On the Performance of the LISP Beta Network

Florin Coras* Damien Saucez† Luigi Iannone ‡ Benoit Donnet§
*Universitat Politècnica de Catalunya (Spain) †INRIA (France) ‡Telecom ParisTech (France) §Université de Liège (Belgium)

***Abstract*—The future Internet has been a hot topic during the past decade and many approaches towards this future Internet, ranging from incremental evolution to complete clean slate ones, have been proposed. One of the proposition, LISP, advocates for the separation of the identifier and the locator roles of IP addresses to reduce BGP churn and BGP table size. Up to now, however, most studies concerning LISP have been theoretical and, in fact, little is known about the actual LISP deployment performance. In this paper, we fill this gap through measurement campaigns carried out on the LISP Beta Network. More precisely, we evaluate the performance of the two key components of the infrastructure: the control plane (i.e., the mapping system) and the interworking mechanism (i.e., communication between LISP and non-LISP sites). Our measurements highlight that performance offered by the LISP interworking infrastructure is strongly dependent on BGP routing policies. If we exclude misconfigured nodes, the mapping system typically provides reliable performance and relatively low median mapping resolution delays. Although the bias is not very important, control plane performance favors USA sites as a result of its larger LISP user base but also because European infrastructure appears to be less reliable.**

I. INTRODUCTION

During the last decade, the Internet has strongly evolved. Its natural growth combined with factors such as multihoming and interdomain traffic engineering has lead to an increase of the BGP routing tables size ([1], [2]) and so-called *churn*, i.e., the traffic necessary to keep those tables up to date [3].

To cope with those problems, several solutions have been proposed. Most of them are based on the *locator/identifier separation* paradigm, which relies on the existence of two different address types [4]: the *identifiers* and the *locators*. An identifier is used on end-systems to identify a connection endpoint, while a locator refers to the attachment point in the Internet topology. The proposals can be classified in two main categories: those associating locators directly to end-systems (e.g., HIP, SHIM6, … [5]) and those associating locators to routers (e.g., LISP [6], [7]). It is worth to notice that, in the current Internet architecture, end-systems' IP addresses have a dual semantic, being at the same time identifier of communication endpoints and locator of these endpoints in the Internet topology.

Among the different proposals, the most widely deployed is the *Locator/IDentifier Separation Protocol* (LISP [6]). In LISP, an identifier is used to identify a connection endpoint and is only locally routable. In contrast, a locator refers to a node attachment point in the Internet topology and is globally routable. Like other solutions, LISP needs an indirection mechanism, commonly named *mapping system*, to bind identifiers to locators or, stated differently, to glue the identifier addressing space to the locator addressing space. Note that LISP does not introduce any new address format. Rather, it splits the current IP addressing space into two sub-spaces where addresses have one clear single semantic (i.e., either identifier or locator).

Since 2008, LISP has been deployed in the wild through the *LISP Beta Network* [8]. Nevertheless, despite the existence of such a worldwide real playground, nothing (or little) is known about its actual behavior and performance. Most of the research on LISP, up to now, has been based either on theoretical models or simulations. Such research works focused mostly on specific aspects of the overall LISP architecture, like for instance, the LISP Cache [9], [10], [11], the dynamics of locators [12], mobility [13], security [14], [15], [16], or traffic engineering techniques [17], [18] and the mapping system [19].

More recently, Saucez et al. [20] analyzed the LISP Beta Network by assessing the state of its deployment and how it has evolved over years. However, they proposed only a first look based on a very high level analysis. In this paper, we propose a much deeper study and a more thorough analysis of the LISP Beta Network. In particular, the evaluation focuses on the two key components of the infrastructure, namely: the *control plane* (i.e., the mapping system), and the *interworking* (i.e., communication between LISP and non-LISP sites). To this end, raw data has been collected through different approaches: (*i*) existing datasets from LISPmon [21]; (*ii*) active LIG (LISP Internet Groper [22]) measurement campaigns; (*iii*) traceroute measurement campaigns towards specific LISP nodes.

Concerning the control-plane evaluation, our findings show that the resolution delay is rather small and stable over time, but the mapping system lacks reliability in the European region. Concerning the interworking evaluation, we provide lower bounds for the stretch introduced by the use of proxies as interworking technology. Additionally, we observe that the BGP routing infrastructure encourages the selection of proxies that are geographically far apart, with an important negative impact on the performance.

The remainder of this paper is organized as follows: Sec. II provides the required background on LISP; Sec. III describes the current LISP Beta network deployment; Sec. IV and Sec. V detail the measurements and analysis carried out as well as discussing the main findings; finally, Sec. VI concludes the paper by summarizing its main achievements and sketching further research directions.

II. LISP ARCHITECTURE

This section provides the necessary background on the overall LISP architecture and its functioning. We first briefly introduce how the LISP Data Plane (i.e., packet forwarding) works (Sec. II-A), before describing the LISP Control Plane (Sec. II-B). Finally, we summarize how a LISP-enabled Internet can interwork with the legacy Internet (Sec. II-C).

ISBN 978-3-901882-58-6 © 2014 IFIP

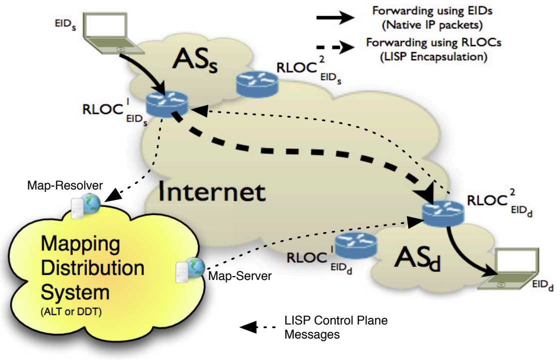

Fig. 1. LISP data and control plane packet forwarding sequence.

A. *LISP Data Plane*

The *Locator/IDentifier Separation Protocol* (LISP) ([6], [7] separates the identification and localization roles of IP addresses by introducing two logical addressing spaces: *(i)* the *Routing LOCator space* (RLOC), that is globally routable; *(ii)* the *Endpoint IDentifier space* (EID), that is only locally routable. With this separation, the Internet core, also known as *Default Free Zone* (DFZ), handles RLOCs addresses like it is done today, i.e., maintaining routes so that packets can be forwarded between any router within the DFZ. Stub networks use instead the EID addressing space that has only local scope.[1] The implication of such a separation lays in stub networks not needing anymore a full knowledge of the Internet routing information, whereby the DFZ does not need anymore to advertise the EID space in its routing infrastructure. Nonetheless, in order to provide end-to-end communication, another level of indirection is required.

In the data plane, LISP provides this level of indirection through a tunneling mechanism over the DFZ, as shown in Fig. 1. More specifically, any communicating host generates regular IP packets using its EID as source address and the destination EID as destination address. Forwarding towards the border router is done as usual in the local domain (see the solid line in the upper left corner of Fig. 1). The border router, now called *Ingress Tunnel Router* (ITR), will encapsulate the packets using the RLOC addressing space, i.e., using its RLOC address as source address and the destination RLOC as destination address in the tunnel header encapsulating the original packet [6]. The encapsulated packets can now be forwarded over the DFZ (see the dashed line in Fig. 1). The border router at the destination site, now called *Egress Tunnel Router* (ETR), will decapsulate the LISP packets so that the original packet can be forwarded to its final EID destination (see the solid line in the bottom right corner of Figure 1).

B. *LISP Control Plane*

In order to perform the data plane operations, mainly related to encapsulation and decapsulation as described in the previous section, tunnel routers need to be able to associate EIDs to RLOCs. The binding between the two addressing spaces is named *mapping*. Such a *mapping* enables a tunnel router (generally referred to as an *xTR*) pursues two objectives. First, retrieving the RLOCs associated to a given EID and to be used in the outer header when encapsulating. Second, performing consistency and security checks when decapsulating. In particular, a mapping consists of an EID prefix associated with a list of <*RLOC, Priority, Weight*> tuples. When selecting an RLOC, the one having the highest priority is preferred. In the case of several RLOCs having the same priority, the weight is used for load balancing flows among them. Mappings are stored in two data structures present on xTRs: the *LISP Database* and the *LISP Cache*.

The LISP Database stores the mappings for EID prefixes for which the xTR is an RLOC. On ITRs, this allows one to select the source RLOC to be used in encapsulation. While on ETRs, it is used to verify that the tunnel router is actually the RLOC of the destination EID in the inner header, hence, being able to deliver the original inner packet to its final destination. The LISP Database is populated by configuration.

The LISP Cache stores mappings for EID prefixes used in ongoing communication towards/from distant sites. On ITRs, this allows the selection of the RLOC to be used in encapsulation to reach the destination EID. While on ETRs, it is used to perform a basic anti-spoof verification, checking whether the encapsulating router (outer header source address) is actually an RLOC for the source EID (inner header source address).

Differently from the LISP Database, the LISP Cache is populated on demand. The procedure is triggered by the first packet of a new flow that does not find among the stored mappings one matching its destination EID. As the name suggests, the LISP Cache only temporarily stores the mappings used in ongoing communications; its entries expire and are purged from the data store when not used for a certain amount of time ([9], [23], [10], [11]).

Because of its LISP Cache element, LISP introduces an important change in the Internet routing infrastructure. Indeed, while current BGP-based routing relies on a push model, i.e., pushes all routing information to the whole Internet, LISP-based routing relies on a pull model, i.e., pulling routing information only when actually needed. The key point of this new approach is *how to make routing information available on an on-demand fashion?* To this end, the LISP control plane introduces a new system, the *Mapping Distribution System (MDS)* in order to provide a lookup infrastructure from where mappings can be retrieved upon explicit query (cfr., Fig. 1).[2]

From an abstract point of view, the MDS works as follows. The ITR that needs a mapping for a new flow first sends a query, consisting of a `Map-Request` message to a *Map-Resolver* [24]. The query is forwarded by the Map-Resolver inside the mapping distribution system according to the specific protocol/architecture used, to reach the *Map-Server* where the site using the requested EID has registered the mapping. The Map-Server then forwards the query to the xTR that registered the mapping for the EID for which a mapping is requested. The

---

[1] A stub network refers to a network that does not provide transit services to others (i.e., being only source or destination of its outbound and inbound traffic). For availability and robustness, stub networks are usually multi-homed (i.e., connected to more than one provider).

[2] The terms *Mapping Distribution System*, *MDS*, and *mapping system* are used interchangeably in this paper.

xTR will in turn send the reply, consisting of a `Map-Reply` message containing the requested mapping, directly to the ITR that, in first place, sent the query. Fig. 1 shows this process as dotted lines. The Map-Resolver and Map-Server elements represent respectively where to ask for a mapping and where to register a mapping so to make it available to other LISP sites. They provide a general front-end for any mapping system, "hiding" the specific mapping system in use to the LISP tunnel routers.

Several mapping systems have been proposed ([19], [25]), however, only two have been deployed (cfr., Section III): *LISP Alternative Topology* (LISP+ALT [26]) and *LISP Delegated Database Tree* (LISP-DDT [27]).

C. Interworking with the Legacy Internet

Interworking between LISP enabled sites and non-LISP sites (also referred to as *legacy Internet*) is possible through the use of two types of proxies [28]: *Proxy Ingress Tunnel Routers* (Proxy-ITRs or PITRs) and *Proxy Egress Tunnel Routers* (Proxy-ETRs or PETRs). We equally refer to them with the acronym *PxTR*.

On the one hand, Proxy-ITRs allow non-LISP sites to send packets to LISP sites without any changes to protocols or equipment at the non-LISP site. A Proxy-ITR acts as an ITR. It advertises in BGP the EID-prefix space on behalf of the LISP sites so that non-LISP sites can reach them. This way, a non-LISP site just sends IP packets using the destination EID as destination IP address. Such a packet will be routed to the Proxy-ITR that advertises the prefix the destination EID belongs to. Once the packet reached the Proxy-ITR, the latter will also take care of encapsulating it into a LISP packet and forwards it towards the destination RLOC.

On the other hand, Proxy-ETRs allow LISP sites to send packets to non-LISP sites, acting as normal ETRs for everything that is in the legacy Internet. A LISP site sends LISP-encapsulated packets using the Proxy-ETR as destination RLOC. Then, the Proxy-ETR decapsulates the packets and forwards them natively in the legacy Internet. Since the destination address is not an EID, the prefix it belongs to is announced in the DFZ and the packet can be forwarded to the destination.

III. THE LISP BETA NETWORK

Since a few years, LISP is deployed in the Internet [8]. Started as an experimental testbed, the LISP Beta network is experiencing a steady growth that is driving it out from being an experimental deployment [20]. The members of this joint effort are academics and research laboratories, startups offering LISP related services, but also major companies (e.g., Microsoft, Facebook, and Verisign) and operators (e.g., Level3). As shown in Fig. 2, participants of this network are located in 27 different countries, with the higher concentration being in Europe and North America.

The network uses two main EID address spaces, namely 153.16.0.0/16 for IPv4 and 2610:00D0::/32 for IPv6. However, there exist also EIDs in different address ranges. Further, other experimental and anycast prefixes are considered. Until the beginning of March 2012, the mapping system used by the LISP Beta network was LISP+ALT [26], a lookup infrastructure based on the deployment of BGP sessions over a virtual topology built on GRE tunnels. Since March 14th 2012, the mapping system in use is the more flexible LISP-DDT [27], based on a hierarchically distributed database, which is conceptually similar to the Domain Name System.

Fig. 2. Current LISP Beta Network participants geographical location.

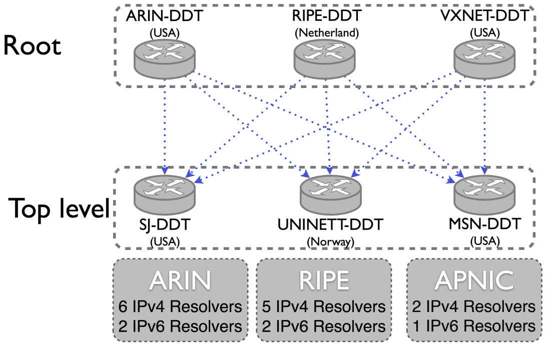

Fig. 3. Current LISP-DDT Mapping System deployment.

While switching mapping system has considerably reduced the configuration and maintenance burden, a slight drop in the performance has been observed [20]. Fig. 3 provides a snapshot of the current deployment of the LISP-DDT mapping infrastructure.

The LISP Beta Network, while still growing, remains limited compared to the pervasive presence of the Internet. Nevertheless, its size is now sufficiently large to allow reasonable measurements, whose results provide insights on its behavior and performance. The outcome of such a work, hence the results we present hereafter, will potentially drive the future evolution of the LISP technology and its deployments model.

IV. CONTROL PLANE PERFORMANCE EVALUATION

As described in Sec. II-B, ingress packet forwarding at an ITR is conditioned by the presence of an entry in the LISP Cache that covers the packet's destination EID. Because a LISP Cache miss triggers a mapping resolution for the EID that caused the miss and because packets causing misses are dropped, it is critical to ensure fast and reliable mapping distribution system operations. In the following, we study mapping distribution system performance metrics to gain a better understanding of the control plane's performance and investigate how these metrics evolved in time.

A. *Measurement Methodology*

We analyze DDT performance using two datasets obtained over the course of two one-week-long measurement campaigns we performed starting March, 19th 2012, and June, 9th 2013. The goal in collecting the data is to obtain a fine grained approximation of the control plane's speed in resolving identifier to locator bindings and with it, a better understanding of the interaction between the deployed infrastructure and globally distributed clients. To this end, we use LIG [22] to send `Map-Request` messages, to all deployed Map-Resolvers, for all existing EIDs in a periodic fashion. We further store the resulting mapping and additional information, if any, LIG returns for subsequent processing. During the two measurement campaigns the number of available resolvers increased from twelve (in 2012) to thirteen (in 2013). We queried them every 30 minutes, over the span of one week, from a set of different vantage points, part of academic networks, commercial Internet, and PlanetLab, all spread across Europe and USA. During the measurement campaign of 2012 eight vantage points were available, whereas only six were available in 2013. We consider a query successful if the associated answer contains mappings information (an EID prefix and a set of RLOCs) or a negative reply, that is, a confirmation that the requested EID is not reachable via LISP at the time the query was issued. If no answer is received within three seconds, the query is assumed to have failed.

Apart from the mapping, LIG also returns an estimate of the latency needed for the `Map-Request`/`Map-Reply` exchange, namely, the round-trip time (RTT) a packet requires to travel to the Map-Resolver, cross the mapping system, reach the ETR, and return to the point that sent out the request. The second step may, however, be significantly shortened when DDT is used. Cold start operation dictates that DDT Map-Resolvers perform a piecemeal (iterative) search for an EID within the DDT hierarchy, starting at the root. Therefore, resolvers incrementally decrease the scope of their search by querying all nodes on a tree branch, descending from the root, up to when they encounter a Map-Server with more information about the desired mapping, or a node indicating the mapping does not exist. To avoid this slow process, resolvers cache pointers to nodes in the DDT hierarchy and start lookups from the node responsible for the most specific prefix that includes the EID being requested. Typically, in practice, these entries are removed after one day of inactivity. As a result, the time required to traverse the mapping system, as observed by LIG, often consists of only the Map-Resolver to Map-Server delay. An important difference for negative replies is that they may be provided by (i) DDT nodes, when the EID is part of a delegation hole (i.e., no DDT node serves the EID-prefix that matches the requested EID), or (ii) Map-Servers, when ETRs fail to update their registrations. In these cases is recommended to remove the entries after only 15 minutes [27]. Consequently, we generally expect negative replies to consistently have a lower resolution delay.

In selecting the list of destination EIDs, we leverage on the data gathered and published by the LISPmon project [21], a LISP-beta network monitoring platform that periodically crawls the IP addressing space, through the DDT mapping system, to discover all the prefixes used as EID addressing space. In addition to the EID-prefixes registered by LISP enabled domains, the dataset LISPmon assembles also contains the list of DDT delegation holes that have no mapping information associated. We download the aggregate list daily and select for each prefix therein, irrespective of its type, the first address in the corresponding address range, which we finally feed to LIG.

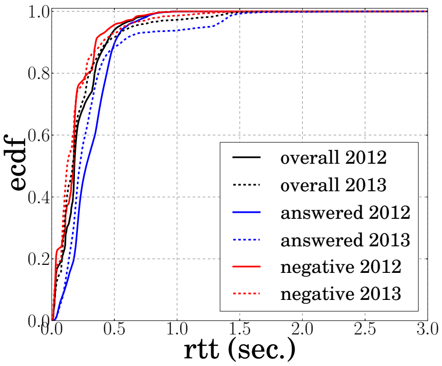

Fig. 4. RTT ECDF comparison between the 2012 and 2013 datasets.

B. *Results*

We have identified in the 2012 and 2013 datasets 318 and, respectively, 615 EID-prefixes for which at least one `Map-Reply` has been received; an important increase over the course of one year. However, when looking at EIDs with one or more locators, the growth is not as large, increasing from 203 to 247 prefixes. For brevity, we refer to the two types of replies as either *solved* or *answered* (as opposed to negatively answered).

Fig. 4 illustrates the Empirical Cumulative Distribution Functions (ECDFs) of the RTT for the two datasets, aggregated over all vantage points. In addition to the overall RTT, we also look at two of its disjoint components, the time needed to solve complete mappings and that needed to obtain negative replies, given their different resolution processes. In the case of the former, a `Map-Request` must reach an ETR, which finally provides an authoritative answer to the inquiring ITR, whereas in case of the latter, the resolution stops at a node configured with a negative record. Overall, RTT values tend to be fairly small. Namely, we found the median to be less than 200ms and only 10% of the measurements to exceed 500ms. Although these values are several orders of magnitude larger than inter-packet delays in high speed links, packet forwarding should not be affected as mappings are only seldom requested, due to the use of LISP Caches [9].

Then, instead of forwarding delays, a more relevant comparison would be to consider the *TCP retransmission timeout*, namely, the time the network stack of an operating system waits prior to retransmitting a TCP `SYN` packet, if no acknowledgment is received. When the delay in retrieving the mapping necessary for an EID exceeds such a timeout, connections initiated to said EID experience an important slowdown, since they have not only their first packet dropped but the retransmitted one as well. However, given that operating systems

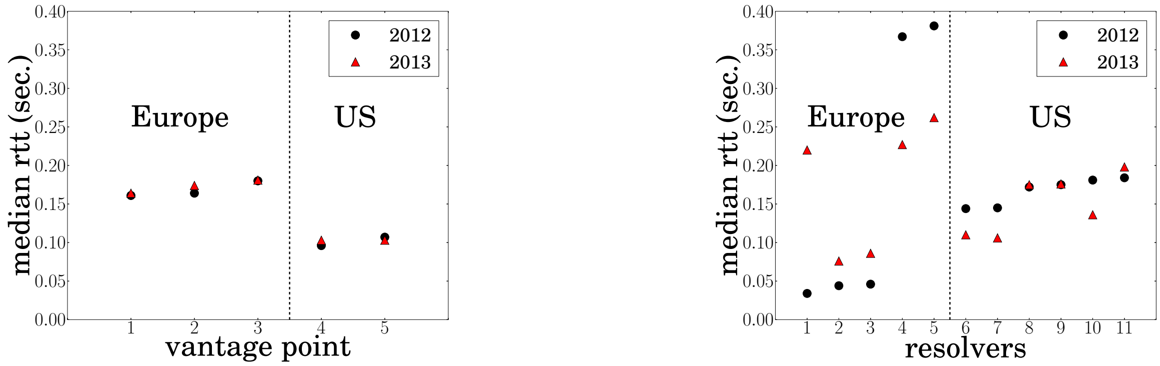

(a) Vantage point

(b) Map-Resolver

Fig. 5. Median RTT per vantage point and Map-Resolver comparison between the 2012 and 2013 datasets.

are often configured to send the first retransmission after roughly three seconds [29], the low RTT values we find in our results shows that LISP Cache misses should not have a very detrimental effect on flow start-up. Even in the worst cases, delays we observed seldom exceed 1.5 seconds, i.e., half the retransmission timeout.

Focusing on the difference between the answered and the negatively answered replies, we see that the latter are retrieved faster, as anticipated. On the one hand, this confirms our assumption that negative replies have a shorter resolution chain. On the other hand, we see that an important part, around 10%, are obtained after a surprisingly high delay. Unfortunately, DDT resolvers originate negative replies themselves when DDT-based lookups end in negative results, hence, since we do not have access to DDT internal traffic, we cannot identify the sources in our measurements. Nevertheless, in light of their high values, we believe this group of negative replies are not provided early on in the resolution process, like in the case of mapping holes, rather they are provided by distant (in the DDT hierarchy) leaf DDT nodes or Map-Servers. It is worth to remark that such delays could be the result of misconfigurations or upgrades within the DDT hierarchy itself. Nevertheless, we suspect that more often than not these EIDs either lack associated mappings, by configuration, or have their authoritative ETR powered off/disconnected (may be for maintenance reasons).

Comparing the results for the two years, we find that overall performance largely stayed the same, despite the growth of the addressing space. The most pronounced difference can be found in the tails of the distributions, where the 2013 dataset exhibits a larger number of high latency points, which, however, is very similar to the increase observed for the answered replies. On closer inspection we found the majority (over 75%) of these points to pertain to two Map-Resolvers that experienced reliability issues over the course of our measurements. Therefore, they are not the result of the changes within DDT nor an indication of mapping retrieval times deterioration. In fact, up to 500ms (about 90% of the answers), the 2013 distribution grows faster, so the mapping system performed slightly better for the larger part of the queries.

To better understand the year-over-year changes, we look in greater detail at the differences in median RTT between our datasets, focusing on (i) the vantage point observed performance, as an indication of the latency to be experienced by any LISP ITR, and (ii) the Map-Resolver performance, to identify possible deployed infrastructure issues. We choose to compare medians due to the slight skew of the RTT distributions seen in Fig. 4. In order to avoid biasing the results, we perform the comparison only over the vantage points and using the Map-Resolvers the two datasets have in common. Fig. 5 depicts our findings, where the different vantage points are just indicated by numbers from 1 to 5.

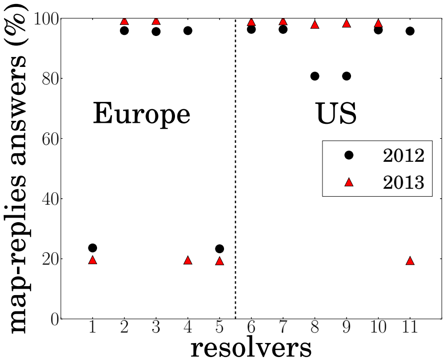

Fig. 6. Percent of answered Map-Replies for the two datasets.

Fig. 5(a) confirms that there is no significant difference in mapping resolution times between the two datasets, even when looking from vantage point perspective. We also find that European (EU) vantage points have a slightly larger delay, with respect to those situated in the US, but we discuss this in more details later. Importantly, the result shows a consistent behavior (no year-over-year change) for all our vantage points. Although to be expected in normal operation, as each vantage point queries all Map-Resolvers, an important but isolated change in latency would have been a good indicator of bogus underlying BGP connectivity to either part of the Map-Resolvers or the ITRs.

Fig. 5(b) provides a comparison of Map-Resolver performance for our datasets. Surprisingly, we observe a considerable variation in performance but mainly for resolvers located within Europe. For those located within the US, performance changed very little and when that happened, it mainly slightly improved. Still, on closer inspection, we found the biggest changes to have occurred for the European resolvers with

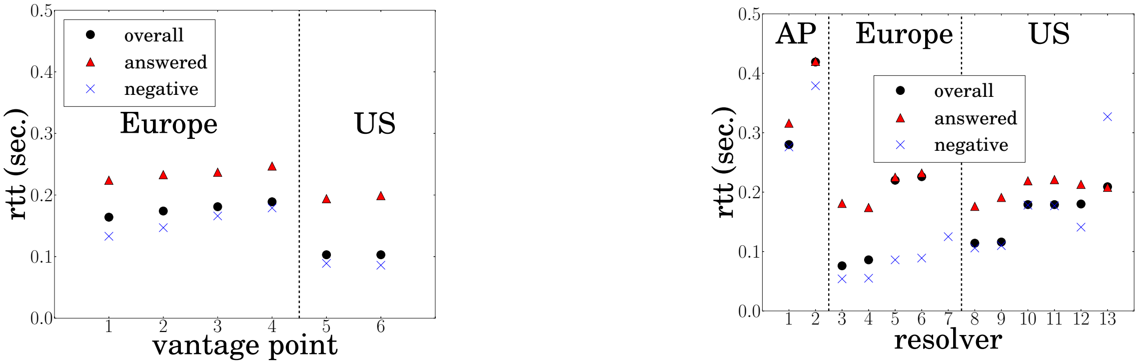

(a) Vantage point

(b) Map-Resolver

Fig. 7. Median RTT per vantage point and Map-Resolver for the 2013 datasets.

sporadic reliability, namely those labeled “1” and “2” in Fig. 5(b). Therefore, they not only incur a larger than average latency increase but also offer unstable availability. In fact, if we consider the resolver reliability results (presented in Fig. 6 where the different Map-Resolvers are indicated by a number from 1 to 11), all of the Map-Resolvers, with the exception of those previously mentioned, appear to have improved their reliability (now being close to 100% availability) over the course of one year. Out of the US set, only the Map-Resolver labeled “11” experienced reliability problems. Putting these last two results in perspective, we may conclude that the US LISP infrastructure has now reached a mature level of stability whereby it offers consistent performance. Unfortunately, the same conclusion does not hold for its European counterpart. Although two resolvers offer good performance, three out of five resolvers are unreliable and incur larger than average latency increase.

Finally, we study the absolute performance exhibited by our latest dataset, considering all vantage points and Map-Resolvers used for the measurements. Fig. 7 illustrates the overall, answered, and negative reply median latency, broken down over all resolvers and vantage points. One of the first noticeable differences, as illustrated in Fig. 7(a) and previously mentioned, is that the European vantage points are at a slight disadvantage when compared to their US peers, independently of the vantage point and of the metric considered. However, two observations are in place. First, North America has a larger number of typically active xTRs compared to Europe and, second, as previously discussed, three of the EU resolvers have intermittent availability. Thereby, requests from US vantage points are directed more often to geographically closer Map-Resolver and receive replies more often by geographically closer ITRs. These positional advantages are better perceivable when comparing the latencies of negative replies, that, for mapping holes, provide an estimate of the vantage point to resolver RTT. In this sense, even the best connected EU vantage point has a higher latency that those in US. The variations within the European group can also be attributable to differences in geographical position or Internet connectivity. For instance, all the vantage points are part of the European academic network, but the first vantage point, the best connected vantage point of the group, is situated in Belgium, closer to the trans-Atlantic connection points to US, whereas the forth is in Spain.

In spite of the slight geographical bias, we observe a couple of encouraging results. First, the latency for answered replies is typically less than double that of the negative replies. This is surprisingly low given that DDT has the potential to introduce important delays if EIDs allocation is not done such that Map-Servers and xTRs are at least within the same continent. But most encouraging, absolute differences for answered replies are close to negligible, even when the variation around the median is accounted for, although this is not shown here. In practice, this means that, for the time being, the position of the vantage point introduces almost no bias in the mapping resolution process of EIDs with at least one locator associated. Nonetheless, we do find a more important influence in the case of negative map replies.

It is important to observe that the result does not suggest the independence of latency from the Map-Resolver and vantage point pair. In fact, Fig. 7(b) shows that the converse actually holds. As the majority of our vantage points are located within Europe, the two reliable European resolvers offer the best performance, followed by the US located ones and finally those from Asia-Pacific.

V. INTERWORKING PERFORMANCE EVALUATION

As explained in Sec. II-C, LISP interworking refers to non-LISP sites exchanging traffic with LISP sites. This is possible with the addition of an intermediate network element, or proxy, between both sites that is in charge of converting LISP (resp. non-LISP) traffic into non-LISP (resp. LISP) traffic. This intermediate element comes in two flavors: ingress (Proxy-ITR) and egress (Proxy-ETR). Proxy tunnel routers are also usually referred to with the acronym *PxTR*. These intermediate elements are of the highest importance as the switch from a non-LISP Internet to a fully LISP Internet should be done incrementally. In this section, the performance of LISP when PxTRs are required is evaluated.

Mimicking the structure of Sec. IV, in the following we first present the measurement methodology and the metric we consider before presenting and discussing the obtained results.

A. *Measurement Methodology*

We measure the interworking performance through delay (i.e., round-trip time) and topological distance (i.e., number

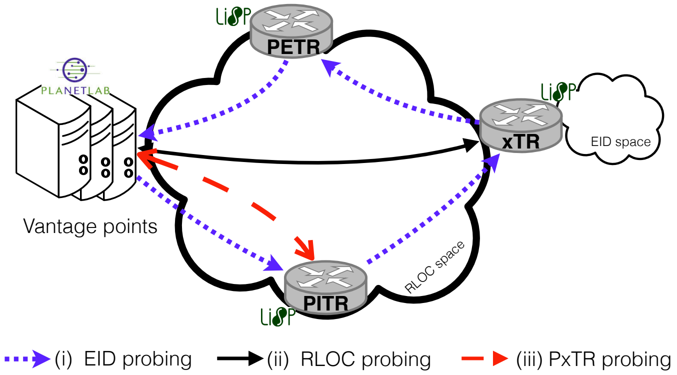

Fig. 8. Breakdown of components influencing interworking measurements.

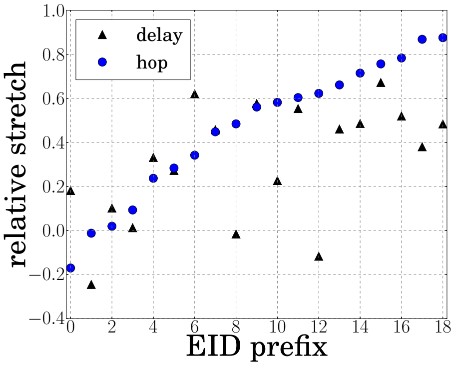

Fig. 9. Relative stretch caused by interworking for LISP sites with both EIDs and RLOCs measurements.

of hops). Measurements have been done using standard traceroute [30] from 200 PlanetLab vantage points. The vantage points were located mostly in North America then Asia, and, finally, Europe. None of the vantage points belongs to the LISP Beta Network. Measurements were done between November, 4th and November, 7th 2013. As targets for our traceroute campaigns, we used three types of address: (i) the EIDs, (ii) the locators, and (iii) the PxTRs.

First, we directly measure each of the existing LISP sites via one of their EID addresses. In particular, we selected the first IP address in the EID prefix (i.e., typically, in the current deployment, an ETR loopback address) of the targeted LISP site. A key element when measuring directly the EID is that it allows us to have information covering the path from (i) the PlanetLab vantage point towards the topologically closest PITR, (ii), this PITR to the targeted EID, and (iii) the path back from the EID towards the PlanetLab vantage point via the PETR of the LISP measured site. Proceeding this way, we are able to produce traffic between a non-LISP site (the PlanetLab vantage point) and a LISP site.

Second, we compare the performance that a site would experience when connecting to the existing LISP sites if it were also using LISP. For that, we consider our vantages points as if they were xTRs and use them to perform traceroutes to each RLOC associated to every EIDs of existing LISP sites. That way, the vantage points measure the delay and number of hops to reach the xTRs without passing by any PxTR.

Finally, we measured each PxTR to estimate a lower bound on the performance for traffic between LISP and non-LISP sites. The performance of the end-to-end path (i.e., the path between the non-LISP target and the LISP target) should be, at least, equal to the performance to the PxTR. Doing so, we do not have information about the actual end-to-end performance between sites, which, however, is not possible to collect, but our results highlight the fact that PxTRs will, for sure, reduce the overall performance than if both sites were LISP capable (i.e., not using PxTR).

We retrieved the list of EIDs and their RLOCs from the *LISPmon* project [21], on May, 17th 2013. Out of this list we only kept the 116 EIDs that had at least one RLOC in the corresponding mapping. We obtained the list of PxTR addresses directly from the LISP Beta network website [8]. Fig. 8 summarizes the paths followed by traceroute probes for the three measurements cases described here above.

As the path followed by packets from a site to another site depends, all other things being equal, on whether or not the PxTR is used, we expect to observe different delays and number of hops in the different situations. More precisely, when interworking is employed, we expect to observe an increase of topological metric values (e.g., delay, number of hops) as the traffic is deviated from its shortest path.

To quantify this increase of topological metric values, we define the *relative stretch* $\rho_{i,j,p}$ between sites i and j when using the PxTR p as follows:

$$\rho_{i,j,p} = \frac{\hat{d}_{i,j,p} - d_{i,j}}{d_{i,j}}, \qquad (1)$$

where $d_{i,j}$ refers to the metric of interest (e.g., delay or number of hops) between i and j by following the shortest path on the RLOC space without the intervention of a PxTR and $\hat{d}_{i,j,p}$ is the one obtained if the traffic transits through PxTR p instead. A positive (resp. negative) value of $\rho_{i,j,p}$ indicates that the metric value obtained by transiting via the PxTR is longer (resp. shorter) than without. A performance degradation (i.e., a positive relative stretch) is expected due to interworking.

B. Results

We must identify two types of LISP sites (assuming one EID prefix corresponds to one site) from our measurement campaign: LISP sites that answered traceroute probes through both their EIDs and their RLOCs, and LISP sites that answered traceroute probes only through their RLOCs. Among the 116 measured EID prefixes, we were able to traceroute at the same time an EID and all its RLOCs for 19 prefixes, while we were able to measure the RLOCs but not the EIDs in 52 cases.

Fig. 9 depicts the impact of interworking for the first case where traceroutes succeeded for both the EID and the RLOCs. This situation is the one that gives the most information. In this study, we consider two topological metrics (i) the round-trip delay, annotated `delay`, and the number of hops between the vantage point and the measured point, annotated `hop`. Fig. 9 gives, for the 19 EID prefixes, the average relative stretch computed according to Eqn. 1 where $\hat{d}_{i,j,p}$ is the value of the

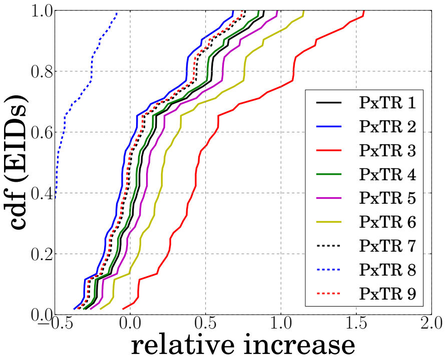

Fig. 10. Approximated relative delay stretch due to interworking.

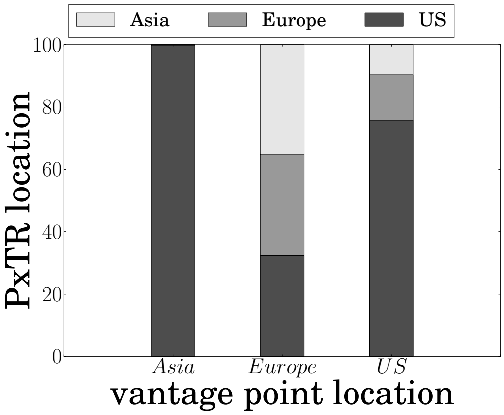

Fig. 11. Breakdown of PxTR location with respect to vantage point location.

topological metric obtained with traceroute probes sent to the EID while $d_{i,j}$ is the value obtained with traceroute probes sent to the RLOC of the same LISP site. As our vantage points are all non-LISP sites, the path followed by probes sent to EIDs have to pass by a PITR and back via a PETR and the selection of the PxTRs is dictated by BGP meaning that the used PxTR can differ between vantage points.

Fig. 9 confirms the intuition that interworking has globally a negative impact on topological metrics. As a matter of fact, the vast majority of measurements show that the delay and number of hops is larger when interworking is used. Indeed, it shows an increase in the delay by at least 20% for 70% of the EID prefixes. However, drawing clear conclusions is harder for the number of hops as by the definition the intermediate routers traversed from a PITR to an ETR or from an ITR to a PETR are invisible, hence paths between PxTRs and xTRs appear like a single hop. Despite this particularity, we see that the number of hops also increases for the vast majority of the measured EID prefixes.

In the following, we generalize the discussion to EID prefixes for which only RLOCs have been successfully measured. In this case, computing the actual stretch is impossible. However, we can compare topological metrics from the vantage points to the RLOCs with the ones obtained from the vantage points to the PITR. Since the PITR is the first point where the traffic is deviated from its normal route when interworking is used, this comparison gives a lower bound on the relative stretch (i.e., the actual stretch is always higher than this value). Therefore, if the relative stretch obtained considering only the partial path is positive, it means that indubitably the actual traffic would be negatively impacted by the interworking technique.

Fig. 10 depicts the relative delay stretch for the nine PITRs deployed in the LISP Beta network. Fig. 10 gives the cumulative distribution of the relative stretch. Fig. 10 shows no less than 43% of the EID prefixes would certainly experience a delay increase when using interworking, regardless of the PxTR that is selected by the underlying BGP routing. Interestingly, we can see that PxTR3 behaves significantly worst than the other PxTRs with 95% of the EID prefixes suffering from a delay increase when using interworking. The reason why delay is different between the different PxTRs comes from the geographical spread of the PxTR that are deployed worldwide.

Fig. 10 shows aggregated results for all vantage points, independently of their location. As PxTRs are replicated and spread worldwide, one could expect that each non-LISP site will use its closest PxTR. Unfortunately, while it is true that BGP will make sure that packets will always be captured by their closest PITR, the notion of closest point in BGP is driven by the economical relationship between autonomous systems, not their geographical position. In the following, we determine whether or not the default PxTR selected by BGP is close to the vantage points or not. For that, we have determined the regional location of each vantage point from PlanetLab information and located the regional position of PxTR using the LISP Beta network information. With the traceroute campaign, we establish the PxTR that is used by default for each measurement and compare the region of the vantage point with the region of the default selected PxTR. From the total of 22,828 traceroutes giving such an information, we draw the geographical spread of the selected PxTR as a function of the vantage point position.

Fig. 11 summarizes the information. The cumulated stack-bars in the figure give the PxTR regional position according to the location of the vantage points, aggregated by region (i.e., Asia, America, and Europe). The first observation is that most of vantage points use PxTRs in Asia even if they are located in Europe or America. Notably, the most popular PxTR in our measurements is PxTR3 and it has been selected 79% of the times. Interestingly, only four out of the nine available PxTRs appear in our traceroutes. Also, it is worth remembering that PxTR3 is the PxTR that presents the largest stretch on average (see Fig. 10). These results pinpoint the problem of selecting where to deploy PxTRs given that the route followed by packets is finally decided by BGP and routing policies. Interestingly, Europe presents a more uniform spread between regions. This can be explained by the fact that Europe has a large diversity of ISPs and IXPs giving more peering possibilities.

VI. CONCLUSION

Despite its ever increasing popularity and deep changes in term of usage with the advent of multimedia distribution and mobility, the core Internet protocols have not evolved that much. The changes are so intense that these protocols have shown their limits and maintaining the Internet infrastructure as efficient as in the past is done at the expense of complexity.

A factor that dramatically affects the Internet architecture is the complementary role of IP addresses that play at the same time a role of identification and of localization. In the quest for a better Future Internet, the Locator/Identifier Separation Protocol (LISP) advocates to split the two complementary roles of addresses and decomposes the addresses space in two spaces: the Routing LOCators space (RLOC) and the Endpoint IDentifiers space (EID). To make the glue between the two spaces while staying compatible with the current infrastructure, LISP uses encapsulation in the data plane and a indirection mechanism, called the mapping distribution system, in the control plane. Theoretical studies have shown the potential of LISP in terms of scalability, traffic engineering, or even mobility. However, even if LISP is deployed since 2008, little is known about its real performance.

In this paper, we have measured the LISP Beta network to estimate the performance of the two LISP key components, i.e., the mapping distribution system and the interworking mechanism in charge of connecting LISP and non-LISP networks together.

Our analysis of the mapping distribution system shows that mapping resolution delays are rather small and stable over time and that the system is able to cope with a significant growth of the addressing space. On the downside, we found a large part of the European infrastructure to be unreliable. More importantly, control plane performance is influenced by how the infrastructure and userbase develop. For the time being, differences between geographically spread points are not that important but there exists the danger of evolution becoming unilateral. Our traceroute based measurements of the interworking mechanism show that despite a good geographical spread of the interworking devices, communications between LISP and non-LISP network have lower performance because BGP does not account for geographical position when it selects routes to advertise.

ACKNOWLEDGMENTS

Research work presented in this paper is supported by Cisco University Research Program (CG N. 574857), the ANR-13-INFR-0009 LISP-Lab Project (www.lisp-lab.org), the Spanish Ministry of Education under scholarship AP2009-3790, research project TEC2011-29700-C02, Catalan Government under project 2009SGR-1140, the EIT KIC ICT-Labs SOFNETS, and the Networks for Future Media Distribution activities.

REFERENCES

[1] G. Huston, “BGP routing table analysis reports,” 2004, see http://bgp.potaroo.net.

[2] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, “IPv4 address allocation and the BGP routing table evolution,” *ACM SIGCOMM Computer Communcation Review*, vol. 35, no. 1, pp. 71–80, 2005.

[3] D. Meyer, L. Zhang, and K. Fall, “Report from the IAB workshop on routing and addressing,” Internet Engineering Task Force, RFC 4984, September 2008.

[4] T. Li, “Recommendation for a routing architecture,” Internet Research Task Force, RFC 6115, February 2011.

[5] P. Savola, “A survey of IPv6 site multihoming proposals,” in *Proc. IEEE Internet Conference of Telecommunications (ConTel 2005)*, June 2005.

[6] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, “The locator/ID separation protocol (LISP),” Internet Engineering Task Force, RFC 6830, January 2013.

[7] D. Saucez, L. Iannone, O. Bonaventure, and D. Farinacci, “Designing a deployable Internet, the locator/IDentifier separation protocol,” *IEEE Internet Computing*, vol. 16, no. 6, pp. 14–21, November/December 2012.

[8] “LISP Beta Network,” see http://www.lisp4.net.

[9] L. Iannone and O. Bonaventure, “On the cost of caching Locator/ID mappings,” in *Proc. ACM SIGCOMM CoNEXT*, December 2007.

[10] H. Zhang, M. Chen, and Y. Zhu, “Evaluating the performance on IDLoc mapping,” in *Proc. IEEE Global Communications Conference (GLOBECOM)*, November 2008.

[11] F. Coras, A. Cabellos-Aparicio, and J. Domingo-Pascual, “An analytical model for the LISP cache size,” in *Proc. IFIP Networking*, May 2012.

[12] D. Saucez and B. Donnet, “On the dynamics of locators in LISP,” in *Proc. IFIP Networking*, May 2012.

[13] M. Menth, D. Klein, and M. Hartmann, “Improvements to LISP mobile node,” in *Proc. International Teletraffic Congress (ITC)*, September 2010.

[14] A. Mahdi, A. Al-Nemrat, and D. Preston, “Securing address registration in Location/ID split protocol using ID-based cryptography,” in *Proc. International Conference on Wired/Wireless Internet Communications (WWIC)*, June 2013.

[15] A. Mahdi, “A novel security protocol for resolving addresses in the Location/ID split architecture,” in *Proc. International Conference on Network and System Security (NSS)*, June 2013.

[16] D. Saucez, L. Iannone, and O. Bonaventure, *Solutions for Sustaining Scalability in Internet Growth*. IGI Global, 2013, ch. The Map-and-Encap Locator/Identifier Separation Paradigm: a Security Analysis, pp. 148–163.

[17] D. Saucez, B. Donnet, L. Iannone, and O. Bonaventure, “Interdomain traffic engineering in a locator/identifier separation context,” in *Proc. Internet Network Management Workshop (INM)*, October 2008.

[18] S. Secci, K. Liu, G. K. Rao, and B. Jabbari, “Resilient traffic engineering in a transit-edge separated internet routing,” in *Proc. IEEE International Conference on Communications (ICC)*, June 2011.

[19] M. Hoefling, M. Menth, and M. Hartmann, “A survey on mapping systems for locator/identifier split internet routing,” *IEEE Communications Surveys and Tutorials*, vol. 15, no. 4, pp. 1842–1858, November 2013.

[20] D. Saucez, L. Iannone, and B. Donnet, “A first measurement look at the deployment and evolution of the locator/id separation protocol,” *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 1, pp. 37–43, April 2013.

[21] “LISPmon,” see http://www.lispmon.net.

[22] D. Farinacci and D. Meyer, “The locator/ID separation protocol Internet groper (LIG),” Internet Engineering Task Force, RFC 6835, January 2013.

[23] J. Kim, L. Iannone, and A. Feldmann, “A deep dive into the LISP cache,” in *Proc. IFIP Networking*, May 2011.

[24] V. Fuller and D. Farinacci, “Locator/ID separation protocol (LISP) map-server interface,” Internet Engineering Task Force, RFC 6833, January 2013.

[25] L. Jakab, A. Cabellos-Aparicio, F. Coras, D. Saucez, and O. Bonaventure, “LISP-TREE: a DNS hierarchy to support the LISP mapping system,” *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 8, pp. 1332 – 1343, October 2010.

[26] V. Fuller, D. Farinacci, D. Meyer, and D. Lewis, “Locator/ID separation protocol alternative logical topology (LISP+ALT),” Internet Engineering Task Force, RFC 6836, January 2013.

[27] V. Fuller, D. Lewis, and A. Ermagan, V. Jain, “LISP Delegated Database Tree,” Internet Engineering Task Force, Internet Draft, Work in Progress draft-ietf-lisp-ddt-01.txt, March 2013.

[28] D. Lewis, D. Meyer, D. Farinacci, and V. Fuller, “Interworking between locator/id separation protocol (LISP) and non-LISP sites,” Internet Engineering Task Force, RFC 6832, January 2013.

[29] N. Seddigh and M. Devetsikiotis, “Studies of TCP’s retransmission timeout mechanism,” in *Proc. IEEE International Conference on Communications (ICC)*, June 2001.

[30] V. Jacobson et al., “traceroute,” UNIX, man page, 1989, see source code: ftp://ftp.ee.lbl.gov/traceroute.tar.gz.