

# HAIR: Hierarchical Architecture for Internet Routing

Anja Feldmann

T-Labs/TU Berlin  
anja@net.t-labs.tu-berlin.de

Luca Cittadini

Università Roma Tre  
luca.cittadini@gmail.com

Wolfgang Mühlbauer

T-Labs/TU Berlin  
wolfgang@net.t-labs.tu-berlin.de

Randy Bush

IJJ  
randy@psg.com

Olaf Maennel

Loughborough University  
olaf@maennel.net

## ABSTRACT

In the light of recent interest in re-designing the Internet, we introduce HAIR, a routing architecture that tackles the problem of routing table growth, restricts the visibility of routing updates, and inherently supports traffic engineering, mobility, and multipath.

HAIR separates locators from identifiers. The routing and mapping system rely on an hierarchical scheme that leverages the structure of today’s Internet. Contrary to proposals such as LISP [5] and shim6 [15], we use a hybrid *edge-based* approach where only some lightweight functionality is added within the network, while the majority of tasks are performed as close to the end hosts as possible.

We estimate the potential benefits of HAIR and present a working proof-of-concept implementation.

## 1. INTRODUCTION

While the research community is still discussing the fundamental limits to routing scalability [11], routing tables in the default-free zone of today’s Internet already contain some 300,000 prefixes and continue to grow super-linearly. This goes along with a steady increase in the rate of changes to the routing tables [8]. Moreover, features such as mobility, security, support for traffic engineering or multi-homing are lacking; IP addresses are a scarce resource, and customers are hindered by provider lock-in.

Given all these problems, suggesting a major overhaul of today’s Internet has become a popular research topic. We, in this paper, introduce HAIR, a routing architecture that tackles the problem of routing table growth, restricts the visibility of routing updates, and inherently provides means for traffic engineering and mobility.

Despite the multitude of existing proposals for future Internet architectures, there is wide consensus among scientists to separate the *identifier* from the *locator* functionality (Loc/ID split), e.g., [5, 10, 15, 20]. Currently, both are mangled together within IP addresses. Loc/ID split allows for persistent identifiers and at the same time for aggregation in the locator namespace. There exist multiple proposals,

amongst others *host-based* (e.g., shim6 [15]) and *network-based* (e.g., LISP [5]) approaches. The main advantage of host-based solutions is that they can be self-deployed without cooperation of the network operators. Network-based approaches, however, are capable of transparently supporting legacy hosts and of reducing routing table size in the core. We note that these alternatives have different objectives and at least in part complement one another. Therefore, we argue for a hybrid *edge-based* solution that transfers to end hosts tasks such as translating identifiers to locators, while keeping some elements of the network-based approaches to address scalability limits and migration issues.

Another key to the design of both the routing *and* the mapping system of HAIR is to leverage the structure, which is inherently present in today’s Internet [3, 17]: the existence of a stable “core” formed by large transit providers (*CORE*), and a more dynamic edge, consisting of smaller access or enterprise networks (*INTERMEDIATE*, short *INT*) and *EDGES*, e.g., Ethernet domains. The goal is to prevent both routing updates (due to routing changes) and mapping updates (due, e.g., to mobility or traffic engineering) from being globally visible.

Before sending a packet, a host asks a mapping service for the corresponding locator(s) of a given identifier. The mapping service is implemented as a distributed system where mappings are stored locally, i.e., at the authorities that own the mappings. A locator encodes a loose source route by specifying one possible exit point from the *CORE* area to the *INT*; one possible exit point from the *INT* to the *EDGE*; and finally the identifier of the destination host. In principle, every host can have multiple locators. According to our edge-based paradigm, the responsibility to add the locators to the packet headers is pushed to end hosts, keeping the gateway devices at the exit points simple.

The rest of this paper is structured as follows. First we present our architecture in Section 2. Then in Section 3 we discuss its advantages. Section 4 presents a proof-of-concept implementation. Finally, we summarize related work in Section 5 and conclude in Section 6.

## 2. ARCHITECTURE

### 2.1 Overview

The motivation behind HAIR is to tackle the problem of routing table growth, to restrict the visibility of routing updates and to provide inherent support for traffic engineering and mobility. To achieve our goals, the design of HAIR follows three key ideas: (i) *Locator/identifier separation*, (ii) *hierarchical scheme* for both the routing and the mapping system that is needed to translate identifiers into locators and (iii) a hybrid *edge-based approach*. In the following, each of these principles will be explained in more detail.

#### 2.1.1 Hierarchical scheme

Work from the past (e.g., [3, 17]) suggests that today’s Internet consists of a stable “core”, formed by large transit providers, and a more dynamic “edge”, consisting of enterprise networks or small access providers. The typical Internet cost structure together with peering policies ensures that the number of links between an “edge” network and the “core” is limited. Most “edge” networks have a small number of upstream providers in the “core” and typically, due to costs, do not have too many upstream links to a single provider. Accordingly, we propose to group “core” networks into level 1 of a hierarchy, “edge” or *intermediate* networks into levels 2 to  $n - 1$ , and local area networks into level  $n$ , obtaining a hierarchy consisting of  $n$  levels (see Figure 1):

**EDGE (Level  $n$ ):** A *EDGE* is an access network which contains end hosts and servers, e.g., an Ethernet LAN.

**INT (Levels 2 to  $n - 1$ )** An *INTERMEDIATE (INT)* network consists only of routers. It provides routing between attached *EDGES* or *INTs* of the next higher level. Hence, routing tables within a *INT* need entries for all routers within the *INT*, routes to the attached *EDGES/INTs*, and default routes to the *INTs* of the next higher level or the *CORE*. In the current Internet, *INTs* may correspond to access providers, enterprise networks, or content distribution networks.

**CORE (Level 1):** The *CORE* ensures global reachability by routing packets between the *INTERMEDIATE* networks of level 2. Routers in the *CORE* are grouped into administrative domains, each under the control of a single transit ISP. Routing tables in the *CORE* contain entries for all routers within the *CORE* as well as routes to the directly attached *INTs*.

According to Figure 1, the hierarchy levels *CORE*, *INTs* or *EDGES* are connected via *attachment points*: a “level  $k$ ” attachment point ( $L_k$ AP) connects a level  $k$  routing domain to a level  $k + 1$  one. Interconnections within the same hierarchical layer (“peerings”) are also possible.

Unlike most previous proposals, HAIR’s hierarchy splits routers from what is now a single AS into different levels

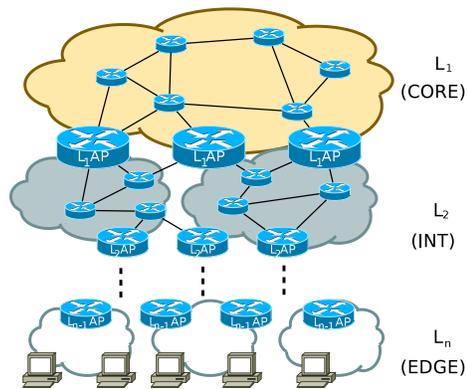


Figure 1: N-layer hierarchical network structure.

of the hierarchy according to their function. One way to map today’s Internet to a 3-layer hierarchy is to group all the routers in the backbone portion of transit ASs within the *CORE*. The policy-based BGP can be used for routing in the *CORE* area according to the policies that ISPs configure on the set of  $L_1$  APs they own. *INTs* such as enterprise networks or access provider can run their preferred routing protocol, e.g., OSPF or IS-IS. Thus HAIR preserves the autonomy of providers. Due to the limited scope of routing domains, HAIR is able to restrict the visibility of routing updates and to tackle the problem of routing table growth, see Section 3.

#### 2.1.2 Locator/Identifier separation

Decoupling locators from identifiers provides inherent support for (end-host) mobility and avoids issues such as provider lock-in. However, a new architectural component – a *mapping service* is needed: Given a certain identifier it returns the current locator(s), see Section 2.3.

The current design of HAIR does not yet specify identifiers for end hosts. However, in the following we will tacitly assume that identifiers are organized in a global flat namespace. After all, this promises to avoid problems such as provider lock-in and renumbering.

In contrast to identifiers the namespace of locators has a local scope. They are managed by the individual *INTs*. A locator for an end host is similar to a *loose source route* from the *CORE* towards an end-host: It consists of a sequence of attachment points that need to be traversed to forward a packet from the *CORE* to the host, see Section 2.2. To support an arbitrary number of hierarchical layers, locators can have variable length. Since hosts can have multiple locators, traffic engineering, multi-homing and multi-path can easily be achieved by tweaking the mapping service.

#### 2.1.3 Edge-based approach

Past evolution in the Internet has shown that it is easier to introduce innovation at the “edge” rather than in the “core”. With HAIR, we propose an *edge-based hybrid so-*

lution. While some lightweight functionality is added to a limited set of devices in the network (i.e., to routing domain borders), most tasks, e.g., querying the mapping system for the locator, is done by the end hosts.

## 2.2 Packet delivery

Sending a packet from a source host to a destination host consists of three steps: (i) forwarding the packet up to the *CORE*, (ii) forwarding within the *CORE* and (iii) sending the packet from the *CORE* down to the destination host. For (i), the source *INT* can either use a default route, leverage the loose source route encoded in the source locator, or combine these approaches. Routing within the *CORE* is based on the destination *CORE* attachment point, that is, the first address that appears in the destination locator. Finally, routing from the *CORE* towards the destination, part (iii), follows the loose source route encoded in the destination locator: Whenever one of the *APs* specified in the destination locator is traversed, forwarding is continued based on the next address in the sequence.

Figure 2 illustrates packet forwarding from source *A* (identifier  $ID_A$ ) to destination *B* (identifier  $ID_B$ ). We assume a 3-layer hierarchy where host *A* is reachable via a *CORE AP* ( $CAP_A$ ) and an *INT AP* ( $IAP_A$ ). *A*'s locator then is  $LOC_A = CAP_A|IAP_A$ , and host *B* has locator  $LOC_B = CAP_B|IAP_B$ . First *A* has to find its identifier  $ID_B$ , e.g., via DNS. Next it queries the mapping system to determine *B*'s locator. The composition of  $CAP_B|IAP_B|ID_B$  is the destination address of the packet. In addition, *A* includes its own source address, namely  $CAP_A|IAP_A|ID_A$ .

Using, e.g., a default route, this packet is forwarded to the *CORE* as *A* and *B* are not in the same *INT*. If *A* and *B* are in the same *INT*, i.e., if  $CAP_A == CAP_B$ , the next step is skipped. Within the *CORE* routing is based on the *CAP* portion of the destination address:  $CAP_B$ . Once the packet reaches  $CAP_B$ , it is handed over to the *INT*. Routing is now based on the *INT* attachment point:  $IAP_B$ . Finally, the packet reaches the *EDGE*. It is now  $IAP_B$ 's responsibility to resolve the identifier,  $ID_B$ , to a layer-2 address and forward the packet to destination *B*.

We point out that direct “peerings” between two *INTs*, e.g. *I1* and *I2*, can be achieved by having *I2* export the addresses of (some of) its *CAPs* directly to *I1*, and vice versa. In this case, traffic does not need to traverse the *CORE*.

## 2.3 Mapping system

Our hierarchical mapping system mirrors the structure of the routing architecture. It consists of a *global directory service*, provided by the *CORE*, and a set of *Intermediate Network Mapping Service (IMSs)* maintained by one or a set of *INTs*. The global directory stores for all identifiers a pointer to the *IMS* which currently holds the mapping, whereas the actual mappings are kept in the *IMSs*. Since *IMSs* are ad-

ministered by *INT* networks, the control over the mappings remains with the *INTs* who manage the *IMSs*.

While HAIR proposes the use of a DHT to resolve identifiers to *IMSs*, we do not put any restrictions on how to implement the *IMS* service. In principle, any distributed directory service (e.g., DNS) can be used for both the global directory at Level 1 as well as the *IMSs*. To recruit enough servers for the global directory, registries that assign resources (such as AS numbers or IP addresses) may require each AS to dedicate resources to the global DHT. Participation in the global DHT directory requires authentication and authorization by a third party, e.g., routing registries or IANA.

Our motivation for a hierarchical scheme is twofold: First, by controlling their own mappings via the *IMS* component, *INTs* are able to perform effective inbound traffic engineering. Second, the hierarchical structure of the mapping system allows us to keep some of the updates local to a single *IMS* or to a handful of cooperating *IMSs*. As an example, whenever a host moves between *EDGES* attached to the same *INT* (e.g., a laptop moving inside an office department or between hotspots of the same provider in airports, coffee shops etc.), it is sufficient to change the mappings in the *IMS* without updating the global directory.

## 2.4 Dynamics

**Link/Router failure:** A failure happening inside a routing domain (e.g., a *INT* or the *CORE*) never generates routing updates that exceed the boundary of the routing domain itself. It suffices that the routing protocol finds an alternative route between all pairs of attachment points. This is not always true in today's Internet, as intra-domain routing protocols can (e.g., due to hot-potato routing) generate BGP updates in response to local link or router failures.

**Failing or unreachable attachment point:** If an attachment point fails or is disconnected from the network, all locators which include it have to be updated. This implies updating portions of the mapping system. The *IMS* can monitor its attachment points and, if they change, update its bindings accordingly. Ongoing sessions can be easily handled if end hosts detect that the currently used locator is not working. Since *IMSs* can return multiple locators for a given identifier, after detecting that the other side's locator is not working, the end host can switch to an alternative locator.

**Change of locator:** Locator changes occur when a node moves to another network or when an *IMS* changes the mapping, e.g., for traffic engineering purposes. In the latter case, the updates automatically propagate to the hosts that do not have an entry in their cache or that update their cached entries. The former case differs from the attachment point failure case, as the host is usually aware of the change. Therefore a host can notify end points of active transport sessions by sending them a packet with its identifier and its new locator, as soon as it learns its new locator.

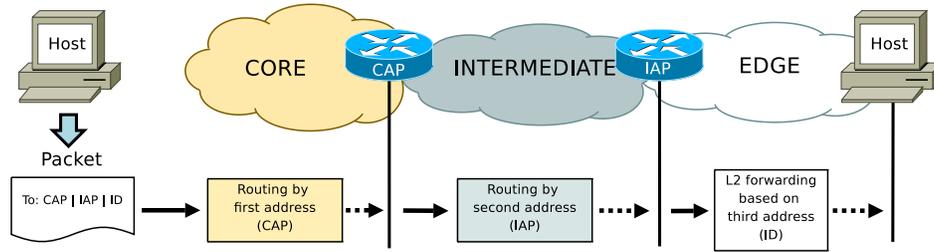


Figure 2: Packet forwarding in HAIR.

### 3. EVALUATION

#### 3.1 Requirement evaluation

**Scalability of the routing system:** Routing is done on a local scope inside *EDGES*, *INTs* and *CORE*. We expect that large ISPs in the *CORE* will have a limited number of *CAPs*, e.g., a handful for each Point of Presence (PoP). Given that the number of PoPs inside today’s top tier providers is generally limited and stable over time [18], we do not expect any scalability problems with respect to routing table size. Due to the physical redundancy that is inherent to the *CORE* [12], the need for routing updates is reduced.

**Scalability of the mapping system:** Since each *INT* operates its own mapping service, the number of entries per *IMS* is limited. Scaling the global directory can be achieved by relying on a DHT. Moreover, the hierarchical design of HAIR ensures that updates to the mapping system are often kept local to one *IMS*.

**Mobility:** Support for mobility is inherent in HAIR due to separation of locators and identifiers. For ongoing transport sessions, a remote endpoint learns about the new locator of its counterpart as soon as the first end-to-end data packet (carrying the new locator) is received. For newly initiated sessions, the mapping system needs to be updated.

**Traffic engineering (TE):** By choosing which locator(s) to assign to which destination, ISPs have a new knob for traffic engineering. Each *INT* can influence where traffic enters its network by changing the locators in its *IMS*. Such updates do not need time to converge, as no routing protocol is involved. Hence, each *INT* can do inbound traffic engineering at a host granularity and prevent other hosts from interfering with its network-wide TE policies.

**Easy migration:** To achieve incremental deployment, we basically need to provide the *CAP* functionality for transit providers and to deploy special gateways inside each *EDGE* that translate packets, sent from legacy hosts, into HAIR packets and vice versa. For this purpose, we propose to leverage existing devices such as firewalls, gateways or proxies that already exist anyway in today’s LANs. First tests with a proof-of-concept implementation of such middle-boxes are promising.

#### 3.2 Estimation of benefits

We estimate the potential benefits to the routing system if HAIR were deployed in today’s Internet. For our analysis we rely on two data sources: (i) BGP updates and table dumps from RIPE [16], and (ii) classification of the ASs [3] according to their type of business into Large Transit Provider (LTP), Small Transit Provider (STP), Enterprise Customer (EC) and Content/Access/Hosting Provider (CAHP). While the former reveals information about routing table size and update churn in today’s Internet, the latter allows us to distinguish “core” from “edge” ASs. Most ASs (31,704) are EC networks, followed by STP (1,663), CAHP (979) and LTP (30) domains.

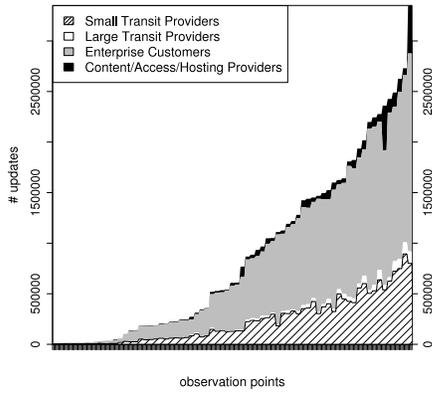
To evaluate the benefits of HAIR<sup>1</sup> on the routing table size, we have to identify the pieces of the current Internet that would form the 3 layers. For this we assume that the *CORE* includes the backbone routers of large and small transit providers (LTPs and STPs) while CAHP and EC form the *INTs*. Given the 1,700 STP and LTP ASs that are then part of the *CORE* and based on the assumption that every STP/LTP operates less than 100 PoPs on average [18], we estimate that the total number of locators that need to be routable inside the *CORE* is less than  $1,700 \cdot 100$ . As such we now have  $\approx 170,000$  locators rather than 300,000 prefixes. Overall, this suggests a considerable improvement over the current status.

Next, we study how effectively HAIR can keep updates local. For this we rely on the updates collected each year in November from 2001 to 2008. For each update trace, we check for each update if it affects a prefix originated by a LTP, STP, EC or CAHP domain. Our assumption is that an update will not be visible in the *CORE*, formed by LTP and STP domains, if it affects a prefix from EC or CAHP.

Figure 3 shows the number of updates observed at our observation points during the first week of November 2008<sup>2</sup>. As some observation points only propagate updates for a small subset of prefixes, we sort the observation points by the number of prefixes they receive along the *x*-axis. The stacked area plot partitions the number of updates, seen at each of our observation points, into the four categories LTP,

<sup>1</sup>For simplicity we consider only a 3-layer deployment.

<sup>2</sup>The results for other years are very similar.



**Figure 3: Number of updates per type across multiple observation points, first week of November, 2008.**

STP, EC or CAHP. Hence, the y-value represents the total number of updates, seen at an observation point.

Figure 3 reveals that a large fraction of updates – for most observation points more than 60% – are due to prefixes originated by EC networks. This number is significantly lower for STP ( $\approx 30\%$ ) and LTP networks ( $\approx 5\%$ ), suggesting a considerable reduction in updates rates for CORE routers if HAIR were deployed in today’s Internet. Since the Internet is growing much faster at the edge than in the core [3], we do not expect scalability problems as in today’s Internet.

#### 4. POC IMPLEMENTATION

We now present a proof-of-concept implementation of our architecture. The code is available online<sup>3</sup>.

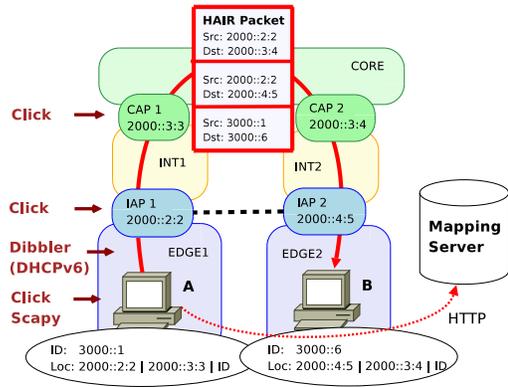
For our test setup, see Figure 4, we implement *all forwarding elements*: Two end hosts *A* and *B*, the *IAP* (*IAP1* and *IAP2*), and the *CAP* (*CAP1* and *CAP2*) devices. Major stumbling blocks in our implementations are transparency to the transport protocol layer and use of standard IP forwarding and existing network applications. Therefore, we decide to use IPv6 addresses for both locators and identifiers<sup>4</sup>. To facilitate future setups of HAIR in testbeds such as Planet-Lab, we try to put new functionality in the *user space* as far as possible. Finally, our implementation incorporates means for easy *bootstrapping*: When an end host enters a *EDGE*, it automatically obtains its *CAPs* and *IAPs* from a DHCPv6 server (*Dibbler*) and thus composes its locator.

To send packets from *A* to *B*, we have implemented three components: End hosts, *IAP*, and *CAP*. At *end hosts* all packets follow a default route to a *tun* interface and are passed to userspace. If the locator for the destination identifier in the packet is unknown, an HTTP-like query is sent to a well-known mapping server and resolved mappings are cached locally<sup>5</sup>. Both the *IAP* and the *CORE* run instances

<sup>3</sup><http://sites.google.com/site/hairarchsite>

<sup>4</sup>Locator and identifier namespaces are nevertheless disjoint.

<sup>5</sup>For now the mapping service only keeps mappings in a text file.



**Figure 4: HAIR: Proof-of-Concept implementation**

of *Click* software router, and decide based on the locators of the incoming packet where to forward the packet. Our implementation also supports direct peerings such as the dashed thick link between *IAP1* and *IAP2*.

We estimate the latency that packets from *A* to *B* experience in HAIR. For this, we rely on `ping6` and compare against the latency that we observe with native IPv6 in an otherwise unchanged setup. There is at least an additional delay of 8 ms compared to the “native” IPv6 setup, with higher latencies for the first ICMP requests due to lookups to the mapping system. Closer scrutiny reveals that a large portion of additional delay is due to generating the HAIR packet in userspace at the end host.

In addition to latency, we analyzed throughput relying on `iperf`. Again, the main bottleneck in forwarding is composing the HAIR packets at the end host in user space. Compared to “native” IPv6 setup, throughput is reduced due to higher CPU utilization at the end host. Moreover, we measure with `iperf` the throughput in a mobility scenario, see Figure 5. After starting the data transfer at time 0:05, host *A* moves<sup>6</sup> to *EDGE2* at 0:15. Host *A* requests the new gateway and *IAP* at time 00:37 (bootstrapping, marked as BP). The connection recovers when host *B* has acknowledged the retransmission of host *A*. Due to the exponential backoff algorithm of TCP, this is not until 01:14.

In principle, higher performance can be achieved by integrating HAIR functionality into the kernel. Nonetheless, our proof-of-concept demonstrates the general feasibility of our approach and highlights that HAIR only requires changes to a small number of devices.

#### 5. RELATED WORK

Many proposals (e.g., [5, 9, 14]) decouple the identifier from the locator functionality within IP addresses. Certain solutions such as Mobile IP or HIP [14] are not designed to reduce routing table size. Contrary to this, HLP [19] does

<sup>6</sup>For this, we only need to change VLANs in our setup

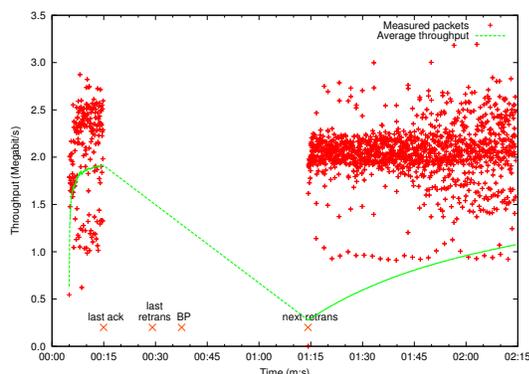


Figure 5: Throughput while A moving to *EDGE2*.

tackle routing table growth and high update churn but without Loc/ID split. Ahlgren et al. [1] adopt a hybrid approach with a hierarchy of heterogeneous networks and support for mobility and communication across them. More radical approaches e.g., [2] route completely based on flat labels.

Previous work, e.g., [5,9,20] frequently adopts a *core/edge separation* approach. For example LISP [5] assigns to end hosts IP addresses with limited scope and translates end host identifiers into global-routable IP addresses, owned by the transit providers, before sending a packet. While LISP uses encapsulation to tunnel packets through the Internet backbone, the Six/One router [20] applies a NAT-style translation between identifier and locator addresses.

Core/edge separation proposals are generally network-based: Gateways in the access network need to look up the locator for incoming packets, which potentially requires to buffer packets, and then to translate/tunnel packets through the core. HAIR borrows from host-based techniques such as shim6 [15], but tackles the problem of routing table growth or update churn in a similar way as network-based solutions.

The mapping service can be implemented in different ways. Some solutions (e.g., [14]) extend DNS and rely on DNS lookups to determine locators for a given identifiers. LISP proposes to use overlays but leaves the choice of whether to use LISP-ALT [4], which is based on an overlay of BGP routers, LISP-DHT [13], relying on DHT techniques, or any other solution that understands LISP queries. [6, 7] adopt an hierarchical scheme for the mapping system.

## 6. CONCLUSION

We introduce HAIR, a routing architecture for the future Internet. HAIR separates locators from identifiers and relies on an hierarchical scheme for both routing *and* the required mapping system that leverages the hierarchical structure of today's Internet. Contrary to other proposals [5, 15], we use a hybrid *edge-based* approach where only some lightweight functionality is added within the network, while the majority of tasks are performed as close to the end hosts as possible.

We present a working proof-of-concept implementation of HAIR and demonstrate that routing table size as well as update load could be substantially reduced. Our future focus will be on the implementation and evaluation of the mapping service and on a security model for HAIR.

## 7. ACKNOWLEDGMENT

The research results presented herein are partly funded by a grant from Deutsche Telekom Laboratories and from G-Lab, a research project (support code 01 BK 0805) funded by the Federal Ministry of Education and Research of the Federal Republic of Germany.

## 8. REFERENCES

- [1] AHLGREN, B., ÅRKKO, J., EGGERT, L., AND RAJAHALME, J. A Node Identity Internetworking Architecture. In *Proc. IEEE Global Internet* (2006).
- [2] CAESAR, M., CONDIE, T., KANNAN, J., LAKSHMINARAYANAN, K., AND STOICA, I. ROFL: Routing On Flat Labels. In *Proc. ACM SIGCOMM* (2006).
- [3] DHAMDHERE, A., AND DOVROLIS, C. Ten Years in the Evolution of the Internet Ecosystem. In *Proc. ACM IMC* (2008).
- [4] FARINACCI, D., FULLER, V., AND MEYER, D. LISP Alternative Topology (LISP+ALT). Internet draft, draft-ietf-lisp-alt-01.txt, 2009.
- [5] FARINACCI, D., FULLER, V., ORAN, D., MEYER, D., AND BRIM, S. Locator/ID Separation Protocol (LISP). Internet draft, draft-ietf-lisp-04.txt, 2009.
- [6] HANKA, O., KUNZMANN, G., SPLEISS, C., EBERSPÄCHER, J., AND BAUER, A. HiiMap: Hierarchical Internet Mapping Architecture. In *Proc. of Future Information Networks* (2009).
- [7] HANKA, O., SPLEISS, C., KUNZMANN, G., AND EBERSPÄCHER, J. A novel DHT-based network architecture for the Next Generation Internet. In *Proc. ICN* (2009).
- [8] HUSTON, G. BGP Routing Table Analysis Reports. <http://bgp.potaroo.net/>.
- [9] JEN, D., MEISEL, M., MASSEY, D., WANG, L., ZHANG, B., AND ZHANG, L. APT: A Practical Transit Mapping Service. Internet draft, draft-jen-apt-01.txt, 2007.
- [10] JEN, D., MEISEL, M., YAN, H., MASSEY, D., AND ZHANG, L. Towards a New Internet Routing Architecture: Arguments for Separating Edges from Transit Core. In *Proc. of HotNets-VII* (2008).
- [11] KRIUKOV, D., CLAFFY, K., FALL, K., AND ARTHUR, B. On Compact Routing for the Internet. *ACM CCR* 37, 3 (2007), 41–52.
- [12] LABOVITZ, C., AHUJA, A., AND JAHANIAN, F. Experimental Study of Internet Stability and Backbone Failures. In *Proc. International Symposium on Fault-Tolerant Computing* (1999).
- [13] MATHY, L., IANNONE, L., AND BONAVENTURE, O. LISP-DHT: Towards a DHT to Map Identifiers onto Locators. February 2008.
- [14] MOSKOWITZ, R., AND NIKANDER, P. Host Identity Protocol (HIP), RFC 4423, 2006.
- [15] NORDMARK, E., AND BAGNULO, M. Shim6: Level 3 Multihoming Shim Protocol for IPv6, RFC 5533. 2009.
- [16] RIPE's Routing Information Service. <http://data.ris.ripe.net/>.
- [17] SIGANOS, G., FALOUTSOS, M., FALOUTSOS, P., AND FALOUTSOS, C. Powerlaws and the AS-level Internet Topology. *ACM/IEEE Transactions on Networking* 11, 4 (2003), 514–524.
- [18] SPRING, N., MAHAJAN, R., AND WETHERALL, D. Measuring ISP Topologies with Rocketfuel. In *Proc. ACM SIGCOMM* (2002).
- [19] SUBRAMANIAN, L., CAESAR, M., EE, C., HANDLEY, M., MAO, M., SHENKER, S., AND STOICA, I. HLP: A Next Generation Inter-domain Routing Protocol. In *Proc. ACM SIGCOMM* (2005).
- [20] VOGT, C. Six/One Router: A Scalable and Backwards Compatible Solution for Provider-Independent Addressing. In *Proc. MobiArch* (2008).