggyback resolutions for future queries as part of the response message for an initial query. The approach improves the cache hit rate as well as reducing the total queries and responses. The approach is particularly attractive because it can be implemented with no changes to the existing DNS protocol.

Trace-base simulations show more than 50% of cache misses can be reduced in the best case while straightforward policies, using frequency and relevancy data for an ADNS, reduce cache misses by 25-40%. These percentages improve if we focus the policies on resource records with smaller ATTLS. We also show improved performance for hybrid approaches that combine the approach with renewal-based approaches.

Keywords: Domain Name System, Correlation

1 Introduction

The study is motivated by our concurrent research on correlations among network data flows¹. We find many flows have a close relationship with each other. They happen concurrently or sequentially in a short period of time. Many application flows such as streaming, Web or Instant Messaging are preceded by a DNS (Domain Name System) flow. Applications like the Web also generate many concurrent HTTP flows between a client and server. These correlations exist due to either inherent behaviors of applications or user access patterns.

Our overall work seeks to exploit relationships between flows for improved performance. Some relationships have been exploited to improve application performance. The persistent connection mechanism specified in HTTP/1.1 [6] is motivated by the observation of many short concurrent or sequential connections existing between two end hosts. Motivated by the same problem, a previous study from the same authors uses an approach that bundles multiple objects in one response [15]. Krishnamurthy et al. proposed a DNS-enabled Web approach that use DNS messages to piggyback Web content [9], which seeks to exploit the relationship between DNS and HTTP flows. Another study [16] presents the methodology to use DNS queries to infer the relative popularity of any Internet server that can be identified by a name. It exploits the relationship that DNS lookups foreshadow the access of those Internet servers. Unlike these previous works the work presented in this paper is focused on exploiting the relationships that exist within a DNS flow.

As part of our research on studying flow relationships, we found that a local domain name server (LDNS) frequently sends more than one query to the same authoritative domain name server (ADNS) for different names within a short period of time. Using network flow data obtained for the WPI campus network in July 2003, we observed that over 40% of flows involving the DNS protocol for name server lookups result in multiple packet exchanges between client and server.

As a means to reduce multiple-packet DNS flows between a local DNS server and an authoritative DNS server, we hypothesize that in many cases the authoritative DNS server can predict subsequent requests by a local DNS server based on knowledge of site usage and history of its DNS accesses. For example, the content of Web pages at busy Web sites is often served by multiple servers at the site, each with distinct names. Similarly, streaming or Instant Messaging applications use their own servers and are often combined with access to Web servers.

If the ADNS can piggyback resolutions of those related names in the response to the first query, it will save the LDNS from sending further queries. We call this approach *Piggybacking Related Names* (PRN). It benefits end-users as they experience less latency for DNS lookups. It also benefits ADNSs as they receive fewer DNS requests. Assuming that the piggybacked resolutions do not require additional packets then the approach reduces the number of packets needing to be routed through the Internet resulting in less congestion.

In order to better understand the PRN approach, we briefly overview the DNS mechanism. DNS is a distributed database providing mappings between addresses and names [11, 12]. The domain name space is a tree-like structure, where each node has a label. Associated with each node, there is a set of resource records (RRs) that comprise the domain database. The whole database is divided into non-overlapping zones that are distributed among name servers. Each zone has one or more delegated name servers, which have the complete information for the zone.

¹ A TCP flow is defined as all packets having the same (srcIP, dstIP, prot, srcPort, dstPort) tuple between SYN and FIN packets. A UDP flow is also bounded by the same tuple, but expires after a small amount of idle time (e.g. 5 minutes).

These name servers are called authoritative domain name servers (ADNSs) for the zone. There are many types of RRs. The most common one is type A that gives the mapping from a domain name to an IP address. For each RR, there is an associated time to live (TTL) parameter that indicates how long the RR can be cached.

The most important function of DNS is to return IP addresses for a given domain name. The process is typically initiated by a local application that calls an underlying resolver routine and passes the name as a parameter. The resolver sends a query to the local domain name server (LDNS), which in turn sends queries to responsible domain name servers if RRs for the name are not cached locally. The LDNS iteratively communicates with a list of name servers from top level domain servers to the authoritative domain name server for the name. Once the LDNS obtains the resolution, it returns it to the caller application.

A domain name lookup is required for any application that identifies servers by names instead of IP addresses. The latency incurred by the lookup procedure can influence the application's performance as a whole. In a previous study [17] we found that the median and mean lookup time for non-cached domain names is on the order of several hundred milliseconds. About 20% of the lookups took more than one second. Cohen et al. [3] indicates that DNS lookup time exceeds three seconds for over 10% of Web servers. This result is consistent with the measurement conducted by Chandranmenon and Varghese [2]. Jung et al. showing that 10-20% of DNS lookups take more than one second based on traces collected in MIT and KAIST [8]. A more recent study shows the average latency for resolving non-cached domain names ranges from 0.95 seconds to 2.31 seconds for a variety of clients [10]. All of these measurements suggest that the DNS lookup time for non-cached domain names can influence the performance of interactive applications particularly when congestion and packet loss occur.

Caching is effective for reducing user perceived latency and adopted by most DNS implementations. Previous studies [17, 8, 7, 4, 2] have shown cache hit rate varies from 50% to 90% from different traces. However with the increasing number of networked applications, there are many DNS requests generated for a single application and many of the cached copies of DNS mappings have a short time-to-live value in the cache. The result is that many requests for non-cached or stale DNS entries still exist. For example, in one day of WPI network flow data we observed over 40,000 DNS flows per hour. Thus further reduction of cache misses is still necessary for improving application performance.

The PRN mechanism reduces the cache miss rate by ADNSs predicting and piggybacking DNS resolutions. We compare our PRN approach with related approaches [4, 7] that aim to reduce cache misses as well. In addition, we examine how the PRN approach can be used in combination with these other approaches.

The remainder of the paper begins with a discussion of our approach in Section 2 and how it compares with other works in Section 3. We investigate the potential usefulness for PRN in Section 4. Implementation and policy issues are discussed in Section 5. We evaluate the performance of different PRN policies in Section 6. Section 7 compares performance among PRN, other related approaches, and hybrid ones. We conclude with a summary of our findings and directions for future work in Section 8.

2 PRN Approach

The *Piggybacking Related Names* (PRN) approach is motivated by the observation that many applications and related applications generate a sequence of DNS requests to the same ADNS for domain names within the same DNS zone of authority. The approach exploits the observation that most DNS packets are smaller than the allowed size of the UDP packet in which they are carried and hence there is potential for ADNSs to include “Additional Records” in response to a client’s request. In RFC1034, it says the “Additional Records” response field “carries RRs which may be helpful in using the RRs in the other sections.” In our work we propose that this field can also be used to contain additional RRs that the ADNS expects the client to subsequently request based upon the current request. As long as the additional records that are included do not cause the maximum allowable size of a DNS packet to be exceeded then these additional records are delivered to the client at essentially no additional cost on the packet-switched Internet.

Figure 1 illustrates the approach with the query/response dialogue between a LDNS client and an ADNS for resolution of multiple server names. In this example, names from a1.b.c to a5.b.c belong to zone b.c. The first query (in this example is a1.b.c.) triggers a response that includes resolutions for the additional names in the zone. Once it obtains the response, the LDNS caches all included entries. Subsequent queries for names a2.b.c, a3.b.c, a2.b.c and a4.b.c will be cache hits. We assume the interval between T6 and T1 is bigger than the ATTL for a2.b.c. So at T6, the entry for a2.b.c is already stale and a query for it causes a cache miss. A cache miss causes a new query to be sent to the ADNS by the LDNS.

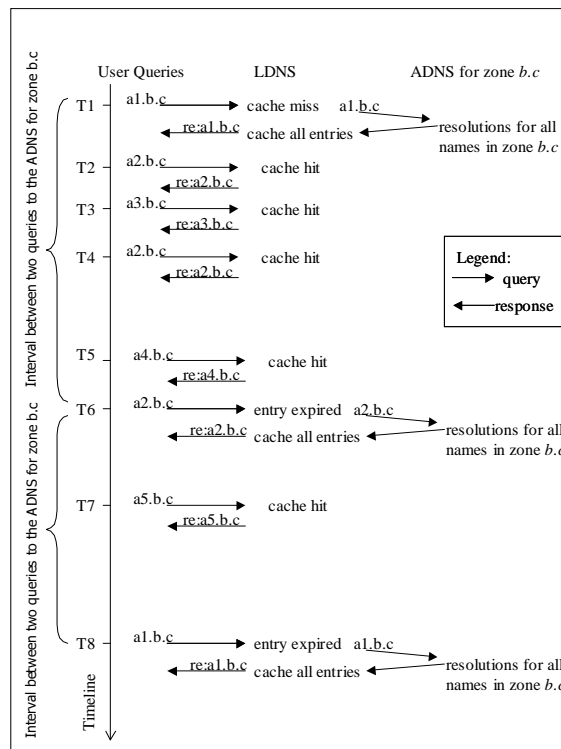


Figure 1: Illustration of Piggybacking Mechanism

It is obviously not practical for an ADNS to piggyback resolutions for all the names in a zone. In reality, an ADNS only needs to piggyback resolutions it expects to be used in the interval before the next query is needed. In the example both “T1 to T6” and “T6 to T8” are such intervals. In Section 5 we study the expected number of these additional records as well as the amount of available room in a DNS response.

A clear advantage of the *Piggybacking Related Names* (PRN) approach is that it requires no changes to the existing DNS protocol while reducing the amount of DNS traffic for local, authoritative and intermediate DNS servers. However the approach does require changes to the implementation of LDNSs and ADNSs. An ADNS must determine which names to be piggybacked and add them to the Additional Records section of an existing response message. This determination can be based on existing DNS queries as well as from knowledge of the site contents. A LDNS must extract the additional records and store them in its cache, which could be a problem in unnecessarily filling up the DNS cache, but in practice DNS cache records are small and cache space is not expected as a limitation. In our experiments, we assume all piggybacked records can be cached and will not be evicted before they expire. In the situation when cache space is limited, those piggybacked records can be tagged and be the first to be replaced if the cache is full.

One potential security issue with including resource records in the additional records field is a DNS-based attack called “cache poisoning” that is caused by allowing non-authoritative RRs to be cached by local DNS [1]. Our approach does not lead to this problem because a ADNS only piggybacks RRs for which it is the authoritative server.

3 Related Approaches

Previous research has examined other approaches for improving the cache hit rate at a LDNS. This section discusses three proposed approaches and compares them with the PRN approach.

One approach for improved DNS performance is for clients to pre-resolve server names [3]. This approach requires applications such as Web browsers to predict based on Web content which DNS lookups will be required and to issue those lookups before the content is retrieved. While this type of predictive policy is similar to the server-side predictions of our PRN approach it requires changes in applications and allows predictions to be made based only on client information.

A second approach is to use separate DNS queries to renew stale DNS cache entries [4]. This approach has the advantage that these queries are done outside of the critical path of an application and will improve the performance of these applications. The problem is that this approach can generate many additional DNS queries that must be handled by ADNSs and routers in the Internet.

The third approach is to piggyback requests for stale entries onto a needed request to an ADNS [7]. This “renewal using piggybacking” (RUP) approach causes no additional DNS packets to be generated, but requires each LDNS to organize all resource records according to their zone. This approach also requires that the DNS protocol support more than one request in a message.

We compare the approaches in terms of the types of cache misses that they avoid. Using the terminology of [4], cache misses can be divided into two types: “first-seen” (FS) misses, indicating the first lookup of a DNS name; and “previously-seen” (PS) misses, indicating entries that have been previously seen, but have expired. The two renewal approaches only reduce PS misses because they can only renew entries that have been seen before. The pre-resolving approach of [3] can reduce both FS and PS misses, but will only do so based on the immediate needs of the

application. The PRN approach not only reduces FS misses based on server knowledge, but if those entries are already cached, it can be used to restore these entries to their full TTL duration, thus reducing PS misses.

4 Potential Impact

Before looking at the details of implementing the PRN approach, an important question is to examine its potential impact. It is known that LDNS caches satisfy over 50% of DNS requests received from local applications. With these hit rates, an argument can be made that DNS performance is not a problem. However, a number of factors justify the need to further reduce the number of non-cached lookups. First, despite the high hit rates, a substantial number of DNS queries must still be satisfied by contacting the appropriate ADNS and as previously mentioned 40% of the DNS traffic in WPI flow data indicate multiple DNS requests. Second, recent studies have shown the average latency for these lookups to be on the order of one second. Third, in the presence of a dropped packet the delay is much larger as LDNSs use a three-second timeout. Fourth, more applications leads to more domain names at a site that must be looked up and many of these names carry shorter ATTLS to allow flexible load balancing.

We used three logs summarized in Table 1 to study the performance of DNS and examine the potential impact of improvements. The first log is of data from WPI’s primary DNS server, which serves as both a LDNS and ADNS for the campus. For our purposes the log was filtered to only consider queries from WPI clients that are handled by the server in its role as a LDNS. We augmented these data by fetching the ATTL and the ADNS(s) for each unique name in the log.

Table 1: Summary of Trace Logs Used

Name	Queries	Date	Dur.	From
WPI	1169569	Apr '03	28 hrs	WPI DNS
RTP	1041275	Oct '03	7 days	NLANR
SJ	457070	Oct '03	7 days	NLANR

The other two logs are generated from two NLANR Web traces [13] as done in [4]. Each entry of the Web trace is a request to an object identified by a uniform resource locator (URL). We extract the host name part from a URL as a query in a DNS trace. Because many browsers themselves cache name-to-IP address mappings for a short time, we make the same assumption as [4] that there is no DNS lookup incurred if a same name is requested again within a 60-second window.

We used these three logs along with the augmented data to determine the performance of DNS using a trace-driven simulation assuming the cache is empty at the beginning of the simulation. The simulation mimics the regular behavior of a DNS cache as well as an ideal behavior where whenever an ADNS receives a query, it returns resolutions for all the names in its zone. Subsequent queries that belong to this zone are satisfied locally as long as these entries are still fresh.

Results in Figure 2 show that the 26% of requests in the WPI DNS log result in misses and this percentage can be potentially reduced to 10% for a relative improvement of over 60%. Similar results are shown in Figure 3 for the RTP log where the percentage of total misses is over 45%

with a potential reduction to under 25% for a relative improvement of about 50%. Similar results were obtained for the SJ log and not shown.

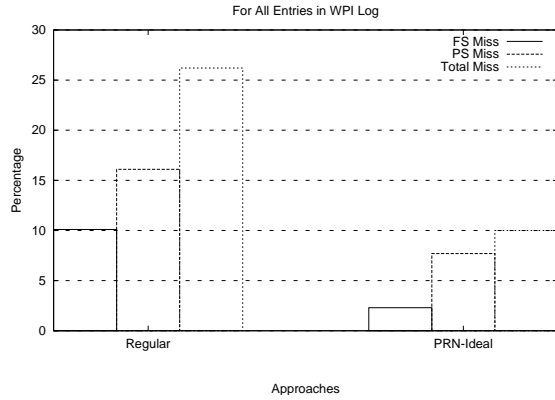


Figure 2: Potential Performance Improvement for Ideal PRN Policy with WPI Log

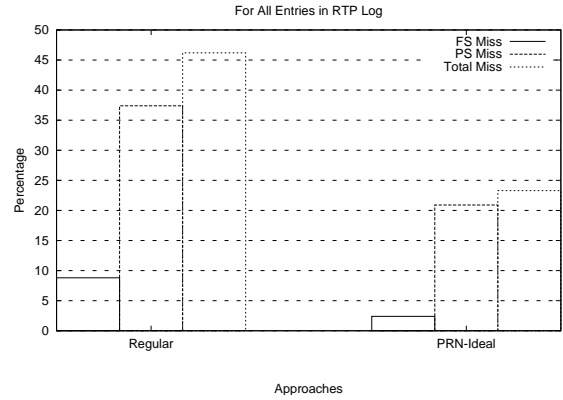


Figure 3: Potential Performance Improvement for Ideal PRN Policy with RTP Log

The collective results show significant reductions are possible in reducing both first-seen and previously-seen misses. Prediction of first-seen requests are not possible in renewal-based approaches while prediction of previously-seen requests extend the lifetime of the corresponding cached entries.

5 Implementation and Policy Issues

Having established the potential usefulness of the PRN approach, in this section we discuss specific implementation issues regarding the number of resource records that need to and can be piggybacked on a DNS response. We also describe specific policies for an ADNS to make decisions on what records to piggyback and what information the ADNS must maintain for these policies.

5.1 Piggybacked Responses

We used the data from the WPI DNS log to determine the number of responses that would ideally be piggybacked on a response. We used the request intervals in Figure 1 to define DNS bundles for a zone. A DNS bundle includes all unique server names for the zone that occur in a request interval. The size of this bundle determines the number of DNS responses that would be useful for the ADNS of the zone to return.

Using this definition, we found about 20,000 DNS bundles in the DNS log. Figure 4 shows the cumulative distribution function (CDF) for the number of names inside each bundle. As shown, about half of bundles have only one name—the response itself while the other half have two or more names. The results show that over 85% of the bundles have no more than five entries and over 95% of the bundles have no more than 15 entries. These results are encouraging for the PRN approach as they indicate it is useful for half of the bundles and that the number of names that need to be piggybacked is not large.

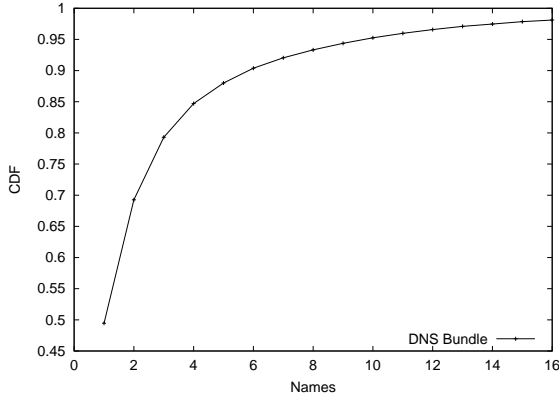


Figure 4: CDF of Size of DNS Bundles

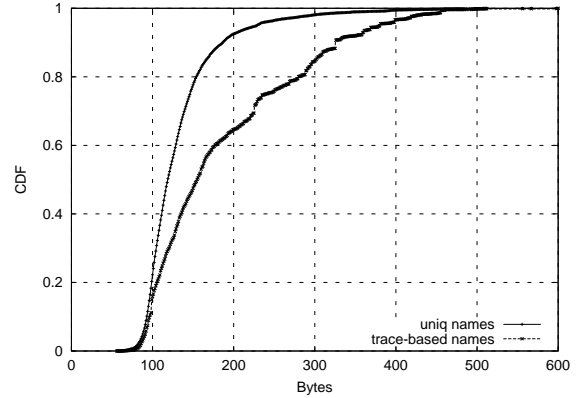


Figure 5: CDF of Sizes for DNS Response Packets on a Unique Name and a Trace-Based Set

5.2 DNS Response Message Capacity

DNS messages are limited to 512 bytes in size when sent over UDP [12], however DNS extension mechanisms [14] extend the limit to 1280 bytes. These mechanisms are supported in the latest 9.0 version of the widely-used BIND software [5]. We checked the sizes of DNS response packets for the 164K unique domain names collected from the three logs in Table 1. The CDF for the response message size for the unique names as well as for message sizes based on access patterns are shown in Figure 5. The trace-based distribution shows packet sizes are bigger than for unique names, which is as expected. With respect to the trace-based statistic, most responses are 100-300 bytes, which affords 200-400 remaining bytes if we use the traditional 512B limit and many more bytes if we use the limit for extended DNS.

Given this amount of available room, how many records can be added? The records to be piggybacked are mappings from domain names to IP addresses, which are typically type “A” RRs and possibly type “CNAME” RRs. If we only consider type “A” RRs in IPv4, the size of all its fields are fixed except the name. While a domain name can be long, it is not necessary to put the full name in that field. DNS provides a mechanism that enables domain names to share their common suffix. As all those names come from the same zone, it is common they share their suffix. By the same trace-based statistics, we observe that over 90% of names have a first label less than 10 characters while the median and average are between 4 and 5. The statistics excludes all names whose first label is “www” as they are likely to be the original queries. By putting those factors together, the length for a piggybacked record is most likely between 18 and 27 bytes (14 bytes for all fields with fixed length, 2 bytes for the pointer to the common suffix, 1 byte for the length count for the first label, and 1-10 bytes for the first label itself). With available space of 200-400 bytes, the total RRs that can be piggybacked are between 7 (200/27) and 22 (400/18). For extended DNS, the range is between 36 and 65.

We considered the situation when one domain name maps to multiple IP addresses, which requires multiple RRs for one name. We find that over 90% of domain names have less than five associated IP addresses while 72% have only one or two. Taking this factor into account, the total names that can be piggybacked are 1-22 for the traditional DNS length and 7-65 for the extended

DNS length.

5.3 Piggyback Policies

The previous two sets of results indicate that a sizeable percentage of the records that could be piggybacked will fit in the additional space of a DNS response. In cases where there are more potential names than can be piggybacked, an ADNS needs to have a policy to decide which names to include. In addition to the *ideal* policy, which we described in Section 4, we define two practical policies: *Most Frequently Queried (MFQ) First* and *Most Related Query (MRQ) First*. These policies are described as follows.

MFQ(n): The ADNS selects names in the order of their requested frequencies and the total count is bounded by n . For this policy, the ADNS needs to track query frequencies for each name in its zone and maintain them in a frequency ordered list (FOL).

MRQ(n, r): The ADNS selects names in the order of their relevancy to the current query. Each name has a relevancy value associated with another and it reflects the possibility that the former name will be queried after the latter. A value of “1” indicates a certain relationship while a value of “0” indicates no known relationship. It is practical to maintain a relevancy list for each name, which includes all names who has relation with it. The list will be organized in the order of their relevancy values and we refer it as relevancy ordered list (ROL). ROL(a) denotes the relevancy list for domain name “ a ”. For the MRQ policy, the total count of names can be piggybacked is still bounded by n . But rather than choosing from FOL, it first checks the ROL associated with the current query. All names in the ROL whose relevancy values are bigger than r are the first names to be considered. If there is still remaining space after they are filled in, names from the FOL are added.

```
1: twdns-01.ns.aol.com. : #zone cnn.com, identified by its first ADNS
2: i.cnn.net(4379) www.cnn.com(1494) sportsillustrated.cnn.com(588) money.cnn.com(263) ... fyi.cnn.com(1) # FOL
3: www.cnn.com(723) i.cnn.net(0.78) money.cnn.com(0.07) ... # ROL(www.cnn.com)
4: sportsillustrated.cnn.com(271) i.cnn.net(0.42) www.cnn.com(0.09) ... #ROL(sportsillustrated.cnn.com )
... ..
51: www.cnnfn.com(7) i.cnn.net(1.00) www.cnn.com(0.29) money.cnn.com(0.14) ... #ROL(www.cnnfn.com)
... ..
```

Figure 6: FOL and ROLs for zone “cnn.com.”

Figure 6 shows an example of a FOL and ROLs from the WPI DNS log. The first line contains the name of first ADNS (in sorted order) for the zone “cnn.com.” The second line is the FOL for the zone and has all names in the order of their query frequency. All subsequent lines are the ROLs for each name. The first element on each line is the name and the remaining element are its related names in the order of their relevancy values.

5.4 Maintenance of Information

As piggyback decisions are all made on the server side, it needs to maintain data structures to facilitate the functionality. Figure 6 is a typical structure that needs to be maintained by an ADNS

to support piggybacking. The FOL is simply a list whose each element includes a name and a counter recording its queried times. ROLs are normally organized as a hash table indexed by queries. It is not necessary to maintain a ROL for a query that does not have any related names. In the combination of the logs in Table 1, we observed the maximal number of names in a zone is 1650 and the maximal number of ROLs is 627.

Both the FOL and ROLs can either be set up manually by administrators who know the internal connections among names, or automatically by following query patterns. The example shown in Figure 6 is generated by analyzing the query patterns for the zone “cnn.com.” The FOL is created by counting queries to each name. The generating process for ROLs are more complicated. We group queries from the same client to the same ADNS that occur within a short period of time (5 minutes in our experiment). Whenever a name happens to be the first query in a group, the counter for its corresponding ROL is increased by 1. For all other names (after removing duplicates) in the group, each is counted once in the ROL for the first query. The relevancy value from name “a” to name “b” is calculated by dividing the counter of “b” in ROL(“a”) by the counter of ROL(“a”). For instance, in Figure 6, line 3 is the ROL for query “www.cnn.com”. The following number “723” is the count for the ROL and indicates there are 723 times “www.cnn.com” is the first query in a group. The number “0.78” following “i.cnn.net” is the relevancy value from “www.cnn.com” to “i.cnn.net”, which indicates out of 723, 78% of times “i.cnn.net” follows “www.cnn.com”. The relationship table is created based on the WPI DNS trace. It would be more accurate if reference patterns could be combined from a greater number of clients.

6 Evaluation

6.1 Methodology

We evaluate the PRN approach by trace-drive simulations of the ideal policy as described in Section 4 as well as the MRQ and MFQ policies described in Section 5.3. We use the relative decrease in the cache miss percentage as the metric to evaluate each policy. The regular policy is used as the baseline to compare effects of other policies with all results shown as the relative decrease in misses compared to the total number of first-seen (FS) and previously-seen (PS) misses for the regular DNS policy. Results for the ideal policy from Section 4 are shown for reference.

We studied the MFQ and MRQ policies with fixed upper bounds. We choose 5 and 15 as two upper bounds based on results from Section 5. In addition to these bounds, the MRQ policy is tested with relevancy values of 0.5 and 0 for a total of four MRQ policy combinations. Note that relevancy bound equal to 0 means the relevancy value should be bigger than 0, hence a qualified name must have some relevancy, even if weak.

6.2 Results

We used the first half of a log to generate relevancy tables and conducted the simulation on the second half of the log. The results are shown in Figure 7 and Figure 8 for the WPI and RTP log respectively. Simulation on the SJ log produced similar results as the RTP log and they are not shown here. In the figures, each set of bars corresponds to a policy. The first bar in a set shows the relative (to the total misses for the regular DNS approach) decrease of first-seen misses,

the second bar indicates the relative decrease of previously-seen misses, and the third bar is the relative decrease of the total misses.

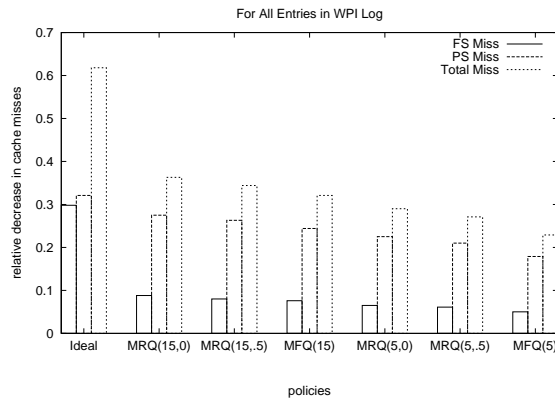


Figure 7: Relative Decrease in Cache Misses Over Different Policies on WPI Log

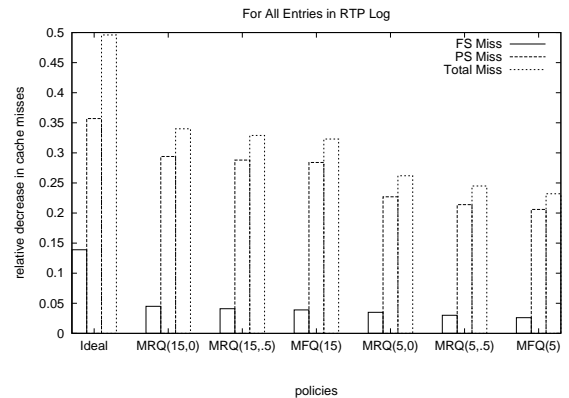


Figure 8: Relative Decrease in Cache Misses over Different Policies on RTP Log

The results show that both the FS misses and PS misses are significantly reduced when any piggyback policy is in use. The MRQ and MFQ policies reduce the total misses in the range from 25% to close to 40%. The results are significant because they are obtained by also reducing the number of queries by the same amount.

In terms of the policies, MRQ policies consistently outperform MFQ policies when they have same bound constraints. Among MRQ policies, those having smaller relevancy bound performs better. As $MFQ(n)$ is similar to $MRQ(n, 1)$, we can summarize the performance relationship among the policies as $MFQ(n) < MRQ(n, .5) < MRQ(n, 0)$. These results indicate names with relevancy, even weak, should be given higher preference than simply piggybacking popular names. Increasing the bound helps reduce cache misses, but even the smaller bound results in a 25% reduction in cache misses.

The same methodology is used for the RTP log with the results shown in Figure 8. The relative decrease in cache misses for this log varies between 25% and 35% with similar variation between the policies as we found with the WPI log.

6.3 Focused Results

As a means to test the PRN approach for resource records with relatively short ATTLS, we filtered the log for queries for ATTLS of 30 minutes or less. This filter removed roughly half of the original DNS requests. These records must be requested more frequently by a LDNS and we hypothesize that the PRN approach would be relatively more effective at reducing the number of cache misses. Results for this analysis for the WPI log are shown in Figure 9 where the total miss rate for regular DNS is 32% as compared to 26% in Figure 2. The results in Figure 9 show relative decreases in cache misses from nearly 30% to over 40%.

We pushed this analysis further and filtered the log to include only entries with an ATTL of 5 minutes or less. This filter removed roughly 80% of the log entries with 46% of requests for these entries resulting in a cache miss. As shown in Figure 10, the PRN policies reduce the cache miss

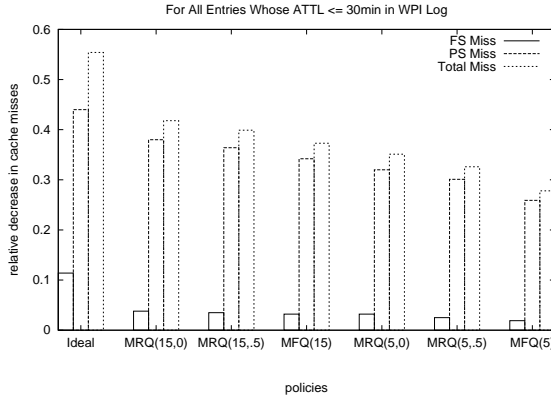


Figure 9: Relative Decrease in Cache Misses over Different Policies on WPI Log Entries with $ATTL \leq 30min$.

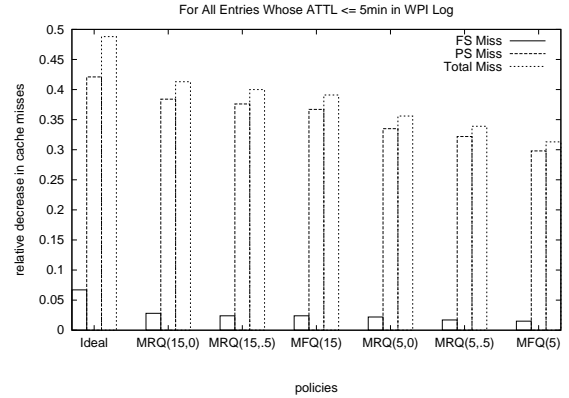


Figure 10: Relative Decrease in Cache Misses over Different Policies on WPI Log Entries with $ATTL \leq 5min$.

rate by over 40%. We found a similar tone of results when we did the same analysis for the RTP log. The results indicate this approach is more useful as the ATTLs grow shorter in duration.

7 Comparison and Combination with Other Approaches

Our final analysis was to compare the performance between our approach and others proposed to reduce cache misses. Furthermore, our approach is compatible with the others, which makes a combination with these approaches possible. We evaluate these hybrid approaches on all the three logs with results for the WPI log, using all log entries, shown.

7.1 Performance Comparison Among Approaches

The proactive caching approach proposed in [4] has several policies. Among them, R-LFU is one of the better and more straightforward policies. We implemented R-LFU(r) for comparison purpose. The renewal using piggyback (RUP) approach proposed in [7] also has several policies. We implemented RUP-MFU, which performs best among all practical approaches. Among our PRN policies, MRQ performs better than MFQ. We choose MRQ(15,0) for the comparisons between approaches. We refer it as PRN-MRQ(15,0).

Figure 11 shows the relative decrease in cache miss percentages for all the three approaches relative to the regular approach. The reduction in total cache misses is close for all the three approaches. When considering FS misses and PS misses separately, RUP-MFU and R-LFU policies behave almost the same, where FS misses are untouched and the reduction rates for PS misses are close. While PRN-MRQ does not reduce PS misses as much as the other two, its reduction on FS misses compensates for the difference. Despite the fact that the performance gains among the three approaches are similar, their costs are different. For the variation of the R-LFU policy we studied, it introduces 56% more queries as well as responses than the regular approach. The PRN-MRQ and RUP-MFU policies do not produce additional queries. Instead, by reducing the total misses, the total queries and responses are reduced as well.

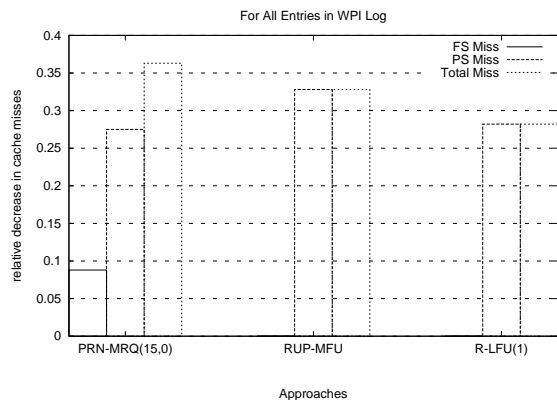


Figure 11: Performance Comparison among Approaches on WPI Log

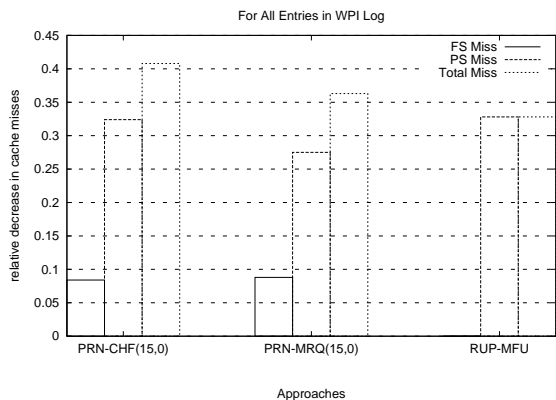


Figure 12: Performance Comparison among PRN-CHF and its Two Component Approaches on WPI log

7.2 Combination of PRN and RUP

In Figure 11 we observe that PRN-MRQ reduces more FS misses while RUP-MFU reduces more PS misses. This result encourages us to consider the possibility of combining the two approaches. Both approaches use piggybacking, but one makes the decision on the server side while the other does on the LDNS client side.

To combine these policies we define a new policy called “piggyback related names with client hint first” (PRN-CHF), where the client DNS server piggybacks its stale names in the query message and the ADNS uses these hints as well as its own relevancy table. The policy is described as:

PRN-CHF(n, r): The total number of names that can be piggybacked is bounded by n , but instead of first looking at the corresponding ROL, the ADNS gives priority to the names piggybacked in the query message. If there is still extra space left, the ROL and FOL are checked in turn.

We show the performance of PRN-CHF and its two component approaches in Figure 12. As we expected, PRN-CHF has the same FS miss rate as PRN-MRQ and the same PS miss rate as RUP-MFU, which makes it performs best among the three.

7.3 Combination of PRN and R-LFU

We also studied the combination of these policies with active renewal. As the R-LFU approach is initiated by the LDNS cache and the PRN approach is initiated by an ADNS, the two approaches can complement each other.

We show performance of the various hybrid approaches along with R-LFU in Figure 13. In order to distinguish our original PRN approaches from their hybrid versions with R-LFU, we refer to those three hybrid approaches with a prefix “R-” to their original names. As with the R-LFU approach, each approach is tested with different aggressiveness in prefetching. As aggressiveness increases, the cache misses are further reduced, but more queries are generated.

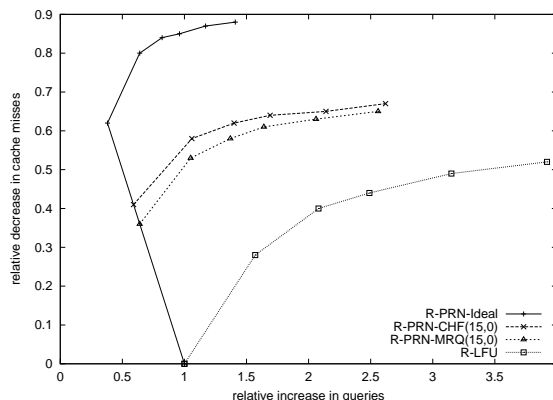


Figure 13: Performance for Hybrid Approaches on WPI Log

The hybrid approaches show significant performance gains in Figure 13 compared with either R-LFU or their original PRN approaches. With about the same number of queries, R-PRN-MRQ performs much better than R-LFU. For instance, having 1.5 times queries as the regular approach, R-PRN-MRQ reduces 54% of total misses while R-LFU reduces 28%. R-PRN-CHF performs slightly better than R-PRN-MRQ as it is the combination of the PRN, RUP and R-LFU approaches. Renewal also benefits the original PRN approach at the expense of more queries.

8 Summary and Future Work

This work is motivated by research on studying the relationships among network data flows. We found many cases where a local DNS server sends multiple DNS queries to the same authoritative DNS server within a short period of time. If the ADNS can predict these subsequent near-future queries once it receives the first one then it can send answers for all of them with the first response. We call this the piggybacking related names (PRN) approach. It helps reduce local cache misses and therefore reduces user-perceived DNS lookup latency. By piggybacking multiple answers in one response packet, the total queries and responses are also reduced, which alleviates the workload on both LDNSs and ADNSs as well as the intermediate network.

Compared with other approaches that also address improving local cache hit rate, our approach is novel. We explicitly use the relationships among queries and allow an ADNS to push resolutions for predicted names to the LDNS. The PRN approach reduces both first-seen misses as well as previously-seen misses while other approaches reduce just the latter. The cost of PRN is also low as it reduces the number of response packets while requiring no changes to the DNS protocol.

Trace-base simulations show more than 50% of cache misses can be reduced if prediction is perfect and response packet space is plentiful. Realistic policies, using frequency and relevancy data for an ADNS, reduce cache misses by 25-40%. These percentages improve if we focus the policies on resource records with smaller ATTLS. We also show improved performance by combining the PRN approach with renewal-based approaches to create hybrid approaches that perform significantly better than their component approaches.

An obvious direction for future work is to examine alternate policies such as ones to consider

the ATTL for an entry. Policies should also be tested with additional logs. Another direction of future work is to deploy the PRN approach at an ADNS. We expect it should perform better than our simulation because the ADNS has more complete knowledge of the site contents and it can also aggregate reference patterns from a greater number of client sites. A final direction is to explore the different types of sites and contents for which it is most useful. Sites with few servers and long authoritative TTLs likely do not need improvement in DNS performance while we expect more dynamic sites would be the first to benefit from this approach.

9 Acknowledgements

We thank the National Laboratory for Applied Network Research for making proxy logs available. Data collection at the NLANR is supported by National Science Foundation grants NCR-9616602 and NCR-9521745.

References

- [1] CERT/CC. Vulnerability Note VU109475.
<http://www.kb.cert.org/vuls/id/109475>.
- [2] Girish P. Chandranmenon and George Varghese. Reducing web latency using reference point caching. In *Proceedings of IEEE Infocom 2001*, pages 1607–1616, 2001.
- [3] Edith Cohen and Haim Kaplan. Prefetching the means for document transfer: A new approach for reducing web latency. In *Proceedings of the IEEE Infocom 2000 Conference*, Tel-Aviv, Israel, March 2000. IEEE.
- [4] Edith Cohen and Haim Kaplan. Proactive caching of DNS records: Addressing a performance bottleneck. In *Proceedings of the Symposium on Applications and the Internet*, pages 85–94, San Diego-Mission Valley, CA, USA, January 2001. IEEE-TCI.
- [5] Internet Software Consortium. BIND DNS Server.
<http://www.isc.org/products/BIND/>.
- [6] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. Hypertext Transfer Protocol – HTTP/1.1, June 1999. RFC 2616.
- [7] Baekcheol Jang and Kilnam Chon. DNS Resolution with Renewal Using Piggyback. In *Proceedings of the Twelfth International World Wide Web Conference (Poster)*, Budapest, Hungary, May 2003.
- [8] Jaeyeon Jung, Emil Sit, Hari Balakrishnan, and Robert Morris. DNS Performance and the Effectiveness of Caching. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop '01*, San Francisco, California, November 2001.
- [9] Balachander Krishnamurthy, Richard Liston, and Michael Rabinovich. Dew: Dns-enhanced web for faster content delivery. In *Proceedings of the Twelfth International World Wide Web Conference*, Budapest, Hungary, May 2003.
- [10] Richard Liston, Sridhar Srinivasan, and Ellen Zegura. Diversity in dns performance measures. In *Proceedings of the second ACM SIGCOMM Workshop on Internet measurement workshop*, pages 19–31, 2002.
- [11] P. Mockapetris. Domain Names - Concepts and Facilities, November 1987. RFC 1034.
- [12] P. Mockapetris. Domain Names - Implementation and Specification, November 1987. RFC 1035.
- [13] NLANR. network traffic packet header traces.
<http://pma.nlanr.net/Traces/>.
- [14] P. Vixie. Extension Mechanisms for DNS (EDNS0), August 1999. RFC 2671.

- [15] Craig E. Wills, Mikhail Mikhailov, and Hao Shang. N for the price of 1: Bundling web objects for more efficient content delivery. In *Proceedings of the Tenth International World Wide Web Conference*, Hong Kong, May 2001.
- [16] Craig E. Wills, Mikhail Mikhailov, and Hao Shang. Inferring relative popularity of internet applications by actively querying dns caches. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Miami, Florida, November 2003.
- [17] Craig E. Wills and Hao Shang. The contribution of dns lookup costs to web object retrieval. Technical Report WPI-CS-TR-00-12, Worcester Polytechnic Institute, July 2000.
<http://www.cs.wpi.edu/~cew/papers/tr00-12.ps.gz>.