

On the scalability of LISP mappings caches



Florin Coras*, Jordi Domingo-Pascual, Albert Cabellos-Aparicio

Departament d'Arquitectura de Computadors, Universitat Politècnica de Catalunya, BarcelonaTech (UPC), Barcelona 08034, Spain

ARTICLE INFO

Article history:

Received 4 October 2014

Revised 27 July 2015

Accepted 20 August 2015

Available online 29 August 2015

Keywords:

Locator/ID Separation Protocol

LISP

Cache scalability

Temporal locality

ABSTRACT

The Locator/ID Separation Protocol (LISP) limits the growth of the Default-Free Zone routing tables by creating a highly aggregatable and quasi-static Internet core. However, LISP pushes the forwarding state to edge routers whose timely operation relies on caching of location to identity bindings. In this paper we develop an analytical model to study the asymptotic scalability of the LISP cache. Under the assumptions that (i) long-term popularity can be modeled as a Generalized Zipf distribution, independent of Internet and LISP site growth and (ii) temporal locality is predominantly determined by long-term popularity, we find that LISP cache miss rate scales $O(1)$ with respect to the amount of prefixes (Internet growth) and users (LISP site growth). We validate the model and discuss the accuracy of our assumptions using several one-day-long packet traces.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

The growth of the Default-Free Zone (DFZ) routing tables [1] and associated churn observed in recent years has led to much debate as to whether the current Internet infrastructure is architecturally unable to scale. Sources of the problem were found to be partly organic, generated by the ongoing growth of the topology, but also related to operational practices which seemed to be the main drivers behind prefix deaggregation within the Internet's core. Diverging opinions as to how the latter could be solved triggered a significant amount of research that finally materialized in several competing solutions (see [2] and the references therein).

In this paper we focus on location/identity separation type of approaches in general, and consider the Locator/ID Separation Protocol (LISP) [3] as their particular instantiation. LISP semantically decouples identity from location,

currently overloaded by IP addresses, by creating two separate namespaces that unambiguously address end-hosts (identifiers) and their Internet attachment points (locators). This new indirection level has the advantage that it supports the implementation of complex operational practices (e.g., traffic engineering mechanisms) but at the same time enables the locator space to remain quasi-static and highly aggregatable [4].

Although generally accepted that such solutions alleviate the scalability limitations of the DFZ, they also introduce new network elements, chiefly a mapping-system, and querying mechanisms needed for obtaining dynamic bindings (mappings) between the two new namespaces. To avoid growing router memory requirements with the size of the identity namespace, as is the case today, routers retrieve mappings according to user traffic. Therefore, to speed-up packet forwarding and to avoid generating floods of resolution requests, routers must store in use mappings in map-caches. This then begs the question: *does the newly introduced LISP edge cache scale?*

A considerable number of studies have empirically evaluated map-cache performance, however these results cannot be extrapolated to provide insight into what traffic parameters affect cache performance nor quantify their impact

* Corresponding author. Tel.: +0034 644350770.

E-mail addresses: fcoras@ac.upc.edu (F. Coras), jordi.domingo@ac.upc.edu (J. Domingo-Pascual), acabello@ac.upc.edu (A. Cabellos-Aparicio).

[5,6,7,8,9]. Additionally, in [10] and [11] we showed how the working-set [12] can be leveraged to estimate temporal locality of real network traces and finally to build a model that links miss rate and cache size. Nevertheless, despite its predictive power, the model is also unable to uncover locality sources and, ultimately, unable to predict by itself how map-cache performance is to scale.

This paper complements our work in [11] and provides an answer to the previous question by analyzing map-cache miss rate scalability with respect to Internet and LISP site growth. To this end we leverage well known results that characterize temporal locality of reference strings [13,14] to show that the relation between cache size and miss rate depends only on the popularity distribution of destinations. Notably, the result holds only if (i) long-term popularity can be modeled as a constant Generalized Zipf [15] and (ii) temporal locality is predominantly determined by long-term popularity. This further enables us to conclude that, for a given miss rate, cache size should scale constantly, $O(1)$, with respect to the growth of the Internet and LISP site size, if popularity is independent of the two. We believe this result is crucial for LISP's deployment since it ensures that if a map-cache is provisioned for a certain performance point, its performance should not degrade over time. If, however, the property does not hold, then the miss rate scales linearly, $O(N)$, with respect to the number of destinations in the Internet. To support our results, we analyzed the popularity distribution of destination prefixes in several one-day-long, real-world packet traces, from two different networks and spanning a period of 3.5 years. All assumptions were empirically validated using the traces.

The rest of the paper is structured as follows. We provide a brief overview of LISP in Section 2. In Section 3 we derive the cache model under a set of assumptions and thereafter discuss its predictions and implications for LISP. In Section 4 we present empirical evidence that supports our assumptions and evaluate the model, while in Section 5 we discuss the related work. Finally, we conclude the paper in Section 6.

2. LISP background

LISP [3] belongs to the family of proposals that implement a location/identity split in order to address the scalability concerns of the current Internet architecture. The protocol specification has recently undergone IETF standardization [16], however development and deployment efforts are still ongoing. They are supported by a sizable community spanning both academia and industry and rely for testing on a large experimental network, the LISP-beta network [17].

The goal of splitting location and identity is to insulate core network routing that should ideally only be aware of location information (locators), from the dynamics of edge networks, which should be concerned with the delivery of information based on identity (identifiers). To facilitate the transition from the current infrastructure, LISP numbers both namespaces using the existing IP addressing scheme, thus ensuring that routing within both core and stub networks stays unaltered. However, as locators and identifiers bear relevance only within their respective namespaces, a form of conversion from one to the other must be performed. LISP




Fig. 1. Example packet exchange between EID_{SRC} and EID_{DST} with LISP. Following intra-domain routing, packets reach xTR_A which obtains a mapping binding EID_{DST} to $RLOC_{B1}$ and $RLOC_{B2}$ from the mapping-system (steps 1–3). From the mapping, xTR_A chooses $RLOC_{B1}$ as destination and then forwards toward it the encapsulated packets over the Internet's core (step 4). xTR_B decapsulates the packets and forwards them to their intended destination.

makes use of encapsulation [18] and a directory service to perform such translation.

Prior to forwarding a host generated packet, a LISP router maps the destination address, or Endpoint Identifier (EID), to a corresponding destination Routing Locator (RLOC) by means of a LISP specific mapping system [8,19]. Once a mapping is obtained, the border router tunnels the packet from source edge to corresponding destination edge network by means of an encapsulation with a LISP-UDP-IP header. The outer IP header addresses are the RLOCs pertaining to the corresponding border routers (see Fig. 1). At the receiving router, the packet is decapsulated and forwarded to its intended destination. In LISP parlance, the source router, that performs the encapsulation, is called an Ingress Tunnel Router (ITR) whereas the one performing the decapsulation is named the Egress Tunnel Router (ETR). One that performs both functions is referred to as an xTR .

Since the packet throughput of an ITR is highly dependent on the time needed to obtain a mapping, but also to avoid overloading the mapping-system, ITRs are provisioned with map-caches that store recently used EID-prefix-to-RLOC mappings. Stale entries are avoided with the help of timeouts, called *time to live* (TTL), that mappings carry as attributes. Whereas, consistency is ensured by proactive LISP mechanisms through which the xTR owner of an updated mapping informs its peers of the change. Intuitively, the map-cache is most efficient in situations when destination EIDs present high temporal and/or spatial locality and its size depends on the diversity of the visited destinations. As a result, performance depends entirely on map-cache provisioned size, traffic characteristics and the eviction policy set in place.

3. Cache model

We start this section by discussing some of the fundamental properties of network traffic that may be exploited to gain a better understanding of cache performance. Then, assuming these properties are characteristic to real network traces we devise a cache model. Finally we analyze and discuss the predictions of the model. Table 1 gives an overview of the main parameters and variables used in defining the cache model.

Table 1
Cache model variables and parameters.

Variable	Description
N	Length of the reference string (number of packets in a trace)
D	Set of destination IP-prefixes
r_i	Reference at the i th unit of time
t	Time, measured in packets
ν	Frequency
p_ν	Probability of object with frequency ν in reference string
$d(t, \nu)$	Number of times the inter-reference distance for an object with frequency ν is t
r and q	Exponents that control GZipf power law slopes in the two regimes
λ and μ	Constants that control the frequency for which GZipf undergoes the crossover
$f(t)$	Inter-reference distance density function
$m(t)$	Working-set miss rate
$s(t)$	Average working-set size
$H(n, m)$	The generalized harmonic number of order n of m
$\Gamma(n, z)$	The incomplete Gamma function
$\zeta(s, a)$	The Hurwitz Zeta function

3.1. Sources of temporal locality in network traffic

We consider the following formalization of packet level traffic throughout the rest of the paper. Let D be a set of objects (i.e., destination IP-prefixes). Then, we consider traffic to be a strings of references $r_1, r_2, \dots, r_i \dots$ where $r_i = o \in D$ is a reference at the i th unit of time that has as destination, or requests, object o . Generally, we consider the length of the reference string to be N . Also, note that we use object and destination interchangeably.

Two of the defining properties of reference strings, important in characterizing cache performance, are the heavy tailed *popularity distribution* of destinations and the *temporal locality* exhibited by the requests pattern. We discuss both in what follows.

Popularity distribution: Copious amounts of studies in fields varied as linguistics [15,20], Web traffic [13,21], video-on-demand [22], p2p overlays [23] and flow level traffic [24] found the probability distribution of objects to have a positive skew. Generally, such distributions are coined Zipf-like, i.e., they follow a power law; whereby the probability of reference is inversely proportional to the rank of an object. Generally, the relation is surmised as: $\nu(k) = \frac{\Omega}{k^\alpha}$ where ν is the frequency, or number of requests observed for an object, k is the rank, $\Omega = 1/H(n, \alpha)$ is a normalizing constant and $H(n, \alpha)$ is the n^{th} generalized harmonic number.

It is interesting to note that although Zipf's law has its origins in linguistics, it was found to be a poor fit for the statistical behavior of words frequencies with low or mid-to-high values of the rank variable. That is, it does not fit the head and tail of the distribution. Furthermore, its extension due to Mandelbrot (often called the Zipf–Mandelbrot law) only improves the fitting for the head of the distribution. Such discrepancies were also observed for Web based and p2p reference strings. Often the head of the distribution is flattened, i.e., frequency is less than the one predicted by the law, or the tail has an exponential cutoff or a faster power law decay

[15,23]. But these differences are usually dismissed on the basis of poor statistics in the high ranks region corresponding to objects with a very low frequency.

Nevertheless, Montemurro solved recently the problem in linguistics by extending the Zipf–Mandelbrot law such that for high ranks the tail undergoes a crossover to an exponential or larger exponent power-law decay. Surprisingly, he found this features, i.e. deviations from the Zipf-like behavior, to hold especially well when very large corpora [15] are considered. We further refer to this model as the Generalized Zipf law or GZipf and, in light of these observations, we assume the following:

Assumption 1. The popularity distribution of destination IP-prefix reference strings can be approximated by a GZipf distribution.

Temporal locality: It can be informally defined as the property that a recently referenced object has an increased probability of being re-referenced. One of the well established ways of measuring the degree of locality of reference strings is the inter-reference distance distribution.

Breslau et al. found in [13] that strings generated according to the Independent Reference Model (IRM) from a popularity distribution, that is, assuming that references are independent and identically distributed random variables, have an inter-reference distance distribution similar to that of the original string. Additionally, they inferred that the probability of an object being re-referenced after t units of time is proportional to $1/t$. Later, Jin and Bestavros proved that in fact temporal locality emerges from both long-term popularity and short-term correlations. However, they found that the inter-reference distance distribution is mainly induced through long-term popularity and therefore is insensitive to the latter. Additionally, they showed that by ignoring temporal correlations and assuming a Zipf-like popularity distribution, an object's re-reference probability after t units of time is proportional to $1/t^{(2-1/\alpha)}$. These observations then lead to our second assumption:

Assumption 2. Temporal locality in destination IP-prefix reference strings is mainly due to the prefix popularity distribution.

We contrast the two assumptions with the properties of several packet-level traces in Section 4. In what follows we are interested in characterizing the inter-reference distribution of a GZipf distribution and further on the cache miss rate using the two statements as support.

3.2. GZipf generated inter-reference distribution

In this section we compute the inter-reference distance distribution for a GZipf popularity. The result is an extension of the one due to Jin and Bestavros for a Zipf-like popularity [14]. As a first step we compute the inter-reference distribution for a single object and then by integration obtain the average for the whole reference string, which we denote by $f(t)$.

If ν is the normalized frequency, namely, the number of reference to an object divided by the length of the reference string N , then, as shown in [15] the probability of observing objects with frequency ν in the reference string is:

$$p_\nu(\nu) \propto \frac{1}{\mu\nu^r + (\lambda - \mu)\nu^q} \quad (1)$$

where $1 \leq r < q$ are the exponents that control the slope of the power laws in the two regimes and μ and λ are two constants that control the frequency for which the tail undergoes the crossover.

From Assumption 2 it follows that references to an object are independent whereby the inter-reference distance t is distributed exponentially with expected value of $1/\nu$. Then, if we denote by $d(t, \nu)$ the number of times the inter-reference distance for an object with frequency ν is t , we can write:

$$d(t, \nu) \sim (\nu N - 1)\nu e^{-\nu t} \quad (2)$$

If ν_{\min} and ν_{\max} are the minimum and respectively the maximum normalized frequency observed for the reference string, we can compute the inter-reference distance for the whole string as:

$$\begin{aligned} f(t) &\sim \int_{\nu_{\min}}^{\nu_{\max}} p_\nu(\nu) d(t, \nu) d\nu \\ &= \int_0^1 \frac{(\nu N - 1)\nu e^{-\nu t}}{\mu\nu^r + (\lambda - \mu)\nu^q} d\nu \end{aligned} \quad (3)$$

Unfortunately, the integral is unsolvable, nevertheless, we can still characterize the properties of $f(t)$ in the two regimes of the GZipf distribution. In the high frequency region, where term having q as exponent dominates the denominator we can write:

$$\begin{aligned} f_q(t) &\sim \int_{\nu_k}^1 \frac{\nu^2 e^{-\nu t}}{\nu^q} d\nu \\ &= \frac{\Gamma(3 - q, \nu_k t)}{t^{3-q}} \end{aligned} \quad (4)$$

where, $\Gamma(n, z) = \int_z^\infty x^{n-1} e^{-x} dx$ is the incomplete Gamma function. $\nu_k = (\mu/(\lambda - \mu))^{1/(q-r)}$ is the frequency for which the two terms that make up the denominator are equal. It is useful to note that for low t values that correspond to high frequencies the nominator presents a constant plateau that quickly decreases, or bends, at the edge as $t \rightarrow 1/\nu_k$. Therefore, we can approximate:

$$f_q(t) \sim \frac{1}{t^{3-q}} \quad (5)$$

Similarly, it may be shown that for low frequencies, that is, in the region where term with r as exponent dominates:

$$f_r(t) \sim \frac{1}{t^{3-r}} \quad (6)$$

Finally, we conclude that the inter-reference distance distribution can be approximated by a piece-wise power-law. Our result is similar to the single sloped power-law obtained by Jin under the assumption of Zipf distributed popularity or the empirical observations by Breslau et. al in [13] for Web reference strings. However, due to its general form it should be able to capture the properties of more varied workloads. In the following section we use the inter-reference distance distribution together with the working-set theory to deduce the miss rate of an LRU cache.

3.3. A cache model

Denning proposed the use of the working-set as a tool to capture the set of pages a program must store (cache) in memory such that it may operate at a desired level of efficiency [12]. The idea is to estimate a program's locality, or in-use pages, with the help of a sliding window of variable length looking into the past of the reference string. In their seminal work characterizing the properties of the working-set [25], Denning and Schwartz showed that the average inter-reference distance is the slope of the average miss rate, which at its turn is the slope of the average working-set size, both taken as functions of the window size. The result is of particular interest as it provides a straightforward link between the properties of the reference string and the performance of a cache that uses the least recently used (LRU) eviction policy, but whose size is variable. To understand the latter, consider that the working-set evicts entries that have not been recently referenced and that the size of the working-set for a given window size depends on the variable number of unique destinations within the window. Given the constraint that the reference string is obtained using IRM, the working-set size is in fact normally distributed and has a low variance that we can approximate as being zero. Finally, because of the traffic's statistical properties, the cache modeled by the working-set becomes an LRU of fixed size.

We leverage in what follows the result above to deduce miss rate of an LRU cache when fed by a reference string obtained using IRM and a GZipf popularity distribution. The miss rate for the upper part of $f(t)$ is:

$$m_q(t) = - \int \frac{C}{t^{3-q}} dt = -C \frac{t^{q-2}}{q-2} \quad (7)$$

where, $t < 1/\nu_k$, $1 < q < 2$ and C is a normalizing constant which ensures that $\sum_{t=1}^{N-1} C f(t) = 1$. We can further compute the average working-set size as:

$$s_q(t) = \int C \frac{t^{q-2}}{q-2} dt = -C \frac{t^{q-1}}{(q-1)(q-2)} \quad (8)$$

To obtain the miss rate as a function of the cache size, not of the inter-reference distance, we take the inverse of s_q and replace it in (7). For $s < s_q(1/\nu_k)$ we get:

$$\begin{aligned} m_q(s) &= C \frac{1}{q-1} (2-q) \frac{1}{q-1} (q-1) \frac{q-2}{q-1} \frac{q-2}{s^{q-1}} \\ &\propto s^{1-\frac{1}{q-1}} \end{aligned} \quad (9)$$

This suggests that the asymptotic miss rate as a function of cache size is a power law of the cache size with an exponent dependent on the slope of the popularity distribution. Similarly, for large inter-reference distances, when $s > s_r(1/\nu_k)$:

$$m_r(s) \propto s^{1-\frac{1}{r-1}} \quad (10)$$

Then, for a reference string whose destinations have a GZipf popularity distribution and where the references to objects are independent, we find that the miss rate presents two power-law regimes with exponents only dependent on the

exponents of the popularity distribution and the cache size. We test the ability of the equations to fit empirical observations in Section 4.4.

3.4. Cache performance analysis

We now investigate how cache size varies with respect to the parameters of the model if the miss rate is held constant. By inverting (9) and (10) we obtain the cache size as a function of the miss rate:

$$s(m) = \begin{cases} g(q) m^{1-\frac{1}{2-q}}, & m \leq m_k \\ g(r) m^{1-\frac{1}{2-r}}, & m > m_k \end{cases} \quad (11)$$

$$\text{with } g(x) = -C^{\frac{1}{2-x}} \frac{(2-x)^{\frac{x-1}{x-2}}}{2-3x+x^2}, \quad m_k = \frac{C}{v_k^{r-2}(2-r)},$$

$$v_k = \left(\frac{\mu}{\lambda - \mu}\right)^{q-r} \text{ and } 0 < m < 1.$$

We see that $s(m)$ is independent of both the number of packets N and the number of destinations D and is sensible only to changes of the slopes of the popularity distribution q, r and the frequency at which the two slopes intersect, v_k . We do note that C does depend analytically on N as it can be seen by considering C 's defining expression (see discussion of (7)): $1/C = H(1/v_k, 3-q) - \zeta(3-r, N) + \zeta(3-r, 1/v_k)$ where $H(n, m) = \sum_{k=1}^n 1/k^m$ is the generalized

harmonic number of order n of m and $\zeta(s, a) = \sum_{k=0}^{\infty} 1/(k+a)^s$ the Hurwitz Zeta function. However, the first and last terms of the expression depend only on popularity parameters while the middle one quickly converges to a constant as N grows. Whereby it is safe to assume C constant with respect to N and consequently that the number of packets does not influence $s(m)$.

On the other hand, if the parameters of the popularity distribution are modified, some interesting dependencies can be uncovered. For brevity, we explore only the case when q and r vary but still respect the constraint that $1 < r < q < 2$. When both exponents jointly change, the cache size required to maintain the miss rate will qualitatively vary as depicted in Fig. 2. Specifically, as their value approach 1, that is, when the popularity distribution is strongly skewed, cache size asymptotically goes to a low value constant, whereas when the exponent approaches 2, the required cache size grows very fast, notice the superlinear growth in the log-log scale. Despite




Fig. 2. Cache size as a function of a GZipf exponent for a fixed miss rate.

not being indicated by (11), $s(m)$ is defined when q or r are 2, that is, it does not grow unbounded. The expression can be obtained if we replace q by 2 in (7) and recompute all equations:

$$s(m) = (C + m) e^{-\frac{m}{C}} \quad (12)$$

3.5. Discussion of asymptotic cache performance and impact

Using the results of the analysis performed in the previous section we are now interested to characterize the asymptotic scalability of the LISP cache size with respect to (i) the number of users in a LISP site (ii) the size of the EID space and (iii) the parameters of the popularity distribution. To simplify the discussion, we assume there are no interactions between the first two and the third:

Assumption 3. The destination prefix popularity distribution is independent of the number of users in a LISP site and the size of the EID space.

Whereby (i) contemplates the variation of the number of packets, N (ii) the variation of the number of destinations D and (iii) the variation of the GZipf parameters q, r, μ and λ , independently. We acknowledge that the popularity distribution may be influenced by a multitude of factors, and in particular by the growth of the users generating the reference string. Nonetheless, we argue that our assumption does make practical sense. For instance, a typical LISP router is expected to serve hundreds to thousands of clients so fluctuations proportional to the size of the user set should not affect overall homogeneity and popularity distribution. Additionally, although user interest in content quickly changes, the same is not necessarily true for the content sources, i.e., prefixes from where the content is served, which the user cannot typically select. This split between content and its location can result in relatively stable popularity distribution of the prefixes despite the dynamic popularity of actual content. We show an example network where this assumption holds in Section 4.2.

In the previous section we found that when the parameters of the popularity distribution are held constant, the cache size is independent of both the number of packets and destinations. As a result, cache size scales constantly, $O(1)$, with the number of users within a LISP site and the size of EID-prefix space for a fixed miss rate. This observation has several fundamental implications for LISP's deployment. First, caches for LISP networks can be designed and deployed for a desired performance level which subsequently does not degrade with the growth of the site and the growth of the Internet address space. Second, splitting traffic between multiple caches (i.e., routers) for operational purposes, within a large LISP site, does not affect cache performance. Finally, signaling, i.e., the number of map-request exchanges, grows linearly with the number of users if no hierarchies or cascades of caches are used. This because the number of resolution requests is $m(s)N$.

If the previous assumption does not hold, then, in the worst case, the cache size scales linearly, $O(N)$, with $|D|$. This follows if we consider that, as the growth of N and D flatten the distribution, thus leading to a uniform popularity, the cache size for a desired miss rate becomes proportional

Table 2
Datasets statistics.

	UPC 2009	UPC 2011	UPC 2012	CESCA 2013
Date	2009-05-26	2011-10-19	2012-11-21	2013-01-24
Packets	6.5B	4.05B	5.57B	20B
Average packets/s	75.3k	46.9k	64.4k	232k
Prefixes	92.8k	94.9k	109.4k	143.7k
Average prefixes/s	2.3k	1.95k	2.1k	2.56k

to the $|D|$. This then explains the very high miss rates observed in [11] for cache polluting attacks, since scanning attacks, whereby an attacker enumerates all possible destinations at a high packet rate, finally result in a flat popularity distribution.

4. Empirical evidence of temporal locality

In this section we verify the accuracy of our assumptions regarding the popularity distribution of destination prefixes and the sources of locality in network traffic. We also verify the accuracy of the predictions regarding the performance of the LISP cache empirically. But first, we present our datasets and experimental methodology.

4.1. Packet traces and cache emulator

We use four one-day packet traces that only consist of egress traffic for our experiments. Three were captured at the 2 Gbps link that connects our University's campus network to the Catalan Research Network (CESCA) and span a period of 3.5 years, from 2009 to 2012. The fourth was captured at the 10 Gbps link connecting CESCA to the Spanish academic network (RedIris) in 2013. UPC campus has about 36k users consisting generally of students, academic staff and auxiliary personnel while CESCA provides transit services for 89 institutions that include the public Catalan schools, hospitals and universities. The important properties of the datasets are summarized in Table 2. It should be noted that this traffic traces have also been used to evaluate the map-cache model proposed in [11]. For convenience, we reproduce here as well some of the most important statistics.

At the time of this writing there exists no policy as to how EID-prefixes are to be allocated. However, it is expected and also the practice today in the LISP-beta network to allocate EIDs in IP-prefix-like blocks. Consequently we performed our analysis considering EID-prefixes to be of BGP-prefix granularity. For each packet within a trace we find the associated prefix using BGP routing tables downloaded from the RouteViews archive [26] that match the trace's capture date. We filtered out the more specific prefixes from the routing tables as they are generally used for traffic engineering and LISP offers a more efficient management of these operational needs. Table 3 gives an overview of the original (BGP_{RT}), and filtered BGP_{ϕ} routing table sizes as well as the ratio (ρ) between the filtered routing table size and the number of prefixes observed within each trace. Both UPC and CESCA visit daily more than half of the prefixes within BGP_{ϕ} .

Apart from the popularity and temporal locality analysis we also implemented an LISP ITR emulator to estimate LRU

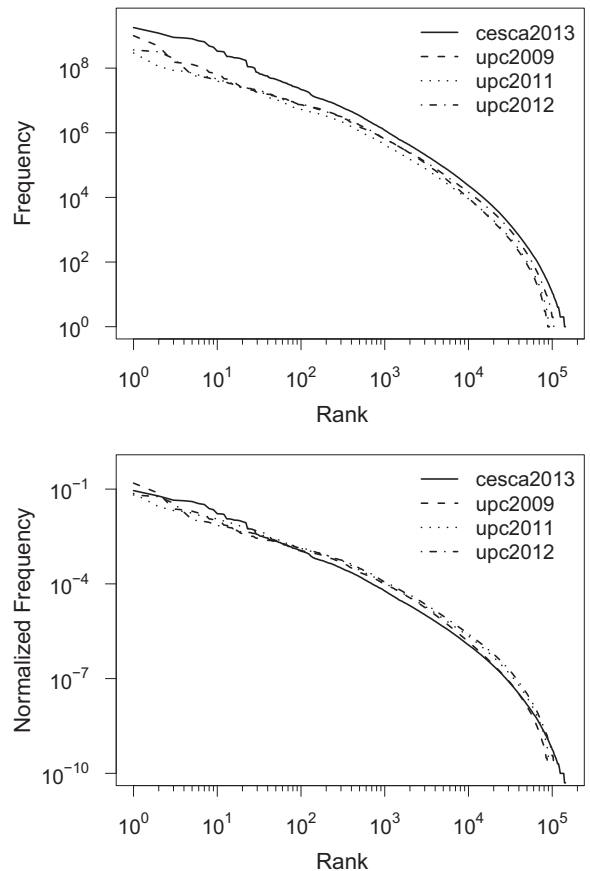
Table 3
Routing tables statistics.

	UPC 2009	UPC 2011	UPC 2012	CESCA 2013
BGP_{RT}	288k	400k	450k	455k
BGP_{ϕ}	142k	170k	213k	216k
ρ	0.65	0.55	0.51	0.66

map-cache performance using the traces and the routing tables as input. We compare the predictions of our cache model with the empirical results in Section 4.4.

4.2. Popularity distribution

Fig. 3 presents the frequency-rank distributions of our datasets for both absolute and normalized frequency. A few

**Fig. 3.** Destination prefix popularity distribution.

observations are in place. First, although clearly not accurately described by Zipf's law, they also slightly deviate from a GZipf. Namely, the head of the distribution presents two power-law regimes followed by a third that describes the tail as it can be seen in Fig. 3 (down). This may be either because a one day sample is not enough to obtain accurate statistics in the Zipf–Mandelbrot head region, or because popularity for low ranks follows a more complex law. Still, we find that for all traces the frequencies of higher ranks (above 2000) are accurately characterized by two power-law regimes (see Fig. 5).

Secondly, the frequency-rank curves for the UPC datasets are remarkably similar. Despite the 50% increase of BGP_ϕ (i.e., D), changes in the Internet content provider infrastructure over a 3.5 years period, and perhaps even changes in the local user set, the popularity distributions are roughly the same.

Finally, the normalized frequency plots for all traces are similar, in spite of the large difference in number of packets between CESCA and UPC datasets. These observations confirm our assumption that growth of the number of users within the site or of the destination space do not necessarily result in a change of the popularity distribution.

To confirm that these results are not due to a bias of popularity for larger prefixes sizes, that is, larger prefixes are more probable to receive larger volumes of traffic because they contain more hosts, we checked the correlation between prefix length and frequency. However, we did not find any evidence in support of this hypothesis (not shown here).

4.3. Prefix inter-reference distance distribution

We now check if knowledge about the popularity distribution suffices to accurately characterize the inter-reference distance distribution or if short-term correlations must also be taken into account. To achieve this, we use a methodology similar to the one used in [14] for Web page traffic. We first generate random versions of our traces according to the IRM model, i.e., by considering only the popularity distribution and geometric inter-reference times, and then compare the resulting inter-reference distance distributions to the originals. Results are shown in Fig. 4. We find that for all traces, popularity alone is able to account for the greater part of the inter-reference distance distribution, like in the case of Web requests. The only disagreement is in the region with distances lower than 100 where short-term correlations are important and IRM traces underestimate the probability by a significant margin.

A rather interesting finding is that the short-term correlations in all traces are such that the power-law behavior observed for higher distances ($t > 100$) is extended up to distance 1. In this region, the exact inter-reference distance equation (4) is a poor fit to reality as it follows the IRM curve. However, the empirical results are appropriately described by our approximate inter-reference model (5) which avoids IRM's bent by assuming equation (4)'s numerator constant.

4.4. Cache performance

Having found that our assumptions regarding network traffic properties hold in our datasets we now if the cache




Fig. 4. Empirical and IRM generated inter-reference distance distributions for the four traces.

model (see (9) and (10)) is able predict real world LRU cache performance.

As mentioned in Section 4.2 and as it may be seen in Fig. 5, the head of the popularity distribution exhibits two power-law regimes instead of one. Then, two options arise, we can either use the model disregarding the discrepancies or adapt it to consider the low rank region behavior. For completeness, we choose the latter in our evaluation. This only consists in approximating $p_\nu(\nu)$ (see (1)) as having three regions, each dominated by an exponent α_i . Recomputing (10) we get that the miss rate has three regions, each characterized by an α_i . Choosing the first option would only result in an overestimation of cache miss rates for low cache sizes.

To contrast the model with the empirical observations, we performed a linear least squares fit of the three regions of the popularity distribution. This allowed us to determine the exponents α_i , computed as $1 + 1/s_i$ where s_i is the slope of the i th segment, and to roughly approximate the frequencies ν_{k1} and ν_{k2} at which the segments intersect. Using them as input to (9) we get a cache miss rate estimate as shown in Fig. 7. Generally, we see that the model is a remarkably good fit for the large cache sizes but constantly underestimates the miss rate for sizes lower than 1000. This may be due to the poor fit of the popularity for low ranks. Nevertheless a more elaborate fitting of ν_{k1} and ν_{k2} should provide better results as it may be seen in Fig. 6 where we performed a linear least squares fit of the three power law regions of the cache miss rate. Knowing that the slope of the cache miss rate




Fig. 5. Frequency-rank distribution of destination prefixes and a linear least squares fit of the three power-law regimes. $\alpha_i = 1 + 1/s_i$, where s_i is the slope of the i th segment.

is $s_i = 1 - 1/(\alpha_i - 1)$ (see (7)), we computed the exponents as depicted in the figure. Comparison with those computed in Fig. 5 shows they are very similar. Overall, we can conclude that the model accurately predicts cache performance.

5. Related work

Denning was first to recognize the phenomenon of temporal locality in his definition of the working-set [12] and together with Schwartz established the fundamental properties that characterize it [25]. Although initially designed for the analysis of page caching in operating systems, the ideas were later reused in other fields including Web page and route caching.

In [13] Breslau et al. argued that empirical evidence indicates that Web requests popularity distribution is Zipf-like of exponent $\alpha < 1$. Using this finding and the assumption that temporal locality is mainly induced through long-term popularity, they showed that the asymptotic miss rates of an LRU cache, as a function of the cache size, is a power law of exponent $1 - \alpha$. In this paper we argue that GZipf with exponents greater than 1 is a closer fit to real popularity distributions and obtain a more general LRU cache model. We further use the model to determine the scaling properties of the cache.

Jin and Bestavros showed in [14] that the inter-reference distribution is mainly determined by the long-term popularity and only marginally by short-term correlations. They

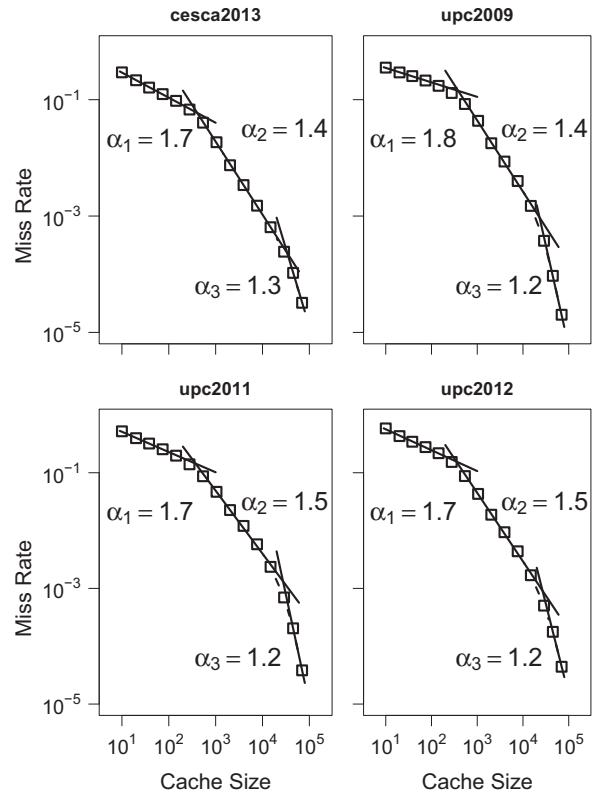


Fig. 6. Empirical miss rate with cache size and a linear least-squares fit of the exponent for the three power-law regions. Notice the similarity with the exponents of the three regions of the popularity distribution in Fig. 5.

also proved that the inter-reference distribution of a reference string with Zipf-like popularity distribution is proportional to $1/t^{2-1/\alpha}$. We build upon their work but also extend their results by both considering a GZipf popularity distribution and by using them to deduce an LRU cache model.

In the field of route caching, Feldmeier [27] and Jain [28] were among the first to evaluate the possibility of performing destination address caching by leveraging the locality of traffic in network environments. Feldmeier found that locality could be exploited to reduce routing table lookup times on a gateway router while Jain, discovered that deterministic protocol behavior limits the benefits of locality for small caches. The works, though fundamental, bear no practical relevance today as they were carried two decades ago, a time when the Internet was still in its infancy.

Recently, Kim et al. [7] performed a measurement study within the operational confinement of an ISP's network and showed the feasibility of route caching. They show by means of an experimental evaluation that LRU cache eviction policy performs close to optimal and better than LRU. Also, they found that prefix popularity distribution is very skewed and that working-set size is generally stable with time. These are in line with our empirical findings and provide practical confirmation for our assumption that the popularity distribution can be described as a GZipf.

Several works [5,29] have previously looked at cache performance in loc/id split scenarios considering LISP as a

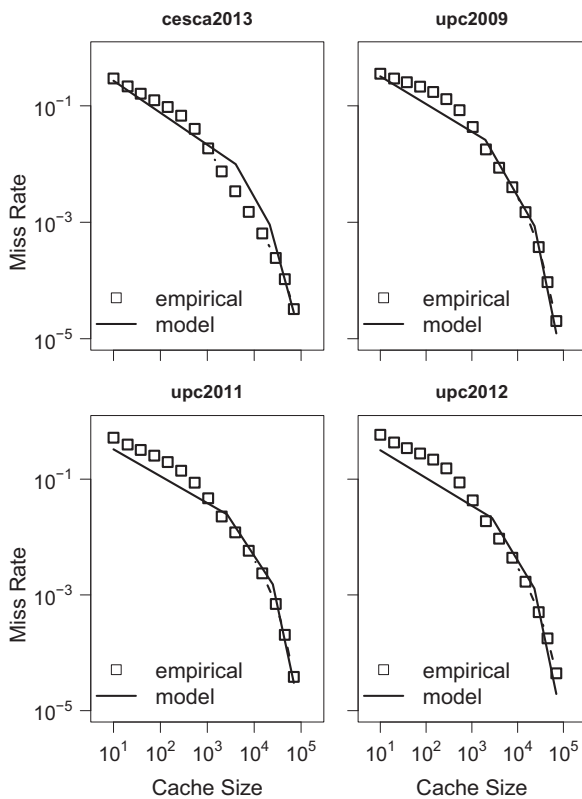


Fig. 7. Empirical miss rate with cache size together with a fit by (9) and (10).

reference implementation (for a complete set of references see [11] and the references therein). Although methodologies differ between the various papers, in all cases the observed LISP cache miss rates were found to be relatively small. This, again, indirectly confirms the skewness of the popularity distribution and its stability at least for short time scales.

Finally, in [10] and [11] we show how the working-set can be leveraged to estimate network traffic temporal locality and thereafter how it can be leveraged to link cache size and miss rate. We further exploit the result to study the vulnerability to scanning attacks and to perform an in-depth, over time analysis of cache performance. In contrast, here we seek to understand what fundamental traffic parameters influence the average working-set size, instead of empirically measuring it. Ultimately our goal is to find the link between popularity distribution, size of LISP site and EID namespace, on the one hand, and map-cache miss rate, on the other.

6. Conclusions

LISP offers a viable solution to scaling the core routing infrastructure of the Internet by means of a location/identity split. However this forces edge domain routers to cache location to identity bindings for timely operations. In this paper we answer the following question: does the newly introduced LISP edge cache scale?

Our findings show that the miss rate scales constantly, $O(1)$, with the number of users as well as with the number of destinations. The result is crucial for operational

environments and should remove all doubts about LISP's scalability, since it implies map-cache provisioned miss rate should not degrade over time. To reach this conclusion, we start from two assumptions: (i) the popularity of destination prefixes is described by a GZipf distribution and (ii) temporal locality is predominantly determined by long-term popularity. Fundamentally, these assumptions are often observed to hold in the Internet [7,24] but also in other fields such as web traffic [13], on-demand video [22] or even linguistics [20]. Arguably, they are inherent to human nature and, as such, are expected to hold in the foreseeable future. In addition, we show that the result is valid only if, as observed for our traffic traces, (iii) popularity is not influenced by the number of users and the number of destinations. Nevertheless, we also show that if the assumptions do not hold, then cache size scales linearly, $O(N)$, with the number of destinations.

At the time of this writing there is an open debate on how the Internet should look like in the near future and in this context, it is important to analyze the scalability of the various future Internet architecture proposals. This paper fills this gap, particularly for the Locator/ID split architecture. Furthermore, our results show that edge networks willing to deploy LISP will not face map-cache size scaling issues—as long as the assumptions hold—even if the edge network itself becomes larger (i.e., more users) or the Internet grows (i.e., more prefixes).

Acknowledgements

The authors would like to express their gratitude to Damien Saucez and Chadi Barakat for their insightful comments. This work has been partially supported by the Spanish Ministry of Education under scholarship AP2009-3790, research project TEC2011-29700-C02, Catalan Government under project 2009SGR-1140 and a Cisco URP Grant.

References

- [1] D. Meyer, L. Zhang, K. Fall, Report From the IAB Workshop on Routing and Addressing, 2007, Internet Engineering Task Force (IETF) (RFC 4984 (Informational)).
- [2] T. Li, Recommendation for a Routing Architecture, 2011, Internet Engineering Task Force (IETF) (RFC 6115 (Informational)).
- [3] D. Saucez, L. Iannone, O. Bonaventure, D. Farinacci, Designing a deployable Internet: The locator/identifier separation protocol, *IEEE Internet Comput.* 16 (2012) 14–21.
- [4] L. Jakab, A. Cabellos-Aparicio, F. Coras, J. Domingo-Pascual, D. Lewis, Locator/Identifier Separation Protocol (LISP) Network Element Deployment Considerations, 2014, Internet Engineering Task Force (IETF) (RFC 7215 (Experimental)).
- [5] L. Iannone, O. Bonaventure, On the cost of caching locator/ID mappings, in: *Proceedings of the 3rd International Conference on Emerging Networking EXperiments and Technologies (CoNEXT'07)*, ACM, 2007, pp. 1–12.
- [6] J. Kim, L. Iannone, A. Feldmann, Caching locator/id mappings: An experimental scalability analysis and its implications, *Comput. Netw.* 57 (4) (2013) 897–909.
- [7] C. Kim, M. Caesar, A. Gerber, J. Rexford, Revisiting route caching: The World should be flat, in: *Proceedings of the 10th International Conference on Passive and Active Network Measurement (PAM'09)*, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 3–12.
- [8] L. Jakab, A. Cabellos-Aparicio, F. Coras, D. Saucez, O. Bonaventure, LISP-TREE: A DNS hierarchy to support the LISP mapping system, *IEEE J. Select. Area Commun.* 28 (8) (2010) 1332–1343.
- [9] H. Zhang, M. Chen, Y. Zhu, Evaluating the performance on ID/Loc mapping, in: *Global Telecommunications Conference (GLOBECOM'08)*, IEEE, 2008, pp. 1–5.

- [10] F. Coras, A. Cabellos-Aparicio, J. Domingo-Pascual, An analytical model for the LISP cache size, in: NETWORKING 2012, Springer, 2012, pp. 409–420.
- [11] F. Coras, J. Domingo-Pascual, D. Lewis, A. Cabellos-Aparicio, An analytical model for Loc/ID mappings caches, IEEE/ACM Trans. Netw. (2014) <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6975256&queryText=Coras,%20F&newsearch=true>.
- [12] P.J. Denning, The working set model for program behavior, Commun. ACM 11 (5) (1968) 323–333.
- [13] L. Breslau, P. Cao, L. Fan, G. Phillips, S. Shenker, Web caching and zipf-like distributions: evidence and implications, in: IEEE Proceedings INFOCOM, vol. 1, IEEE, 1999, pp. 126–134.
- [14] S. Jin, A. Bestavros, Sources and characteristics of web temporal locality, in: Proceedings of the International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, IEEE, 2000, pp. 28–35.
- [15] M.A. Montemurro, Beyond the zipf–mandelbrot law in quantitative linguistics, Phys. A: Stat. Mech. Appl. 300 (3) (2001) 567–578.
- [16] D. Farinacci, V. Fuller, D. Meyer, D. Lewis, The Locator/ID Separation Protocol (LISP), 2013, Internet Engineering Task Force (IETF) (RFC 6830 (Experimental)).
- [17] LISP Testbed, n.d. <http://www.lisp4.net/>
- [18] R. Hinden, New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG, 1996, Internet Engineering Task Force (IETF) (RFC 1955 (Informational)).
- [19] V. Fuller, D. Farinacci, D. Lewis, LISP Delegated Database Tree (LISP-DDT), 2011, Internet Engineering Task Force (IETF) (draft-ietf-lisp-ddt-00). Work in progress.
- [20] G.K. Zipf, Human behavior and the principle of least effort, 1949, Addison-Wesley Press.
- [21] A. Mahanti, C. Williamson, D. Eager, Traffic analysis of a web proxy caching hierarchy, IEEE Netw. 14 (3) (2000) 16–23.
- [22] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, S. Moon, I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system, in: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, ACM, 2007, pp. 1–14.
- [23] G. Dán, N. Carlsson, Power-law revisited: Large scale measurement study of p2p content popularity., in: IPTPS, USENIX, 2010, p. 12.
- [24] N. Sarrar, S. Uhlig, A. Feldmann, R. Sherwood, X. Huang, Leveraging zipf's law for traffic offloading, ACM SIGCOMM Comput. Commun. Rev. 42 (1) (2012) 16–22.
- [25] P.J. Denning, S.C. Schwartz, Properties of the working-set model, Commun. ACM 15 (3) (1972) 191–198.
- [26] RouteViews Project, University of Oregon, n.d. <http://www.routeviews.org>
- [27] D. Feldmeier, Improving gateway performance with a routing-table cache, in: Proceedings of the INFOCOM'88. Networks: Evolution or Revolution. Seventh Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE, 1988, pp. 298–307.
- [28] R. Jain, Characteristics of destination address locality in computer networks: A comparison of caching schemes, Comput. Netw. ISDN Syst. 18 (4) (1990) 243–254.
- [29] J. Kim, L. Iannone, A. Feldmann, A deep dive into the LISP cache and what ISPs should know about it, in: Proceedings of the 10th International IFIP TC 6 Conference on Networking: Vol. Part I (NETWORKING'11), Springer-Verlag, Berlin, Heidelberg, 2011, pp. 367–378.



Florin Coras received the diploma engineer degree from the Technical University of Cluj-Napoca, Romania in telecommunication engineering in 2009 and the M.Sc. degree in computer networks from the Technical University of Catalonia, Spain in 2011. Since October 2009, he has been a Ph.D. candidate at the Technical University of Catalonia. His research interests lie in the fields of network routing and new Internet architectures with a focus on Locator/ID Separation Protocol.



Jordi Domingo-Pascual received the engineering degree in telecommunication and the Ph.D. degree in computer science. He is a full professor at the Computer Architecture Department, Universitat Politècnica de Catalunya (UPC), Barcelona (BarcelonaTECH). Since 1983 he has been a lecturer at the Computer Architecture Department and currently he is teaching at Computer Science School (FIB) and at the Telecommunications Engineering School (ETSETB). He is the head of the Broadband Communications and Services research group (CBA) of UPC that is a qualified research group by the Catalan Government. He is one of the promoters and founders of the inter-departmental research center for Advanced Broadband Communications (CCABA) that is recognized as distinguished research center by the UPC. He was director of the Computer Architecture Department of the UPC (2005–2011). He has been advisor of 11 Ph.D. thesis and has published more than 130 papers in international journals and conferences.



Albert Cabellos-Aparicio received the B.Sc., M.Sc., and Ph.D. degrees in Computer Science Engineering from the Technical University of Catalonia, in 2001, 2005, and 2008, respectively. He is also assistant professor of the Computer Architecture Department and researcher of the Broadband Communications Group since 2005. In 2010 he joined the NaNoNetworking Center in Catalunya (<http://www.n3cat.upc.edu>) where he is the scientific director. He is an editor of the Elsevier Journal on Nano Computer Network and founder of the ACM NANOCOM conference, the IEEE MONACOM workshop and the N3Summit. He has also founded the LISPmob open-source initiative (<http://lispmob.org>) along with Cisco. He has been a visiting researcher at Cisco Systems and Agilent Technologies and a visiting professor at the Royal Institute of Technology (KTH) and the Massachusetts Institute of Technology (MIT). He has given more than 10 invited talks (MIT, Cisco, INTEL, MIET, Northeastern University, etc.) and co-authored more than 15 journal and 40 conference papers. His main research interests are future architectures for the Internet and nano-scale communications.