A First Measurement Look at the Deployment and Evolution of the Locator/ID Separation Protocol

Damien Saucez
Inria
France
damien.saucez@inria.fr

Luigi Iannone
Telecom ParisTech
France
luigi.iannone@telecom-paristech.fr

Benoit Donnet
Université de Liège
Belgium
benoit.donnet@ulg.ac.be

ABSTRACT

During the last decade, we have seen the rise of discussions regarding the emergence of a Future Internet. One of the proposed approaches leverages on the separation of the identifier and the locator roles of IP addresses, leading to the LISP (Locator/Identifier Separation Protocol) protocol, currently under development at the IETF (Internet Engineering Task Force). Up to now, researches made on LISP have been rather theoretical, *i.e.*, based on simulations/emulations often using Internet traffic traces. There is no work in the literature attempting to assess the state of its deployment and how this has evolved in recent years. This paper aims at bridging this gap by presenting a first measurement study on the existing worldwide LISP network (`lisp4.net`). Early results indicate that there is a steady growth of the LISP network but also that network manageability might receive a higher priority than performance in a large scale deployment.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Internet

General Terms

LISP, Measurement

Keywords

mapping system, flexibility, Internet architecture

1. INTRODUCTION

Motivated by the BGP routing tables size increase [1, 2] and the BGP churn [3], since a few years, the Internet research community is examining the Internet architecture and, in particular, how it could be improved. This tough process aims at redesigning the Internet. Some of the emerging ideas are rather theoretical [4, 5, 6, 7], while others [8, 9, 10, 11] are much more practical with think-tanks heavily discussing within the IETF (Internet Engineering Task Force) and the IRTF (Internet Research Task Force). One of the most successful proposals is *LISP* [10, 11], the Locator/IDentifier Separation Protocol.

The idea behind LISP is to separate the *identifier* and *locator* roles of an IP address. An identifier is used to identify a connection endpoint and is only locally routable. On the contrary, a locator refers to a node attachment point in the Internet topology and is globally routable. LISP provides a mapping system allowing to associate a given identifier to a set of locators. Each locator in the set provides a different path between the two identifiers.

Nevertheless, little is known about LISP out in the wild, under real conditions. This knowledge is of the highest importance, particularly in the perspective of large-scale deployments. So far, no previous work has attempted to report advances and evolution of LISP deployment. We believe this lack is mainly due to the absence of extensive measurements campaigns.

In this paper, we aim at bridging this gap by providing a first measurement study of a real LISP deployment behavior. We rely on the worldwide LISP network, namely `lisp4.net`, for performing an intensive measurement campaign between March and April 2012. We couple active measurements based on the use of the LISP Internet Groper (*LIG* [12]) with the analysis of long-term mapping dataset obtained from the *LISPmon Project* [13]. Furthermore our campaign covers the adoption, by the LISP network, of a new mapping system, allowing us to evaluate the impact and improvement brought to the network.

Our study shows two main results. On the one hand, performance is not the main criterion for the adoption of a new mapping system. Indeed, we observe a slight performance decrease after the mapping system switchover. However, it has to be put in perspective with the much better manageability provided by the new system. On the other hand, while LISP deployment is consistently growing, its traffic engineering capabilities are still under exploited.

The remainder of this paper is organized as follows: Sec. 2 presents the background on LISP and the `lisp4.net` network required for this paper; Sec. 4 describes our measurement campaign and discusses some early observations obtained from our measurements; finally, Sec. 5 concludes this paper and plans future research directions.

2. LISP: A MAP-AND-ENCAP PROTOCOL

In this section, we provide the required background for an easy reading of the remainder of this paper. We first briefly overview the LISP protocol (Sec. 2.1) and, next, the mapping system (Sec. 2.2).

2.1 LISP: Protocol

The *Locator/Identifier Separation Protocol* (LISP) [10, 11], as the name indicates, separates the identifier and the locator roles of IP addresses, introducing so two independent address spaces. The *Endpoint IDentifier space* (EID) identifies end-systems and consists of IP addresses that are only locally routable. The *Routing LOCator space* (RLOC) locates

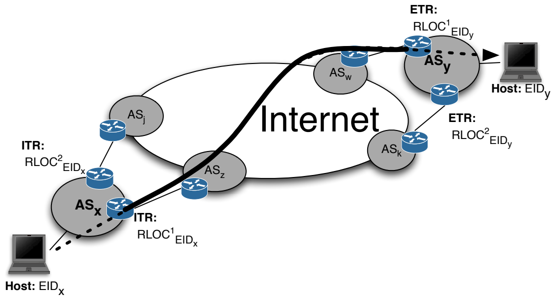

Figure 1: LISP topology example with a communication between host EID$_x$ and host EID$_y$. Dashed lines refer to forwarding between the EID and the RLOC, while the plain line refers forwarding between the RLOCs.

Figure 2: LISP network nodes worldwide (source: www.lispmon.net)

EIDs in the Internet topology and consists of IP addresses that are globally routable. RLOCs are handled by routers in the core Internet like it is today, maintaining routes so that packets can be forwarded between any router. Stub domains, on the contrary, use EIDs, and since they are only locally routable, routers in the core Internet do not need to maintain routes towards EIDs. The main objectives behind this separation are to reduce BGP routing tables size and the BGP churn. An example of how LISP is deployed and where the different addressing spaces are used is depicted in Fig. 1.

To enable the communication among EIDs of different domains, LISP tunnels packets in the core Internet from the RLOC of the source EID to the RLOC of the destination EID.[1] When a packet has to be sent from a source EID to a destination EID, the sender initially creates a standard IP packet, using EIDs as source and destination addresses, that is forwarded to a border router of the source domain for tunneling (this is illustrated by dashed lines in Fig. 1). The border router, also called *Ingress Tunneling Router* (ITR), performs a lookup (locally or through a distributed system – the so-called *Mapping System*) for obtaining a mapping binding the destination EID to its RLOC, that is the border router of the destination domain (also called *Egress Tunneling Router* – ETR). Once the mapping has been obtained, the ITR encapsulates the packet using RLOCs as source and destination IP addresses. The encapsulated packet is then forwarded as usual towards the ETR (this is illustrated by the plain line in Fig. 1). Upon reception of the packet, the ETR decapsulates it and then delivers the original packet to the destination EID.

The flexible usage of tunnel routers offered by LISP does not only allow to achieve routing tables size reduction (LISP's original motivation), but also offer interesting traffic engineering capabilities [11, 14].

2.2 LISP: Mapping System

As explained above, ITRs acquire mappings binding EIDs to a set of RLOCs for ongoing communications via a mapping system that is a key element of LISP. A mapping associates an EID prefix to a list of <*RLOC*, *priority*, *weight*> tuples. The priority and weight, associated to each RLOC, help the ITR in selecting the RLOC to use for reaching a given EID. RLOCs with the highest priority are preferred. When RLOCs have the same priority, the weight is used for load balancing flows among them. So far, several mapping systems have been proposed for LISP [15]. However, only two have been deployed: *LISP Alternative Topology* (LISP+ALT [16]) and *LISP Delegated Database Tree* (LISP-DDT [17]).

LISP Alternative Topology (LISP+ALT) was the initial mapping system for LISP and relied on a BGP overlay [16]. In LISP+ALT, ETRs store mappings they are authoritative for. The overlay is constructed by connecting ETRs together via tunnels, for example GRE tunnels [18]. This ETRs' overlay is called the *Alternative Logical Topology (ALT)* where routers are called *ALT routers*. Any ALT router maintains a BGP session with its neighbor and announces the EID prefixes it is authoritative for, making the EIDs routable in the ALT. At this point, it is worth to notice that BGP is only used to build the ALT, not to announce mappings. To get a mapping, an ITR sends a `Map-Request` for the EID on the ALT topology. The source address for the `Map-Request` is the ITR RLOC and the destination the EID. The `Map-Request` eventually reaches an originator ETR for the EID prefix that matches the destination EID. This ETR resolves the EID and sends a `Map-Reply` directly to the ITR RLOC. `Map-Replies` are not sent on the ALT.

After a few years of experimentation and operation, it has been noticed that LISP+ALT was cumbersome to manage [19, 20]. The *LISP Delegated Database Tree* (LISP-DDT) [17] mapping system has thus been designed with manageability and isolation in mind. LISP-DDT is a DNS-like system with a hierarchy of LISP-DDT Servers queried by *Map Resolvers* and by *Map Servers* [21].[2] The Map Resolver accepts `Map-Requests` from an ITR and resolves the mapping using a DNS-like distributed database. The Map

[1] Note that several RLOCs can be associated to a given EID.

[2] *Map Servers* and *Map Resolvers* are a general front-end for any mapping system, allowing to "hide" the specific mapping system in use to the LISP Tunnel Routers, which now deal only with these two type of servers. Such a front-end can be used also in the context of ALT. However, because such a technology has been developed after the design of ALT, it has been less deployed in the ALT context.

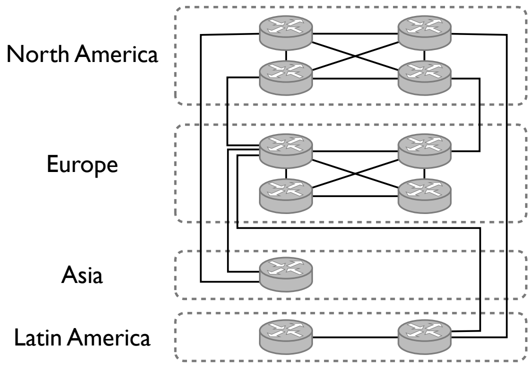

Figure 3: LISP+ALT topology as used in the LISP network. Boxes refer to regions of the ALT topology, each region being made up to several routers (in full mesh). Each region is connected to Europe and North America regions.

Server learns mappings from an ETR and publishes it in such database. The hierarchy is maintained as a tree where each node is responsible for a part of the EID space. A child node is responsible for a portion of the EID space of its parent. Mapping information is only stored at the tree leaves. Intermediate nodes only maintain pointers to their children.

When a mapping must be retrieved for a given EID, a root server is queried first. By definition, a root server is responsible for the entire EID space, this space being divided into several portions, each one being managed by one of the root's children. The root replies with a pointer (*i.e.,* a referral) to its children responsible for the EID prefix to resolve. The process is recursively repeated with the returned child considered as the root of the sub-tree where a mapping can be retrieved for the EID. This recursive process is stopped when a leaf has been reached. In LISP-DDT, leaves are made of *Map Servers*. Each Map Server maintains a list of ETR authoritative for the different EID prefixes registered to it (at least one matching the requested EID). Thus, when the leaf has been reached, the mapping is retrieved by sending a `Map-Request` to one of the ETR authoritative for the matching EID prefix. To simplify the operation, mapping resolutions are performed by Map Resolvers on behalf of ITRs.

It is worth to notice that, like for DNS, in order to speed up the mapping process, it is possible to cache the mapping results [22, 23, 24, 25].

3. THE LISP NETWORK

Since a few years, LISP is actually deployed in the Internet and is experiencing a steady growth that is driving it out from being a simple testbed. Participants in the testbed are startups offering LISP related services but also major companies (*e.g.*, Microsoft, Facebook, and Verisign) and operators (*e.g.*, Level3).[3] As shown in Fig. 2, participants of this network are located in 27 different countries, most of them being in Europe and USA, and consist of both academic institutions and companies. 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.

[3]See `http://www.lisp4.net`

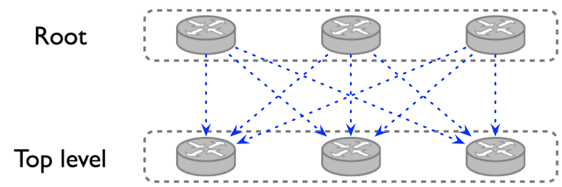

Figure 4: LISP-DDT topology as used in the LISP network. The topology is made of two levels, each level being composed of three servers (all are reachable in anycast).

Until the beginning of March 2012, the whole LISP network was using LISP+ALT as mapping system. However, on March 14th 2012, it switched to LISP-DDT. This was mainly due to the growing overhead in maintaining LISP+ALT. As we will describe in the next section, from a research viewpoint, this gave us the possibility to perform measurements before, during, and after such an important switchover.

As depicted in Fig. 3, the ALT topology, in its very last setup, was decomposed in four regions (North America, Europe, Asia Pacific, and Latin America). The different ALT routers in a given region are connected in full-mesh with BGP over GRE tunnels and every region is connected to North America and Europe regions for a total count of eleven ALT routers. At the time of writing this paper, the LISP-DDT hierarchy is composed of six servers composing a two levels topology, as illustrated in Fig. 4. The root is made of three servers running on in Europe and USA and reachable via anycast addresses. The hierarchy is composed of only one level below the root, composed of three logical servers, also reachable in anycast.

4. LISP: STATE OF THE UNION

In this section, we describe our initial observations of the LISP network. We first describe our measurement methodology (Sec. 4.1), then, we present and discuss the early measurement results obtained for the mapping system (Sec. 4.2) as well as the mappings themselves (Sec. 4.3).

4.1 Measurements Campaign

Our campaign has been carried out from March 7th, 2012 to April 2nd, 2012 following two methodologies.

Firstly, we actively collected data using *LIG* (LISP Internet Groper [12]), querying the LISP mapping system every 15 minutes to obtain all known mappings. This is achieved by sending `Map-Request` for the targeted EIDs to the set of existing Map Resolvers [21]. At the time of writing this paper, twelve different Map Resolvers are deployed in the LISP Network. Therefore, for each targeted EID, all twelve Map Resolvers have been queried every 15 minutes. Each query is successful when a `Map-Reply` is received, containing the mapping (*i.e.*, EID prefix and RLOC set). Furthermore, LIG provides an estimate of the latency for the `Map-Request`/`Map-Reply` exchange. We run these periodic LIG requests from several vantage points, scattered in Eu-

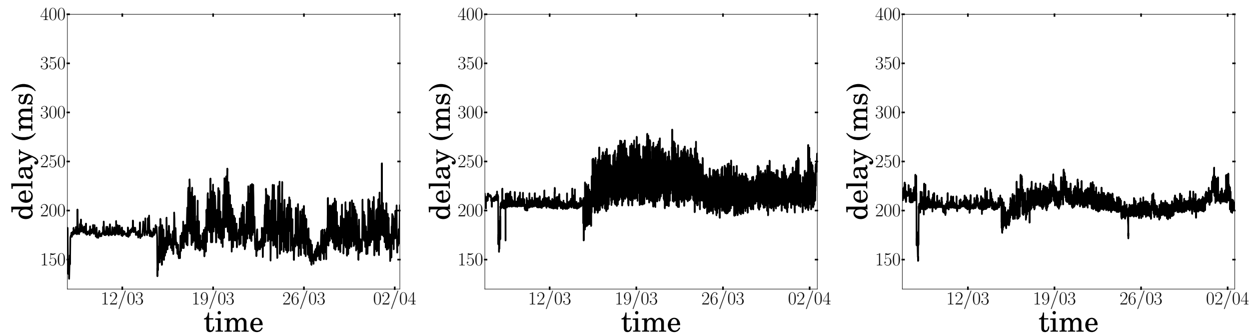

(a) PlanetLab (b) Internet (c) EID

Figure 5: Average Mapping System query latency evolution from three representative and different vantage points.

rope and North America, and either part of PlanetLab, the commercial Internet, or in the EID space. Sec. 4.2 discusses results obtained from this probing campaign.

Secondly, we rely on data gathered by the LISPmon project, a `lisp4.net` monitoring website maintained by the Universitat Politècnicà de Catalunya in Barcelona [13]. In particular, as we are interested in mapping information, we downloaded the daily list of mappings available and used them to feed our active LIG probing. Moreover, we downloaded the daily mapping snapshots from January 6$^{\text{th}}$, 2010 (the oldest snapshot available) to May 19$^{\text{th}}$, 2012 (the time of this writing), in order to have a long-term view on the evolution of the LISP network.[4] This dataset is discussed in Sec. 4.3.

As previously mentioned, during our measurement campaign, the mapping system underwent an important update to replace the LISP+ALT mapping system [11], based on BGP, by LISP-DDT [17] that has a hierarchical DNS-like design. Hence, the measurements cover as well the mapping system switchover.

4.2 Mapping System Performance

In the present section, we provide a glimpse of the dynamic aspect of the LISP mapping system, more specifically looking at the latency in obtaining mappings. The purpose is to characterize what is the performance in retrieving existing mappings, also considering the technology switchover occurred during our measurement campaign. An analysis of the mappings themselves will then be provided in Sec. 4.3.

Fig. 5 shows the evolution of the query latency (*i.e.*, the delay between sending a `Map-Request` and obtaining the corresponding `Map-Reply`) all along our measurements campaign, from three representative vantage points towards every EID. The three vantage points differ not only in their geographical position, but also in the type of network they are in. More specifically, the first vantage point is part of the PlanetLab infrastructure (hosted in USA), the second vantage point is an host located in France and connected to the commercial Internet, while the third one is actually a node part of the LISP network and belongs to the EID addressing space (hosted in Germany). The horizontal axis represents the time of the various LIG runs, while the vertical axis represents the average latency to obtain the different mappings.

While the delays are different for each vantage point, we can observe the same trend in their evolution: delays are quite stable, with relatively small variations, until March 15$^{\text{th}}$. Starting from that moment, while remaining in the same order of magnitude, just slightly higher on average, we can observe that the delay varies more from time to time. Interestingly, this is particularly true for the PlanetLab and the Internet vantage points (see Fig. 5(a) and Fig. 5(b)). However, the vantage point located in the EID addressing space (Fig. 5(c)) experiences a less severe impact with smaller variations.

The change in the dynamics and stability of the mapping query delay coincides with the LISP network mapping system switchover, from LISP+ALT to LISP-DDT. In LISP-DDT, every query normally traverses the entire hierarchy. However, to avoid this overhead, caches are used [22, 23, 24, 25]. We believe that the observed higher dynamics is related to this caching feature. It is worth to remark that, despite the performance degradation, LISP-DDT remains a better solution compared to LISP+ALT as the BGP overlay that was the cornerstone of LISP+ALT has proven to be extremely cumbersome to maintain as the network was growing [19, 20]. Hence, from an operational point of view, the improved manageability of LISP-DDT is a more important benefit than latency reduction.

4.3 LISP Network Evolution

The present section gives an overview on how mappings have evolved since January 2010 in number, type, and RLOC sets. Note that in the following results, we consider no changes happening in the middle of 2011, for which no snapshots are available (*cfr.* Sec. 4.1), thus observing flat curves in the corresponding period of the graphs.

In Fig. 6, we can observe the daily evolution of the absolute number of mappings. The number of mapping in January 2010 was as low as 20, and has consistently grown in the last two years to reach 80 mappings, four times the initial number, with a regular growth.

Actually, in the LISP infrastructure, beside what we simply call *mappings*, which are the ones binding EIDs to a set of RLOCs, there exist as well the so-called *negative* map-

[4]Due to technical problems of the LISPmon project, snapshots from the beginning of July 2011 to the beginning of August 2011 are missing.

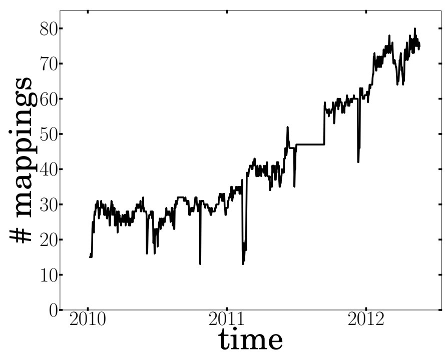

Figure 6: Evolution of the number of Mappings since January 2010

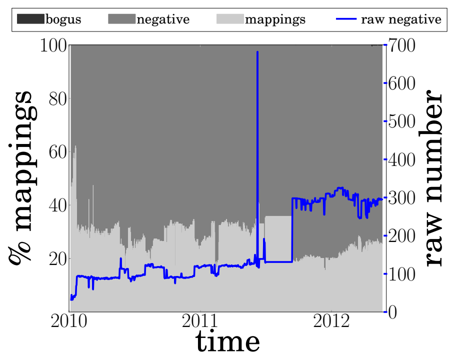

Figure 7: Evolution of the percentage breakdown of the type of mapping since January 2010

pings. These mappings have a different function. Given an IP prefix, if a negative mapping exists, it means that the IP prefix of the mapping is not part of the EID space and no RLOC set is associated to the prefix. The basic function of such a mapping is to tell the ITRs that the address is not part of the EID space, hence, not part of a LISP-enabled site, and that packets with destination address in this prefix have to be forwarded natively, without being encapsulated. Ideally the set of negative mappings should cover all the existing addressing space, which is not part of the EID space.

Fig. 7 shows the evolution of such a kind of mappings. In particular, the line represents the evolution of the raw number of negative mappings, to be read using the right-hand side vertical axis.[5] Fig. 7 also shows the percentage breakdown of the different type of mappings, indicated by the left-hand side vertical axis and the different gray areas. During the second part of 2011, the number of negative mappings doubled within two months from about 150 to about 300. Such an increase, which we are further exploring, can be related to a reorganization of the set of mappings, creating a more fragmented EID space. In the latest period of our measurement campaign we also observed malformed mappings, more specifically mappings that are not negative but have an empty RLOC set. We call these mappings *bogus*. The origin of these bogus mappings is still not clear and under further investigation. At this point, we are not able

[5]The reason of the peak in June 2011 remains unknown, but, we believe it was just a temporary misconfiguration.

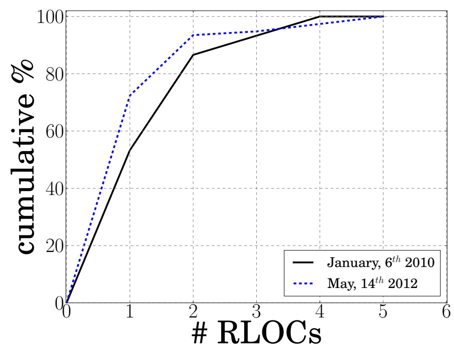

Figure 8: CDF of the number of RLOCs per mapping in January 2010 and May 2012

to state whether bogus mappings are due to buggy software deployed in the mapping system, or are the consequences of measurement software (that has not been developed by us) that is misreading some replies having a peculiar structure. In the former case, the issue may be critical, leading to important random lack of connectivity. In the latter case, it is a measurement bias that can be easily fixed. Nevertheless, all in all, we never observed so far more than a few bogus mappings in the same day, representing less then 0.5% of the total number of mappings (even if hard to see they are represented in the top right corner of Fig. 7).

As a last type of measurement, we looked at the distribution of the size of the RLOC set, *i.e.*, how many RLOCs are, on average, associated to an EID prefix. In particular, in order to provide evidence of its evolution, we calculated such a distribution for January 2010 and May 2012, obtaining the results shown in Fig. 8. Interestingly, despite the fact that the Internet is increasingly using multi-homing [26], the evolution of the number of RLOCs per EID prefix is contradictory. On the one hand, the percentage of mappings using two or less RLOCs has increased, *i.e.* more mappings use fewer RLOCs. On the other hand, we can observe that today there exists mappings with four RLOCs, which were not present in January 2010.

5. CONCLUSION

This paper provides a first measurement study of the Locator/Identifier Separation Protocol (LISP) deployment, assessing its evolution and performance. We have carried out measurements by querying the mapping system and analyzing available datasets. A first interesting take away is the fact that, when deploying a new protocol in the wild, sometimes manageability is more important than performance. This is demonstrated by the adoption of the more manageable LISP-DDT mapping system compared to the better performing LISP+ALT. Another important take away is the fact that the LISP network is constantly growing roughly at a rate of 100% every year. However, its traffic engineering features (*i.e.*, the possibility to use several RLOCs) remains largely unused. Future works aim at looking in more details to the collected datasets in order to provide a more thorough analysis, especially concerning the use of RLOCs. Are RLOCs of a single site geographically distant? Are they used for load balancing or backup? E.g., How they are exploited to take advantage of multi-homing? How much is reachability improved? These are a small sample of questions we plan to answer with our forthcoming works. This will allow evaluating the increasing usage of LISP for flexible traffic engineering. Furthermore, we expect to study LISP deployed in data centers, for traffic engineering but also for virtual machines mobility.

6. 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] A. Feldmann, "Internet clean-slate design: What and why?" *ACM SIGCOMM Computer communication Review*, vol. 37, no. 3, pp. 59–64, July 2007.

[5] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling innovation in campus networks," *ACM SIGCOMM Computer communication Review*, vol. 38, no. 2, pp. 69–74, March 2008.

[6] A. Feldmann, L. Cittadini, W. Muhlbauer, R. Bush, and O. Maennel, "HAIR: Hierarchical Architecture for Internet Routing," in *Proc. ACM CoNEXT ReArch Workshop*, December 2009.

[7] D. Clark, R. Braden, A. Falk, and V. Pingali, "FARA: Reorganizing the addressing architecture," in *Proc. ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, August 2003.

[8] R. Moskowitz and P. Nikander, "Host identity protocol (HIP) architecture," Internet Engineering Task Force, RFC 4423, May 2006.

[9] E. Nordmark and M. Bagnulo, "Shom6: Level 3 multihoming shim protocol for IPv6," Internet Engineering Task Force, RFC 5533, June 2009.

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

[11] 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.

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

[13] J. Domingo-Pascual, A. Cabellos-Aparicio, F. Coras, and L. Jakab, "LISPmon," see http://www.lispmon.net.

[14] B. Quoitin, L. Iannone, C. de Launois, and O. Bonaventure, "Evaluating the benefits of the locator/identifier separation," in *Proc. ACM/IEEE International Workshop on Mobility in the Evolving Internet Architecture (MOBIARCH)*, August 2007.

[15] 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.

[16] 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.

[17] 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-00.txt, July 2012.

[18] D. Farinacci, T. Li, S. Hanks, d. Meyer, and P. Traina, "Generic routing encapsulation (GRE)," Internet Engineering Task Force, RFC 2784, March 2000.

[19] T. Li, "Recommendation for a routing architecture," Internet Research Task Force, RFC 6115, February 2011.

[20] V. Fuller, "LISP-DDT: A new mapping database for LISP," June 2012, nanog 55 Talk.

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

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

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

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

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

[26] A. Akella, S. A., and R. Sitaraman, "A measurement-based analysis of multihoming," in *Proc. ACM SIGCOMM*, August 2003.