A Survey of Mapping Systems for Locator/Identifier Split Internet Routing

Michael Hoefling, *Graduate Student Member, IEEE,* Michael Menth, *Senior Member, IEEE,* and Matthias Hartmann

***Abstract*—The locator/identifier split is a core principle of many recently proposed routing architectures for a scalable future Internet. It splits the function of today's IP addresses into two separate pieces. End-hosts are addressed using identifiers which are not globally routable while network attachment points have globally routable locators assigned. In most architectures, either the sending host or an intermediate node has to query a mapping system to obtain locators for identifiers. Such a mapping system must be fast, reliable, secure, and may be able to relay data packets. In this paper, we propose requirements and a general taxonomy for mapping systems and use it to provide a survey on recent proposals. We address general aspects of mapping systems and point out remaining research opportunities.**

***Index Terms*—Locator/identifier split, scalable Internet routing, mapping system.**

I. Introduction

TODAY, IP addresses serve a double purpose. They are not only names that identify communication endpoints, but also routing locators that describe an endpoint's Internet attachment point. The coupling of both functions currently causes multiple problems in the Internet. An end-user network usually is assigned IP addresses from the IP number space of its Internet service provider (ISP). If the end-user network changes the ISP, it has to release the previously assigned IP addresses and obtain new addresses from a different IP number space of the new ISP. Thus, expensive renumbering of customer equipment is required [1]. If the network keeps its IP addresses while changing ISPs, the changed location or network attachment point in the Internet must be reflected in inter-domain routing. Multi-homing and traffic engineering also require additional entries and updates. This leads to increased BGP signaling rates, fragmented IP number space, and increased BGP routing tables. These issues raise scalability concerns for Internet routing.

The locator/identifier (Loc/ID) split principle addresses these problems and, in particular, the scaling issues in the Internet [2], [3]. Loc/ID split separates the above mentioned functions of current IP addresses into two different parts: the routing locator (RLOC) and the endpoint identifier (EID). A mapping system binds both address spaces.[1] When a node or a whole network changes its point of attachment to the Internet or performs traffic engineering [6], the mapping system is updated with the new EID-to-RLOC information.

From the application layer's point of view, only EIDs are visible and used to address packets. Depending on the routing architecture, either the source node or an intermediate node, e.g., a gateway router, queries the mapping system via a new layer in the IP stack to obtain an RLOC for the destination EID. This mapping node adds the RLOC to the packets to make them Internet routable, and may cache EID-to-RLOC mappings locally to avoid unnecessary future lookups.

If Loc/ID split becomes part of future Internet routing, mapping systems become a vital part of the Internet architecture, and must be secure and resilient to outages. When the mapping is performed by an intermediate mapping node, it may buffer or drop subsequent packets addressed to an EID that arrive before the EID-to-RLOC mapping is returned and stored in the cache. To avoid extensive delay and packet loss, the mapping system should provide a packet relaying function which temporarily forwards unresolved EID-addressed packets over the mapping system to the correct destination.

To the best of our knowledge, this paper is the first suggesting a taxonomy of mapping systems and providing a comprehensive survey on currently proposed mapping architectures, especially with regard to the above mentioned properties. We use the nomenclature of LISP (see Section II-B) where possible to unify the terminology. Table I provides an overview on often used abbreviations in this paper to ease comprehension. Additionally, the acronyms used in the taxonomy are summarized in Table II in Section IV-C to provide a quick lookup.

The paper is organized as follows. In the next section, we present a few examples of routing proposals implementing the Loc/ID split. Section III briefly analyzes requirements for mapping systems. Section IV proposes a taxonomy for mapping systems. We use it to classify and review recent mapping systems in Section V. In Section VI we address general aspects of mapping systems and point out remaining research opportunities. Section VII concludes this work.

Manuscript received 29 February 2012; revised 18 September 2012. This work was funded by the Federal Ministry of Education and Research of the Federal Republic of Germany (support code 01 BK 0800, G-Lab) and by the Deutsche Forschungsgemeinschaft (DFG) under grant ME2727/1-1. The authors alone are responsible for the content of the paper.

M. Hoefling and M. Menth are with the University of Tuebingen, Department of Computer Science, Chair of Communication Networks, Germany (e-mail: {hoefling,menth}@informatik.uni-tuebingen.de). M. Hoefling is the formal Corresponding Author for this submission.

M. Hartmann is with the University of Wuerzburg, Department of Computer Science, Chair of Communication Networks, Germany (e-mail: hartmann@informatik.uni-wuerzburg.de).

Digital Object Identifier 10.1109/SURV.2013.011413.00039

[1]There are other Loc/ID split solutions which do not require an additional mapping service, such as Shim6 [4], [5], a host-based multihoming technique for IPv6. As this paper focuses on mapping systems, these architectures are out of scope.

TABLE I
FREQUENTLY USED ABBREVIATIONS

Short form	Long form
Loc/ID split	locator/identifier split
EID	endpoint identifier
RLOC	routing locator
LISP	Locator/Identifier Separation Protocol
ETR	egress tunnel router
ITR	ingress tunnel router

II. ROUTING ARCHITECTURES BASED ON LOC/ID SPLIT

We briefly introduce the most prominent Loc/ID split routing architectures and categorize them according to the location of the mapping node. In the first class, the end-nodes themselves perform the mapping lookup. This usually implies upgrades to the hosts' network stack, but allows that mappings can be obtained before the communication flow is established. The second class consists of approaches where an intermediate node queries the mapping service and binds the RLOC to the EID.

A. Loc/ID Split with Mapping Lookup in Hosts

The Identifier/Locator Network Protocol (ILNP) [7], [8] implements the Loc/ID split in hosts. It may be based on both IPv4 and IPv6 as underlying network architecture. EIDs in ILNP can be either 32-bit or 64-bit. In case of ILNPv6, EIDs and RLOCs are directly embedded in the 128-bit IPv6 addresses through splitting the IPv6 addresses into two separate fields. The high-order bits serve as RLOC and the low-order bits as EID. In ILNPv4, this embedding is not possible because IPv4 addresses themselves are 32-bit. Thus, the RLOC is used in the IPv4 header and the EID is carried encapsulated in an additional header [9, Section 2.3]. ILNP assumes that nodes have upgraded network stacks and that applications use only DNS names to designate other devices. To communicate with them, a node queries the DNS for the EIDs and the RLOCs for DNS names, which is described in detail in Section V-D3.

The clean-slate Hierarchical Architecture for Internet Routing (HAIR) [10] implements Loc/ID split in hosts, too. It uses hierarchical addresses with three levels. The source node queries the mapping system and equips the packet with the complete RLOC information which can be seen as a kind of source routing. HAIR does not need middleboxes and is useful for mobility purposes.

With Hierarchical IP [11], hosts query the mapping information, insert it into packets, and intermediate nodes only rearrange the order of the information in the headers.

The Host Identity Protocol (HIP) [12]–[14] also implements the Loc/ID split. It introduces a Host Identifier Tag (HIT) as a location-independent designator of a node, which is used instead of its IPv6 address as identifier on the transport layer. This makes multihoming and mobility possible, and allows changes to the IPv6 address without any impact on the transport layer. A mapping service maps HITs to IP addresses. Note that the main focus of HIP is secure connection establishment and not routing scalability.

The Mobility and multihoming supporting Identifier Locator Split Architecture (MILSA) [15], [16] is another alternative Loc/ID architecture where hosts are responsible for the mapping lookup. The architecture comes with its own mapping system called Realm Zone Bridging Server (RZBS) which is presented in this paper as well. Interworking mechanisms are described in [17].

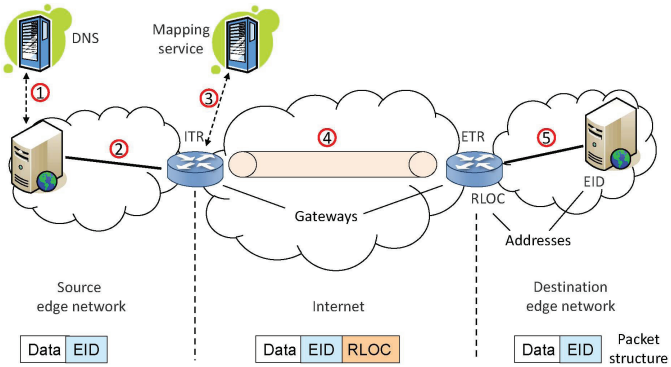

Fig. 1. Packet flow and destination addresses in LISP.

In all these protocols, the RLOC information is added in the end-hosts. This has the advantage that hosts can wait until the EID-to-RLOC mapping is available, e.g., until it is returned from the mapping system, and only then packets are sent. This is similar to the resolution of domain names to IP addresses before first packets of a flow are sent. Therefore, the DNS itself or DNS technology could be reused for the mapping system when hosts query the mapping information [18].

B. Loc/ID Split with Mapping Lookup in Intermediate Nodes

The Loc/ID Separation Protocol (LISP) [19], [20] implements the Loc/ID split without requiring host updates. LISP uses a special gateway for nodes within a LISP domain that performs the mapping lookup and makes packets globally routable. The DNS returns regular IP addresses as LISP-EIDs, which are routable only within LISP domains, but not in the global Internet. Two nodes in the same LISP domain can communicate with each other just like in the current Internet.

The communication between nodes in two different LISP domains is illustrated in Fig. 1. LISP domains are connected to the Internet through gateways that act as so-called ingress and egress tunnel routers (ITRs, ETRs). The xTRs (ITRs or ETRs) have globally routable addresses which are used as RLOCs for the EIDs hosted in their LISP domains. Before a hosts starts sending packets, it first resolves the host name of the destination host using DNS (step 1). When a host sends a packet to a locally unknown EID, the packet is default-forwarded to an ITR (step 2). The ITR retrieves the EID-to-RLOC mapping from its cache. In case of a cache miss, the ITR retrieves the mapping from the mapping system by sending a map-request and receiving a map-reply message (step 3). If the mapping is available, the ITR encapsulates the packet towards the respective RLOC (step 4). This is denoted as "map-and-encaps" operation. The packet can then be sent to the ETR of the destination LISP domain, where it is decapsulated, and eventually forwarded according to the EID to the destination node (step 5). Interworking techniques with the non-LISP Internet exist and are documented in [21]. In

this context, new mobility solutions have also been presented [22]–[24].

It is desirable to decouple the underlying routing architecture from the implementation specifics of the mapping system, especially because there could be several different mapping systems in the beginning of a Loc/ID split deployment phase. In LISP, ETRs are the only authoritative source for mappings. That is, any mapping system for LISP has to interact with ETRs to provide mappings. LISP-MS [25] proposes a LISP-specific interface for mapping systems so that they can interact with ITRs and ETRs in a standardized way. LISP-MS provides a map-resolver (MR) interface for ITRs and a map-server (MS) interface for ETRs that might be implemented in a separate box. ITRs send map-requests to MRs which inject them into the mapping system. ETRs register the EID prefixes they are responsible for at MSs. The MSs receive the map-requests and forward them to the ETRs. The ETRs send map-replies either directly to ITRs or to MRs. In the latter case, a MR can cache the mappings and also respond to map-requests with map-replies from the cache.

LISP is currently under standardization in the IETF as experimental standard and already deployed in pilot networks. The majority of other Loc/ID split proposals also uses intermediate nodes to add RLOCs to EID-addressed packets: Six/One router [26], GLI-Split [27], APT [28], a Novel DHT-Based Network Architecture for the Next Generation Internet [29], the Node Identity Architecture [30], RANGI [31], IVIP [32], [33], and IRON [34], [35].

In the remainder of the paper, we use ITRs as mapping nodes, but hosts may request mappings as well. The proposed mapping systems may be used for both types of architectures. We decided to use this nomenclature, along with the LISP terminology, to present the mapping systems in a unified way.

III. Requirements Analysis

Depending on whether the mapping lookup is performed in hosts or intermediate nodes, the requirements for a mapping system are slightly different. Therefore, a brief requirement analysis for mapping systems is given in this section.

A. Mapping Structure

The underlying future Internet routing architecture mandates the mapping structure to be supported by the mapping system. The granularity of EIDs and the way they are aggregated has a deep impact on the architecture. We will discuss these structures in Section IV-A.

B. Scalability

The Internet has been growing fast in the last three decades in terms of hosts and in terms of networks. We expect this process to continue due to the increasing ubiquity of Internet applications. Therefore, the number of EIDs and/or EID-prefixes will also increase in the future. As we could not foresee the tremendous growth of the Internet in the past, mapping systems must be able to handle a similar growth in the future.

C. Resilience

Today's Internet already uses a mapping service: the DNS. If the DNS fails, the reachability of other end systems is strongly compromised unless the user knows their IP addresses in addition to their DNS names. If the mapping service for a Loc/ID split routing architecture fails, there is mostly no way for the user to reach other end systems as the user knows only their EIDs. Therefore, the EID-to-RLOC mapping system must not fail. As a consequence, resilience requirements for the EID-to-RLOC mapping system are clearly stricter than for the DNS.

D. Security

For the same reasons as above, an EID-to-RLOC mapping system has to provide at least the same security as the existing DNS. It has to withstand direct and indirect security threats. An example of direct threats is denial-of-service attacks. These types of attacks are difficult to prevent. Examples of indirect threats are takeover of authoritative mapping sources and cache poisoning which happens when a mapping node accepts wrong mappings. A carefully designed security model should address as many threats as possible to minimize targets for attackers, and it should be extendable.

E. Relaying

ITRs may locally cache mappings including an expiration date to reduce the request loads [36], [37] for any type of mapping system. During regular operation of a mapping system, cache misses can occur. A cache miss means that requested information for a certain mapping has not been fetched from the authoritative source, yet. It is advantageous if the mapping system offers a relay service so that mapping nodes can forward packets with missing RLOCs over the mapping system to another node that can forward the packet to the destination. This avoids packet loss and extensive delay, but can lead to packet reordering when subsequent packets are forwarded directly.

IV. Taxonomy of Mapping Systems

The architecture of a potential mapping system depends on the structure of EIDs and mappings. We explain the different structures of both, propose a taxonomy of mapping systems and discuss the properties of the proposed classes on an abstract level.

A. EID and Mapping Structure

EIDs in the Loc/ID split context should be globally unique. Their uniqueness can be achieved through administrative or statistical means. IP or Ethernet addresses are examples for the first category. Numbers authorities like IANA or IEEE assign EID address prefixes to organizations that may further partition their obtained address space and re-assign sub-prefixes to other organizations or directly to nodes. This leads to hierarchically structured EIDs. As an alternative, EIDs may be randomly created like in HIP [12]. If they are sufficiently long, the probability for the creation of the same EIDs is very small. These EIDs are unstructured and we call their address space flat. Also semi-structured addresses have been proposed,

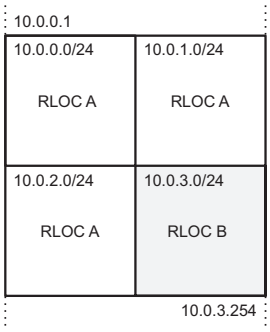

Fig. 2. Four EID-prefixes covered by $10.0.0.0/22$ as an example for aggregation.

which combine hierarchically assigned prefixes and random suffixes [14], [31], [38].

A set of unstructured EIDs cannot be aggregated by a common prefix. Therefore, each EID in the set needs its own EID-to-RLOC mapping, even if all of them are located behind the same RLOC. When structured EIDs with a common prefix have the same RLOC, the mapping information can easily be aggregated to an EID-prefix-to-RLOC mapping. This is attractive as it saves memory and communication overhead.

There are four types of mapping aggregation. No mapping aggregation, normal mapping aggregation, mapping aggregation with exceptions, and mapping aggregation with more specifics. Fig. 2 shows four IPv4-like EID-prefixes which are covered by the shorter EID-prefix $10.0.0.0/22$. The prefixes $10.0.0.0/24$, $10.0.1.0/24$, and $10.0.2.0/24$ are located behind RLOC A. The prefix $10.0.3.0/24$ is located behind RLOC B. If no mapping aggregation is used, the mapping system has to store four mappings.

If normal aggregation is used, then the prefixes $10.0.0.0/24$ and $10.0.1.0/24$ are aggregated to $10.0.0.0/23$. Normal aggregation means that there is no overlap between two stored prefixes. As a consequence, three EID-to-RLOC mappings are needed to cover the mappings given in the example.

The authors of [39] propose to maximize EID-block aggregation to maximum aggregatable EID blocks. That means, blocks are aggregated if most of their EIDs share the same RLOC and if only a few sub-blocks of them exhibit different or no RLOCs. The mapping of the maximum aggregatable EID block is accompanied by exceptions that indicate these holes. This approach requires two mappings in our example. The three prefixes sharing RLOC A are aggregated to $10.0.0.0/22$ with exception $10.0.3.0/24$ and another mapping assigns RLOC B to $10.0.3.0/24$.

Mapping aggregation with more specifics also aggregates EID blocks if most of the contained EIDs share the same mappings. Instead of providing exceptions for the EIDs with another RLOC, mappings with more specific prefixes are provided. That means, when a mapping is requested for a certain EID, several mappings may be returned by the mapping system but only the one with the longest prefix match is valid for that EID [19, Section 6.1.5.]. In our example, this approach maps the prefix $10.0.0.0/22$ to RLOC A and the more specific prefix $10.0.3.0/24$ to RLOC B so that two mappings are needed.

When more specifics are allowed, the mapping system returns the most specific mapping for a requested EID. In addition, it returns all other more specifics that fall into the prefix range of the returned mapping. We show in an example that this is needed to correctly support caching. Assume that the mapping for an EID of the $10.0.0.0/24$ block has been requested. According to the above rule, $10.0.0.0/22 \rightarrow$ RLOC A and $10.0.0.3/24 \rightarrow$ RLOC B are returned and stored in the cache. When another EID from the $10.0.0.3/24$ address space also requires a mapping, its most specific mapping is already available in the cache. Without delivering and caching $10.0.0.3/24 \rightarrow$ RLOC B, the ITR needs to issue another map-request for another EID of $10.0.0.3/24$. Without the above policy, one could never be sure in case of a cache hit that the most specific mapping is found so that another map-request needs to be issued which essentially erases the advantage of caching. Another consequence of more specifics are that they must not be purged from the cache before least specific mappings. This makes cache management more complex.

B. Mapping System Structure

We define a map-base (MB) as a node or distributed system that is the authoritative source of EID-to-RLOC mappings. In general, MBs must have globally reachable RLOC addresses so that they can be contacted without another mapping lookup.

1) Full Knowledge vs. Partial Knowledge: There are basically two options to implement MBs. One option is to implement a central MB to store the EID-to-RLOC mappings for all existing EIDs in a single MB with full knowledge (MBFK). As an alternative to MBFKs, EID-to-RLOC mappings may be stored in distributed MBs each of which holds only partial knowledge (MBPK).

MBFKs may be replicated to multiple mirrors for resilience and load-balancing, and to bring the MBs closer to the ITRs. An ITR needs to be configured with the address of at least one MBFK. It sends map-requests (step 1) and receives replies (step 2) directly from that MB, as visualized in Fig. 3(a). Changes of EID-to-RLOC mappings lead to a large amount of update traffic and frequent changes in the MBs as any change affects any MB. Therefore, frequent mapping changes should be avoided. Packet relaying with MBFKs is simple. In case of a cache miss, ITRs can immediately relay packets to the MBFK. The receiving MB can serve as a proxy ITR and encapsulate the received data packets towards their destinations.

In contrast, MBPKs help to keep mapping updates local and possibly facilitate mobility support using Loc/ID split [22], [23]. They may be operated on behalf of EID owners or on behalf of networks or autonomous systems (ASes) and are responsible only for their EIDs. This approach also has the advantage that only local MB operators have control over the mappings.

2) Discovery Options for MBPKs: When an ITR needs a mapping for a certain EID, it needs to know which MBPK to query. There are several discovery options to find the correct MBPK.

a) Local Lookup (MBPK-LL): The fastest option to find the appropriate MB is to configure ITRs with a table that

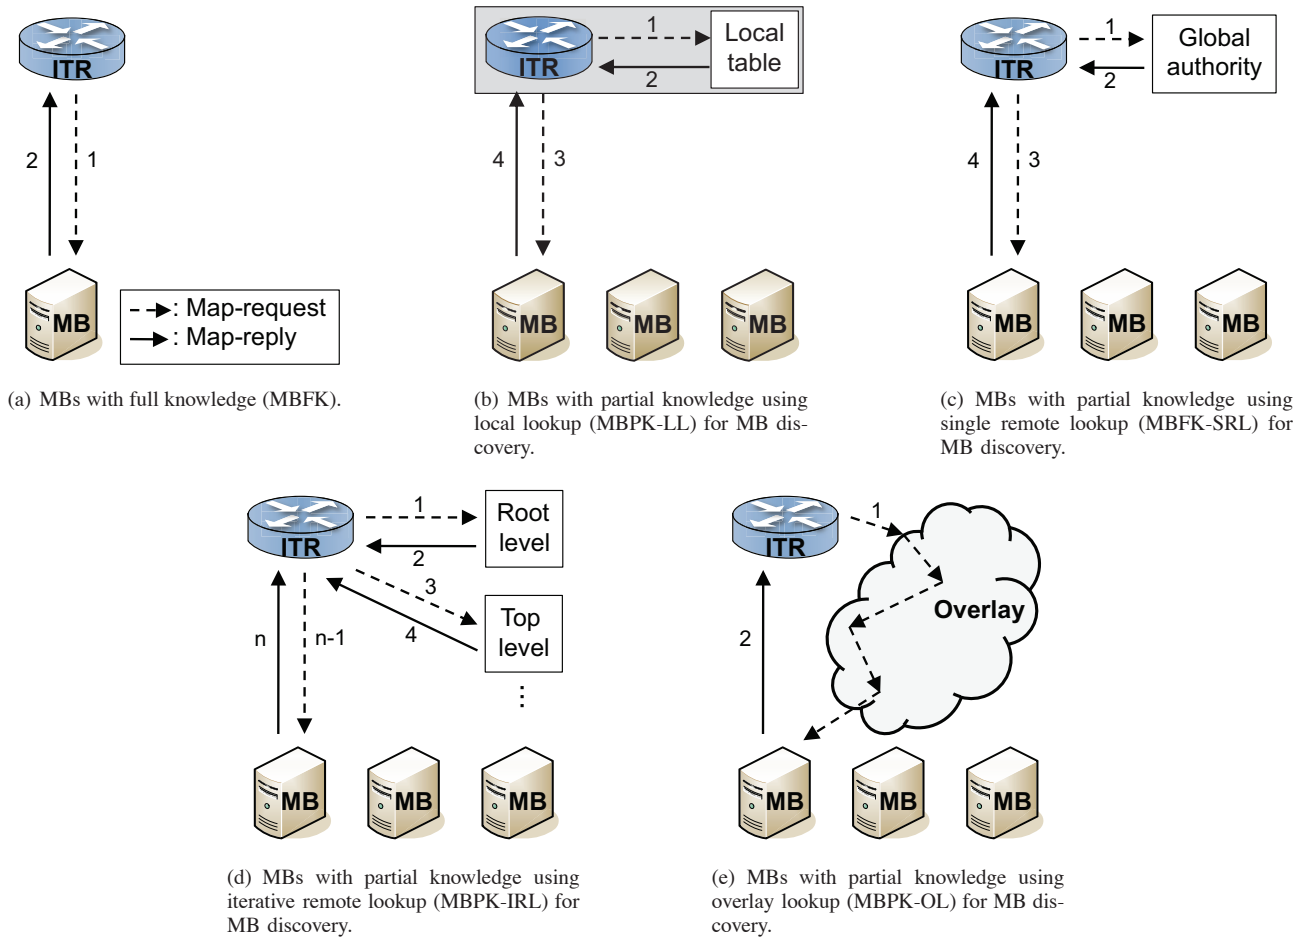

(a) MBs with full knowledge (MBFK).

(b) MBs with partial knowledge using local lookup (MBPK-LL) for MB discovery.

(c) MBs with partial knowledge using single remote lookup (MBFK-SRL) for MB discovery.

(d) MBs with partial knowledge using iterative remote lookup (MBPK-IRL) for MB discovery.

(e) MBs with partial knowledge using overlay lookup (MBPK-OL) for MB discovery.

Fig. 3. Structure and operation of different mapping system classes.

stores EID(-prefix)-to-MB information. It points to the MB that stores EID-to-RLOC mappings for certain EID-prefixes. This is visualized in Fig. 3(b). Only one local lookup suffices to find an appropriate MB (step 1 and 2). The actual EID-to-RLOC mapping is obtained with a single query over a direct path between ITR and MB which keeps the mapping delay short (step 3 and 4).

The EID(-prefix)-to-MB information in the local tables of ITRs is relatively stable so that the update load is low [40]. The challenge is the construction and the distribution of the EID(-prefix)-to-MB tables, and to keep them up to date. To keep the table small, mapping aggregation is crucial, i.e., EID-to-RLOC mappings for EIDs from a common prefix should be provided by the same MB.

With MBPK-LLs, a packet relaying service can easily be implemented. After a cache miss, the ITR locally looks up the address of the appropriate MB to send a map-request. The ITR may then relay packets to this MB which stores the matching EID-to-RLOC mapping. The MB can then forward the packet like an ITR.

b) Single Remote Lookup (MBPK-SRL): Instead of a local table, a global authority can be used. In this approach, the knowledge about which MB is responsible for which EIDs is maintained and stored by a global authority as shown in Fig. 3(c). To retrieve an EID-to-RLOC mapping for a specific EID, an ITR must first query the global authority for an EID(-prefix)-to-MB mapping (step 1 and 2) and then send a map-request to the respective MB (step 3 and 4).

Single remote lookups for MB discovery introduce more communication overhead than local lookups, but ITRs may store the obtained EID(-prefix)-to-MB information in a cache to reduce communication overhead and delay for future requests. This is similar to the caching of EID-to-RLOC information. EID(-prefix)-to-MB information is expected to change less frequently than EID-to-RLOC information so that this information can be cached for a longer time.

With MBPK-SRLs, the implementation of a packet relaying service is problematic because packets need to be stored or dropped until the ITR has retrieved the EID-to-MB information from the global authority. As soon as the ITR knows the appropriate MB, it can relay packets to the MB that may forward the packets to their destination.

c) Iterative Remote Lookup (MBPK-IRL): The MB that is responsible for a certain EID may be found iteratively, i.e., in a similar way like the Domain Name System (DNS) finds authoritative name servers. The abstract architecture is shown in Fig. 3(d). A level-0 authority (root level in DNS) returns a EID(-prefix)-to-level-1 mapping to the ITR (step 1 and 2). A level-1 authority (top level in DNS) returns a EID(-prefix)-to-level-2 mapping to the ITR (step 3 and 4). This procedure continues iteratively until, eventually, the EID-to-level-n mapping designates the actual MB (step $n-1$ and n). Again, caching of entries is possible to save communication overhead for future map-requests.

Packet relaying is difficult with MBPK-IRLs for the same reasons as for MBPK-SRLs. A slow, remote query is required to find out the appropriate MB. During this time, packets need to be stored or dropped by the ITR.

d) Overlay Lookup (MBPK-OL): The last proposed class uses overlay networks to find appropriate MBs. Fig. 3(e) visualizes the basic concept. A map-request is sent into an overlay network where it is forwarded to the appropriate MB (step 1) which responds to the ITR with a map-reply (step 2). Each ITR must be configured with at least one entry node in the overlay network. There are many different implementation options for the overlay network, which are classified in the next section (Section IV-B3).

The resolution delay of MBPK-OL is rather large by design for two reasons. Map-requests are sent over possibly multiple hops within an overlay network to the appropriate MB instead of using the direct path as other mapping systems do. Moreover, most MBPK-OL implementations require that overlay nodes process the map-requests on the application layer to forward them to the appropriate next hop which is time-consuming. As a consequence, transport over the overlay network is usually much slower than over a direct path.

In most MBPK-OLs, a packet relaying service can be implemented. After a cache miss, ITRs may send EID-addressed packets through the overlay network where they eventually reach a MB that can forward the packets to the destination. When first packets are relayed over the overlay network and subsequent packets are tunneled by the ITR, packet reordering can occur due to the large difference in the transportation time over the different paths. This may cause problems for some applications.

The overlay network is a vital part of the mapping system and should be run on a trusted infrastructure. Operators of nodes participating in an MBPK-OL may control transiting sensitive traffic from other participants. This could be a threat to participants whose map-requests are carried over nodes that do not have at least indirect business relations with them. They require that these nodes reliably process and forward their map-requests. Carrying traffic from participants with whom they do not have business relations is also a burden for the operator of an overlay node. The node is expected to process and forward the requests without receiving revenues from them although the data rate may be high, especially if also packets are relayed. Thus, new business models are needed for the deployment of MBPK-OLs.

Parts of the overlay network can fail or be attacked. Since customers cannot simply increase the availability of the overlay network by replication, MBPK-OLs require special backup concepts to avoid service degradation in failure cases.

Despite these shortcomings, several MBPK-OLs have been proposed in the past. Many of them do not need an infrastructure that is managed by a global authority, which is very appealing especially in the prototype stage.

3) Lookup Overlays for Discovery of MBPKs: In the following, we describe different approaches to implement MBPK-OL.

a) Hierarchically Structured Overlay (MBPK-HSO): In a hierarchically structured overlay network, nodes represent EID-prefixes and are arranged in a hierarchical manner with regard to these EID-prefixes. The MBs are connected to the nodes with the most-specific EID-prefixes. Fig. 4(a) visualizes the forwarding of a map-request in such a overlay network. At first, map-requests travel up the hierarchy if needed (step 1 and 2) and then down towards the nodes with the most specific EID-prefixes over which they finally reach the appropriate MB (step 3 through $n-1$). In step n, the MB answers the query to the ITR. Forwarding of map-requests in MBPK-HSOs may be done iteratively, recursively or in a hybrid fashion. If map-request forwarding is strictly iterative, MBPK-HSOs are similar to MBPK-IRLs.

b) Distributed Hash Tables (MBPK-DHT): A distributed hash table (DHT) consists of connected MBs and uses a hash function that determines at which MB the EID-to-RLOC mapping for a special EID is stored. An example for such an architecture is visualized in Fig. 4(b). When map-requests are injected into the DHT (step 1), each node knows how to forward them to neighboring nodes (step 2 through $n-2$) so that the map-request eventually reaches the appropriate MB (step $n-1$). Again, in step n, the MB answers the query to the ITR.

c) Multicast Overlay (MBPK-MCO): In a multicast overlay network, the EID space is partitioned and multicast groups are created for each of the EID-subsets. MBs with EID-to-RLOC mappings subscribe to all groups which cover one of their EIDs. This is illustrated in Fig. 4(c). The ITR is configured with all multicast groups and sends map-requests to the multicast group the EID for which the RLOC is needed belongs to (step 1). Thus, the map-request is carried to all MBs containing EIDs for the same multicast group (step 2). The MBs with the matching mapping send a map-reply back to the ITR (step 3).

TABLE II
ABBREVIATIONS USED IN THE TAXONOMY

Short form	Long form
MB	map-base (authoritative source for a mapping)
MBFK	map-base with full knowledge
MBPK	map-base with partial knowledge
MBPK-LL	MBPK with local lookup
MBPK-SRL	MBPK with single remote lookup
MBPK-IRL	MBPK with iterative remote lookup
MBPK-OL	MBPK with overlay lookup
MBPK-HSO	MBPK with hierarchically structured overlay
MBPK-DHT	MBPK with distributed hash table
MBPK-MCO	MBPK with multicast overlay

C. Summary

Fig. 5 shows a summary of our proposed taxonomy and Table II provides an easy to use overview on the acronyms. Mapping systems consist of MBs with either full or partial knowledge. For the latter, we distinguished the way how ITRs determine the appropriate MB for a given EID. The options are local lookup, single remote lookup, iterative remote lookup, and overlay lookup. The overlay network may be a hierarchically structured overlay, a DHT, or a multicast overlay. This taxonomy does not claim to be complete, but it classifies well the existing proposals that we review in the next section.

V. COMPARISON OF MAPPING SYSTEMS

We review mapping systems that were presented in the context of Loc/ID split based Internet routing. We classify

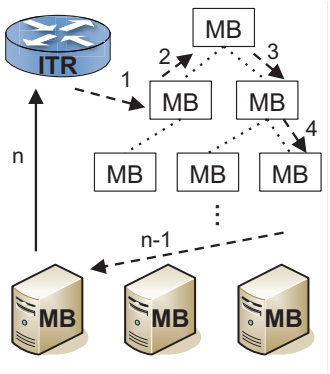

(a) MBs with partial knowledge using hierarchically structured overlays (MBPK-HSO) for MB discovery.

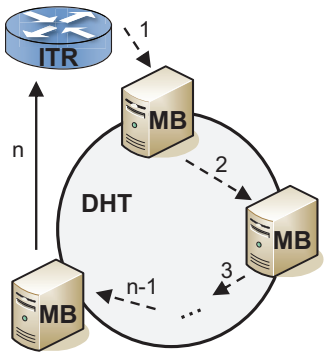

(b) MBs with partial knowledge using distributed hash tables (MBPK-DHT) for MB discovery.

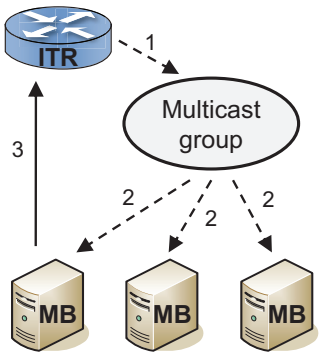

(c) MBs with partial knowledge using multicast overlays (MBPK-MCO) for MB discovery.

Fig. 4. Special cases of lookup overlays.

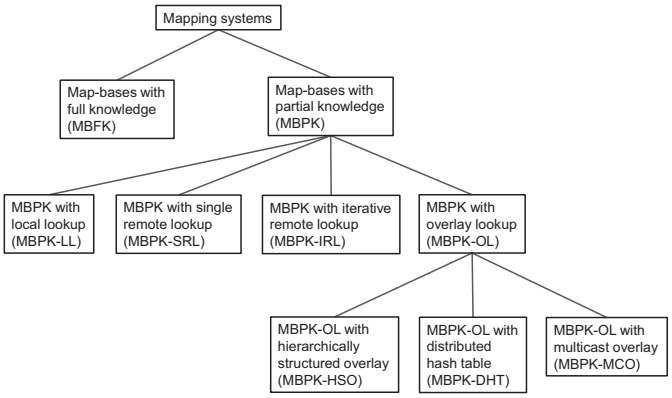

Fig. 5. Hierarchical taxonomy of mapping systems.

them into the categories presented in Section IV and discuss their properties.

A. Map-Bases with Full Knowledge (MBFK)

1) LISP-NERD: The "Not-so-novel EID to RLOC Database" for LISP (LISP-NERD) [41] is a mapping system that is primarily designed to avoid packet drops. To achieve this, mapping information is distributed to all ITRs in advance. One or several authorities assign EIDs to organizations and run a MB (called NERD) with authoritative mappings. An ITR is configured with the addresses of possibly several authoritative NERDs and pulls the entire mapping information from them upon system start. To facilitate incremental updates, changes to the NERD are associated with a version number and a change file. ITRs regularly poll the NERDs for their latest version numbers and download and apply the change files to their local database if needed. All information sent from the NERDs to the ITRs is digitally signed using X.509 certificates. As all mappings are locally available at the ITRs, cache misses cannot occur. Querying delay and packet loss cannot happen so that a packet relay service is not needed.

The initial assumption while constructing LISP-NERD was that all EIDs of an assigned prefix have the same EID-to-RLOC mapping so that NERDs store in fact EID-prefix-to-RLOC mappings. If finer mapping granularity is needed, e.g., due to mobile nodes, then the NERDs need to store significantly more mapping entries which raises scalability concerns.

2) APT: "A Practical Tunnelling architecture" (APT) [28] uses tunnelling from ITRs to ETRs, a mapping distribution system [42], and handles failures of ETRs.

Like in LISP, EID prefixes are mapped to RLOCs. APT's mapping system assumes that each ISP has a default mapper (DM), i.e., a mirror containing the global mapping information. DMs of neighboring ASes know each other and exchange mapping information via a mapping dissemination protocol using signed messages. The prefix owners inject the mapping information into the DMs of their ISPs. Whenever new information is available, DMs push it to their neighboring DMs. When an ITR encounters a cache miss for a packet destined to an unknown EID, the ITR sends the packet to the DM of its own domain. The DM chooses a single RLOC, returns the appropriate mapping to the ITR, and tunnels the packet to an appropriate ETR. Thus, APT provides a packet relay service.

APT provides a reroute mechanism for the failure of ETRs. The assumption is that ETRs reside within the AS and do not serve as border routers. The DM in the destination network announces a high-cost route towards every ETR it serves via IGP so that traffic is rerouted to the DM if an ETR fails. The DM re-encapsulates the traffic towards an alternative ETR so that the traffic eventually reaches its destination. In addition, the DM of the destination network notifies the DM of the source network about the ETR failure so that the source DM can suggest its ITRs to encapsulate further packets towards alternative ETRs.

B. Map-Bases with Partial Knowledge using Local Lookup (MBPK-LL)

FIRMS is a mapping system for future Internet routing [40].

It assumes that the EID space is partitioned by IANA and delegated to the five regional Internet registries (RIRs) which further partition and delegate it to local Internet registries (LIRs) from which organizations (prefix owners) receive an EID prefix. The basic structure and operation of FIRMS is shown in Fig. 6. Each prefix owner runs a MB or mandates

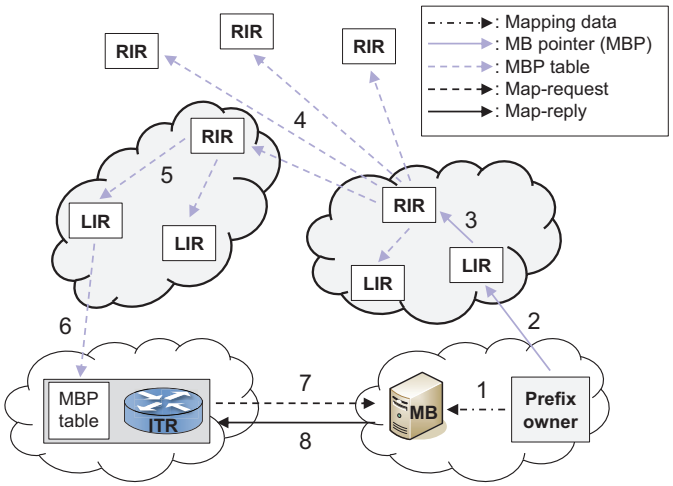

Fig. 6. In FIRMS, the global MBP table is incrementally pushed to ITRs which use it to send map-requests immediately to responsible MBs.

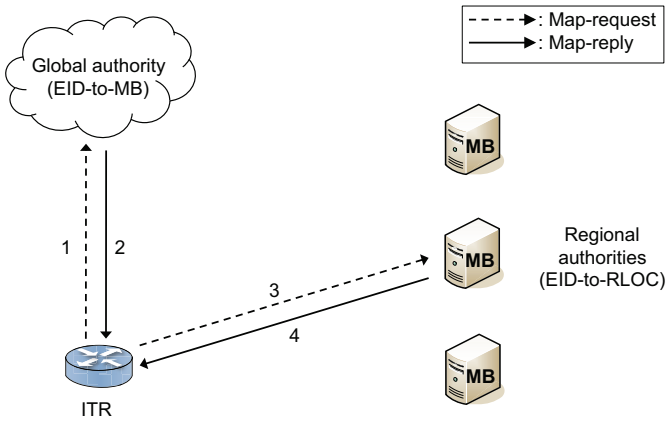

Fig. 7. In HiiMap, the ITR first queries the global authority if the required EID-to-MB mapping is not in its cache; then it queries the regional authority for the EID-to-RLOC mapping.

a company to run it on its behalf. ITRs have a table with the global EID-prefix-to-MB information to look up the RLOCs of appropriate MBs for a query. This table is filled in a distributed way. The prefix owner registers EID-prefix-to-RLOC information at the MB (step 1) and additionally a MB pointer (MBP) at the LIR from which it receives its EID prefix (step 2). This MBP essentially contains the EID-prefix-to-MB information. The LIR propagates this information to its RIR (step 3), the RIR pushes it to the other RIRs (step 4), and they push it to their subordinate LIRs (step 5). Thus, LIRs have a global MBP table. ITRs subscribe to LIRs to download a copy of the MBP table and to receive updates of it (step 6).

FIRMS can relay packets as described in Section IV-B2a. It achieves resilience by replicating system components and using backups when primary elements fail. The security concept is based on X.509 resource certificates similar to LISP-NERD and an additional public-key infrastructure (PKI) for MBs. An important property is that ITRs can validate map-replies locally and do not need to verify a trust chain before they can use an obtained mapping. Thus, the mapping lookup in FIRMS is fast since ITRs can immediately issue a map-request that is carried on the direct path to the MB (step 7 and 8), and the security features of FIRMS add only little and predictable delay to validate the map-reply.

C. Map-Bases with Partial Knowledge using Single Remote Lookup (MBPK-SRL)

HiiMap [43] is a "Hierarchical Internet Mapping Architecture" and the only mapping system that falls in this category. HiiMap assumes that EIDs are under the control of a region which may be, e.g., a country. Fig. 7 shows that a single global authority stores an EID-to-MB mapping, a so-called "regional prefix" for each EID. It points to a regional authority which has a MB that stores all EID-to-RLOC mappings for all EIDs under its control. An ITR queries the global authority for the "regional prefix" (step 1) and after its reception (step 2) the ITR queries the regional authority for the EID-to-RLOC mapping (step 3 and 4).

HiiMap can support flat EID spaces as it does not take advantage of EID prefix aggregation. However, this feature makes it only as scalable as MBFKs. The global authority stores a regional prefix per EID which leads to large storage requirements. Furthermore, the global authority may be a performance bottleneck for updates and requests. If the ITR has no regional prefix for an EID in its local cache, it must query both the global and the regional authority to obtain the EID-to-RLOC mapping which leads to increased lookup delay.

D. Map-Bases with Partial Knowledge using Iterative Remote Lookup (MBPK-IRL)

DNS has been proven to be a powerful and scalable architecture, but it has not been secure. Security has recently been added [44] and clients trust the received data when they are signed by the authoritative DNS server. However, if the client does not trust the public key of the authoritative DNS, it must first validate that key before it can validate the actual data. Thus, the client needs to iteratively validate the trust chain up to a common trust anchor. This adds delay to the mapping lookup. None of the following MBPK-IRLs support packet relaying.

1) One-Phase Lookup Using Reverse DNS: The one-phase lookup as presented in [45, Scheme 1] assumes that the reverse DNS reply contains both a pointer resource record (PTR-RR) and an A-RR for the ETR of the queried IP number, i.e., an RLOC for the requested EID. Normally, A-RRs provide IPv4 addresses for the DNS names. In this context, they provide the IPv4 RLOCs of the ETRs for the EIDs. The prefix owner can set up an authoritative DNS server returning the A-RR with the RLOC information for his EID prefix and register the address of this delegation server with the authority from which it has received his EID prefix. Thereby, the prefix owner has still control over the mappings. This idea has been sketched for the LISP context in [45] and in [46]. However, it did not prevail since the existing DNS infrastructure should not be burdened with another heavy service. Moreover, for many people this approach did not seem sufficiently robust and powerful to be applied as a mapping system in a Loc/ID split context where intermediate nodes query the mappings.

2) Two-Phase Lookup Using Reverse DNS: Intermediate nodes may perform a two-phase lookup to retrieve an EID-to-RLOC mapping. An ITR first makes a reverse lookup to get the DNS name for an EID (which must be an IP address), and then it makes a forward lookup for the RLOC of that DNS name. This requires the use of reverse DNS, a PTR-RR that provides a DNS name for each EID, and the definition of a new locator resource record (L-RR). However, this obvious solution is quite slow since it requires two lookups, which is a problem for the lookup of EID-to-RLOC mappings by intermediate nodes. In addition, the approach raises security concerns. An attacker can register "BadGuy.com" in the DNS for an EID from "GoodGuy.com" together with a bad RLOC. The reverse lookup for the EID yields both "BadGuy.com" and "GoodGuy.com". When "BadGuy.com" is selected, the bad RLOC is returned and the EID of "GoodGuy.com" can be hijacked.

3) Mapping Lookup using only DNS Names: The Identifier/Locator Network Protocol (ILNP) [7], [47] defines four new resource record (RR) types for DNS [48]: ID, L32, L64, and LP. The ID-RR stores an EID for a special DNS name. RLOCs are stored as L32-RRs and L64-RRs, depending on whether ILNPv4 or ILNPv6 locators are to be retrieved. LP-RR stands for locator pointer RR and is used to provide multihoming and mobility in ILNP [49].

Hosts use the fully qualified domain name to retrieve the ID-RR from the DNS. With the ID-RR, hosts further query the DNS to retrieve the corresponding L32-RR or L64-RR. With this information, the hosts compose the appropriate IP numbers. In case of multihoming or mobility, the ID-RR is used to retrieve the LP-RR which is further used and resolved to L32-RRs or L64-RRs.

This principle works well with ILNP as the host adds locator information to destination addresses. Other routing architectures where intermediate nodes add locator information cannot adopt this principle as the lookup requires a DNS name which is not contained in the IP packet header.

4) Use of DNS for HIT-to-IP Lookup in HIP: HIP requires a mapping system to find an IP address for a given HIT. The authors of [50] propose to use the DNS system to find the IP addresses for HITs. For reverse DNS, the authors of [51] postulate a "hit-to-ip.arpa" domain in which HITs are denoted like IPv6 addresses within "ipv6.arpa". Since HITs are not hierarchically structured, all HITs need to be known by top-level servers that are run by authorities. The authors give evidence that DNS servers are powerful enough for their purpose. Since improved mobility is an objective of HIP, HIT-to-IP mappings are likely to change often. As updates of DNS records take orders of magnitude longer than retrievals, a two-level hierarchy is introduced. The entries in the top-level DNS servers just refer to second-level DNS servers. These entries are likely to stay the same for long time. As a result, top-level servers experience fewer updates which reduces the infrastructure expenses for authorities. This also provides direct control over the actual HIT-to-IP mapping to the HIT owner which is important to support mobility.

5) LISP-TREE: LISP-TREE [52] makes use of DNS technology and stores EID-to-RLOC mappings in a DNS-like fashion, but runs on a different infrastructure. It assumes that

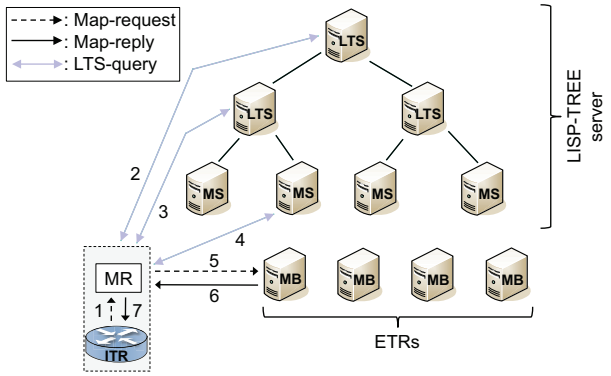

Fig. 8. In LISP-TREE, the ITR resolves the address of the responsible MB like in DNS and then sends a map-request.

the EID address space is partitioned among regional EID registrars (RERs) which allocate parts of their EID space to local EID registrars (LERs). LERs further allocate EID space to other LERs or customers. To be compliant with LISP, EID-to-RLOC mappings are stored by authoritative ETRs which serve as MBs for EID-prefixes.

LISP-TREE uses a tree-like overlay structure of LISP-TREE servers (LTSs) as illustrated in Fig. 8. They run a DNS service for EIDs and assist ITRs to find the authoritative MB for a given EID. The root LTSs are run by the RERs and store pointers to the LTSs for their /8 prefixes (at most 256). Lower level LTSs are managed by the corresponding LERs and hold pointers from more specific EID-prefixes to lower level LTSs that are responsible for a smaller subset of the EID address space defined by the delegated prefix. LISP-TREE uses the generic LISP-MS mapping system interface so that MSes constitute the leaves of the LTS tree.

ITRs are configured with the root LTSs and iteratively or recursively query LTSs to eventually receive the address of the correct MB (step 1 through 4). Then, they query it to get the EID-to-RLOC mapping from the MB (step 5, 6, and 7). Intermediate results from LTSs are cached so that the ITR must query the root LTSs only rarely. The authors have shown that only the iterative mode is scalable. Security in LISP-TREE is provided by the use of DNSSEC [44]. The layered mapping system (LMS) discussed in [53] and presented in [54] is very similar to LISP-TREE and therefore not further discussed in this paper.

6) LISP Delegated Database Tree: LISP delegated database tree (LISP-DDT) [55] is a hierarchical mapping system and the logical successor of LISP-TREE. In LISP-DDT, LTS are called DDT nodes and do not use DNS technology. However, the abstract structure of LISP-TREE is preserved, i.e., the DDT nodes form a hierarchical EID-prefix tree. At the lowest level of the database tree, DDT map-servers hold EID-prefix-to-MB mappings. DNS security mechanisms cannot be used directly because LISP-DDT does not use DNS technology. However, security is provided using pre-shared keys between DDT nodes, which works similar to DNSSEC and LISP-SEC [56] mechanisms. LISP-DDT in combination with LISP-MS is currently the preferred mapping system for LISP.

7) Distributed Real Time Mapping System for IVIP and LISP: IVIP [32], [33] is an alternative to LISP and has its own "distributed real time mapping system" (DRTM) [57]. The EID space is partitioned in mapped address blocks (MABs) by MAB operating companies (MABOCs). The resulting MABs are assigned to user organizations. The user organizations can further partition their MABs into micronets which are arbitrarily long EID prefixes.

MABOCs store EID(-prefix)-to-RLOC mappings on behalf of the prefix owners on authoritative query servers (QSAs). Each QSA is only authoritative for a subset of all MABs. Resolver query servers (QSRs) use a DNS-based mechanism to resolve the EID(-prefix)-to-QSA mapping and then query the appropriate QSA for the EID(-prefix)-to-RLOC mapping.

ITRs communicate with QSRs directly or use a cascade of caching query servers (QSCs) to speedup consecutive lookups. QSAs store internally the last mapping requesters. In case of a mapping change, this enables QSAs to flush the cache entries in the QSRs, QSCs, and ITRs to force a new mapping lookup, and to reduce signaling complexity.

In contrast to other approaches, only one RLOC is stored per micronet. This is possible since IVIP assumes that edge networks hire third parties to effect real-time updates to the mapping system to take advantage of multihoming for inbound traffic engineering and service restoration in case of ITR/ETR failures. The author of [57] states further that resolution process of DRTM is fast enough so that ITRs can buffer initial packets of a flow without experiencing buffer overflow.

8) ID Mapping System: The ID mapping system (IDMS) [18] uses an extended version of DNS. IDs in IDMS are hierarchically structured and have a similar format like electronic mail addresses, i.e., *hostname@authority*. EID-suffix-to-MB mappings are stored in the extended DNS, i.e., new RRs are defined for DNS to enable the lookup process. MBs are implemented as so-called ID mapping servers and provide EID(-suffix)-to-RLOC mappings.

Each authority provides ID mapping servers and updates the EID-to-RLOC mapping in real-time. Mappings stored in DNS are stable while mappings in ID mapping servers may change frequently due to host mobility. The lookup works as described in Section IV-B2c.

The scalability of the system is limited by DNS and the ID mapping server implementations. The first is a general limitation for mapping systems based on DNS. The latter gives local authorities the freedom to choose a scalable solution for their own EID space, i.e., each local authority runs its own ID mapping server implementation. Local authorities are also responsible to implement appropriate resilience mechanisms for their servers. Security is provided through PKI and digital signatures.

9) Mapping Lookup for Intercepted DNS Queries: Before a host starts communication with a remote system, it mostly resolves a DNS name to an IP address. ITRs may intercept these DNS queries, query a new EID-to-RLOC mapping for the contained DNS names, and store the DNS reply in their cache so that packets sent to corresponding EIDs do not encounter a cache miss. This idea is not a mapping system but a DNS scheme for LISP [45, Scheme 2], i.e., a prefetching technique.

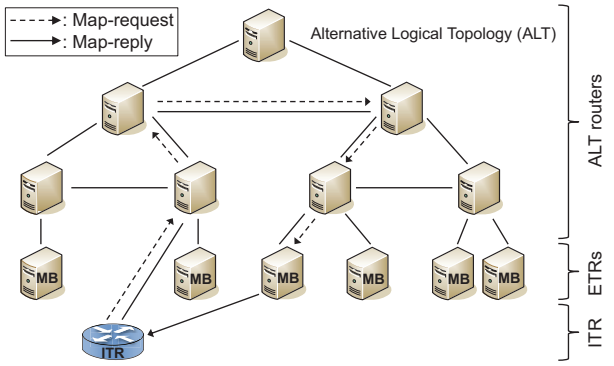

Fig. 9. In LISP+ALT, ALT routers are connected in a semi-hierarchical way and forward map-requests to responsible MBs through the "Alternative Logical Topology".

E. Map-Bases with Partial Knowledge using Overlay Lookup (MBPK-OL)

1) Map-Bases with Partial Knowledge using Hierarchically Structured Overlay (MBPK-HSO):

a) LISP+ALT: In LISP Alternative Logical Topology (LISP+ALT) [58] so-called ALT routers build a semi-hierarchically structured overlay network: the ALT. Fig. 9 gives an impression of the architecture. ALT routers are associated with EID prefixes and connected in a semi-hierarchical manner with respect to these prefixes. Shortcuts are possible on the same hierarchy level. A leaf ALT router connects to all MBs that store EID-to-RLOCs for its EID-prefix. Even though the architecture is strongly aggregation oriented, there are no root nodes. ALT routers communicate to neighboring ALT routers via BGP and exchange aggregated EID prefixes that can be reached through them. In contrast to regular BGP, ALT routers possibly aggregate prefixes received via BGP before forwarding them, which makes the ALT scalable.

The ITR addresses map-requests to the queried EID and sends them to an ALT router. The map-request is forwarded through the ALT overlay based on the EID. Eventually, the map-request reaches the appropriate MB which returns a map-reply directly to the ITR.

The operation of LISP+ALT is very efficient. The ALT routers are directly connected over tunnels using generic routing encapsulation (GRE). ALT routers simply forward packets addressed to EIDs according to their routing tables that are composed with the help of BGP on the basis of the EID prefixes associated with the ALT routers. Thus, ALT routers do not need to process the packets on the application layer like it is done in other MBPK-OL proposals. In case of a cache miss at the ITR, packets can also be carried over the ALT, but this is not recommended in the current LISP proposal [58, Section 3.3].

b) LISP-CONS: The "content distribution overlay network service for LISP" (LISP-CONS) [59] was a predecessor to LISP+ALT. LISP-CONS does not necessarily use BGP for communication between nodes of the hierarchy. Map-replies are returned from the ETRs back to the ITRs over the overlay network which is also different from LISP+ALT. LISP-CONS also allows carrying mapping requests and packets over the overlay network.

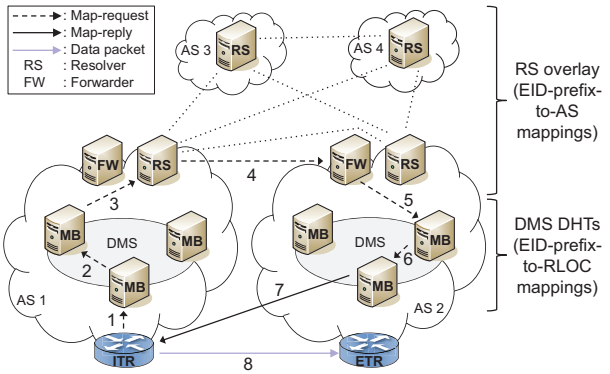

Fig. 10. In LISP-HMS, map-requests are served by a AS-local MB. If the MB cannot find the requested EID in its own database, it forwards the map-request over the DMS to another MB and possibly over the resolver overlay network to an appropriate MB in another AS.

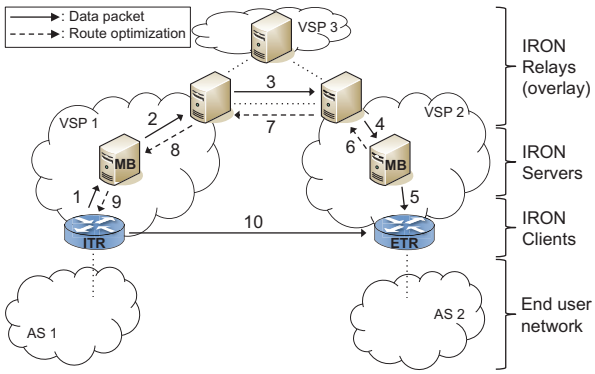

Fig. 11. When a cache miss occurs at an IRON ITR, packets are relayed over the overlay until the ITR is informed by an IRON server about an appropriate EID-to-RLOC mapping.

c) A Hierarchical Mapping System for LISP: LISP-HMS [60] is a hierarchical mapping system which combines BGP and one-hop DHTs. Fig. 10 shows the structure of LISP-HMS. MBs are called mapping servers and are responsible for a pre-defined mapping domain, i.e., a set of EID-prefixes. The MBs form a one-hop DHT called destination mapping server (DMS) which stores all mapping information of an AS. Note that there may be more than one DMS per AS. To provide mappings between ASes, forwarders aggregate the mappings of the DMSs even further, and eventually resolvers exchange EID-prefix-to-AS mapping information using BGP. Thus, the resolvers form an overlay network for inter-AS mapping information.

The resolution process works as follows. ITRs are configured with MBs. If the MB knows the mapping, it replies directly to the ITR. Otherwise, the map-request is forwarded to the associated DMS (step 1). In case the requested EID belongs to the same AS, the appropriate MB in the DMS replies directly to the ITR. Otherwise, the map-request is forwarded to the AS' resolver (step 2 and 3) which then forwards the request to the correct AS (step 4), DMS (step 5), and MB (step 6) which replies to the ITR (step 7). Then, the ITR can start sending packets (step 8).

When a new device joins an AS, it registers at its ITR. The ITR registers the EID-to-RLOC mapping at its MB. The MB updates the mappings at its DMS and reports its aggregated mappings to the forwarder. The forwarder aggregates its information base further and reports the mappings to its resolver. Then, the resolver propagates the new mapping information to the other ASes using BGP. Depending on the granularity of the aggregation, mapping information at upper levels stays stable. As BGP is used between ASes, security mechanisms of BGP can be utilized to protect the dissemination of mapping information.

d) ID/Locator Distributed Mapping Server: The mapping and relaying system presented in [61] has AS-specific MBs that keep the EID-to-RLOC mappings for the EIDs hosted in the AS. The MBs are implemented as DHTs. ITRs query their local MB in case of cache misses. Map-requests that cannot be served directly are forwarded to a so-called border server. Border servers of different ASes exchange the EID-prefixes under their control via BGP and build a hierarchically structured overlay for which EID aggregation is prerequisite for scalability.

e) IRON: The "Internet routing overlay network" (IRON) [34], [35] assumes that the global EID space is partitioned among virtual service providers (VSPs) in aggregated prefixes (AP) which are further partitioned in client prefixes (CP) and eventually delegated to end-hosts.

A VSP forms an IRON instance which comprises of IRON agents, i.e., IRON servers, IRON clients and at least one IRON relay. Fig. 11 shows the basic architecture and the basic information flow in IRON. IRON clients fulfill the function of ITRs and ETRs. IRON servers store the EID(-prefix)-to-RLOC information and announce their stored mappings to their IRON relays using eBGP, i.e., they are the MBs. Each IRON relay connects to the Internet as an AS using BGP and forms an overlay network with the other IRON relays. They aggregate the mapping information of all MBs of their respective IRON instance and internally store the EID(-prefix)-to-MB information. On the IRON relay overlay network, EID(-prefix)-to-IRON-relay information is exchanged between the IRON relays using iBGP.

When an ITR encounters a cache miss, it relays the packet without locator information to the MB of its VSP (step 1) which forwards it to one of its IRON relays (step 2). The IRON relay natively forwards the packet to the IRON relay of the destination VSP (step 3) which further tunnels it to the appropriate MB (step 4) that tunnels it to the ETR of the destination network (step 5).

Explicit map-requests do not exist. However, the MB responsible for the destination EID sends a "route optimization" message to the ITR (step 6 through 9) so that the ITR can tunnel further packets directly to the ETR (step 10). Thus, route optimization messages are similar to map-replies.

f) Realm Zone Bridging Server: RZBS is the mapping system of MILSA [15], [16]. Apparently, RZBS designates both the name of the mapping system architecture and a MB. The system shares some similarities with DNS. First of all, IDs in RZBS are structured like URIs, e.g., a valid ID in RZBS would be *bob.x.foo.com*. The ID space is partitioned in domains, which are called realms in RZBS. Each realm can be further partitioned in subrealms, subsubrealms and so on. MBs are responsible to store EID-suffix-to-MB and EID-to-RLOC

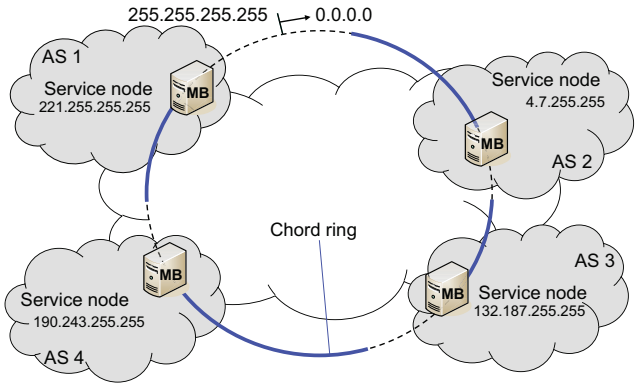

Fig. 12. In LISP-DHT, map-requests are forwarded over a Chord ring to responsible MBs.

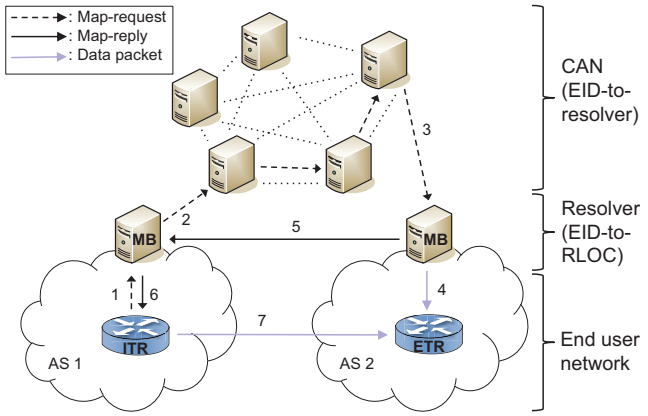

Fig. 13. In DHT-MAP, map-requests are served by an AS-local resolver (MB). If the resolver cannot find the requested EID in its own database, it forwards the map-request over the CAN to an appropriate resolver in another AS.

information. Thus, MBs form a hierarchically structured forest of realm trees. To connect the realm trees, DNS is used. The lookup works as described in Section IV-B3a.

RZBS achieves scalability and resilience through replication of MBs at different hierarchical levels. Security is realized through trust relationships between realm trees, i.e., each subtree trusts its root node and may trust its neighboring tree directly. The trust relationships influence how signaling messages are forwarded inside the overlay network.

2) Map-Bases with Partial Knowledge using Distributed Hash Tables (MBPK-DHT):

a) LISP-DHT: LISP-DHT stores mappings in a distributed hash table (DHT) [62]. Fig. 12 illustrates that MBs join a Chord ring as so-called service nodes to build a DHT. They have an ID that determines their position within the ring structure. Some modifications are applied to standard Chord. The ID of a service node is the highest number in the EID prefix for which it is responsible. Thus, service node IDs and EIDs are taken from the same number space. The un-hashed EIDs of map-requests are used for message forwarding in the DHT. Thus, a map-request is carried within the DHT over several hops to the service node with the smallest ID that is at least as large as the requested EID. These changes ensure that map-requests are forwarded to the service nodes that are responsible for them so that they can answer a map-reply to the requesting ITR.

An important feature of LISP-DHT is that prefix owners keep control over the mappings as they are kept local in the service nodes. If a service node is responsible for several EID prefixes, it has several IDs and is connected to the Chord ring at several positions. To prevent malicious nodes from EID prefix hijacking, joining service nodes must be authenticated as the rightful owners of their EID prefixes. For that purpose, the use of X.509 resource certificates [63] is proposed. To inject map-requests, ITRs join the Chord ring as stealth nodes which do not participate in message forwarding or other critical tasks. To address resilience concerns, LISP-DHT uses backup nodes providing the same mappings like normal service nodes.

b) ER+MO: In [64], a mapping and relaying system is presented which combines techniques similar to LISP+ALT and LISP-DHT. A customer network stores the mappings for its EIDs in a MB which is part of a mapping overlay (MO) very similar to LISP-DHT. However, Kademlia is used instead of Chord as DHT, and mappings are stored per EIDs instead of per EID prefix. Thus, a MB joins the DHT as a service node once for each EID under its control. This induces significant management overhead when multiple nodes join or leave. The relaying system works similarly as LISP+ALT. It consists of EID routers (ER) which learn EID-prefixes from ETRs via BGP and relay packets if needed.

c) DHT-MAP: In contrast to many other approaches, DHT-MAP [65] supports a flat identifier space. Fig. 13 shows the structure of DHT-MAP. Each AS operates a MB – called resolver – which stores the AS-specific EID-to-RLOC mappings for the EIDs supported within the AS. Resolvers of different ASes are connected to a content addressable network (CAN) which is a special type of DHT in which EID-to-MB mappings are stored.

ITRs of an AS are connected to a resolver. When an ITR encounters a cache miss, it sends a map-request including the packet to the resolver (step 1). If the resolver knows the EID-to-RLOC mapping, it tunnels the packet to the ETR and returns a map-reply to the ITR (step 6); otherwise, it sends the map-request including the packet into the CAN (step 2). The CAN node that is responsible for the requested EID may have different EID-to-MB mappings, chooses one of them, and forwards the map-request to that resolver (step 3). This resolver has an appropriate EID-to-RLOC mapping, tunnels the packet to the ETR (step 4), and sends a map-reply to the requesting resolver (step 5) which forwards it to the requesting ITR (step 6). All subsequent packets are tunneled directly to the destination ETR (step 7).

We briefly explain how the mappings are registered in DHT-MAP. When a new device joins an AS, it registers at the ITR. The ITR registers the new EID-to-RLOC mapping at the resolver and the resolver registers the EID-to-resolver mapping in the CAN. Therefore, DHT-MAP's resolvers know only the RLOCs of one AS. Since the CAN node that is responsible for the EID forwards map-requests only to a single resolver, the ITRs receives only the RLOCs of a single AS for an EID. This is a strong limitation for multihoming.

DHT-MAP can support flat EID spaces. It is able to relay

packets, but the extensive path stretch for first packets effects that they face longer delays and that mapping lookups take relatively long. This causes more relayed traffic than for short mapping lookups so that substantial forwarding capacity is needed in the CAN. DHT-MAP relies on the resilience features of the DHT to carry the map-request to backup nodes in the DHT and on resilience features in the AS to carry map-requests to backup resolvers.

d) HIP-DHT: HIP-DHT [66] proposes to use a DHT for looking up HIP-related information based on a HIT (see Section V-D4). The authors specify how their concept works with OpenDHT. However, the authors also list security concerns pointing out potential map-reply spoofing attacks leading to stale information or mapping pollution since authentication is not required to register new or already existing mappings in the system.

e) RANGI: In the "routing architecture for the next generation Internet" (RANGI) [31], the name space of host identifiers (HI) is partitioned by prefixes among administrative domains (ADs). HIs consist of two parts: the globally unique AD ID which is assigned by a central authority like IANA and a cryptographical part that is generated as a hash containing the AD ID and a public key value like in HIP. An AD takes care that the HIs under its control are unique. RANGI uses a hierarchical DHT to map HIs to RLOCs. A top-level DHT guides map-requests to bottom-level DHTs using the AD ID in the HI. The bottom-level DHTs uses the unstructured cryptographical part of the HI to resolve the mapping and send map-replies to ITRs.

f) CoDoNS: CoDoNS stands for cooperative domain name system [67]. It is proposed as a substitute for the DNS and it is implemented based on a DHT called Pastry and enhanced using a proactive caching layer called Beehive. CoDoNS replicates mapping information across the DHT to achieve an access time of practically $O(1)$. Large organizations should participate in CoDoNS with at least two nodes. These nodes store data from other organizations and the organization's own data are probably stored on nodes of other organizations. This property is hard to accept in practice which is also a general argument against the straightforward use of DHTs as a mapping system.

g) LISP-SHDHT: The main objectives of LISP single-hop DHT mapping overlay (LISP-SHDHT) are fast lookup and load balancing [68]. MBs are called SHDHT nodes and form a single-hop DHT. The system internally uses two different namespaces: node IDs and partition IDs. Node IDs designate MBs. Partition IDs correspond to the hashed EID space and designate end-hosts. Each MB has at least one partition ID assigned which indirectly defines the partition ID range the MB is responsible for. IDs in a partition range are by definition called resource IDs. Each MB knows all mappings between partition IDs and node IDs, i.e., each MB can resolve EID-to-MB mappings in one hop.

When a new mapping is registered in LISP-SHDHT, an ETR sends a map-register message to a known MB. The MB generates the resource ID from the to-be-registered EID using a hash function. The resource ID is matched to the closest MB using the MB's internal node routing table. If the current node is responsible for the resource ID, it stores the mapping.

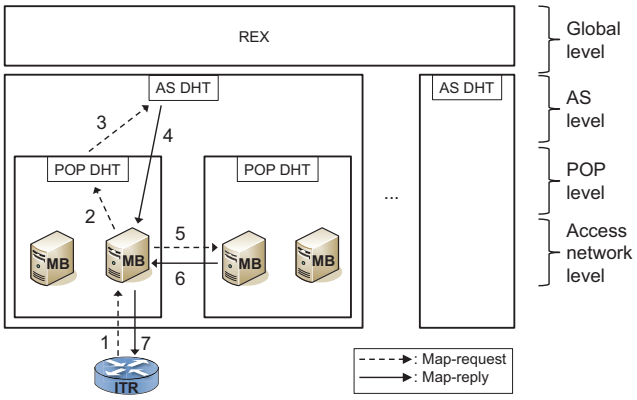

Fig. 14. In MDHT, map-requests are served by an access-network-local resolver (MB). If the resolver cannot find the requested EID in its own DHT, it forwards the request to the next higher DHT until the appropriate MB is found, sends the map-request to the authoritative MB and returns the mapping information to the ITR.

Otherwise, the current MB forwards the map-register message to the responsible MB which then registers the mapping.

The lookup procedure in LISP-SHDHT works similarly. An ITR sends a map-request message to a known MB. The MB generates the resource ID from the requested EID using a hash function. The resource ID is matched to the closest MB using the MB's internal node routing table. If the current node is responsible for the resource ID, it replies the mapping. Otherwise, the current MB forwards the map-request message to the responsible MB which then replies the mapping.

By the time of writing, LISP-SHDHT is still under development and security mechanisms have not been discussed, yet. Resilience is realized through replication of MBs. Packet relaying is not supported.

h) MDHT: MDHT stands for multi-level DHT. It has been proposed as a name resolution service for information-centric networks in [69] and maps flat object identifiers to network addresses. In theory, it could be reused for EID-to-RLOC mappings. Therefore, we describe MDHT with different nomenclature. In MDHT, all EIDs located in a specific access network are stored in a MB called access node. MBs are grouped in a nested, hierarchical structure of DHT areas, e.g., a point of presence (POP) DHT holds EID-to-MB pointers for all EIDs stored within its domain, and an AS DHT holds EID-to-MB pointers for all EIDs stored within its domain. Fig. 14 shows the basic structure and the basic information flow in MDHT.

Mapping retrieval in MDHT works as follows. An ITR queries its MB for the mapping (step 1). If the MB holds the requested mapping, it can locally retrieve the mapping, otherwise it queries its peers in the DHT area. If the mapping is not retrievable in the own DHT area, the query is recursively forwarded to the next higher DHT area until it can be answered (step 2 and 3). The answering DHT area returns an EID-to-MB mapping to the requesting MB (step 4). Finally, the requesting MB sends a map-request to the authoritative MB (step 5 and 6) and returns the mapping to the ITR (step 7).

The presented mapping system works only within a single AS. A global DHT is unlikely to scale. To solve that problem, the paper suggests object identifier prefixes and a global

resolution exchange system (REX).

3) Map-Bases with Partial Knowledge using Multicast Overlay (MBPK-MCO): EMACS-LISP stands for "EID mappings multicast across cooperating systems" for LISP [70]. MBs join multicast groups for all EID prefixes they are responsible for. If that prefix is X.Y.A.B/16, the address of the corresponding multicast group is, e.g., 238.1.X.Y. In case of a cache miss for EID X.Y.A.B, the ITR sends the data packet to the corresponding multicast group so that all MBs of that group receive it. All MBs having appropriate mappings for the requested EID can respond with a map-reply. However, when data packets are relayed over this structure, only one of these MBs should deliver the packet to avoid duplicates at the destination. This approach has several drawbacks. Up to 2^{16} multicast groups need to be maintained in BGP and a large amount of unnecessary extra traffic is generated through multicast delivery.

VI. Discussion

In this section we address general aspects of mapping systems and point out remaining research opportunities.

A. Fulfillment of Requirements

In this work we have surveyed a large number of mapping systems for future Internet routing based on Loc/ID split. We summarize their most important properties in Table III. There is not a single mapping system that meets all requirements listed in Section III. Hence, there is room for additional work. However, each future Internet routing architecture comes with its own specific requirements that impact the choice of the most appropriate mapping system architecture. For instance, HIP depends on a flat namespace so that only mapping systems are eligible that support this feature. As another example, routing architectures where end hosts perform the mapping lookup do not need a mapping system that performs packet relaying. In a similar way, workarounds may be created for routing architectures with intermediate mapping nodes to cope without packet relaying.

B. Economical and Political Aspects

The operation of a future Loc/ID based Internet routing architecture will depend on the correct operation of its mapping system. Therefore, the organization operating the mapping system controls the Internet. Furthermore, the operation of the infrastructure is costly and expenses need to be refunded in some way by those benefitting from this service. Care must be taken that the owners of all components of the mapping system are obliged or have incentives to forward all map-requests, map-replies or relayed packets to achieve proper network operation for all participants.

C. Provisioning of Mapping Systems

Once the mapping system is chosen, the infrastructure must be provisioned. The number and placement of components need to be determined to guarantee smooth operation even under heavy load and in failure cases.

D. Engineering of Mapping Systems for LISP

As LISP is the Loc/ID split routing architecture that presently sees major deployment, engineering mapping systems for LISP is a valid issue. LISP+ALT has been replaced by LISP-DDT as the current mapping system in the LISP pilot network. A major shortcoming of LISP-DDT is the missing capability to relay packets for which EID-to-RLOC mappings are not available in the ITRs. Therefore, LISP-DDT will remain the preferred mapping system only until a better mapping system architecture is available and adopted.

E. Performance Measurements and Improvements

To the best of our knowledge, only theoretical performance studies and simulations of future mapping systems were conducted [3], [36], [37]. However, they cannot replace measurement studies of operational systems which give insights in the actual behavior of mapping systems under load. After the deployment of LISP-DDT in the LISP pilot network a preliminary performance measurement study was presented at the IETF 84 meeting [71]. However, more work needs to be carried out in this area to detect potential bottlenecks and find solutions.

VII. Conclusion

We presented requirements and a taxonomy for mapping systems for Loc/ID split Internet routing architectures. We provided a comprehensive review of recently proposed mapping systems and classified them into our proposed categories. We discussed the different approaches especially with regard to the requirements and pointed out similarities and differences.

Constructing a mapping system is a complex task and it is hard to fulfill all requirements at once. There are still open problems which are not solved satisfactorily. Therefore, mapping systems for Loc/ID split routing still provide interesting research topics for the next years.

Acknowledgments

The authors thank Jeffrey Ahrenholz, Marcelo Bagnulo, Roland Bless, Scott Brim, Brian Carpenter, Anja Feldmann, Vince Fuller, Oliver Hanka, Luigi Iannone, Dominik Klein, Eliot Lear, Tony Li, David Meyer, Erik Nordmark, David Oran, Benno Overeinder, Oleg Ponomarev, Christopher Spleiss, Steve Uhlig, Christian Vogt, Robin Whittle, Lixia Zhang, Mark Schmidt, Li Cheng, Alfons Martin, and Cynthia Mills for valuable input and stimulating discussions.

References

[1] B. Carpenter, R. Atkinson, and H. Flinck, "Renumbering Still Needs Work," RFC5887, May 2010.
[2] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," RFC4984, Sep. 2007.
[3] B. Quoitin, L. Iannone, C. de Launois, and O. Bonaventure, "Evaluating the Benefits of the Locator/Identifier Separation," in *ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, Kyoto, Japan, Aug. 2007.
[4] E. Nordmark and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6," RFC5533, Jun. 2009.
[5] C. de Lanois and M. Bagnulo Braun, "The Paths Toward IPv6 Multihoming," *IEEE Commun. Surveys & Tutorials*, vol. 8, no. 2, 2006.

TABLE III
SUMMARY OF THE SURVEYED PROPOSALS FOR MAPPING SYSTEMS FOR LOC/ID SPLIT ROUTING

Category	Name	Mapping Structure	Scalability	Resilience	Security	Relaying
MBFK	LISP-NERD	hierarchical	MBs store partition of mappings, ITRs assemble complete mapping table	replication of MBs	X.509 certificates	yes
	APT	hierarchical	DMs know all mapping information	replication of DMs	digital signatures	yes
MBPK-LL	FIRMS	flat and hierarchical	ITRs/MRs know global MBP table	replication of all components	X.509 certificates, PKI	yes
MBPK-SRL	HiiMap	flat and hierarchical	one regional prefix per EID, large storage requirements	replication, DHTs	PKI	no
MBPK-IRL	One-Phase Lookup Using Reverse DNS/DNSMAP	flat	uses DNS infrastructure	relies on DNS	DNSSEC	no
	Two-Phase Lookup Using Reverse DNS	flat	uses DNS infrastructure	relies on DNS	DNSSEC	no
	ILNP-DNS	flat	uses DNS infrastructure	relies on DNS	DNSSEC	no
	Use of DNS for HIT-to-IP Lookup in HIP	flat	uses DNS infrastructure	relies on DNS	DNSSEC	no
	LISP-TREE	flat and hierarchical	uses DNS infrastructure, optionally own physical infrastructure based on DNS software	relies on DNS	DNSSEC	no
	LISP-DDT	flat and hierarchical	based on DNS software	similar to DNS	similar to DNSSEC	no
	IVIP DRTM	hierarchical	aggregation of mapping information, load balancing between components	replication of all components	none	no
	IDMS	flat and hierarchical	uses DNS infrastructure, IDMS implementation	DNS, replication of MBs	PKI, digital signatures	no
MBPK-HSO	LISP+ALT	hierarchical	BGP aggregation and limitations, complex configuration	replication of all components	BGP security	optionally
	LISP-CONS	hierarchical	strict aggregation hierarchy	replication of all components, redundant toplogy	similar to BGP and DNSSEC	yes
	LISP-HMS	hierarchical	strong aggregation of mapping information, DHT	replication of all components	BGP security	no
	ID/Locator Distributed Mapping Server	hierarchical	BGP and DHT, aggregation	DHTs	none	no
	IRON	hierarchical	aggregation	replication of all components	mutual authentication between components	yes
	RZBS	hierarchical	similar to DNS	replication of all components	trust relationships between subrealms	no
MBPK-DHT	LISP-DHT	hierarchical	DHT	replication of all components	X.509 certificates	no
	ER+MO	hierarchical	BGP and DHT, aggregation	replication of all components	BGP security, Kademlia security	yes
	DHT-MAP	flat	DHT	replication of all components	digital signatures	no
	HIP-DHT	flat	DHT	DHTs	none	no
	RANGI	hierarchical	DHT	DHTs	digital signatures	no
	CoDoNS	flat	DHT	replication of all components	none	no
	MDHT	flat	multilevel DHT	replication of all components	none	yes
	LISP-SHDHT	flat and hierarchical	DHT	replication of all components	none	no
MBPK-MCO	EMACS-LISP	hierarchical	number of multicast groups, unnecessary multicast traffic	replication of MBs	none	no

[6] K. Li, S. Wang, S. Xu, and X. Wang, "ERMAO: An Enhanced Intradomain Traffic Engineering Approach in LISP-Capable Networks," in *IEEE Globecom*, Houston, TX, USA, Dec. 2011, pp. 1–5.

[7] R. Atkinson, S. Bhatti, and S. Hailes, "ILNP - Identifier/Locator

Network Protocol," http://ilnp.cs.st-andrews.ac.uk/, 2009.
[8] ——, "ILNP: Mobility, Multi-Homing, Localised Addressing and Security through Naming," *Telecommunication Systems*, vol. 42, no. 3–4, pp. 273 – 291, Dec. 2009.
[9] R. Atkinson, "ILNP Concept of Operations," draft-rja-ilnp-intro-11, Jul. 2011.
[10] A. Feldmann, L. Cittadini, W. Mühlbauer, R. Bush, and O. Maennel, "HAIR: Hierarchical Architecture for Internet Routing," in *Re-Architecting the Internet (ReArch)*, Rome, Italy, Dec. 2009.
[11] P. Frejborg, "Hierarchical IPv4 Framework," RFC6306, Jul. 2011.
[12] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," RFC4423, May 2006.
[13] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, "Host Identity Protocol," RFC5201, Apr. 2008.
[14] S. Jiang, "Hierarchical Host Identity Tag Architecture," draft-jiang-hiprg-hhit-arch-04, May 2010.
[15] J. Pan, S. Paul, R. Jain, and M. Bowman, "MILSA: A Mobility and Multihoming Supporting Identifier Locator Split Architecture for Naming in the Next Generation Internet," in *IEEE Globecom*, New Orleans, LA, Nov. 2008.
[16] J. Pan, R. Jain, S. Paul, M. Bowman, X. Xu, and S. Chen, "Enhanced MILSA Architecture for Naming, Addressing, Routing and Security Issues in the Next Generation Internet," in *IEEE International Conference on Communications (ICC)*, Dresden, Germany, Jun. 2009.
[17] J. Pan, S. Paul, R. Jain, and X. Xu, "Hybrid Transition Mechanism for MILSA Architecture for the Next Generation Internet," in *IEEE Globecom*, Honolulu, Hawaii, Dec. 2009.
[18] J. H. Wang, Y. Wang, M. Xu, and J. Yang, "Separating Identifier from Locator with Extended DNS," in *IEEE International Conference on Communications (ICC)*, Ottawa, Canada, Jun. 2012, pp. 2780 – 2784.
[19] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "Locator/ID Separation Protocol (LISP)," draft-ietf-lisp-24.txt, Nov. 2012.
[20] D. Meyer, "The Locator Identifier Separation Protocol (LISP)," *The Internet Protocol Journal*, vol. 11, no. 1, pp. 23–36, Mar. 2008.
[21] D. Lewis, D. Meyer, D. DinoFarinacci, and V. Fuller, "Interworking LISP with IPv4 and IPv6," draft-ietf-lisp-interworking-06.txt, Mar. 2012.
[22] D. Farinacci, D. Lewis, D. Meyer, and C. ChrisWhite, "LISP Mobile Node," draft-meyer-lisp-mn-08.txt, Oct. 2012.
[23] M. Menth, D. Klein, and M. Hartmann, "Improvements to LISP Mobile Node," in *International Teletraffic Congress (ITC)*, Amsterdam, The Netherlands, Sep. 2010.
[24] D. Klein, M. Hartmann, and M. Menth, "NAT Traversal for LISP Mobile Node," in *Re-Architecting the Internet (ReArch)*, Philadelphia, PA, USA, Nov. 2010.
[25] V. Fuller and D. Farinacci, "LISP Map Server Interface," draft-ietf-lisp-ms-16.txt, Mar. 2012.
[26] C. Vogt, "Six/One Router: A Scalable and Backwards Compatible Solution for Provider-Independent Addressing," in *ACM International Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, Seattle, WA, USA, Aug. 2008.
[27] M. Menth, M. Hartmann, and D. Klein, "Global Locator, Local Locator, and Identifier Split (GLI-Split)," University of Würzburg, Technical Report, No. 470, Apr. 2010.
[28] D. Jen, M. Meisel, D. Massey, L. Wang, B. Zhang, and L. Zhang, "APT: A Practical Tunneling Architecture for Routing Scalability," http://fmdb.cs.ucla.edu/Treports/080004.pdf, UCLA Computer Science Department, Tech. Rep. 080004, Mar. 2008.
[29] O. Hanka, C. Spleiss, G. Kunzmann, and J. Eberspächer, "A Novel DHT-Based Network Architecture for the Next Generation Internet," in *International Conference on Networking (ICN)*, Cancun, Mexico, Mar. 2009.
[30] S. Schuetz, R. Winter, L. Burness, P. Eardley, and B. Ahlgren, "Node Identity Internetworking Architecture," draft-schuetz-nid-arch-00, Sep. 2007.
[31] X. Xu, "Routing Architecture for the Next Generation Internet," draft-xu-rangi-04, Aug. 2010.
[32] R. Whittle, "Ivip - a scalable routing and mobility architecture for the IPv4 and IPv6 Internets," www.firstpr.com.au/ip/ivip/, 2011.
[33] ——, "Ivip (Internet Vastly Improved Plumbing) Architecture," draft-whittle-ivip-arch-04, Mar. 2010.
[34] F. Templin, "The Internet Routing Overlay Network (IRON)," RFC6179, Mar. 2011.
[35] ——, "The Intradomain Routing Overlay Network (IRON)," draft-templin-ironbis-12.txt, Oct. 2012.
[36] L. Iannone and O. Bonaventure, "On the Cost of Caching Locator/ID Mappings," in *ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Dec. 2007.
[37] J. Kim, L. Iannone, and A. Feldmann, "A deep dive into the lisp cache and what isps should know about it," in *IFIP* 10^{th} *International Conference on Networking*, vol. 6640, Berlin / Heidelberg, Germany, May 2011, pp. 267–278.
[38] T. Narten, R. Draves, and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," RFC4941, Sep. 2007.
[39] K. Sriram, Y.-T. Kim, and D. Montgomery, "Enhanced Efficiency of Mapping Distribution Protocols in Scalable Routing and Addressing Architectures," in *IEEE International Conference on Computer Communications and Networks (ICCCN)*, Zurich, Switzerland, Aug. 2010.
[40] M. Menth, M. Hartmann, and M. Hoefling, "FIRMS: A Mapping System for Future Internet Routing," *IEEE J. Sel. Areas Commun., Special Issue on Internet Routing Scalability*, vol. 28, no. 8, Oct. 2010.
[41] E. Lear, "NERD: A Not-so-novel EID to RLOC Database," draft-lear-lisp-nerd-09.txt, Apr. 2012.
[42] D. Jen, M. Meisel, D. Massey, L. Wang, B. Zhang, and L. Zhang, "APT: A Practical Transit Mapping Service," draft-jen-apt-01, Nov. 2007.
[43] O. Hanka, G. Kunzmann, C. Spleiß, J. Eberspächer, and A. Bauer, "HiMap: Hierarchical Internet Mapping Architecture," in *ICFIN*, Beijing, China, Oct. 2009.
[44] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC4033, Mar. 2005.
[45] D. Farinacci, D. Oran, V. Fuller, and J. Schiller, "Locator/ID Separation Protocol (LISP2) [DNS-based Version]," http://www.dinof.net/ dino/ietf/lisp2.ppt, Nov. 2006.
[46] C. Vogt, "DNS Map – A DNS-Based Resolution System for IP Address Mappings," Ericsson, Technical Report, Feb. 2008.
[47] R. Atkinson and S. Bhatti, "An Introduction to the Identifier-Locator Network Protocol (ILNP)," in *London Communications Symposium (LCS)*, London, UK, Jul. 2006.
[48] R. Atkinson and S. Rose, "DNS Resource Records for ILNP," draft-rja-ilnp-dns-11, Jul. 2011.
[49] R. Atkinson and S. Bhatti, "ILNP Architectural Description," draft-irtf-rrg-ilnp-arch-06, Jul. 2012.
[50] P. Nikander and J. Laganier, "Host Identity Protocol (HIP) Domain Name System (DNS) Extensions," RFC5205, Apr. 2008.
[51] O. Ponomarev and A. Gurtov, "Embedding Host Identity Tags Data in DNS," draft-ponomarev-hip-hit2ip-04, Mar. 2009.
[52] L. Jakab, 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., Special Issue on Internet Routing Scalability*, vol. 28, no. 8, Oct. 2010.
[53] S. Letong, "Layered Mapping System," http://www.ietf.org/mail-archive/web/rrg/current/msg05491.html, Dec. 2009.
[54] S. Letong, Y. Xia, W. Z. Liang, and W. Jianping, "A Layered Mapping System for Scalable Routing," http://tinyurl.com/LeXiLi09, Dec. 2009, Tsinghua University.
[55] V. Fuller, D. Lewis, V. VinaErmagan, and A. Jain, "LISP Delegated Database Tree," draft-fuller-lisp-ddt-04.txt, Sep. 2012.
[56] F. Maino, V. Ermagan, A. AlbertCabellos-Aparicio, D. Saucez, and O. Bonaventure, "LISP-Security (LISP-SEC)," draft-ietf-lisp-sec-04.txt, Oct. 2012.
[57] R. Whittle, "DRTM - Distributed Real Time Mapping for Ivip and LISP," draft-whittle-ivip-drtm-01, Mar. 2010.
[58] V. Fuller, D. Farinacci, D. Meyer, and D. Lewis, "LISP Alternative Topology (LISP+ALT)," draft-ietf-lisp-alt-10.txt, Dec. 2011.
[59] S. Brim, N. Chiappa, D. Farinacci, V. Fuller, and D. Lewis, "LISP-CONS: A Content distribution Overlay Network Service for LISP," draft-meyer-lisp-cons-04, Apr. 2008.
[60] H. Zhang and Z. Zhang, "A Hierarchical Mapping System for LISP," draft-zhang-lisp-hms-01.txt, Dec. 2012.
[61] F. Hu and J. Luo, "ID/Locator Distributed Mapping Server," draft-hu-lisp-dht-00, Oct. 2009.
[62] L. Mathy and L. Iannone, "LISP-DHT: Towards a DHT to Map Identifiers onto Locators," in *Re-Architecting the Internet (ReArch)*, Madrid, Spain, Dec. 2008.
[63] G. Huston, "Resource Certification," *The Internet Protocol Journal*, vol. 12, no. 1, pp. 13–26, Mar. 2009.
[64] G. Chen et al., "An Incremental Deployable Mapping Service for Scalable Routing Architecture," draft-chen-lisp-er-mo-01, Jul. 2009.
[65] H. Luo, Y. Qin, and H. Zhang, "A DHT-Based Identifier-to-Locator Mapping Scheme for a Scalable Internet," *IEEE Trans. Parallel Distrib. Syst.*, vol. 20, no. 10, Oct. 2009.
[66] J. Ahrenholz, "HIP DHT Interface," draft-ahrenholz-hiprg-dht-06, Nov. 2009.
[67] V. Ramasubramanian and E. G. Sirer, "The Design and Implementation of a Next Generation Name Service for the Internet," in *ACM SIGCOMM*, Portland, OR, USA, 2004.

[68] L. Cheng and J. Wang, "LISP Single-Hop DHT Mapping Overlay," draft-cheng-lisp-shdht-02.txt, Oct. 2012.
[69] M. D'Ambrosio, C. Dannewitz, H. Karl, and V. Vercellone, "Fast, Effective and Stable IP Recovery using Resilient Routing Layers," in *ICN'11 Proc. ACM SIGCOMM workshop on Information-centric networking*, New York, NY, USA, Aug. 2011.
[70] S. Brim, D. Farinacci, D. Meyer, and J. Curran, "EID Mappings Multicast Across Cooperating Systems for LISP," draft-curran-lisp-emacs-00, Nov. 2007.
[71] D. Saucez, L. Iannone, and B. Donnet, "A brief look at the Mapping System," in *Proc. IETF 84*, Vancouver, BC, Canada, Jul. 2012.

Michael Hoefling is researcher at the Department of Computer Science at the University of Tuebingen/Germany and pursuing his Ph.D. at the Chair for Communication Networks. Prior, he received a B.Sc. and a M.Sc. degree in computer science from the University of Umeå/Sweden in 2009, and a German diploma in computer science with minor in physics (Dipl.-Inform.) from the University of Wuerzburg/Germany in 2010. His current research focuses on current and future Internet addressing and routing, smart grid, information-centric networking, data center networks, grid and cloud computing, as well as energy efficient routing.

Michael Menth is professor at the Department of Computer Science at the University of Tuebingen/Germany and chairholder of Communication Networks. He received a diploma and a PhD degree in 1998 and 2004 from the University of Wuerzburg/Germany. Prior he was studying computer science at the University of Texas at Austin and worked at the University of Ulm/Germany. His special interests are performance analysis and optimization of communication networks, resource management, resilience issues, smart grids, and future Internet. He holds numerous patents and received various scientific awards for innovative work.

Matthias Hartmann studied computer science and mathematics at the University of Wuerzburg/Germany, the University of Texas at Austin, and at the Simula Research Laboratory in Oslo/Norway. He received his Diploma degree in computer science in 2007. Currently, he is a researcher at the Institute of Computer Science in Wuerzburg and pursuing his Ph.D. His current research focuses on performance and resilience analysis as well as on future Internet routing.