

An Analytical Model for Loc/ID Mappings Caches

Florin Coras, Jordi Domingo-Pascual, Darrel Lewis, and Albert Cabellos-Aparicio

Abstract—Concerns regarding the scalability of the interdomain routing have encouraged researchers to start elaborating a more robust Internet architecture. While consensus on the exact form of the solution is yet to be found, the need for a semantic decoupling of a node's location and identity is generally accepted as a promising way forward. However, this typically requires the use of caches that store temporal bindings between the two namespaces, to avoid hampering router packet forwarding speeds. In this article, we propose a methodology for an analytical analysis of cache performance that relies on the working-set theory. We first identify the conditions that network traffic must comply with for the theory to be applicable and then develop a model that predicts average cache miss rates relying on easily measurable traffic parameters. We validate the result by emulation, using real packet traces collected at the egress points of a campus and an academic network. To prove its versatility, we extend the model to consider cache polluting user traffic and observe that simple, low intensity attacks drastically reduce performance, whereby manufacturers should either overprovision router memory or implement more complex cache eviction policies.

Index Terms—Cache modeling, location/identity separation, Locator/ID Separation Protocol (LISP), network traffic locality, working-set model.

I. INTRODUCTION

THE GROWING consensus within the network community today is that a semantic separation of the *identity* and *location* information, currently overloaded in IP addresses, would solve the limitations of the routing infrastructure [1], [2]. Such architectural update involves network infrastructure upgrades in few selected network points like border routers, where the linking of the two namespaces must be performed. However, more importantly, it also entails the more challenging implementation of a *mapping system* [3] able to distribute dynamic location-identity bindings (*mappings*) to globally distributed routers.

In this context, routers generally retrieve mappings on user demand, as opposed to proactively fetching them. This is done

such that the amount of memory a router requires to participate in the system does not grow with identifier space, as is the case today, but is instead dependent on the packet-level traffic the router processes. As a result, to diminish retrieval times, increase packet forwarding speed, and protect the mapping system from floods of resolution requests, routers are provisioned with mappings caches (*map-caches*) that temporarily store in use bindings.

Although caches placed between processor and main memory in operating systems or in Web proxies are well studied [4]–[6], route and mappings caches have yet to be thoroughly analyzed. A considerable number of experiments have empirically evaluated map-cache performance, however they are mainly focused on providing a circumstantial description of cache behavior, that is, for particular cache configurations and network traffic traces, as opposed to a general one [7]–[11]. Typically, these results yield accurate estimates for cache performance but unfortunately cannot be extrapolated to provide rough projections or bounds on cache behavior for workloads with different characteristics, nor can they provide insight into what traffic properties influence cache performance and to what degree. Answering such questions would not only be a first important step toward understanding the overall performance of the mapping system, but would also provide a quick way of gauging the expected map-cache performance of any network domain.

In this paper, we present an analytical model that, to the best of our knowledge, constitutes the first theoretical framework for map-cache performance analysis and provisioning. The model relies on coarse traffic parameters and aims to be applicable to a wide range of scenarios. In particular, we first show how the working-set theory [12] may be used to estimate simple parameters that characterize the intrinsic locality of network traffic and thereafter explain how they can be leveraged to link cache size and miss rate. The key assumption that enables the analysis is that the stochastic mechanism underlying the generation of traffic is stationary, i.e., the distributions characterizing packet arrivals (such as interreference distance distribution), is time-translation invariant. We find stationarity to hold for real network traffic, and to facilitate the use of the model, we propose a simple methodology that tests for it in network traces. Finally, we validate the result by emulation, using packet traces collected at the edges of a campus and an academic network. Part of these results have been previously presented in [13].

Another contribution of this paper is that we exploit the model to: 1) perform an in-depth, over-time analysis of cache performance for our datasets; and 2) study the security of the map-cache by evaluating its vulnerability to scanning attacks. Our findings show that miss rate decreases at an accelerated pace with cache size and eventually settles to a power-law decrease. As a result, even for a relative small size, 10% of the Internet's

Manuscript received November 14, 2013; revised July 18, 2014; accepted October 21, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Sen. This work was supported in part by the Spanish Ministry of Education under Scholarship AP2009-3790, Research Project TEC2011-29700-C02, the Catalan Government under Projects 2014SGR-1427 and 2009SGR-1140, and a Cisco URP Grant.

F. Coras, J. Domingo-Pascual, and A. Cabellos-Aparicio are with the DAC, Universitat Politècnica de Catalunya—BarcelonaTech (UPC), 08034 Barcelona, Spain (e-mail: fcoras@ac.upc.edu; jordi.domingo@ac.upc.edu; acabello@ac.upc.edu).

D. Lewis is with Cisco Systems, San Jose, CA 95134 USA (e-mail: darlewis@cisco.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2014.2373398

aggregated routing table, the performance is still acceptable, below 0.2% packet miss rate. On the other hand, we also find that even simple and low-intensity attacks may compromise performance, and that increasing the cache size does not alleviate the problem, while the map-cache cannot accommodate the whole destination address space. Thereby, to achieve acceptable forwarding speeds, manufacturers should either over-provision router memory or design efficient cache management algorithms.

For the sake of clarity, we focus our analysis on Locator/ID Separation Protocol (LISP) map-cache performance (see Section II). Nevertheless, the results are relevant for other architectures inspired by the location/identity separation paradigm, including those like ILNP [14] that use DNS as their mapping system, since the equations could be used to approximate DNS resolver caching performance. Moreover, the cache models could be applied to route-caching and scalability techniques that focus on shrinking routing tables to extend router lifetimes [15].

The rest of the paper is structured as follows. Section II provides an overview of the location/identity split paradigm and one of its most successful implementations, LISP. In Section III, we first introduce the cache model problem and then present our analytical model and its extension that accommodates for cache pollution attacks. Section IV describes the methodology used and the results that validate the two models. Section V discusses the predictions of our results when applied to our datasets and some design guidelines to improve cache performance. Section VI presents the related work. Finally, Section VII concludes the paper.

II. LOCATION/IDENTITY SPLIT BACKGROUND

Prompted by the growing BGP routing tables and associated churn, researchers in both academia and operational community have spent recent years (arguably the last couple of decades) trying to improve the scalability of the Internet's architecture. The factors supporting the routing table's growth are partly organic in nature, as new domains are continuously added to the topology, however they are also closely related to current operational practices [2]. In this sense, multihoming, traffic engineering, and allocations of provider-independent prefixes are the main drivers behind the prefix deaggregation that drives the increase of the routing tables. On the one hand, this effect is detrimental because it leads to engineering limitations, as the growth rate of the routing table surpasses that of its supporting technology, but it is also harmful to network operators' business since the added capital expenses result in less cost-effective networks [2].

It has been long hinted, and today it is generally accepted, that the origin of this architectural inflexibility can be tracked down to the semantic overloading of the IP addresses with both *location* and *identity* information [16]. Consequently, their separation, typically referred to as a *loc/id split*, has been proposed by many solutions aiming to mitigate the routing problems [17], but also by those aiming to integrate new protocol mechanisms [18], currently not supported by the routing system. Apart from infrastructure upgrades, these architectures also require the deployment of mapping systems for linking the two new namespaces




Fig. 1. Example packet exchange between EID_{SRC} and EID_{DST} with LISP. Following intradomain routing, packets reach TR_A , which, not having a prior map-cache entry, obtains a mapping binding EID_{DST} to $RLOC_B$ from the mapping system (steps 1–3). TR_A stores the mapping in the map-cache (step 4) and then encapsulates and forwards the packets to $RLOC_B$ over the Internet's core (step 5). TR_B decapsulates the packets and forwards them to their intended destination.

and often rely on caches to limit router memory requirements and improve packet forwarding speeds.

Given that the location/identity separation paradigm only outlines a set of principles, for completeness, in what follows we use the LISP [19], [20] to illustrate the operation of a fully fledged architecture. LISP is one of the most successful loc/id split implementations to date and counts with the support of a sizable community [21] dedicated to its development.

Prior to forwarding a host generated packet (see Fig. 1), potentially part of a flow, a LISP router maps the destination address, an end-host identifier (EID), to a corresponding destination routing locator (RLOC). This is achieved by first looking up the destination in the local mappings cache, the map-cache, and if this fails, by requesting the mapping to a mapping system [10], [22], [23]. An overview of recent mapping system proposals can be found in [3]. Thereby, instead of proactively storing bindings for all the possible identifiers in the Internet, LISP routers request and store them in accordance to the packet-level traffic they process. This ensures map-cache size is independent of the identity namespace size and only dependent on local user needs. Once the mapping is obtained, the border router encapsulates and tunnels [24] the packets in the flow from source edge to corresponding destination edge tunnel router (TR), where they are decapsulated and forwarded to their intended destination.

Stale map-cache entries are avoided with the help of time-outs, called *time to live* (TTL), that mappings carry as attributes, while consistency is ensured by proactive LISP mechanisms that allow the owner of an updated mapping to inform its peers of the change.

The map-cache is most efficient in situations when destination EIDs present high temporal locality, and its size depends on the visited destinations set size. As a result, performance depends entirely on map-cache provisioned size, traffic characteristics, and the eviction policy set in place. We dedicate the remainder of the paper to finding and evaluating the relation

between parameters like cache size, performance, and locality of network traffic.

III. CACHE PERFORMANCE

This section presents the theoretical background and analytical methodology used to model map-cache performance. After briefly introducing the working-set theory, we formalize the cache modeling problem and show that the working-set is suitable for the analysis of real network traces. The results enable us to derive a cache model and, by extension, one that accounts for cache polluting attacks.

A. Problem Definition

For operating systems, a general resource allocation treatment was possible after it was observed that programs often obey the so-called *principle of locality*. The property arises from the empirical observation that programs favor only a subset of their information at a given time, so they may be efficiently run only with a fraction of their total data and instruction code. It was shown that if the subset, called the program's *locality*, can be computed, the policy of keeping the locality in memory is an optimal or near-optimal memory management policy [25].

Based on previous results [7]–[9], we argue that prefix-level network traffic roughly obeys the same principle because of, among others, skewed destination popularity distributions and flow burstiness in time, as they give rise to temporal locality, but also due to aggregation, i.e., multiplexing of a large number of user flows, which leads to a form of geographical locality. We therefore evaluate the feasibility of in-router *prefix/mappings caching* by analyzing the locality of network traffic. Since we want to avoid any assumptions regarding the structure of the process generating the network traffic, we opt in our evaluation for the working-set model of locality. For a list of other locality models, see [25].

Next, we provide a summary of the working-set terminology. For brevity, we use the term *unit of reference*, or simply *unit*, as a substitute for the referenced object (e.g., prefixes); and *reference set* to represent the set of all referenced units. Then, considering a *reference set* N , we define a *reference string* as the sequence $\rho = r_1 r_2 \dots r_i \dots$ where each unit $r_i \in N$. If t is a measure of time in *units*, then we can state the following.

Definition 1: Given a reference string, the working-set $W(t, u)$ is the set of distinct *units* that have been referenced among the u most recent references, or in the interval $[t - u + 1, t]$.

A graphical depiction can be found in Fig. 2. In accordance with [26], we refer to u as the *window size* and denote the number of distinct units in $W(t, u)$, the *working-set size*, as $w(t, u)$. The *average working-set size*, $s(u)$, measures the growth of the working-set with respect to the size of the window u , extending in the past, but independent of absolute time t , and is defined as

$$s(u) = \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=1}^k w(t, u). \quad (1)$$

It can be proved that the *miss rate*, $m(u)$, which measures the number of units of reference entering or returning to the




Fig. 2. Working-set model. At time t , $W(t, u)$ is the set of units recently referenced in the time window u , i.e., during the interval $[t - u + 1, t]$.

working-set per time unit, is the derivative of the previous function. This can be intuitively explained by the fact that the slope, or the growth speed, of $s(u)$ represents on average the number of unique references not observed among the last u references and therefore those that generate a working-set miss. Similarly, the sign inverted slope of the miss rate function, the second slope of $s(u)$, represents the *average interreference distance density function*, $f(u)$. For a broader scope discussion of these properties and complete proofs, the interested reader is referred to [26].

It is important to note that $s(u)$ and $m(u)$, if computable, provide estimates on the minimum average size of a cache able to hold the working-set, i.e., the prefixes in the active locality, and its corresponding miss rate with respect to the number of references processed. Our goal in this paper is to determine if the miss rate as a function of cache size exists for real network traffic and if it can be modeled as a simple function without a considerable loss of precision.

B. Network Traffic Locality

As explained by Denning in [26], a working-set analysis of reference strings may be performed only if based on three constraints that provide for a more rigorous definition for *locality of reference*.

- 1) Reference strings are unending.
- 2) The stochastic mechanism underlying the generation of a reference string is stationary, i.e., independent of absolute time origin.
- 3) For $t > 0$, r_t and r_{t+x} become uncorrelated as $x \rightarrow \infty$.

The first, though obviously not fulfillable, introduces an insignificant error because the reference strings generated by practical programs or traces are long from a statistical standpoint. The third requires that references become uncorrelated as the distance goes to infinity and can usually be asserted as being true in practice. The most restrictive of the three is the second, which limits the analysis to a locality where all three constraints, including stationarity, must be satisfied.

Given that in practice network traffic reference strings may consist of sequences of localities, we are foremost interested in establishing the possible limitations posed by the second constraint on real network traffic analysis. It should be observed that in such scenarios, derived working-set results are only valid within their respective localities and are not to be extrapolated for the whole trace. To test for this possibility and to identify the reference string segments having different generating processes, we devise the simple experimental methodology that follows.

First, since computing $w(t, u)$ for all acceptable integer combination of t and u is intractable, to enable our analysis, we define the *working-set curve* to be $w(t, u)$ as a function of u but

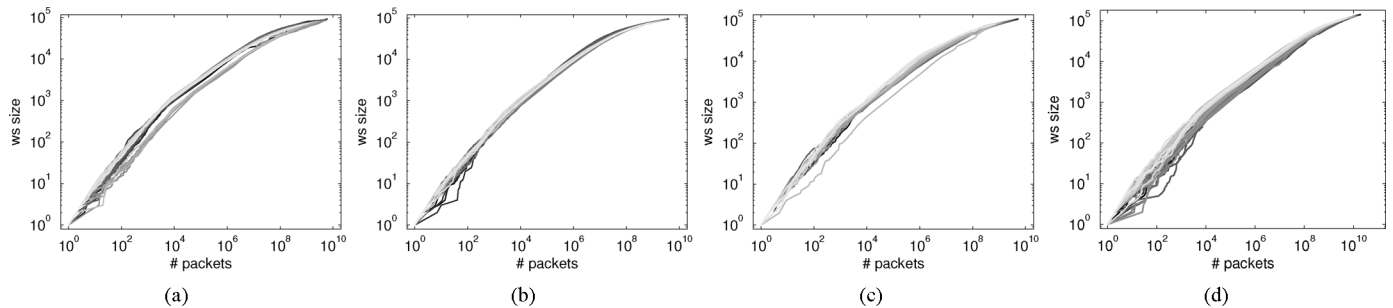


Fig. 3. Empirical working-set curves of four network traces captured at the egress points of two large academic networks over the course of 3 years. Starting times are spaced by 30 min and are evenly distributed over a day. Notice the clustering of the curves that indicates the stationarity of their generating process. We describe the datasets in Section IV-A, but provide their names (location and capture date) for future reference. (a) *upc* 2009. (b) *upc* 2011. (c) *upc* 2012. (d) *cesca* 2013.

additionally constrained to a fixed past starting point, i.e., constrained to $t - u = k$, with k a constant. For instance, considering Fig. 2, $w(t, u)$ and $w(t+1, u+1)$ are consecutive points on a working-set curve with points of type $w(t+n, u+n)$, $n > 0$ and with start time $t - u + 1$. Second, for a given trace, we compute multiple empirical destination prefix working-set curves with start times spanning one day and spaced by a fixed interval. Intuitively, one would expect the clustering patterns of the curves to distinguish between the multitude of generating processes. That is, curves with close start times should have a similar growth shape (cluster) because they follow a similar sequence of localities, whereas those separated by larger time lags should behave differently. More formally stated, we have the following proposition.

Proposition 1: The clustering of the working-set curves, under which $\forall u, w(t, u)$ is normally distributed, is equivalent to the stationarity of the process generating the reference string.

We provide a proof for the proposition in the Appendix. In addition, we leverage this result to determine the stationary processes embedded in four real network traffic traces. Details regarding the traffic captures can be found in Section IV-A.

For each network trace, we computed working-set curves spaced by half an hour. Fig. 3 presents the working-set curves of the four datasets in log-log plots, where it can be observed that all traces exhibit a strong clustering and sublinear growth due to temporal locality. Although the number of samples does not allow for an accurate enough testing, we could also confirm that $\forall u, w(t, u)$ is close to normally distributed by manual inspection. Therefore, in light of the previous proposition, we conclude that each trace can be considered as generated by a single stationary process.

To validate this finding, we also independently tested the stationarity of the processes generating the reference strings by applying the augmented Dickey–Fuller unit root test [27] to the interreference distance time series. Due to the large size of the datasets, we first aggregated the series by computing the mean for each window of 10 k data points and subsequently applied the test to the resulting mean series. For all traces, the null hypothesis that the series has a unit root, with the alternative that the series is stationary, was rejected ($p < 0.01$).

Undoubtedly, the most surprising result of the analysis is the stationarity of the processes generating the four traces. It follows

that the average working-set size and the other metrics derived from it characterize a trace in its entirety. This might seem to be somewhat counterintuitive if one considers the nonstationarity of network traffic when analyzed at short timescales. However, in large aggregate traces, we believe stationarity to be the result of flow multiplexing, whereby, as shown in the case of Web traffic in [28], the nonstationarities due to short-term correlations are canceled out by the influence of destination popularity, typically stable over longer time spans.

Similarly, Kim *et al.* observed in [9] that the working-set size for prefixes tends to be highly stable with time for traffic pertaining to a large ISP. We would like to stress that we do not assume all Internet traffic is generated by stationary processes, and thereby possesses an approximate time translation invariance of the working-set curves like the one observed in Fig. 3. In fact, we require that the model we develop further be applied only to traces that have this property.

C. Analytical Cache Model

Because the exact form of the average working-set size, as defined in (1), is rather cumbersome to work with, a shorter but approximate representation is desirable. With hindsight, one can recognize that, besides clustering, the empirical working-set curves are also piecewise linear when depicted in log-log scale (see Fig. 3). By exploiting this property, we can approximate the average working-set size, $s(u)$, for each trace with respect to the number of packets u by means of a piecewise linear fit of the log-log scale plot. As a result, we obtain estimates of both the slope, α , and the y -intercept, β , for all segments. In our analysis, we limit the number of segments to just four, however more segments may be used if an improved accuracy is desirable. Through conversion to linear scale, the average working-set-size equation becomes piecewise power law of the type

$$s(u) = e^{\beta(u)} u^{\alpha(u)} \quad (2)$$

where $\alpha(u)$, $0 < \alpha(u) \leq 1$, and $\beta(u) \geq 0$ are piecewise constant, decreasing, and respectively increasing functions obtained through fitting. Reusing the working-set properties presented in Section III-B, we can estimate the miss rate for a trace by computing the derivative of $s(u)$ like

$$m(u) = e^{\beta(u)} \alpha(u) u^{\alpha(u)-1}. \quad (3)$$

Next, by taking the inverse of (2) and replacing it in (3), we obtain an analytical relation that links cache size and the estimated miss rate

$$m(c) = e^{\beta^*(c)/\alpha^*(c)} \alpha^*(c) c^{1-1/\alpha^*(c)} \quad (4)$$

where c represents the cache size in number of entries and $\alpha^*(c)$ and $\beta^*(c)$ are piecewise constant functions with knees dependent on c . This equation accurately predicts cache performance over long time spans if $s(u)$ is relatively stable in the considered time frame (we further discuss this in Section V).

Interestingly, we find that miss rate only depends on cache size and the trace's temporal locality, and it is otherwise independent of instantaneous traffic parameters (e.g., number of packets/s). Moreover, we find that the pair $(\alpha(u), \beta(u))$ provides a very compressed characterization of temporal locality with respect to time, measured in number of packets, that is also easily computable, so it can provide quick cache performance gauging for specific network traffic loads.

Regarding the type of cache eviction policy modeled, it is useful to note that the working-set, $W(t, u)$, generally models a cache that always contains the $w(t, u)$ most recently referenced units. However, given that stationarity implies that $\forall u, w(t, u) \sim N(\sigma, \mu^2)$, as explained in [26], $W(t, u)$ actually models a cache of normally distributed size, dependent on u but independent of t . In particular, when the variance σ is small, or goes to zero, it behaves as if the size were fixed. Consequently, because the implicit eviction policy requires entries not referenced in a window of length u to be discarded, for low σ values, as those observed in our traces, the working-set closely approximates a cache of fixed size with a Least Recently Used (LRU) eviction policy. We validate the model and evaluate the efficiency of LRU in Section IV-C.

D. Cache Model for Cache Pollution Attacks

The model presented in Section III-C can be extended to account for situations when intradomain users perform EID space scanning attacks that significantly alter the working-set curves and therefore the shared map-cache's performance. In this work, we focus on assessing the damage users can inflict through data-plane attacks and do not consider control-plane attacks like those described in [29].

We define a scanning attack as the situation when one or multiple users, acting jointly, send packets over a large period of time (e.g., hours) to destinations having a high probability of not being found in the cache. The goal would be to either generate cache misses, resulting in control plane overload, or, if the cache is not large enough, to generate cache evictions, which would affect ongoing flows. For instance, having a list of EID prefixes, an attack would consist in sending packets with destinations enumerating all prefixes in the set in a random order, at a certain packet rate. Once all destinations are exhausted, the enumeration would start over.

In what follows, we formally define the parameters of the attack (see Table I for a quick overview). Let Ω be the EID-prefix set used in the attack, and Ψ the network's visited EID-prefix set. We define the relative attack intensity ρ as the ratio between the attack packet rate and the legitimate traffic packet

TABLE I
SCANNING ATTACK PARAMETERS

Parameter	Description
Ω	EID-prefix set employed by an attack
Ψ	A network's visited and non-overlapping EID-prefix set
ρ	Relative attack intensity
δ	Attack overlap
τ	Proportion of legitimate traffic within a trace

rate. Additionally, let the attack overlap δ be the ratio between the number of prefixes common to Ω and Ψ and the cardinality of Ψ , thus $\delta = |\Omega \cap \Psi|/|\Psi|$.

If a network trace with average working-set $s(u)$ is augmented by a scanning attack of relative attack intensity ρ and overlap δ , the resulting average working-set becomes

$$s_a(u + \rho u) = \begin{cases} s(u) + \rho u - \frac{\delta s(u_k)}{u_k} u, & u < u_k \\ s(u) + |\Omega| - \delta s(u) & \end{cases} \quad (5)$$

where $u_k = |\Omega|/\rho$, and it represents the number of legitimate packets after which the attack exhausts all $|\Omega|$ destinations and the scan restarts. The aggregate working-set has three components. The first is due to legitimate traffic, $s(u)$, and the second due to the attack packets, ρu . However, because the two may overlap, a third component subtracts the number of shared prefixes. For simplicity, we approximate the probability of having normal and attack traffic destinations overlap to be uniform. Thereby, the growth of the overlap is linear with u up to u_k , where it reaches a maximum of $\delta s(u_k)$ and afterwards linear with $s(u)$.

After a change of variable and denoting the ratio of legitimate traffic in the trace by $\tau = 1/(1 + \rho)$, then $u_k = \frac{|\Omega|}{1-\tau}$, and the previous equation becomes

$$s_a(u) = \begin{cases} s(\tau u) + \left(1 - \tau - \frac{\tau \delta s(u_k)}{u_k}\right) u, & u < u_k \\ (1 - \delta) s(\tau u) + |\Omega| & \end{cases} \quad (6)$$

Then, by derivation, the miss as a function of the number of processed packets is

$$m_a(u) = \begin{cases} \tau m(\tau u) + \left(1 - \tau - \frac{\tau \delta s(u_k)}{u_k}\right), & u < u_k \\ \tau(1 - \delta) m(\tau u) & \end{cases} \quad (7)$$

Unfortunately, in this case, the miss rate cannot be represented analytically as a function of the cache size since $s_a^{-1}(u)$ is not expressible in terms of standard mathematical functions. It can, though, be computed numerically as a function u , when $s(u)$ is known. Then, given that both $s_a(u)$ and $m_a(u)$ are known, they suffice to understand the cache's miss rate as a function of the cache size. The resulting model predicts overall cache misses, not only those due to legitimate traffic, and therefore provides an estimate of the control plane overload, not a data plane performance estimate for legitimate traffic. We empirically validate the result in Section IV-C.

IV. MODEL VALIDATION

In this section, we present the evaluation methodology and the results that validate our models. We start by describing our datasets and then pass on to present the ITR emulator employed

TABLE II
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
Av. pkt/s	75312	46936	64484	232063
Ψ prefs ^a	92801	94964	109451	143775
Av. pref/s	2323	1952	2123	2560

TABLE III
ROUTING

	upc 2009	upc 2011	upc 2012	cesca 2013
BGP _{RT} ^a	288167	400655	450796	455647
BPG _{ϕ} ^b	142154	170638	213070	216272
Ψ /BGP _{ϕ}	0.65	0.55	0.51	0.66

^aBGP routing table as reported by routeviews^bFiltered BGP routing table with no prefix overlaps

^aThe maximum number of non-overlapping prefixes visited within a trace

in the empirical evaluation of the cache. Last, we compare the empirical results to those predicted by the analytical models.

A. Datasets

Four one-day packet traces were used for the purpose of our experiments. Three of them have been captured at the 2-Gb/s link connecting several of our university's campus networks to the Catalan Research Network (CESCA) and consist only of egress traffic. The Universitat Politècnica de Catalunya (UPC) campus, Barcelona, Spain, has more than 36 k users. The first trace dates back to May 26, 2009, while the other two are from October 19, 2011, and November 21, 2012. The fourth trace was captured on January 24, 2013, at CESCA's 10-Gb/s link connecting it to the Spanish academic network (RedIris) and consists of egressing academic traffic. In 2011, CESCA was providing transit services for 89 institutions, including the greater part of public Catalan schools, hospitals, and universities.

Table II summarizes some of the important properties of our datasets. First of all, it can be seen that *cesca 2013*, being an aggregate trace, is about 3.6 times larger than the most recent UPC trace in terms of number of packets and packet rate. However, it only contains 1.3 times more prefixes. This shows that although the number of users and packets exchanged is considerably higher, the diversity of the destinations is only slightly incremented. Out of the UPC traces, *upc 2009* exhibits a surprisingly high number of packets, but this is explained by the very large packet rates seen during the active hours of the day. In fact, the average packet rate during the peak hours was 4.7 times higher than for the rest of the day. Again, this difference did not reflect in the number of unique prefixes observed in a 1-s window as, on average, we observed just 1.3 times more prefixes in the peak hours than during the remainder of the day. These two observations suggest that higher packet rates, either resulting from larger user sets or from higher throughput flows, do not increase destination diversity (as illustrated in Fig. 3) but instead reinforce temporal locality. In addition, these properties also explain why the working-set curves for *upc 2009* present a time-of-day behavior [see Fig. 3(a)].

B. Map-Cache Emulator

To evaluate the two models and the effectiveness of the working-set as a tool for cache performance prediction, we implemented a packet trace-based emulator that mimics basic ITR functionality.

Both for computing the working-sets in Section III-B and for the cache performance evaluation, destination IP addresses had to be mapped to their corresponding prefixes. We considered EID-prefixes to be of BGP-prefix granularity. For each traffic trace, we linked IP addresses to prefixes using BGP routing tables (BGP_{RT}) downloaded from the RouteViews archive [30] that matched the trace's capture date. In particular, we used collector *route-views4* situated at the University of Oregon, Eugene, OR, USA. The only preprocessing we performed was to filter out more specific prefixes. Generally, they are used for traffic engineering purposes, but LISP provides mechanisms for a more efficient management of these operational needs that do not require EID-prefix deaggregation. The filtering also avoids the so called *cache-hiding* problem when the presence of a coarse (parent) prefix prevents the installment of more specific (child) ones as all child prefix misses are wrongfully classified as parent hits. We refer to the resulting list as BGP _{ϕ} . Table III shows the size of the routing tables used for each trace and provides the proportion of prefixes seen within each trace out of the total registered in the filtered routing table, Ψ /BGP _{ϕ} . It may be seen that this ratio is always higher than 0.5, so more than half of the possible destination prefixes are visited in one day for all traces.

For each packet processed, the emulator maps the destination IP address to a prefix in BGP _{ϕ} . If this prefix is already stored in the ITR's cache, its cache entry is updated, and the emulator continues with the next packet. Should the prefix not yet be stored in the cache, two possibilities arise. First, if the cache is not full, the destination prefix is stored in and the processing proceeds to the next packet. Second, if the cache is full, an entry is evicted, the new prefix is stored in, and then the emulator moves to the next packet. The entry to be evicted is chosen according to the LRU eviction policy. We use LRU because, as explained in Section III-C, its performance should be close to optimal due to the stationarity of the trace generating process. Accordingly, the performance of the cache should be appropriately described by (4).

Scanning attacks are emulated by uniform insertion of attack packets in between those pertaining to the processed traffic trace, according to the relative attack intensity ρ . The number of attack addresses generated depends on δ , the overlap between Ω and Ψ , and the number of destinations in BGP _{ϕ} . Note that $\Omega \subseteq \text{BGP}_\phi$ and $\Psi \subseteq \text{BGP}_\phi$. Therefore, presuming the attack maximizes the number of addresses used to increase effectiveness, we have that $|\Omega| = |\text{BGP}_\phi - \Psi| + \delta|\Psi|$. In particular, when no overlap exists, we generate $|\Omega| = |\text{BGP}_\phi - \Psi|$ new destination addresses, while for a full overlap, the attack consists of

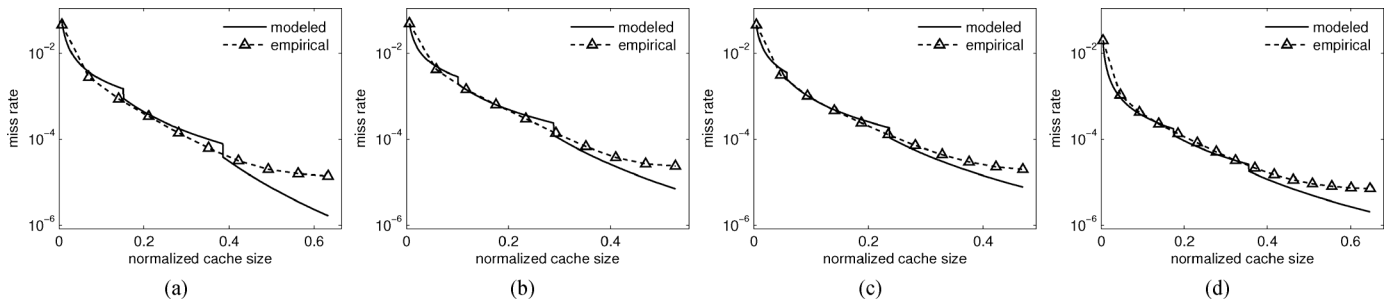


Fig. 4. Comparison between empirical and modeled LRU miss rates [obtained using (4)] of the four traces. For each trace, cache size is normalized with respect to the size of the EID namespace, BGP_ϕ (see Table III). (a) *upc* 2009. (b) *upc* 2011. (c) *upc* 2012. (d) *cesca* 2013.

$|\Omega| = |BGP_\phi|$ addresses. If $\delta \neq 0$, the addresses used in the attack and part of Ψ are uniformly distributed among those part of $BGP_\phi - \Psi$.

C. Comparison of Analytical and Empirical Results

To validate the models, we use the emulator to estimate empirical cache miss rate for several cache sizes. Fig. 4 presents a comparison of the empirical and predicted results for normal traffic, where cache size is normalized with corresponding BGP_ϕ routing table size (see Table III). It may be seen that, typically, the absolute error is negligible despite the discontinuities of the model due to the piecewise fitting. Furthermore, the equation appropriately predicts performance staying acceptable even for small cache sizes, a fact also observed by [7], [8], and [10]. Note that our emulator does not rely on TTL to remove stale entries (in fact, we assume TTL to be infinite), but instead uses a fixed cache size and LRU eviction policy to manage stored content. Nevertheless, combining LRU and TTL should only slightly increase miss rate as frequently used entries may be dismissed on TTL expiry. This has been experimentally observed for both LISP [8] traffic and DNS [31] queries.

The figures present the cache miss rate only up to a fraction of BGP_ϕ , the one associated to $|\Psi|$, because the growth of $s(u)$ cannot be extrapolated after this point. Since the working-set grows slower for larger u , as $\alpha(u)$ is a strictly decreasing function, potentially much longer traces would be needed to enable inference about larger cache sizes. In fact, given that only part of the whole prefix space may be visited by the clients of a stub network, even for longer traces the analysis may be limited to a cache size lower than $|BGP_\phi|$.

There are two limitations to the precision of our analysis for large u values. First, as cache size increases and approaches $|\Psi|$, the accuracy of the prediction diminishes. This is explained by the 24-h length of the traces, whereby there are few working-set curves that span close to a whole day and thus grow to reach the maximum number of destination prefixes. Recall that the start times for the working-set curves span the whole trace and are spaced by 30 min, so the last curves consist of few packets. As a consequence, $s(u)$ is estimated using a reduced number of points, i.e., with a lower precision, at the higher end. To counter this effect, we compute average working-sets of slightly diminished length. The second limitation is the bias of our emulation results for large cache sizes. Caches whose sizes are close to $|\Psi|$ fill only once the traces are processed almost in their entirety. Due to this cold-start effect, cache are exposed to a low number

of hits up to the end of the traces. As these hits do not manage to outweigh the misses generated during the cache fill up, the miss rate of the emulator for large cache sizes is slightly overestimated. Nevertheless, despite these limitations, it may be seen that the results still yield a good fit for large u .

We also validate our cache model that accounts for scanning attacks considering as scenarios a complete attack overlap or one that is zero. Thereby, the attackers may use as attack prefix set either the whole EID-prefix space or just the part not visited by the attacked network's clients. In the latter case, if the cache is not large enough to hold all prefixes, all attack packets would generate a cache miss. However, note that building such a prefix set would require full knowledge about the network's traffic. In the former case, some packet destinations may generate cache hits, but $|\Omega|$ may potentially be much larger, and this could prove beneficial to the attacker. In light of these properties, we consider the two attacks as *worst-case* scenarios, from the attacked network's perspective, for the situations when attackers respectively have or do not have knowledge about the attacked network's traffic.

Fig. 5 compares the analytical and empirical results for the cache miss rate, when $\delta \in \{0, 1\}$ and $\rho \in \{0.01, 0.1, 0.5\}$. We present the results just for *upc 2012* since those for the other three traces are similar. It may be observed that for both no overlap and complete overlap, the results exhibit little absolute error. In particular, for $\delta = 0$ when cache size is larger and miss rates are less than 10^{-4} , the errors are more significant. As in the case of the cache model, this is explained by the trace length and the reduced number of points used in estimating the higher end of $s(u)$. The effect is not noticeable for $\delta = 1$ because $s_a(u)$ reaches its maximum (saturates) for low values of u due to the larger attack set. Consequently, the fit is very good along the whole spectrum of cache sizes. Once the cache size reaches $|BGP_\phi|$, the miss rate becomes 0 since there are no more destination prefixes outside those already present in the cache to generate a miss.

V. DISCUSSION

In this section, we discuss the results and predictions of our models regarding map-cache performance for both normal and malicious traffic. We also discuss possible avenues to diminishing the effect of cache attacks.

The results we obtained are relevant only when reasoned about jointly with the traffic traces used in the analysis. However, the diversity of our data sets and previous results from

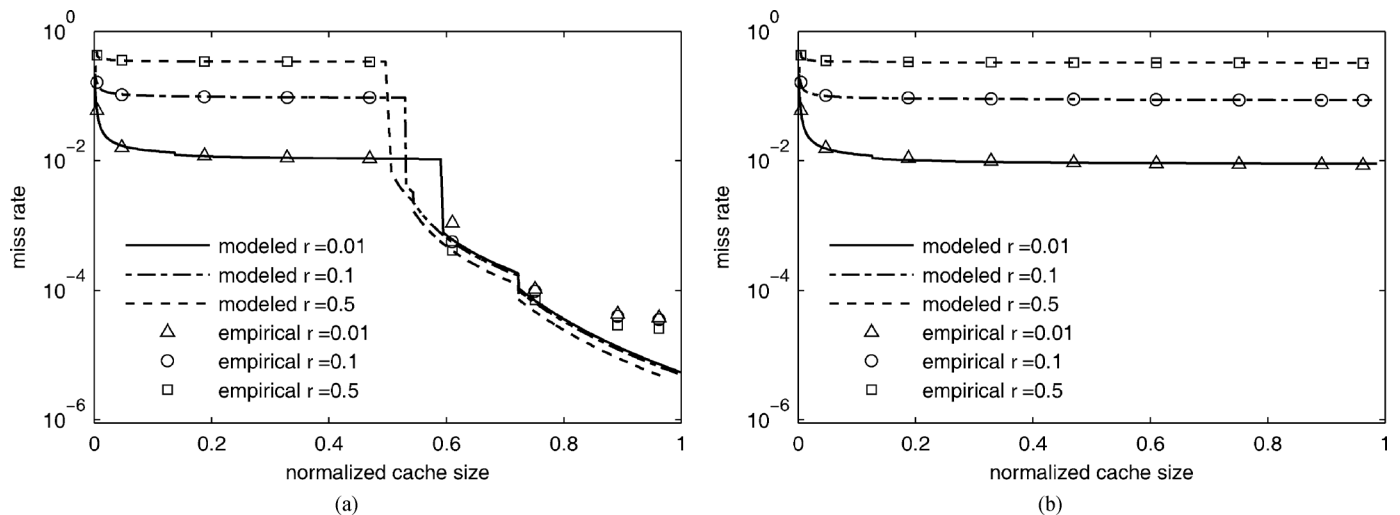


Fig. 5. Comparison between empirical and modeled miss rates, as estimated by (6) and (7), under scanning attacks with attack intensity $\rho \in \{0.01, 0.1, 0.5\}$ for trace *upc 2012*. Interestingly, nonoverlapping attacks are less damaging than overlapping ones, which do not yield acceptable performance if cache size is less than BGP_ϕ (b). Cache size is normalized with respect to the size of the EID namespace, BGP_ϕ (see Table III). (a) No overlap ($\delta = 0$), $\Omega = BGP_\phi - \Psi$. (b) Complete overlap ($\delta = 1$), $\Omega = BGP_\phi$.

Kim [9], Iannone [7], and Kim [8] suggest that the properties uncovered are not the expression of isolated user and network behavior.

A. Cache Model Results and Predictions

Under the condition of a stationary generating process or, equivalently, the approximate time translation invariance of the working-set curves, our methodology enables the estimation of the time-invariant piecewise functions $\alpha(u)$ and $\beta(u)$ that characterize the locality of a network traffic trace from the average working-set size $s(u)$. This further facilitates the following two findings. First, due to the low variance of $s(u)$ and experimentally proven good performance, we can now recommend the use of the LRU eviction policy for LISP caches. Second, in such a situation, (4) may be used to dimension the cache sizes in operational environments, according to the desired miss rate. The prediction of its mathematical expression, considering that $\alpha^*(c) \rightarrow 0$ when c increases, is that miss rate decreases at an accelerated pace with cache size and finally settles to a power-law decrease. This may also be observed in Fig. 4, where at each discontinuity point, the function switches to a faster decreasing curve. Of course, the speed of the decrease depends on the degree of locality present in the trace. Overall, the equation indicates that cache sizes need not be very large for obtaining good performance. For instance, having a cache of size 10% of BGP_ϕ , about 14 k–21 k entries for UPC traces and 21 k for the CESCA one, would result in a miss rate of approximately 0.09%–0.2% and, respectively, 0.03%.

In this context, an important point would be to determine the extent of time over which the results and predictions hold. Fig. 6 provides a coarse answer for the particular case of the traffic used in our analysis. First, considering the UPC traces, it may be observed that over a span of 3 and a half years, the average working-set is rather stable when size is less than about 50 k prefixes. In fact, the $s(u)$ curves of *upc 2011* and *upc 2012* are very similar, independent of u value, while the one for *upc 2009* exhibits a lower slope for $u > 100$ M packets. This might

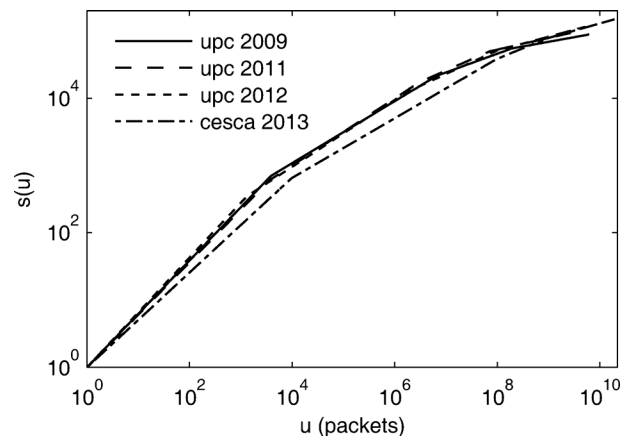


Fig. 6. Evolution of $s(u)$ over the years for *upc* traces compared to *cesca*. Given that $s(u)$ for *upc* appears to be very stable over the course of more than 3 years, the comparison to the larger *cesca* trace suggests that $s(u)$ might be influenced by a network's number of users and destination popularity distribution.

appear to be in agreement with the relative, year-over-year, increase of BGP_ϕ . However, the relative differences between the values of $s(u)$ and the increase of BGP_ϕ are not directly related since the growth from 2009 to 2011 was smaller than the one from 2011 to 2012. This is to be expected because many of the new prefixes, resulting from the sustained growth of the Internet's edge [32], may never become destinations for users of other edge networks. Thus, a direct relation between the increment of the routing table size and that of $s(u)$ should not necessarily be expected. Instead, given the good overlap for lower values of u , prefix popularity distribution should be more relevant to the shape of the working-set than the absolute number of destination. Second, the comparison between *cesca* 2013 and the *UPC* traces reveals that larger user sets, and implicitly higher traffic rates, result in a slightly slower growing $s(u)$. In fact, the only noticeable difference is at short timescales, where the larger trace has a smaller slope. This could be explained by a change in the destination popularity distribution, as *cesca* 2013

aggregates more types of user profiles, but also by a shift in the short-term temporal correlation. However, considering the large number of users, their synchronization at short timescales seems rather unlikely.

Then, although apparently stable over relatively long time spans, the shape of $s(u)$ seems to be influenced by nontrivial interactions between the number of clients the prefix popularity distribution and possibly other unexplored factors. We are not ready to model these interactions here. Thus, our cautious inference is that the average working-set should be stable over time if the number of clients and the popularity distribution are relatively stable.

Despite not being indicated by our measurements, it may be finally proven that the variability with time of $s(u)$ is highly dependent on properties of the network being measured, themselves time-dependent. Should this be the case, the methodology we develop in this paper is still valuable for the analysis of cache performance if not for long-term provisioning of caches.

B. Cache Poisoning and Management Strategies

We compared the effect of scanning attacks with zero or complete overlap to identify the one most damaging and to quantify their effect on cache performance. Associated results were presented in Fig. 5.

If cache sizes are small, both attacks results in very high cache miss rate, including for $\rho = 0.01$, when the attack has a rate of only 644 packets per second. In this range, miss rate is almost independent of attack overlap, only slightly higher for $\delta = 0$ due to the informed selection of attack address space. However, for $\delta = 0$, the cache performance is much improved after a certain threshold is passed, whereas for $\delta = 1$, it barely changes up to when cache size becomes BGP_ϕ and miss rate drops to 0. In other words, the nonoverlapping attack may be absorbed with larger cache sizes, while not the overlapping one. Perhaps counterintuitively, overlapping attacks are more damaging against a map-cache. They are easier to generate, as they do not require prior knowledge about the attacked network. However, they are also harder to defend against since, after a certain point and for a wide range of values, increasing the cache sizes does not yield much improved performance if less than BGP_ϕ .

Arguably, the most worrisome result we observe is the rather high miss rate, which barely drops under 0.01, even for $\rho = 0.01$ and only when the normalized cache size is higher than 0.2. As a comparison, under normal operation, this miss rate would be obtained with a normalized cache size of about 0.02, an order of magnitude less. Therefore, considering the high packet throughput of border routers, some more complex cache management strategies should be set in place to avoid hundreds to thousands of packet drops per second.

One possible first step to circumventing the effect of cache polluting attacks would be to detect them prior to taking action. This may be achieved with $s(u)$ if a ground truth estimation of its shape exists. Thus, if $s(u)$ is known beforehand, an estimate of the average miss rate for the cache size used can be computed. Then, if the instantaneous miss rate surpasses this estimate by more than a configured threshold, a cache protecting action could be taken. For instance, the top, or most recently used, part of the cache could be “protected” against eviction.

Other measures could include the implementation of a cache hierarchy or the limiting of user request rate for new destinations. In the former case, evicted entries should be stored in a larger but higher access time cache, while in the latter, some of the network elements should monitor per-user traffic and filter out attack attempts. Since it is not within the scope of this paper, we do not explore other aspects related to the implementation of such tools.

VI. RELATED WORK

Feldmeier [33] and Jain [34] were among the first to evaluate the possibility of performing destination address caching by exploiting the locality of traffic in network environments. Feldmeier analyzed traffic on a gateway router and showed that caching of both flat and prefix addresses can significantly reduce routing table lookup times. Jain, however, performed his analysis in an Ethernet environment. He was forced to concede that despite the locality observed in the traffic, small cache performance was struggling due to traffic generated by protocols with deterministic behavior. Both works were fundamental to the field, however their results were empirical and date back to the beginning of the 1990s, years when the Internet was still in its infancy.

Recently, Kim *et al.* [9] showed the feasibility of route caching after performing a measurement study within the operational confinement of an ISP's network. They show by means of an experimental evaluation that LRU cache eviction policy performs close to optimal and that working-set size is generally stable with time. We also observe the stability of the working-set for our data sets, but we further leverage it to build an LRU model instead of just empirically evaluating its performance.

Several works have previously looked at cache performance in location/identity split scenarios considering LISP as a reference implementation. Iannone *et al.* [7] performed an initial trace-driven study of the LISP map-cache performance. Instead of limiting the cache size by using an eviction policy, their cache implementation evicted stale entries after a configurable timeout value. Furthermore, Kim *et al.* [8] have both extended and confirmed the previous results with the help of two ISP traces and by considering LISP security aspects in their evaluation as well. They found that securing the control plane generates a cache size overhead in the order of 20% and that diminishing TTL values barely affect cache performance but do considerably reduce cache size. Ignoring control plane security concerns, which we did not consider, and despite differences regarding the use of TTL for cache eviction, the results of these last two works seem to be in agreement with ours. In contrast to the control plane, we found that securing the data plane against scanning attacks either requires more resources (a cache capable of holding all the EID namespace) or more complex eviction policies. Zhang *et al.* [11] performed a trace-based mappings cache performance analysis assuming an LRU eviction policy. They used two 24-h traffic traces captured at two egressing links of the China Education and Research Network backbone network and concluded that a small cache can offer good results. Other works that looked at map-cache performance are [35] and [36], however they are mainly focused on tuning TTL to

ensure cache consistency. Finally, Jakob *et al.* [10] analyzed the performance of several LISP mapping systems and, without focusing on a cache analysis, also observed very low miss rates for a cache model similar to that used in [7].

Our work confirms previous LISP cache analysis results, however it also tries to provide a better understanding of the reasons behind the relatively good performance of map-caches. In this sense, it introduces an analytical model that could be used to theoretically evaluate or dimension the caching performance for operational needs. Moreover, to the best of our knowledge, it is also the first work to perform an analysis and propose an analytical model for the map-cache performance when under scanning data-plane attacks.

VII. CONCLUSION

The implementation of a location/identity split at network level has been recently recommended as a viable solution to the Internet's routing scalability problems. However, unless deployed with care, the cure may prove worse than the disease, so a good understanding of the newly introduced components, like the map-cache, is paramount.

In this paper, we propose a methodology to evaluate map-cache performance. Our model is built by exploring the link between cache performance and parameters that approximate the intrinsic locality of packet-level user traffic. To this end, we advance the use of the working-set model as a tool to capture said properties, but also as a performance predictor. Accordingly, we define a framework wherein to perform the analysis and find that the clustering of the working-set curves is the only condition needed to ensure the accuracy of the model. We empirically validate our result by emulation, using traffic traces collected at the edges of a campus and an academic network.

Besides the possibility of using the model for cache dimensioning or detecting attacks in operational environments, we believe the equation may also be used as part of more complex models that evaluate the scalability of loc/id separation-based architectures. To stress these points, we show the versatility of our methodology by characterizing map-cache performance for our datasets and by building an extension that accounts for cache pollution attacks. Our observations indicate that increasing cache size quickly diminishes miss rates in normal conditions, but has little to no effect under simple cache pollution attacks. In the latter case, we advise that more complex management strategies be devised and set in place. In the future, we plan to fully investigate the sources of locality in prefix-level network traffic to better understand their impact on cache performance.

APPENDIX

PROOF OF PROPOSITION 1

Proof: If $\forall u$ the working-set size $w(t, u)$ is normally distributed, and therefore independent of t , it follows that the process generating $w(t, u)$ is stationary. This in turn implies the stationarity of the process generating the reference string and, as a result, necessity is proven. It remains to be proved that if the working-set curves are generated by the same stationary process, then they will tend to cluster. In [26], it is shown that for a certain window size u , the distribution of $w(t, u)$

converges to a normal distribution if the locality conditions hold. This proves sufficiency.

ACKNOWLEDGMENT

The authors want to express their gratitude to D. Saucez, L. Iannone, and P. Barlet for their insightful comments.

REFERENCES

- [1] G. Huston, "BGP report," 2014 [Online]. Available: <http://bgp.potaroo.net/>
- [2] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," Internet Engineering Task Force, RFC 4984 (Informational), 2007.
- [3] M. Hoefling, M. Menth, and M. Hartmann, "A survey of mapping systems for locator/identifier split Internet routing," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 4, pp. 1842–1858, 4th Quart., 2013.
- [4] A. Agarwal, J. Hennessy, and M. Horowitz, "An analytical cache model," *Trans. Comput. Syst.*, vol. 7, no. 2, pp. 184–215, May 1989.
- [5] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," in *Proc. IEEE INFOCOM*, 1999, vol. 1, pp. 126–134.
- [6] L. Rizzo and L. Vicisano, "Replacement policies for a proxy cache," *IEEE/ACM Trans. Netw.*, vol. 8, no. 2, pp. 158–170, Apr. 2000.
- [7] L. Iannone and O. Bonaventure, "On the Cost of Caching Locator/ID Mappings," in *Proc. 3rd CoNEXT*, Dec. 2007, pp. 1–12.
- [8] J. Kim, L. Iannone, and A. Feldmann, "Caching locator/id mappings: An experimental scalability analysis and its implications," *Comput. Netw.*, vol. 57, no. 4, pp. 897–909, 2013.
- [9] C. Kim, M. Caesar, A. Gerber, and J. Rexford, "Revisiting route caching: The world should be flat," in *Proc. 10th PAM*, Berlin, Germany, 2009, PAM '09, pp. 3–12.
- [10] L. Jakob, A. Cabellos-Aparicio, F. Coras, D. Saucez, and O. Bonaventure, "LISP-TREE: A DNS hierarchy to support the LISP mapping system," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1332–1343, Oct. 2010.
- [11] H. Zhang, M. Chen, and Y. Zhu, "Evaluating the performance on ID/loc mapping," in *Proc. IEEE GLOBECOM*, 2008, pp. 1–5.
- [12] P. J. Denning, "The working set model for program behavior," *Commun. ACM*, vol. 11, no. 5, pp. 323–333, 1968.
- [13] F. Coras, A. Cabellos-Aparicio, and J. Domingo-Pascual, "An analytical model for the LISP cache size," in *Proc. NETWORKING*, Berlin, Germany, 2012, vol. 7289, Lecture Notes in Computer Science, pp. 409–420.
- [14] R. Atkinson and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) architectural description," Internet Engineering Task Force, RFC 6740 (Experimental), 2012.
- [15] H. Ballani, P. Francis, T. Cao, and J. Wang, "Making routers last longer with Viaggre," in *Proc. USENIX NSDI*, 2009, vol. 9, pp. 453–466.
- [16] J. Saltzer, "On the naming and binding of network destinations," Internet Engineering Task Force, RFC 1498 (Informational), 1993.
- [17] T. Li, "Recommendation for a routing architecture," Internet Engineering Task Force, RFC 6115 (Informational), 2011.
- [18] X. Misseri, J.-L. Rougier, and D. Saucez, "Internet routing diversity for stub networks with a map-and-encap scheme," in *Proc. IEEE ICC*, 2012, pp. 2861–2866.
- [19] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "The Locator/ID Separation Protocol (LISP)," Internet Engineering Task Force, RFC 6830 (Experimental), 2013.
- [20] D. Saucez, L. Iannone, O. Bonaventure, and D. Farinacci, "Designing a deployable Internet: The locator/identifier separation protocol," *IEEE Internet Comput.*, vol. 16, no. 6, pp. 14–21, Nov.–Dec. 2012.
- [21] "LISP testbed," [Online]. Available: <http://www.lisp4.net/>
- [22] V. Fuller, D. Farinacci, and D. Lewis, "LISP Delegated Database Tree (LISP-DDT)," Internet Engineering Task Force, draft-ietf-lisp-ddt-04, work in progress, 2012.
- [23] M. Menth, M. Hartmann, and M. Hofling, "FIRMS: A future Internet mapping system," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1326–1331, Oct. 2010.
- [24] R. Hinden, "New scheme for internet routing and addressing (ENCAPS) for IPNG," Internet Engineering Task Force, RFC 1955 (Informational), 1996.
- [25] J. R. Spirin and P. J. Denning, "Experiments with program locality," in *Proc. AFIPS*, New York, NY, USA, Dec. 5–7, 1972, pp. 611–621.
- [26] P. J. Denning and S. C. Schwartz, "Properties of the working-set model," *Commun. ACM*, vol. 15, no. 3, pp. 191–198, 1972.

- [27] A. Banerjee, J. J. Dolado, J. W. Galbraith, and D. Hendry, *Co-Integration, Error Correction, and the Econometric Analysis of Non-Stationary Data*. Oxford, U.K.: Oxford Univ. Press, 1993.
- [28] S. Jin and A. Bestavros, "Sources and characteristics of web temporal locality," in *Proc. Int. Symp. Model., Anal. Simulat. Comput. Telecommun. Syst.*, 2000, pp. 28–35.
- [29] D. Saucez, L. Iannone, and O. Bonaventure, "LISP threats analysis," Internet Engineering Task Force, draft-ietf-lisp-threats-10.txt, work in progress, 2014.
- [30] University of Oregon, Eugene, OR, USA, "RouteViews project," 2005 [Online]. Available: <http://www.routeviews.org>
- [31] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "Dns performance and the effectiveness of caching," *IEEE/ACM Trans. Netw.*, vol. 10, no. 5, pp. 589–603, Oct. 2002.
- [32] L. Cittadini *et al.*, "Evolution of Internet address space deaggregation: Myths and reality," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1238–1249, Oct. 2010.
- [33] D. Feldmeier, "Improving gateway performance with a routing-table cache," in *Proc. 7th Annu. IEEE INFOCOM*, 1988, pp. 298–307.
- [34] R. Jain, "Characteristics of destination address locality in computer networks: A comparison of caching schemes," *Comput. Netw. ISDN Syst.*, vol. 18, no. 4, pp. 243–254, 1990.
- [35] X. Li, H. Zhou, F. Qiu, H. Zhang, and I. You, "A lease algorithm to maintain strong mapping cache consistency," *Wireless Pers. Commun.*, vol. 76, no. 4, pp. 713–730, 2014.
- [36] S.-C. Chen, C.-J. Chang, and Y.-H. Chu, "A fuzzy cache lifetime determination scheme for LISP mapping systems," in *Proc. ICOIN*, Jan. 2013, pp. 297–302.



Florin Coras received the Diplom Engineer degree in telecommunication engineering from the Technical University of Cluj-Napoca, Cluj-Napoca, Romania, in 2009, and the M.Sc. degree in computer networks from Universitat Politècnica de Catalunya—BarcelonaTech (UPC), Barcelona, Spain, in 2011, and is currently pursuing the Ph.D. degree in computer science at UPC.

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 telecommunications and the Ph.D. degree in computer science from Universitat Politècnica de Catalunya—BarcelonaTech (UPC), Barcelona, Spain, in 1982 and 1987, respectively.

He is a Full Professor with the Computer Architecture Department, UPC. Since 1983, he has been a Lecturer with the Computer Architecture Department, and currently he is teaching with the Computer Science School (FIB) and the Telecommunications Engineering School (ETSETB). He is the head of the

Broadband Communications and Services research group (CBA) of UPC that

is an accredited research group by the Catalan Government. He is one of the promoters and founders of the interdepartmental research center for Advanced Broadband Communications (CCABA) that is recognized as a distinguished research center by the UPC. He was Director of the Computer Architecture Department of the UPC from 2005 to 2011. He has been advisor of 11 Ph.D. dissertations and has published over 130 papers in international journals and conferences.



Darrel Lewis has over 20 years of experience as an Engineer for routing infrastructure vendors and network service providers. In his current position as LISP Technical Leader with Cisco Systems, San Jose, CA, USA, he has been instrumental as an architect in support of the development of the Locator/ID Separation Protocol (LISP). As one of the original coauthors of RFC 6830 and other LISP RFCs, he supports the development of the protocol, its implementation, and its production deployment around the globe. Prior to Cisco, he worked for

Riverhead Networks (which was acquired by Cisco and became the Cisco Guard DDoS mitigation system), where he was the lead Consulting Engineer.

Mr. Lewis is active in the North American Network Operators Group (NANOG), the Internet Engineering Task Force (IETF), and is a noted instructor in the fields of both Internet routing and security.



Albert Cabellos-Aparicio received the B.Sc., M.Sc., and Ph.D. degrees in computer science engineering from Universitat Politècnica de Catalunya—BarcelonaTech (UPC), Barcelona, Spain, in 2001, 2005, and 2008, respectively.

In 2004, he was awarded with a full scholarship to carry out Ph.D. studies with the Department of Computer Architecture, Technical University of Catalonia (UPC), Barcelona, Spain. In 2005, he became an Assistant Professor with the Computer Architecture Department and a Researcher with the

Broadband Communications Group. In 2010, he joined the NaNoNetworking Center in Catalunya, Barcelona, Spain, where he is the Scientific Director. He has been a Visiting Researcher with Cisco Systems and Agilent Technologies, and a Visiting Professor with the Royal Institute of Technology (KTH), Stockholm, Sweden, and the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. He has participated in several national (Cicyt), EU (FP7), USA (NSF), and industrial projects (Samsung and Cisco). He has given more than 10 invited talks (MIT, Cisco, INTEL, MIET, Northeastern University, etc.) and coauthored more than 15 journal and 40 conference papers. His main research interests are future architectures for the Internet and nano-scale communications.

Dr. Cabellos-Aparicio is an editor of *Nano Computer Networks* and founder of the ACM NANOCOM conference and the IEEE MONACOM workshop and the N3Summit. He has also founded the LISPmob open-source initiative along with Cisco.